



System z9

Processor Resource/Systems Manager Planning Guide

SB10-7041-03





System z9

Processor Resource/Systems Manager Planning Guide

SB10-7041-03

Note

Before using this information and the product it supports, be sure to read the general information under “Safety and Environmental Notices” on page xi and Appendix C, “Notices,” on page C-1.

Fourth Edition (May 2008)

- | This edition, SB10-7041-03, applies to the IBM® System z9® Servers.
- | This edition replaces SB10-7041-02. A technical change to the text or illustration is indicated by a vertical line to the left of the change.

There may be a newer version of this document in **PDF** format available on **Resource Link™**.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

© **Copyright International Business Machines Corporation 2005, 2008. All rights reserved.**

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

Contents

Figures	vii
Tables	ix
Safety and Environmental Notices	xi
Safety Notices	xi
World Trade Safety Information	xi
Laser Safety Information.	xi
Laser Compliance	xi
Environmental Notices	xi
Product Recycling and Disposal	xi
Refrigeration.	xiii
Battery Return Program	xiii
Flat Panel Display	xiv
Monitors and Workstations	xiv
About This Publication	xvii
What is Included in this Publication	xix
Related Publications	xx
z/Architecture	xx
Enterprise Systems Architecture/390 (ESA/390)	xx
Hardware	xx
Software	xx
How to Send Your Comments	xxii
Summary of changes	xxiii
Chapter 1. Introduction to Logical Partitions.	1-1
Overview	1-2
Prerequisites for Operation	1-2
PR/SM	1-2
zSeries Parallel Sysplex Support	1-4
Guest Coupling Simulation	1-5
Control Program Support in a Logical Partition	1-5
Input/Output Configuration Program (IOCP) Support	1-14
Hardware Support	1-14
Operator Training	1-14
Logical Partitions	1-15
Characteristics	1-15
Potential Applications.	1-18
Compatibility and Migration Considerations.	1-20
Device Numbers	1-20
Multiple Subchannel Sets (MSS).	1-20
Control Programs	1-20
CPU IDs and CPU Addresses	1-22
HSA Allocation	1-24
HSA Estimation Tool	1-24
TOD Clock Processing	1-25
No Sysplex Timer Attached and Server Time Protocol Not Enabled	1-25
Sysplex Timer Attached	1-25
Server Time Protocol Enabled	1-25
Sysplex Testing Without a Sysplex Timer and Server Time Protocol Not Enabled	1-26

Synchronized Time Source and the Coupling Facility	1-26
Extended TOD-Clock Facility	1-26
Chapter 2. Planning Considerations	2-1
Overview	2-3
Planning the I/O Configuration	2-3
Planning Considerations	2-3
Maximum Number of Logical Partitions.	2-5
Managing Logical Paths for ESCON and FICON Channels	2-7
Managing the Establishment of Logical Paths	2-11
Shared Channel Overview	2-21
Unshared ESCON or FICON Channel Recommendations	2-27
Dynamically Managed CHPIDs	2-27
IOCP Coding Specifications	2-28
Coupling Facility Planning Considerations	2-39
Test or Migration Coupling Configuration.	2-39
Production Coupling Facility Configuration	2-40
Internal Coupling Facility (ICF)	2-41
Dynamic Internal Coupling Facility (ICF) Expansion.	2-42
Enhanced Dynamic ICF Expansion Across ICFs	2-43
System-Managed Coupling Facility Structure Duplexing	2-44
Single CPC Software Availability Sysplex	2-44
Coupling Facility Nonvolatility	2-45
Coupling Facility Mode Setting	2-45
Coupling Facility LP Definition Considerations.	2-46
Coupling Facility LP Storage Planning Considerations.	2-47
Dump Space Allocation in a Coupling Facility	2-48
Coupling Facility LP Activation Considerations.	2-49
Coupling Facility Shutdown Considerations.	2-49
Coupling Facility LP Operation Considerations	2-49
Coupling Facility Control Code Commands	2-50
Coupling Facility Level (CFLEVEL) Considerations	2-50
Coupling Facility Resource Management (CFRM) Policy Considerations	2-54
Coupling Facility Channels	2-54
Considerations when Migrating from ICMF to ICs	2-59
Linux Operating System Planning Considerations	2-59
Integrated Facility for Linux (IFL)	2-59
z/VM Version 5 Utilizing IFL Features	2-60
System z9 Application Assist Processor (zAAP)	2-60
IBM System z9 Integrated Information Processor (zIIP)	2-61
Concurrent Patch	2-61
CFCC Enhanced Patch Apply.	2-62
Dynamic Capacity Upgrade on Demand	2-62
PR/SM Shared Partitions	2-63
Mixed Shared and Dedicated PR/SM Partitions	2-63
Multiple Dedicated PR/SM Partitions	2-64
Shared Internal Coupling Facility	2-64
Dynamic Capacity Upgrade on Demand Limitations.	2-65
Concurrent Memory Upgrade	2-66
Capacity Backup Upgrade (CBU) Capability	2-66
Enhanced Book Availability.	2-67
Preparing for Enhanced Book Availability	2-67
Customer Initiated Upgrade (CIU)	2-70
Concurrent Processor Unit Conversion	2-70
Planning for Hot Plugging Crypto Features	2-71

Chapter 3. Determining the Characteristics of Logical Partitions	3-1
Planning Overview	3-4
Performance Considerations	3-4
Recovery Considerations	3-5
Determining the Characteristics	3-5
Control Program Support	3-5
IOCDS Requirements	3-6
Logical Partition Identifier	3-6
Mode of Operation	3-7
Storage Configurations	3-7
Central Storage	3-8
Expanded Storage	3-10
Dynamic Storage Reconfiguration	3-12
Number of Central Processors	3-28
Central Processor Recommendations for Intelligent Resource Director (IRD)	3-29
Processor Considerations for Linux-Only LPs	3-30
Processor Considerations for Coupling Facility LPs	3-30
Processor Considerations for LPs with Multiple CP Types	3-35
Dedicated Central Processors	3-35
Shared Central Processors	3-36
Enforcement of Processing Weights	3-39
Defining Shared Channel Paths	3-48
Dynamic CHPID Management (DCM) Considerations	3-50
I/O Priority Recommendations	3-50
Security-Related Controls	3-50
Dynamic I/O Configuration	3-52
Assigning Channel Paths to a Logical Partition	3-52
Automatic Load for a Logical Partition	3-55
Defining Logical Partitions	3-55
Global Reset Profile Definitions	3-57
General	3-60
Processor Characteristics	3-63
Security Characteristics	3-68
Establishing Optional Characteristics	3-71
Storage Characteristics	3-73
Load Information	3-75
Cryptographic Characteristics	3-77
Enabling Input/Output Priority Queuing	3-81
Changing Logical Partition Input/Output Priority Queuing Values	3-81
Moving Unshared Channel Paths	3-83
Moving Unshared Channel Paths from a z/OS System	3-83
Moving a Channel Path from the Hardware Console	3-83
Releasing Reconfigurable Channel Paths	3-83
Configuring Shared Channel Paths	3-84
Deconfiguring Shared Channel Paths	3-84
Removing Shared Channel Paths for Service	3-84
Changing Logical Partition Definitions	3-84
Changes Available Dynamically to a Running LP	3-84
Changes Available at the Next LP Activation	3-85
Changes Available at the Next Power-On Reset (POR)	3-86
Chapter 4. Operating Logical Partitions	4-1
Overview	4-2
Available Operator Controls	4-2
Operator Controls Not Available	4-4
Operator Tasks	4-5

Editing Activation Profiles	4-5
Activating a CPC	4-5
Activating an LP	4-5
Performing a Load on an LP or Activating a Load Profile	4-5
Deactivating an LP	4-6
Locking and Unlocking an LP	4-6
Deactivating a CPC	4-7
Chapter 5. Monitoring the Activities of Logical Partitions.	5-1
Overview	5-2
Monitoring Logical Partition Activity	5-2
Reviewing Current Storage Information.	5-2
Reviewing and Changing Current Channel Status	5-3
Reviewing and Changing Current Logical Partition Controls	5-3
Reviewing and Changing Current Logical Partition Group Controls	5-3
Reviewing and Changing Current Logical Partition Security	5-5
Reviewing Current Logical Partition Cryptographic Controls	5-6
Reviewing Current System Activity Profile Information	5-6
Reviewing and Changing Logical Partition I/O Priority Values	5-7
Logical Partition Performance	5-8
RMF LPAR Management Time Reporting	5-8
Dedicated and Shared Central Processors	5-9
CPENABLE	5-9
Start Interpretive Execution (SIE) Performance	5-9
Recovery Strategy	5-10
Operation Considerations	5-10
Application Preservation	5-11
Transparent Sparing	5-11
Appendix A. Coupling Facility Control Code Support	A-1
Legend	A-1
Appendix B. Developing, Building, and Delivering a Certified System	B-1
Creating Common Criteria-Based Evaluations	B-1
Functional Characteristics	B-2
Trusted Configuration	B-2
System z9 PR/SM Characteristics	B-4
Central and Expanded Storage	B-4
I/O Security Considerations	B-5
IOCDS Considerations.	B-5
Operational Considerations	B-6
Input/Output Configuration Data Set (IOCDS)	B-8
LPAR Input/Output Configurations	B-8
Activation	B-9
Control Authority	B-9
Reconfiguring the System	B-10
Trusted Facility Library	B-14
Appendix C. Notices	C-1
Trademarks and Service Marks	C-2
Electronic Emission Notices.	C-3
Glossary	D-1
Index	X-1

Figures

I	1-1.	Characteristics of Logical Partitions	1-17
	1-2.	Migration of Four Production Systems to LPs	1-19
	1-3.	Support for Three XRF Systems	1-19
	1-4.	CPU ID Format	1-22
	1-5.	CPU Identification Number Format	1-22
	2-1.	An ESCON Configuration that Can Benefit from Better Logical Path Management	2-9
	2-2.	A Shared ESCON Configuration that Can Benefit from Better Logical Path Management	2-10
	2-3.	Deactivating Unneeded Logical Partitions	2-14
	2-4.	Configuring Offline Unneeded Channels or Shared Channels on an LP Basis.	2-15
	2-5.	Defining Devices to a Subset of Logical Partitions.	2-17
	2-6.	Defining Devices to a Subset of Logical Partitions.	2-18
	2-7.	Using the ESCD to Manage Logical Paths by Prohibiting Dynamic Connections.	2-20
	2-8.	Progression of Busy Condition Management Improvements	2-22
	2-9.	Consolidating ESCON Channels and ESCON Control Unit Ports	2-24
	2-10.	Consolidating ESCON Channels and ESCD Ports.	2-25
	2-11.	Consolidating ESCON Channels Used for ESCON CTC Communications	2-26
	2-12.	Shared Devices Using Shared ESCON Channels	2-31
	2-13.	Physical Connectivity of Shared Device 190	2-32
	2-14.	Logical View of Shared Device 190	2-33
	2-15.	LPAR Configuration with Duplicate Device Numbers	2-34
	2-16.	Duplicate Device Numbers for Console.	2-35
	2-17.	Two Examples of Duplicate Device Number Conflicts	2-36
	2-18.	Example of a Prepare for Enhanced Book Availability Results Panel	2-68
	2-19.	Reassign Non-Dedicated Processors Panel	2-70
	3-1.	Central Storage Layout	3-14
	3-2.	Reconfigured Central Storage Layout	3-15
	3-3.	Initial Central Storage Layout	3-16
	3-4.	Central Storage Layout Following Reconfiguration.	3-17
	3-5.	Initial Central Storage Layout	3-18
	3-6.	Central Storage Layout Following Reconfiguration.	3-19
	3-7.	Expanded Storage Layout	3-20
	3-8.	Reconfigured Expanded Storage Layout	3-21
	3-9.	Expanded Storage Layout	3-22
	3-10.	Initial Central Storage Layout	3-23
	3-11.	Initial Expanded Storage layout	3-23
	3-12.	Central Storage Layout Following Reconfiguration.	3-24
	3-13.	Expanded Storage Layout Following Reconfiguration.	3-25
	3-14.	Backup Partition Layout Before Nonspecific Deactivation	3-27
	3-15.	Backup Partition Layout After Nonspecific Deactivation	3-27
	3-16.	Options Page, Reset Profile	3-57
I	3-17.	Partitions Page, Reset Profile	3-59
	3-18.	General Page, Image Profile	3-60
	3-19.	Time Offset, Image Profile	3-61
	3-20.	ESA Mode Logical Partition with shared Central Processors (CPs), Integrated Facilities for Applications (IFAs) and System z9 Integrated Information Processors (zIIPs).	3-63
	3-21.	Customization for a Linux-Only Mode Logical Partition with shared Integrated Facilities for Linux (IFLs). There can be both an initial and reserved specification for the IFLs.	3-64
	3-22.	Customization for a Coupling Facility Mode Logical Partition with Dedicated Internal Coupling Facilities (ICFs) and shared Central Processors. There can be both an initial and reserved specification for the ICFs and the Central Processors.	3-65
	3-23.	Security Page, Image Profile	3-69
	3-24.	Options Page, Image Profile.	3-71
	3-25.	Storage Page, Image Profile.	3-73

3-26.	Load Page, Image Profile	3-75
3-27.	Crypto Page, Image Profile	3-77
I 3-28.	Enabling I/O Priority Queuing	3-81
3-29.	Change LPAR I/O Priority Queuing	3-81
5-1.	Storage Information Task	5-2
5-2.	Change Logical Partition Controls Page	5-3
5-3.	Change LPAR Group Controls	5-4
5-4.	Change Logical Partition Security Page	5-5
5-5.	View LPAR Cryptographic Controls Page	5-6
5-6.	Change LPAR I/O Priority Queuing Page	5-7
5-7.	ETR Increasing with CPU Utilization	5-8

Tables

1.	Terminology Used in This Publication	xvii
1-1.	CPU IDs for a z9 EC	1-23
2-1.	HCD Function Support	2-3
2-2.	z/VM Dynamic I/O Support for MIF and the Coupling Facility	2-4
2-3.	Logical Path Summary by Control Unit	2-8
2-4.	MIF Maximum Channel Requirements	2-23
2-5.	Nonvolatility Choices for Coupling Facility LPs	2-45
2-6.	Coupling Facility Mode Setting	2-46
2-7.	Coupling Facility LP Storage Definition Capabilities	2-48
2-8.	CPC Support for Coupling Facility Code Levels.	2-51
3-1.	Control Program Support on z9	3-6
3-2.	Central Storage Granularity	3-8
3-3.	PR/SM LPAR Processor Weight Management with Processor Resource Capping	3-40
3-4.	PR/SM LPAR Processor Weight Management without Processor Resource Capping	3-41
3-5.	Example of Maintaining Relative Weight of a Capped Logical Partition	3-42
3-6.	Example Selection of Crypto Numbers	3-79
3-7.	LPAR & Crypto Assignments	3-79
A-1.	Coupling Facility Limits at Different Coupling Facility Code Levels	A-1
B-1.	Trusted Facility Library for PR/SM	B-14

Safety and Environmental Notices

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 national languages. 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 safety information in your national language 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 booklet. 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)


Environmental Notices

Product Recycling and Disposal

This unit must be recycled or discarded according to applicable local and national regulations. IBM encourages owners of information technology (IT) equipment to responsibly recycle their equipment when it is no longer needed. IBM offers a variety of product return programs and services in several countries to assist

equipment owners in recycling their IT products. Information on IBM product recycling offerings can be found on IBM's Internet site at <http://www.ibm.com/ibm/environment/products/index.shtml>.

Esta unidad debe reciclarse o desecharse de acuerdo con lo establecido en la normativa nacional o local aplicable. IBM a los propietarios de equipos de tecnología de la información (TI) que reciclen responsablemente sus equipos cuando éstos recomiendan ya no les sean útiles. IBM dispone de una serie de programas y servicios de devolución de productos en varios países, a fin de ayudar a los propietarios de equipos a reciclar sus productos de TI. Se puede encontrar información sobre las ofertas de reciclado de productos de IBM en el sitio web de IBM <http://www.ibm.com/ibm/environment/products/index.shtml>.



Notice: This mark applies only to countries within the European Union (EU) and Norway.

Appliances are labeled in accordance with European Directive 2002/96/EC concerning waste electrical and electronic equipment (WEEE). The Directive determines the framework for the return and recycling of used appliances as applicable throughout the European Union. This label is applied to various products to indicate that the product is not to be thrown away, but rather reclaimed upon end of life per this Directive.

In accordance with the European WEEE Directive, electrical and electronic equipment (EEE) is to be collected separately and to be reused, recycled, or recovered at end of life. Users of EEE with the WEEE marking per Annex IV of the WEEE Directive, as shown above, must not dispose of end of life EEE as unsorted municipal waste, but use the collection framework available to customers for the return, recycling, and recovery of WEEE. Customer participation is important to minimize any potential effects of EEE on the environment and human health due to the potential presence of hazardous substances in EEE. For proper collection and treatment, contact your local IBM representative.

注意: このマークは EU 諸国およびノルウェーにおいてのみ適用されます。

この機器には、EU 諸国に対する廃電気電子機器指令 2002/96/EC(WEEE) のラベルが貼られています。この指令は、EU 諸国に適用する使用済み機器の回収とリサイクルの骨子を定めています。このラベルは、使用済みになった時に指令に従って適正な処理をする必要があることを知らせるために種々の製品に貼られています。

Remarque : Cette marque s'applique uniquement aux pays de l'Union Européenne et à la Norvège.

L'étiquette du système respecte la Directive européenne 2002/96/EC en matière de Déchets des Equipements Electriques et Electroniques (DEEE), qui détermine les dispositions de retour et de recyclage applicables aux systèmes utilisés à travers l'Union européenne. Conformément à la directive, ladite étiquette précise que le produit sur lequel elle est apposée ne doit pas être jeté mais être récupéré en fin de vie.

Refrigeration

This system contains one or more modular refrigeration units with R-134A refrigerant and a polyol ester oil. This refrigerant must not be released or vented to the atmosphere. Skin contact with refrigerant may cause frostbite. Wear appropriate eye and skin protection. Modular refrigeration units are hermetically sealed and must not be opened or maintained.

This notice is provided in accordance with the European Union (EU) Regulation 842/2006 on fluorinated greenhouse gases. This product contains fluorinated greenhouse gases covered by the Kyoto Protocol. Per Annex I, Part 1, of EU Regulation 842/2006, the global warming potential of R-134A is 1300. Each unit contains 1.22 kg of R-134A.

Battery Return Program

This product may contain sealed lead acid, nickel cadmium, nickel metal hydride, lithium, or lithium ion battery(s). Consult your user manual or service manual for specific battery information. The battery must be recycled or disposed of properly. Recycling facilities may not be available in your area. For information on disposal of batteries outside the United States, go to <http://www.ibm.com/ibm/environment/products/index.shtml> or contact your local waste disposal facility.


In the United States, IBM has established a return process for reuse, recycling, or proper disposal of used IBM sealed lead acid, nickel cadmium, nickel metal hydride, and other battery packs from IBM Equipment. For information on proper disposal of these batteries, contact IBM at 1-800-426-4333. Please have the IBM part number listed on the battery available prior to your call.

For Taiwan:



Please recycle batteries 廢電池請回收

For the European Union:



Notice: This mark applies only to countries within the European Union (EU)

Batteries or packaging for batteries are labeled in accordance with European Directive 2006/66/EC concerning batteries and accumulators and waste batteries and accumulators. The Directive determines the framework for the return and recycling of used batteries and accumulators as applicable throughout the European Union. This label is applied to various batteries to indicate that the battery is not to be thrown away, but rather reclaimed upon end of life per this Directive.

Les batteries ou emballages pour batteries sont étiquetés conformément aux directives européennes 2006/66/EC, norme relative aux batteries et accumulateurs en usage et aux batteries et accumulateurs usés. Les directives déterminent la marche à suivre en vigueur dans l'Union Européenne pour le retour et le recyclage des batteries et accumulateurs usés. Cette étiquette est appliquée sur diverses batteries pour indiquer que la batterie ne doit pas être mise au rebut mais plutôt récupérée en fin de cycle de vie selon cette norme.

バッテリーあるいはバッテリー用のパッケージには、EU 諸国に対する廃電気電子機器指令 2006/66/EC のラベルが貼られています。この指令は、バッテリーと蓄電池、および廃棄バッテリーと蓄電池に関するものです。この指令は、使用済みバッテリーと蓄電池の回収とリサイクルの骨子を定めているもので、EU 諸国にわたって適用されます。このラベルは、使用済みになったときに指令に従って適正な処理をする必要があることを知らせるために種々のバッテリーに貼られています。

In accordance with the European Directive 2006/66/EC, batteries and accumulators are labeled to indicate that they are to be collected separately and recycled at end of life. The label on the battery may also include a chemical symbol for the metal concerned in the battery (Pb for lead, Hg for mercury, and Cd for cadmium). Users of batteries and accumulators must not dispose of batteries and accumulators as unsorted municipal waste, but use the collection framework available to customers for the return, recycling, and treatment of batteries and accumulators. Customer participation is important to minimize any potential effects of batteries and accumulators on the environment and human health due to the potential presence of hazardous substances. For proper collection and treatment, contact your local IBM representative.

For Spain:

This notice is provided in accordance with Royal Decree 106/2008. The retail price of batteries, accumulators, and power cells includes the cost of the environmental management of their waste.

For California:

Perchlorate Material - special handling may apply. See <http://www.dtsc.ca.gov/hazardouswaste/perchlorate>.

The foregoing notice is provided in accordance with California Code of Regulations Title 22, Division 4.5, Chapter 33. Best Management Practices for Perchlorate Materials. This product, part, or both may include a lithium manganese dioxide battery which contains a perchlorate substance.

Flat Panel Display

The fluorescent lamp or lamps in the liquid crystal display contain mercury. Dispose of it as required by local ordinances and regulations.

Monitors and Workstations

New Jersey – For information about recycling covered electronic devices in the State of New Jersey, go to the New Jersey Department of Environmental Protection Web site at http://www.state.nj.us/dep/dshw/recycle/Electronic_Waste/index.html.

Oregon – For information regarding recycling covered electronic devices in the State of Oregon, go to the Oregon Department of Environmental Quality Web site at <http://www.deq.state.or.us/lq/electronics.htm>.

Washington – For information about recycling covered electronic devices in the State of Washington, go to the Department of Ecology Web site at <http://www.ecy.wa.gov/programs/swfa/eproductrecycle>, or telephone the Washington Department of Ecology at 1-800Recycle.

About This Publication

This publication is intended for system planners, installation managers, and other technical support personnel who need to plan for operating in logically partitioned (LPAR) mode on the IBM System z9 Enterprise Class (z9 EC) and IBM System z9 Business Class (z9 BC).

This publication assumes previous knowledge of the characteristics and functions of the installed central processor complex (CPC).

To improve readability, we refer to the different CPCs using the following terminology whenever possible:

Table 1. Terminology Used in This Publication

Terminology	Central Processor Complex (CPC)
z9 BC	IBM System z9 Business Class Model R07 Model S07
z9 EC	IBM System z9 Enterprise Class Model S08 Model S18 Model S28 Model S38 Model S54

Some features, panels, and functions are model-dependent, engineering change (EC) level-dependent, machine change level-dependent (MCL-dependent), or control program-dependent. For this reason, not all of the functions discussed in this publication are necessarily available on every CPC.

Sample tasks and panels explained in this publication reference tasks and panels available from the support element console. However, detailed procedures for operator tasks and accurate task and panel references are explained in the *Support Element Operations Guide*.

Some illustrations and examples in this publication describe operation with as few as 2 logical partitions (LPs), although up to 60 LPs can be defined.

Hardware Management Console operators or support element console operators should use the appropriate operations guide for instructions on how to perform tasks. Control program operators should refer to the appropriate control program publication for information on control program commands.

For information about PR/SM™ LPAR mode on prior models, see the following publications:

- *System z9 109 Processor Resource/Systems Manager Planning Guide*, SB10-7041-00
- *zSeries 890 and 990 Processor Resource/Systems Manager Planning Guide*, SB10-7036
- *zSeries 800 and 900 Processor Resource/Systems Manager Planning Guide*, SB10-7033
- *S/390® Processor Resource/Systems Manager Planning Guide*, GA22-7236

- *ES/9000[®] Processor Resource/Systems Manager Planning Guide*, GA22-7123

However, for the most current coupling facility control code information for all models, use this publication.

What is Included in this Publication

The information presented in this publication is organized as follows:

- Chapter 1, “Introduction to Logical Partitions” describes the prerequisites for establishing and using LPAR, the general characteristics and some potential applications for LPs.
- Chapter 2, “Planning Considerations” presents considerations and guidelines for I/O configuration planning and coupling facility planning.
- Chapter 3, “Determining the Characteristics of Logical Partitions” includes a list of the panels, provides guidelines for determining the CPC resources, and describes the operator tasks used to define the characteristics of LPs.
- Chapter 4, “Operating Logical Partitions” describes how to operate the Hardware Management Console and the support element console, and describes the procedure for initializing the system.
- Chapter 5, “Monitoring the Activities of Logical Partitions” describes the panels and operator tasks used to monitor LP activity.
- Appendix A, “Coupling Facility Control Code Support,” on page A-1 lists and explains the support provided at different levels of coupling facility control code Licensed Internal Code (LIC).
- Appendix B, “Developing, Building, and Delivering a Certified System,” on page B-1 provides guidance in setting up, operating, and managing a secure consolidated environment using System z9 PR/SM.
- Appendix C, “Notices,” on page C-1 contains electronic emission notices, legal notices, and trademarks.

Related Publications

The following publications provide information about the functions and characteristics of the different CPCs and the related operating systems that run on them.

z/Architecture®

- *z/Architecture Principles of Operation*, SA22-7832

Enterprise Systems Architecture/390® (ESA/390)

- *Enterprise Systems Architecture/390 Principles of Operation*, SA22-7201

Hardware

System z9 BC

- *System z9 Business Class System Overview*, SA22-1083
- *Input/Output Configuration Program User's Guide*, SB10-7037
- *Stand-alone IOCP User's Guide*, SB10-7152
- *Hardware Management Console Operations Guide*, SC28-6859
- *Support Element Operations Guide*, SC28-6860

System z9 EC

- *System z9 Enterprise Class System Overview*, SA22-6833
- *Input/Output Configuration Program User's Guide*, SB10-7037
- *Stand-alone IOCP User's Guide*, SB10-7152
- *Hardware Management Console Operations Guide*, SC28-6859
- *Support Element Operations Guide*, SC28-6860

ESCON® Concepts

- *Introducing Enterprise Systems Connection*, GA23-0383
- *ESCON and FICON Channel-to-Channel Reference*, SB10-7034

FICON®

- *ESCON and FICON Channel-to-Channel Reference*, SB10-7034

Crypto Features

The following publications provide additional information on the Crypto features:

- *Support Element Operations Guide*, SC28-6860
- *zSeries User Defined Extensions Reference and Guide*, website:
<http://www.ibm.com/security/cryptocards> (Select a crypto card, and then click **Library**)
- *IBM System z9 109 Technical Guide*, SG24-7124

Software

z/OS®

zSeries Parallel Sysplex: The following publications provide additional information about the zSeries in the z/OS environment:

- *z/OS Parallel Sysplex Overview*, SA22-7661
- *z/OS Parallel Sysplex Application Migration*, SA22-7662
- *z/OS MVS Setting Up a Sysplex*, SA22-7625
- *z/OS MVS Programming: Sysplex Services Guide*, SA22-7617
- *z/OS MVS Programming: Sysplex Services Reference*, SA22-7618

Multiple Image Facility: The following publications provide additional information about multiple image facility in the z/OS environment:

- *z/OS Hardware Configuration Definition: User's Guide*, SC33-7988

Dynamic I/O Configuration: The following publication provides information about dynamic I/O configuration in the z/OS environment:

- *z/OS Hardware Configuration Definition Planning*, GA22-7525

Dynamic Storage Reconfiguration: The following publications provide additional information on the commands, functions, and capabilities of dynamic storage reconfiguration in the z/OS environment:

- *z/OS MVS Initialization and Tuning Reference*, SA22-7592
- *z/OS MVS Recovery and Reconfiguration Guide*, SA22-7623
- *z/OS MVS System Commands*, SA22-7627

Crypto Features: The following publications provide additional information on the Crypto features:

- *z/OS ICSF Administrator's Guide*, SA22-7521
- *z/OS ICSF System Programmer's Guide*, SA22-7520
- *z/OS ICSF Trusted Key Entry PCIX Workstation User's Guide*, SA22-7524

Sysplex Failure Manager: The following publication provides an overview of SFM and practical information for implementing and using SFM in the z/OS environment:

- *z/OS MVS Setting Up a Sysplex*, SA22-7625

LPAR Management Time: The following publication provides information about the RMF™ Partition Data Report that includes LPAR Management Time reporting in a z/OS environment:

- *z/OS Resource Measurement Facility User's Guide*, SC33-7990

Intelligent Resource Director (IRD): The following publication provides information about Intelligent Resource Director in a z/OS environment:

- *z/OS Intelligent Resource Director*, SG24-5952

z/VM®

Hardware Configuration Definition (HCD): The following publication provides information about the Hardware Configuration Definition (HCD):

- *z/VM I/O Configuration*, SC24-6100

Hardware Configuration Manager: The following publication provides information about the Hardware Configuration Manager:

- *z/OS and z/VM Hardware Configuration Manager User's Guide*, SC33-7989

Dynamic I/O Configuration: The following publication provides information about dynamic I/O configuration:

- *CP Planning and Administration*, SC24-6083
- *z/VM I/O Configuration*, SC24-6100

Guest Operating Systems: The following publication provides information about running **guest** operating systems:

- *z/VM Running Guest Operating Systems*, SC24-6115

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. 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 or table number).

Summary of changes

Summary of Changes for SB10-7041-03:

This revision contains editorial changes and the following new technical changes:

New Information

- InfiniBand® channel paths (TYPE=CIB), host channel adapter (HCA), InfiniBand coupling links for Parallel Sysplex, channel path information has been added.
- Control Program support information has been updated.

Summary of Changes for SB10-7041-02:

This revision contains editorial changes and the following new technical changes:

New Information

- Some Image/Activation Profile Panels have been re-designed, resulting in performance improvements.
- Change LPAR Group Controls - new task added to the Hardware Management Console and the Support Element Console. This task allows groups of LPARs to have a capacity value that applies to the entire group, and not just the individual image.

Summary of Changes for SB10-7041-01a:

This revision contains editorial changes and the following new technical changes:

New Information

- Information regarding Server Time Protocol (STP) was added.

Summary of Changes for SB10-7041-01:

This revision contains editorial changes and the following new technical changes.

New Information

- This book was updated to include information for both the IBM System z9 Enterprise Class (z9 EC) and the IBM System z9 Business Class (z9 BC).
- Support was added for the IBM System z9 Integrated Information Processor, or zIIP. This is the latest customer-inspired specialty engine for System z9.

Summary of Changes for SB10-7041-00a:

This revision contains editorial changes and the following new technical changes.

Changed Information

- Updates were made to Appendix B, “Developing, Building, and Delivering a Certified System.”

Chapter 1. Introduction to Logical Partitions

Overview	1-2
Prerequisites for Operation	1-2
PR/SM	1-2
Logical Partitioning	1-2
Central Storage	1-2
Expanded Storage	1-2
Central Processors	1-3
Multiple Image Facility	1-3
Crypto Features	1-3
Configurable Single Processor Crypto Express2	1-4
CP Crypto Assist Functions	1-4
zSeries Parallel Sysplex Support	1-4
Guest Coupling Simulation	1-5
Control Program Support in a Logical Partition	1-5
z/OS	1-5
z/VM	1-9
z/VSE	1-11
TPF (Transaction Processing Facility)	1-11
Linux	1-13
Hardware Configuration Definition (HCD)	1-13
z/VM Dynamic I/O Configuration	1-14
Input/Output Configuration Program (IOCP) Support	1-14
Hardware Support	1-14
Operator Training	1-14
Logical Partitions	1-15
Characteristics	1-15
Potential Applications	1-18
Examples of Logical Partition Applications	1-19
Compatibility and Migration Considerations	1-20
Device Numbers	1-20
Multiple Subchannel Sets (MSS)	1-20
Control Programs	1-20
z/OS	1-20
EREP	1-21
CPU IDs and CPU Addresses	1-22
CPU ID Fields	1-22
Examples of CPU ID Information	1-23
HSA Allocation	1-24
HSA Estimation Tool	1-24
TOD Clock Processing	1-25
No Sysplex Timer Attached and Server Time Protocol Not Enabled	1-25
Sysplex Timer Attached	1-25
Server Time Protocol Enabled	1-25
Sysplex Testing Without a Sysplex Timer and Server Time Protocol Not Enabled	1-26
Synchronized Time Source and the Coupling Facility	1-26
Extended TOD-Clock Facility	1-26

Overview

This chapter introduces the characteristics of logical partitioning and migration and compatibility considerations. Processor Resource/Systems Manager™ (PR/SM) is standard on all z9 EC and z9 BC models.

Prerequisites for Operation

The prerequisites for operation are:

- Programming compatibility
- Hardware support
- Operator training
- Programming support
 - Control program support
 - Input/Output Configuration Program (IOCP) support

PR/SM

PR/SM enables logical partitioning of the central processor complex (CPC).

Logical Partitioning

PR/SM enables the logical partitioning function of the CPC. The operator defines the resources that are to be allocated to each logical partition (LP). Most resources can be reconfigured without requiring a power-on reset. After an ESA/390, ESA/390 TPF, or Linux®-Only LP is defined and activated, you can load a supported control program into that LP. If a coupling facility logical partition is defined and activated, the coupling facility control code is automatically loaded into the LP.

Central Storage

Central storage is defined to LPs before LP activation. When an LP is activated, the storage resources are allocated in contiguous blocks. These allocations can be dynamically reconfigured. Sharing of allocated central storage among multiple LPs is not allowed.

On System z9, all storage is defined as central storage. Allocation of storage to logical partitions can be made as either central storage or expanded storage. Any allocation of expanded storage to an LP reduces the amount of storage available for allocation as central storage. See “Single Storage Pool” on page 3-7. If no storage is allocated as expanded storage and no other LP is currently active, an individual LP can have a central storage amount equal to installed storage minus the size of the hardware system area (HSA). The sum total of all LP central and expanded storage cannot exceed the amount of installed storage - HSA. Storage allocated to LPs never includes HSA. Central storage in excess of 2 GB should only be allocated in an LP that has an operating system capable of utilizing the new 64-bit z/Architecture.

Expanded Storage

Optional expanded storage is defined to LPs before LP activation. When an LP is activated, the storage resources are allocated in contiguous blocks. These allocations can be dynamically reconfigured. Sharing of allocated expanded storage among multiple LPs is not allowed.

See “Expanded Storage” on page 3-10.

Central Processors

Central processors (CPs) can be dedicated to a single LP or shared among multiple LPs. CPs are allocated to an LP when the LP is activated. You can use operator tasks to limit and modify the use of CP resources shared between LPs while the LPs are active.

Multiple Image Facility

The multiple image facility (MIF) is available on all CPCs discussed in this publication. MIF allows channel sharing among LPs. For information about accessing devices on shared channel paths and defining shared channel paths, see “Defining Shared Channel Paths” on page 3-48.

MCSS: Multiple Logical Channel Subsystems (CSS) are available on all CPCs discussed in this publication. Each CSS supports a definition of up to 256 channels.

Channel Paths: Active LPs can share channels. Shared channels require that the channel subsystem establish a logical path for each channel image corresponding to an active LP that has the channel configured online. CNC, CTC, OSC, OSD, OSE, OSN, CBP, CFP, CIB, ICP, FC, FCV, FCP, and IQD channel path types can be shared. CVC and CBY channel paths **cannot** be shared.

For information about accessing devices on shared channel paths and defining shared channel paths, see “Defining Shared Channel Paths” on page 3-48.

Crypto Features

An optional feature, Crypto Express2, is available, which is designed for Federal Information Processing Standard (FIPS) 140-2 Level 4 Certification. Crypto Express2 contains two PCI-X adapters, each of which can be configured independently as either a coprocessor or an accelerator. The default is coprocessor. Crypto Express2 can be configured using the Cryptographic Configuration panel found under the Configuration tasks list.

Crypto Express2 Coprocessor (CEX2C) replaces the functions of the PCI X Crypto Coprocessor (PCIXCC), which was available on some previous zSeries processors, such as the z890 and z990.

Crypto Express2 Accelerator (CEX2A) replaces the functions of the PCI Crypto Accelerator (PCICA), which was available on some previous zSeries processors such as the z890 and z990.

Crypto Express2 Coprocessor is used for secure key encrypted transactions, and is the default configuration. CEX2C:

- Supports highly secure cryptographic functions, use of secure encrypted key values, and User Defined Extensions (UDX).
- Is designed for Federal Information Processing Standard (FIPS) 140-2 Level 4 certification.

Crypto Express2 Accelerator is used for SSL acceleration. CEX2A:

- Supports clear key RSA acceleration.
- Offloads compute-intensive RSA public-key and private-key cryptographic operations employed in the SSL protocol.

Configurable Single Processor Crypto Express2

Available exclusively for the z9 BC, the Crypto Express2-1P feature has one PCI-X adapter. The PCI-X adapter can be defined as either a Coprocessor or an Accelerator. A minimum of two features must be ordered:

- Crypto Express2-1P Coprocessor - for secure-key encrypted transactions (default).
 - Supports highly secure cryptographic functions, use of secure encrypted key values, and User Defined Extensions (UDX).
 - Is designed for Federal Information Processing Standard (FIPS) 140-2 Level 4 certification.
- Crypto Express2-1P Accelerator - for Secure Sockets Layer (SSL) acceleration.
 - Supports clear key RSA acceleration.
 - Offloads compute-intensive RSA public-key and private-key cryptographic operations employed in the SSL protocol.

CP Crypto Assist Functions

CP Crypto Assist Functions (CPACF), supporting clear key encryption, offers the following on every Processor Unit (PU) identified as a Central Processor (CP), Integrated Facility for Linux (IFL), System z9 Integrated Information Processor (zIIP), or System z9 Application Assist Processor (zAAP):

- Data Encryption Standard (DES)
- Triple Data Encryption Standard (TDES)
- Advanced Encryption Standard (AES)
- SHA-1
- SHA-256
- Pseudo Random Number Generation (PRNG)

zSeries Parallel Sysplex® Support

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 running in a sysplex and sharing data in a coupling facility.

Parallel sysplex allows you to manage a transaction processing workload, balanced across multiple z/OS images running on multiple 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 sysplex using a coupling facility by making it possible to dynamically reapportion the workload across the remaining active sysplex participants. Additionally, you can dynamically add processing capacity (CPCs or LPs) during peak processing without disrupting ongoing workload processing.

CPC support consists of having the capability to do any or all of the following:

- Install coupling facility channels
- Define, as an LP, a portion or all of the CPC hardware resources (central processors, storage, and coupling facility channels) for use as a coupling facility that connects to z/OS images for data sharing purposes
- Connect to a coupling facility to share data
- Installing a z9 BC or z9 EC with only ICFs (1 or more) operates as a stand-alone coupling facility; it **cannot** run z/OS or any other operating system.

For more information on the coupling facility including z/OS and CPC support for coupling facility levels, see “Coupling Facility Planning Considerations” on page 2-39.

Guest Coupling Simulation

Guest coupling simulation is available with z/VM.

z/VM guest coupling simulation allows you to simulate one or more complete parallel sysplexes within a single system image, providing a test environment for parallel sysplex installation. The simulated environment is not intended for production use since its single points of failure diminish the availability advantages of the parallel sysplex environment. There are no special hardware requirements – external coupling facility channels, external coupling facilities, and Sysplex Timers are neither necessary nor supported. Guest operating systems within a simulated sysplex can only be coupled (through simulated coupling facility channels) to coupling facilities also running as guests of the same z/VM system. You can have up to 32 virtual machines running z/OS within a simulated sysplex, with each z/OS virtual machine coupled to up to 8 virtual machines running as coupling facilities.

There is no system-imposed limit to the number of guest parallel sysplex environments that z/VM can simulate. However, practical limits on the number of guests that can be supported by a particular hardware configuration will also constrain the number of simulated parallel sysplex environments.

Control Program Support in a Logical Partition

Control programs require certain characteristics. Before planning or defining LP characteristics, call your installation management to determine which control programs are in use or planned for operation.

Notes:

1. Use IBM Service Link to view the appropriate PSP bucket subset ID for hardware and software maintenance information.
2. For more detailed information about support for coupling facility levels (including hardware EC, driver, and MCL numbers and software APAR numbers), see “Coupling Facility Level (CFLEVEL) Considerations” on page 2-50.

z/OS

z/OS Release 9 Support

- z/OS Management Console
- CHPID type OSA performance enhancements
- Enhancements to On/Off Capacity on Demand
- Enhancements to z9 BC Capacity Backup
- Group Capacity Limit for LPAR
- LDAP Support for HMC user authentication
- FICON Express4-2C SX on the z9 BC
- CFLEVEL 15 support (with applicable PTF)
- System-initiated CHPID reconfiguration
- Multipath IPL
- OSA-Express Network Traffic analyzer
- QDIO Diagnostic Synchronization
- Power monitoring
- Hardware Decimal Floating Point facilities

- Support for Server Time Protocol (STP)
- Support for z9 BC and z9 EC
- Multiple Subchannel Sets (MSS) for ESCON (CHPID type CNC) and FICON (CHPID type FC)
- Support for FICON Express2 (CHPID type FC), including Channel-To-Channel (CTC)
- OSA Express2 large send (also called TCP segmentation offload)
- Support for zAAP
- Support for zIIP
- Support for up to 54 CPs on z9 models
- System-Managed CF Structure Duplexing support
- Support for multiple LCSS's
- XES Coupling Facility cache structure architecture extensions for Batch Write, Castout, and Cross-Invalidate functions
- z/Architecture 64-bit addressing
- 1 TB of central storage on z9 models
- 256 channel paths
- Coupling facility level 14
- Support for Message Time Ordering providing enhanced scalability of parallel sysplex
- Support z ELC software pricing structure
- Support for zSeries PCI Cryptographic Accelerator feature by the Integrated Cryptographic Services Facility and System SSL functions
- Support for zSeries PCI X Cryptographic Coprocessor feature by the Integrated Cryptographic Services Facility and System SSL functions
- Support Crypto Express2 and Crypto Express2-1P (only available on the z9 BC) and provides access to the secure key and System SSL functions via the Integrated Cryptographic Services Facility.
- Intelligent Resource Director (IRD)
- WorkLoad Pricing
- Peer mode channels (ICP, CFP, CBP)
- WorkLoad Manager (WLM) Multisystem Enclaves support
- XES Coupling Facility List Structure Architecture Extensions [also known as Message and Queuing (MQ) Series]
- Logical partition time offset
- Internal coupling facility channels
- System Managed Rebuild

z/OS Release 8 Support

- CHPID type OSA performance enhancements
- Enhancements to On/Off Capacity on Demand
- Enhancements to z9 BC Capacity Backup
- Group Capacity Limit for LPAR
- LDAP Support for HMC user authentication
- FICON Express4-2C SX on the z9 BC
- CFLEVEL 15 support (with applicable PTF)
- System-initiated CHPID reconfiguration

- Multipath IPL
- OSA-Express Network Traffic analyzer
- QDIO Diagnostic Synchronization
- Power monitoring
- Hardware Decimal Floating Point facilities
- Support for Server Time Protocol (STP)
- Support for z9 BC and z9 EC
- Multiple Subchannel Sets (MSS) for ESCON (CHPID type CNC) and FICON (CHPID type FC)
- Support for FICON Express2 (CHPID type FC), including Channel-To-Channel (CTC)
- OSA Express2 large send (also called TCP segmentation offload)
- Support for zAAP
- Support for zIIP
- Support for up to 32 CPs
- System-Managed CF Structure Duplexing support
- Support for multiple LCSS's
- XES Coupling Facility cache structure architecture extensions for Batch Write, Castout, and Cross-Invalidate functions
- z/Architecture 64-bit addressing
- 1 TB of central storage on z9 models
- 256 channel paths
- Coupling facility level 14
- Support for Message Time Ordering providing enhanced scalability of parallel sysplex
- Support z ELC software pricing structure
- Support for zSeries PCI Cryptographic Accelerator feature by the Integrated Cryptographic Services Facility and System SSL functions
- Support for zSeries PCI X Cryptographic Coprocessor feature by the Integrated Cryptographic Services Facility and System SSL functions
- Support for Crypto Express2 and Crypto Express2-1P (only available on the z9 BC) and provides access to the secure key and System SSL functions via the Integrated Cryptographic Services Facility.
- Intelligent Resource Director (IRD)
- WorkLoad Pricing
- Peer mode channels (ICP, CFP, CBP)
- WorkLoad Manager (WLM) Multisystem Enclaves support
- XES Coupling Facility List Structure Architecture Extensions [also known as Message and Queuing (MQ) Series]
- Logical partition time offset
- Internal coupling facility channels
- System Managed Rebuild

z/OS Release 7 Support

- Multipath IPL
- System-initiated CHPID reconfiguration
- CFLEVEL 15 support (with applicable PTF)

- FICON Express4-2C SX on the z9 BC
- Hardware Decimal Floating Point facilities
- Support for Server Time Protocol (STP)
- Support for z9 BC and z9 EC
- Multiple Subchannel Sets (MSS) for ESCON (CHPID type CNC) and FICON (CHPID type FC)
- Support for FICON Express2 (CHPID type FC), including Channel-To-Channel (CTC)
- OSA Express2 large send (also called TCP segmentation offload)
- Support for zAAP
- Support for zIIP
- Support for up to 32 CPs
- System-Managed CF Structure Duplexing support
- Support for multiple LCSS's
- XES Coupling Facility cache structure architecture extensions for Batch Write, Castout, and Cross-Invalidate functions
- z/Architecture 64-bit addressing
- 128 GB of central storage
- 256 channel paths
- Coupling facility level 14
- Support for Message Time Ordering providing enhanced scalability of parallel sysplex
- Support z ELC software pricing structure
- Support for zSeries PCI Cryptographic Accelerator feature by the Integrated Cryptographic Services Facility and System SSL functions
- Support for zSeries PCI X Cryptographic Coprocessor feature by the Integrated Cryptographic Services Facility and System SSL functions
- Support for Crypto Express2 and Crypto Express2-1P (only available on the z9 BC) and provides access to the secure key and System SSL functions via the Integrated Cryptographic Services Facility.
- Intelligent Resource Director (IRD)
- WorkLoad Pricing
- Peer mode channels (ICP, CFP, CBP)
- WorkLoad Manager (WLM) Multisystem Enclaves support
- XES Coupling Facility List Structure Architecture Extensions [also known as Message and Queuing (MQ) Series]
- Logical partition time offset
- Internal coupling facility channels
- System Managed Rebuild

z/OS Release 6 Support

- System-initiated CHPID reconfiguration (with applicable PTF)
- Multipath IPL (with applicable PTF)
- CFLEVEL 15 support (with applicable PTF)
- FICON Express4-2C SX on the z9 BC (with applicable PTF)
- Hardware Decimal Floating Point facilities (with applicable PTF)
- Support for z9 BC and z9 EC (with applicable PTFs)

- Support for FICON Express2 (CHPID type FC), including Channel-To-Channel (CTC)
- OSA Express2 large send (also called TCP segmentation offload)
- Support for zAAP
- Support for zIIP
- Support for > 16 CPs (up to 32)
- System-Managed CF Structure Duplexing support
- Support for multiple LCSS's
- XES Coupling Facility cache structure architecture extensions for Batch Write, Castout, and Cross-Invalidate functions
- z/Architecture 64-bit addressing
- 128 GB of central storage
- 256 channel paths
- Coupling facility level 14
- Support for Message Time Ordering providing enhanced scalability of parallel sysplex
- Support z ELC software pricing structure
- Support for zSeries PCI Cryptographic Accelerator feature by the Integrated Cryptographic Services Facility and System SSL functions
- Support for zSeries PCI X Cryptographic Coprocessor feature by the Integrated Cryptographic Services Facility and System SSL functions
- Support for Crypto Express2 and Crypto Express2-1P (only available on the z9 BC) with ICSF 64-bit Virtual Support for z/OS V1.6 and z/OS.e V1.6 provides access to the secure key and System SSL functions via the Integrated Cryptographic Services Facility.
- Intelligent Resource Director (IRD)
- WorkLoad Pricing
- Peer mode channels (ICP, CFP, CBP)
- WorkLoad Manager (WLM) Multisystem Enclaves support
- XES Coupling Facility List Structure Architecture Extensions [also known as Message and Queuing (MQ) Series]
- Logical partition time offset
- Internal coupling facility channels
- System Managed Rebuild

z/VM

z/VM Version 5 Release 3 Support

- System-managed Coupling Facility structure duplexing, for z/OS guests
- Dedicated OSA port to an operating system
- CHPID type OSA performance enhancements, for z/OS guests
- CHPID type FCP performance enhancements
- FICON Express4-2C SX on the z9 BC
- Hardware Decimal Floating Point facilities
- z/VM integrated systems management
- Support for z9 BC and z9 EC
- Support for FICON Express2 (CPHID type FC), including Channel-To-Channel (CTC)
- Support for FICON Express2 (CHPID type FCP) for native support of SCSI disks

- 32 CPs
- z/Architecture 64-bit addressing
- 256 GB of central storage
- 128 GB of expanded storage
- 256 channel paths
- Coupling facility level 15
- Crypto Express2 and Crypto Express2-1P on the BC models (guest use)
- Peer mode channels (ICP, CFP, CBP)
- Guest Coupling simulation
- Dynamic I/O configuration support through the CP configurability function
- Subspace group facility (guest use)
- Able to use IFL engines (for Linux workloads)
- Performance assist via pass-through of adapter I/O operations and interruptions for FCP, IQD, and OSD CHPID types
- zIIP and zAAP Simulation on CPs
- Support for zIIP
- Support for zAAP

z/VM Version 5 Release 2 Support

- CHPID type OSA performance enhancements, for z/OS guests
- CHPID type FCP performance enhancements
- FICON Express4-2C SX on the z9 BC
- Hardware Decimal Floating Point facilities
- Support for z9 BC and z9 EC
- System-managed Coupling Facility structure duplexing, for z/OS guests
- Support for FICON Express2 (CPHID type FC), including Channel-To-Channel (CTC)
- Support for FICON Express2 (CHPID type FCP) for native support of SCSI disks
- 24 CPs
- z/Architecture 64-bit addressing
- 128 GB of central storage
- 128 GB of expanded storage
- 256 channel paths
- Coupling facility level 15
- Crypto Express2 and Crypto Express2-1P on the BC models (guest use)
- Peer mode channels (ICP, CFP, CBP)
- Guest Coupling simulation
- Dynamic I/O configuration support through the CP configurability function
- Subspace group facility (supported for guest use)
- Able to use IFL engines (for Linux workloads)
- Performance assist via pass-through of adapter I/O operations and interruptions for FCP, IQD, and OSD CHPID types

z/VM Version 5 Release 1 Support

- System-managed Coupling Facility structure duplexing, for z/OS guests
- CHPID type OSA performance enhancements, for z/OS guests
- CHPID type FCP performance enhancements

- FICON Express4-2C SX on the z9 BC
- Support for z9 BC and z9 EC (with applicable PTFs)
- Support for FICON Express2 (CHPID type FC), including Channel-To-Channel (CTC)
- Support for FICON Express2 (CHPID type FCP) for native support of SCSI disks
- ESA/390 mode
- z/Architecture 64-bit addressing
- 24 CPs
- 128 GB of central storage
- 128 GB of expanded storage
- 256 channel paths
- Crypto Express2 and Crypto Express2-1P on the BC models (guest use)
- Dynamic I/O configuration support through the CP configurability function
- Subspace group facility (supported for guest use only)
- Guest coupling simulation
- Able to use IFL engines (for Linux workloads)
- Performance assist via passthrough of adapter interruptions for FCP, IQD, and OSD CHPID types

z/VSE

z/VSE Version 4 Release 1 Support

- z/Architecture mode and 64-bit real addressing
- 8 GB of processor storage
- MWLC pricing option

z/VSE Version 3 Release 1 Support

- ESA/390 mode
- 2 GB of processor storage
- Up to 10 CPs
- 256 channel paths
- Support for z9 EC and z9 BC
- OSA Express2 family
- HiperSockets
- Support for FICON Express2 / FICON Express4 (CHPID type FC), including Channel-To-Channel (CTC)
- Support for FICON Express2 / FICON Express4 (CHPID type FCP) for native support of SCSI disks
- CHPID type FCP performance enhancements
- CP Assist for Cryptographic Functions (CPACF)
- Configurable Crypto Express2 (zVSE offers support for clear-key SSL transactions only)
- Crypto Express-1P on the z9 BC

TPF (Transaction Processing Facility)

z/TPF Version 1 Release 1 Support

- Support for FICON Express2 (CHPID type FC), including Channel-To-Channel (CTC)
- ESA/390 or ESA/390 TPF mode
- Up to 38 CPs (either shared or dedicated LP)
- 512 GB's of central storage
- Coupling facility level 9 with APAR support
- 256 channels

TPF Version 4 Release 1 Support

- Support for FICON Express2 (CHPID type FC), including Channel-To-Channel (CTC)

- ESA/390 or ESA/390 TPF mode
- 16 CPs (either shared or dedicated LP)
- 2048 MB of central storage
- Coupling facility level 9 with APAR support
- Shared multiple processor LPs with PUT 03
- 256 channels

Linux

Linux for S/390

- Support for FICON Express2 (CHPID type FC), including Channel-To-Channel (CTC)
- Support for FICON Express2 (CHPID type FCP) for support of SCSI disks
- OSA Express2 large send (also called TCP segmentation offload)
- Linux-Only mode
- 16 CPs
- 2048 MB of central storage
- 262144 MB of expanded storage
- 256 channels
- PKCS #11 API support for zSeries PCI Cryptographic Accelerator Feature
- WLM Management of shared logical processors
- Performance assist via passthrough of adapter interruptions for FCP, IQD, and OSD CHPID types
- PCIxCC support
- Support for zSeries Crypto Express2 feature only for SSL functions (available on release 7 only). Support for clear key RSA operations only.

Hardware Configuration Definition (HCD)

You can use HCD's interactive panels to define configuration information both to the CPC and to the operating system.

Note: HCD and Hardware Configuration Manager (HCM) are also available on z/VM.

HCD, running on z/OS allows you to dynamically change the current I/O configuration of the CPC. HCD allows you to dynamically change the current I/O configuration of the operating system and to create a new IOCDs and make it the active IOCDs.

HCD is required to define the I/O configuration to the operating system. HCD is also the recommended way to define hardware configurations.

The HCD component of z/OS allows you to define the hardware and software I/O configuration information necessary for a parallel sysplex solution environment, including the capability to define:

- peer-mode channel paths (CFP, CBP, CIB, and ICP) to connect z/OS systems to coupling facility images, and
- peer-mode channel paths (CFP, CBP, CIB, and ICP) to connect coupling facility images to one another, in support of System-Managed CF Structure Duplexing.

Additionally, HCD allows you to remotely write IOCDs from one support element CPC to another support element CPC as long as both support elements are powered-on, LAN-attached, enabled for remote licensed internal code (LIC) update, and defined to the same Hardware Management Console.

Dynamic I/O configuration does **not** support the following:

- Adding/deleting MIF image IDs from the initial ones in the IOCDs. (However, you can change which LPs go to which MIF image IDs).

When using HCD in z/OS, you can define and control the configuration of the CPC affecting all LPs. Those LPs that run with HCD or z/VM can dynamically change their software configuration definitions. Other LPs may require an IPL in order to use the new configuration.

When you use HCD you must install, in the LP, the appropriate version of IOCP. Throughout the remainder of this publication, all the capabilities or restrictions documented regarding the IOCP program, also apply to definitions entered and controlled through HCD.

For more information about dynamic I/O configuration on z/OS see:

- *z/OS Hardware Configuration Definition Planning*, GA22-7525
- *z/OS Hardware Configuration Definition: User's Guide*, SC33-7988

z/VM Dynamic I/O Configuration

You can dynamically change the current I/O configuration of the CPC. You can also change the current I/O configuration of the operating system and create a new IOCDS and make it the active IOCDS.

Dynamic I/O configuration does **not** support the following:

- Adding/deleting MIF image IDs from the initial ones in the IOCDS. (However, you can change which LPs go to which MIF image IDs).

You can define and control the configuration of the CPC affecting all LPs. Those LPs that run z/VM can dynamically change their software configuration definitions.

Input/Output Configuration Program (IOCP) Support

To perform a power-on reset you must use an LPAR IOCDS. To generate an LPAR IOCDS you need to use the ICP IOCP program.

PTFs for supported IOCP versions must be applied and can be obtained from the IBM Software Support Center. For more information on ICP IOCP, see *Input/Output Configuration Program User's Guide*, SB10-7037.

Hardware Support

LPs operate independently but can share access to I/O devices and CPC resources. Each active LP must have sufficient channel paths and storage to meet the particular requirements of that LP. Additional central storage, expanded storage, channel paths, consoles, and other I/O devices might be necessary for the planned configuration of LPs.

Operator Training

A general knowledge of z/Architecture and ESA/390 is useful and, in some cases, required of all technical support personnel, LPAR planners, and operators.

Generally, the operator performs the following tasks:

- Edit activation profiles
 - Reset profiles
 - Select an IOCDS
 - Optionally specify LP activation sequence
 - Image profiles
 - Define LP characteristics
 - Optional automatic load specification
 - Load profiles
- Performing a CPC activation
- Activating an LP
- Performing a load on an LP or activating a load profile
- Deactivating a CPC

- Deactivating an LP

Logical Partitions

This section provides an overview of LP characteristics. Some of the characteristics described in this section are model-dependent, EC-level dependent, MCL-dependent, LP mode dependent, or control-program dependent. For this reason, all of the characteristics described here are not necessarily available on all CPCs.

The resources of a CPC can be distributed among multiple control programs that can run on the same CPC simultaneously. Each control program has the use of resources defined to the logical partition in which it runs.

You can define an LP to include:

- One or more CPs
- Central storage
- Channel paths
- Optional expanded storage
- Two or more optional Crypto Express2 Coprocessors or Crypto Express2 Accelerators

An LP can be defined to include CPs, zAAPs, and zIIPs (in combination, so long as one CP is defined) or IFLs.

You can also define an LP to be a coupling facility running the coupling facility control code.

Characteristics

LPs can have the following characteristics. For more information or details about exceptions to any of the characteristics described below, see “Determining the Characteristics” on page 3-5).

- The maximum number of LPs you can define on a z9 EC is 60. The z9 BC S07 supports 30 LPs. The BC R07 only supports the activation of 15 LPs but you can power-on reset the R07 with an IOCDS containing up to 30 LPs. You cannot define an IOCDS with more than 30 LPs for a z9 BC.
- LPs can operate in ESA/390, ESA/390 TPF, Linux-Only, or coupling facility mode.
- The storage for each LP is isolated. Central storage and expanded storage cannot be shared by LPs.
- Using dynamic storage reconfiguration, an LP can release storage or attach storage to its configuration that is released by another LP.
- All channel paths can be defined as reconfigurable. Channel paths are assigned to LPs. You can move reconfigurable channel paths between LPs using tasks available from either the Hardware Management Console or the support element console. If the control program running in the LP supports physical channel path reconfiguration, channel paths can be moved among LPs by control program commands without disruption to the activity of the control program.
- MIF allows channel paths to be shared by two or more LPs at the same time. Only CNC, CTC, CBP, CIB, CFP, ICP, OSC, OSD, OSE, OSN, FC, FCV, FCP, and IQD channel paths can be shared.
- LPs can be defined to have as many as 54 logical CPs. CPs can be dedicated to LPs or shared by them. CPs that you define as dedicated to an LP are not available to perform work for other active LPs. The resources of shared CPs are allocated to active LPs as needed. You can cap (limit) CP resources, if required.

You cannot define a mix of shared and dedicated CPs for a single LP (except for a coupling facility LP using the internal coupling facility feature. See “Internal Coupling Facility (ICF)” on page 2-46). CPs for an LP are either all dedicated or all shared. However, you can define a mix of LPs with shared CPs and LPs with dedicated CPs and activate them concurrently.

For security purposes, you can:

- Reserve reconfigurable channel paths for the exclusive use of an LP (unless overridden by the operator)
- Limit the authority of an LP to read or write any IOCDS in the configuration and limit the authority of an LP to change the I/O configuration dynamically
- Limit the authority of an LP to retrieve global performance data for all LPs in the configuration
- Limit the authority of an LP to issue certain control program instructions that affect other LPs

Figure 1-1 shows some of the characteristics that can be defined for an LP. You can view each LP as a CPC operating within the physical CPC.




Figure 1-1. Characteristics of Logical Partitions

Potential Applications

The use of LPs allows multiple systems, including the I/O for the systems, to be migrated to a single CPC while maintaining the I/O performance, recovery, and multipathing capability of each system, and with minimum impact to the system generation procedures.

LPs are suitable for consideration in the following environments:

Consolidation

Multiple production system images can be consolidated onto 1 CPC without having to merge them into one image.

Migration

Control programs or applications can be migrated by running the old and new systems or applications in independent LPs that are active on the same CPC at the same time.

Production and Test

Multiple production and test systems can run on the same CPC at the same time.

Coupling Facility

A coupling facility enables high performance, high integrity data sharing for those CPCs attached to it and configured in a sysplex.

Coupled Systems

Multiple instances of the same workload can be run in multiple LPs on one or more CPCs as part of a sysplex configuration that takes advantage of the centrally accessible, high performance data sharing function provided by the coupling facility.

Extended Recovery Facility (XRF)

Primary and alternate XRF systems can run on 1 CPC. Multiple and alternate XRF systems can run on 1 CPC.

Communications Management Configuration (CMC)

The communications management configuration (CMC) machine, usually run on a separate CPC, can be run as an LP on the same CPC.

Departmental Systems

Multiple applications can be isolated from one another by running each in a separate LP.

Constrained Systems

Those systems that cannot fully use a large system because of storage constraints can alleviate the problem by using LPs to define multiple system images on the same CPC.

Diverse Workloads

Interactive workloads such as the Customer Information Control System (CICS®) and time-sharing option (TSO) can be isolated by running each in a separate LP.

Examples of Logical Partition Applications

Figure 1-2 and Figure 1-3 show examples of how LPs on a CPC can be used to support multiple systems.

Figure 1-2 represents the migration of two z/OS, one z/VM and one stand-alone coupling facility systems onto a single CPC. These systems, which were running on four separate IBM 9672 CPCs, now operate as four LPs on a single z9 EC. Two LPs are each running a z/OS system in ESA/390 mode, one is running a z/VM system in ESA/390 mode, and one is running in coupling facility mode.




Figure 1-2. Migration of Four Production Systems to LPs

Figure 1-3 represents a CPC with three LPs: an active XRF system, an alternate XRF system, and an alternate for another primary XRF system.




Figure 1-3. Support for Three XRF Systems

Compatibility and Migration Considerations

This section provides migration and compatibility information for the System z9.

Device Numbers

When multiple systems are migrated to a System z9 CPC, the combination of systems could include different devices or shared devices with identical device numbers. Each system can operate in an LP without changing the device numbers as long as identical device numbers do not occur in the same LP. However, duplicate device numbers can exist in the same LP if these device numbers are in different subchannel sets.

Duplicate device number conflicts can occur when the I/O configuration is reconfigured. For example, if a reconfigurable channel path is reassigned to another LP and devices attached to the channel path have device numbers that are already assigned in the receiving LP to other online channel paths, a conflict results. When IOCP generates an LPAR IOCDS, the initial configuration contains no duplicate device number conflicts in an LP.

Device number conflicts are also detected when operator tasks change the I/O configuration (channel path tasks from the Hardware Management Console or support element console; or control program configuration command) or during LP activation.

Duplicate device number conflicts are also detected when a dynamic I/O configuration change is made.

Multiple Subchannel Sets (MSS)

A new Multiple Subchannel Sets (MSS) structure for z9 EC allows increased device connectivity for Parallel Access Volumes (PAVs). 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, exclusive to the System z9, is supported by ESCON (CHPID type CNC), FICON when configured as CHPID type FC, and z/OS.

Control Programs

PTFs for supported control programs must be applied and can be obtained from the IBM Software Support Center. A supported control program operates in an LP as it does in one of the basic modes, with the following exceptions:

z/OS

- Physical reconfiguration, either offline or online, of CPs is not supported on the System z9. Logical CP reconfiguration, either offline or online, is supported in an LP. This does not affect the online/offline status of the physical CPs. To reconfigure a logical CP offline or online, use the following z/OS operator command:

```
CF CPU(x),<OFFLINE/ONLINE>
```

- Physical reconfiguration, either offline or online, of central and expanded storage is supported.

To reconfigure a central storage element offline or online, use the following z/OS operator command:

```
CF STOR(E=1),<OFFLINE/ONLINE>
```

Additionally you can use the following command to reconfigure smaller amounts of central storage online or offline:

```
CF STOR(nnM),<OFFLINE/ONLINE>
```

To reconfigure an expanded storage element offline or online, use the following z/OS operator command:

```
CF ESTOR(E=x),<OFFLINE/ONLINE>
```

Reconfigurable Storage Unit (RSU) Considerations: The RSU parameter should be set to the same value that you specified in the central storage Reserved field. See *z/OS MVS Initialization and Tuning Reference* for appropriate RSU parameter syntax.

- Reconfiguration, either offline or online, of channel paths by z/OS operator commands is supported on a System z9. This capability also allows channel paths to be moved among LPs using z/OS operator commands.
- Preferred paths to a device are supported on a System z9. If the preferred path parameter is specified in an LPAR IOCDS, it is accepted.
- Specifying SHAREDUP for devices is not recommended. If used, z/OS treats the device as a SHARED device.
- Each z/OS LP can run the Resource Measurement Facility (RMF). RMF enhancements for PR/SM allow a single LP to record system activities and report them in the Partition Data Report. To enable this, use the Change LPAR Security task available from the CPC Operational Customization task list and select Performance Data Control for the LP.

For z/OS, RMF reporting includes LPAR Management Time.

- RMF provides enhanced reporting for coupling facility configurations.
- RMF with APAR support identifies which logical and physical CPs are of each type when any combination of general purpose, IFA, IFL, zIIP and ICF processors are present in the configuration on its partition data report. You can have up to 32 virtual machines running z/OS within a simulated sysplex, with each z/OS virtual machine coupled to up to 8 virtual machines running as coupling facilities

EREP

Each control program operating in an LP has its own environmental recording, editing, and printing (EREP) processing. EREP records for ICP channel paths go to the z/OS logs of the z/OS systems attached to a coupling facility LP.

CPU IDs and CPU Addresses

Application packages and software products that are licensed to run under specific CPU identification (CPU ID) information should be checked because they may need to be updated.


CPU ID information is system-generated for each logical CP in the LP during the LP activation. It consists of a version code for the CPC machine type, a CPU identification number that is unique for each logical partition, a model number for the CPC machine type, and a value of X'8000'.

The Store CPU ID (STIDP) instruction stores the CPU ID for each logical CP in storage in the following format (Figure 1-4):

0	7 8	31 32	47 48	63
Version Code	CPU Identification Number	Machine Type	8000	

Figure 1-4. CPU ID Format

Figure 1-5 shows the format of the CPU identification number (bits 8 through 31 of the CPU ID format).



Legend :

- P Logical partition identifier
- n Digit derived from the serial number

Figure 1-5. CPU Identification Number Format

CPU ID Fields

The CPU identification number, with the version code and the machine type, permits a unique CPU ID for each logical partition.

- The **version code** for the System z9 is always zero and is not affected by the operating mode.
- The **CPU identification number** for each logical CP (see Figure 1-5) consists of a two-digit LP identifier, and digits derived from the serial number of the CPC.
 - The **logical partition identifier** is specified using the Partition identifier field on the General page in either the reset or image profile used by the LP and must be unique for each active LP in the configuration.
- The following machine types (CPC model numbers) are returned as indicated below:

Machine Type

2094

2096

Models

System z9 EC

(S08, S18, S28, S38, S54)

System z9 BC

(R07, S07)

Note: STIDP is provided for purposes of backward compatibility. IBM highly recommends that the Store System Information instruction (STSI) be used rather than STIDP. STSI is the preferred means to obtain all CPU information including machine serial number. When a unique logical CPU address is all that is required, use the Store CPU Address (STAP) instruction.

Examples of CPU ID Information

The following examples show the format and contents of the CPU ID information stored by the STIDP instruction for logical CPs in active LPs.

Table 1-1 shows the CPU ID information for a z9 EC with 3 active LPs.

Table 1-1. CPU IDs for a z9 EC

LP Name	LP Identifier	Number of CPs Defined	CPU ID Returned by STIDP
ZVSE	1	1	00 019999 2094 8000
ZOSTEST	2	1	00 029999 2094 8000
ZOSPROD	3	8	00 039999 2094 8000 00 039999 2094 8000 00 039999 2094 8000 00 039999 2094 8000 00 039999 2094 8000 00 039999 2094 8000 00 039999 2094 8000 00 039999 2094 8000

HSA Allocation

This section addresses hardware system area (HSA) allocation.

HSA Estimation Tool

The HSA Estimation Tool, available on Resource Link at <http://www.ibm.com/servers/resource link>, estimates the HSA size for the specified configuration and product.

The HSA allocation can vary according to the I/O configuration. The following system I/O configurations require the given HSA size on a z9 EC.

- **Small system I/O configuration (S08): 1344 MB**
 - 1 CSS
 - 2 LPs
 - 96 physical control units
 - 4096 devices
 - 64 channels defined in the IOCDS
 - Dynamic I/O configuration not enabled
 - Concurrent code change not authorized
- **Large system I/O configuration (S38): 4288 MB**
 - 4 CSSs
 - 60 LPs
 - 65280 devices per LCSS for subchannel set 0
 - 65535 devices per LCSS for subchannel set 1
 - Dynamic I/O configuration enabled
 - Concurrent code change authorized

Intermediate HSA sizes exist for configuration sizes in between the given small and large I/O configuration examples.

The HSA range for the z9 BC is approximately 1-2 GB.

TOD Clock Processing

The CPC TOD clocks of all the CPs are automatically set during CPC activation. The time reference used depends on whether or not a Sysplex Timer[®] is attached to the CPC or Server Time Protocol (STP) is enabled. When STP is enabled, a CPC can participate in: an ETR network, a Mixed Coordinated Timing Network (CTN), or an STP-only CTN. In an ETR network or Mixed CTN, the Sysplex Timer provides the timekeeping information. In an STP-only CTN, the Sysplex Timer is no longer required. In this case the Current Time Server for the STP-only CTN provides the time information.

No Sysplex Timer Attached and Server Time Protocol Not Enabled

During PR/SM initialization, the CPC TOD clocks for each CP are set to the TOD value of the support element. Each LP starts out with this CPC TOD value at the completion of LP activation. The operating system running in an LP can set a TOD value for itself and this will be the only TOD reference it will see. Setting the TOD clock for one logical CP in the LP sets the TOD clock for all logical CPs in that LP, but does not affect the logical CPs in any other LP. The TOD clock value is used for the duration of the LP activation, or until a subsequent Set Clock instruction is issued in the LP.

Sysplex Timer Attached

The attachment and use of an IBM Sysplex Timer is supported. Also, during PR/SM initialization, when a Sysplex Timer is attached, the CPC TOD clocks for each CP are set to the TOD value from the Sysplex Timer.

The operating system in each LP can independently choose whether or not to synchronize to the Sysplex timer if one is present. Operating systems in LPs that do synchronize to the Sysplex Timer will all be running with identical TOD values. Operating systems in LPs that do not synchronize to the Sysplex Timer do not need to be aware of the presence of a Sysplex Timer and can set their TOD values independently of all other LPs.

Note that z/OS does not allow you to change the value of the TOD setting when it is using the Sysplex Timer (ETRMODE=YES in the CLOCKxx parmlib member).

The System z9 models support the specification of a logical partition time offset. When all members of a sysplex are in logical partitions on these supported models, the Logical partition time offset can be used for:

- Different local time zone support in multiple sysplexes using the same sysplex timer. Many sysplexes have the requirement to run with a LOCAL=GMT setting in a sysplex (ETRMODE=YES) where the time returned from a Store Clock instruction yields local time. To do this, the time returned by the sysplex timer must be local time. Prior to Logical partition time offset support, this could only be accomplished for one sysplex for a given sysplex timer. With Logical partition time offset support, multiple sysplexes can each have their own local time reported to them from the sysplex timer if desired. For instance, the sysplex timer can be set to GMT, one set of sysplex partitions could specify a Logical partition time offset of minus 5 hours, and a second set of sysplex partitions could specify a Logical partition time offset of minus 6 hours.

Server Time Protocol Enabled

The enablement of STP is supported. Also, during PR/SM initialization, when STP is enabled, the CPC TOD clocks for each CP are set to the TOD value from STP.

The operating system in each LP can independently choose whether or not to synchronize to the current time source for STP, if present. Operating systems in LPs that do synchronize to STP will all be running with identical TOD values. Operating systems in LPs that do not synchronize to STP do not need to be aware of the presence of STP and can set their TOD values independently of all other LPs.

Note that z/OS does not allow you to change the value of the TOD setting when synchronized to STP (STPMODE=YES in the CLOCKxx parmlib member).

The System z9 models support the specification of a logical partition time offset. When all members of a sysplex are in logical partitions on these supported models, the Logical partition time offset can be used for:

- Different local time zone support in multiple sysplexes using the same CTN. Many sysplexes have the requirement to run with a LOCAL=GMT setting in a sysplex (STPMODE=YES) where the time returned from a Store Clock instruction yields local time. To do this, the time returned by STP must be local time. With Logical partition time offset support, multiple sysplexes can each have their own local time reported to them from STP if desired. For instance, STP can be set to GMT, one set of sysplex partitions could specify a Logical partition time offset of minus 5 hours, and a second set of sysplex partitions could specify a Logical partition time offset of minus 6 hours.

Sysplex Testing Without a Sysplex Timer and Server Time Protocol Not Enabled

You can do sysplex testing without a Sysplex Timer or Server Time Protocol enabled by setting up a test sysplex of several z/OS LPs in multiple LPs in the same LPAR configuration. Use the SIMETRID keyword in the CLOCKxx parmlib member for z/OS to synchronize the members of the sysplex in the LPs.

Synchronized Time Source and the Coupling Facility

Improved processor and coupling facility link technologies inherent on System z9 models necessitate more rigorous time synchronization tolerance for members of a parallel sysplex hosted by those models. To help ensure that any exchanges of time-stamped information between members of a sysplex observe the correct time ordering, time stamps are now included in the message-transfer protocol between the systems and the coupling facility.

Consequently, a coupling facility hosted by any System z9 model requires connectivity to the same synchronized time source as the other z/OS systems in its parallel sysplex. If a member of its parallel sysplex is on the same server as the coupling facility, required connectivity is already provided to the synchronized time source. However, when a coupling facility is a resident of a System z9 model, which **does not** include a member of the coupling facilities parallel sysplex, connectivity attached to the synchronized time source must be implemented.

Extended TOD-Clock Facility

The extended TOD-clock facility provides an extended form TOD clock and a TOD programmable register. The extended form TOD clock is a 128-bit value that extends the current basic form by appending 8 bits on the left and 56 bits on the right. The extended form TOD clock is returned by a new problem-program instruction, STORE CLOCK EXTENDED (STCKE). The contents of the TOD programmable register are stored into the rightmost portion of the extended form TOD value when the TOD clock is inspected by STCKE. A TOD programmable

register exists for each CPU and contains the TOD programmable field in bits 16-31. The TOD programmable register is set by a new privileged instruction, SET TOD PROGRAMMABLE FIELD (SCKPF). The leftmost byte of the extended form TOD clock is the TOD Epoch Index (TEX), and is stored as zeros in machines running ESA/390.

The extended TOD clock facility satisfies three main objectives:

- Relieve constraints that exist in the current 64-bit TOD clock
- Extend the TOD-clock architecture to multi-system configurations
- Help ensure sysplex-wide uniqueness of the STCKE TOD values

The TOD Programmable Field (TODPF) is a 16-bit quantity contained in bit positions 16-31 of the TOD programmable register. The contents of the register can be set by the privileged instruction SET TOD PROGRAMMABLE FIELD. The contents of the register can be stored by the instruction STORE CLOCK EXTENDED, which stores the TOD programmable field in the last 16 bits of the extended form TOD clock. The contents of the register are reset to a value of all zeros by an initial CPU reset.

Chapter 2. Planning Considerations

Overview	2-3
Planning the I/O Configuration	2-3
Planning Considerations	2-3
Control Program Support	2-3
Hardware Configuration Definition (HCD) Support	2-3
z/VM Dynamic I/O Configuration Support	2-4
Input/Output Configuration Program (IOCP) Support	2-4
Characteristics of an IOCDs	2-4
Maximum Number of Logical Partitions	2-5
Determining the Size of the I/O Configuration	2-5
Maximum Size of the I/O Configuration	2-5
Guidelines for Setting Up the I/O Configuration	2-5
Recovery Considerations	2-6
Managing Logical Paths for ESCON and FICON Channels	2-7
Definition	2-7
Control Unit Allocation of Logical Paths	2-7
Why Manage Logical Paths?	2-7
Managing the Establishment of Logical Paths	2-11
Logical Path Considerations	2-11
Recommendations	2-12
Shared Channel Overview	2-21
MIF Requirements	2-21
Understanding ESCON and MIF Topologies	2-21
MIF Performance Planning Considerations	2-22
Unshared ESCON or FICON Channel Recommendations	2-27
Dynamically Managed CHPIDs	2-27
IOCP Coding Specifications	2-28
IOCP Statements for ICP	2-28
Shared Devices Using Shared Channels	2-31
Shared Devices Using Unshared Channels	2-31
Duplicate Device Numbers for Different Physical Devices	2-33
Coupling Facility Planning Considerations	2-39
Test or Migration Coupling Configuration	2-39
CFCC Enhanced Patch Apply	2-39
Production Coupling Facility Configuration	2-40
Production Coupling Facility Configuration for Full Data Sharing	2-40
Production Coupling Facility Configuration for Resource Sharing	2-40
Internal Coupling Facility (ICF)	2-41
Dynamic Coupling Facility Dispatching (DCFD)	2-42
Dynamic Internal Coupling Facility (ICF) Expansion	2-42
Enhanced Dynamic ICF Expansion Across ICFs	2-43
System-Managed Coupling Facility Structure Duplexing	2-44
Single CPC Software Availability Sysplex	2-44
Coupling Facility Nonvolatility	2-45
Nonvolatility Choices	2-45
Setting the Conditions for Monitoring Coupling Facility Nonvolatility Status	2-45
Coupling Facility Mode Setting	2-45
Coupling Facility LP Definition Considerations	2-46
Internal Coupling Facility (ICF)	2-46
Coupling Facility LP Storage Planning Considerations	2-47
Structures, Dump Space, and Coupling Facility LP Storage	2-47
Estimating Coupling Facility Structure Sizes	2-48

Dump Space Allocation in a Coupling Facility	2-48
Coupling Facility LP Activation Considerations.	2-49
Coupling Facility Shutdown Considerations	2-49
Coupling Facility LP Operation Considerations	2-49
Coupling Facility Control Code Commands	2-50
Coupling Facility Level (CFLEVEL) Considerations	2-50
CPC Support for Coupling Facility Code Levels	2-50
Level 15 Coupling Facility	2-52
Level 14 Coupling Facility	2-52
Level 13 Coupling Facility	2-52
Level 12 Coupling Facility	2-52
Level 11 Coupling Facility	2-53
Level 10 Coupling Facility	2-53
Level 9 Coupling Facility	2-53
Level 8 Coupling Facility	2-53
Coupling Facility Resource Management (CFRM) Policy Considerations	2-54
Coupling Facility Channels	2-54
Internal Coupling Channel	2-55
Integrated Cluster Bus Channels	2-56
High-Performance Coupling Facility Channels (HiPerLinks)	2-56
InfiniBand host channel adapter (HCA)	2-56
InfiniBand coupling links for Parallel Sysplex	2-56
Coupling Facility Channels (TYPE=CFP, TYPE=CBP, TYPE=CIB, or TYPE=ICP)	2-57
Defining Internal Coupling Channels (TYPE=ICP)	2-58
I/O Configuration Considerations	2-58
Considerations when Migrating from ICMF to ICs	2-59
Linux Operating System Planning Considerations	2-59
Integrated Facility for Linux (IFL)	2-59
z/VM Version 5 Utilizing IFL Features	2-60
System z9 Application Assist Processor (zAAP)	2-60
IBM System z9 Integrated Information Processor (zIIP)	2-61
Concurrent Patch	2-61
CFCC Enhanced Patch Apply.	2-62
Dynamic Capacity Upgrade on Demand	2-62
PR/SM Shared Partitions	2-63
Mixed Shared and Dedicated PR/SM Partitions	2-63
Multiple Dedicated PR/SM Partitions	2-64
Shared Internal Coupling Facility	2-64
Dynamic Capacity Upgrade on Demand Limitations.	2-65
Concurrent Memory Upgrade	2-66
Capacity Backup Upgrade (CBU) Capability	2-66
Enhanced Book Availability.	2-67
Preparing for Enhanced Book Availability	2-67
Getting the System Ready to Perform Enhanced Book Availability	2-67
Reassigning Non-Dedicated Processors	2-69
Customer Initiated Upgrade (CIU)	2-70
Concurrent Processor Unit Conversion	2-70
Planning for Hot Plugging Crypto Features	2-71

Overview

This chapter describes planning considerations for I/O configuration and for coupling facility LPs.

Planning the I/O Configuration

This section describes the planning considerations and guidelines for creating an IOCDS. It assumes you understand the IOCP configuration and coding requirements described in the *Input/Output Configuration Program User's Guide*, SB10-7037.

Planning Considerations

PR/SM is standard on all System z9 models.

Control Program Support

The maximum number of supported devices is limited by the control program. In planning an I/O configuration, installation management should be aware of the maximum number of devices supported by the control program run in each LP. Please see the user's guide for the respective operating systems.

Hardware Configuration Definition (HCD) Support

HCD supports definition of the I/O configuration for an entire installation. It is required for parallel sysplex and LPAR clusters. A single I/O data file is created for the installation and used for multiple machines and I/O configuration data sets.

HCD supports:

- Up to 60 logical partitions (LPs) per central processing complex (CPC)
- Coupling facility configurations
- Multiple Image Facility (MIF)
- Dynamic CHPID Management (DCM) channel paths.
- MCSS
- Assigning reserved logical partitions a meaningful name

Table 2-1. HCD Function Support

HCD Function	z/OS	z/VM 5.3
Define 60 logical partitions?	Yes	Yes
Define shared channel paths?	Yes	Yes
Define coupling facility channel paths?	Yes	Yes (Note 2)
Define dynamically managed channel paths?	Yes	Yes (Note 2)
Write IOCDs remotely?	Yes (Note 1)	No
Access I/O devices on shared channel paths?	Yes	Yes
Use software-only dynamic I/O?	Yes	Yes
Notes: 1. HCD, running on z/OS allows you to remotely write IOCDs from one CPC to another CPC that is powered on, LAN-attached, enabled for remote LIC update, and defined to the same Hardware Management Console. 2. HCD, running on z/VM, allows you to define coupling facility channel paths or managed channel paths for a z/OS LP but z/VM does not support coupling facility channel paths or dynamically managed channel paths for use by z/VM or guest operating systems.		

For more information on using HCD with Multiple Image Facility, see

- *z/OS Hardware Configuration Definition: User's Guide*, SC33-7988
- *z/OS Hardware Configuration Definition Planning*, GA22-7525
- *z/VM I/O Configuration*, SC24-6100.

z/VM Dynamic I/O Configuration Support

z/VM Support for the Coupling Facility: z/VM allows you to define configurations that use the coupling facility. However, z/VM does **not** support the coupling facility itself. Instead, z/VM's dynamic I/O configuration capability allows you to define resources that can be used by a z/OS system in another LP. For a summary of z/VM's support of dynamic I/O configuration, see Table 2-2.

z/VM Support for the Multiple Image Facility (MIF): You can use z/VM to define shared channel paths. For a summary of z/VM support of dynamic I/O configuration, see Table 2-2.

Table 2-2. z/VM Dynamic I/O Support for MIF and the Coupling Facility

z/VM Function	Release
	5.3, 5.2, 5.1
Define shared channel paths?	Yes
Define coupling facility channel paths?	Yes (Note)
Write IOCDs remotely?	No
Access I/O devices on shared channel paths?	Yes
Use software-only dynamic I/O?	Yes
Use hardware and software dynamic I/O?	Yes
Use shared ESCON CTCs?	Yes
Note: z/VM can define coupling facility channel paths for a z/OS LP but does not support real coupling facility channel paths for use by z/VM or guest operating systems.	

Input/Output Configuration Program (IOCP) Support

You can create up to four IOCDs. ICP IOCP is the required supported version. You can define as many as 60 LPs. For more information on ICP IOCP, see *Input/Output Configuration Program User's Guide*, SB10-7037.

Characteristics of an IOCDs

The definitions for channel paths, control units, and I/O devices are processed by the IOCP and stored in an IOCD. During initialization of the CPC, the definitions of a selected IOCDs are transferred to the hardware system area (HSA). The IOCDs is used to define the I/O configuration data required by the CPC to control I/O requests.

Channel paths in an IOCDs are assigned to one or more LPs. The characteristics of an IOCDs are:

- Using the IOCP RESOURCE statement, you define logical channel subsystems (CSSs) and the logical partitions that have access to the channel paths in a CSS.
- Using the IOCP RESOURCE statement, you can name logical partitions and assign MIF image ID numbers to them. MIF image ID numbers are necessary for ESCON CTC and FICON CTC definitions. You can also reserve unnamed logical partitions associated with specific MIF image ID numbers that can later be dynamically named.

- Using the IOCP CHPID statement, you can assign a channel path as reconfigurable or shared.
- Using the IOCP CHPID statement, you can specify a Dynamic CHPID Management (DCM) channel path and the cluster to which the CHPID belongs. The CHPID is shareable among active LPs that have become members of the specified cluster.
- You can duplicate device numbers within a single IOCP input file, but the device numbers cannot be duplicated within an LP. See “IOCP Coding Specifications” on page 2-28.

Maximum Number of Logical Partitions

The maximum number of LPs you can define on a z9 EC is 60. The z9 BC S07 supports 30 LPs. The BC R07 only supports the activation of 15 LPs but you can power-on reset the R07 with an IOCDS containing up to 30 LPs. You cannot define an IOCDS with more than 30 LPs for a z9 BC. See “Determining the Size of the I/O Configuration” before making decisions about the number of LPs to define.

Determining the Size of the I/O Configuration

To determine the size of the current I/O configuration (number of control unit headers and devices), review the IOCDS Totals Report for the current IOCDS.

Maximum Size of the I/O Configuration

Limits within an I/O configuration exist for the following:

- Devices
- Control unit headers
- Physical control units

System z9 Models:

- The maximum number of control unit headers (CUHs) is 4096 per logical channel subsystem (CSS).
- The maximum number of physical control units is 8192.
- The maximum number of devices is 65280 per CSS for subchannel set 0.
- The maximum number of devices is 65535 per CSS for subchannel set 1.

I/O Configuration Size and Central Storage Resources: Central storage reserved for the HSA cannot be used by LPs. Increases in the size of the I/O configuration can decrease the amount of central storage available for use by LPs.

After creating an IOCDS that increases the size of the I/O configuration, make sure that all of the planned combinations of LPs can be activated.

To determine the amount of central storage available for assignment to LPs, refer to the Available Storage field under Central Storage Allocation in the Storage Information task (see Figure 5-1 on page 5-2). If the amount of central storage needed to activate the LPs is more than the amount of storage available, change the necessary LP definitions in the image profiles.

Guidelines for Setting Up the I/O Configuration

Follow these guidelines when setting up an I/O configuration.

1. Determine the number of LPs and in which logical channel subsystem (CSS) they should exist.
2. For dynamic I/O configurations, include any logical partitions for which you do not yet have a meaningful name. These logical partitions will be reserved until a subsequent dynamic I/O configuration change is made to assign them a name.

3. Determine if you want to move any channel paths among LPs. If you do, then these channel paths must be defined as reconfigurable in the IOCP CHPID statement. You cannot move a channel path from an LP in one CSS to an LP in another CSS.
4. Determine if you want to share any channel paths among LPs in the same CSS. If you do, then specify these channel paths as SHARED in the IOCP CHPID statement. Doing this, helps reduce the number of channel paths configured to a physical control unit and device. Make sure that the channel path type supports being shared.
5. Determine if you want to share any channel paths among LPs in different CSSs. If you do, then define these channel paths as spanned by specifying multiple CSS IDs in the PATH keyword of the IOCP CHPID statement. Doing this further helps reduce the number of channel paths configured to a physical control unit and device. Make sure that the channel path type supports being spanned.
6. Within each LP, configure primary and backup paths from separate channel adapter cards.
7. Within each LP, configure primary and backup paths from separate self-timed interfaces (STIs).

Recovery Considerations

When planning for recovery, consider the following I/O configuration guidelines.

- Assign channel paths to LPs as described in “Guidelines for Setting Up the I/O Configuration” on page 2-5.
- Review the recoverability characteristics of the I/O configuration described in the section “Shared Devices” on page 2-31.

Managing Logical Paths for ESCON and FICON Channels

This section describes logical paths, explains overall system considerations, and makes specific configuration recommendations.

Definition

A logical path is a logical connection between a control unit and an ESCON channel (TYPE=CNC or TYPE=CTC), a FICON channel (TYPE=FCV) attached to an ESCON Director, or a FICON channel (TYPE=FC). (FICON Express2 and FICON Express4 features do not support TYPE=FCV channels.) Logical paths are important because each sharing LP on a CPC requires that a logical path be established between an ESCON or FICON channel and a control unit for I/O operations to occur.

Logical paths do **not** exist for ESCON channels attached to a 9034 ESCON Converter Model 1 (TYPE=CBY or TYPE=CVC).

Logical paths do **not** exist for coupling facility channel paths (TYPE=CFP), integrated cluster bus channel paths (TYPE=CBP), InfiniBand internal coupling channel paths (TYPE=ICP), Open Systems Adapter channel paths (TYPE=OSC, TYPE=OSD or TYPE=OSE, TYPE=OSN), internal queued direct communication (HiperSockets) channel paths (TYPE=IQD), or fibre channel protocol channel paths (TYPE=FCP).

Control Unit Allocation of Logical Paths

Control units allocate logical paths to channels dynamically on a first-come-first-served basis. Control units do not manage the allocation of logical paths but instead allow channels to compete for logical paths until all of the control unit's logical paths are used.

Why Manage Logical Paths?

The ESCON and FICON environments (the use of ESCON channels, FICON Express channels, FICON Express2 and FICON Express4 channels, and ESCON and FICON Directors) greatly enhances the connectivity potential for control units. In addition, you can define shared channels that can request additional logical paths. However, control units can only allocate a limited number of logical paths (see Table 2-3 on page 2-8) in relation to the number of logical paths that channels can request. In configurations where channels request more logical paths than a control unit can allocate, you must manage logical paths to help ensure that the I/O operations you intend take place.

The FICON Express2 and FICON Express4 SX and LX features support all of the functions of FICON Express, with the exception of support for channel path type FCV. FICON Express2 and FICON Express4, however, offers increased connectivity in the same amount of physical space, and offers the possibility of increased performance. Up to 240 FICON Express2 and FICON Express4 channels on a z9 EC (and 112 channels on a z9 BC) can be employed to greatly expand connectivity and throughput capability. The FICON connectivity solution is based on industry-standard Fibre Channel technology and leverages our exclusive native FICON architecture. For detailed information, see *Input/Output Configuration Program User's Guide*, SB10-7037.

Table 2-3. Logical Path Summary by Control Unit

Control Unit	Maximum Physical Links	Logical Paths per Link	Maximum Logical Paths
2105 Storage Control 2105	32	64	2048
3990 Storage Control 3990-2 3990-3 3990-6	16 16 16	Note 1 Note 1 Note 2	16 16 128
9343 Storage Controller 9343 D04 9343 DC4	4 4	64 64	64 64
3590 Magnetic Tape Subsystem 3590	2	64	128
3490 Magnetic Tape Subsystem 3490 D31, D32, D33, D34 3490 C10, C11, C12 3490 A01 3490 A10 3490 A02 3490 A20	1 1 2 4 4 8	16 16 16 16 16 16	16 16 32 64 64 128
2074 Console Support Controller 2074 002	2	32	64
3172 Interconnect Controller 3172-1	2	Note 3	16
3174 Establishment Controller 3174 12L (SNA only) 3174 22L (SNA only)	1 1	Note 3 Note 3	8 8
Notes: 1. 3990-2 and 3990-3 Storage Controls can have eight logical paths per side, two sides, with a minimum of one logical path per adapter. 2. The 3990-6 Storage Control can have 64 logical paths per side, two sides, with a maximum of 32 logical paths per adapter. 3. The 3172 Interconnect Controller and the 3174 Establishment Controller can have one logical path per control unit header (CUH) with a maximum of eight CUHs per link.			

ESCON Example: Figure 2-1 shows a situation where 20 ESCON channel paths routed through an ESCON Director (ESCD) each attempt to establish a logical path to a 3990 Storage Control. The 3990 will allocate all 16 of its logical paths among the 20 channel paths dynamically on a first-come-first-served basis. Clearly, some of the ESCON channel paths will not operate in this configuration. Managing logical paths can provide a way to alleviate this situation.




Figure 2-1. An ESCON Configuration that Can Benefit from Better Logical Path Management

MIF Example: Figure 2-2 shows an ESCON shared channel configuration on an MIF-capable CPC. In this example, all five LPs share each of four ESCON channels attached to a 3990. Each shared ESCON channel represents five channel images corresponding to the five LPs. Each channel image requests a logical path to the 3990. As in Figure 2-1 on page 2-9, a total of 20 logical paths is requested from the 3990, which can only satisfy 16 of the requests. Again, you can avoid this situation by managing logical paths.




Figure 2-2. A Shared ESCON Configuration that Can Benefit from Better Logical Path Management

Managing the Establishment of Logical Paths

You can manage the establishment of logical paths between channels and control units. With proper planning, you can create I/O configuration definitions that allow control units in the configuration to allocate logical paths for every possible request made by channels in either of the following ways:

- Create a one-to-one correspondence between the logical path capacity of all control units in the physical configuration and the channels attempting to request them.
- Create I/O configurations that can exceed the logical path capacity of all or some of the control units in the physical configuration but, at the same time, provide the capability to selectively establish logical connectivity between control units and channels as needed.

This capability can be useful or even necessary in several configuration scenarios. See “Recommendations” on page 2-12.

Logical Path Considerations

You can better understand how to manage the establishment of logical paths by understanding the following:

- Control unit considerations
- Connectivity considerations
- Channel configuration considerations
- 9035 ESCON Converter considerations

Control Unit Considerations: Consider the following factors concerning the allocation of logical paths by control units:

- Control units allocate logical paths dynamically on a first-come-first-served basis. Control units do not manage the allocation of logical paths but instead allow channels to compete for logical paths until all of the control unit’s logical paths are used.
- Control units vary in the number of logical paths they support. See Table 2-3 on page 2-8.

Connectivity Considerations: ESCON and FICON system hardware, CPCs, and ESCON and FICON Directors significantly affect the volume of logical path requests to a control unit as follows:

- Control units can attach to one or more ports on a Director or to additional ports on other Directors. Each Director port can dynamically connect to many other ports to which channels requesting logical paths are attached.
- For CPCs, each logical partition attaching to the same control unit will compete for the control unit’s logical paths.
- In a configuration where control units are shared by different CPCs, I/O configuration definitions for individual control units are not coordinated automatically among the IOCDs of the different CPCs. Each CPC competes for a control unit’s logical paths.
- Shared channels require the establishment of a logical path for each channel image corresponding to an active LP sharing the channel. This can significantly increase the number of logical paths that a single channel requests.

Channel Configuration Considerations: The following configuration rules determine how logical paths are established for ESCON and FICON channels.

- A channel initially attempts to establish logical paths:

- If you perform POR, only those channels configured online to LPs that are activated at POR will attempt to establish logical paths. Shared channels will attempt to establish logical paths only for those activated LPs with the channel configured online.
- When the LP is activated.
- When configured online (if previously configured offline).
- A channel cannot establish a logical path or will have its logical path removed when:
 - An LP is deactivated. A shared channel will continue to operate for any other remaining activated LPs to which it is defined. Logical paths to those LPs will remain established.
- A shared channel **cannot** establish a logical path to a control unit for an LP that **cannot** access any of the I/O devices on the control unit. In IOCP, the PARTITION or NOTPART keyword on the IODEVICE statement specifies which LPs can access a device.
- A channel that cannot initially establish a logical path can reattempt to establish a logical path if the channel detects or is notified of:
 - A change in the state of a control unit
 - A change in the state of a link or port
 - A dynamic I/O configuration change that frees previously allocated logical paths
- A channel cannot establish a logical path or will have its logical path removed if:
 - The Director that connects the channel to the control unit blocks either the channel port or control unit port used in the path.
 - The Director that connects the channel to the control unit prohibits the dynamic connection or communication between the channel port and the control unit port used in the path.
 - A link involved in the path fails or is disconnected. When a shared channel is affected by a port being blocked, a dynamic connection or communication being prohibited, or a link failing or being disconnected, each LP sharing the channel will be equally affected and all logical paths using the port or link (regardless of which LP they are associated) will be removed.
 - The channel is configured offline. When a shared channel is configured offline for an LP, it will continue to operate for any other LP that has the channel configured online. Logical paths to these other logical partitions will remain established.
 - Power to the channel, control units, or Directors in the configuration is turned off.

9035 ESCON Converter Considerations: For the 3990, 9035 configuration requires special Vital Product Data (VPD) specifications. When an ESCON Director blocks a port supporting a 9035 or prohibits a dynamic connection between the channel attached to the 9035 and the control unit, the logical path is removed from the 3990. This path is now available to other channels, either to other 9035s configured in the VPD or to channels not attached to a 9035, that attempt to establish a logical path. However, the 3990 will always attempt to establish a logical path to the 9035s configured in the VPD before it allows a channel not attached to a 9035 to have a logical path.

Recommendations

Creating I/O configuration definitions where channels could request more logical paths to control units than the control units could support can be useful in the following scenarios:

- **Workload balancing**

When a system image becomes overloaded, you may need to reassign a workload and the necessary logical paths (for example, its tape or DASD volumes, a set of display terminals, or a set of printers) to another system image that has available capacity.

- **Backup**

When an outage occurs, you can move the critical application set (the program and associated data) and the necessary logical paths to a backup or standby CPC. This process is simple if the CPCs have identical I/O configurations.

In I/O configurations where channels can request more logical paths to control units than the control units can support, you can manage how logical paths are established by:

- Deactivating unneeded LPs.
- Configuring offline unneeded channels. For shared channels, configure offline unneeded channels on an LP basis.
- Limiting the number of LPs that can access the I/O devices attached to a control unit when the control unit attaches to shared channels. In IOCP, specify the PARTITION or NOTPART keyword on the IODEVICE statement for every I/O device attaching to a control unit so that 1 or more LPs **cannot** access any of the I/O devices.
- Using the Director to block ports or prohibit dynamic connections or communication between ports.
- Combinations of the above.

To better understand how you can manage logical paths using these methods, consider the following examples.

Deactivating Unneeded Logical Partitions: Deactivating unneeded LPs can prove useful for managing how logical paths are established on CPCs in some situations.

The system establishes logical paths only when an LP is activated. Deactivating an LP results in removal of those logical paths associated with the LP. This can greatly reduce the number of logical paths requested by the system at any given time.

In Figure 2-3 on page 2-14, if all five of the LPs each share all four of the ESCON channels and all of the LPs are activated, the 3990 would be requested to establish five logical paths for each of the four shared ESCON channels (or a total of 20 logical paths). Because the 3990-3 only supports 16 logical paths, you will need to manage how logical paths are established to help ensure the I/O connectivity you require.

For example, if you used LP4 and LP5 as test LPs that did not need to be active concurrently, you could reduce the number of logical paths requested by four by not activating either LP4 or LP5. In this case, four LPs (LP1, LP2, LP3, and LP4 or LP5) configured to four shared ESCON channels would request a total of 16 logical paths. Later, you could transfer logical paths between LP4 and LP5 by first deactivating one LP to remove its logical paths, then activating the other LP to use the freed logical paths.




Figure 2-3. Deactivating Unneeded Logical Partitions

Configuring Offline Unneeded Channels or Shared Channels on an LP Basis:

You can configure offline unneeded channels or shared channels on an LP basis to manage how logical paths are established. In Figure 2-4 on page 2-15, all five LPs need to be active concurrently. If all five LPs had each of the four shared ESCON channels configured online, 20 logical paths (four logical paths for each of the five LPs) would be requested, exceeding the 3990's logical path capacity.

However, if LP4 or LP5 (both test LPs) did not require four channel paths each to the 3990, you could configure offline two of the four channel images used by LP4 and two of the four channel images used by LP5, reducing the total number of logical paths requested from 20 to 16 and matching the 3990's logical path capacity.




Figure 2-4. Configuring Offline Unneeded Channels or Shared Channels on an LP Basis

Note: Because the 3990-3 supports only eight logical paths per cluster, you would need to configure offline the channel images so that the number of logical paths requested from each cluster remains at eight.

It is also possible to manage how logical paths are established by using IOCP or Hardware Configuration Definition (HCD) to create I/O configuration definitions that:

- Define a subset of LPs that will have their corresponding channels configured online at power-on reset (POR) (CHPID access list)
- Allow LPs to configure online their channels at a later time (CHPID candidate list)

Use IOCP or HCD to define the access lists and candidate lists for channel paths to determine the configurability of a channel to an LP. This capability exists for both unshared and shared channels and can help automate and establish the configuration in Figure 2-4. Additionally, HCD allows you to dynamically change the access list and candidate list for a channel path.

Defining Devices to a Subset of Logical Partitions: You can limit I/O device access from LPs to I/O devices assigned to shared channels by using IOCP or HCD to specify device candidate lists. By defining devices attached to a control unit to a subset of LPs, you can manage which LPs will attempt to establish logical paths to the control unit through a shared channel.

If you define no devices to a control unit from a particular LP, the shared channel associated with the LP will not attempt to establish a logical path. However, if there is at least one device defined to the control unit for the shared channel associated with a particular LP, the shared channel for the LP will attempt to establish a logical path to the control unit for the LP.

In Figure 2-5 on page 2-17, LP access to a series of 3174s is managed through use of the device candidate lists for the I/O devices attached to the control units. The shared ESCON channel will attempt to establish only one logical path to each of the 3174s even though 5 LPs share the channel. This is useful because the 3174 in non-SNA mode only supports one logical path.

In the example, the channel only attempts to establish a logical path for LP1 to the 3174 defined as control unit 10 because only LP1 has a device defined to that control unit. Similarly, only LP2 can access the 3174 defined as control unit 11, only LP3 can access the 3174 defined as control unit 12, only LP4 can access the 3174 defined as control unit 13, and only LP5 can access the 3174 defined as control unit 14.

Partial IOCP Deck for the Configuration: The following is a partial IOCP deck for the example in Figure 2-5 on page 2-17.

```
CHPID PATH=30,SHARED

CNTLUNIT CUNUMBR=10,PATH=30
IODEVICE ADDRESS=VVVV,CUNUMBR=10,PART=LP1

CNTLUNIT CUNUMBR=11,PATH=30
IODEVICE ADDRESS=VVVV,CUNUMBR=11,PART=LP2

CNTLUNIT CUNUMBR=12,PATH=30
IODEVICE ADDRESS=VVVV,CUNUMBR=12,PART=LP3

CNTLUNIT CUNUMBR=13,PATH=30
IODEVICE ADDRESS=VVVV,CUNUMBR=13,PART=LP4

CNTLUNIT CUNUMBR=14,PATH=30
IODEVICE ADDRESS=VVVV,CUNUMBR=14,PART=LP5
```




Figure 2-5. Defining Devices to a Subset of Logical Partitions

In Figure 2-6, a 3174 in SNA mode is defined as five control unit headers (CUHs). Because each 3174 CUH supports a maximum of one logical path, it is equally important in this example that the shared channel only attempts to establish a single logical path to each 3174 CUH.




Figure 2-6. Defining Devices to a Subset of Logical Partitions

Even though 5 LPs share the ESCON channel, the channel will only attempt to establish a logical path for LP1 to CUH 0 because only LP1 has a device defined on that CUH. Similarly, only LP2 can access CUH 1, only LP3 can access CUH 2, only LP4 can access CUH 3, and only LP5 can access CUH 4.

Partial IOCP Deck for the Configuration: The following is a partial IOCP deck for the example in Figure 2-6.

CHPID PATH=30,SHARED

CNTLUNIT CUNUMBR=10,PATH=30,CUADD=0
IODEVICE ADDRESS=VVVV,CUNUMBR=10,PART=LP1

CNTLUNIT CUNUMBR=11,PATH=30,CUADD=1
IODEVICE ADDRESS=WWW,CUNUMBR=11,PART=LP2

CNTLUNIT CUNUMBR=12,PATH=30,CUADD=2
IODEVICE ADDRESS=XXXX,CUNUMBR=12,PART=LP3

CNTLUNIT CUNUMBR=13,PATH=30,CUADD=3
IODEVICE ADDRESS=YYYY,CUNUMBR=13,PART=LP4

CNTLUNIT CUNUMBR=14,PATH=30,CUADD=4
IODEVICE ADDRESS=ZZZZ,CUNUMBR=14,PART=LP5

Using a Director to Block Ports or Prohibit Dynamic Connections or Communication: When ESCON or FICON Directors are used in an I/O configuration, you can prevent channels from establishing logical paths or can remove established logical paths by either blocking a Director port or by prohibiting a dynamic connection or communication between two Director ports.

In terms of logical path removal, blocking a Director port connected to a channel produces a similar outcome to configuring offline a channel or all channel images of a shared channel. Blocking a Director port connected to a control unit prevents any logical path from being established to the attached control unit port.

You can more selectively prevent logical paths from being established by prohibiting a dynamic connection between two ESCON Director ports or dynamic communication between two FICON Director ports instead of blocking a Director port. By prohibiting a dynamic connection or communication between two Director ports, you can control which channels have connectivity to a control unit port rather than blocking all connectivity to the control unit port.

Prohibiting a dynamic connection or communication between two Director ports affects all channel images of a shared channel. The system will not establish any logical paths to the attached control unit port from any of the LPs that share the ESCON or FICON channel.

You can prohibit dynamic connections or communication between Director ports by modifying the active configuration table. The active configuration table specifies the connectivity status of a port relative to the other ports on the Director. When a Director is first installed, it has a default configuration that allows any-to-any connectivity (every port can dynamically connect or communicate with every other port). If you require a different configuration than this, you can define and designate a different table to be the default configuration used at power-on of the Director. This table will allow only those dynamic connections or communication necessary to establish the logical paths the configuration requires. Dynamic connections or communication necessary to establish other logical paths (for example, those necessary for backup configurations) would be prohibited by the default configuration of the Director.

Figure 2-7 on page 2-20 shows an example of prohibiting dynamic connections. CPC1, CPC2, CPC3, and CPC4 are all production systems and CPC5 is a backup system to be used only if one of the other CPCs fail. If the default configuration used by the ESCON Director (ESCD) prohibits all dynamic connections between CPC5 and the 3990, the 3990 will only be requested to establish a total of 16 logical paths from the channels on CPC1, CPC2, CPC3, and CPC4. If one of four

production CPCs fails, you could transfer the logical paths from the failing CPC to the backup CPC by prohibiting the dynamic connection to the failed CPC and allowing the dynamic connection to the backup CPC.

If a control unit is connected to more than one Director, it is necessary to coordinate allocation of the control unit's logical paths across all of the Directors. You can use the System Automation for z/OS (SA z/OS) to dynamically manage the Directors and logical paths by sending SA z/OS commands to reconfigure one or more Directors. SA z/OS then sends the appropriate operating system Vary Path requests. SA z/OS can also provide coordination between operating systems when logical paths are removed from one system and transferred to another system as a result of blocking Director ports or prohibiting Director dynamic connections or communication.




Figure 2-7. Using the ESCD to Manage Logical Paths by Prohibiting Dynamic Connections

Shared Channel Overview

MIF allows channels to be shared among multiple LPs. Shared channels are configured to an LP giving the LP a channel image of the shared channel that it can use. Each channel image allows an LP to independently access and control the shared channel as if it were a physical channel assigned to the LP.

By providing the logical equivalent of multiple physical channels dedicated to multiple LPs, a shared channel can reduce hardware requirements without a corresponding reduction in I/O connectivity. This reduction in hardware requirements can apply to physical channels, Director ports, and control unit ports, depending on the configuration.

MIF Requirements

To take advantage of MIF, you need:

- An ESCON-capable operating system:
 - z/OS
 - z/VM
 - z/VSE
 - AIX/ESA®
 - TPF
 - z/TPF
- ESCON channels operating in native mode (CNC) or channel-to-channel mode (CTC), FICON channels attached to an ESCON director (FCV), FICON channels (FC), coupling facility channels (ICP, CBP, CIB, or CFP), open systems adapter channels (OSC, OSD, OSE or OSN), internal queued direct communication (HiperSockets) channels (IQD), and fibre channel protocol channels (FCP).

Note: ESCON channels that attach to a 9034 ESCON Converter Model 1 (CVC or CBY) cannot be shared among LPs.

- IBM ESCON-capable or FICON-capable control units, or fibre channel switches.

Understanding ESCON and MIF Topologies

This section describes the following I/O topologies:

- Point-to-point topology
- Switched point-to-point topology
- MIF channel sharing topology

Point-to-Point Topology: The traditional point-to-point topology requires a unique path between any two points that communicate. A channel and CU adapter are required between each communication point. In a point-to-point configuration, a channel can communicate with only one control unit. A control unit that communicates with more than one channel requires a separate control unit adapter interface to each channel. The configuration limitations in a point-to-point topology can lead to inefficient use of channels and control units. See Figure 2-8 on page 2-22 for an example of the point-to-point topology.

Switched Point-to-Point Topology: Switched point-to-point topologies eliminate the disadvantages of point-to-point and multidrop topologies. Switched point-to-point topology enables switching between any channel path and control unit adapter when using an ESCON or FICON Director. By using the Director, paths can also be shared by multiple points. The Director enables one channel to switch to multiple control units, or one control unit to switch to multiple channels. Paths necessary to satisfy connectivity requirements in a point-to-point topology are not required in the switched point-to-point topology. This can reduce the number of channels and

control unit interfaces required without a corresponding reduction in I/O activity. See Figure 2-8 for an example of switched point-to-point topology.

MIF Channel Sharing Topology: MIF further improves control unit connection topologies for CPCs with multiple LPs. MIF enables many LPs to share a physical channel path. MIF can reduce the number of channels and control unit interfaces required without a corresponding reduction in I/O connectivity. See Figure 2-8 for an example of MIF channel sharing topology.

MIF Performance Planning Considerations

Your installation can take advantage of MIF performance enhancements offered by:

- Understanding and utilizing I/O-busy management enhancements
- Planning for concurrent data transfer
- Understanding examples of MIF consolidation

Understanding and Utilizing I/O-Busy Management Enhancements: This section shows how the various ESCON or FICON and MIF topologies offer improvements in managing I/O-busy conditions. Figure 2-8 compares the point-to-point, switched point-to-point, and MIF channel sharing topologies.




Figure 2-8. Progression of Busy Condition Management Improvements

Point-to-Point Topologies: Concentrate I/O attempts at the control unit level and are distance dependent. At the time of a control unit busy encounter, the control unit must present control unit busy status to the channel. Once the control unit is free, it presents a control unit no longer busy status to the channel. This process of presenting status to the channel requires control unit processing and many trips over the control unit to channel link.

Switched Point-to-Point Topologies: Concentrate I/O attempts within the Director and therefore encounter switch port busies. The processing of switch port busies

does not require any control unit involvement. Busies are handled by the ESCON or FICON Director. Therefore, the control unit is effectively relieved of handling busy conditions and is able to handle more I/O requests. Because switch port busies require fewer trips over the ESCON or FICON connection link, they are less sensitive to increased distances than control unit busy encounters.

MIF Channel Sharing: Moves busy management back into the channel subsystem, providing the most efficient management of busy conditions. Because multiple LPs access the same physical channel, I/O attempts are concentrated at the channel level. A CU busy or a switch port busy is handled as a channel busy.

Planning for Concurrent Data Transfer: Before you can consolidate channels, you must be aware of the channel requirements of the particular control units you are configuring. Table 2-4 shows the maximum recommended ESCON channels for a single system going to a control unit. The number of channels needed is independent of the number of LPs on a system. The number of channels is based on the number of concurrent data transfers the control unit is capable of. Although the recommended number of channels satisfies connectivity and performance requirements, additional channels can be added for availability.

Table 2-4. MIF Maximum Channel Requirements

Control Unit Type	Maximum Channels (See note)
9343 Storage Controller	
9343 D04	4
9343 DC4	4
3990 Storage Control	
3990-2	4
3990-3	4
3990-6	4
3490 Magnetic Tape Subsystem	
3590	2
3490 D31, D32, D33, D34	1
3490 C10, C11, C12	1
3490 A01	1
3490 A10	1
3490 A02	2
3490 A20	2
2074 Console Support Controller	
2074 002	2
3174 Establishment Controller	
3174 12L (SNA only)	1
3174 22L (SNA only)	1
Note: The recommended maximum channels given in this table are based on performance and connectivity requirements and do not include channels for availability.	

Understanding Examples of MIF Consolidation: The following examples provide some general guidelines to show how MIF can help you consolidate and use hardware resources more efficiently:

ESCON Configurations: Figure 2-9 on page 2-24 shows how four shared ESCON channels can replace 16 unshared (dedicated or reconfigurable) ESCON channels and use 12 fewer control unit ports.




Figure 2-9. Consolidating ESCON Channels and ESCON Control Unit Ports

ESCD Configurations: Figure 2-10 on page 2-25 shows how shared ESCON channels can reduce ESCD port requirements. In this example, two shared ESCON channels replace 10 unshared (dedicated or reconfigurable) ESCON channels and use eight fewer ESCD ports without a reduction in I/O connectivity.




Figure 2-10. Consolidating ESCON Channels and ESCD Ports

ESCON CTC Configurations: Figure 2-11 shows how shared ESCON channels can reduce the ESCON channel requirements for ESCON CTC configurations. In this example, the CPC requires CTC communications among all its LPs.




Figure 2-11. Consolidating ESCON Channels Used for ESCON CTC Communications

By using two shared ESCON CTC/CNC pairs (4 shared ESCON channels), you can:

- Replace five unshared ESCON CTC/CNC pairs (10 unshared ESCON channels) and the ESCD used to connect them
- Provide full redundancy

I/O connectivity is maintained while hardware requirements (channels and an ESCD) are reduced.

In situations where ESCON CTC communication is required among LPs that exist on two or more CPCs, shared channels can reduce even further channel and other hardware requirements and their associated cost.

ESCON CTC configurations are well-suited to take advantage of the consolidation benefits associated with shared channels. CTC/CNC and CTC/FCV pairs used for ESCON CTC communications have no limitation on the number of logical paths that can be established between them. The only limitations are the number of control units that can be defined for an ESCON CTC channel and the performance expectations you determine for your configuration.

Infrequently Used ESCON or FICON Control Units: ESCON or FICON control units not frequently used can make use of shared channels. You can attach such a control unit to a shared channel that is also attached to other, more frequently used control units without adding greatly to the channel utilization of the shared channel. The control unit within the Director is a good example of this.

Notes:

1. You cannot define a control unit (or multiple control units with common I/O devices) to a mixture of shared and unshared channel paths in the same IOCDs.
2. You cannot define more than one control unit with the same CUADD to the same link on a Director (or point-to-point) if the attaching CHPIDs are shared.

Unshared ESCON or FICON Channel Recommendations

Not all ESCON or FICON configurations benefit from the use of shared channels. There are some configurations where use of an unshared channel is more appropriate. Consider the following:

- **Logical Path Limitations of the Control Unit**

While many ESCON control units can communicate with multiple LPs at a time using multiple logical paths, there are some ESCON-capable control units that can only communicate with one LP at a time. For example, consider the 3174 Establishment Controller (Models 12L and 22L). When configured in non-SNA mode, the 3174 establishes only one logical path at a time. A shared channel would offer no connectivity benefit in this situation. However, if you defined an unshared, reconfigurable channel to the 3174, it would allow you to dynamically reconfigure the channel for any LP that had to communicate with the 3174 at a given time.

- **Channel Utilization**

Typically, the channel utilization of shared channels will be greater than unshared channels.

If you use shared channels to consolidate channel resources, you must consider the channel utilization of all the channels you consolidate. The channel utilization of a shared channel will roughly equal the sum of the channel utilizations of each unshared channel that it consolidates. If this total channel utilization is capable of decreasing performance, you should consider using unshared channels or a different configuration of shared and unshared channels to meet your connectivity needs.

Dynamically Managed CHPIDs

A key aspect of the Intelligent Resource Director (IRD) provided by z/OS's WLM component is Dynamic CHPID Management (DCM). DCM provides the ability to have the system automatically move the available channel bandwidth to where it is

most needed. CHPIDs identified as managed in the IOCDs (via CHPARM and IOCLUSTER keywords) are dynamically shared among z/OS images within an LPAR cluster.

Prior to DCM, available channels had to be manually balanced across I/O devices in an attempt to provide sufficient paths to handle the average load on every controller. Natural variability in demand means that some controllers at times have more I/O paths available than they need, while other controllers possibly have too few. DCM attempts to balance responsiveness of the available channels maximizing the utilization of installed hardware. Fewer overall channels are required because the DCM CHPIDs are more fully utilized. RMF provides a report showing the average aggregate utilization for all managed channels.

By using DCM, you now only have to define a minimum of one nonmanaged path and up to seven managed paths to each control unit (although a realistic minimum of two nonmanaged paths are recommended), with dynamic channel path management taking responsibility for adding additional paths as required. For more information on defining and using DCM, including detailed examples, see the *z/OS Intelligent Resource Director*, SG24-5952.

IOCP Coding Specifications

ICP IOCP can only generate an LPAR IOCDs. No IOCP invocation parameter is required to generate an LPAR IOCDs.

IOCP Statements for ICP

The RESOURCE statement is used to specify all the logical partition names defined in a machine configuration. To plan for growth in the number of logical partitions in the configuration, one or more asterisks (*) may be used to specify that one or more logical partitions are to be reserved along with their associated CSS and MIF image IDs. A reserved LP can only be specified for a dynamic-capable IOCDs. A dynamic-capable IOCDs is built when using HCD on z/OS or z/VM or specifying IOCP CMS utility option DYN for z/VM. Space in the hardware system area (HSA) is allocated for reserved LPs but cannot be used until a dynamic I/O configuration change is made to assign a name to the LP. The following rules apply when specifying reserved LPs:

- A reserved LP must have a user-specified MIF image ID
- A reserved LP cannot have any channel paths assigned to it
- An IOCDs cannot contain only reserved LPs. At least one LP must be defined with a name.

Dynamic CHPID Management (DCM) channel paths defined for a given LPAR cluster are shareable among all active LPs that have joined that cluster. Other than DCM channel paths, you must assign each channel path to a logical partition in an LPAR IOCDs. For each DCM channel path, ICP requires the CHPARM keyword have a value of 01 and the IOCLUSTER keyword on a CHPID statement. All other channel paths require the PARTIPARTITION, NOTPART, or SHARED keyword on all CHPID statements unless a channel path is defined as spanned by specifying multiple CSS IDs in the PATH keyword of the IOCP CHPID statement.

Use the CHPARM and IOCLUSTER keywords on the CHPID statement to specify channel paths reserved for the use of a particular LPAR cluster. A DCM channel path becomes available to a candidate logical partition when the LP is activated and joins the specified cluster.

Use the CHPID PARTITION, NOTPART, and SHARED keywords to determine which:

- Channel paths are assigned to each LP
- Devices and control units are shared among LPs
- Channel paths are reconfigurable
- Channel paths are shared

Use the CHPID CPATH keyword to connect two internal coupling channels.

Use the CHPID PATH keyword to define a channel path as spanned to multiple CSSs. Spanned channel paths are also shared channel paths.

DCM channel paths are implicitly shared. Use of the IOCLUSTER keyword implies a null access list (no logical partition has the channel path brought online at activation) and a candidate list of all defined logical partitions. The IOCLUSTER keyword is mutually exclusive with the PARTITION and NOTPART keywords.

All LP names that you specify in the CHPID statements must match those specified in the RESOURCE statement. An IOCDS must have at least one LP name defined.

```
PARTITION={ (CSS(cssid), {name|0} [,REC]) |  
            (CSS(cssid), access list) |  
            (CSS(cssid), (access list) [, (candidate list)] [,REC]) |  
            ((CSS(cssid), (access list) [, (candidate list)]), ...) }
```

```
NOTPART={ (CSS(cssid), access list) |  
          ((CSS(cssid), (access list) [, (candidate list)]), ...) }
```

```
IOCLUSTER=cluster_name
```

```
SHARED
```

```
CPATH=(CSS(cssid), chpid number)
```

Where:

name specifies the name of the LP that has authority to access the CHPID. The LP name is a 1–8 alphanumeric (0–9, A–Z) character name that must have an alphabetic first character. Special characters (\$, #, @) are not allowed. A reserved LP cannot have any channel paths assigned to it.

The following words are reserved and you cannot use them as LP names:

```
PHYSICAL  
REC  
SYSTEM  
PRIMnnnn (where nnnn are digits)
```

ICP IOCP supports a maximum of 60 LP names for the CPC.

cluster_name specifies the name of an LPAR cluster that has authority to access the specified DCM CHPID. The name of the LPAR cluster is a one-to eight- alphanumeric character name (0-9, A-Z) that must have an alphabetic first character. Special characters (\$, #, @) are not allowed.

REC specifies that the CHPID is reconfigurable. A reconfigurable CHPID must have an initial access list of one LP name. It's candidate list must consist of one or more LP names.

access list specifies the LPs that have initial access to the CHPID at the completion of the initial power-on reset. An LP name may only appear once in an access list.

You can specify that no LPs will access the channel path following LP activation for the initial POR of an LPAR IOCDS. Specifying 0 indicates a null access list.

candidate list specifies the LPs that have authority to access the CHPID. Any LP that is not in a CHPID's candidate list cannot access the CHPID.

You may specify as many LP names as your CPC supports. However, the number of unique LP names specified in both the access list and candidate list may not exceed the number of LPs your CPC supports.

If you specify the candidate list, you do not need to specify again the LP names specified in the initial access list. The initial access list is always included in the candidate list.

An LP name can only appear once in a candidate list. If the candidate list is not specified, it defaults to all LPs in the configuration for reconfigurable and shared channels.

Note: It is highly recommended that a peer mode CHPID (CBP, CFP, CIB, or ICP) have at most one coupling facility LP specified in its initial access list in order to avoid confusion on subsequent LP activations. A peer mode CHPID can be online to only one coupling facility LP at a time.

Using the SHARED keyword specifies that the channel paths on the CHPID statement are shared. More than one LP, at the same time, can access a shared CHPID. When CHPIDs are not shared, only one LP can access it. Although you can dynamically move a reconfigurable CHPID between LPs, it can only be accessed by 1 LP at any given time. Only CNC, CTC, CBP, CIB, CFP, ICP, OSC, OSD, OSE, OSN, FC, FCV, FCP, and IQD channel paths (TYPE keyword) can be shared.

The CPATH keyword is only valid for ICP and CIB channel paths (TYPE keyword) and required for all ICP and CIB definitions. CPATH specifies the connection between 2 ICPs at either end of a coupling link:

```
PATH=FE,TYPE=ICP,CPATH=FF,...  
PATH=FF,TYPE=ICP,CPATH=FE,...
```

specifies that ICP channel path FF connects to ICP channel path FE. Every ICP channel path of a coupling facility must be connected to an ICP channel path of a z/OS LP. The connection needs to be specified for each channel path. ICP channel paths cannot connect to each other if they both have candidate lists with the same, single logical partition. Note that this prevents the definition of internal coupling channels in a LPAR configuration with only one logical partition. Also, an ICP channel path cannot connect to itself.

The CPATH value for a CIB CHPID specifies the CSS and CHPID number this CIB CHPID connects with on the target system. For example:

```
PATH=C0,TYPE=CIB,CPATH=(CSS(1),D0),...
```

Defines a CIB CHPID, C0, on this system that connects with CIB CHPID D0 in CSS 1 on the remote system.

Shared Devices Using Shared Channels

MIF allows you to use shared channels when defining shared devices. Using shared channels reduces the number of channels required, allows for increased channel utilization, and reduces the complexity of your IOCP input.

Note: You cannot mix shared and unshared channel paths to the same control unit or device.

The following is an example of an IOCDs with a shared device.




Figure 2-12. Shared Devices Using Shared ESCON Channels

The following is the IOCP coding for the above figure.

```
CHPID PATH=(30),TYPE=CNC,SWITCH=03,SHARED . . .
CHPID PATH=(34),TYPE=CNC,SWITCH=03,SHARED . . .
CNTLUNIT CUNUMBR=000,PATH=(30,34),UNITADD=((90)),LINK=(E0,E0),UNIT=xxx . . .
IODEVICE ADDRESS=(190),CUNUMBR=000,UNIT=xxx . . .
```

Shared Devices Using Unshared Channels

When coding an IOCP input file, the following specifications are allowed:

- Duplicate device numbers can be specified within a single IOCP input file, if device numbers are not duplicated within an LP.
- You can assign a maximum of eight channel paths from each LP to a device.
Device sharing among LPs is accomplished by attaching multiple channel paths from each LP to a device.

The following section illustrates IOCP coding for IOCDs when shared devices on unshared channels and duplicate device numbers are specified.

Shared Devices: The following examples illustrate this concept by showing the physical connectivity of an I/O configuration for multiple LPs and the IOCP coding for the same configuration.

Using Channels: Figure 2-13 on page 2-32 shows an example of an I/O configuration with a device shared by each of the four logical partitions. In this representation of a shared device, each logical partition views device 190 as part of

its own I/O configuration. Notice the recoverability characteristics of this configuration: each logical partition has two channel paths to the shared device, each attached to a different storage director.




Figure 2-13. Physical Connectivity of Shared Device 190

The following example shows the IOCP statement for Figure 2-13.

```

CHPID PATH=(10),PART=(A,REC)
CHPID PATH=(14),PART=(A,REC)
CHPID PATH=(18),PART=(B,REC)
CHPID PATH=(1C),PART=(B,REC)
CHPID PATH=(20),PART=(C,REC)
CHPID PATH=(24),PART=(C,REC)
CHPID PATH=(28),PART=(D,REC)
CHPID PATH=(2C),PART=(D,REC)
CNTLUNIT CUNUMBR=0001,PATH=(10,18,20,28),UNITADD=((90)) . . .
CNTLUNIT CUNUMBR=0002,PATH=(14,1C,24,2C),UNITADD=((90)) . . .
IODEVICE ADDRESS=(190),CUNUMBR=(0001,0002) . . .
  
```

If 8 or less channels attach to the device, this method of defining the IOCP input provides greater flexibility because it allows you to move CHPIDs from one LP to another and eliminates possible conflicts (see Figure 2-16 on page 2-35).

Figure 2-14 on page 2-33 shows an alternative method of defining the configuration. This method is required if there are greater than 8 paths to the device. Note that this logical representation has the same recoverability characteristics as the physical connectivity:

- Each LP has two channel paths to the shared device
- Each LP is attached to a different storage director

However, paths cannot be moved between the LPs.




Figure 2-14. Logical View of Shared Device 190

The following example shows the IOCP statement for Figure 2-14.

```
CHPID PATH=(10),PARTITION=(A), . . .
CHPID PATH=(14),PARTITION=(A), . . .
CNTLUNIT CUNUMBR=0001,PATH=(10),UNITADD=((90)) . . .
CNTLUNIT CUNUMBR=0002,PATH=(14),UNITADD=((90)) . . .
IODEVICE ADDRESS=(190),CUNUMBR=(0001,0002) . . .

CHPID PATH=(18),PARTITION=(B), . . .
CHPID PATH=(1C),PARTITION=(B), . . .
CNTLUNIT CUNUMBR=1001,PATH=(18),UNITADD=((90)) . . .
CNTLUNIT CUNUMBR=1002,PATH=(1C),UNITADD=((90)) . . .
IODEVICE ADDRESS=(190),CUNUMBR=(1001,1002) . . .

CHPID PATH=(20),PARTITION=(C) . . .
CHPID PATH=(24),PARTITION=(C) . . .
CNTLUNIT CUNUMBR=2001,PATH=(20),UNITADD=((90)) . . .
CNTLUNIT CUNUMBR=2002,PATH=(24),UNITADD=((90)) . . .
IODEVICE ADDRESS=(190),CUNUMBR=(2001,2002) . . .

CHPID PATH=(28),PARTITION=(D), . . .
CHPID PATH=(2C),PARTITION=(D), . . .
CNTLUNIT CUNUMBR=3001,PATH=(28),UNITADD=((90)) . . .
CNTLUNIT CUNUMBR=3002,PATH=(2C),UNITADD=((90)) . . .
IODEVICE ADDRESS=(190),CUNUMBR=(3001,3002) . . .
```

Duplicate Device Numbers for Different Physical Devices

Figure 2-15 on page 2-34 illustrates a configuration where duplicate device numbers are used to represent a console (110) and a printer (00E) within each of four logical partitions.




Figure 2-15. LPAR Configuration with Duplicate Device Numbers

The following example shows the IOCP statement for Figure 2-15. This IOCP coding example groups the input statements by logical partition. When coding IOCP, view the I/O devices from a logical partition perspective.

```

CHPID PATH=(10),PARTITION=(A), . . .
CHPID PATH=(14),PARTITION=(A), . . .
CNTLUNIT CUNUMBR=0011,PATH=(10),UNITADD=(10), . . .
CNTLUNIT CUNUMBR=0012,PATH=(14),UNITADD=(0E), . . .
IODEVICE ADDRESS=(110),CUNUMBR=(0011), . . .
IODEVICE ADDRESS=(00E),CUNUMBR=(0012), . . .

CHPID PATH=(18),PARTITION=(B), . . .
CHPID PATH=(1C),PARTITION=(B), . . .
CNTLUNIT CUNUMBR=0013,PATH=(18),UNITADD=(10), . . .
CNTLUNIT CUNUMBR=0014,PATH=(1C),UNITADD=(0E), . . .
IODEVICE ADDRESS=(110),CUNUMBR=(0013), . . .
IODEVICE ADDRESS=(00E),CUNUMBR=(0014), . . .

CHPID PATH=(20),PARTITION=(C), . . .
CHPID PATH=(24),PARTITION=(C), . . .
CNTLUNIT CUNUMBR=0015,PATH=(20),UNITADD=(10), . . .
CNTLUNIT CUNUMBR=0016,PATH=(24),UNITADD=(0E), . . .
IODEVICE ADDRESS=(110),CUNUMBR=(0015), . . .
IODEVICE ADDRESS=(00E),CUNUMBR=(0016), . . .

CHPID PATH=(28),PARTITION=(D), . . .
CHPID PATH=(2C),PARTITION=(D), . . .
CNTLUNIT CUNUMBR=0017,PATH=(28),UNITADD=(10), . . .
CNTLUNIT CUNUMBR=0018,PATH=(2C),UNITADD=(0E), . . .
IODEVICE ADDRESS=(110),CUNUMBR=(0017), . . .
IODEVICE ADDRESS=(00E),CUNUMBR=(0018), . . .

```

Eight IODEVICE statements are used, one for each console and one for each printer that has a duplicate device number. Device numbers 110 and 00E occur four times each; however, they are not duplicated within a logical partition. When coding an IOCP input file, remember that the unique device number rule applies for logical partitions in an IOCDs.

Figure 2-16 shows another example of a logical partition configuration in which the device number for a console (110) is duplicated for all four logical partitions.




Figure 2-16. Duplicate Device Numbers for Console

The following example shows the IOCP coding for the previous configuration. Four IODEVICE and four CNTLUNIT statements are used, one each for the console within each logical partition that has a duplicate device number.

```
CHPID PATH=(10),PARTITION=(A), . . .
CHPID PATH=(14),PARTITION=(A), . . .
CNTLUNIT CUNUMBR=0001,PATH=(10,14),UNITADD=((10)), . . .
IODEVICE ADDRESS=(110),CUNUMBR=(0001), . . .
```

```
CHPID PATH=(18),PARTITION=(B), . . .
CHPID PATH=(1C),PARTITION=(B), . . .
CNTLUNIT CUNUMBR=0002,PATH=(18,1C),UNITADD=((10)), . . .
IODEVICE ADDRESS=(110),CUNUMBR=(0002), . . .
```

```
CHPID PATH=(20),PARTITION=(C), . . .
CHPID PATH=(24),PARTITION=(C), . . .
CNTLUNIT CUNUMBR=0003,PATH=(20,24),UNITADD=((10)), . . .
IODEVICE ADDRESS=(110),CUNUMBR=(0003), . . .
```

```
CHPID PATH=(28),PARTITION=(D), . . .
CHPID PATH=(2C),PARTITION=(D), . . .
CNTLUNIT CUNUMBR=0004,PATH=(28,2C),UNITADD=((10)), . . .
IODEVICE ADDRESS=(110),CUNUMBR=(0004), . . .
```

Duplicate Device Number Conflicts: IOCP allows duplicate device numbers in an IOCDS only if the duplicate device numbers do not occur in the same logical partition. Therefore, IOCP allows systems to use different logical partitions to integrate a processor complex without changing device numbers.

IOCP requires a unique device number for each device within a logical partition. When IOCP completes without error, the initial configuration contains no duplicate device number conflicts within a logical partition.

Conflicts can occur when the I/O configuration is modified. If a channel path is configured to a logical partition and devices attached to the channel path have device numbers that are already assigned in the receiving logical partition to other online channel paths, a conflict results.

When an I/O configuration is dynamically modified so the logical partition can gain access to a device not previously accessible, a device conflict can occur. The conflicts are detected when commands are processed that change the I/O configuration or when you attempt to activate the logical partition which has the device number conflict. A message displays identifying the error.

The identified device cannot be accessed while a conflict exists. Two types of conflict are possible:

- Conflicts between device numbers for the same device (a shared device)
- Conflicts between device numbers for different devices (unshared devices)

Activation fails if a duplicate device number conflict exists.

Examples of Duplicate Device Number Conflicts: Figure 2-17 provides two examples of duplicate device number conflict.




Figure 2-17. Two Examples of Duplicate Device Number Conflicts

The following example shows the IOCP statement for Figure 2-17. Both examples use identical IOCP statements.

```
CHPID PATH=(00),PARTITION=(ZOSPROD,REC)
CHPID PATH=(04),PARTITION=(ZOSPROD,REC)
CNTLUNIT CUNUMBR=0001,PATH=(00,04),UNITADD=80
IODEVICE ADDRESS=180,CUNUMBR=0001
```

```
CHPID PATH=(10),PARTITION=(ZOSTEST)
CNTLUNIT CUNUMBR=0002,PATH=(10),UNITADD=80
IODEVICE ADDRESS=180,CUNUMBR=0002
```

Channel path 04 is reassigned from ZOSPROD to ZOSTEST in each example. This creates a duplicate device number conflict for device number 180 when the devices are connected to two different control units. This occurs because a device numbered 180 already exists on the original channel path 10. If such conflicts occur, the operator must know what configuration is desired.

Shared Device

In the example on the left, the duplicate device numbers refer to the same device from different logical partitions (a new path to the same device has been moved to ZOSTEST). This may result in a performance problem because the control program in logical partition OS390 cannot access the device from channel path 4.

Unshared Device

In the example on the right, the duplicate device numbers refer to a different device from each logical partition (a new device has been moved to ZOSTEST). This may result in a data integrity problem because the control program in logical partition ZOSTEST cannot access the correct device from channel path 04.

Resolving Duplicate Device Number Conflicts: Consider options A, B, and C when planning the I/O configuration and the reconfigurability of channel paths. You can resolve duplicate device number conflicts by choosing one of the options:

A Use the original channel path:

If the receiving logical partition does not need a new path to a shared device or does not need the new (unshared) device, take no action. The conflict is resolved by using only the original path (shared device) or the original device. (Access is still allowed to any nonconflicting devices on the newly configured channel path.)

In Figure 2-17 on page 2-36 ZOSTEST can access device 180 only through channel path 10 if the operator takes no action in response to the conflict message.

B Deconfigure the original channel path:

If the logical partition must have the reassigned channel path to a shared device or access to a new (unshared) device, the conflict is resolved by substituting the reassigned channel path for the original channel path. Do the following:

1. Configure offline the original channel path (CHP 10 in Figure 2-17).
2. Configure offline and then online the reassigned channel path (CHP 04 in Figure 2-17).
3. If necessary, configure online the original channel path (CHP 10 in Figure 2-17). Another conflict message is issued because a new conflict has been created. The operator then ignores this conflict as described in option A. (Access is still allowed to any nonconflicting devices on the original channel path.)

In Figure 2-17, ZOSTEST can access device 180 only through channel path 04 if the preceding steps are performed in response to the conflict message.

C Change the I/O Configuration:

Only option C provides a permanent resolution to a device number conflict.

If the logical partition must have access to all devices over the original channel path and the reassigned channel path (shared devices), or to a new device and the original device (unshared devices), do one of the following:

- Create a new configuration with unique device numbers, if they are unshared devices.
- For shared devices, define a single device with access to all of the channel paths attached to the physical control units.
- For a shared device assigned to unshared channel paths, change the channel paths to shared and consolidate the control units and device definitions to one each.
- If the device is assigned to shared channel paths, control access to the devices using their device candidate list.

The configuration can be activated by performing a POR or by performing a dynamic I/O configuration.

In Figure 2-17 (shared device), ZOSTEST can access device 180 through CHP 04 and CHP 10 if CHP 04 is defined to ZOSTEST in the IOCDS.

In Figure 2-17 (unshared device), ZOSTEST can access either device 180 (unshared device) if one or both of the devices are assigned a new device number in the IOCDS.

When a device number conflict exists, logical partitions will fail to activate. This will happen when one of the following occurs:

- The receiving logical partition was deactivated when a channel path is reassigned
- The receiving logical partition is deactivated after a channel path is reassigned

Failure to activate can result if options A or B are used. If a logical partition fails to activate, use option B or C to resolve the conflict and to activate the logical partition.

In Figure 2-17, if ZOSTEST is not active when CHP 04 is reassigned, or ZOSTEST is deactivated and then activated after CHP 04 is reassigned, ZOSTEST does not activate until the conflict over device 180 is resolved.

If you resolve the conflict by using option B do the following steps:

1. Establish the correct configuration by configuring offline one of the channel paths (CHP 04 or CHP 10)
2. Configure offline and then online the other channel path

If it's necessary to have access to other devices on the first channel path, the operator can configure online the first channel path while the LP is active. Ignore the messages issued at the hardware console.

The following IOCP statement example shows coding that removes duplicate device number conflicts for shared devices.

```
CHPID PATH=(00),PARTITION=(ZOSPROD,REC), . . .
CHPID PATH=(04),PARTITION=(ZOSPROD,REC), . . .
CHPID PATH=(10),PARTITION=(ZOSTEST), . . .
CNTLUNIT CUNUMBR=0001,PATH=(00,04),UNITADD=80
CNTLUNIT CUNUMBR=0002,PATH=(10),UNITADD=80
IODEVICE ADDRESS=180,CUNUMBR=(0001,0002)
```

Coupling Facility Planning Considerations

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

The coupling facility allows applications, running on multiple z/OS images that are configured in a sysplex, to efficiently share data so that a transaction processing workload can be processed in parallel across the sysplex.

PR/SM LPAR allows you to define the coupling facility, which is a special logical partition (LP) that runs coupling facility control code. Coupling facility control code is Licensed Internal Code (LIC).

At LP activation, coupling facility control code automatically loads into the coupling facility LP from the support element hard disk. No initial program load (IPL) of an operating system is necessary or supported in the coupling facility LP.

Coupling facility control code runs in the coupling facility LP 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 LPs to avoid unnecessary operator activity.

Coupling facility channel hardware provides the connectivity required for data sharing between the coupling facility and the CPCs directly attached to it. Coupling facility channels are point-to-point connections that require a unique channel definition at each end of the channel. See “Coupling Facility Channels” on page 2-54.

Test or Migration Coupling Configuration

You can run a test or migration coupling facility to test and develop data sharing applications. You can define a test or migration coupling facility LP on the same CPC where other LPs are:

- Running z/OS images connected to the coupling facility
- Running noncoupled production work

A single CPC configuration has the following considerations:

- Simultaneous loss of the coupling facility and any z/OS images coupled to it (a more likely possibility in a single CPC configuration) can potentially cause extended recovery times

You can define a test or migration coupling facility with or **without** coupling facility channel hardware. See “Defining Internal Coupling Channels (TYPE=ICP)” on page 2-58 for information on how to define a test or migration facility **without** coupling facility channel hardware.

CFCC Enhanced Patch Apply

With the CFCC Enhanced Patch Apply, you can perform a disruptive install of new CFCC code on a *test* CF and run it, while a *production* CF image in the same CEC remains at the base CFCC code level. Then when the test CF is successful, the new CFCC code can be installed on the production CF. Both installs can be done without a Power On Reset (POR) of the CEC.

Production Coupling Facility Configuration

IBM recommends that you run your production applications on a sysplex that uses a production coupling facility configuration.

A properly configured production coupling facility configuration can reduce the potential for extended recovery times, achieve acceptable performance, and maximize connectivity to the coupling facility.

For production configurations, the use of one or more dedicated Coupling Facility engines is recommended; shared Coupling Facility engines are strongly discouraged. For more information, see the Important Note on page 3-31.

Production Coupling Facility Configuration for Full Data Sharing

The preferred solution for a full data sharing (IMS™, DB2®, VSAM/RLS) production parallel sysplex is a coupling facility configuration that consists of:

- One stand-alone coupling facility running as a single dedicated coupling facility LP to provide large capacity shared storage and maximum coupling facility channel connectivity (up to 64 coupling facility channels).
- A second stand-alone coupling facility, similarly configured, to reduce the possibility of a single point of failure. A second stand-alone coupling facility improves application subsystem availability by allowing fast recovery from one coupling facility to the other in the event of a coupling facility outage. Alternatively, an Internal Coupling Facility (ICF) feature can be used to provide the backup coupling facility. See “Internal Coupling Facility (ICF)” on page 2-41.

Notes:

1. The backup CF in the configuration must provide sufficient storage, processor, and connectivity resources to assume the workload of the other production CF in the event of its failure.
2. With the use of System-Managed CF Structure Duplexing for all relevant data sharing structures, it is possible to have a production data-sharing configuration that uses only 2 or more internal CFs, because duplexing avoids the “single point of failure” failure-isolation issue.

Production Coupling Facility Configuration for Resource Sharing

A viable solution for a resource sharing (XCF Signaling, Logger Operlog, RACF®, BatchPipes®, Logger Logrec, Shared Tape, GRS, WLM Enclave Support, LPAR Clusters) production level parallel sysplex is a coupling facility configuration that consists of:

- One dedicated ICF provides reduced cost of ownership without compromising sysplex availability or integrity.
- A second dedicated ICF reduces the possibility of a single point of failure. A second ICF improves application subsystem availability by allowing fast recovery from one coupling facility to the other in the event of a coupling facility outage.

Note: The backup CF in the configuration must provide sufficient storage, processor, and connectivity resources to assume the workload of the other production CF in the event of its failure.

These configurations offer the best performance, the best reliability, availability, and serviceability (RAS).

Internal Coupling Facility (ICF)

You can purchase and install one or more ICF features for use in coupling facility LPs. With this feature, the coupling facility runs on special ICF CPs that no customer software can utilize. This allows the coupling facility function to be performed on the CPC without affecting the model group and thus without impacting software licensing costs for the CP resources utilized by the coupling facility. See “Considerations for Coupling Facilities using Internal Coupling Facility (ICF) CPs” on page 3-31.

These features are ordered separately, and are distinguished at the hardware level from any Integrated Features for Linux (IFLs), Integrated Facility for Applications (IFAs) features, and zIIPs. ICFs, IFLs, IFAs, or zIIPs are perceived by the system as multiple resource pools.

With the CFCC Enhanced Patch Apply process, you can perform a disruptive install of new CFCC code on an ICF image by simply deactivating and then reactivating the CF image, without the much greater disruption of a Power On Reset (POR) of the entire CEC that contains the CF image. This greatly improves the availability characteristics of using ICFs.

ICFs are ideal for coupling resource sharing sysplexes (sysplexes that are not in production data sharing with IMS, DB2 or VSAM/RLS). You can simplify systems management by using XCF structures instead of ESCON CTC connections.

IBM does not recommend use of ICFs for most coupling facility structures involved in data sharing because of the possibility of “double outages” involving the simultaneous loss of an ICF image and one or more z/OS system images that are using the ICF for data sharing. Depending on the structure, a double outage can result in a significantly more involved recovery than a single outage of either an ICF image or a z/OS image in isolation from one another.

With the use of System-Managed CF Structure Duplexing for all relevant data sharing structures, it is possible to have a production data-sharing configuration that uses only 2 or more internal CFs, because duplexing avoids the “single point of failure” failure-isolation issue.

ICFs on stand-alone coupling facilities need configuration planning to account for storage and channels. ICFs will likely increase storage requirements for the CPC with an ICF installed, especially if software exploits the coupling facility to provide additional function not available except when running a coupling facility in a parallel sysplex.

The following table indicates the maximum number of ICFs that you can install on a given model.

CPC Model	Maximum Number of ICF's Supported
System z9 BC	
Model R07	6
Model S07	7
System z9 EC	
Model S08	8
Model S18	16
Model S28	16
Model S38	16
Model S54	16

Dynamic Coupling Facility Dispatching (DCFD)

With DCFD, the coupling facility uses CP resources in a shared CP environment efficiently, making a coupling facility using shared CPs an attractive option as a back-up coupling facility. Without DCFD, a coupling facility using shared CPs would attempt to get all the CP resource it could even when there was no real work for it to do. With dynamic coupling facility dispatching, the coupling facility monitors the request rate that is driving it and adjust its usage of CP resource accordingly. If the request rate becomes high enough, the coupling facility reverts back to its original dispatching algorithm, constantly looking for new work. When the request rate lowers, the coupling facility again becomes more judicious in its use of CP resource.

This behavior plays well into a hot-standby back-up coupling facility role. In back-up mode, the coupling facility will have a very low request rate so it will throttle back to very low CP usage. In this mode, the requests themselves will experience some elongation in response time but this will not adversely affect the performance of the overall system. Since the coupling facility is not consuming more CP resource than it needs to, you can now set the processor weights for the coupling facility to a value high enough to handle the load if the coupling facility was to take over for a failing primary coupling facility. If the primary coupling facility does fail, the requests can be moved immediately to the back-up coupling facility which can then get the CP resource it needs automatically with properly defined LP weights.

Dynamic coupling facility dispatching is particularly useful in configurations where less than one CP of capacity is needed for the back-up coupling facility use. Dynamic coupling facility dispatching is automatically enabled for any coupling facility LP that uses shared CPs, except for stand-alone coupling facilities. To enable dynamic coupling facility dispatching on stand-alone coupling facilities, use the DYNDISP coupling facility control code command. See “Coupling Facility Control Code Commands” on page 2-50. It is recommended that anytime a shared processor is used for a coupling facility, whether standalone or ICF processors, that DCFD be enabled by insuring that the DYNDISP coupling facility control command be set to ON.

Both the ICF and dynamic coupling facility dispatching enhancements represent real production configuration options at significantly reduced cost. You will need to examine the RAS trade-offs for each potential configuration to determine their acceptability for your environment.

Dynamic Internal Coupling Facility (ICF) Expansion

Important Note

Dynamic ICF Expansion is not recommended for coupling facilities that are used for primary production workload. All processors assigned to the coupling facility partition should be dedicated to that logical partition if it is used for primary production workload.

Dynamic internal coupling facility (ICF) expansion allows an ICF LP to acquire additional processing power from the LPAR pool of shared general purpose CPs being used to run production and/or test work on the system.

This is very useful when you are using an ICF LP to back up another coupling facility. In this scenario, the ICF, using dynamic ICF expansion, can acquire additional processing capacity to handle the full coupling facility workload.

Dynamic ICF expansion combines the benefits of ICFs and dynamic CF dispatching. Dynamic ICF expansion allows you to define a coupling facility LP that uses a combination of dedicated ICF CPs and shared general purpose CPs. The dedicated ICF CPs run a continuous polling loop looking for requests to process while the shared general purpose CPs use a dynamic coupling facility dispatch algorithm. This provides the capability to allow peak workload expansion into the shared pool of general purpose CPs.

For each ICF coupling facility LP, you define the number of ICF features (maximum varies by model and what has been installed) that are dedicated to that LP and the amount of additional capacity you want to be made available by specifying a number of shared general purpose CPs the ICF LP can acquire. You specify a processing weight for the shared CPs as you would for any other LP using shared CPs. Since the coupling facility is not using more CP resource than it requires on the shared CPs, you must set the processing weight to a high value but this will not affect other work during normal operation.

IBM recommends that the processing weight you specify for the shared logical CPs of a coupling facility using dynamic ICF expansion be set higher than the weight for all other logical CPs that are sharing the same physical CPs. On occasions when the shared CP is needed by the coupling facility, you can achieve the best response time when this CP is dispatched as quickly as possible. To help ensure that this occurs, you should assign the shared coupling facility CP a weight which is much higher than that of the other LPs' shared CPs.

Note that when you define and activate a dynamic ICF configuration for a coupling facility LP, dynamic CF dispatching will be forced on any shared CPs in that coupling facility. You cannot use the DYNDISP command to turn it off.

Enhanced Dynamic ICF Expansion Across ICFs

Important Note

Enhanced Dynamic ICF Expansion across ICFs is not recommended for coupling facilities that are used for primary production workload. All processors assigned to the coupling facility partition should be dedicated to that logical partition if it is used for primary production workload.

An alternative dynamic ICF expansion configuration is supported where a CF logical partition can be defined with a mixture of both dedicated and shared ICF processors. This has all the same operating characteristics of the original dynamic ICF expansion configuration except that the shared processors are in the ICF pool rather than the general purpose processor pool. These shared ICF processors may be used by other CF logical partitions using only shared ICF processors and/or other CF logical partitions using a mixture of dedicated and shared ICF processors.

The following table lists the desired actions for all configurations involving coupling facility control code. It covers two areas: forcing of dynamic CF dispatching and defaulting of dynamic CF dispatching. It is recommended that anytime a shared processor is used for a coupling facility, whether standalone or ICF processors, that DCFD be enabled by insuring that the DYNDISP coupling facility control command be set to ON. See "Coupling Facility Control Code Commands" on page 2-50.

Coupling Facility Configuration	All Dedicated CPs or ICF(s)	Dedicated ICF(s) and Shared CP(s)	Dedicated ICF(s) and Shared ICF(s)	Shared ICF	Shared CP	Shared ICF
CF Model	—	—	—	System z9 with only ICF processors	System z9 with 1 or more general purpose CPs	System z9 with 1 or more general purpose CPs
Dynamic Dispatching Default Value	Off	On	On	Off	On	Off
Dynamic Dispatching Forced?	No	Yes	Yes	No	No	No

System-Managed Coupling Facility Structure Duplexing

A set of parallel sysplex architectural extensions is provided for support of *system-managed* duplexing of coupling facility structures for high availability. All three structure types, cache, list, and locking, can be duplexed using this architecture.

Benefits of system-managed CF structure duplexing include:

- **Availability:** Faster recovery of structures by having the data already in the second CF.
- **Manageability and Usability:** A consistent procedure to set up and manage structure recovery across multiple exploiters
- **Cost Benefits:** Enables the use of non-standalone CFs (for example, ICFs) for all resource sharing and data sharing environments.

Preparations for CF duplexing include the requirement to connect coupling facilities to one another with coupling links. The required CF-to-CF connectivity is bi-directional, so that signals may be exchanged between the CFs in both directions. A single peer-mode coupling link between each pair of CFs can provide the required CF-to-CF connectivity; however, for high availability at least two peer-mode links between each pair of CFs are recommended.

Note that while peer-mode CHPIDs cannot be shared between multiple coupling facility images, they can be shared between a single coupling facility image and one or more z/OS images. With this sharing capability, a single peer-mode ISC-3 or ICB link can serve both as a z/OS-to-CF link and as a CF-to-CF link, and thus provide the connectivity between z/OS and CF partitions as well as the connectivity between CFs that is required for CF Duplexing. Of course, at least two such links are recommended for high availability. In addition IFB links can provide the ability to actually share the same physical link between multiple CF images. By defining multiple CHPIDs on the same physical IFB link, the individual CHPIDs can be defined for a single CF image while the physical link is being shared by multiple CF images.

Single CPC Software Availability Sysplex

For single CPC configurations, System z9 models can utilize an ICF to form a single CPC sysplex, providing significant improvement in software continuous operations characteristics when running two z/OS LPs in data-sharing mode versus one large z/OS image. For these configurations, overall RAS is improved over that provided by a single z/OS image solution. Hardware failures can take down the entire single CPC sysplex, but those failures are far less frequent than conditions

taking down a software image, and planned software outages are the predominant form of software image outages in any case. Forming a single CPC sysplex allows software updates to occur in a “rolling” IPL fashion, maintaining system availability throughout. An LPAR cluster is one example of a single CPC sysplex which has significantly improved system availability over a single LP. For additional benefits provided by an LPAR cluster using IRD technology, see redbook *z/OS Intelligent Resource Director*, SG24-5952.

Coupling Facility Nonvolatility

Continuous availability of the transaction processing workload in a coupling facility configuration requires continuous availability of the shared structures in the coupling facility. To help ensure this, you must provide an optional backup power supply to make coupling facility storage contents nonvolatile across utility power failures.

Nonvolatility Choices

The following table indicates the optional nonvolatility choices available and their capabilities:

Table 2-5. Nonvolatility Choices for Coupling Facility LPs

Nonvolatility Choices	z9 BC	z9 EC
Uninterruptible power supply (UPS) (See Note 1)	Yes	Yes
Internal Battery Feature (IBF)	Yes	Yes
Local Uninterruptible Power Supply (LUPS) (See Note 2)	Yes	Yes
Notes: 1. Optional uninterruptible power supply (UPS) provides a secondary power source for use during extended utility power outages allowing continuous coupling facility operation. 2. The optional Local Uninterruptible Power Supply supports 0 to 18 minutes of full power operation.		

Setting the Conditions for Monitoring Coupling Facility Nonvolatility Status

In addition to installing an optional backup power supply to help ensure continuous availability, you **must** also set the conditions by which the coupling facility determines its volatility status. Software subsystems with structures defined in the coupling facility can monitor this status. Use the coupling facility control code MODE command as follows:

- MODE NONVOLATILE sets coupling facility volatility status to nonvolatile and should be used if a floor UPS is available to the CPC. Coupling facility control code does **not** monitor the installation or availability of UPS but maintains a nonvolatile status for the coupling facility.
- MODE VOLATILE sets coupling facility volatility status to volatile and should be used if no backup power supply is installed and available. Coupling facility control code maintains volatile status for the coupling facility even if a backup power supply is installed and available.

The coupling facility MODE setting is saved across power-on reset and activation of the coupling facility. You can use online help from the Operator Messages panel to get additional information on coupling facility control code commands.

Coupling Facility Mode Setting

The following table summarizes the relationship between the coupling facility MODE setting, and the resulting conditions you can expect if utility power fails at your site.

Table 2-6. Coupling Facility Mode Setting

CF MODE Setting	Local UPS or IBF Installed	Results on Utility Power Failure
VOLATILE	Yes	Ride out utility power failure on UPS/IBF.
VOLATILE	No	Machine down unless alternate floor level UPS/IBF provided.
NONVOLATILE	Yes	Ride out utility power failure on UPS/IBF.
NONVOLATILE	No	Machine down unless alternate floor level UPS/IBF provided. Note: This is the recommended setting when providing floor-wide UPS/IBF backup.
Note: Reflects the real time status of the power (volatile or nonvolatile).		

Coupling Facility LP Definition Considerations

You can define coupling facility mode for an LP at the Hardware Management Console or Support Element console using the Activation Profiles task available from the CPC Operational Customization Tasks Area.

You can define coupling facility LPs with shared or dedicated CPs on all System z9 models. Coupling facility LPs must be defined with at least 128 MB of central storage. See Table 3-2 on page 3-8.

Coupling facility LPs do **not** support some LP definition controls typically available to other LPs. For coupling facility LPs, you **cannot** define:

- Reserved central storage (coupling facility LPs do **not** support dynamic storage reconfiguration)
- Expanded storage initial or reserved amounts or expanded storage origin
- Crypto Express2 Coprocessors or Crypto Express2 Accelerators
- Automatic load
- Automatic load address
- Automatic load parameters

Internal Coupling Facility (ICF)

You can install one or more internal coupling facility (ICF) features. See “Considerations for Coupling Facilities using Internal Coupling Facility (ICF) CPs” on page 3-31.

Coupling Facility LP Storage Planning Considerations

You must define at least 128 MB of central storage for a coupling facility LP to activate.

This storage is reserved for coupling facility control code use and **cannot** be used for other purposes. Minimum storage size for coupling facilities is **primarily** a function of the coupling facility control code level. This implies that, over time, the minimum storage size required by a coupling facility on a particular machine can grow as new coupling facility control code updates are applied.

You must also define additional storage to accommodate the shared structures and dump space used by software subsystems using the coupling facility.

If a hot standby back-up coupling facility is defined, you must allocate to the back-up coupling facility all the storage required by that coupling facility to handle the structures it may need to hold in an active take-over role. Storage resources for coupling facility LPs are static and must be allocated at LP activation time of the coupling facility LP. This includes coupling facility LPs using ICF or dynamic coupling facility dispatching, or dynamic ICF expansion.

Structures, Dump Space, and Coupling Facility LP Storage

A coupling facility allocates storage for structures and for dump space based on values specified in the SIZE, INITSIZE, and DUMPSPACE parameters of the coupling facility resource management (CFRM) policy used for the coupling facility.

Structures consist of control objects and data elements. The control objects include entries and a variety of other control structures used to manipulate the entries. The data elements store user data associated with structure entries.

Dump space is storage in the coupling facility set aside for use as a dump table when a structure dump is taken. Dump tables are used for application development and problem determination purposes.

The following table indicates the maximum central storage possible for the System z9 models:

Table 2-7. Coupling Facility LP Storage Definition Capabilities

Models	Maximum Central Storage (GB)
System z9 BC	
Model R07	64
Model S07	64
System z9 EC	
Model S08	128
Model S18	256
Model S28	384
Model S38	512
Model S54	512

Estimating Coupling Facility Structure Sizes

Estimating coupling facility structure sizes is useful to system programmers to help ensure that there is enough coupling facility storage to meet application needs. The estimation of the minimum central storage requirements of your List or Cache structure has been superseded by the following two commands used by the CFSizer Utility to make space calculations in the coupling facility:

- **Compute List-Structure parameters:** For computing the minimum central storage requirements for a List structure. (What z/OS calls *lock structures* are actually a special form of list structure.)
- **Compute Cache-Structure parameters:** For computing the minimum central storage requirements for a Cache structure.

IMPORTANT NOTE

When implementing a new CFLEVEL in your configuration, redetermine the size of List and Cache structures using the CFSizer Utility and update the CFRM policy with the newly acquired values. For more information on the CFSizer Utility, see the following link:

<http://www-1.ibm.com/servers/eserver/zseries/cfsizer/>

Dump Space Allocation in a Coupling Facility

Dump space is storage you define using the DUMPSPACE parameter in the coupling facility resource management (CFRM) policy. It is set aside for the creation of dump tables. Dump tables are portions or snapshots of a structure typically saved for application development or problem determination purposes. The coupling facility allocates dump space in multiples of the coupling facility storage increment.

Dump tables for several different structures can exist in dump space at the same time. The amount of storage in any one dump table depends on the following factors:

- Amount of information you want to save to a dump table
The software subsystem can request the portions of a structure that are to be captured in the dump table. For example, lock tables, lists within a list structure, or directory entries belonging to particular storage classes or castout classes.
- Free dump space

Structures share dump space. If a structure is using some of the dump space, other structures cannot use that portion of the dump space until it is released.

- Characteristics of the structure saved to a dump table

When saving structure objects to a dump table, the amount of dump space used depends on the parameters specified for the structure. For example, list entry size, the number of list entries in a specified list, the number of directory entries in a castout class, whether adjunct data is included, etc.

The coupling facility can return the maximum requested dump space value. This indicates the largest amount of dump space requested by the software subsystems using the coupling facility. This value allows you to adjust the amount of allocated dump space to better match actual usage.

Coupling Facility LP Activation Considerations

At LP activation, coupling facility control code automatically loads into the coupling facility LP from support element hard disk. No initial program load (IPL) of an operating system is necessary or supported in the coupling facility LP.

All coupling facility channel path types targeted to be brought online will automatically be configured online.

Note: All channel paths types that are targeted to be brought online will automatically be configured online if the coupling facility LP is redefined as an ESA mode LP.

Coupling Facility Shutdown Considerations

IMPORTANT NOTE

It is important to properly remove all structures from a coupling facility that will be permanently taken out of the sysplex prior to shutting down the coupling facility. Failure to remove all structures may result in a pending condition (or transitioning structure) when attempting to allocate structures in a new coupling facility. Reactivation of the old coupling facility may be required to resolve the condition.

A running coupling facility LP can contain structures and data important to a sysplex. Make sure that you take proper precautions to preserve these structures and data **prior to** a power off, power-on reset (POR), LP deactivation, or shutdown (using the coupling facility control code SHUTDOWN command) of the coupling facility LP. Coupling facility structures and data will **not** survive any of these actions. Loss of coupling facility structures and data will occur regardless of the volatility state of the coupling facility including situations where a battery backup is in place.

For more information on removing, replacing, or shutting down a coupling facility, see *z/OS MVS Setting Up a Sysplex*, SA22-7625.

Coupling Facility LP Operation Considerations

Coupling facility control code runs in the coupling facility LP with minimal operator intervention. Operator activity is confined to the Operating Systems Messages task. PR/SM LPAR limits the hardware operator controls usually available for LPs to avoid unnecessary operator activity.

Coupling facility LPs only support the following tasks typically available to the hardware console operator:

- Activate
- Deactivate
- Operating System Messages

Coupling Facility Control Code Commands

Coupling facility control code does provide a limited set of hardware operator controls unique to the coupling facility LP. These controls are available from the Operating System Messages panel. From this panel, you can enter the HELP command to display coupling facility control code command syntax.

Coupling facility control code provides the following commands for use in the coupling facility LP:

- CONFIGURE (configure coupling facility channel paths online or offline)
- CP (configure a central processor online or offline)
- DISPLAY (display coupling facility resource information)
- DYNDISP (turn dynamic coupling facility dispatching on or off for a coupling facility LP). See “Dynamic Coupling Facility Dispatching (DCFD)” on page 2-42.
- HELP (display coupling facility control code command syntax)
- MODE (define coupling facility volatility mode)
- SHUTDOWN (shutdown coupling facility operation)
- TIMEZONE (sets timezone offset from Greenwich Mean Time for a coupling facility)
- TRACE (sets or modifies tracing options)

Notes:

1. Support for the CP and HELP coupling facility control code commands is available on all System z9 models.
2. The settings established using DYNDISP, MODE and TIMEZONE commands are recorded in the policy file for the coupling facility. As a result, all values are persistent across resets, deactivations, and reactivations.

Coupling Facility Level (CFLEVEL) Considerations

To support migration from one coupling facility level to the next, you can run different levels of the coupling facility concurrently as long as the coupling facility LPs are running on different CPCs. CF LPs running on the same CPC share the same coupling facility control code EC level.

When migrating CF levels, lock, list, and cache structure sizes might need to be increased to support new function. The amount of space needed for the current CFCC levels must be redetermined by visiting the CFSizer tool at <http://www-1.ibm.com/servers/eserver/zseries/cfsizer/>.

CPC Support for Coupling Facility Code Levels

The following table summarizes CPC support for the different coupling facility code levels.

Table 2-8. CPC Support for Coupling Facility Code Levels

CPC Models	Coupling Facility Code Level							
	Level 15	Level 14	Level 13	Level 12	Level 11	Level 10	Level 9	Level 8
2096 z9 BC	EC G40953 (Ver 2.9.2) MCL 006	EC J99670 (Ver 2.9.1) MCL 008	NA	NA	NA	NA	NA	NA
2094 z9 EC	EC G40953 (Ver 2.9.2) MCL 006	EC J99670 (Ver 2.9.0) MCL 008	NA	NA	NA	NA	NA	NA
2084 z990	NA	EC J13481 (Ver 1.8.2) MCL 025	EC J13481 (Ver 1.8.2) MCL 003	EC J12555 (Ver 1.8.0) MCL 009	NA	NA	NA	NA
2086 z890	NA	EC J13481 (Ver 1.8.2) MCL 025	EC J13481 (Ver 1.8.2) MCL 003	NA	NA	NA	NA	NA
2066 z800	NA	NA	EC J11207 (Ver 1.7.3) MCL 029	EC J11207 (Ver 1.7.3) MCL 022	NA	NA	NA	NA
2064 z900	NA	NA	EC J11207 (Ver 1.7.3) MCL 029	EC J11207 (Ver 1.7.3) MCL 022	NA	EC J10633 (Ver 1.7.2) MCL 012	EC H25496 (Ver 1.7.1) MCL 005 EC H25106 (Ver 1.7.0) MCL 002	NA
9672 X7/Z7	NA	NA	NA	NA	EC F99908 (Ver 1.6.2) MCL 020	NA	EC F99908 (Ver 1.6.2) MCL 018	EC F12017 (Ver 1.6.0) MCL 007
9672 R06 9672 R6	NA	NA	NA	NA	EC F99908 (Ver 1.6.2) MCL 020	NA	EC F99908 (Ver 1.6.2) MCL 018 EC F12017 (Ver 1.6.0) MCL 007	EC F12017 (Ver 1.6.0) MCL 006

Notes:

1. For Coupling Facility Code Levels 0 through 6, refer to Table 2-9 in the *S/390 Processor Resource/Systems Manager Planning Guide*, order no. GA22-7236.
2. Previous machine types that support the coupling facility are included for completeness and sysplex connectivity purposes.
3. For 9672 Model R06, EC numbers refer to the support element (SE) CFCC EC number. For all other models, the EC number refers to the SE D partition EC number.
4. The (Ver n.n.n) refers to the version of code installed in the support element. You can verify what version code you have by looking at the title bar on the Workplace window of your support element.
5. All MCLs cited in the table are current as of the publication date of this edition and can be superseded by later MCLs. The MCLs cited in the table are the latest service levels as of the current publication date of this edition. They do not necessarily reflect the minimum service level required for a particular

CFLEVEL but rather the recommended service level. For the latest EC and MCL information, use IBM Service Link to view the appropriate PSP bucket subset ID for hardware and software maintenance information. The most current CPC Support for CFCC table can be found online at <http://www1.s390.ibm.com/products/psocftable.html>.

Level 15 Coupling Facility

A level 15 coupling facility (CFLEVEL=15) provides the following enhancement:

- Support for 112 tasks in the CF
- Storage increment size change to 512K

Note: Considerations for structure re-sizing should be taken into account when moving to a CFLevel 15 Coupling Facility. It is highly recommended to use the CFSIZER tool for this.

CPC Support: See Table 2-8 on page 2-51 for a listing of the CPCs that support a level 15 coupling facility.

Level 14 Coupling Facility

A level 14 coupling facility (CFLEVEL=14) provides the following enhancement:

- CFCC Dispatcher Enhancements

CPC Support: See Table 2-8 on page 2-51 for a listing of the CPCs that support a level 14 coupling facility.

Level 13 Coupling Facility

A level 13 coupling facility (CFLEVEL=13) provides the following enhancement:

- Performance enhancements for Castout

CPC Support: See Table 2-8 on page 2-51 for a listing of the CPCs that support a level 13 coupling facility.

Software Corequisites: For a list of the software levels that exploit the function and levels that can coexist with CFLEVEL=13, see the “Summary of CFLEVEL Functions” section of the *z/OS MVS Setting Up a Sysplex* document.

Level 12 Coupling Facility

Enhancements: A level 12 coupling facility (CFLEVEL=12) provides the following enhancements:

- 64-bit CFCC support; removal of the 2 GB control store limit, and removal of the distinction between control and non-control store
- 48 concurrent tasks
- Support for Message Timing Ordering providing enhanced scalability of parallel sysplex.
- System-Managed Coupling Facility (CF) Duplexing
- Cache structure architecture extensions for Batch Write, Castout, and Cross-Invalidate functions
- Support for > 15 LPs
- Support for MCSS

CPC Support: See Table 2-8 on page 2-51 for a listing of the CPCs that support a level 12 coupling facility.

Software Corequisites: For a list of the software levels that exploit the function and levels that can coexist with CFLEVEL=12, see the “Summary of CFLEVEL Functions” section of the *z/OS MVS Setting Up a Sysplex* document.

Level 11 Coupling Facility

A level 11 coupling facility (CFLEVEL=11) provides the following enhancements:

- System-Managed Coupling Facility (CF) Duplexing
- Support for >15 LPARS.

CPC Support: See Table 2-8 on page 2-51 for a listing of the CPCs that support a level 11 coupling facility.

Software Corequisites: For a list of the software levels that exploit the new function and levels that can coexist with CFLEVEL=11, see the “Summary of CFLEVEL Functions” section of the *z/OS MVS Setting Up a Sysplex* document.

Level 10 Coupling Facility

Enhancements: A level 10 coupling facility (CFLEVEL=10) provides the following enhancements:

- z900 GA2 support.

CPC Support: See Table 2-8 on page 2-51 for a listing of the CPCs that support a level 10 coupling facility.

Software Corequisites: For a list of the software levels that exploit the new function and levels that can coexist with CFLEVEL=10, see the “Summary of CFLEVEL Functions” section of the *z/OS MVS Setting Up a Sysplex* document.

Level 9 Coupling Facility

Enhancements: A level 9 coupling facility (CFLEVEL=9) provides the following enhancements:

- Support for WorkLoad Manager (WLM) Multisystem Enclaves. This provides the capability to manage and report on parallel work requests that are executed on multiple z/OS images.
- XES Coupling Facility (CF) List Structure Architecture Extensions [also known as Message and Queuing (MQ) Series].

CPC Support: See Table 2-8 on page 2-51 for a listing of the CPCs that support a level 9 coupling facility.

Software Corequisites: For a list of the software levels that exploit the new function and levels that can coexist with CFLEVEL=9, see the “Summary of CFLEVEL Functions” section of the *z/OS MVS Setting Up a Sysplex* document.

Level 8 Coupling Facility

Enhancements: A level 8 coupling facility (CFLEVEL=8) provides the following enhancements:

- A system-managed rebuild capability for list, cache, and lock structures whose users do not support user-managed rebuild, with the object of moving a structure’s contents from one structure instance to another as quickly as possible. Also, two new commands (Compute List-Structure parameters and Compute Cache-Structure parameters) have been added to replace the coupling facility storage allocation formulas.

CPC Support: See Table 2-8 on page 2-51 for a listing of the CPCs that support a level 8 coupling facility.

Software Corequisites: For a list of the software levels that exploit the new function and levels that can coexist with CFLEVEL=8, see the “Summary of CFLEVEL Functions” section of the *z/OS MVS Setting Up a Sysplex* document.

Coupling Facility Resource Management (CFRM) Policy Considerations

To define how to manage z/OS images and coupling facilities in the sysplex, you must specify hardware configuration information in the coupling facility resource management (CFRM) policy as follows:

- Coupling facility node descriptor information

You must identify each coupling facility in the sysplex and the processor complex on which it is running. To do so, you must specify the following information in the CFRM policy:

CFRM Parameter	Description
PLANT	Plant of manufacture
SEQUENCE	Machine sequence number
SIDE	Machine side
TYPE	Machine type
MFG	Manufacturer
CPCID	CPC identifier

This information is available on the CPC Details panel. You can access the CPC Details panel by opening the CPC object that is running the coupling facility LP.

- LP information for the coupling facility

For a coupling facility residing on a System z9 model, the partition ID specified on the activation profile for the CF image on the support element or hardware master console must match the number specified in the PARTITION keyword of the CF statement in the policy information defined in the CFRM policy. IBM recommends that the LP names for the CF LPs in IOCP input files match the names used in the NAME keyword in the CF statement in the CFRM policy.

You can find the LP names in either the IOCP or HCD reports.

Coupling Facility Channels

Coupling facility channels are System z9 channels that use either fiber optic cables (CFP channel paths) or the channel subsystem's self-timed interface (STI) cables (CBP channel paths), an InfiniBand host channel adapter (HCA) (CIB channel paths), or internal memory bus (ICP channel paths) to provide the connectivity for data sharing between a coupling facility and the central processor complexes (CPCs) or logical partitions (LPs) directly attached to it.

The class of CHPIDs, known as peer mode channels, provide both sender and receiver capability on the same link. Peer mode links come in three varieties: Internal Coupling Channels (TYPE=ICP), Integrated Cluster Bus Channels (TYPE=CBP), InfiniBand (TYPE=CIB), and HiPerLinks (TYPE=CFP). Each ICP, CBP, and CFP channel can be configured as an unshared dedicated channel path to a single coupling facility and at the same time, as a shared channel path among several z/OS images and one coupling facility image.

Note: The following bulleted items only describe z/OS to coupling facility connections. However, they also apply to coupling facility duplexing connections (CF to CF).

Coupling facility channels:

- Require a point-to-point connection (direct channel attach between a CPC or LP and a coupling facility). Internal Coupling channels can only be used to connect a coupling facility and LPs on the same CPC.
- Can be used to connect a coupling facility to other LPs on the same CPC when a coupling facility is one of multiple LPs running on a single CPC. Internal coupling channels should be used for these connections.
- Can be redundantly configured (two or more coupling facility channels from each CPC involved in coupling facility data sharing) to enhance availability and avoid extended recovery time. This does not apply to Internal Coupling channels.
- Require ICP, CBP, CIB, or CFP channel path definition at the coupling facility end of a coupling facility channel connection.
- Require ICP, CBP, CIB, or CFP channel path definition at the z/OS (and, for System-Managed CF Structure Duplexing, the coupling facility) end of a coupling facility channel connection.
- Require a CBP, CFP, CIB, or ICP channel path be connected to a CBP, CFP, CIB, or ICP channel path, respectively. You must define the ICP channel paths in pairs and you must connect each pair. You connect an ICP channel path to an ICP channel path by specifying the CPATH keyword on the CHPID statement for every ICP channel path.

Internal Coupling Channel

The Internal Coupling channel emulates the coupling facility functions in LIC between images within a single system. Internal Coupling channel implementation is totally logical requiring no channel or even cable hardware. However, a CHPID number must be defined in the IOCDS. Internal Coupling channels cannot be used for coupling connections to images in external systems.

Partitions with Internal Coupling channels can also have coupling facility channels which allow external system coupling. ICs, which utilize the system bus, are extremely fast (approximately 6 GB/second).

Internal Coupling channels will have channel path type of ICP (Internal Coupling Peer). Internal Coupling channels are identified by 2 CHPIDs representing the two ends of a coupling link. The rules that apply to the ICP CHPID type are the same as those which apply to external coupling link types, with the exception that the following functions are not supported:

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

Internal coupling channels have improved coupling performance over ICBs and coupling facility channels.

Integrated Cluster Bus Channels

Integrated cluster bus channels (ICBs) are coupling facility channels that use the channel subsystem's self-timed interface (STI) cables to provide the connectivity for data sharing between a coupling facility and the CPCs or LPs directly attached to it. This allows an integrated cluster bus channel to use most of the high bandwidth of the STI bus.

Additionally, the improved data rate made possible by the integrated cluster bus channel's superior bandwidth enhances performance by reducing the latency of synchronous send messages and CP idle time which results in improved workload throughput.

CPCs that use integrated cluster bus channels can be no more than approximately 7 meters apart.

Integrated cluster bus channels:

- Require CBP channel path definition (coupling facility peer channel capability) at the coupling facility end of a coupling facility channel connection. CBP channel paths are defined for ICB-3 and ICB-4 links.
- Require CBP channel path definition (coupling facility peer channel capability) at the z/OS (and, for System-Managed CF Structure Duplexing, the coupling facility) end of a coupling facility channel connection. CBP channel paths are defined for ICB-3 and ICB-4 links.
- Each CBP channel path must be connected to another CBP channel path. ICB-3 and ICB-4 links are both defined as CBP channel paths. However, it is further required that an ICB-3 link be connected to another ICB-3 link and that an ICB-4 link be connected to another ICB-4 link.

Summary of CPC Support for Integrated Cluster Bus Channels: Integrated cluster bus channels are optional on System z9 models.

High-Performance Coupling Facility Channels (HiPerLinks)

System z9 high-performance coupling facility channels (HiPerLinks) are improved coupling facility channels (TYPE=CFP channel paths). HiPerLinks provide improved channel efficiency and response times in processing coupling facility requests. With HiPerLinks, coupling facility channel capacity is improved on average by up to twenty percent.

Existing CPCs that do not have HiPerLinks as their coupling facility channels can connect to a CPC that uses HiPerLinks.

InfiniBand host channel adapter (HCA)

An InfiniBand host channel adapter (HCA) can be used for a coupling link between a z/OS image and a coupling facility (CF) image. The IOP wraps the coupling messages in an InfiniBand packet and uses industry-standard InfiniBand link architecture to send the message. The IOP at the other end unwraps the coupling message and delivers it to the intended receiver.

InfiniBand coupling links for Parallel Sysplex

InfiniBand coupling link technology provides increased bandwidth at greater cable distances. The IBM System z10 supports a 12x (12 lanes of fibre in each direction) InfiniBand-Double Data Rate (IB-DDR) coupling link which is designed to support a total interface link data rate of 6 GigaBytes per second (GBps) in each direction. The maximum distance for this point-to-point link over fibre optic cabling is 150

meters (492 feet). This InfiniBand coupling link provides improved performance over the current ISC-3 coupling link in data centers where systems are less than 150 meters apart.

A 12x InfiniBand-Single Data Rate (IB-SDR) coupling link is available on System z9 EC and System z9 BC S07 servers configured as internal coupling facilities (ICFs) only. This coupling link is designed to support a total interface link data rate of 3 GigaBytes per second (GBps) in each direction. This InfiniBand coupling link provides improved performance over the current ISC-3 coupling link in data centers where systems are less than 150 meters apart. When a System z10 server is connected to a System z9 server using point-to-point InfiniBand cabling, the link auto-negotiates to the highest common data rate of 3 GBps.

InfiniBand coupling links also provide the ability to define up to 16 CHPIDs on a single HCA adapter, allowing physical coupling links to be shared by multiple sysplexes. This also provides additional subchannels for Coupling Facility communication, improving scalability, and reducing contention in heavily utilized system configurations. It also allows for one CHPID to be directed to one CF, and another CHPID directed to another CF on the same target server, using the same physical link. For detailed information about InfiniBand coupling links, see *Getting Started with InfiniBand on System z10 and System z9*, SG24-7539, available on the Redbooks Web site: <http://www.redbooks.ibm.com/>.

InfiniBand coupling links:

- Require CIB channel path definition (coupling facility peer channel capability) at the coupling facility end of a coupling facility channel connection. CIB channel paths are defined for IFB links.
- Require CIB channel path definition (coupling facility peer channel capability) at the z/OS (and, for System-Managed CF Structure Duplexing, the coupling facility) end of a coupling facility channel connection. CIB channel paths are defined for IFB links.
- Each CIB channel path must be connected to another CIB channel path. IFB links are defined as CIB channel paths. However, it is further required that an IFB link be connected to another IFB link.
- When defining CIB CHPIDs, you must specify both the name of the system (CSYSTEM keyword) this CHPID will connect to and the CSS and CHPID (CPATH keyword) on that system that this CHPID will connect to.

Coupling Facility Channels (TYPE=CFP, TYPE=CBP, TYPE=CIB, or TYPE=ICP)

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

- An unshared dedicated channel path to a single LP
- An unshared reconfigurable channel path that can be configured to only one LP at a time but which can be dynamically moved to another LP by channel path reconfiguration commands
- A shared channel path that can be shared between at most one coupling facility image and one or more z/OS images (CFP, CBP, and ICP only).

Shared Coupling Facility Channel Path Recommendations: IBM recommends the following:

1. For shared coupling facility channel paths, make sure that only LPs that need to use the channel path have it configured online. This eliminates unnecessary traffic on the channel path from those systems that have it online but do not have the attached coupling facility in the active CFRM policy.

2. These channel paths can result in 'Path Busy' conditions when another LP is using the path. This can result in delays in getting requests to the coupling facility on this path. The number of 'Path Busy' conditions can be found in the RMF CF Subchannel Activity report in the BUSY COUNTS column labelled PTH. As a guideline, if this count exceeds 10% of the total requests, you should consider **not** sharing the channel path or adding additional coupling facility channel paths.

Defining Internal Coupling Channels (TYPE=ICP)

Internal Coupling channels 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.

It is suggested that you define a minimum of internal coupling channels. For most customers, IBM suggests defining just one pair of ICP channel paths for each coupling facility logical partition (LP) in your configuration. For instance, if your general purpose configuration has several ESA LPs and one CF LP, you would define one pair of connected ICP CHPIDs shared by all the LPs in your configuration. If your configuration has several ESA LPs and two CF LPs, you still would only define one connected pair of ICP CHPIDs, but one ICP CHPID should be defined as shared by the ESA LPs and one CF LP while the other ICP CHPID is defined as shared by the ESA LPs and the other CF LP. Both of these examples best exploit the peer capabilities of these coupling channels by using the *sending* and *receiving* buffers of both channels.

You must define ICP CHPIDs in pairs, and you must connect each pair. A connected pair of ICP CHPIDs is called an *internal coupling link*. Both ends of the internal coupling link must be specified by defining which ICPs are to communicate with each other. Use the CPATH keywords in the CHPID statement to connect internal coupling CHPIDs (see page 2-29).

Maximum Recommended Number of ICP CHPIDs: Real CPU resources are used to implement the link function of connected ICP CHPIDs. Production environments should limit the maximum number of internal coupling links that are defined for a CPC to optimize the internal coupling link function utilization of CPU resources. This maximum number of internal coupling links is based on the number of available physical CPs on the CPC. This maximum number of internal coupling links can be calculated by taking the number of CPs in the CPC that are used for general purpose CPs and for ICF CPs, and subtracting one from that total. For example: a CPC that consists of four general purpose CPs and two ICF CPs would have a maximum five $[(4 + 2) - 1 = 5]$ internal coupling links recommended. This represents a maximum total of ten ICP CHPIDs being defined.

I/O Configuration Considerations

ICP IOCP supports coupling facility channel path definition on the System z9 CPC.

With z/OS, HCD provides controls for defining coupling facility channels. HCD also automatically generates the control unit and device definitions associated with CFP, CBP, CIB, or ICP channel paths.

IBM recommends that you use the Hardware Configuration Definition (HCD), when possible, to define the coupling facility channel configuration to the channel subsystem.

Considerations when Migrating from ICMF to ICs

On previous CPCs, PR/SM LPAR provided an Integrated Coupling Migration Facility (ICMF) which is not supported on System z9.

- Use HCD to:
 - Define one shared ICP CHPID that is available for all the ESA logical partitions that were communicating with the *first* or only coupling facility logical partition using ICMF.
 - If there is a second coupling facility using ICMF being migrated, make this ICP CHPID available to the *second* coupling facility logical partition as well.
 - If there is **not a second** coupling facility, make this ICP CHPID available to the first or only coupling facility as well.
 - Define a second shared ICP CHPID that is available for the *first* or only coupling facility partition using ICMF.
 - If there is a second coupling facility using ICMF that is being migrated, make this second ICP CHPID available to all the ESA logical partitions that were communicating with the *second* coupling facility logical partition using ICMF as well.
 - If there is **not a second** coupling facility, make this second ICP CHPID available to all the ESA logical partitions that were communicating with the *first* or only coupling facility logical partition using ICMF as well. For more information, please see “Defining Internal Coupling Channels (TYPE=ICP)” on page 2-58.)
 - Connect the two ICP CHPIDs together.

For the case above where there is only one coupling facility logical partition using ICMF being migrated, the result is that both ICP CHPIDs will be shared by all the ESA logical partitions and also the coupling facility logical partition that they are to communicate with. This is the correct configuration for a peer coupling link connected within the same CPC.

Linux Operating System Planning Considerations

Linux is an open operating system with a wealth of applications which, in most cases, can run on a System z9 with a simple recompile. The System z9 includes features that provide an extremely cost-effective environment in which to run Linux.

Integrated Facility for Linux (IFL)

On System z9 models, you can purchase and install one or more IFL features exclusively for Linux workloads (a single Linux image, or z/VM Version 5 with only Linux users) with no effect on the System z9 model designation. Consequently, no additional IBM operating system or middleware charges are incurred with the addition of this capacity unless that software is actually running in that additional capacity.

These features are ordered separately, and are distinguished at the hardware level from any ICFs, IFAs, or zIIPs. ICFs, IFLs, IFAs and zIIPs are perceived by the system as multiple resource pools.

With this feature, Linux, or z/VM Version 5 with only Linux users, runs on IFLs, which cannot be used to run other IBM operating systems such as z/OS, VSE, or TPF. Only logical partitions specified as Linux-Only Mode in their Activation profiles can be allocated IFLs. IFLs can be allocated as either dedicated or shared.

z/VM Version 5 Utilizing IFL Features

z/VM Version 5 utilizing IFL features provides an easy-to-use high-performance hypervisor that operates within a logical partition. It has the capability to create a significant number of Linux images. z/VM Version 5 creates and manages Linux images quickly and easily, providing the ability to share resources, and supports an arbitrary number of internal networks that can be used for high-speed communication among Linux images.

System z9 Application Assist Processor (zAAP)

The Application Assist Processor (zAAP) is available on the System z9. This specialized processing unit provides an economical Java execution environment for customers who desire the traditional Qualities of Service and the integration advantages of the System z9 platform. The System z9 Application Assist Processor is also known as an IFA (Integrated Facility for Applications). The term IFA appears in panels, messages, and other z/OS information relating to the System z9 Application Assist Processor, including this publication.

When configured with general purpose processors within logical partitions running z/OS, zAAPs may help increase general purpose processor productivity and may contribute to lowering the overall cost of computing for z/OS Java technology-based applications. zAAPs are designed to operate asynchronously with the general processors to execute Java programming under control of the IBM Java Virtual Machine (JVM). This can help reduce the demands and capacity requirements on general purpose processors which may then be available for reallocation to other System z9 workloads.

The IBM JVM processing cycles can be executed on the configured zAAPs with no anticipated modifications to the Java application(s). Execution of the JVM processing cycles on the zAAP is a function of the Software Developer's Kit (SDK) 1.4.1 for zSeries, z/OS 1.6 or later, and PR/SM.

The amount of general purpose processor savings will vary based on the amount of Java application code executed by zAAP(s). This is dependent upon the amount of Java cycles used by the relevant application(s) and on the zAAP execution mode selected by the customer.

Execution of the Java applications on zAAPs, within the same z/OS logical partition as their associated database subsystems, can also help simplify the server infrastructure and improve operational efficiencies. For example, use of zAAPs could reduce the number of TCP/IP programming stacks, firewalls, and physical interconnections (and their associated processing) that might otherwise be required when the application servers and their database servers are deployed on separate physical server platforms.

IBM does not impose software charges on zAAP capacity. Additional IBM software charges will apply when additional general purpose CP capacity is used.

Customers are encouraged to contact their specific ISVs/USVs directly to determine if their charges will be affected.

On System z9 models, you can purchase and install one or more Integrated Facility for Applications (IFA) features with no effect on the model designation. Consequently, no additional IBM operating system or middleware charges are incurred with the addition of this capacity unless that software is actually running in

that additional capacity. IFA features are ordered separately and are distinguished at the hardware level from any Internal Coupling Facility (ICF) processors, Integrated Features for Linux (IFLs), and System z9 Integrated Information Processor (zIIPs).

z/VM 5.3 can run in a logical partition that includes zAAPs and can dispatch z/OS guest virtual zAAPs on the real zAAPs. z/VM 5.3 can also simulate zAAPs for z/OS guests, dispatching virtual zAAPs on real general-purpose processors (CPs).

IBM System z9 Integrated Information Processor (zIIP)

The IBM System z9 Integrated Information Processor, or zIIP, is the latest customer-inspired specialty engine for System z9. It provides a cost-effective DB2 workload redirection environment. The zIIP is 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.

The IBM zIIP's execution environment will accept eligible work from z/OS V1.6 or 1.7, which will manage and direct the work between the general purpose processor and the zIIP. DB2 for z/OS V8 will exploit the zIIP capability for eligible workloads. IBM zIIPs are designed to free up general purpose capacity which may be utilized by other workloads.

In addition to improving the utilization of existing System z9 resources, the zIIP may help you to leverage the z/OS and DB2 for z/OS qualities of service for data access and information management across your enterprise. It does this by making direct access to DB2 more cost effective and potentially reducing the need for many local copies of the data and the complexity that brings.

On System z9 models, you can purchase and install one or more Integrated Facility for Applications (IFA) features with no effect on the model designation. Consequently, no additional IBM operating system or middleware charges are incurred with the addition of this capacity unless that software is actually running in that additional capacity. IFA features are ordered separately and are distinguished at the hardware level from any Internal Coupling Facility (ICF) processors, Integrated Features for Linux (IFLs), and System z9 Application Assist Processor (zAAPs).

z/VM 5.3 can run in a logical partition that includes zIIPs and can dispatch z/OS guest virtual zIIPs on the real zIIPs. z/VM 5.3 can also simulate zIIPs for z/OS guests, dispatching virtual zIIPs on real general-purpose processors (CPs).

Concurrent Patch

Concurrent patch is available on all System z9 models. It is possible to apply BPC, UPC, support element (SE), Hardware Management Console, channel Licensed Internal Code (LIC), LPAR, coupling facility control code, I390, and PU patches nondisruptively and concurrent with system operation. There can still be situations where a small percentage of patches are disruptive; however, with all System z9 models, all major LIC components now support concurrent patch.

Additionally, there is also support for multiple EC streams (one per major LIC component) further minimizing the number of disruptive patch sessions. On previous models, a single EC stream contained all major LIC components and provided a mandatory sequence for patch application. This could lead to situations where a disruptive patch belonging to one LIC component, for example, PU, could

prevent you from applying all nondisruptive SE patches if one or more of the SE patches came after this channel patch in the patch application sequence.

Patches for each major LIC component have their own sequence for patch application. This means that disruptive patches belonging to one LIC component no longer stand as an obstacle to the concurrent application of nondisruptive patches belonging to another LIC component as long as the patches in question are not otherwise defined as corequisite for each other.

CFCC Enhanced Patch Apply

The CFCC patch apply process is enhanced to eliminate the need for a Power On Reset (POR) of the System z9 to apply a disruptive CFCC patch. CFCC enhanced patch apply provides you the ability to:

- Selectively apply the new patch to one of possibly several CFs running on a System z9. 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 System z9, you 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 System z9 where a disruptive CFCC patch will be applied to continue to run without being impacted by the application of the disruptive CFCC patch.


The enhanced patch apply 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.

Note: Any POR of the servers will result in the disruptive patch being applied to CF partitions.

Dynamic Capacity Upgrade on Demand


The System z9 includes a function to dynamically increase the number of CPs, ICFs, IFLs, IFAs or zIIPs without an intervening IPL. A logical partition (LPAR) may be defined with both an initial and reserved amount of logical CPs. This enables a definition where the number of logical CPs for a logical partition is greater than the number of physical CPs installed on the model. These reserved CPs are automatically in a deconfigured state at partition activation and can be brought online at any future time via the SCP operator command if the requested resource is available. To prepare for a nondisruptive upgrade, a Logical Partition simply needs to be defined and activated in advance with an activation profile indicating reserved CPs. This helps ensure that any planned logical partition can be as large as the possible physical machine configuration, nondisruptively.

The following example assumes a nondisruptive concurrent CP upgrade from an 8-Way to a 9-Way Server.




PR/SM Shared Partitions

PR/SM configurations supporting multiple partitions, all sharing CPs, support concurrent CP upgrades. PR/SM code, once signaled that one or more central processors have been made available to the configuration, will vary them online automatically into the “shared pool” of physical CPs and begin full utilization of the added capacity. In the following example, three partitions sharing eight physical CPs are able to share the increased capacity resulting from a nondisruptive upgrade to a 9-way server without any other configuration changes. In the second upgrade scenario, Partition 1 is additionally changed from a 4-way to a 5-way partition nondisruptively. The preparation for this is straightforward and easy. Simply define and activate logical Partition 1 with four initial and one reserved logical CPs (see Figure 3-20 on page 3-63 for a similar example). At the time of the concurrent CP upgrade, the SCP operator command can be used in Partition 1 to configure a fifth CP online to the shared partition without interruption to any logical partition.




Mixed Shared and Dedicated PR/SM Partitions

As with configurations in which all PR/SM partitions share available CPs, those shared partitions within a mixed configuration also support concurrent CP upgrade. CPs are added, without disruption to any of the partitions, to the pool of physical CPs shared amongst the shared CP LPs. In addition, partitions configured with dedicated CPs in a mixed environment can add new CP capacity while both the shared CP partitions and the dedicated CP partition run uninterrupted. To prepare for this ability, in the following example simply define Partition 3 as a 3-way dedicated partition with two initial logical CPs and one reserved logical CP (see Figure 3-20 on page 3-63 for a similar example). The reserved logical CP is offline automatically at the time of partition activation. At the time of the concurrent CP upgrade, the SCP operator command can be used in the dedicated partition to configure a third CP online to the dedicated partition without interruption to any logical partition.




Multiple Dedicated PR/SM Partitions

Configurations in which all PR/SM partitions use dedicated CPs, where there is more than one dedicated partition, also support concurrent CP upgrade. CPs are added to the configuration, without disruption to any of the partitions, and can be brought online to a dedicated partition without an interruption. In the following example, all ICFs in the configuration are dedicated. The partition (ICF 1) is defined for the ability to be upgraded dynamically. To prepare for this ability, simply define ICF 1 as a 2-way dedicated partition with one initial and one reserved logical CP (see Figure 3-20 on page 3-63 for a similar example). At the time of the concurrent CP upgrade, the CF operator command can be used in the ICF 1 dedicated partition to configure a second ICF CP online to the dedicated partition without interruption to any logical partition. Note the same type of procedure can be accomplished with logical partitions using dedicated general purpose CPs.



Shared Internal Coupling Facility

Concurrent upgrade can be used to add a PU to a shared pool of PUs supporting existing Internal Coupling Facilities. In the following example, Partition 1 is defined for the ability to be upgraded dynamically. To prepare for this ability, simply define Partition 1 as a 7-way shared partition with six initial and one reserved logical CPs (see Figure 3-20 on page 3-63 for a similar example). At the time of the concurrent CP upgrade, the CF operator command can be used in Partition 1 shared partition to configure a seventh ICF CP online to the shared partition without interruption to any logical partition. Note that Partition 2 in this case could have also been defined as a 7-way shared ICF partition with one ICF configured offline. This would allow Partition 2 to grow concurrently without an outage as well.



Dynamic Capacity Upgrade on Demand Limitations

1. Valid upgrade towers are as follows starting at the smallest model in the pathway:

z9 EC Models

S08, S18, S28, S38, S54

z9 BC Models

R07, S07

Inactive (spare) PUs can be added concurrently, dynamically providing nondisruptive upgrade of processing capability. They can be characterized, in any combination, as CPs, ICFs IFLs, IFAs, or zIIPs. Neither the number of zAAPs nor zIIPs can exceed the number of CPs.

2. For shared CP PR/SM configurations, added CPs are brought into the pool of shared CPs, effectively increasing the number of physical CPs to be shared amongst partitions. To increase the number of logical CPs online to an LP, simply define the LP with both initial and reserved logical CPs. Then all you need to do is configure on the extra logical CP(s) after the concurrent upgrade.
3. All models have the ability to define logical partitions with as many reserved CPs as necessary. With thoughtful planning, there is never a need for a disruptive increase in the number of logical CPs
4. The maximum initially online logical CP width that logical partitions can be defined and activated with at any point in time is as follows (the maximum defined logical CP width is as great as the total number of CPs achievable via concurrent CPU upgrade):
 - Maximum initial logical CPs defined for a dedicated partition **equals** physical CPs online for the current model **minus** physical CPs currently dedicated and online to other dedicated partitions **minus** the maximum online number of shared CPs among all the activated logical partitions that are using shared CPs.
 - Maximum initially online logical CPs defined for a shared partition **equals** physical CPs online for the current model **minus** physical CPs currently dedicated and online to other partitions using dedicated CPs.
 - Maximum **total** logical CPs (including initial and reserved) for any partition **equals** the number of physical CPs achievable through concurrent CP upgrade.
 - When a logical partition is defined to use ICFs, IFLs, zIIPs or IFAs, these rules are applied against the installed processors of that type. In the case of IFAs and zIIPs, logical partitions can have a combination of general purpose processors and IFAs/zIIPs defined. These rules apply to each processor type defined to the partition independently. The total number of all logical processors defined to the logical partition cannot exceed the maximum supported for a logical partition by the CEC (32), independent of processor type.

Concurrent Memory Upgrade

The System z9 includes a function to dynamically increase the amount of configured storage. Concurrent Memory Upgrade allows for a memory upgrade without changing hardware or experiencing an outage, provided there is enough spare memory already existing on the memory cards. An IML is not required to use the previously unavailable storage. It is immediately available for allocation to logical partitions as either central or expanded storage. The new storage can be allocated to either newly activated logical partitions or to already active logical partitions by using dynamic storage reconfiguration (see “Dynamic Storage Reconfiguration” on page 3-12). In planning for a concurrent memory upgrade, logical partition activation profiles should be defined with storage (central and/or expanded) specifications that include a reserved as well as an initial storage amount. Following the completion of the concurrent memory upgrade operation, issue the z/OS command `CF STOR(E=1),ONLINE` or the z/OS command `CF ESTOR(E=1),ONLINE` to bring the new memory online to already active LPs for which a reserved central storage amount, or a reserved expanded storage amount (respectively) was specified.

Capacity Backup Upgrade (CBU) Capability

This orderable feature, available on z9 EC Models, can be integrated with Geographically Dispersed Parallel Sysplex™. This should reduce disaster recovery times by automating the Capacity Backup Upgrade process at a remote site.

Operationally, the planning considerations and limitations for Capacity Backup Upgrade, for mode purposes, are similar to those for Dynamic Capacity Upgrade on Demand. Planning is simple because reserved logical CPs can be specified for logical partitions on the backup system. The logical CPs can then be brought online, quickly and nondisruptively.

Concurrent Undo CBU is provided to dynamically remove from the configuration processors that had been added via Capacity Backup Upgrade. The server for disaster recovery with Capacity Backup Upgrade activated can now be restored to its base configuration without requiring a system outage.

When a disaster occurs, Capacity Backup Upgrade is intended to provide the extra capacity without disruption. When the disaster is over and normalcy is restored, Concurrent Undo CBU is intended to allow the system to be returned to its previous configuration without disruption.

Annual testing of a Capacity Backup Upgrade system is highly recommended. Invoking the Capacity Backup Upgrade configuration is nondisruptive and now returns to the original configuration, after the test, without disruption as well.

When a typical Capacity Backup Upgrade occurs, usually logical processors are configured online so that the system closely resembles the failing system. To prepare for Concurrent Undo CBU, the most expedient method is to simply configure offline all those logical processors that were configured online in support of the failing system.

In order for Concurrent Undo CBU to proceed, the restored original configuration must have a physical processor for each online dedicated logical processor. Theoretically, the only additional requirement is that at least one nondedicated physical processor remains that matches each type (general purpose, ICF, IFL, IFA or zIIP) of online shared logical processors. However, it is highly recommended that

shared logical processors be configured offline so that the highest number of online shared logical CPs for any active LP does not exceed the number of nondedicated physical CPs remaining. For further guidelines, see the rules governing the number of CPs that can be specified for an activating logical partition on page 3-28.

Enhanced Book Availability

The z9 EC is designed to allow a single book, in a multi-book server, to be concurrently removed from the server and reinstalled during an upgrade or repair action, while continuing to provide connectivity to the server I/O resources using a second path from a different book. To help minimize the impact on current workloads and applications, you should ensure that you have sufficient inactive physical resources on the remaining books to complete a book removal.

Enhanced book availability may also provide benefits should you choose not to configure for maximum availability. In these cases, you should have sufficient inactive resources on the remaining books to contain critical workloads while completing a book replacement. Contact your IBM representative to help you determine the appropriate configuration. With proper planning, you may be able to avoid planned outages when utilizing enhanced book availability.

Enhanced driver maintenance, exclusive to the z9 EC, is another step in reducing the duration of a planned outage. One of the greatest contributors to downtime during planned outages is Licensed Internal Code (LIC) updates performed in support of new features and functions. When properly configured, the z9 EC is designed to support activating select new LIC level concurrently. Concurrent activation of the select new LIC level is only supported at specific sync points (points in the maintenance process when LIC may be applied concurrently - MCL service level). Sync points may exist throughout the life of the current LIC level. Once a sync point has passed, you will be required to wait until the next sync point supporting concurrent activation of a new LIC level. Certain LIC updates will not be supported by this function.

Preparing for Enhanced Book Availability

This option will determine the readiness of the system for the targeted book. The configured processors and the in-use memory will be evaluated for evacuation from the targeted book to the unused resources available on the remaining books within the system configuration. In addition, the I/O connections associated with the targeted book will be analyzed for any Single Path I/O connectivity.

There are three states which can result from the prepare option:

- The system is ready to perform the Enhanced Book Availability for the targeted book with the original configuration.
- The system is not ready to perform the Enhanced Book Availability due to conditions noted from the prepare step. See “Getting the System Ready to Perform Enhanced Book Availability” for more details.
- The system is ready to perform the Enhanced Book Availability for the targeted book. However, processors were reassigned from the original configuration in order to continue. “Reassigning Non-Dedicated Processors” on page 2-69 for details.

Getting the System Ready to Perform Enhanced Book Availability

Review the conditions that are preventing the Enhanced book Availability option from being performed. There are tabs on the resulting panel for Processors, Memory, and for various Single Path I/O conditions. Only the tabs that have

conditions preventing the perform option from being executed will be displayed. Each tab will indicate what the specific conditions are and possible options to correct the conditions. Following is an example panel:




Figure 2-18. Example of a Prepare for Enhanced Book Availability Results Panel

In the example shown in Figure 2-18, the Processor tab displays, indicating the corrective actions suggested for the processor configuration. Following is a list of tabs that may appear for your particular configuration:

Processors

Use this tab to view the corrective actions required for the processor configuration conditions that are preventing the Perform Enhanced Book

Availability option from being performed for the targeted book. You may need to deactivate partitions or deconfigure processors to meet requirements as indicated by the panel data.

Memory

Use this tab to view the corrective actions required for the memory configuration conditions that are preventing the Perform Enhanced Book Availability option from being performed for the targeted book. You may need to deactivate partitions to meet requirements as indicated by the panel data. The In-Use memory must be less than or equal to the available memory on the remaining books within the system.

Single Path I/O

Use this tab to view the corrective actions required for the single I/O configuration conditions that are preventing the Perform Enhanced Book Availability option from being performed for the targeted book. You will need to deconfigure all of the PCHIDs that are indicated by the panel data.

Single Domain

Use this tab to view the corrective actions required for the single I/O domain configuration conditions that are preventing the Perform Enhanced Book Availability option from being performed for the targeted book. You will need to change the alternate path to a different book or deconfigure the PCHID.

Single Alternate Path

Use this tab to view the corrective actions required for the single I/O configuration conditions that are preventing the Perform Enhanced Book Availability option from being performed for the targeted book. You will need to correct the alternate path error condition or deconfigure the PCHIDs.

Reassigning Non-Dedicated Processors

Following is an example showing the Reassign Non-Dedicated Processors panel. Use this panel to change or accept the system processor assignments that are generated during the processing of the Prepare for Enhanced Book Availability option. The processor values that are entered from this panel will be the processor configuration utilized during the Perform Enhanced Book Availability processing.

WARNING

Note: The values should never be altered without approval from the system programmer.

Processor Type	Dedicated Count	Non-Dedicated Count	Processor Totals	LICCC Count
CPU	0	16	16	22
ICF	0	1	1	2
IFL	0	1	1	2
IFA	0	1	1	2
SAP	5	0	5	6
Available to use		0	0	
Remaining Book Totals	5	19	24	

Figure 2-19. Reassign Non-Dedicated Processors Panel

Customer Initiated Upgrade (CIU)

This feature is designed to allow timely response to sudden increased capacity requirements by downloading and automatically applying a System z9 processor and/or memory upgrade via the Web, using IBM Resource Link and the Remote Support Facility.

Operationally, the planning considerations and limitations for CIU are similar to those for Dynamic Capacity Upgrade on Demand and Concurrent Memory Upgrade. Planning is simple because reserved logical CPs can be specified for logical partition (LP)s in the original configuration. The logical CPs can then be brought online, quickly and nondisruptively at the completion of the Concurrent Upgrade. Similarly, a reserved central and/or expanded storage amount can be specified for participating LPs and brought online non-disruptively following the completion of the concurrent memory upgrade.

Concurrent Processor Unit Conversion

The System z9 supports concurrent conversion of different Processor Unit (PU) types. This capability is extended to Central Processor (CPs), Integrated Facility for Linux (IFLs), Integrated Facility for Applications (IFAs), System z9 Integrated Information Processors (zIIPs), and Internal Coupling Facility (ICFs) providing flexibility in configuring a System z9 to meet the changing business environments.

For example, a z9 EC Model S08 configured with eight CPs could be reconfigured to seven CPs and one IFL by ordering the appropriate PU conversion. This order will generate a new LIC CC which can be installed concurrently in two steps. First, one CP is removed from the z9 EC configuration. Second, the newly available PU is activated as an IFL. This two-step process assumes that one CP remains active during the conversion of other processor units.

Planning for Hot Plugging Crypto Features

In order to make the addition of Crypto Express2 Coprocessors (CEX2C) or Crypto Express2 Accelerators (CEX2A) nondisruptive, logical partitions must be configured with the appropriate domain indices and Cryptographic numbers when the partitions are activated (see the Cryptographic Candidate List information under “Parameter Descriptions” on page 3-77).

To maximize concurrent upgrade possibilities for Cryptographic functions, some preparations need to be made (These can be done even if there are no CEX2Cs or CEX2As currently installed on the machine):

- In the Image Activation Profile, cryptographic numbers should be defined as candidates for those logical partitions that may require access to the CEX2C or CEX2A. Specifically, the Usage Domain Index and the Cryptographic Candidate List need to be set up in the Image Activation Profiles to reflect the Crypto feature that is planned to be nondisruptively added.

Use the View LPAR Cryptographic Controls panel to display the current information on the activated partition.

- If the customer plans to use a TKE workstation, the Control Domain Index must be also set up in the Image Activation Profiles to reflect all the usage domains that can be managed from a TKE active in this logical partition.

Use the View LPAR Cryptographic Controls panel to display the information on the activated partition.

- If the customer plans to use ICSF or the optional cryptographic hardware, the CP Crypto Assist functions (CPACF DES/TDES) must be enabled. Many IBM products will take advantage of the cryptographic hardware using ICSF, so enabling CPACF is recommended. Use the CPC Details panel (double-click the CPC icon from the CPC Work Area) to determine if the CPACF feature is installed.

For more detailed information, see the *Support Element Operations Guide* and the *z/OS ICSF Trusted Key Entry PCIX Workstation User's Guide*.

It is important to remember that when hot plugging Crypto Express2, the default configuration of the card is coprocessor (CEX2C). To change the configuration of Crypto Express2, use the Cryptographic Configuration Panel.

Once the Crypto Express2 feature has been installed, and the previous preparations made, the CEX2C and/or CEX2A will be available to all partitions specifying their assigned Cryptographic numbers in the Candidate List. To bring the Cryptographic numbers online, perform either:

- Config On of the Crypto (nondisruptive), or
- An Activate of the partition (disruptive).

Note: The Cryptographic Online List needs to be set up in the Image Activation Profile to reflect the CEX2C and/or CEX2A that you plan to bring online during the partition activation. If the Cryptographic Candidate List does not reflect the CEX2C or CEX2A that you plan to use, then these Cryptos are not available to the logical partition even after this partition is reactivated. If the Crypto Online List is set up properly and the partition is reactivated, a Config On and an Activate from the ICSF Option 1, Coprocessor Management panel will cause the Cryptos to come online to the partition.

Chapter 3. Determining the Characteristics of Logical Partitions

Planning Overview	3-4
Performance Considerations.	3-4
Dedicated and Shared Central Processors (CPs)	3-4
Dedicated and Shared Channel Paths	3-4
ITR Performance	3-4
Capped Logical Partitions.	3-5
Recovery Considerations	3-5
Determining the Characteristics	3-5
Control Program Support	3-5
IOCDS Requirements	3-6
Logical Partition Identifier.	3-6
Mode of Operation	3-7
Storage Configurations.	3-7
Storage Resources	3-7
Single Storage Pool.	3-7
Central Storage	3-8
Initial Central Storage	3-8
Reserved Central Storage	3-9
Central Storage Origin	3-9
Expanded Storage	3-10
Initial Expanded Storage	3-10
Reserved Expanded Storage	3-11
Expanded Storage Origins	3-11
Dynamic Storage Reconfiguration	3-12
Central Storage Dynamic Storage Reconfiguration Examples	3-14
Expanded Storage Dynamic Storage Reconfiguration Examples	3-20
Recommendations for Storage Map Planning	3-25
CPCs with the Sysplex Failure Manager (SFM)	3-26
TARGETSYS(ALL) Examples	3-26
Number of Central Processors	3-28
Maximum Number of Central Processors	3-28
Workload Requirements	3-29
Central Processor Recommendations for Intelligent Resource Director (IRD)	3-29
Processor Considerations for Linux-Only LPs	3-30
Processor Considerations for Coupling Facility LPs.	3-30
Coupling Facility LPs Using Dedicated Central Processors (CPs) or Dedicated Internal Coupling Facility (ICF) CPs.	3-31
Coupling Facility LPs Using Shared Central Processors (CPs) or Shared Internal Coupling Facility (ICF) CPs.	3-31
Considerations for Coupling Facilities Running on Uniprocessor Models	3-33
Coupling Facility LPs Using Dynamic ICF Expansion	3-34
Processor Considerations for LPs with Multiple CP Types	3-35
Dedicated Central Processors	3-35
Suitable Workloads	3-36
Shared Central Processors.	3-36
Suitable Workloads	3-36
Processing Weights	3-37
Enforcement of Processing Weights	3-39
Processing Weight Management.	3-39
Processing Weight Management Examples.	3-40
Maintaining the Same Relative Percentages of CPU Resources	3-42

Processor Running Time	3-43
Wait Completion.	3-44
Workload Manager LPAR CPU Management of Shared CPs	3-45
Workload Charging by Soft-Capping to a Defined Capacity	3-46
Recommendations on Setting Up an LPAR Cluster	3-46
Enabling Management of Linux Shared Logical Processors by WLM's LPAR CPU Management Component	3-47
Defining Shared Channel Paths	3-48
Channel Path Access and Candidate Lists	3-48
I/O Device Candidate List	3-49
Procedure for Defining Shared Channel Paths	3-49
Communicating by Means of ESCON or FICON CTC	3-49
Dynamic CHPID Management (DCM) Considerations	3-50
I/O Priority Recommendations	3-50
Security-Related Controls	3-50
Logical Partition Isolation	3-50
I/O Configuration Control Authority	3-51
Global Performance Data Control Authority	3-51
Cross-Partition Authority	3-51
Dynamic I/O Configuration	3-52
Managing Dynamic I/O Configuration	3-52
Planning and Operation Considerations	3-52
Assigning Channel Paths to a Logical Partition	3-52
ESA/390 Logical Partitions	3-53
Coupling Facility Logical Partitions	3-53
Channel Path Reconfiguration and Logical Partition Activation	3-53
Dynamic I/O Configuration Effects on Channel Path Reconfiguration	3-53
Automatic Channel Path Reconfiguration	3-54
Automatic Load for a Logical Partition.	3-55
Defining Logical Partitions	3-55
Global Reset Profile Definitions	3-57
Options Page Definitions	3-57
Parameter Descriptions	3-58
Partitions Page Definitions	3-59
Parameter Descriptions	3-59
General	3-60
Parameter Descriptions	3-60
Parameter Descriptions	3-61
Processor Characteristics	3-63
Processor Page Definitions.	3-63
Parameter Descriptions	3-65
Security Characteristics	3-68
Parameter Descriptions	3-70
Establishing Optional Characteristics	3-71
Parameter Descriptions	3-72
Storage Characteristics	3-73
Central Storage Parameter Descriptions	3-73
Expanded Storage Parameter Descriptions	3-74
Load Information	3-75
Parameter Descriptions	3-75
Cryptographic Characteristics	3-77
Parameter Descriptions	3-77
Enabling Input/Output Priority Queuing	3-81
Changing Logical Partition Input/Output Priority Queuing Values	3-81
Parameter Descriptions	3-82
Moving Unshared Channel Paths	3-83

Moving Unshared Channel Paths from a z/OS System	3-83
Moving a Channel Path from the Hardware Console	3-83
Releasing Reconfigurable Channel Paths	3-83
Configuring Shared Channel Paths	3-84
Deconfiguring Shared Channel Paths	3-84
Removing Shared Channel Paths for Service	3-84
Changing Logical Partition Definitions	3-84
Changes Available Dynamically to a Running LP.	3-84
Changes Available at the Next LP Activation	3-85
Changes Available at the Next Power-On Reset (POR)	3-86

Planning Overview

This chapter provides a planning overview for defining logical partitions (LPs). Support for features, functions, and panels can differ depending on machine type, engineering change (EC) level, or machine change level (MCL). During IOCP execution, the names and numbers of the LPs are specified and channel paths are assigned to the LPs. Sample tasks and panels are explained in this publication, including references to tasks and panels available from the support element console. For detailed procedures for these operator tasks, see the *Hardware Management Console Operations Guide* and the *Support Element Operations Guide*.

Performance Considerations

The performance of an LP is controlled by:

- The number of logical CPs online
- Whether CPs are dedicated to or are shared by the LP
- The processing weight assigned to each LP
- The processor running time interval

The total number of logical CPs across all active LPs is one of the factors used to determine the dynamic dispatch interval. See “Processor Running Time” on page 3-43.

Use the RMF Partition Data Reports to determine the effective dispatch time for LPs. For more information about this RMF reporting enhancement see “RMF LPAR Management Time Reporting” on page 5-8.

The greater the number of active logical CPs relative to the number of physical CPs configured, the smaller the dispatch interval.

Dedicated and Shared Central Processors (CPs)

LPs can have CPs dedicated to them, or they can share CPs with other active LPs. Because the use of dedicated or shared CPs in an LP affects performance in several ways, the characteristics, limitations, and advantages of each should be carefully studied. Coupling facilities using dynamic ICF expansion can have a mix of dedicated ICF and shared general purpose CPs or a mixture of dedicated ICF and shared ICF CPs.

If IFAs or zIIPs are defined for a logical partition, they are defined in the same way as the general purpose CPs for the logical partition. If the general purpose CPs for a logical partition are shared, the IFAs/zIIPs are shared. If the general purpose CPs for a logical partition are dedicated, the IFAs/zIIPs are dedicated.

Dedicated and Shared Channel Paths

A configuration defining shared channel paths offers additional capabilities over an equivalent configuration containing unshared channel paths, while maintaining comparable system performance.

ITR Performance

The best ITR performance is achieved with dedicated LPs. To achieve optimal ITR performance in sharing LPs, keep the total number of logical CPs online to a minimum. This reduces both software and hardware overhead.

Capped Logical Partitions

It is recommended that LPs be defined as capped LPs at the Support Element/Hardware Management Console only when needed to support planned requirements. When a capped LP does not obtain needed CP resources, because it has reached its cap, activity for that LP is similar to a system running out of CP resources. Response time can be slower on systems which operate at their cap. For this reason, interactive response times can suffer when there is a mix of interactive and CP-intensive work in the same capped LP.

Recovery Considerations

Resources should be defined to LPs so that any hardware failure has a minimal impact on the remaining active LPs.

For example, the failure of a physical CP can cause the temporary loss of any logical CP that was dispatched on the physical CP. In many instances, recovery of a logical CP that was running on a failed physical CP will take place automatically when an available spare physical CP is dynamically brought into the configuration. Also, PR/SM LPAR is often able to transparently re-dispatch a shared logical CP on a different physical CP even when no spares are available. If a logical CP is still lost, the LP owning the logical CP can continue operating if it was running on an LP with at least two CPs dispatched on different physical CPs, and if the control program that is active in the LP can recover from CP failures.

Determining the Characteristics

The information in this section should help you determine the type and amount of CPC resources you need for each LP.

The total amount of resources that you can define for all LPs can exceed the configured resources.

Individual LP definitions are checked against the total resources installed. The actual allocation of these resources takes place only when the LP is activated. This design characteristic allows considerable flexibility when defining and activating LPs.

Control Program Support

Table 3-1 on page 3-6 summarizes the characteristics of the control programs that can be supported in an LP. See “Control Program Support in a Logical Partition” on page 1-5 for more information.

Some control programs require specific LP characteristics. For this reason consider all control programs before planning or defining LP characteristics.

Table 3-1. Control Program Support on z9

Control Program	Control Program Operating Mode	Maximum Number CPs	Maximum Central Storage	Maximum Expanded Storage	Maximum Number Channels
z/OS 1.9	ESA/390	54	1 TB	NA	256
z/OS 1.8	ESA/390	32	1 TB	NA	256
z/OS 1.7	ESA/390	32	128 GB	NA	256
z/OS 1.6	ESA/390	32	128 GB	NA	256
z/OS 1.5	ESA/390	16	128 GB	NA	256
z/OS 1.4	ESA/390	16	128 GB	NA	256
z/VM 5.3	Linux-Only or ESA/390	32	256 GB	128 GB	256
z/VM 5.2	Linux-Only or ESA/390	24	128 GB	128 GB	256
z/VM 5.1	Linux-Only or ESA/390	24	128 GB	128 GB	256
z/VSE 4.1	z/Architecture mode	10	8 GB	NA	256
z/VSE 3.1	ESA/390	10	2 GB	NA	256
z/TPF 1.1	ESA/390 or ESA/390 TPF	38	512 GB	NA	256
TPF 4.1	ESA/390 or ESA/390 TPF	16	2048 MB	NA	256
Linux for S/390	Linux-Only or ESA/390	16	2048 MB	256 GB	256

Notes:

1. On all z/OS versions, ESA/390 mode is the initial mode at IPL time; however, during IPL z/OS switches to z/Architecture 64-bit addressing.
2. z/VM is initially in ESA/390 mode at IPL time; however, during IPL it switches to z/Architecture mode.
3. z/VSE 4.1 is initially in ESA/390 mode at IPL time; however, during IPL it switches to z/Architecture mode. z/VSE 4.1 supports 64-bit real addressing. It does not support 64-bit addressing mode for user applications.

IOCDs Requirements

You must use IOCP or BUILD IOCDs with HCD to create an LPAR IOCDs. You can specify the LP names and MIF image ID numbers in an LPAR IOCDs.

Logical Partition Identifier

The logical partition identifier (ID) is used as the third and fourth hexadecimal digits of the operand stored by the Store CPU ID instruction for each CP in the LP. Even though at most 60 (decimal) logical partitions can be defined, valid identifiers for LPs are X'00' through X'3F'. The LP identifier must be unique for each active LP.

Use the Customize/Delete Activation Profiles task available from the CPC Operational Customization Tasks list to open a reset or image profile to define the LP identifier for an LP. The Partition identifier field is located on the General page for the LP. (See Figure 3-20 on page 3-63).

Mode of Operation

The mode of an LP depending on the model can be ESA/390, ESA/390 TPF, Linux-Only, or Coupling Facility.

The mode of an LP must support the mode of the control program loaded into it. ESA/390 LPs support ESA/390 control programs, Coupling Facility LPs support the coupling facility control code, and Linux-Only LPs support Linux, or z/VM Version 5.

Use the Customize/Delete Activation Profiles task available from the CPC Operational Customization Tasks list to open a reset or image profile to define the mode for an LP. The Mode field is located on the General page for the LP. (See Figure 3-20 on page 3-63).

Storage Configurations

The definition of central storage and expanded storage to the LP depends on the size of the I/O configuration, the storage limits of the control program loaded into the LP and on the storage requirements of the applications running in the LP.

Storage Resources

Use standard capacity-planning procedures to assess CPC storage requirements for LPs.

With dynamic storage reconfiguration (see “Dynamic Storage Reconfiguration” on page 3-12), an appropriately defined LP can dynamically add storage to its configuration that is released when another active LP removes it from its configuration or is deactivated. Additionally, an LP can dynamically add storage to its configuration following an increase in configured storage via concurrent memory upgrade (see “Concurrent Memory Upgrade” on page 2-66).

Single Storage Pool

With this function, all physical storage is dynamically designated by LPAR as either central storage or expanded storage as requirements of active logical partitions dictate.

As a result, the need to predesignate configured storage as either central storage or expanded storage prior to IML on the Storage Page of the Customize Activation Profile panel is not necessary and is not provided. The system programmer now has greater flexibility when planning the division of storage in order to satisfy anticipated logical partition definitions. Single storage pool streamlines the planning effort because LPAR automatically provides the correct storage designation for any configuration as the need arises. This feature is especially useful as it simplifies planning for migration to z/OS a 64-bit capable operation system.

All configured storage is defined, at IML, as central storage. As logical partitions with initial expanded storage are activated (and when reserved expanded storage is configured on-line), LPAR automatically allocates backing physical storage from the available pool as expanded storage.

A feature of this function is that the central storage addressability ranges used to back expanded storage allocated to logical partitions are communicated on the logical partition storage allocation display of the Storage Information task.

Central Storage

Central storage is defined to LPs before LP activation. When an LP is activated, storage resources are allocated in contiguous blocks. These allocations can be dynamically reconfigured. Sharing of allocated central storage among multiple LPs is not allowed.

Granularity of initial and reserved central storage amounts is dependent on the largest central storage amount (CSA) for the LP as follows:

Table 3-2. Central Storage Granularity

Largest Central Storage Amount	Storage Granularity
64 MB < CSA ≤ 32 GB	64 MB
32 GB < CSA ≤ 64 GB	128 MB
64 GB < CSA ≤ 128 GB	256 MB
128 GB < CSA ≤ 256 GB	512 MB
256 GB < CSA ≤ 512 GB	1024 MB

This granularity applies across the central storage input fields (Initial and Reserved fields). Use the larger of the initial and reserved central storage amounts to calculate storage granularity. For example, on a z9 EC, for an LP with an initial storage amount of 30 GB and a reserved storage amount of 48 GB, the central storage granularity of initial and reserved central storage fields is 128 MB, using the larger reserved storage amount to determine storage granularity. Central storage granularity for the initial storage origin input field is fixed at 64 MB.

Additionally, when migrating to these models, you must convert the existing central storage amounts that are not in multiples of the storage granularity supported.

For example, an LP with an initial central storage size of 10208 MB on a previous model with 16 GB of storage installed would require conversion of the initial central storage size to a multiple of 64 MB (32 MB granularity was valid on this prior model configuration).

Check your central storage definitions and consult your system programmer to ensure that they meet your processing needs.

Use the Customize/Delete Activation Profiles task available from the CPC Operational Customization Tasks list to open a reset or image profile to define the central storage for an LP. The Initial, Initial storage origin, and Reserved fields are located on the Storage page for the LP. See Figure 3-25 on page 3-73.

Initial Central Storage

- The initial amount of central storage represents the amount of storage allocated to an LP when it is activated.
- You must specify a nonzero number for the initial amount of central storage to activate the LP.
- If no origin is specified for the initial value, the LP will be treated as an LP that owns its own **addressability range**. This means that no other LP can map over it.

Reserved Central Storage

- Only ESA/390, ESA/390 TPF, and Linux-Only LPs can have nonzero amounts for this parameter. Coupling facility LPs cannot have reserved central storage amounts.
- The reserved amount of central storage defines the additional amount of central storage that can become available to an LP when no other activated LP has this reserved storage online. Reserved storage amounts are always offline after LP activation.
- Only LPs that specify a central storage origin can be allocated storage within the reserved central storage defined for another LP. Reserved central storage can only overlap LPs that specify an origin.
- For LPs that do not specify central storage origins, the LP's reserved storage is available to be brought online to the LP whenever there is any storage that is not being used by any other LPs.
- For LPs that do specify a central storage origin, if some of the reserved central storage is not available, the reserved central storage that is available can still be configured online starting at the reserved storage origin.
- Specifying a zero for the reserved size indicates the LP's central storage cannot get any larger than the initial size for the duration of the activation.

Central Storage Origin

- There must be enough contiguous central storage addressability to fit the reserved amounts, but the reserved amounts can be in use by other LPs (that specified origins for storage) at the time the LP is activated.
- The origin for central storage defines the starting megabyte where the central storage **addressability range** begins for the LP. If enough storage is not available to satisfy the initial central storage request starting at this origin, LP activation will fail.
- For central storage, the specification of the origin parameter provides the only way to overlap storage definitions. For example, the reserved storage definition for one LP can overlap the storage definition of another LP when the origin parameter is specified for both LPs.

All **installed** storage is configured at IML as central storage. Storage will be configured as expanded storage when a logical partition with initial expanded storage is activated or when a reserved expanded storage element of an already active logical partition is configured online.

The total amount of central storage **addressability** that can be used to map central storage for LPs is **at least** twice the amount of **installed** storage. You can review current LP storage allocations by using the Storage Information task available from the CPC Operational Customization Tasks lists.

- Since the specification of a central storage origin is optional, there are two implementations of dynamic storage reconfiguration for reserved central storage:

1. If you do not specify a central storage origin, the reserved central storage is available for the LP when there is sufficient physical storage available to meet the reserved central storage request.

Storage can be reassigned between LPs that have noncontiguous address ranges. Reconfiguration is not limited to growing into an adjacent LP. LPs are eligible to use noncontiguous central storage if they have unique (not overlapping) storage range definitions. Specifying an origin is not required when defining a reserved storage amount for an LP, and by not specifying an origin, the LP is reserved an addressable storage range for its entire storage configuration (initial plus reserved amount). Not specifying an origin ensures no overlap. The reserved central storage is available for LPs that do not specify an origin when there is sufficient physical storage to meet the reserved central storage request.

2. If you specify a central storage origin, the entire reserved central storage amount is only available for the LP when:
 - No other LP has this central storage address range, to which the reserved central storage is mapped, online.
 - There is sufficient physical storage available to meet the reserved central storage request.

- If no origin is specified for an LP, the system assigns the storage to the LP using a top-down first-fit algorithm. Storage that is part of reserved central storage for another activated LP is not available for any LPs that do not specify a central storage origin.

Expanded Storage

Optional expanded storage is defined to LPs before LP activation. When an LP is activated, storage resources are allocated in contiguous blocks. These allocations can be dynamically reconfigured. Sharing of allocated expanded storage among multiple LPs is not allowed.

Expanded storage granularity is fixed at 64 MB. Expanded storage granularity applies across the expanded storage input fields (Initial, Initial storage origin, and Reserved fields).

When migrating to these models, you must convert any existing expanded storage amounts that are not in multiples of the storage granularity supported.

For example, an LP with an initial expanded storage size of 10208 MB on a previous model with 16 GB of storage installed would require conversion of the initial expanded storage size to a multiple of 64 MB (32 MB granularity was valid on this prior model configuration)

Use the Customize/Delete Activation Profiles task available from the CPC Operational Customization Tasks list to open a reset or image profile to define the expanded storage for an LP. The Initial, Initial storage origin, and Reserved fields are located on the Storage page for the LP. See Figure 3-25 on page 3-73.

Check your expanded storage definitions and consult your system programmer to ensure that they meet your processing needs.

Initial Expanded Storage

- The initial amount of expanded storage represents the amount of storage allocated to an LP when it is activated. Expanded storage cannot be defined for coupling facility mode LPs.

- Initial expanded storage is reconfigurable (can be configured offline or online by the control program) for ESA/390 LPs that have no reserved expanded storage defined.
- The initial amount of expanded storage can be zero when the reserved expanded storage is nonzero.
- A request for an initial expanded storage amount can be satisfied only when both sufficient expanded storage addressability and sufficient physical storage are available.

Also, there must be a large enough contiguous central storage **addressability range** to back the initial expanded storage amount. This range will be chosen by the system in a top-down first-fit algorithm. It is not user-specifiable. If a large enough central storage **addressability range** is not available to satisfy the initial expanded storage request, LP activation will fail.

Reserved Expanded Storage

- Only LPs that specify a reserved expanded storage origin can be allocated storage within the reserved expanded storage for another LP. Reserved expanded storage can only overlap LPs that have an expanded storage origin specified.
- Expanded reserved storage is contiguous to and above the initial amount of expanded storage.
- ESA/390 and Linux-Only LPs can have nonzero amounts for this parameter. Coupling facility LPs cannot have reserved expanded storage amounts.
- The reserved amounts of expanded storage defines the additional amounts of storage that can become available to an LP when no other activated LP has this reserved storage online. Reserved storage amounts are always offline after LP activation.
- For reserved expanded storage, all of the reserved expanded storage must be available to configure any reserved expanded storage online.
- Reserved expanded storage for an LP can be specified with or without specifying an origin.

Also, there must be a large enough contiguous central storage **addressability range** to back the reserved expanded storage amount. This range will be chosen by the system in a top-down first-fit algorithm. It is not user-specifiable. If a large enough central storage **addressability range** is not available to satisfy the reserved storage request, the element cannot be brought online.

Expanded Storage Origins

- The origin for expanded storage defines the starting megabyte where the expanded storage **addressability range** begins for the LP. If enough expanded storage is not available to satisfy the initial request starting at this origin, LP activation will fail. If there is no initial amount of expanded storage defined, the reserved amount of expanded storage starts at this origin.
- There must be enough contiguous expanded storage **addressability** to fit the reserved amounts, but the reserved addressability can be in use by other LPs (that have expanded storage origins specified) at the time the LP is activated.
- For expanded storage, the specification of the origin parameter provides the only way to overlap storage definitions. For example, the reserved storage definition for one LP can overlap the storage definition of another LP when the origin parameter is specified for both LPs. The total amount of expanded storage **addressability** that can be used to map expanded storage for LPs is **at least**

twice the amount of **installed** storage. You can review current LP storage allocations by using the Storage Information task available from the CPC Operational Customization Tasks list.

- The availability of reserved expanded storage depends on whether or not you specify an expanded storage origin for an LP. By making the specification of an expanded storage origin optional, there are two implementations of dynamic storage reconfiguration for reserved expanded storage:
 1. If you do not specify an expanded storage origin, the reserved expanded storage is available for the LP when there is sufficient physical storage available to meet the reserved expanded storage request.

Expanded storage can be reassigned between LPs that have noncontiguous address ranges. Reconfiguration is not limited to growing into an adjacent LP. LPs are eligible to use noncontiguous expanded storage if they have unique (not overlapping) storage range definitions. Specifying an origin is not required when defining a reserved storage amount for an LP and by not specifying an origin the LP is reserved an addressable storage range for its entire storage configuration (initial plus reserved amount). Not specifying an origin ensures no overlap.

The reserved expanded storage is available for LPs that do not specify an origin when there is sufficient physical storage to meet the reserved expanded storage request.
 2. If you specify an expanded storage origin, the reserved expanded storage is only available for the LP when:
 - No other LP has this expanded storage address range online
 - There is sufficient physical expanded storage available to meet the reserved expanded storage request.
- If no origin is specified for an LP, the system assigns the storage to the LP using a top-down-first-fit algorithm.

Dynamic Storage Reconfiguration

Dynamic storage reconfiguration allows both central and expanded storage allocated to an LP to be changed while the LP is active. It is supported in ESA/390 LPs running z/OS. Dynamic storage reconfiguration is not supported in coupling facility or Linux-Only LPs.

Dynamic storage reconfiguration provides the capability to reassign storage from one LP to another without the need to POR the CPC or IPL the recipient LP. Every LP has a storage range definition consisting of an origin, initial, and reserved amounts. The reserved value determines how much additional storage can be acquired using dynamic storage reconfiguration.

Storage is released when an LP is either deactivated or its reserved storage element is deconfigured. Additionally, you can release central storage in amounts smaller than the defined storage element size.

With dynamic storage reconfiguration, ESA/390 LPs can have reserved amounts of central storage and expanded storage. This storage can become available to the LP if no other active LP has this reserved storage online. Reserved central storage and reserved expanded storage can be made available to the LP by commands from the operating system console.

If the operating system running in the LP supports physical storage reconfiguration, use operating system commands to make the reserved storage available to the LP without disrupting operating system activities.

For z/OS, use the following command format to reconfigure central storage:

```
CF STOR(E=1),<OFFLINE/ONLINE>
```

Dynamic storage reconfiguration on the System z9 for central storage enables central storage to be reassigned between LPs that have noncontiguous address ranges. In this case, LPAR can allocate a “hole” or some set of central storage addresses for which there is no backing physical storage assigned. Later, LPAR is able to configure storage online to the LP by assigning some physical central storage to the hole.

Similarly, dynamic storage reconfiguration for expanded storage enables expanded storage to be reassigned between LPs that have noncontiguous expanded storage address ranges.

For central storage, you can reconfigure central storage in amounts equal to the storage granularity supported by the CPC. For z/OS, use the following command format:

```
CONFIG STOR(nnM),<OFFLINE/ONLINE>
```

Central Storage Dynamic Storage Reconfiguration Examples

Figure 3-1 shows an example of central storage with dynamic storage reconfiguration capability. This figure shows LP-A, LP-B, and LP-C.




Figure 3-1. Central Storage Layout

To reconfigure central storage, deactivate LP-B to free up storage directly above LP-A. Figure 3-2 on page 3-15 shows how you may reconfigure central storage in this case.




Figure 3-2. Reconfigured Central Storage Layout

Figure 3-3 is another example of dynamic reconfiguration of central storage. For this example, assume the amount of installed storage is 32 GB. The amount of physical central storage used by hardware system area (HSA) is 2 GB in this example. The storage granularity is 64 MB so the HSA addressable storage size is 2048 MB. This leaves 1022 GB of central storage addressability to be allocated to LPs.

LP-A and LP-B are defined with an initial amount of 14 GB each of central storage, a reserved amount of 2 GB, and system determined central storage origins. LP-A and LP-B are activated and IPLed. At the completion of IPL, LP-A has its reserved central storage configured online by entering **CF STOR(E=1),ONLINE** from the z/OS software console for LP-A. Figure 3-3 shows the resulting storage layout following these actions.




Figure 3-3. Initial Central Storage Layout

At a later point in time, the reserved storage from LP-A can be reconfigured to LP-B. Note that both LPs should specify an RSU value of at least 32 (2048/64) for reconfigurations of storage to work. From the z/OS software console for LP-A, enter **CF STOR(E=1),OFFLINE**. Next, from the z/OS software console for LP-B, enter **CF STOR(E=1),ONLINE**. Figure 3-4 shows the resulting storage layout following these actions. The reserved storage is fully reconfigured without an outage to either LP. Also, the procedure can be reversed without an outage as well.




Figure 3-4. Central Storage Layout Following Reconfiguration

Figure 3-5 on page 3-18 and Figure 3-6 on page 3-19 are examples of dynamic reconfiguration of central storage. For this example, assume the amount of installed storage is 32 GB. The amount of physical central storage used by hardware system area (HSA) is 2 GB in this example leaving 30 GB of physical storage for allocation to LPs. The storage granularity is 64 MB so the HSA addressable storage size is 2048 MB. This leaves 1022 GB of central storage addressability to be allocated to LPs.

LP-A is defined with an initial amount of 14 GB of central storage, a reserved amount of 2 GB, and a system determined central storage origin. LP-B is defined with an initial amount of 14 GB of central storage, with a system determined central storage origin and an initial amount of 2 GB of expanded storage. The central storage addressability used to back online expanded storage is always system determined and in this case is coincidentally but not necessarily contiguous with the addressability used to back the logical partition's central storage. LP-A and LP-B are activated and IPLed. Figure 3-5 on page 3-18 shows the resulting storage layout following these actions.




Figure 3-5. Initial Central Storage Layout

At a later point in time, LP-B deconfigures its expanded storage and the physical storage that was used for its expanded storage is reconfigured to LP-A. Note that LP-A should specify an RSU value of at least 32 (2048/64) for reconfigurations of storage to work. First, from the z/OS software console for LP-B, enter **CF ESTOR(E=0),OFFLINE**. Next, from the z/OS software console for LP-A, enter **CF STOR(E=1),ONLINE**. Figure 3-6 shows the resulting storage layout following these actions. The reserved storage is fully reconfigured without an outage to either LP-A or LP-B.




Figure 3-6. Central Storage Layout Following Reconfiguration

Expanded Storage Dynamic Storage Reconfiguration Examples

To dynamically reconfigure expanded storage, the system programmer may plan which LPs will donate storage and which LPs will receive storage. To do this, the system programmer may specify origins that overlap the reserved (RSVD) storage of the receiver with the initial (INIT) storage (and possibly some RSVD storage) of the donor.

Figure 3-7 shows an example of expanded storage with dynamic storage reconfiguration capability. This figure shows LP-A, LP-B, and LP-C.




Figure 3-7. Expanded Storage Layout

To reconfigure expanded storage, deactivate LP-B to free up storage directly above LP-A. Figure 3-8 on page 3-21 shows how you might reconfigure expanded storage in this case.




Figure 3-8. Reconfigured Expanded Storage Layout

Figure 3-9 on page 3-22 shows an example of expanded storage with dynamic storage reconfiguration capability, where the system determines logical partition origins.

In Figure 3-9, the expanded storage **addressability** above 1023 GB and below 1019 GB is unused and considered a gap. For the purpose of this example, assume that LP-C was activated with 1 GB of storage and was then deactivated to arrive at this storage layout. The storage range between between 1019 GB and 1021 GB is assigned to LP-A for as long as the LP is active.




Figure 3-9. Expanded Storage Layout

Notice that the total amount of expanded storage **addressability** that can be used to map expanded storage for LPs is at least **twice** the amount of **installed** storage.

Figure 3-10 on page 3-23, Figure 3-11 on page 3-23, Figure 3-12 on page 3-24, and Figure 3-13 on page 3-25 describe examples of dynamic reconfiguration of expanded storage. For this example, assume the amount of installed storage is 32 GB. The amount of physical central storage used by hardware system area (HSA) is 2GB in this example leaving 30 GB of physical storage for allocation to LPs. The storage granularity is 64 MB so the HSA addressable storage size is 2048 MB. This leaves 1022 GB of central storage addressability to be allocated to LPs.

LP-A is defined with an initial amount of 14 GB of central storage, an initial amount of 2 GB of expanded storage and a reserved amount of 2 GB of expanded storage, and system determined central storage origins. The origin of central storage used to back online expanded storage is always system determined and in this case is coincidentally but not necessarily contiguous with the addressability used to back the logical partition's initial central storage amount. LP-B is defined with an initial amount of 14 GB of central storage, and a reserved amount of 2 GB of central storage, with a system determined central storage origin. LP-A and LP-B are activated and IPLed. At the completion of IPL, LP-B has its reserved central storage configured online by entering **CF STOR(E=1),ONLINE** from the z/OS software console for LP-B. Figure 3-10 on page 3-23 shows the resulting storage layout following these actions.




Figure 3-10. Initial Central Storage Layout




Figure 3-11. Initial Expanded Storage layout

At a later point in time, the reserved central storage from LP-B can be reconfigured to LP-A as expanded storage. Note that LP-B should specify an RSU value of at least 32(2048/64) for reconfigurations of storage to work. From the z/OS software console for LP-B, enter **CF STOR(E=1),OFFLINE**. Next from the z/OS software console for LP-A, enter **CF ESTOR(E=1),ONLINE**. Figure 3-12 and Figure 3-13 on page 3-25 show the resulting central storage and expanded storage layouts following these actions. The reserved storage expanded storage is fully reconfigured without an outage to LP-A or LP-B.




Figure 3-12. Central Storage Layout Following Reconfiguration




Figure 3-13. Expanded Storage Layout Following Reconfiguration

Recommendations for Storage Map Planning

Planning storage maps as described below will help avoid storage fragmentation and remove dependencies on the order of activation of LPs.

For more information and examples of storage map planning, see the publications listed for dynamic storage reconfiguration in “About This Publication” on page xvii.

- It is possible to define expanded storage to LPs in such a way that many LPs use the same expanded storage addressability (one LP at a time). If you want this, no more than one LP should have this expanded storage defined as its initial expanded storage. Instead, define it as reserved expanded storage.

Note: If you do not observe this convention, it can inhibit automatic activation of LPs.

- Map all LPs that require an origin from the bottom up (that is, start with 0 MB and build upward). When both central and expanded storage require origins, map the storage in the same order for both central and expanded storage.

If no origin is specified for an LP, the system assigns the storage to the LP using a top-down first-fit algorithm for both central and expanded storage.

Operation Considerations for Dynamic Storage Reconfiguration:

- Initial central storage for an LP is allocated on logical storage element 0 and any reserved central is allocated on logical storage element 1.
- Whenever a load clear or system reset clear is performed on an LP, it forces the reserved central storage element offline. This allows z/OS to flag the reserved storage element as reconfigurable to allow it to be later deconfigured from the LP.

The reserved expanded storage element will remain online if it was already online from the prior IPL.

Note: When z/OS is IPLed immediately after the LP is activated, both the reserved central storage element and the reserved expanded storage element are offline.

- Whenever z/OS is re-IPLed in an LP that has a reserved central storage element, a load clear or system reset clear followed by load normal should be performed to force the reserved central storage element offline.
- When a standalone dump is to be performed on an LP, perform a load normal (not a load clear) on that LP to keep the reserved storage element online and preserve the storage contents.

CPCs with the Sysplex Failure Manager (SFM)

The Sysplex Failure Manager (SFM) allows you to reset and reconfigure one or more logical LPs and their related storage. SFM allows workload redistribution from the failed primary system to the backup system without operator intervention. For a detailed description of how SFM works, see the z/OS Setting Up a Sysplex publication.

To allow an LP to initiate these functions, use the Customize/Delete Activation Profiles task to open a reset profile to authorize an LP to issue instructions to other LPs. The Cross partition authority check box is located on the Security page for the LP.

The following functions exist for SFM:

- Cross Partition System Reset
This function causes a specified LP to be reset. The reset is accomplished via the RESETTIME(nnn) keyword in the SYSTEM statement of the z/OS SFM policy.
- Cross Partition Deactivation
This function causes a specified LP to be deactivated. The deactivation is accomplished via the DEACTTIME(nnn) keyword in the SYSTEM statement of the SFM policy, and also, the RECONFIG statement in the SFM policy with a specific TARGETSYS(sysname) specified.
- Cross Partition Nonspecific Deactivation
This function causes all logical partitions which are currently using any portion of the reconfigurable central or expanded storage of the issuing partition to be deactivated. The issuing partition is not deactivated. The nonspecific deactivation is accomplished via the RECONFIG statement in the SFM policy with a non-specific TARGETSYS(ALL) issued.

The Automatic Reconfiguration Facility (ARF) function is a hardware/LPAR function that is part of the cross-partition authority control setting. ARF functions are used by SFM policy functions within z/OS, when RESETTIME, DEACTTIME, or the RECONFIG statement is coded in the SFM policy.

TARGETSYS(ALL) Examples

Specifying an Origin: Assume that the backup partition has specified an origin, minimal initial storage, and a large amount of reserved storage. Since the backup system does not own its complete addressable range, two other partitions, are defined in the reserved storage of the backup partition. See Figure 3-14 on page 3-27

3-27 for the storage layout before nonspecific deactivation.




Figure 3-14. Backup Partition Layout Before Nonspecific Deactivation

Assume that the backup LP has been given cross partition authority. See Figure 3-15 for the storage layout at the completion of a takeover by the backup LP.




Figure 3-15. Backup Partition Layout After Nonspecific Deactivation

Number of Central Processors

The CPs defined to an LP are called logical CPs. The total number of initial logical CPs for each LP cannot be greater than the number of physical CPs installed. CPs can be dedicated to LPs or shared by LPs (*sharing* LPs).

The number of CPs defined for an LP represents the number of logical CPs on which the control program will dispatch work and is determined by several factors as described in “Maximum Number of Central Processors.”

On a System z9, you can optionally install one or more Internal Coupling Facility (ICF) features for use by a coupling facility LP and/or one or more Integrated Facility for Linux (IFL) features for use by a Linux-Only LP. See “Coupling Facility LPs Using Dedicated Central Processors (CPs) or Dedicated Internal Coupling Facility (ICF) CPs” on page 3-31.

You can also optionally install one or more Integrated Facility for Applications (IFA) features for use by an LP. The total number of initial logical IFAs for each LP cannot be greater than the number of physical IFAs installed.

You can also optionally install one or more Integrated Information Processors (zIIP) features for use by an LP. The total number of initial logical zIIPs for each LP cannot be greater than the number of physical zIIPs installed.

Use the Customize/Delete Activation Profiles task available from the CPC Operational Customization Tasks list to open a reset or image profile to define the number of CPs for an LP. The Number of processors field is located on the Processor page for the LP.

Maximum Number of Central Processors

The maximum number of CPs that can be defined depends on:

- The number of CPs that are available.

The maximum number of logical CPs available for definition in a single LP is the total number of CPs achievable via concurrent CPU upgrade, or 32, whichever is less.

Notes:

1. The maximum initial logical CPs defined for a dedicated partition **equals** physical CPs online for the current model **minus** physical CPs currently dedicated and online to other dedicated partitions **minus** the maximum online number of shared CPs among all the activated logical partitions that are using shared CPs.
 2. The maximum initially online logical CPs defined for a shared partition **equals** physical CPs online for the current model **minus** physical CPs currently dedicated and online to other partitions using dedicated CPs.
 3. The preceding rules for maximum CPs apply independently to each type (general purpose, ICF, IFL, IFA or zIIP) of processor defined to the logical partition. For instance, when defining a logical partition to use a combination of general purpose CPs and IFAs/zIIPs, the rules for the maximum initially online of each type of processor is calculated independently against what is currently installed and in use for that type of physical processor.
- The number of CPs that are supported by the required control program.
Some control programs support as many as 32 CPs. The number of CPs defined for an LP should not exceed the number supported by the control program used

for that LP. The stated maximum supported processors for a particular control program is applied against the sum of the counts of all processor types defined to the partition.

Workload Requirements

The number of logical CPs defined also depends on the workload requirements and the ability of the control program or application program to effectively use multiple logical CPs.

- The number of CPs required to meet the peak demands of the LP
When a sharing LP is activated, it should be assigned enough CPs to meet its peak demands and any immediate growth requirements.

Note: Too few CPs could limit the number of potential transactions, and too many active logical CPs could affect performance. In addition to a number of initially online logical CPs sufficient to meet the current peak demands of the LP, the definition should include the number of reserved logical CPs required for possible growth requirements.

For an LP that uses dedicated CPs, the number of CPs required to meet peak demand should be determined.

- The number of CPs shared by an LP
The physical CPs used by a sharing LP can be limited by the number of CPs defined to it. For example, on a six-way CPC, a two-way LP could never get more than a third of the CP resources.
On most machines, there are many possible configurations. For example, if a two-way CPC is to be used by three sharing LPs, configuration options include all two-way LPs, all one-way LPs, or a mix of one-way and two-way LPs.
Three two-way LPs should only be considered if all LPs have peak demands exceeding the capacity of a one-way CPC. In this instance, the average requirements during peak periods should be examined carefully to ensure that the total does not exceed the capacity of the CPC.
Three 1-way LPs are optimal for ITR performance and should be used if no LP has peaks exceeding half of the capacity of a two-way CPC.

Central Processor Recommendations for Intelligent Resource Director (IRD)

It is highly recommended that the number of initial CPUs specified be maximized for logical partitions whose CPU resources, enabled by IRD clustering technology, are managed by z/OS's WorkLoad Manager (WLM). WLM's LPAR CPU Management component will automatically configure **off** all initially online logical CPs in excess of the minimum required to meet current demands so over-specification has no negative consequences.

Moreover, only those logical CPs defined as initially online, or configured **on** explicitly by operator commands, will subsequently be configured **on** by WLM if and when increases in workload requirements demand. Specification of reserved logical CPs should be used only to accommodate potential increases to the number of installed physical CPs.

WLM CPU management does not currently manage any IFA or zIIP logical processors defined to a WLM managed logical partition. As such, the maximizing of initial CPUs recommendation applies only to the general purpose processor

definition for the logical partition; not the IFA or zIIP definition. Maximizing reserved IFAs and zIIPs is still recommended to allow for concurrently adding logical IFAs or zIIPs to the partition if needed.

WLM managed CPs cannot have dedicated CPs. Though capping of shared CPs is also disallowed, WLM's License Manager component may soft-cap logical CPs when a LPs rolling 4-hour average CPU utilization has reached its defined capacity. (See "Workload Charging by Soft-Capping to a Defined Capacity" on page 3-46.)

Processor Considerations for Linux-Only LPs

A Linux-Only mode LP can be allocated either general purpose CPs or IFLs. For optimum cost-effectiveness, IFLs should be used whenever possible. Choose IFL CPs on the Processor page of the Customize/Delete Activation Profile task. IFLs can be allocated to a Linux-Only mode LP as either dedicated or shared. Unlike Coupling Facility mode LPs, Linux-Only LPs cannot be defined with a mix of dedicated and shared processors or a mix of general purpose and ICF processors.

Utilization of IFLs will be included in the data reported for ICF CPs in RMF Partition Data reports and other similar reports. There is no way to distinguish on such a report between data reported for ICF CPs and IFLs.

A Linux-Only LP, whether allocated IFLs or general purpose CPs, will not support any of the IBM traditional operating systems (such as z/OS, TPF, or z/VSE). Only Linux, or z/VM Version 5 with only Linux users can run in a Linux-Only mode LP. Logical partitions defined as Linux-Only that attempt to load a traditional operating system will be system check-stopped.

Processor Considerations for Coupling Facility LPs

Important Note

Dynamic ICF Expansion and Enhanced Dynamic ICF Expansion across ICFs is not recommended for coupling facilities that are used for primary production workload. All processors assigned to the coupling facility partition should be dedicated to that logical partition if it is used for primary production workload.

You can define a coupling facility to use any one of the following combinations of CP definitions:

- One or more dedicated general purpose CPs
- One or more shared general purpose CPs
- One or more dedicated ICF CPs
- One or more shared ICF CPs
- One or more dedicated ICF CPs and one or more shared general purpose CPs (dynamic ICF expansion definition). Because the coupling facility is not using more CP resource than it requires on the shared CPs, you should set the processing weight to a high value without affecting other work during normal operation.
- One or more dedicated ICF CPs and one or more shared ICF CPs (enhanced dynamic ICF expansion definition). Because the coupling facility is not using more CP resource than it requires on the shared CPs, you should set the processing weight to a high value without affecting other work during normal operation.

Any other combination of processor definitions is not supported.

Coupling Facility LPs Using Dedicated Central Processors (CPs) or Dedicated Internal Coupling Facility (ICF) CPs

Important Note

IBM **strongly** recommends using dedicated CPs or dedicated ICFs for production coupling facility LPs because coupling facility channel paths and requests have critical response time requirements. When the coupling facility is running on a dedicated CP or dedicated ICF CP, an “active wait” polling algorithm is used to look for coupling facility requests. This results in the fastest response time and throughput possible for coupling facility requests.

Considerations for Coupling Facilities using Internal Coupling Facility (ICF)

CPs: The following considerations apply to coupling facility LPs that use ICF CPs.

- ICF CPs are managed separately from general purpose CPs. Their presence and usage does not affect any of the definition or activation rules pertaining to LPs that use general purpose CPs.
- ICF CPs will appear in RMF Partition Data reports and other similar reports only when an LP is activated and assigned to use an ICF CP.
- Non-ICF work is never allowed to be run on an ICF CP. Only PR/SM (to manage the CP) or a logical partition defined to use ICFs is allowed on the ICF CP.

Coupling Facility LPs Using Shared Central Processors (CPs) or Shared Internal Coupling Facility (ICF) CPs

Using shared CPs for coupling facility LPs has the following considerations. The same set of considerations for using shared CPs equally apply to using shared ICF CPs.

- You should **not** cap the processing weights for a coupling facility LP. If you cap processing weights and you are running a coupling facility, LPAR attempts to support the cap but may not be successful and the system can be less responsive.
- A CF with shared CPs may use less CPU resource, but this may come at an expense to performance. The CF LIC runs in a polling loop so it always looks busy, even if there is little CF activity. If the CF is sharing a CP with a z/OS image that has low CPU demands, the CF could use the entire CP. If the CF is sharing CP(s) with another CF, both will look busy and the CPU resource will be split between them. For example, if a production CF and a test CF shared a CP with equal weights, each would only have use of one-half the CP. This would appear on the RMF CF Usage Summary Report as LOGICAL PROCESSORS: DEFINED 1 EFFECTIVE 0.5.

Requests to a CF that is sharing a CP may have to wait if another image (z/OS or CF) is using the shared CP. In fact, the CF's logical CP may have to wait for several logical CPs of other partitions to run before it runs again. This has the net effect of elongating response times. Requests to the CF that could typically take several microseconds in a dedicated environment can be elongated by many milliseconds, easily 12.5 to 25 milliseconds or more. If CF response time is not a consideration (for example, a test CF or a backup CF), the installation can further limit CF usage of CPU resource by enabling the Dynamic Dispatch option. This will periodically pause the polling loop. The number and length of the pauses is determined by the activity in this CF. While this gives more CPU resource to the other partitions that are sharing the CP, it elongates the service time in this CF, in most cases, by a considerable amount.

- Shared ICF and shared general purpose processors are each managed as separate “pools” of physical resources. As such, the processing weights assigned to logical partitions using shared ICF processors are totaled and managed separately from the total weights derived from all of the logical partitions using shared general purpose processors. Similarly, when the processor running time is dynamically determined by the system, the calculation of the dynamic running time is performed separately for each pool. If a user supplied run time is specified, this will apply to both processor pools.
- RMF identifies which logical and physical CPs are of each type when both general purpose and ICF processors are present in the configuration on its partition data report.
- An uncapped coupling facility LP honors its processing weight up to a point. PR/SM LPAR attempts to help ensure that each logical CP defined to a coupling facility LP gets at least 1 run time interval of service per every 100 milliseconds. For example, for a typical 12.5 millisecond run time interval, each logical CP gets 1/8th of a physical CP. This can translate into a response time elongation that is several thousand times as long as a typical CF request using dedicated CPs. With dynamic coupling facility dispatch, the coupling facility will not necessarily consume entire run time intervals at low request rate times. In low request rate times, each coupling facility engine can consume far less than 1/8th of a physical CP. The CP resource consumption can be more in the 1-2% range. At higher request rates (for example, when the coupling facility is actually busy handling requests), the 1/8th minimum will again become effective.

Note: Anticipated processor usage by a coupling facility may spike much higher than what you would intuitively expect given the non-CF workload. For instance, system reset of a system or logical partition that was communicating with (connected to) a coupling facility can temporarily cause a considerable increase in the demands placed on the coupling facility.

- All requests to coupling facilities from z/OS LPs that share CPs with a coupling facility are treated internally to the machine as asynchronous requests. This is true even if the requests are to a coupling facility that the z/OS LP is not sharing CPs with. This conversion is transparent to z/OS but it can result in increased synchronous service times to the coupling facility as reported by RMF. As far as the operating system, RMF, and the exploiter code is concerned, the requests that are initiated synchronously by software are still being processed synchronously, and they show up as such on RMF reports.
- Choose a weight for the coupling facility LP based on the anticipated CP requirements of the coupling facility. When deciding how much CP resource should be given to each coupling facility logical CP, consider the following:
 - When using dynamic CF dispatching, the weight for the coupling facility can safely be set to a value that affords the proper CP resources to the coupling facility in times of the highest volume of requests to the coupling facility. In low request rate periods, the coupling facility will automatically throttle back on its CP usage making the CP resource available for redistribution to the other defined LPs in the configuration. Also, note that at low request rate times, RMF Coupling Facility Activity Reports will show some elongation in response times for requests to the coupling facility. With the low rate of these requests, overall system performance should not be noticeably impacted.
 - CP resource requirements vary depending on the coupling facility exploiter functions and your sysplex hardware and software configuration. As a general guideline, when the anticipated CP requirement of the coupling facility is less than one physical CP, set the weight of the coupling facility LP so that the

coupling facility logical CP has 50% or more of a physical CP resource. If you have less than a full physical CP allocated to the coupling facility logical CP, **this will result in elongation of response times.**

Examine the RMF Coupling Facility Activity Report and tune the coupling facility LP weight until your performance requirements are met.

Note: When examining the RMF Coupling Facility Activity Report, you may see elongated average response times. Usually, these are accompanied with high standard deviations as well. This indicates most requests are in the expected range with an occasional very elongated response that throws the average off.

- Give more weight to the coupling facility LP for functions that have more stringent responsiveness requirements. For example, you can set the coupling facility weight higher for coupling facility exploiter functions such as IMS/IRLM. For IMS/IRLM, you may want to set the weight of the coupling facility LP so that each coupling facility logical CP runs almost dedicated. For example, you may want to set a weight that will give each coupling facility logical CP 95% or more of physical CP resources. In another case, if the CF contains structures which are using System Managed Duplexing, you should set the weight of the coupling facility LP so that the coupling facility CP has at least 95% of a physical CP. If the coupling facility has less than 95% of a physical CP, there is a possibility that response from the partner of the duplexed structure will timeout and the duplexing of that structure will cease.
- Less weight may be required for coupling facility exploiter functions, such as the JES2 Checkpoint data set. If your coupling facility is being used exclusively as a JES2 checkpoint, your coupling facility responsiveness requirements may be less stringent. If this is the case, try decreasing the weight of the coupling facility LP so that each coupling facility logical CP receives, for example, 40-50% of physical CP resources.
- As the total traffic (requests per second) to the coupling facility increases, there is a greater need for real coupling facility CP time. To a point, the increase in traffic may not require an increase in the coupling facility LP weight. This is because coupling facility active wait time turns into coupling facility busy time. You must monitor coupling facility utilization and adjust the LP weight to help ensure that your performance requirements are being met.
- Even in a test environment, the above guidelines should be followed. Optimally, a weight resulting in approximately 50% or more of a processor should be made for each coupling facility logical processor. ***Failure to provide sufficient weight to a coupling facility may result in degraded performance, loss of connectivity, and possible loss of coupling links due to time-outs.*** Dynamic CF dispatching must not be disabled. Dynamic CF dispatching is automatically enabled by default for coupling facility LPs that are using shared general purpose CPs. Additionally, processor resource capping of the coupling facility logical partition(s) logical CPs must **not** be enabled.

Considerations for Coupling Facilities Running on Uniprocessor Models

On a uniprocessor or smaller machine, IBM strongly recommends that coupling facility LPs should not share general purpose CPs with non-CF workloads (for example, z/OS), even in a test environment. ***Such a configuration carries with it significant risks of degraded performance, loss of connectivity, and possible loss of coupling links due to time-outs and must not be used for a production level system.*** For production configurations, IBM strongly recommends using one or more internal coupling facility (ICFs) CPs for the coupling facility LPs.

For a test configuration only, if this option is not available and coupling facility LPs must share a general purpose CP with non-CF workloads, diligent adherence to the preceding recommendations for coupling facilities using shared CPs will minimize undesirable consequences.

For greater detail on how to plan for and set-up a parallel sysplex in a shared CP environment, please see one or more of the following:

Dynamic ICF Expansion (<http://www.ibm.com/support/techdocs/atsmastr.nsf/PubAllNum/W98028>)

MVS/ESA™ Parallel Sysplex Performance (<http://www.ibm.com/support/techdocs/atsmastr.nsf/PubAllNum/W9731A>)

MVS/ESA Parallel Sysplex Performance - LPAR Performance Considerations for Parallel Sysplex Environments (<http://www.ibm.com/support/techdocs/atsmastr.nsf/PubAllNum/W9609>)

Coupling Facility LPs Using Dynamic ICF Expansion

Important Note

Dynamic ICF Expansion and Enhanced Dynamic ICF Expansion across ICFs is not recommended for coupling facilities that are used for primary production workload. All processors assigned to the coupling facility partition should be dedicated to that logical partition if it is used for primary production workload.

Using dedicated ICF CPs and shared general purpose CPs for coupling facility LPs has the following considerations:

- The dedicated ICF CPs have the same considerations as listed earlier under “Coupling Facility LPs using dedicated Central Processors (CPs) or Dedicated Internal Coupling Facility (ICF) CPs”.
- RMF reports the data for a dynamic ICF expansion partition on two lines; one line to report the dedicated ICF CPs and a second line to report the shared general purpose or shared ICF CPs.
- The processing weight specified for a coupling facility LP using dynamic ICF expansion functions very much the same as with a logical partition using only shared CPs. The only difference is the processing weight is divided among the shared general purpose or shared ICF CPs in the coupling facility rather than all the CPs in the logical partition.
- You should **not** cap the processing weight for a coupling facility LP using dynamic ICF expansion. If you do, LPAR attempts to support the cap but may not be successful and the system can be less responsive.
- All shared general purpose or shared ICF CPs of a dynamic ICF expansion coupling facility LP will use dynamic coupling facility dispatching on those shared CPs for their dispatching. The DYNDISP coupling facility command cannot be issued for a dynamic ICF expansion coupling facility. The value will be forced ON in this configuration.
- Each shared general purpose or shared ICF CPs defined and online in a coupling facility LP using dynamic ICF expansion will incur some overhead from a system perspective. Each logical CP runs as an independent server of coupling facility requests to the coupling facility using dynamic CF dispatching. Since the resources used by each shared general purpose CP in this configuration subtract from what could otherwise be used to run customer workloads, the number of shared logical CPs defined to a coupling facility LP should be minimized.

- Synchronous requests to coupling facilities from z/OS LPs that share general purpose CPs with a coupling facility that uses dynamic ICF expansion are **not** treated internally to the machine as asynchronous requests as is the case with a coupling facility with only shared CPs. Since there are dedicated ICFs available in the dynamic ICF expansion coupling facility, there is no automatic internal conversion of synchronous to asynchronous requests as the result of this configuration.
- It is recommended that the processing weight for the shared logical CPs of a coupling facility using dynamic ICF expansion be set significantly higher (25 times higher) than the weight for all other logical processors that are sharing the same physical processors. On the occasions when the shared CP is needed by the CF, the best response time is achieved when this CP is dispatched as quickly as possible. To help ensure that this occurs, the shared coupling facility CP should be assigned a weight which is much higher than that of any other logical partition's shared CPs. Since the shared CPs in the coupling facility will automatically throttle back on its CP usage, defining a higher weight should not have a negative impact on normal operations. Should the shared CPs be needed for a takeover scenario or workload peak, you would want the coupling facility to be able to respond quickly, effectively running as a dedicated CP for the duration of the required workload peak if the peak grows enough to warrant this.

For example, on a z9 EC with 1 ICF installed, one z/OS LP is defined with 8 shared CPs and the LP processing weight is 80 (10 per CP). Another coupling facility LP is defined with 1 dedicated ICF and one shared general purpose CP with a weight of 250. This allows the z/OS LP to normally get almost all of the 8 shared CPs until such time that the request rate to the coupling facility causes the shared CP of the coupling facility to start processing commands.

Processor Considerations for LPs with Multiple CP Types

You can define a logical partition to use one or more Integrated Facility for Applications (IFAs) and/or Integrated Information Processors (zIIPs) with z/OS 1.6 or higher with either of the following combinations:

- One or more dedicated general purpose CPs and one or more dedicated IFAs/zIIPs
- One or more shared general purpose CPs and one or more shared IFAs/zIIPs

The mode specified for the logical partition must be set to "ESA/390" in order to allow the definition of IFAs or zIIPs to the logical partition as stated above.

Dedicated Central Processors

An LP can have CPs dedicated to it. When an LP that uses dedicated CPs is activated, a physical CP is assigned to each defined logical CP. The LP then has exclusive use of its physical CPs.

The physical CPs that belong to an LP that uses dedicated CPs are always available for its use, but the capacity that is not used cannot be used by other LPs.

For coupling facility considerations, see "Processor Considerations for Coupling Facility LPs" on page 3-30.

Use the Customize/Delete Activation Profiles task available from the CPC Operational Customization Tasks list to open a reset or image profile to define dedicated CPs for an LP. The Dedicated field is located in the Logical processor assignment group box on the Processor page for the LP.

Suitable Workloads

Workloads best suited for logical partitions that use dedicated processors are those that maintain a fairly even external throughput rate (ETR) while using most of the capacity of the logical partition. Logical partitions with timing dependencies might require dedicated processors.

The installation goals for a logical partition that uses dedicated processors should be similar to the goals of the processor complex. For example, if the goal is an average processor utilization rate of 70%–85%, then the same goal should be used for a logical partition that uses dedicated processors.

Shared Central Processors

LPs can share CPs with other LPs. A sharing LP does not have exclusive use of the physical CPs. There is no correspondence between the logical CPs in a sharing LP and the physical CPs on which the logical CPs are dispatched (except on a one-way CPC). A logical CP can be dispatched on any physical CP and, on successive dispatches, the logical CP can be dispatched on different physical CPs.

The number of CPs available for use by sharing LPs is determined by adding the number of CPs already assigned to active, dedicated LPs and subtracting that sum from the total number of physical CPs available.

The total of all logical CPs for all sharing LPs can be larger than the number of physical CPs serving the sharing LPs. For example, if four LPs are active on a six-way CPC and each LP is defined to have four CPs, the total number of online logical CPs is 16.

For coupling facility considerations, see “Processor Considerations for Coupling Facility LPs” on page 3-30.

Use the Customize/Delete Activation Profiles task available from the CPC Operational Customization Tasks list to open a reset or image profile to define shared CPs for an LP. The Not dedicated field is located in the Logical processor assignment group box on the Processor page for the LP.

Suitable Workloads

Workloads best suited for sharing logical partitions are those that have a widely fluctuating ETR or would not fit well into the capacity of a dedicated logical partition. If a workload can use only a small portion of the capacity of a dedicated logical partition, redefine the logical partition to use shared processors to free the available capacity for use by other logical partitions.

A workload with a widely fluctuating ETR would experience peaks and valleys in its processor-utilization curve. Such fluctuations can occur over extremely short periods of time (minutes or seconds). This type of workload could take advantage of the time- and event-driven dispatching available. With event-driven dispatching, a sharing logical partition receives the resources required as needed and leaves the capacity free for other logical partitions when not needed. See “Wait Completion” on page 3-44.

When combining workloads on a processor complex by means of logical partitions, examine their average and peak requirements carefully. If the workloads fluctuate over very short intervals, the total capacity of the system must meet the sum of the average requirements for each workload. If processor utilization fluctuates over

longer periods, and the peak utilization periods for these workloads occurs simultaneously, then the total capacity of the logical partitions must meet the sum of the peak requirements for each workload.

Sharing logical partitions that use event-driven dispatching are better able to maintain high transaction rates with fluctuating demand while being responsive. However, the ITR for a sharing logical partition is lower than the ITR for a dedicated logical partition.

The capability to limit CPU usage for any or all logical partitions with shared processors is provided by the PR/SM capping function. The capping function enhances LPAR workload balancing controls for environments with a requirement that the CPU resources for a logical partition be limited. Capped logical partitions are recommended for use when CPU resources must be limited for business reasons (in accordance with a contract), or when the impact that one logical partition can have on other logical partitions needs to be limited.

Processing Weights

An LP with dedicated CPs is not affected by processing weights.

Processing weights are used to specify the portion of the shared CP resources allocated to an LP. Although PR/SM always manages sharing LPs according to the specified processing weights, there are times when an LP will receive either more or less than its processing share:

- An LP will receive more than its processing share when there is excess CP capacity, provided it has work to do and other LPs are not using their share.
- An LP will receive less than its processing share when its workload demand drops below the capacity specified by its weight.
- An LP will not receive more than its processing share when the CP resources for that LP are capped.

The recommended procedure is to specify processing weights to satisfy the peak requirements of the LPs.

Use the Customize/Delete Activation Profiles task available from the CPC Operational Customization Tasks list to open a reset or image profile to define the processing weight for a shared LP. The Processing weight field is located in the Logical processor assignment group box on the Processor page for the LP.

Use of Processing Weights: As an example, consider a system with 6 CPs and 3 LPs defined as follows:

LP Name	Logical CPs	Weight
-----	-----	-----
ZVSE	1	300
ZOSTEST	6	100
ZVM	2	900

Processing weights can range from 1 to 999 (weights of less than 2% difference are not considered significant) and are used as follows:

- The processing weights for all active, sharing LPs are added together. This total is considered to be 100% of the processing resource available to shared CPs.
- The share of processing resources for each LP is calculated by dividing the processing weight for each sharing LP by the total processing weight. For example, at peak CP utilization levels, the dispatcher allocates shared processing resources to each of the LPs as follows:

ZVSE	300/1300 = 23.1%
ZOSTEST	100/1300 = 7.7%
ZVM	900/1300 = 69.2%

- The share of processing resource for each online logical CP is calculated by dividing the share for each LP by the number of online logical CPs. The share for each logical CP is as follows:

ZVSE	23.1/1 CP = 23.1%
ZOSTEST	7.7/6 CPs = 1.3%
ZVM	69.2/2 CPs = 34.6%

These percentages are used to determine preemption priority for I/O interruptions. A lower priority logical CP can be preempted when an I/O interruption is pending for a higher priority logical CP when the following occurs:

- The higher priority logical CP is further behind in its share, or
- The higher priority logical CP is not as far ahead of its share as the lower priority logical CP.

For example, the lower priority LP is receiving 15% more than its processing share, and the higher priority LP is receiving 10% more than its processing share.

As long as there is excess CP capacity, processing weights have no effect on the CP resources consumed. Weights affect processing when the number of logical CPs that need processing time is greater than the number of physical CPs available.

Processing Weights and Shared CP, ICF, IFA, IFL and zIIP Processors: Shared general purpose, ICF, IFA, IFL and zIIP processors are each managed as separate “pools” of physical resources. The processing weights assigned to logical partitions using shared ICF, IFA, IFL, zIIP or general purpose processors are totaled and managed together only with the total weights from all of the logical partitions using the same processor type. The calculations shown in previous examples as well as the examples to follow are done independently for general purpose, ICF, IFA, IFL, or zIIP processors, on a machine that has them.

Processing Weights for Logical Partitions with Multiple Shared CP Types:

When a logical partition is defined to use one or more shared general purpose CPs and one or more shared IFAs/zIIPs, each of the types of logical processors is managed independently. The shared general purpose CPs compete with all other shared general purpose CPs defined in other logical partitions in the configuration. The IFAs compete with all other shared IFAs defined in other logical partitions. The zIIPs compete with all other shared zIIPs defined in other logical partitions. General purpose, IFA, and zIIP processors each have a separate processing weight specified. Note that if WLM weight management is being used for such a logical partition, only the weight of the shared general purpose processor portion of the logical partition will be altered by WLM. The specified weight for an LP’s IFA/zIIP processors is unaltered by WLM.

Effects of Processing Weights: Several implications are derived from the rules described above. First, every time a sharing LP is activated, the share of all other active LPs, using the same processor types, changes. This happens because the total of the processing weights has changed.

Because the processing share of LPs can vary, the actual utilization reported by monitors such as RMF can be different from the weights. In systems at less than 100% utilization, some LPs could receive greater than their share if other LPs are not using their share. The number of CPs defined also affects the maximum

resource allocated to that LP. For example, an LP defined to have two CPs on a three-way CPC can never be allocated more than 67% of the CP resources no matter what its processing weight.

Capping Processing Weights: The PR/SM capping function provides the capability of limiting CPU resource usage for one or more LP. The relative processing weight of an LP is its capping value.

A capped LP running at its cap does not have access to the CP resources that are not utilized by other LPs. However, CP resources that are not used by a capped LP can be used by other LPs. Equitable distribution of CP resources is maintained.

Capping values can be dynamically adjusted. The capping function can be turned on and off, and provides the capability of specifying capping for individual LPs without a re-IPL of the LP.

Use the Change LPAR Controls task from the CPC Operational Customization Tasks list to change the Capped setting for the specific partition. Checking the Initial Capping check box turns the capping function on. Select the Save Running System push button to have this option take effect immediately for an active partition. (Save and Change will change the running system and update the partition's profile. Save to Profiles will not change the running system; it will just save the new definition to the partition's profile.) Alternatively, if you do not need to change a running system, the Customize/Delete Activation Profiles task available from the CPC Operational Customization Tasks list can be used to open a reset or image profile to cap processing weight for an LP. This change would take effect when the partition is activated.

An LP's relative weight is divided by the number of shared logical CPs online for the LP to give the share for each logical CP. The goal of the LPAR dispatcher is to give each logical CP its share of the total relative weight. Capping is done on a logical CP basis.

An LP's share of CP resources is determined by its weight. The combined processing weights for all active LPs are considered to be 100% of the available shared CP resources. The activating and deactivating of LPs changes the amount of CP resources available to LPs, making the percentage of CP resources requested for each active LP a relative one, and not a fixed percentage of CP resources.

Note: If an extremely low processing weight is specified for a capped logical partition, tasks such as Reset Clear, Activate, and Deactivate may fail due to a time-out. To prevent this problem, avoid use of capping in conjunction with low processing weight. A preferable solution is specification of a processing weight that results in a 1/10 share (the suggested minimum) or greater of one physical CP for a logical CP. If the extremely low weight cannot be avoided, temporarily turn capping off for the logical partition prior to activating, deactivating, resetting, or loading the logical partition. Restore the cap following completion of these operations.

Enforcement of Processing Weights

Processing Weight Management

PR/SM LPAR enforces LP processing weights as follows:

- For LPs with processor resource capping, PR/SM LPAR enforces the processing weights to within 3.6% of the LP's physical CP share for logical CPs entitled to

1/10 or more of one physical CP. Typically, PR/SM LPAR will manage processing weights to within 1% of the LP's physical CP share. See "Example 1. With Processor Resource Capping."

- For LPs **without** processor resource capping, PR/SM LPAR enforces the processing weights to within 3.6% of the LP's physical CP share for logical CPs entitled to 1/2 or more of one physical CP. Typically, PR/SM LPAR will manage the processing weights to within 1% of the LP's physical CP share. See "Example 2. Without Processor Resource Capping" on page 3-41.
- If a logical CP falls outside the enforceable ranges for logical CPs entitled to less than 1/10 of a physical CP using capping or less than 1/2 of a physical CP not using capping, PR/SM LPAR enforces the processing weights to within 3.6% of the total capacity of the shared physical CP resources. However, PR/SM LPAR should typically manage the processing weights to within 1% accuracy.
- Unused CP cycles to which a logical CP is entitled are made available to other logical CPs in proportion to the weights set for the other logical CPs.
- An uncapped coupling facility LP with shared CPs and coupling facility channels defined honors its processing weight up to a point. PR/SM LPAR attempts to help ensure that each logical CP defined to a coupling facility LP gets at least 1 run time interval of service per every 100 milliseconds. For example, for a typical 12.5 millisecond run time interval, each logical CP gets 1/8th of a physical CP.

Processing Weight Management Examples

For the formulas used in the following two examples, see "Processing Weight Management Formulas" on page 3-41.

Example 1. With Processor Resource Capping: In the following example:

- Six physical CPs are online
- All LPs are capped
- All LPs have sufficient workload demand to use their shares

Table 3-3. PR/SM LPAR Processor Weight Management with Processor Resource Capping

LP	LCPs Online	Weight (1)	Weight per LCP	LCP % of PCP	3.6% Share	3.6% Total	Resulting Weight Range	Resulting Utilization Range
A	6	500	83.3	50	3.0	–	80.3–86.3	48.2%–51.8%
B	3	480	160.0	96	5.8	–	154.2–165.3	46.3%–49.7%
C(2) C(3)	3	20	6.7	4	0.2	– 12	6.5–6.9 0.0–18.7	2.0%– 2.1% 0.0%– 5.6%
Total Capacity of the Shared PCP Resources		1000 / 6 PCPs = 166.7 = PCP capacity						
Legend: LCP Logical CP PCP Physical CP								
Notes: 1. Use the Customize/Delete Activation Profiles task available from the CPC Operational Customization tasks list to open a reset or image profile to set the processing weight for an LP. The Processing weight field is located on the Processor page for the LP. 2. The logical CPs for LP C (2) fall outside the enforceable range because each logical CP's share is only 4% of a physical CP. 3. The LP's range is bounded by the line shown for LP C (3). Typically, though, even this configuration will see results comparable to the line shown for LP C (2).								

Example 2. Without Processor Resource Capping: In the following example:

- Six physical CPs are online
- No LPs are capped
- All LPs have sufficient workload demand to use their shares

Table 3-4. PR/SM LPAR Processor Weight Management without Processor Resource Capping

LP	LCPs Online	Weight (1)	Weight per LCP	LCP % of PCP	3.6% Share	3.6% Total	Resulting Weight Range	Resulting Utilization Range
A	6	500	83.3	50	3.0	—	80.3–86.3	48.2%–51.8%
B	3	300	100.0	60	3.6	—	96.4–103.6	28.9%–31.1%
C (2)	3	200	66.7	40	2.4	—	64.3–69.1	19.3%–20.7%
C (3)						12	54.7–78.7	16.4%–23.6%
Total Capacity of the Shared PCP Resources		1000 / 6 PCPs = 166.7 = PCP capacity						
Legend: LCP Logical CP PCP Physical CP								
Notes: 1. Use the Customize/Delete Activation Profiles task available from the CPC Operational Customization tasks list to open a reset or image profile to set the processing weight for an LP. The Processing weight field is located on the Processor page for the LP. 2. The logical CPs for LP C (2) fall outside the enforceable range because each logical CP's share is only 40% of a physical CP. 3. The LP's range is bounded by the line shown for LP C (3). Typically, though, even this configuration will see results comparable to the line for LP C (2).								

Processing Weight Management Formulas: The following formulas were used to compute the values in the previous two examples:

$$\text{Weight per LCP} = \frac{\text{LPCTL Weight}}{\text{Number of LCPs online}}$$

$$\text{LCP percent of PCP} = \frac{\text{LPCTL Weight}}{\text{Total of LPCTL Weights}} \times \frac{\text{Number of PCPs online}}{\text{Number of LCPs online}}$$

$$3.6 \text{ percent Share} = \text{Weight per LCP} \times 3.6 \text{ percent}$$

$$3.6 \text{ percent Total} = \frac{\text{Total Weight} \times 3.6 \text{ percent}}{\text{Number of LCPs online}}$$

$$\text{Resulting Weight Range} = \text{Weight per LCP} \pm 3.6 \text{ percent Share}$$

or

$$\text{Resulting Weight Range} = \text{Weight per LCP} \pm 3.6 \text{ percent Total}$$

$$\text{Resulting Utilization Range} = \text{Resulting Weight Range} \times \frac{\text{Number of LCPs online}}{\text{Total Weight}}$$

Maintaining the Same Relative Percentages of CPU Resources

To maintain the same relative percentage of CP resources requested for a capped LP, processing weights should be readjusted immediately prior to, or immediately after, the activation or deactivation of an LP.

Processing weight values for use when specific LPs are activated or deactivated should be calculated in advance, and be readily available. Otherwise, it is recommended that a convenient method be developed for changing the processing weight values to readjust relative shares after an LP is activated or deactivated.

For example, if the sum of the weights of the active LPs totals 100, then the sum of the relative weights of the active LPs also totals 100. This provides an easy means for adjusting weights upon the activation or deactivation of LPs. Another good approach to maintaining the desired share for a capped LP is to also readjust the processing weights for LPs with the capping function turned off, as shown in Table 3-5 for the LP ZOSTEST.

Table 3-5. Example of Maintaining Relative Weight of a Capped Logical Partition

LP Name	Four LPs Active			Three LPs Active		
	Status	Weight	Capped	Status	Weight	Capped
ZVSE	A	30	No	D	—	—
ZOSPROD	A	40	No	A	64	No
ZOSTEST	A	20	Yes	A	20	Yes

When the sum of all the relative weights is maintained at 100, it is easy to recalculate weights when an LP is deactivated. After deactivating ZVSE, the weight for ZOSPROD can be changed to 64 to maintain the same relative weight of 20 for ZOSTEST, the capped LP.

Capping in a Single Logical Partition: In order to use capping for an LP on a CPC where there is a need for only one active LP using shared CPs, you must define and activate a second “dummy” LP. The dummy LP must also be defined as using shared CPs. The weights of the two LPs can be adjusted to attain the desired cap for the one LP that will actually be used.

The “dummy” LP does not have to be capped. In most cases, the “dummy” LP does not need to have anything IPLed into it. If nothing is IPLed into the “dummy” LP, the “dummy” LP will not use any CP resources. In some situations where the single capped LP runs an extremely CP-intensive workload, it may be necessary to run a program in the “dummy” LP to smooth distribution of CP resources to the LPs. The program can be a simple branch loop that spins to waste time. Without this program running in the “dummy” LP, the CP-intensive, capped LP can experience a lurching effect with its dispatched times. It will be capped to its weight properly, but it could get all of its allotted CP resource quickly and then wait for a period of time before it can run again.

If nothing is to be IPLed into the “dummy” LP, the “dummy” LP can be defined and activated with no channel paths or devices. A definition with no devices is desirable to prevent control unit logical paths from being established for the “dummy” LP. See “Managing Logical Paths for ESCON and FICON Channels” on page 2-7.

Capping and Disabling Event-Driven Dispatching in a Single LP

Configuration: If you want to disable event-driven dispatching and enable capping in a single LP configuration, you must do **both** of the following:

- Define a “dummy” LP (a “dummy” LP configured with 1 logical CP is sufficient)
- Make sure that the logical CP in the “dummy” LP is in a fully enabled wait state (the first four bytes of the PSW must be X'030E000')

If you have configured the “dummy” LP with more than 1 logical CP (this is **not** recommended), you will have to put each logical CP in a fully enabled wait state. (Alternatively, each logical CP in the “dummy” LP can be running a branch loop to waste time.)

With event-driven dispatching disabled and the “dummy” LP in a fully enabled wait state, PR/SM LPAR will dispatch the logical CPs on the basis of time intervals according to the specified weights.

This situation does **not** apply if either of the following is true:

- Event-driven dispatching is enabled
- Event-driven dispatching is disabled and more than one LP is activated and running an operating system

Use the Customize/Delete Activation Profiles task available from the CPC Operational Customization tasks list to open a reset profile to enable or disable event-driven dispatching for an LP. On the Options page, select **Do not end the time slice if a partition enters a wait state** to disable event-driven dispatching for the LP. To enable event-driven dispatching for the LP, do not select this option.

Processor Running Time

The processor running time is the length of continuous time (determined by the dispatch interval) allowed for the dispatch of a logical CP.

When the processor running time is dynamically determined by the system, the calculation of the default running time is performed separately for general purpose and ICF CPs. All logical processors that are shared, either general purpose or ICF, will be assigned a default running time but the values used for each type of processor may be different. If a user-defined run time is specified, the value applies to all shared general purpose and shared ICF CPs.

The default value is dynamically calculated and changes when the number of active, scheduled logical CPs changes.

The default running time is determined using the formula:

$$\frac{25 \text{ milliseconds} \times (\text{number of physical shared CPs})}{(\text{total number of logical CPs not in stopped state for all sharing LPs})}$$

The default value is used whenever the processor running time is dynamically determined by the system. The run-time value can change whenever an LP is activated or deactivated and when a logical CP stops or starts (for instance, when a logical CP is configured online or offline). The default processor running time is limited to the range of 12.5 to 25 milliseconds.

The logical CP might not use all of its run time because it goes into a wait state. With event-driven dispatching, when a logical CP goes into a wait state, the physical CP is reassigned to another logical CP ready to do work. When a logical CP does not go into a wait state during its run time, it loses the physical CP when it reaches the end of its run time. Therefore, an LP with CP-bound work cannot permanently take control of the physical CPs.

You can choose to set the runtime value yourself. However, when event-driven dispatching is enabled, it is generally recommended that the processor running time be dynamically determined. If event-driven dispatching is disabled, you should consider setting runtime values of 2 to 8 milliseconds. The recommended procedure is to start by using the default processor running time. That value should be acceptable when all sharing LPs have similar proportions of interactive work; for example, two LPs each running 40%–70% of interactive work.

Adjustments to the runtime value might be necessary when one sharing LP contains a large proportion of CP-bound work and the other sharing LPs contain only short, interactive transactions. Degraded response time in one LP can indicate that the runtime value should be reduced to decrease the length of continuous time given to CP-bound work. The run-time value should be decreased by approximately 10% several times over several days while monitoring performance carefully. The processing weights should also be adjusted to favor the interactive LP. See “Processing Weights” on page 3-37.

Use the Customize/Delete Activation Profiles task available from the CPC Operational Customization tasks list to open a reset profile to define processor running time. Select either **Dynamically determined by the system** or **Determined by the user** on the Options page for the LP. See Figure 3-16 on page 3-57.

Wait Completion

Shared CP resources can be divided on either an event-driven basis or a time-driven basis. The RMF Partition Data Report refers to these two forms of dispatching as Wait Completion=No (event-driven dispatching) and Wait Completion=Yes (time-driven dispatching). When the resource is divided on an event-driven basis, switching between users is done as the result of some event, such as a wait or an interruption. When the resource is divided on a time-driven basis, switching is done after a predetermined interval of time. The resource is given to a user for a fixed interval and then given to another user for a subsequent fixed interval. Either of these methods can divide CP resources between sharing LPs.

You can select either of these methods by using the Customize/Delete Activation Profiles task available from the CPC Operational Customization tasks list to open the reset profile for the CPC. On the Options page, you can select **Do not end the time slice if a partition enters a wait state** to turn off event-driven dispatching and cause shared CP resources to be distributed on the basis of time intervals and according to the weights specified. Not selecting it (the default) requests event-driven dispatching. Event-driven dispatching also has an element of time-driven dispatching because the processor running time limits the time a logical CP can be continuously dispatched. This setting does not affect LPs that use dedicated CPs.

Event-driven dispatching has two advantages that are of special interest. The first advantage is a flexibility that allows the system to better meet the peak transaction rates of an LP. The event-driven dispatcher has the flexibility to manage CP resources to compensate for fluctuations in demand among the LPs instead of distributing CP resources in a rigid, predetermined manner.

The second advantage is improved responsiveness over time-driven scheduling through I/O interruption and wait detection. With event-driven dispatching, the time that one logical CP spends in a wait state can be made available to other logical CPs.

When event-driven dispatching is disabled, all logical CPs are given an entire dispatch interval even if they have no work to perform and even if they spend that time in a wait state because LPs are considered to be busy for 100% of their running time. Therefore, processing weights are always in effect.

Generally, response time is much better when event-driven dispatching is enabled and ITR is slightly better when event-driven dispatching is disabled. However, event-driven dispatching is recommended because improved responsiveness and the ability to use wait time more than compensates for a small decrease in ITR.

The wait completion setting applies to both shared general purpose and shared ICF processors. However, event-driven dispatching is always on for shared processors (general purpose or ICF) of a CF logical partition defined to use dynamic ICF expansion independent of the setting on the panel.

Workload Manager LPAR CPU Management of Shared CPs

WLM's LPAR CPU Management component, together with the LPAR clustering technology of the System z9 servers, provides the ability to dynamically manage workloads within an LPAR cluster comprised of multiple logical z/OS images on a single System z9 server. Each LP is assigned a **transaction** goal (desired response time) and an **importance** level. WLM monitors how well each LP is achieving its goals. A donor/receiver approach is utilized to reapportion CPU resources between LPs in the cluster. When WLM LPAR Weight CPU Management decides to change the weight of an LP, it adjusts the receiver LP and the donor LP by a percentage of the current weight of the receiver. WLM takes resources away from an LP that is over-achieving its target or has a workload that is less important (as defined by the installation). Any resource given to a particular LP is taken away from another LP in the LPAR cluster. LPs whose workloads are of the same importance level should all have similar performance indexes (a measure of how closely the workload is meeting its defined goal).

One can think of the entire LPAR cluster as having a total processing weight. The total weight for an LPAR cluster is the sum of all the initial processing weights of all the LPs that have joined the cluster. As a new logical partition joins the cluster, its initial processing weight is added to that of the cluster. Though weights are adjusted from one LP to another, the total weight for the cluster is consistent. When an LP leaves the LPAR cluster, as when it is either system reset, deactivated, or re-IPL'd, its initial processing weight, which it had been contributing to the LPAR cluster, is removed from the total weight available to the cluster. The weight removed from the cluster is not necessarily equal to the current weight for the exiting LP.

The optional minimum and maximum processing weights for an LP govern how much flexibility WLM has in adjusting weights from one LP in the cluster to another. The installation should assign a reasonably wide range of processing weights to each WLM managed LP. Assigning the same value for initial, minimum, and maximum weights effectively disables WLM LPAR CPU Management of processor weights.

Though logical CPs of WLM managed LPs may need to be soft-capped (as for workload pricing, see "Workload Charging by Soft-Capping to a Defined Capacity" on page 3-46), initial capping (traditional hardware capping) of these LPs is disallowed. Similarly, if an LP is WLM managed, its logical CPs must be shared, not dedicated. For more information regarding internals of WLM CPU Management, see redbook *z/OS Intelligent Resource Director*, SG24-5952.

Workload Charging by Soft-Capping to a Defined Capacity

Workload charging introduces the capability to pay software license fees based on the size of the LP the product is running in, rather than on the total capacity of the CPC. The capability is enabled by the LPAR clustering technology of the System z9 servers together with the License Manager component of z/OS. Each LP is assigned a **defined capacity** by the installation in terms of Millions of Service Units (MSUs).

WLM helps ensure that the rolling 4-hour average CPU utilization for the LP does not exceed this amount by tracking the CPU utilization for the logical partition. If the 4-hour average CPU consumption of the LP exceeds the defined capacity of the LP, WLM dynamically activates LPAR capping (soft-capping). When the rolling 4-hour average dips below the defined capacity, the soft-cap is removed.

WLM *will not* dynamically adjust the defined capacity for an LP. This is the responsibility of the installation. If an LP consistently exceeds its defined capacity, the license certificates and the defined capacity of the LP should be adjusted to reduce the amount of time the LP is soft-capped. If you have a configuration where the LP weights move significantly from one LP to another according to shift, then you must license the products in each LP at the highest capacity that will be used by that LP.

Recommendations on Setting Up an LPAR Cluster

- An LPAR cluster is a collection of two or more logical partitions, on a particular CPC, that are part of the same parallel sysplex. LPAR clusters do not span CPCs as do parallel sysplexes. Though the member LPs of an LPAR cluster will all be in the same parallel sysplex, all members of a parallel sysplex might not be members of the same LPAR cluster. A given parallel sysplex can have member LPs that belong to multiple LPAR clusters, each on a different CPC.
- Identify logical partitions on the CPC that will be in the cluster (members of the same parallel sysplex). A single CPC can have several LPAR clusters just as a single CPC can have many LPs, each having membership in a different parallel sysplex.
- IBM recommends allocation of shared CPs and enablement of WLM management for cluster members (see Note 1). The number of initially online CPs should be maximized to provide optimum flexibility to WLM. The number of reserved CPs defined should be the maximum allowed for an LP in your configuration minus the number of initially online CPs. See “Number of Central Processors” on page 3-28 for additional information on central processors.
- Establish an initial weight for each LP in the cluster. This will be the weight for the LP immediately after it is activated (see Note 2). Triple digit values should be used, wherever possible, for initial weights because WLM reapportions weights on a percentage basis. The total weight of the cluster will equal the sum of all the initial weights of its member LPs. Leave the minimum and maximum weights blank or make the range as wide as possible (optimally 1 to 999) to provide WLM maximum flexibility as it distributes CPU resource among the cluster members.
- Enable each LP in the cluster for WLM management.
- To enable DCM of managed channel paths for a logical partition, the name specified on the IOCLUSTER keyword for managed channel paths in the IOCDs must match the sysplex name of the software running in the logical partition. See “Dynamically Managed CHPIDs” on page 2-27 for more information on the IOCLUSTER keyword.
- Calculation to estimate the number of cache structures that can be supported:

The number of cache buffer data items that can be maintained locally in a Logical Partition is directly proportional to the number of on-line central storage pages in the LP. Each cache buffer or data item needs a local cache vector space bit. A heuristic value of 4 bits per on-line central storage 4K page is assigned by the system to each exploiting logical partition.

The number of cache buffers supportable is easily calculated. Multiply the number of on-line central storage pages, in the z/OS logical partition exploiting the cache vector space, by two to get the number of cache vector bytes provided.

For instance, if an exploiting logical partition has 32 GB of central storage on-line the amount of cache vector space provided would be $32 * 1024 \text{ (MB per GB)} * 256 \text{ (pages per MB)} * 2 \text{ (nybbles per byte)} = \text{number of bytes provided}$. For an LP with 32 GB 16777216 bytes, or 16 MB is provided. The size of the cache vector for an LP whose central storage definition includes the capability of using Dynamic Storage Reconfiguration will grow when the reserved storage amount is configured on-line.

Notes:

1. Logical partitions in a sysplex that have dedicated CPs can join a cluster but will not be enabled for WLM LPAR Weight and Vary CPU Management. They can derive I/O management benefits, however, from WLM's Dynamic CHPID Management capability.
2. Though the current weight of a logical partition can be changed (increased or decreased) by WLM once it joins an LPAR cluster, the initial weight will be restored when (on IPL) it leaves the cluster. Software can then rejoin the same or a different cluster, again donating its initial weight to the sum available for the entire cluster.

Enabling Management of Linux Shared Logical Processors by WLM's LPAR CPU Management Component

1. For the target Linux LP, use the **Options** page (see Figure 3-24 on page 3-71) of the Customize Image Profiles panel to specify the CP management cluster name of the intended LPAR cluster. Also select **Enable WorkLoad Manager** on the **Processor** page (see Figure 3-20 on page 3-63) of the Customize Image Profiles panel.
2. In the WLM policy statement, define a new service class giving it a velocity goal and an importance. Also, in the SYSH subsystem (create SYSH subsystem if one does not already exist), define a classification rule with the attribute **SY** (system name). Associate the system name of the target Linux logical partition and the service class you just defined.
3. IPL the Linux 2.4 kernel and provide a system name of the Linux logical partition by executing `insmod hwc_cpi.o system_name=xxxxxxx`, where `xxxxxxx` matches the system name communicated to WLM in Step 2.

Note: WLM is only able to manage Linux shared logical processors running on general purpose CPs. Management of Linux shared logical processors running on Integrated Facility for Linux (IFL) processors is not supported. It is important to realize that Linux mode partitions participating in a CP management cluster should be system reset through the Support Element/Hardware Management Console following a Linux shutdown command to help ensure accurate clean-up by WLM's LPAR CPU Management component.

Defining Shared Channel Paths

Before defining shared channel paths, consider the following:

- Only CTC, CNC, CBP, CFP, CIB, ICP, OSC, OSD, OSE, OSN, FC, FCV, FCP, and IQD channel paths can be shared. CVC and CBY channel paths **cannot** be shared. On CF only models, CBP, CFP, and ICP channel paths **cannot** be shared.
- A failure in a shared channel path or I/O device can affect more than one LP; therefore, critical I/O devices (for example, DASD containing vital data sets) still need multiple channel paths. You can provide multiple shared channel paths (up to 8) to critical I/O devices.
- Using shared channel paths does not reduce the number of logical paths needed at a control unit. A control unit requires a logical path for each active LP that can access I/O devices through a shared channel path.

There are three possible channel path modes:

shared	A channel path that can be configured online to one or more LPs at the same time. One or more LPs can access I/O devices at the same time using this channel path. Spanned channel paths are shared by LPs in multiple logical channel subsystems (CSSs). Unspanned channel paths can only be shared by LPs in the same CSS.
reconfigurable	An unshared channel path you can reconfigure offline from one LP, then online to another, within the same CSS. Only one LP can access I/O devices on this channel path at a time.
dedicated	An unshared and nonreconfigurable channel path. Only one LP can access I/O devices on this channel path.

You cannot mix unshared and shared CHPIDs to a device.

Channel Path Access and Candidate Lists

If a channel path is either shared or reconfigurable, you can specify which LPs have access to that channel path. Use the channel path access list with or without the channel path candidate list.

Channel Path Access List

An LP has **initial** access to a channel path if the LP is on that channel path's access list.

For the first power-on reset with an LPAR IOCDS, the access list defines which LPs will initially have the channel path configured online. Reconfigurable and shared CHPIDs may later be accessed by LPs not in the access list. Subsequent power-on resets with the same LPAR IOCDS will have any reconfigurations applied and the LPs may have access to channel paths other than those that were specified in the initial access lists for the channel paths. See "Channel Path Reconfiguration and Logical Partition Activation" on page 3-53 for more information on assigning channel paths to an LP.

Channel Path Candidate List

An LP can **gain** access to a channel path if the LP is on that channel path's candidate list. An LP is allowed to configure a channel path online if the LP is in that channel path's candidate list.

I/O Device Candidate List

The I/O device candidate list specifies the LPs which can access the device. You can use the I/O device candidate list to restrict LP access to I/O devices on shared channel paths. If you do not specify an I/O device candidate list, all LPs that share the channel paths, to which the device is attached, can have access to the device. For coupling facility devices, the device candidate list is not supported.

Procedure for Defining Shared Channel Paths

To share channel paths, use the following general procedure:

1. Select which I/O devices to access through shared channel paths
2. Select the LPs that will share channel paths
 - Specify the desired channel paths as shared
 - Use the access list and candidate list of the desired shared channel path to select which LPs can access that shared channel path.
3. For each I/O device that will be on a shared channel path, you can use the I/O device's candidate list to restrict which LPs can access that I/O device.
4. Make physical connections to the shared channel paths and the I/O devices.
5. Update the software I/O configuration for the control programs running in LPs that can access devices through shared channels.
6. Use IOCP or HCD to create an IOCDS that defines the I/O configuration, including shared channel paths, to the CPC channel subsystem. The channel subsystem includes all channel paths, control units, and I/O devices accessible by all LPs.

Communicating by Means of ESCON or FICON CTC

You can connect shared or unshared CNC or FCV channel paths to shared or unshared CTC channel paths for the purpose of CTC communication. The connected CNC or FCV and CTC channel paths can be on the same CPC (to communicate between the LPs on that CPC) or on different CPCs.

You can connect shared or unshared FC channel paths to shared or unshared FC channel paths for the purpose of CTC communication. The connected FC channel paths can be on the same CPC (to communicate between the LPs on that CPC) or on different CPCs.

CTC communications involving a shared channel path require the specification of control unit logical addresses. The control unit logical address identifies the MIF image ID number for the LP to which the channel path is to communicate. If the remote channel path is not shared, the logical address must be zero or not specified. If the remote channel path is shared, the logical address must equal the desired MIF image ID number for the LP that the shared channel path can access. If the remote channel path is a shared FC channel path and the target LP has a CSS ID other than zero, then the logical address must equal the combination of the desired CSS ID and the MIF image ID for the LP with which you want to communicate. For example, if the remote channel path is shared within CSS 1 and you want to communicate with the LP that has MIF image ID 5, specify logical address 15. You must define a control unit and a control unit logical address for every LP that you want to communicate with.

For additional information, see *zSeries ESCON or FICON Channel-to-Channel Reference*, SB10-7034.

Dynamic CHPID Management (DCM) Considerations

DCM CHPIDs used by WorkLoad Manager (WLM) to optimize I/O throughput across an LPAR cluster are identified in the IOCDs by specifying CHPARM equals 01 and IOCLUSTER equals name (where name is an 8-byte EBCDIC cluster identifier).

All DCM CHPIDs are inherently shareable by all LPs but reserved for use by the WLM enabled members of the specified cluster. At completion of an LP activation, its DCM CHPIDs will be in a deconfigured state. When an LP joins a cluster, the DCM CHPIDs in that cluster become available for use by WLM but are not brought online until the need for greater throughput arises. System reset of an LP that was a member of a cluster causes each of the DCM CHPIDs that were online to the LP to be deconfigured. For information on how to define DCM CHPIDs, see “IOCP Statements for ICP” on page 2-28. For allocation rationale of DCM CHPIDs, see redbook *z/OS Intelligent Resource Director*, SG24-5952.

I/O Priority Recommendations

Channel subsystem I/O priority queuing is used by z/OS WLM component to dynamically manage the priority of I/O operations for given workloads based on the performance goals for these workloads as specified in the WLM policy.

Channel subsystem I/O priority queueing is used by z/VM to manage the priority of I/O operations performed by guests. The VM Resource Manager adjusts guest I/O priorities based on the performance goals of the associated workloads. It is recommended that you establish a range of I/O priorities for z/VM logical partitions that is sufficient to enable effective discrimination among guest I/O requests of different importance.

In order to provide WLM the greatest flexibility in managing I/O requests across members of an LPAR cluster, it is highly recommended that you establish the same range of priorities for each member of an LPAR cluster. A range of eight priorities (from minimum to maximum) is optimum. If a range greater than eight is specified, only the top eight will be utilized by WLM.

Non-WLM managed LPs should be assigned I/O priorities according to their importance relative to the LPs that are members of LPAR clusters. Unless it is running z/VM (in which case the recommendations above should be followed), a non-WLM managed LP should be assigned equal values for minimum and maximum I/O priorities. For more information and examples of how I/O priorities are used by WLM to prioritize I/O requests across diverse workloads, see redbook *z/OS Intelligent Resource Director*, SG24-5952.

Security-Related Controls

You can define security-related controls for an LP.

Logical Partition Isolation

This control reserves reconfigurable unshared channel paths for the exclusive use of an LP. Channel paths assigned to an isolated LP are not available to other LPs and remain reserved for that LP when they are configured offline.

Use the Release task available from the CHPID Operations task list to release an unshared channel path from an isolated LP.

Use the Customize/Delete Activation Profiles task available from the CPC Operational Customization tasks list to open a reset or image profile to enable isolation for an LP. The Logical partition isolation selection is located on the Security page for the LP.

Using IOCP, you can control access to channel paths using the channel path candidate list. Access to I/O devices on shared channel path can be further controlled through the I/O device candidate list.

I/O Configuration Control Authority

This control can limit the ability of the LP to read or write any IOCDS in the configuration locally or remotely. LPs with control authority for the I/O configuration data can read and write any IOCDS in the configuration, and can change the I/O configuration dynamically.

Additionally, this control allows the OSA Support Facility to control OSA configuration for other LPs.

Use the Customize/Delete Activation Profiles task available from the CPC Operational Customization tasks list to open a reset or image profile to enable I/O configuration control for an LP. The Input/output (I/O) configuration control selection is located on the Security page for the LP.

Global Performance Data Control Authority

This control limits the ability of a LP to view CP activity data for other LPs. LPs with control authority for global performance data can view CP utilization data and Input/Output Processor (IOP) busy data for all of the LPs in the configuration. Additionally, gathering of FICON channel measurements requires selection of this parameter.

With the exception of an LP that is a member of a WLM Cluster, an LP without control authority for the performance data can view only the CP utilization data for that LP.

Use the Customize/Delete Activation Profiles task available from the CPC Operational Customization tasks list to open a reset or image profile to enable global performance data control for an LP. The Global performance data control selection is located on the Security page for the LP.

Note: An LP running a level of RMF that supports FICON requires control authority even if no FICON is installed.

Cross-Partition Authority

This control can limit the capability of the LP to issue certain control program instructions that affect other LPs. LPs with cross-partition authority can issue instructions to perform a system reset of another LP, deactivate any other LP, and provide support for the automatic reconfiguration facility.

The automatic reconfiguration facility permits a backup LP to deactivate a primary LP if a problem is detected in the primary LP. The backup LP can then configure online, storage resources that become available when the primary LP is deactivated. See “CPCs with the Sysplex Failure Manager (SFM)” on page 3-26.

Use the Customize/Delete Activation Profiles task available from the CPC Operational Customization tasks list to open a reset or image profile to enable cross-partition authority for an LP. The Cross partition authority selection is located on the Security page for the LP.

Dynamic I/O Configuration

Dynamic I/O configuration, available with z/OS and z/VM, provides the capability of changing the currently active I/O configuration. With dynamic I/O configuration, channel paths, control units, and devices of the currently active I/O configuration can be added, deleted, or modified without requiring a power-on reset and an IPL for the change to take effect. Changes made to the currently active I/O configuration can be saved and the IOCDS that reflects these changes can be written and made the active IOCDS. You cannot dynamically add/delete to the pool of MIF image IDs defined in the IOCDS but you can change the MIF image ID assigned to an LP.

Use the Customize/Delete Activation Profiles task available from the CPC Operational Customization Tasks list to open a reset profile to enable dynamic I/O configuration for the CPC. On the General page, select a dynamic IOCDS from the IOCDS list, then select the Dynamic page for the IOCDS and select the Allow dynamic changes to the channel subsystem input/output (I/O) definition check box.

Managing Dynamic I/O Configuration

For detailed information about changing the I/O configuration dynamically, refer to *z/OS Hardware Configuration Definition Planning*, GA22-7525, *CP Planning and Administration*, SC24-6083, and *z/VM I/O Configuration*, SC24-6100.

Planning and Operation Considerations

Guidelines for planning and operating in a dynamic I/O environment are detailed in the appropriate z/OS, and z/VM publications. The planning and operation considerations described in this section are additional guidelines that apply to a dynamic I/O environment.

Careful planning will avoid confusion when moving CHPIDs, and eliminate duplicate device situations should a backout be required when a power-on reset is performed.

- Prepare for a backout situation

Before changing the currently active I/O configuration, prepare for a possible backout situation. Record the current CHPID assignments and the planned I/O changes.

It is important to prepare for a backout situation. A backout situation occurs when changes are made to the I/O configuration but the changes are not saved prior to performing a power-on reset. If the changes are not saved, the CHPID assignments prior to the first dynamic I/O change will take effect.

In addition CHPID assignments after a backout will also reflect any changes made by hardware operator tasks or control program commands.

- Avoid CHPID reconfigurations concurrent with dynamic I/O changes

Do not perform hardware operator tasks and system control program commands to cause a CHPID to be moved from one LP to another or to give or take access to a shared channel path while dynamic I/O changes are being made. Use these commands only after dynamic I/O changes are saved.

Assigning Channel Paths to a Logical Partition

Channel paths are defined in the IOCDS. Channel paths that are specified as reconfigurable can be moved among LPs.

Channel paths that are specified as shared can be accessed by one or more LPs at the same time. Unshared channel paths that are defined to a LP as not reconfigurable are dedicated to that LP.

ESA/390 Logical Partitions

Channel paths assigned in the IOCDS to ESA/390 LPs can be shared, reconfigurable, or dedicated.

Coupling Facility Logical Partitions

Channel paths (CFP, CBP, CIB, and ICP) assigned in the IOCDS to coupling facility LPs can be online to only one coupling facility at a time. On System z9 models a coupling facility can be assigned CBP, CFP, CIB, or ICP channel paths.

Note: CBP, CFP, and ICP channel paths that are online to a single coupling facility are shareable by multiple noncoupling facility logical partitions.

Channel Path Reconfiguration and Logical Partition Activation

When a Configure on task or an equivalent system control program command is run successfully for a channel path in an **active** LP, the channel path is configured to the LP at that time.

When a successful Configure on task completes for a channel path in an LP that is **not active**, the channel path is not actually configured online to the LP at that time. Rather, the channel path is **targeted** to be configured online to that LP when the LP is activated.

When an LP is deactivated, all shared channel paths configured to it at the time of deactivation are targeted to be configured online to it when it is subsequently activated. Unshared channels that were last configured to this LP and were not yet reconfigured to another LP are also targeted to be configured online to this LP at its next activation. However, the targeting of a channel path may change prior to the next activation due to channel reconfiguration commands, dynamic I/O configuration changes, or POR.

Channel paths can also be targeted to be configured online to an LP by using dynamic I/O configuration. See “Dynamic I/O Configuration Effects on Channel Path Reconfiguration” for more details.

LPAR manages lists of targeted channel paths for each LP on an IOCDS basis so that all channel paths that are targeted to be configured online to an LP will be automatically configured online when that LP is activated. Exceptions to this rule are:

- Targeted CBP, CFP, CIB, and ICP channel paths online to another coupling facility are not automatically configured online when that LP is activated.

Note: CBP, CFP, CIB, and ICP channel paths can only be online to one active coupling facility LP at a time. If such a channel path is targeted to a coupling facility LP but is already online to another coupling facility LP, then it will be removed (deconfigured) from the activating LP.

- The targeted channel path is in single channel service mode or otherwise broken.

Dynamic I/O Configuration Effects on Channel Path Reconfiguration

If a channel path is dynamically added to the configuration using dynamic I/O configuration, all the LPs in the channel path access list (for a shared channel path) or the one LP in the channel path access list (for an unshared channel path) are

targeted to have this new channel path configured online. The dynamic I/O configuration change does not bring the channel path online.

The channel path will be configured online to the targeted LP when one of the following occurs:

- The system control program running in the targeted LP issues the appropriate reconfiguration command to configure the channel path online. For z/OS, this would be:

CF CHP(nn),ONLINE

For z/VM, this would be:

VARY ON CHPID nn

- The Configure on task is used to configure the channel path online to the targeted LP while the LP is active.
- A power-on-reset is done with the IOCDS that has the new dynamic changes defined in it without the changes being made active using HCD z/OS, or z/VM commands. Subsequent to the POR, activation of the targeted LP will configure the channel path online.

Automatic Channel Path Reconfiguration

LPAR records shared channel path configurations and unshared channel path reconfigurations and uses the information to modify the initial targeting of channel paths that are defined in the IOCDS. This information is maintained on an IOCDS basis.

When a particular IOCDS is used in a POR for the first time after it has been written, the definitions in that IOCDS are used to determine the assignment of channel paths to LPs according to the channel path access lists that are defined. All previous information about channel configuration associated with this IOCDS is discarded. The exception to this rule is when a newly written IOCDS is first used as part of a dynamic I/O change to the system. (For example, the new IOCDS is used as a result of a “Switch IOCDS for Next POR” action by HCD, or the new IOCDS is the target of the ACTIOCDS= parameter of the z/OS ACTIVATE command.) When a new IOCDS is used in this manner, the current state of channel configurations is preserved and immediately associated with the newly written IOCDS.

Over time, the list of channel paths targeted to be configured online to a given LP can be changed by system control program configuration commands, configure on tasks, or dynamic I/O configuration commands issued through HCD or z/VM. Similarly, reconfigurable unshared channel paths can be moved from one LP to another using the same commands; changing the owner of the unshared channel path. For activated coupling facility LPs, you can change the channel paths targeted to be configured online using coupling facility control code commands. Automatic channel path reconfiguration restores all of the latest changes for each POR with the IOCDS automatically.

Automatic channel path reconfiguration does **not** preserve the online/offline state of unshared channel paths (reconfigurable or dedicated). Rather, at POR time, all unshared channel paths are targeted to come online to the LP that last owned it. For dedicated channel paths, this owner never changes but for reconfigurable channel paths the owner can change and is remembered.

Following a POR, a channel path that is targeted to come online to a LP will be physically online to that LP and usable at the completion of activation for the LP.

Automatic Load for a Logical Partition

If you want the control program to be loaded automatically each time the LP is activated, select the Load during activation checkbox. Use the Customize/Delete Activation Profiles task available from the CPC Operational Customization tasks list to open a image profile. The Load during activation selection is located on the Load page for the LP. See Figure 3-26 on page 3-75.

Coupling Facility Logical Partitions

Do not support automatic load because the coupling facility control code is automatically loaded and made operational at LP activation. No IPL of an operating system is necessary.

z/OS LPs

Specify parameter (PARM) information if desired. Byte 2 and byte 3 of the PARM field are used to select an IOCONFIG member for this IPL if you do not want the default (00).

Defining Logical Partitions

Before using this section you should first read “Determining the Characteristics” on page 3-5. This section describes the panels, parameters, and tasks, you can use to define LP definitions.

Sample tasks and panels explained in this section reference tasks and panels available from the support element console of a System z9. However, detailed procedures for operator tasks and accurate task and panel references are left to the appropriate operator guides. See the *Hardware Management Console Operations Guide* and *Support Element Operations Guide*.

LP definitions are saved across a power-on reset and are used during each power-on reset. You can use reset, image, and load profiles to modify LP definitions. Use image and load profiles to modify LP definitions after the CPC has been activated. Use the Customize/Delete Activation Profiles task available from the CPC Operational Customization tasks list to open a reset, image, and load profiles.

You can change reset, image, and load profiles at any time. However, some of these definition parameters cannot affect the running system if the affected LP is currently activated. See “Changing Logical Partition Definitions” on page 3-84.

LP definition parameters fall into the following categories:

Global reset profile definitions

- Processor running time
- Automatic I/O interface reset
- LP automatic activation order
- Enable global input/output I/O priority queuing

General

- Logical partition identifier
- LP mode of operation
- Enable logical partition time offset

Processor characteristics

- Dedicated or shared general purpose CPs
- Internal dedicated or shared coupling facility (ICF) CPs
- Initial, minimum, and maximum weight
- Weight capping

- Number of general purpose, and IFA, IFL, zIIP, or ICF CPs

Security characteristics

- Global performance data control
- Input/output configuration control
- Cross partition authority
- Logical partition isolation

Storage

- Central storage
- Expanded storage

Time Offset

- Time offset

Load information

- Load during activation
- Load address
- Use of dynamically changed address
- Load parameter
- Use of dynamically changed parameter
- Time-out value

Cryptographic characteristics

- Candidate list
- Online list
- Control domain index
- Usage domain index

Note: Coupling Facility partitions cannot take on these characteristics.

Options

- Minimum and maximum input/output (I/O) priority queuing values
- Defined capacity

Global Reset Profile Definitions

Use the Customize/Delete Activation Profiles task available from the CPC Operational Customization tasks list to open a reset profile.

Options Page Definitions

Open the Options page to define the following LP characteristics:

- Processor running time
- Enablement of event driven dispatching
- Enable global input/output (I/O) priority queuing




Figure 3-16. Options Page, Reset Profile

Parameter Descriptions

Enable global input/output (I/O) priority queuing

Select this option to enable I/O priority queuing. Selecting this option causes I/O requests to be prioritized according to the values specified on the Options page of the Image Profile.

Automatic input/output (I/O) interface reset

Enables automatic system reset of a logical partition on behalf of a software initiated recovery action.

Dynamically determined by the system

Select this option to allow the system to dynamically calculate the length of CP time dispatched. The default is selected for this checkbox.

Determined by the user

Select this option to manually set the length of CP time dispatched.

Processor running time

Selecting the **Determined by the user option** requires that you enter a running time amount from 1–100 milliseconds.

Do not end the timeslice if a partition enters a wait state

When the processor running time is determined by the user through this profile, check this box to indicate whether logical processors lose their share of running time when their logical partition enters a wait state.

Partitions Page Definitions

Use the Customize/Delete Activation Profiles task available from the CPC Operational Customization tasks list to open a reset profile. Open the Partitions page to define the following LP characteristics:

- LP automatic activation order

Customize Activation Profiles: P0LXSM10 : D63JRESET : Partitions

P0LXSM10

D63JRESET

General

Storage

Dynamic

Options

CP/SAP

Partitions

LP06

LP02

LP04

LP08

LP10

LP12

LP14

LP16

Specify the order in which the logical partitions will be activated. If no order is specified for a partition, it will be activated.

Partition	Order
LP06	1
LP02	2
LP04	3
LP08	4
LP10	5
LP12	6
LP14	7
LP16	8

Figure 3-17. Partitions Page, Reset Profile

Parameter Descriptions

- Partition

LP name.
- Order

Enter a number indicating when the LP will be activated in the automatic activation order.

General

Use the Customize/Delete Activation Profiles task available from the CPC Operational Customization tasks list to open an image profile for an LP. Open the General page to define the following LP characteristics:

- Logical partition identifier
- LP mode of operation
- Clock type assignment

Customize Image Profiles: P0LXSM09 : LP01 : General

Tree View:

- P0LXSM09
 - LP01
 - General**
 - Processor
 - Security
 - Storage
 - Options
 - Load
 - Crypto

Form Fields:

- Profile name: LP01
- Description: This is the LP01 Image profile.
- Partition identifier: 1
- Mode: ESA/390, ESA/390 TPF, Coupling facility, LINUX only
- Clock Type Assignment:
 - ☒ Standard time of day
 - ☐ Logical partition time offset

Figure 3-18. General Page, Image Profile

Parameter Descriptions

Partition identifier

Enter a hex value (X'00' through X'3F') for the LP. This parameter identifies the LP and is used as the third and fourth hexadecimal digit of the operand stored by the Store CPU ID Instruction for each logical CP in the LP. The partition identifier must be unique for each active LP.

Mode

Select an LP operating modes from this scrollable list - ESA/390, ESA/390 TPF, Linux-Only, or coupling facility mode.

Logical partition time offset

Select this choice to set the logical partition's clock using an offset from the time of day supplied by its time source. Then use the Time Offset window to set the offset.




Figure 3-19. Time Offset, Image Profile

Parameter Descriptions

Offset Type or spin to the number of days, hours, and minutes you want to set as the offset from the time of day supplied by its time source. You can set an offset within the following range:

- 0 to 999 days
- 0 to 23 hours
- 0, 15, 30, or 45 minutes

Decrease system time value by the amount shown

Select this choice to set the logical partition's clock back from the time of day supplied by its time source by the number of days, hours, and minutes in the offset. Use this setting to provide a local time zone WEST of GMT.

Increase system time value by the amount shown

Select this choice to set the logical partition's clock ahead of the

time of day supplied by its time source by the number of days, hours, and minutes in the offset. Use this setting to provide a local time zone EAST of GMT or a date and time in the future.

Processor Characteristics

Processor Page Definitions

Use the Customize/Delete Activation Profiles task available from the CPC Operational Customization tasks list to open a reset or image profile for an LP. Open the Processor page to define the following LP characteristics:

- Dedicated or shared general purpose CPs
- Dedicated or shared internal coupling facility (ICF) CPs
- Number of initial and reserved processors
- Initial processing weight
- Initial weight capping
- Workload manager enablement
- Minimum processing weight
- Maximum processing weight

Customize Image Profiles: T25A : LP1 : Processor

Group Name: Group 1

Logical Processor Assignments

☐ Dedicated processors

Select	Processor Type	Initial	Reserved
☒	Central processors (CPs)	3	2
☒	Integrated facilities for applications (IFAs)	2	0
☒	System z9 integrated information processors (zIIPs)	3	0

Not Dedicated Processor Details for :

☐ CPs ☐ IFAs ☒ zIIPs

zIIPs

zIIP Details

Initial processing weight: 25 (1 to 999) ☐ Initial capping

☐ Enable workload manager

Minimum processing weight: 0

Maximum processing weight: 0

Buttons: Save, Copy Notebook, Paste Profile, Cancel, Help

Figure 3-20. ESA Mode Logical Partition with shared Central Processors (CPs), Integrated Facilities for Applications (IFAs) and System z9 Integrated Information Processors (zIIPs). There can be both an initial and reserved specification for each processor type. Each processor can have it's own processing weight(s). If WLM enablement or Initial Capping is selected for any processor then all processors types are WLM enabled or are initially capped, respectively. WLM enablement and Initial Capping are mutually exclusive.

Customize Image Profiles: T25A : LP1 : Processor

T25A

LP1

General

Processor

Security

Storage

Options

Load

Crypto

Group Name

Logical Processor Assignment

☐ Dedicated central processors
☐ Dedicated integrated facility for Linux
☐ Not dedicated central processors
☒ Not dedicated integrated facility for Linux

Not Dedicated Processor Details

Initial processing weight 1 to 999 ☐ Initial capping

☐ Enable workload manager

Minimum processing weight

Maximum processing weight

Number of processors - Initial Reserved

Save

Copy Notebook

Paste Profile

Cancel

Help

Figure 3-21. Customization for a Linux-Only Mode Logical Partition with shared Integrated Facilities for Linux (IFLs). There can be both an initial and reserved specification for the IFLs.

Customize Image Profiles: T25A : LP1 : Processor

Group Name: Group 1

Logical Processor Assignment

- ☐ Dedicated central processors
- ☐ Dedicated internal coupling facility processors
- ☐ Not dedicated central processors
- ☐ Not dedicated internal coupling facility processors
- ☒ Dedicated internal coupling facility processors and not dedicated central processors

Not Dedicated Processor Details

Initial processing weight: 10 (1 to 999) ☐ Initial capping

(Maximum of 11 initial central processors and/or 1 initial internal coupling facility processors.)
 (Maximum of 16 initial and reserved central processors and/or 16 initial and reserved internal coupling facility processors.)

Processor Type	Initial	Reserved
Dedicated internal coupling facility processors	3	0
Not dedicated central processors	1	0

Buttons: Save, Copy Notebook, Paste Profile, Cancel, Help

Figure 3-22. Customization for a Coupling Facility Mode Logical Partition with Dedicated Internal Coupling Facilities (ICFs) and shared Central Processors. There can be both an initial and reserved specification for the ICFs and the Central Processors.

Parameter Descriptions

Note: Depending on the processor page, (see Figure 3-20, Figure 3-21 and Figure 3-22) some of the following parameters may not be present.

Group Name If you choose to assign the processor to a group, select a defined group from the drop-down list.

Dedicated processors

Select this option if you want to select general purpose CPs, IFAs or zIIPs to be dedicated when the LP is activated. You can then specify the number of initial and reserved processors for each.

Not dedicated processor details

Select the processor (CPs, IFAs, or zIIPs) to display details such as Initial processing weight, Initial capping, and Enable workload manager.

Dedicated central processors

Select this option if you want the general purpose CPs that are allocated for the LP to be dedicated when the LP is activated.

Not dedicated central processors

Select this option if you want the general purpose CPs that are allocated for the LP to be shared when the LP is activated.

Dedicated integrated facility for Linux

Note: All processors assigned to a coupling facility partition should be dedicated to that logical partition if it is used for primary production workload.

If Integrated Facility for Linux (IFL) is supported and installed in the Central Processor Complex (CPC), select Dedicated integrated facility for Linux if you want an IFL processor dedicated to each logical processor.

Not dedicated integrated facility for Linux

If you want the logical processors to share not dedicated integrated facility for Linux (Integrated Facility for Linux (IFL) processors that are not already dedicated to other activated logical partitions when this logical partition is activated), select Not dedicated integrated facility for Linux.

Dedicated internal coupling facility processors

If internal coupling facility processors are supported by and installed in the Central Processor Complex (CPC), select Dedicated internal coupling facility processors if you want one dedicated to each logical processor.

Not dedicated internal coupling facility processors

Note: All processors assigned to a coupling facility partition should be dedicated to that logical partition if it is used for primary production workload.

If you want the logical processors to share not dedicated internal coupling facility processors (internal coupling facility processors that are not already dedicated to other activated logical partitions when this logical partition is activated), select Not dedicated internal coupling facility processors.

Dedicated internal coupling facility processors and not dedicated central processors

Note: Dynamic ICF Expansion and Enhanced Dynamic ICF Expansion across ICFs is not recommended for coupling facilities that are used for primary production workload. All processors assigned to the coupling facility partition should be dedicated to that logical partition if it is used for primary production workload.

If internal coupling facility processors are supported by and installed in the CPC, select Dedicated internal coupling facility processors and not dedicated central processors if you want to assign a combination of dedicated internal coupling facility processors and not dedicated central processors to the logical partition.

Dedicated and not dedicated internal coupling facility processors

Note: Dynamic ICF Expansion and Enhanced Dynamic ICF Expansion across ICFs is not recommended for coupling facilities that are used for primary production workload. All processors assigned to the coupling facility partition should

be dedicated to that logical partition if it is used for primary production workload.

If internal coupling facility processors are supported by and installed in the CPC, select Dedicated and not dedicated internal coupling facility processors and not dedicated central processors if you want to assign a combination of dedicated internal coupling facility processors and not dedicated internal coupling facility processors to the logical partition.

Initial processing weight

Enter a value between 1–999 to set the processing weight for the type of processor for an LP. The default value is 10.

Initial capping

Select this option to cap the CP resources for the type of processor for an LP. Capping has no effect on LPs with dedicated CPs.

Enable WorkLoad Manager

Select this option so that CPU and I/O resources can be managed by WLM using IRD clustering technology.

Minimum processing weight

Select this option to establish a minimum processing weight that WLM will allocate to the LP. Do not specify a value here unless you determine a true need for it in your configuration. Specifying a value here can needlessly constrain what WLM can do to optimize the management of your workload.

Maximum processing weight

Select this option to establish a maximum processing weight that WLM will allocate to this LP. Do not specify a value here unless you determine a true need for it in your configuration. Specifying a value here can needlessly constrain what WLM can do to optimize the management of your workload.

Security Characteristics

Use the Customize/Delete Activation Profiles task available from the CPC Operational Customization tasks list to open a reset or image profile for an LP. Open the Security page to define the following LP characteristics:

- Global performance data control
- Input/output (I/O) configuration control
- Cross partition authority
- Logical partition isolation




Figure 3-23. Security Page, Image Profile

Parameter Descriptions

Global performance data control

Select this option to allow the LP to view the CPU utilization data and the Input/Output Processor (IOP) data for all LPs in the configuration. Not selecting this option only allows the LP to view its own CPU utilization data. Additionally, gathering of FICON channel measurements requires selection of this parameter. The default is selected.

Note: An LP running a level of RMF that supports FICON requires control authority even if no FICON is installed.

Input/output (I/O) configuration control

Select this option to allow the LP to read or write any IOCDS in the configuration and to make dynamic I/O changes. Additionally, this parameter allows the OSA Support Facility for z/OS, z/VM, and z/VSE to control OSA configuration for other LPs. The default is selected.

Cross partition authority

Select this option to allow the LP to issue control program commands that affect other LPs; for example, perform a system reset of another LP, deactivate an LP, or provide support for the automatic reconfiguration facility. The default is not selected.

Logical partition isolation

Select this option to reserve unshared reconfigurable channel paths for the exclusive use of the LP. The default is not selected.

Establishing Optional Characteristics

Use the Customize/Delete Activation Profiles task available from the CPC Operational Customization tasks list to open a reset or image profile for an LP. Open the Options page to define the following LP characteristics:

- Minimum input/output (I/O) priority
- Maximum input/output (I/O) priority
- Defined capacity




Figure 3-24. Options Page, Image Profile

Parameter Descriptions

Minimum input/output (I/O) priority

Enter the minimum priority to be assigned to I/O requests from this logical partition.

Maximum input/output (I/O) priority

Enter the maximum priority to be assigned to I/O requests from this logical partition.

Defined capacity

Enter the upper bound in terms of millions of service units (MSUs) beyond which the rolling 4-hour average CPU utilization cannot proceed.

CP management cluster name

Enter the name of the Sysplex Cluster of which this logical partition will be made a member.

Storage Characteristics

Use the Customize/Delete Activation Profiles task available from the CPC Operational Customization Tasks list to open a reset or image profile for an LP. Open the Storage page to define the following LP characteristics:

- Central storage
- Expanded storage

The screenshot shows a window titled "P00LXLM5: Customize Image Profiles: LP01 : LP01 : Storage". Inside, there's a tree view on the left with "LP01" expanded, showing sub-items: "General", "Processor", "Security", "Storage" (selected), "Options", "Load", and "Crypto". The main area is divided into two sections: "Central Storage" and "Expanded Storage".

Central Storage

Amount (in megabytes)		Storage origin
Initial	4096	☒ Determined by the system
Reserved	2048	☐ Determined by the user
		Origin 0

Expanded Storage

Amount (in megabytes)		Storage origin
Initial	0	☒ Determined by the system
Reserved	0	☐ Determined by the user
		Origin 0

At the bottom, there are buttons: "Save", "Copy Notebook", "Paste Profile", "Assign Profile", "Cancel", and "Help".

Figure 3-25. Storage Page, Image Profile

Central Storage Parameter Descriptions

See "Central Storage" on page 3-8 for a discussion of the appropriate entries for these fields.

- | | |
|-----------------|--|
| Initial | Enter, in MB, the initial amount of central storage to be allocated to the LP at activation. |
| Reserved | Enter, in MB, the amount of additional central storage requested for the LP. The reserved storage space is storage that can be dynamically brought online to the LP at some point after LP activation. Entering 0 limits central storage to the initial amount for |

the duration of the LP activation. Enter a value that is compatible with the storage granularity supported by your CPC.

Storage origin

If **Determined by the user** is selected, enter, in MB, the central storage origin for the LP. When the LP is activated, it is allocated at the origin you specify here. Enter a value that is compatible with the storage granularity supported by your CPC.

Determined by the system

Select this option if you want the system to allocate where the LP storage resides.

Determined by the user

Select this option if you want to allocate where the LP storage resides.

Expanded Storage Parameter Descriptions

See “Expanded Storage” on page 3-10 for a discussion of the appropriate entries for these fields.

Initial

Enter, in MB, the initial amount of expanded storage to be allocated to the LP at activation. Enter a value that is compatible with the storage granularity supported by your CPC.

Reserved

Enter, in MB, the amount of additional expanded storage requested for the LP. The reserved storage space is storage that might be dynamically brought online to the LP at some point after LP activation. Entering 0 limits expanded storage to the initial amount for the duration of the LP activation. Enter a value that is compatible with the storage granularity supported by your CPC.

Storage origin

Enter, in MB, the expanded storage origin for the LP. When the LP is activated, it is allocated at the origin you specify here. Enter a value that is compatible with the storage granularity supported by your CPC.

Determined by the system

Select this option if you want the system to allocate where the LP storage resides.

Determined by the user

Select this option if you want to allocate where the LP storage resides.

Load Information

Use the Customize/Delete Activation Profiles task available from the CPC Operational Customization tasks list to open an image or load profile for an LP. Open the Load page to define the following LP characteristics:

- Load during activation
- Load address
- Load parameter
- Load timeout value




Figure 3-26. Load Page, Image Profile

Parameter Descriptions

Load during activation

Selecting this option allows initial program load of the operating system to occur automatically at LP activation. The default is not selected.

Load address

Enter the hex address of the I/O device containing the operating system to be loaded automatically at LP activation.

Use dynamically changed address

Select this option if you want to use the load address from a dynamically changed I/O configuration. This option and the **Load address** option are mutually exclusive.

Load parameter

Enter a 1 to 8 character, optional IPL load parameter for loading an operating system on each activation of the LP. This is useful for loading z/OS, or z/VSE.

Use dynamically changed parameter

Select this option if you want to use the load parameter from a dynamically changed I/O configuration. This option and the **Load address** option are mutually exclusive.

Time-out value

Enter a time-out value in seconds to limit the amount of time for successful completion of the operating system load.

Cryptographic Characteristics

Use the Customize/Delete Activation Profiles task available from the CPC Operational Customization tasks list to open a reset or image profile for an LP.

Note: Any changes to the Crypto Image Profile Page settings require a DEACTIVATE and ACTIVATE of the LP to have the change take effect. To verify the active settings for the Cryptographic characteristics use the View LPAR Cryptographic Controls task available from the CPC Operational Customization tasks list (For information regarding the View LPAR Cryptographic Controls page, see “Reviewing Current Logical Partition Cryptographic Controls” on page 5-6.)




Figure 3-27. Crypto Page, Image Profile

Parameter Descriptions

Control domain index

The Control Domain Index identifies the Crypto domain index(es) that can be administered from this logical partition being set up as the TCP/IP host for the TKE Workstation. The Control Domain Index selected must also be selected in the Usage Domain Index of the logical partitions to be managed by the TKE Workstation.

If you are setting up the host TCP/IP in this logical partition for communicating with the TKE Workstation, the partition will be used as a

path to this and the other domains' Master Keys. Indicate all the domains you want to access, including this partition's own domain, from this partition as control domains.

For more TKE Workstation information, refer to the *z/OS ICSF TKE Workstation User's Guide*.

Usage domain index

The Usage Domain Index identifies the cryptographic domain(s) assigned to the partition for all cryptographic coprocessors or accelerators that are configured for the partition. The usage domain denotes where a specific partition's secure crypto data resides and where applications running on that partition will be obtaining the cryptographic data.

If running z/OS, the usage domain index(es) selected should match the domain number(s) entered in the Options dataset when starting this partition's instance of ICSF. As of z/OS 1.2 the usage domain specification in the Options dataset is only required if multiple usage domain index(es) are selected.

If running z/VM in a logical partition with guests, such as Linux or z/OS, a range of usage domain indices should be selected when assigning access to the cryptographic accelerator or coprocessor. A range will allow more than one guest to have dedicated or shared access to the cryptographic queues. For further information, see the *z/VM CP Planning and Administration* and *z/VM Running Guest Operating Systems* documents.

The Usage Domain Index in combination with the Cryptographic Number selected in the Candidate List, must be unique across all partitions defined to the CPC. Although the Customize Image Profiles panel allows you to define duplicate combinations of Cryptographic number(s) and Usage Domain Index(es), such logical partitions cannot be concurrently active. This is a useful option, for example, for backup configurations.

Cryptographic Candidate List

The Cryptographic Candidate List identifies the Cryptographic numbers that are eligible to be accessed by this logical partition. Select from the scrollable list the number(s), from 0 to 15, that identify the coprocessor or accelerator to be accessed by this partition.

When the partition is activated, an error condition is not reported if a Cryptographic number selected in the Candidate list is not installed in the system. The Cryptographic number is ignored and the activation process continues.

To plan ahead for a Nondisruptive Hardware Install of any Crypto Express2 feature, the logical partition must be predefined with both the appropriate Usage Domain Index and Cryptographic number selected in its candidate list. It is possible to select all 16 index values (0-15) even before a Crypto Express2 feature is installed. When a new Crypto Express2 feature is installed and its cryptographic number(s) have been previously selected in the Candidate list of an active partition, it can be configured on to the partition from the Support Element using the **Configure On/Off** option in the Crypto Service Operations task list.

Selecting all the values will not cause a problem if you have 16 or fewer partitions in your configurations that will be using the Crypto Express2 feature. If you have more than 16 partitions that will need to use coprocessors or accelerators, you will have to carefully assign the selection

of cryptographic numbers across the partitions. Or, if you are running z/VM in one or more partitions, you may want to assign more than one Cryptographic number to each VM partition so that one or more than one type of Crypto card is available for use by VM guests.

For example, you have 19 logical partitions defined. All will require the use of Crypto Express2. Some partitions will only need Crypto for acceleration of SSL handshakes using CEX2A, and others will need secure key operations from CEX2C. VM partitions may need access to several usage domains on several cards. Crypto Express2 is a single feature (or book) with two PCI-X Adapter Cards. These PCI-X Adapter Cards can be configured as accelerators or coprocessors. Each Crypto Express2 can be configured with two accelerators, or two coprocessors, or one accelerator and one coprocessor. Each of those PCI-X adapter cards will have a unique cryptographic number assigned. These numbers are assigned in sequential order during installation. A CEX2A or a CEX2C can be shared across 16 partitions. So, assume the 19 partitions will share 2 CEX2C features and 2 CEX2A features. The Crypto selection could be as shown here for the following Crypto numbers:

Table 3-6. Example Selection of Crypto Numbers

Feature	Adapter	Crypto Configuration	Type	Crypto Number
Crypto Express2 1	1	Accelerator	A	00
Crypto Express2 1	2	Accelerator	A	01
Crypto Express2 2	1	Coprocessor	X	02
Crypto Express2 2	2	Coprocessor	X	03
Crypto Express2 3	1	Accelerator	A	04
Crypto Express2 3	2	Coprocessor	X	05
Crypto Express2 4	1	Coprocessor	X	06
Crypto Express2 4	2	Accelerator	A	07

Note: For Availability reasons, it is recommended that at least two cryptographic elements of the same type and capability be assigned to each partition that executes cryptographic operations on the adjunct processor. Because CEX2As do not contain any internal security data (master keys), all CEX2As are equivalent. CEX2Cs, on the other hand, will contain master keys and it is recommended that at least two CEX2Cs with the appropriate domains and master keys be assigned to an LPAR that requires secure key operations.

Table 3-7. LPAR & Crypto Assignments

LPAR & Crypto Use	Usage Domain Assignment	Logical Partition Assignment	Backup Required? Specify 2nd Logical Partition
ACME0 Prod SSL only ^{1,2}	UD=0	A00	A04
ACME1 Prod SSL only ¹	UD=1	A00	A04
ACME2 Prod SSL & secure	UD=2	A00 & X02	A04 & X05
ACME3 Prod SSL only ²	UD=0	A01	A07
..... SSL only	UD=3...10	A00	A01
ACMEF Prod SSL only	UD=0	A04	A07

Table 3-7. LPAR & Crypto Assignments (continued)

LPAR & Crypto Use	Usage Domain Assignment	Logical Partition Assignment	Backup Required? Specify 2nd Logical Partition
ACM17 Prod SSL & secure ³	UD=4	A00 & X05	A01 & X06
ACM18 Test SSL & secure ³	UD=5, 2 ⁴	A00 & X02	A04 & X05
ACM19 Test SSL only	UD=6	A00	A04
ACM5VM Prod VM	UD=7, 8, 9, 10	A07 & X05	
Notes: 1. LPAR ACME0 and ACME1 both use Accelerator cards A00 and A04, however, they use two different Usage Domains on these cards. 2. LPAR ACME0 and ACME3 both use Usage Domain 0, but they use them on different accelerator cards, A00/A04 and A01/A07. 3. LPAR ACM17 and ACM18 both use Crypto Coprocessor X05, but they use different Usage Domains on those cards, so there is no conflict. 4. ACM18 has two Usage Domains assigned, but only one can be used at a time. Normally, this TEST LPAR will provide SSL and Secure support for the Test environment using Usage Domain 5 on crypto accelerator cards A00 and A01, and crypto coprocessor cards X05 and X06. By defining this LPAR with access to Usage Domain 2 it can be a backup LPAR for ACME2. If and when there is a problem with LPAR ACME2, that operating system can be IPL'd in this LPAR, with the ICSF started task pointing to UD=2, and it will be able to access the master keys for ACME2, which are stored in Usage Domain 2 on X02 and X05.			

It is very important to make the correct Crypto number assignments in the Cryptographic Candidate List for each of these logical partitions to avoid assignment conflicts.

Installation of a Cryptographic Adapter requires the IBM CP Assist for Cryptographic Functions (CPACF) feature. See the *z/OS ICSF Application Programmer's Guide* and the *z/OS ICSF System Programmer's Guide* for complete information.

Cryptographic Online List

The Cryptographic Online List identifies the Cryptographic numbers that are automatically brought online during logical partition activation. The Cryptographic numbers selected in the Online List must also be selected in the Candidate List.

After partition activation, installed Cryptographic features that are in the partition Cryptographic Candidate list but not in the Cryptographic Online List are in a *configured off* state (Standby). They can be later configured *on* to the partition from the Support Element using the **Configure On/Off** option in the Crypto Service Operations task list.

When the partition is activated, an error condition is not reported if the Cryptographic number selected in the Online list is not installed in the system. The Cryptographic number is ignored and the activation process continues.

If a Cryptographic number selected in the Online list has been configured off to the partition, it is automatically configured back on during the next partition activation.

Enabling Input/Output Priority Queuing

Use the Enable I/O Priority Queuing task available from the CPC Operational Customization tasks list to either enable or disable I/O priority queuing for the entire CPC.




Figure 3-28. Enabling I/O Priority Queuing

Changing Logical Partition Input/Output Priority Queuing Values

Use the Change LPAR I/O priority queuing task available from the CPC Operational Customization tasks list to set the minimum and maximum I/O priority queuing values for logical partitions.

Logical Partition	Active	Minimum Input/Output (I/O) Priority	Maximum Input/Output (I/O) Priority
LP01	Yes	3	4
LP02	Yes	2	13
LP03	Yes	15	15
LP04	Yes	1	14
LP05	No	2	15
LP06	No	1	14
LP07	No	5	15
LP08	No	1	5
LP09	No	2	15
LP10	No	1	12
LP11	No	1	12
LP12	No	2	15
LP13	No	3	12
LP14	No	1	15
LP15	No	2	12
LP16	No	1	15

Total: 16

Save to Profiles Change Running System Save and Change Reset Cancel Help

Figure 3-29. Change LPAR I/O Priority Queuing

Note: Minimum I/O Priority and Maximum I/O Priority should be specified as a range of values that give software some ability to make choices. All logical

partitions in a given LPAR cluster should be given the same range of values so that WorkLoad Manager can optimize I/O throughput across the LPAR cluster.

If the software in the logical partition does not have an understanding of I/O Priority Queuing, the system programmer should set the Minimum and Maximum I/O priorities to the same value. The value chosen will be assigned to that logical partition as a constant priority value relative to all other logical partitions. This way even logical partitions that do not employ IRD technologies can benefit from this support.

Parameter Descriptions

Minimum input/output (I/O) priority

Enter the minimum priority to be assigned to I/O requests from this logical partition.

Maximum input/output (I/O) priority

Enter the maximum priority to be assigned to I/O requests from this logical partition.

Moving Unshared Channel Paths

You can move reconfigurable channel paths owned by one LP to another LP.

Moving Unshared Channel Paths from a z/OS System

1. Select the LP that owns the channel path to display channel path information for the LP.
2. Move the channel path from the z/OS console.
 - a. Enter **CF CHP(nn),OFFLINE**, where *nn* is the number of the desired channel path, from the z/OS operator console that has the reconfigurable channel path online.
 - b. Enter **CF CHP(nn),ONLINE**, where *nn* is the number of the desired channel path, from the z/OS operator console that is the target LP.

Notes:

- 1) If a channel path is configured offline while the LP is isolated and remains offline when you change the LP's isolation status from enabled to disabled, the channel path must be configured offline again. See "Releasing Reconfigurable Channel Paths."
- 2) If the channel path you want to move is currently assigned to a deactivated LP, you must configure it offline from the hardware console.
- 3) If the LP is not running a system control program that supports channel path configuration commands, you can move the channel path from the hardware console.

Moving a Channel Path from the Hardware Console

1. Select the LP that owns the channel path to display channel path information for the LP.
2. Select a reconfigurable channel path from the CHPIDs work area. Drag and drop the CHPID on the Reassign Channel Path task available from the CHPID Operations task list.
3. The Reassign a Channel Path panel displays the targeted channel path, the current owning LP, and a list of target LPs to which you can reassign the channel path. Select the LP to which you want to reassign the channel path.
4. Click on the Reassign push button and confirm the action to release the channel path.
5. When the "Requested operation is complete" message appears, click on the OK push button.

Releasing Reconfigurable Channel Paths

Use this procedure when the owning LP has LP isolation enabled.

1. Select the LP that owns the channel path to display channel path information for the LP.
2. Select a reconfigurable channel path from the CHPIDs work area. Drag and drop the CHPID on the Release task available from the CHPID Operations task list.

A Confirm the Action panel will display warning you that the channel paths will be released and made available for reassignment to other LPs. Confirming the action will release the channel path and a "Requested operation is complete" message will display. Click on the OK push button to complete the task.

Configuring Shared Channel Paths

Verify the status of the channel path for each LP to which you plan to configure the channel path by opening each LP's CHPIDs Work Area.

Enter **CF CHP(nn),ONLINE** (where *nn* is the number of the desired CHPID) from each z/OS operator console to which the CHPID is to be brought online.

If the operating system running in the LP is not z/OS, use the Configure On/Off task available from the CHPID Operations task list to configure the CHPID online. The shared channel path will be physically configured when the first LP configures the channel path online.

Note: Dynamically managed channel paths can be configured **Off** but **cannot** be configured **On** from the support element's CHPID Operations task list.

Deconfiguring Shared Channel Paths

Verify the status of the channel path for each LP to which you plan to configure the channel path by opening each LP's CHPIDs Work Area.

Enter **CF CHP(nn),OFFLINE** (where *nn* is the number of the desired CHPID), from each z/OS operator console from which the CHPID is to be taken offline.

If the operating system running in the LP is not z/OS, use the Configure On/Off task available from the CHPID Operations task list to configure the CHPID offline. The shared channel path will be physically deconfigured when the last LP that had the channel path online configures it offline.

Removing Shared Channel Paths for Service

1. Enter (from the z/OS console that has the shared channel path online) **CF CHP(nn) OFFLINE**, where *nn* is the number of the desired CHPID.
2. Use the Toggle all off option in the Configure On/Off task available from the CHPID Operations task list to remove the CHPID.

The Toggle all off option detaches the CHPID from all LPs that it is currently attached to regardless of the target LP. A CHPID that is shared by multiple LPs is detached from all LPs without forcing you to detach it individually from each LP.

Toggle all off is also valid for unshared CHPIDs and it is mutually exclusive of the Release task.
3. Use the Service On/Off task available from the CHPID Operations task list to remove the CHPID.

Changing Logical Partition Definitions

You can make changes to LP definitions that are available dynamically to a running LP or that are available at the next LP activation.

Changes Available Dynamically to a Running LP

The following changes are available dynamically to a running LP:

- Using the Change Logical Partition Controls task, you can change the following LP definitions:
 - Initial processing weight

- Minimum processing weight
- Maximum processing weight
- WorkLoad Manager enablement
- Processor weight capping
- Processor running time (globally applicable to logical CPs of all shared LPs)
- Global enablement of event-driven dispatching
- Defined capacity
- Using the Change Logical Partition Input/Output (I/O) Priority Queuing Controls task, you can change the following LP definitions:
 - Minimum input/output (I/O) priority
 - Maximum input/output (I/O) priority
- Using the Change Logical Partition Security task, you can change the following LP definitions:
 - Global performance data control
 - Input/Output configuration control
 - Cross partition authority
 - Logical partition isolation

Changes Available at the Next LP Activation

The following changes are available at the next LP activation:

- Use the IOCP RESOURCE statement to specify the MIF image ID numbers assigned to the LPs. The only way to change the specified MIF image ID numbers is by creating a new IOCDs.
- Use the Partitions page in the reset profile to change the:
 - LP automatic activation order
- Use the General page in the image profile to change the:
 - Logical partition identifier
 - Mode of the LP
 - Enablement of the Logical partition time offset
- Use the Processor page in the image profile to change the:
 - Number of CPs, ICFs, IFLs, IFAs, or zIIPs
 - Whether or not CPs are dedicated to an LP
 - Weight capping
 - WorkLoad Manager enablement
 - Initial processing weight
 - Initial capping enablement
 - Minimum processing weight
 - Maximum processing weight
- Use the Options page in the image profile to change the:
 - Minimum input/output (I/O) priority
 - Maximum input/output (I/O) priority
 - Defined capacity
- Use the Security page in the image profile to change the:
 - Global performance data control
 - Input/output configuration control
 - Cross partition authority
 - Logical partition isolation
- Use the Storage page in the image profile to change the:
 - Central storage definitions
 - Expanded storage definitions
- Use the Time Offset page in the image profile to change the:
 - Logical partition time offset
- Use the Load page in the image or load profile to change the:

- Automatic load data
- Use the Crypto page in the image profile to update the:
 - Control domain index
 - Usage domain index
 - Cryptographic Candidate List
 - Cryptographic Online List

Changes Available at the Next Power-On Reset (POR)

The following changes are available at the next power-on reset.

- Use the Options page in the reset profile to change the:
 - Global enablement of input/output (I/O) priority queuing
 - Processor running time interval
 - Enablement of event-driven dispatching

Chapter 4. Operating Logical Partitions

Overview	4-2
Available Operator Controls	4-2
Operator Controls Not Available	4-4
Operator Tasks	4-5
Editing Activation Profiles	4-5
Reset Profiles	4-5
Image Profiles	4-5
Load Profiles	4-5
Activating a CPC	4-5
Activating an LP	4-5
Performing a Load on an LP or Activating a Load Profile	4-5
Deactivating an LP	4-6
Locking and Unlocking an LP	4-6
Deactivating a CPC	4-7

Overview

This chapter provides information on operating the hardware system.

Sample tasks and panels explained in this chapter reference tasks and panels available from the support element console of a System z9. However, detailed procedures for operator tasks and accurate task and panel references are explained in the *Support Element Operations Guide*, SC28-6858.

Available Operator Controls

The following tasks are available when logged on in system programmer mode:

Task	Tasks List
Hardware Messages	All
Operating System Messages	All
Activate	Daily Tasks
Reset Normal	Daily Tasks
Deactivate	Daily Tasks
Grouping	Daily Tasks
Activity	Daily Tasks
Start All	CPC Recovery
Stop All	CPC Recovery
Reset Normal	CPC Recovery
Reset Clear	CPC Recovery
Load	CPC Recovery
Power-on Reset	CPC Recovery
Load from CD-ROM, DVD, or Server	CPC Recovery
Service Status	Service
View Service History	Service
Checkout Tests	Service
Report a Problem	Service
Transmit Service Data	Service
Dump LPAR Data	Service
Delete LPAR Dump Data	Service
Offload Virtual RETAIN Data To HMC DVD	Service
STI Status and Control	Service
Perform Problem Analysis	Service
Change Internal Code	Change Management
Define Clonable Internal Code Levels	Change Management
System Information	Change Management
Force Channel Internal Code Change	Change Management
Authorize Internal Code Changes	Change Management
Authorize Concurrent Internal Code Change	Change Management
Alternate Support Element	Change Management
Query Coupling Facility Reactivations	Change Management
Query Channel/Crypto Configure Off/On Pending	Change Management
Remote Service	CPC Remote Customization
Customer Information	CPC Remote Customization

Task	Tasks List
Customize/Delete Activation Profiles	CPC Operational Customization
Automatic Activation	CPC Operational Customization
Customize Scheduled Operations	CPC Operational Customization
Change LPAR Controls	CPC Operational Customization
Change LPAR Security	CPC Operational Customization
Storage Information	CPC Operational Customization
System Activity Profiles	CPC Operational Customization
Enable/Disable Dynamic Channel Subsystem	CPC Operational Customization
View LPAR Cryptographic Controls	CPC Operational Customization
Export/Import Profile Data	CPC Operational Customization
Enable I/O Priority Queuing	CPC Operational Customization
Change LPAR I/O Priority Queuing	CPC Operational Customization
Change CP/SAP Allocation	CPC Operational Customization
Change LPAR Group Controls	CPC Operational Customization
Perform Model Conversion	CPC Configuration
Transmit Vital Product Data	CPC Configuration
View Frame Layout	CPC Configuration
System (Sysplex) Time	CPC Configuration
Input/Output (I/O) Configuration	CPC Configuration
View Hardware Configuration	CPC Configuration
View CBU Feature Information	CPC Configuration
Channel PCHID Assignment	CPC Configuration
View On/Off CoD Feature Information	CPC Configuration
Cryptographic Configuration	CPC Configuration
Cryptographic Management	CPC Configuration
Display NPIV Config	CPC Configuration
Start	CP Toolbox
Stop	CP Toolbox
Display or Alter	CP Toolbox
PSW Restart	CP Toolbox
Stop on CP Address Match	CP Toolbox
Interrupt	CP Toolbox
Store Status	CP Toolbox
Configure On/Off	CHPID Operations
Release	CHPID Operations
Show LED	CHPID Operations
Channel Problem Determination	CHPID Operations
FCP NPIV On/Off	CHPID Operations
Configure On/Off	Crypto Service Operations
Service On/Off	Crypto Service Operations
Advanced Facilities	Crypto Service Operations
Channel Problem Determination	Crypto Service Operations
Configure On/Off	Channel Operations
Release	Channel Operations
Service On/Off	Channel Operations
Show LED	Channel Operations
Advanced Facilities	Channel Operations
Reassign Channel Path	Channel Operations
Channel Problem Determination	Channel Operations

Operator Controls Not Available

The following tasks are not available when logged on in system programmer mode:

- Enable TOD
- Change Operation Rate
- Stop on I/O Address Match
- Stop on I/O Event
- Stop on PSW Event
- Trace

Operator Tasks

Editing Activation Profiles

You can edit reset, image, and load profiles for configurations using the Edit Activation Profiles task available from the CPC Operational Customization tasks list.

Reset Profiles

Use the reset profile to:

- Select an IOCDS
- Optionally specify an LP activation sequence
- Enable I/O Priority Queuing.

You can select an IOCDS using the General page of the reset profile. To specify an optional LP activation sequence, use the Partitions page of the reset profile.

Image Profiles

Use the image profile to:

- Define LP characteristics
- Optionally specify automatic load settings.

To define LP characteristics, use the General, Processor, Security, Storage, Options, and Crypto pages of the image profile (see “Defining Logical Partitions” on page 3-55). To specify optional automatic load settings, use the Load page of the image profile (see “Load Information” on page 3-75).

Load Profiles

If you are not using the image profile to specify load options for an LP, use the load profile.

Activating a CPC

Activate a CPC by dragging and dropping the CPC icon onto the Activate task available from the Daily Tasks task list. This will activate the hardware system and, if LP activation sequence is enabled for LPs, will activate those LPs in the order specified in the reset profile for the CPC. You can also automatically load the LP's operating system as specified in the image profile for the LP.

Activating an LP

Activate an LP by dragging and dropping the LP icon onto the Activate task available from the Daily Tasks tasks list. This will activate the LP and, if automatic load is specified for an LP, will automatically load the LP's operating system as specified in the image profile for the LP.

Performing a Load on an LP or Activating a Load Profile

Perform a load on an LP or activate a load profile for an LP by dragging and dropping the LP icon for a previously activated LP onto the Edit Activation Profiles task available from the CPC Operational Tasks tasks list.

Select or, if necessary, customize or create a load profile for the LP. Assign the load profile to the LP's activation profile and exit the Edit Activation Profiles task saving your changes. Next, drag and drop the LP icon onto the Activate task available from the Daily Tasks tasks list.

Also, in recovery situations, you can drag and drop the LP icon onto the Load task available from the CPC Recovery tasks list.

Deactivating an LP

Deactivate an LP by dragging and dropping the LP icon onto the Deactivate task available from the Daily Tasks tasks list. This will deactivate the LP and any operating system running in the LP.

Locking and Unlocking an LP

Lock an LP by double-clicking on the CPC image icon representing the LP to open its Detail page. Set the **Lockout disruptive tasks** radio button to **Yes** and click on the **Save** button.

You can use this same procedure to unlock an LP by setting the **Lockout disruptive tasks** radio button to **No** and saving the setting.

Locking an LP can prevent accidentally performing disruptive tasks on an LP. You can lock the following tasks:

- **Daily tasks**
 - Activate
 - Reset Normal
 - Deactivate
- **CPC Recovery tasks**
 - Start all
 - Stop all
 - Reset Normal
 - Reset Clear
 - Load
 - Power-on Reset
 - Load from CD-ROM, DVD, or Server
- **Service tasks**
 - Checkout Tests
 - Dump Machine Loader Data
- **Change Management tasks**
 - Change Internal Code
 - Define Clonable Internal Code Levels
 - Query Coupling Facility Reactivations
- **CPC Operational Customization tasks**
 - Change LPAR Controls
 - Change LPAR Security
 - Change LPAR Group Controls
- **CP Toolbox tasks**
 - Start
 - Stop
 - Display or Alter
 - PSW Restart
 - Interrupt
 - Store Status
- **CHPID Operations tasks**
 - Configure On/Off
 - Release
 - Channel Problem Determination
- **Crypto Service Operations tasks**
 - Configure On/Off
 - Service On/Off
 - Advanced Facilities

- Channel Problem Determination
- **Channel Operations tasks**
 - Configure On/Off
 - Release
 - Service On/Off
 - Advanced Facilities
 - Reassign Channel Path
 - Channel Problem Determination

Deactivating a CPC

Deactivate a CPC by dragging and dropping the CPC icon onto the Deactivate task available from the Daily Tasks tasks list. This will deactivate the CPC and any activated LPs and their associated operating systems.

Chapter 5. Monitoring the Activities of Logical Partitions

Overview	5-2
Monitoring Logical Partition Activity	5-2
Reviewing Current Storage Information.	5-2
Reviewing and Changing Current Channel Status	5-3
Reviewing and Changing Current Logical Partition Controls	5-3
Reviewing and Changing Current Logical Partition Group Controls	5-3
Reviewing and Changing Current Logical Partition Security	5-5
Reviewing Current Logical Partition Cryptographic Controls	5-6
Reviewing Current System Activity Profile Information	5-6
Reviewing and Changing Logical Partition I/O Priority Values	5-7
Logical Partition Performance	5-8
RMF LPAR Management Time Reporting	5-8
Dedicated and Shared Central Processors	5-9
CPENABLE.	5-9
Start Interpretive Execution (SIE) Performance	5-9
Recovery Strategy	5-10
Operation Considerations	5-10
Application Preservation	5-11
Transparent Sparing	5-11

Overview

This chapter describes the panels and tasks that can be used to monitor LP activity. It also provides LP performance information and provides guidelines and suggestions for planning a recovery strategy for operation.

Monitoring Logical Partition Activity

In addition to viewing LP information located in the General, Processor, Security, Storage, Load, and Channels pages of an LP's reset, image, and load profiles, you can use the following tasks to further monitor LP activity:

- Current storage information
- Current channel information
- Current System activity profile information

Reviewing Current Storage Information

Use the Storage Information task available from the CPC Operational Customization tasks list to open the Storage Information task to display LP storage information for LPs currently activated on the CPC.

For this example, assume the amount of installed storage is 32768 (32 GB). The amount of physical central storage used by hardware system area (HSA) is 1280 MB leaving 32488 MB of physical storage for allocation to LPs. The addressable range used for HSA is 2176 MB so the size of addressable storage for LP allocation is 1046400 MB.

Central Storage Allocation						
Name	Origin	Initial	Current	Maximum	Gap	Expanded Storage Element
LP03	1015552	256	256	256	0	
LP08	1014784	512	512	768	0	
LP08	1014528		256		0	0

Expanded Storage Allocation					
Name	Origin	Initial	Current	Maximum	Gap
LP08	1048064	256	256	512	0
LP03	1	0	0	0	1048063

Figure 5-1. Storage Information Task. This example shows two partitions active. LP02 has 4096 MB of initial central storage and 2048 MB of initial expanded storage. LP01 has 1024 MB of initial central storage, 1024 MB of reserved central storage and 1024 MB of initial expanded storage. LP01 was activated first, without any expanded storage. Then, LP02 was activated and LP01 was activated again with expanded storage.

Note: The central storage addressability required for LP01's expanded storage is not necessarily contiguous to the CS addressability for LP01's initial or reserved central storage.

Reviewing and Changing Current Channel Status

Use the tasks available on the CHPID Operations task list to display and modify channel status information in an LP's CHPID work area.

Reviewing and Changing Current Logical Partition Controls

Use the Change LPAR Controls task available from the CPC Operational Customization tasks list to display and modify LP controls for an LP.

Logical Partition	Active	Defined Capacity	WLM	Current Weight	Initial Weight	Min Weight	Max Weight	Current Capping	Initial Capping	Number of Dedicated Processors	Number of Not dedicated Processors
LP01	No	0	☒	0	30	30	40	No	☐	0	1
LP02	No	0	☐	0	10			No	☐	0	1
LP03	Yes	0	☐	130	130			Yes	☒	0	1
LP04	No	0	☐	0	10			No	☐	0	1
LP05	No	0	☒	0	25			No	☐	0	1
LP06	No	0	☐	0	80			No	☒	0	3
LP07	No	0	☐	0	10			No	☐	0	1
LP08	Yes	0	☐	10	10			No	☐	0	1
LP09	No	0	☐	0	10			No	☐	0	1

Processor running time

Attention: It is recommended that you select "Dynamically determined by the system". Selecting "Determined by the user" risks suboptimal use of processor resources.

☒ Dynamically determined by the system

☐ Determined by the user

Running time: 1 to 100 milliseconds

Figure 5-2. Change Logical Partition Controls Page. The Current Weight for an active shared LP will always show nonzero. For non-WLM managed logical partitions, the Current Weight will always be equal to the Initial Processing Weight. The Current Weight for WLM managed logical partitions will be whatever value WLM has currently assigned to the LP. Current Capping will always be indicated for an active LP that has Initial Capping. Current capping will also be indicated for an LP which has been dynamically soft-capped by the Workload Manager (WLM).

Reviewing and Changing Current Logical Partition Group Controls

Change LPAR Group Controls allows you to define LPAR group capacity limits, allowing you to specify one or more groups of LPARs on a server, each with its own capacity limit. This is designed to allow z/OS to manage the groups in such a way that the sum of the LPARs' CPU utilization within a group will not exceed the group's defined capacity. Each LPAR in a group can still optionally continue to define an individual LPAR capacity limit.

LPAR group capacity limit requires that all LPARs to be managed in the group be running z/OS or z/OS.e V1.8 or later. LPAR group capacity limits may help provision a portion of a System z9 server to a group of LPARs allowing the CPU resources to float more readily between those LPARs, resulting in more productive use of "white

space” and higher server utilization.



Change LPAR Group Controls

Edit ▼

Input/Output configuration data set (IOCDS): a1

Group Name	Member Partitions	Group Capacity Value	
DEFAULT	LP13	33	
LENORA	LP02 LP03 LP05 LP06 LP08 LP09 LP10	123	
GROUP1	LP11	33	

Save to Profiles

Change Running System

Save and Change

Reset

Cancel

Help

Figure 5-3. Change LPAR Group Controls

I

I

I

For information about how workload management and workload license charges relate to the Group Capacity setting, see *z/OS MVS Planning: Workload Management*, SA22-7602.

Reviewing and Changing Current Logical Partition Security

Use the Change LPAR Security task available from the CPC Operational Customization tasks list to display and modify LP security controls for an LP.

P00LXLM5: Change Logical Partition Security

Change Logical Partition Security

Input/output configuration data set (IOCDs): a1 PROD123

Logical Partition	Active	Performance Data Control	Input/Output Configuration Control	Cross Partition Authority	Partition Isolation
LP01	Yes	☒	☒	☐	☐
LP02	Yes	☒	☒	☐	☐
LP03	Yes	☒	☒	☐	☐
LP04	Yes	☒	☒	☐	☐
LP05	No	☒	☒	☐	☐
LP06	No	☒	☒	☐	☐
LP07	No	☒	☒	☐	☐
LP08	No	☒	☒	☐	☐
LP09	No	☒	☒	☐	☐
LP10	No	☒	☒	☐	☐
LP11	No	☒	☒	☐	☐
LP12	No	☒	☒	☐	☐
LP13	No	☒	☒	☐	☐
LP14	No	☒	☒	☐	☐
LP15	No	☒	☒	☐	☐
LP16	No	☒	☒	☐	☐

Total: 16

Save to Profiles Change Running System Save and Change Reset Cancel Help

Figure 5-4. Change Logical Partition Security Page

Reviewing Current Logical Partition Cryptographic Controls

Use the View LPAR Cryptographic Controls task available from the CPC Operational Customization tasks list to display the current cryptographic configuration for the logical partition.




Figure 5-5. View LPAR Cryptographic Controls Page

Reviewing Current System Activity Profile Information

Use the System activity profiles task available from the Daily Tasks tasks list to open a System activity information window. Locate the System Activity profile which displays the system activity information you want to view.

Reviewing and Changing Logical Partition I/O Priority Values

Use the the Change LPAR I/O Priority Queuing task available from the CPC Operational Customization tasks list to display and modify I/O priority queuing values for one or more LPs.

P00LXLM5: Change Logical Partition Input/Output (I/O) Priority Qu

 **Change Logical Partition Input/Output (I/O) Priority Queuing** 

Input/output configuration data set (IOCDS): a1
Global input/output (I/O) priority queuing: Enabled
Maximum global input/output (I/O) priority queuing value: 15

Logical Partition	Active	Minimum Input/Output (I/O) Priority	Maximum Input/Output (I/O) Priority
LP01	Yes	3	4
LP02	Yes	2	13
LP03	Yes	15	15
LP04	Yes	1	14
LP05	No	2	15
LP06	No	1	14
LP07	No	5	15
LP08	No	1	5
LP09	No	2	15
LP10	No	1	12
LP11	No	1	12
LP12	No	2	15
LP13	No	3	12
LP14	No	1	15
LP15	No	2	12
LP16	No	1	15

Total: 16

Save to Profiles

Change Running System

Save and Change

Reset

Cancel

Help

Figure 5-6. Change LPAR I/O Priority Queuing Page

Logical Partition Performance

Contact your IBM service representative for assistance in planning configuration capacity requirements.

The performance of LPs depends on:

- The operating systems and their workloads
- The configuration of the LPs
- The configuration of the CPC
- The performance tuning parameters

You can contact an service representative who has access to a proprietary IBM performance planning tool (LPARCE from CPSTOOLS at WSCVM) to assist you in this task.

RMF LPAR Management Time Reporting

RMF processor utilization reporting includes LPAR Management Time on the Partition Data Report. RMF provides information about all shared logical CPs that remain active for the duration of the reporting interval. RMF can also be run on a single LP if additional information is required.

There are two types of LPAR management time reporting: time spent managing the LP, and time spent managing the physical configuration. With LPAR Management Time reporting, the time used to manage an LP can be separated from the time used by the workload in an LP.




Figure 5-7. ETR Increasing with CPU Utilization

Figure 5-7 on page 5-8 is an example of how information from the enhanced RMF LPAR Management Time reports can be used. This example shows the LPAR Management Time and the partition effective dispatch time for each LP. As CPU utilization increases, LPAR management time decreases, and the external throughput rate (ETR) increases.

Using the RMF report, CPU-utilization calculations can be based on an LP's effective dispatch time. The effective dispatch time can be determined by excluding LPAR Management Time from the time that a physical CP is assigned to a logical CP.

Note: For more information about RMF support, the RMF Partition Data Report, and LPAR Management Time see the *Resource Measurement Facility User's Guide*.

Dedicated and Shared Central Processors

Generally, LPs with dedicated CPs require fewer processing resources. The internal throughput rate (ITR) of LPs with dedicated CPs is higher than that of identically defined LPs with shared CPs.

Generally, the capability to recognize CP resources that are not used and to balance workloads can result in an improved ETR for LPs with shared CPs. In particular, for workloads that exhibit fluctuations in processing demands, the consolidation of several systems as LPs with shared CPs on a CPC can improve responsiveness and increase ETR.

CPENABLE

The z/OS CPENABLE parameter can be used to provide selective enablement for I/O interrupts. On System z9 CPCs, the best ITR is achieved when the fewest CPs are enabled for I/O interrupts. Selective enablement for I/O and the CPENABLE parameter are described in the *z/OS MVS Initialization and Tuning Reference*, SA22-7592.

For logical partitions using dedicated processors, IBM recommends a CPENABLE=(10,30) setting to minimize the number of CPs handling I/O interrupts in an LP.

A CPENABLE=(0,0) setting, enabling all CPs to handle I/O interrupts, can cause performance degradation for logical partitions with dedicated CPs.

For logical partitions sharing CP resources, IBM recommends a CPENABLE=(10,30) setting. There is a negligible ITR impact between a CPENABLE=(0,0) setting versus a CPENABLE=(10,30) setting for logical partitions sharing CP resources.

Start Interpretive Execution (SIE) Performance

PR/SM implementation uses CPC hardware mechanisms that are also used to implement the SIE facility when operating in ESA/390 mode. Therefore, if SIE is executed in an LP, SIE performance is reduced relative to operation in ESA/390 mode. The performance reduction depends on the frequency of SIE invocation and should be evaluated for each application.

Recovery Strategy

Recovery planning requires that the appropriate planners and technical support personnel understand the recovery strategy.

In planning for recovery, consider the following guidelines:

- Recovery is considered successful if an LP is able to perform useful work and critical application programs remain operational (even if one or more LPs are disabled) after a failure occurs. Recoverability depends on the extent and location of a failure and the ability of the operating system to handle the failure.
- Develop a recovery strategy for the specific installation that addresses the specific hardware and applications of that installation. For example, if a failure occurs that might disrupt multiple LPs or a critical LP, the operator should be able to determine what must remain active and what can be deactivated, and to perform the appropriate recovery procedures.
- The operator should follow established local procedures for reporting problems and for performing recovery procedures. It is recommended that recovery procedures be ranked from least to most disruptive. The operator should know what procedures to follow if any or all LPs do not respond to the recovery actions directed to it.
- Assign channel paths to LPs as described in the guidelines under “Guidelines for Setting Up the I/O Configuration” on page 2-5.
- Define resources to LPs so that any hardware failure has a minimal impact on any LP that must remain active.

For example, the failure of a physical CP causes the temporary loss of the logical CP that was dispatched on the physical CP. The LP owning that logical CP may continue running if it was running on an LP with at least two CPs and if the operating system can recover from a CP failure.

Operation Considerations

If an individual LP is affected, the following recovery actions (ranked from least to most disruptive) should be considered when planning recovery procedures.

- If an affected LP continues to operate with a problem, allow it to do so.
- If the operating system in an LP remains active but processing is interrupted, consider the applicability of a restart or IPL.
- Perform a stand-alone dump and IPL the operating system.
- Perform an orderly shutdown of the operating system in an affected LP. If the LP is not critical, allow the other LPs to continue.

If all LPs are affected, or if a critical LP is affected and did not respond to the recovery actions directed only to it, the following recovery actions (also ranked from least to most disruptive) should be considered when planning recovery procedures.

- Perform an orderly shutdown of all the operating systems. Activate the LPs and IPL the operating systems.
- Perform an orderly shutdown of all the operating systems. Perform the most appropriate recovery action (for example, in response to a hardware failure). Perform a power-on reset.
- If a power-on reset fails to initialize the LPs, perform a power-on reset to attempt a recovery. IPL the most critical operating system.

Application Preservation

The application preservation facility enhances system availability and provides additional protection from CP failures. This support is available on System z9 CPC models with two or more central processors (CPs). Using application preservation, the system moves an application in process from a failing CP to another operating CP. Both the CP failure and the move are transparent to the application.

There are no software corequisites when running with shared CPs. For LPs using shared CPs, even 1-way LPs can survive a CP failure without experiencing a failure to the application or the LP, providing the 1-way LP is being run on a model with more than one physical CP.

Transparent Sparing

Transparent sparing takes LP recovery one step further by combining the benefits of application preservation and concurrent CP sparing to allow for the transparent recovery of an LP and its applications (including CF LPs). Transparent sparing uses hardware to handle the recovery, requires no software support or operator intervention, and is effective for both shared and dedicated LP environments.

Transparent sparing configures a spare PU (processor unit) to replace a failed CP, ICF, IFL, IFA, zIIP or SAP. Each z9 EC has 2 spare PUs per processor book. The S08 has 2 spare PUs, the S18 has 4, the S28 has 6, the S38 has 8, and the S54 has 8.

Appendix A. Coupling Facility Control Code Support

Coupling facility control code is Licensed Internal Code (LIC) that supports the following coupling facility limits:

Table A-1. Coupling Facility Limits at Different Coupling Facility Code Levels

Coupling Facility Limit	Coupling Facility Code Level								
	Level 15	Level 14	Level 13	Level 12	Level 11	Level 10	Level 9 (Service Level 2.00 and Higher)	Level 9 (Service Level less than 2.00)	Level 8-7
Maximum number of CPs	16	16	16	16	16	16	16	12	12
Storage increment	512 KB	256 KB	256 KB	256 KB	256 KB	256 KB	256 KB	256 KB	256 KB
Structure ID limit	1023	1023	1023	1023	1023	1023	1023	1023	1023
Retry buffer limit	2039	1799	1799	1559	1559	1559	1559	519	519
Facility information	64 bytes	64 bytes	64 bytes	64 bytes	64 bytes	64 bytes	64 bytes	64 bytes	64 bytes
Maximum list element characteristic	4	4	4	4	4	4	4	4	4
Maximum lock table entry characteristic	5	5	5	5	5	5	5	5	5
User identifier limit	32	32	32	32	32	32	32	32	32
Maximum data area element characteristic	4	4	4	4	4	4	4	4	4
Local cache identifier limit	255	255	255	255	255	255	255	255	255
Storage class limit	63	63	63	63	63	63	63	63	63
Castout class limit	1024	1024	1024	1024	1024	1024	1024	1024	1024

Legend

Maximum number of CPs

Indicates the maximum number of CPs that can be used by a coupling facility logical partition.

Storage increment

Indicates that storage allocation requests are rounded up in units of 256 KB (512 KB for CF Level 15).

Structure ID limit

Cache and list structure ID range 1–1023.

Retry buffer limit

Retry buffer range upper limit.

Facility information

This area contains coupling facility control code release and service level information

Maximum list element characteristic

The maximum list element characteristic is 4. The size of a list structure list element in bytes equals $256 * (2^{**} \text{list element characteristic})$, for example, $256 * (2^{**4}) = 4K$.

Maximum lock table entry characteristic

The maximum lock table entry characteristic is 5. The maximum size of a lock table entry is 32. The size of a lock table entry in bytes equals $2^{**} \text{lock table entry characteristic}$.

User identifier limit

The maximum number of users and list notification vectors that can be attached to a list structure.

Maximum data area element characteristic

The maximum data area element characteristic is 4. The size of a data area element in bytes equals $256 * (2^{**} \text{data area element characteristic})$, for example, $256 * (2^{**4}) = 4K$.

Local cache identifier limit

The maximum number of local caches that can be attached to a cache structure.

Storage class limit

Storage classes are in the range 1–63.

Castout class limit

Castout classes are in the range 1–1024.

Appendix B. Developing, Building, and Delivering a Certified System

This appendix is intended to provide guidance in setting up, operating, and managing a secure consolidated environment using System z9 PR/SM. It is primarily for the security administrator, but can also be useful to other involved operations technical support personnel.

Creating Common Criteria-Based Evaluations

In October 1998, after two years of intense negotiations, government organizations from the United States, Canada, France, Germany, and the United Kingdom signed a historic mutual recognition arrangement for Common Criteria-based evaluations. This arrangement, officially known as the Arrangement of the Recognition of Common Criteria Certificates in the field of IT Security, was a significant step forward for government and industry in the area of IT product and protection profile security evaluations. The partners in the arrangement share the following objectives in the area of Common Criteria-based evaluation of IT products and protection profiles:

- To help ensure that evaluations of IT products and protection profiles are performed to high and consistent standards and are seen to contribute significantly to confidence in the security of those products and profiles
- To increase the availability of evaluated, security-enhanced IT products and protection profiles for national use
- To eliminate duplicate evaluations of IT products and protection profiles, and
- To continuously improve the efficiency and cost-effectiveness of security evaluations and the certification/validation process for IT products and protection profiles.

The purpose of this arrangement is to advance those objectives by bringing about a situation in which IT products and protection profiles which earn a Common Criteria certificate can be procured or used without the need for them to be evaluated and certified/validated again. It seeks to provide grounds for confidence in the reliability of the judgement on which the original certificate was based by declaring that the Certification/Validation Body associated with a Participant to the Arrangement shall meet high and consistent standards. The Arrangement specifies the conditions by which each Participant will accept or recognize the results of IT security evaluations and the associated certifications/validations conducted by other Participants and to provide for other related cooperative activities.

The IBM System z9 PR/SM functionality and assurances have been evaluated and certified at an EAL5 level of assurance.

This appendix must be used in conjunction with other pertinent System z9 manuals supplied with the System z9 to give a security administrator all the required information to configure and operate an LPAR-mode system in a secure manner. This appendix provides instruction on the correct use of the system so that a secure consolidation environment is created and maintained. It defines and explains the parameters, settings, and commands recommended by IBM, including references to those sections in the manuals being discussed in “Trusted Facility Library” on page B-14.

Functional Characteristics

Logical partitions are defined, and the I/O resources of the overall physical computing system are pre-allocated by the security administrator. I/O allocation is an integral part of the process of defining a total system configuration, and must be completely performed before that system configuration can be initialized. This pre-allocation is done by executing the Input/Output Configuration Program (IOCP) or Hardware Configuration Definition (HCD) to create a hardware-specific data set, called an Input/Output Configuration Data Set (IOCDS), of the I/O resources and their allocation to specific logical partitions. LPAR allocates an entire resource, such as an I/O channel path or a contiguous region of storage. At no time is any real resource allocated to more than one logical partition. Each complete I/O resource allocation is called a *configuration*. During the period between processor initializations, several IOCDS configurations can be stored, but only one is in effect at any time. The configuration becomes effective as part of the power-on reset sequence. In order to change the active configuration it is necessary to perform an activation of the hardware.

The preceding paragraph deliberately omits any discussion of Dynamic I/O Configuration, Dynamic CHPID management, Reconfigurable channel paths (CHPIDs), I/O resource sharing using Multiple Image Facility (MIF) or Intelligent Resource Director (IRD), because each of them has characteristics that, if inappropriately used, can compromise the secure consolidation capability of PR/SM. Cautions and requirements relating to their use are included throughout this appendix.

The remainder of the logical partition's resources are defined by the security administrator prior to the activation of the logical partition. These resources include storage size, number of logical processors, scheduling parameters, and security controls, which can be specified by the security administrator using the appropriate interfaces on the HMC/SE. Many of the control and security parameters can be changed at any time and will take effect dynamically with few exceptions (e.g., specifying dedicated processors for a partition will only take effect if the partition is not yet activated.) Logical partition definitions take effect at logical partition activation, and generally are static while the partition they pertain to is active.

When a resource is allocated to a logical partition, it is set to its architecturally-defined reset state. Channel paths are reset, main and expanded storage is zeroed.

Trusted Configuration

This section describes the actions the Security Administrator must take to help ensure that the computer system is configured for a secure mode of operation. The contents of this section specify the configuration of the evaluated product. Any deviation from the specified configuration will not be consistent with that of the evaluated product and may result in partitions that do not provide strict separation.

Subsequent sections in this document detail the security related characteristics of the evaluated product as well as security configurations that were not included in the evaluation. These details are provided to explain and highlight the differences between the various security settings. Nevertheless, to insure strict separation of partitions, only the configuration specified in this section should be used.

Note: All configuration requirements listed in subsequent sections are mandatory regardless of whether the term *must* or *should* is used.

- The hardware and any networks used to connect the hardware must be physically secure. Access to I/O devices must be restricted to authorized personnel. The hardware system console must be physically protected from access other than by authorized system administrators.
- The Hardware Management Console User Interface Style selection must specify "Classic Style".
- The remote support facility must be disabled.
- Devices should be configured so that no device is accessible by more than one partition (although they may be accessible by more than one channel path).
- Each I/O (physical) control unit should be allocated to a single partition in the current configuration.
- The Security Administrator should not reconfigure a channel path unless all attached devices and control units are attached to that path only.
- The Security Administrator should help ensure that all devices and control units on a reconfigurable path are reset before the path is allocated to another partition.
- No channel paths should be shared between partitions.
- The amount of reserved storage for a partition should be zero to eliminate the possibility of a covert channel.
- Although the system will help ensure that the total number of dedicated and shared processors are not over allocated, the System Administrator should make sure that the number of processors dedicated to activated partitions is less than the total number available. This is important so that some processors are available for partitions that do not have dedicated access.
- Dynamic I/O Configuration changes should be disabled (i.e. Changes require a power-on reset).
- I/O Priority Queuing should be disabled.
- Workload Manager should be disabled so that CPU and I/O resources are not managed across partitions.
- No partition should be configured to enable hipersockets (Internal Queued Direct I/O).
- Partitions should be prevented from receiving performance data from resources that are not allocated to them (no partition should have Global Performance Data Control Authority).
- At most one partition should have I/O Configuration Control Authority (i.e. No more than one partition should be able to update any IOCDS).
- The Security Administrator should help ensure that write access is disabled for each IOCDS, unless that IOCDS is to be updated (the current IOCDS should not be updated).
- The Security Administrator should verify any changed IOCDS after a power-on reset with that IOCDS, before any partitions have been activated (the Security Administrator may determine whether the IOCDS has been changed by inspecting the date of the IOCDS).
- No partition should have Cross-partition Control Authority (i.e. No partition should be able to reset or deactivate another partition).
- No partition should have coupling facility channels which would allow communication to a Coupling Facility partition.
- Replication of HMC Customizable Data must be disabled.

System z9 PR/SM Characteristics

- There is a Hardware Management Console (HMC) and Support Element (SE) from which the system can be operated. Therefore the system administrators of the system must be cleared for the highest security classification of work being performed on the system.
- Hardware-related operations for each logical partition will be conducted from the HMC or SE. Operations are invoked by selecting the CPC image icon for the desired logical partition, and dragging/dropping it on a task icon for the desired operation.
- For enhanced integrity of execution, locking of partitions is recommended. The partition must then be unlocked before other disruptive operations can be performed on that partition. Lock a logical partition (LP) by double-clicking on the CPC image icon representing the LP to open its Detail page. Set the Lockout disruptive tasks radio button to Yes and click the OK button. You can use this same procedure to unlock an LP by setting the Lockout disruptive tasks radio button to No and saving the setting. Locking an LP can prevent accidentally performing disruptive tasks on an LP.
- When entering values on an HMC or SE panel, values are not recognized by the system until you save the data and confirm the changes appearing on the screen.
- The Security Log records system operator actions and responses for operations that are security relevant. The entries are in chronological order and provide an audit log. Entries also include a user (system administrator) identifier when appropriate.
- The Security Log, when full, will be pruned to 2/3 of its capacity. The pruned records, the oldest 1/3 of the log, are discarded. Care should be taken to periodically off load the security log to insure that no records are lost.
- When the security log is successfully off loaded to removeable media, the active log is pruned so that it does not exceed 20% of it's capacity. If the active security log is below 20%, then no entries are removed. If it is above 20%, then enough active security log entries are removed (from oldest to newest) to reduce the size of the active security log to 20%. The oldest entries are still in the offloaded log.
- The security log on both the HMC and SE is 10 megabytes. Entries range from 40 bytes to 400 bytes.
- Access to the security log is achieved by selecting "View Security Logs" from the list of console actions. This is available on both the SE and HMC. Access to the Security Log from a logical partition is impossible (this capability does not exist).

Central and Expanded Storage

Throughout this document there are statements that announce "Sharing of allocated central/expanded storage among multiple logical partitions is not allowed", "... it becomes available to the logical partition if no other logical partition is using the same storage.", and "No other logical partition has this expanded storage address range on line." This is because PR/SM* has a mechanism that detects conditions where sharing was defined (where address ranges overlap), and rejects such requests. PR/SM* licensed internal code (LIC) and hardware rigidly enforce the "no-sharing" rule at logical partition definition, during logical partition activation, during logical partition reconfiguration, and during logical partition execution. PR/SM* monitors each instruction's storage accesses for validity; accesses outside the logical partition's defined storage are not permitted to proceed. Only storage increments within the logical partition's storage allocation as defined in the activation profile can be placed offline. Storage is varied off and on line by using the

z/OS or MVS/ESA CONFIG (CF) operator command. See *z/OS MVS System Commands* for further detail. While processing this command, MVS must interact with PR/SM, via a service call instruction, to request that the storage be varied. Because storage cannot be varied without PR/SM involvement, no way exists to circumvent the validity checking PR/SM does to confine a partition occupant within the storage limits defined for the logical partition.

I/O Security Considerations

IOCDs Considerations

Chapter 2, “Planning Considerations” contains a very thorough discussion of I/O configuration-related topics. It should be read in its entirety before reading the following security considerations.

When the IOCDs does not specify any sharing, I/O devices are owned solely by the logical partitions that own the channel paths that are attached to them. Even if a channel path has been designated as reconfigurable, that channel path cannot be removed from a logical partition unless the channel path has first been taken offline from within that logical partition. For z/OS System Control Program (SCP) partitions, this is done with the SCP operator command CONFIG (CF). For partitions containing other SCPs, the **Channel Operations** task list must be used. Use the **Configure On/Off** task to configure channel paths that are online. Use the **Release task** to release the channel paths that are isolated. Or, use the **Reassign Channel Path** task to reconfigure a CHPID in one step.

I/O sharing, Dynamic I/O Configuration and Dynamic CHPID Management are available in the trusted product, they should never be used in secure installation. If the IOCDs were to specify I/O sharing, it would be indicated in the Input/Output Configuration Program's Configuration Reports (see the *Input/Output Configuration Program User's Guide*.)

In a secure installation, a channel path must never be defined as shared in the IOCDs. Specification of a shared channel path can compromise the security of the installation. A shared channel path is defined by specifying one of the following on the CHPID statement:

- SHARED keyword
- NOTPART keyword
- PARTITION keyword with more than one logical partition in the access list
- IOCLUSTER keyword
- PATH keyword with more than one CSS ID (i.e. a spanned channel path)

Shared channel paths were not part of the evaluated product. Use of a shared channel path allows the possibility of two partitions having access to the same I/O control units and devices. This is in contradiction to the policy of strict separation. Additionally, the use of shared channels may facilitate some form of covert signaling. However, if covert signaling is not perceived to be a significant threat, it is highly recommended that each use of a shared channel be carefully analyzed for its possible effect on the installations security policy. Although a shared channel path is defined to be shared, **none** of the devices that are connected to it need to be shared among logical partitions. When devices are assigned to a single logical partition, they cannot be accessed by any other logical partition.

Low-speed devices (such as SCP Operator's Consoles) are especially inviting targets for sharing a single channel path using MIF. The following paragraph discusses how to share the channel path, but none of the console devices.

If you choose to share channel paths between logical partitions, and their access to specific devices attached to that channel path must be restricted, I/O Device Candidate Lists are the means for restricting access to devices. The default, if no I/O Device Candidate List is specified, is that all partitions sharing the MIF channel path, also share access to all attached devices. Such free access is incompatible with the concept of a secure consolidation platform that provides disjoint, non-communicating logical partitions, and is therefore not recommended. We recommend that when sharing is specified for a CHPID, all the associated, attached I/O devices (IODEVICE statement) must have a candidate list specified. Following a rule of always specifying a device's partition explicitly prevents unexpected results from defaults being applied. For further details on I/O device candidate list, refer to the discussion of the IODEVICE statement's PARTITION parameter in the IOCP User's Guide.

Sharing of channel paths is controlled by the **SHARED** parameter, and the partition names specified in the **PARTITION** and **NOTPART** parameters for each channel path definition (CHPID statement) in the IOCDS. If the PARTITION parameter specifies multiple partition names, it specifies that this particular CHPID is shared among the named partitions. If a NOTPART parameter is used, it implies the sharing characteristic. However, if a NOTPART parameter excludes all partition names but one, in both access and candidate lists, no sharing is permitted. Devices attached to a shared CHPID are restricted to the partitions included in the device candidate list (specified in the IODEVICE PARTITION parameter). If the IOCDS does not specify sharing, then no sharing of CHPIDs will take place.

Operational Considerations

Global, system-wide control of Dynamic I/O Configuration is provided by the **I/O Configuration Control Authority**. Use of this facility does not result in strict separation of partitions and was not included in the evaluated product. At most, only use the **Customize/Delete Activation Profiles** task from the CPC Operational Customization tasks list to open a reset or image profile to enable I/O Configuration control for a logical partition (LP). The I/O configuration control selection is located on the **Security** page for the LP.

Logical partitions may also be designated as **Isolated**. For such logical partitions, an offline, reconfigurable CHPIDs cannot be assigned to another logical partition unless the CHPID **Release** task is invoked (or the **Reassign** task) by the System Administrator from the SE or HMC. These tasks are available from **Channel Operations**. The CHPID statement's candidate list can be used to limit the "mobility" of a reconfigurable channel. The system will only accept configure on commands for CHPIDS in partitions specified in the candidate list of the target channel path.

All channel path reconfiguration procedures should be specifically described in the secure installation's procedures. Any not described, must not be permitted. While developing these procedures, consideration must be given to the security implications of defining a channel path (and its attached devices) to be reconfigurable. Specifically, which from-to pairs of logical partitions are valid? (when this has been established, candidate lists are the means for implementing this aspect of the installation's security policy). In the process of reconfiguration, could data be passed from one logical partition to another via one of the attached

devices? What other procedural controls must be followed in order that your organization's security policy is maintained? What operator actions are required to reconfigure this specific channel path in a secure manner? Lastly, careful attention to the IOCDS language rules relating to the CHPID REC parameter is necessary to achieve the result desired.

Channel path reassignments which result from executing configure CHPID actions, are remembered by the system by recording these changes on the SE hard disk and associating them with the IOCDS (the IOCDS itself is not changed). These changes to channel path assignments (I/O Configuration) take effect whenever the logical partitions are again activated. If the IOCDS is rewritten (by invoking HCD or IOCP), the channel path reassignments are erased (at the first Activation using that newly rewritten IOCDS).

When a channel path is deconfigured from a logical partition, each subchannel (an internal structure that provides the logical appearance of an I/O device, and is uniquely associated with one I/O device) for which this channel path is the only (remaining) online path, is removed from the logical partition. Before the subchannels are removed, they are drained and disabled. Subsequently the channel path is reset. If the channel path being deconfigured is the last channel path to a device, that device is also reset. Actions directed to a removed subchannel result in a condition code=3 (not operational).

At that very first use of a newly created IOCDS, activation configures all channel paths to the logical partitions as defined by the IOCDS. The subsequent movements of reconfigurable channel paths, from one logical partition to another, is remembered by the system. During subsequent activations, as each logical partition is activated, if a channel path was (previously) moved out of a logical partition, the channel path is taken off-line to that logical partition; if a channel path was moved into a logical partition, the channel path is brought on line to that logical partition. These logical configuration changes can be viewed by performing the following steps:

- Go to the **CPC Configuration** task list
- Select **Input/Output (I/O) Configuration** task
- Select the IOCDS which is marked **Active**
- Select the **View** pulldown
- In the **View** pulldown, select **Channel Path Configuration**
- Select a **PCHID**
- Select the **View** pulldown
- In the **View** pulldown, select **CHPID information**

The Security Administrator can repeat the final three steps shown above to see all defined CHPIDS and determine which partition(s) each is assigned to, whether they are Dedicated, Shared or Reconfigurable, and the type of each CHPID. In a secure installation, CHPIDs must not be shared.

Careful review of installation security guidelines must precede using the **Swap Channel Path** procedure available from the **Channel Operations** task list. All devices attached to the channel path being switched in, may be accessible to the logical partition. This caution does not apply to a truly "spare" channel path, one with no devices currently defined or attached.

Input/Output Configuration Data Set (IOCDS)

An IOCDS defines the logical partitions by name, allocates I/O resources to each of them, and specifies the security characteristics of those I/O resources. The following list describes the security-relevant parameters of each type of IOCDS source statement.

Statement Type	Discussion
ID	No security-relevant parameters.
RESOURCE	Assign logical partition names and MIF image IDs so that explicit control is asserted, and maximum checking of following IOCDS source statements is enabled.
CHPID	• Use PARTITION parameter to specify which logical partition each channel path is allocated to. • Don't use the SHARED parameter. • Don't use REC without study of security implications. • Specify whether the channel path is REConfigurable, and specify which logical partitions are to have access (using logical partition names in the candidate list). • Do not use the IOCLUSTER keyword. Use of this keyword will enable sharing of CHPIDs by partitions within the named cluster.
CNTLUNIT	Specification of the PATH parameter must be accorded care so that a secure configuration results.
IODEVICE	Specification of the CUNUMBR parameter must be accorded care so that a secure configuration results.

LPAR Input/Output Configurations

- In general, I/O devices must not be shared by logical partitions, since they can be used to pass information from one partition to another. There may be special cases, such as an output-only device which an installation may consider sharable after careful review of any related security risks, and defining related security procedures and processes.
- The PCHID Summary Report, Channel Path Identifier (CHPID) Summary Report and I/O Device Report produced by the Input/Output Configuration Program must be thoroughly examined by the Security Administrator for indications of unwanted sharing or reconfigurability of channels and devices.
- A thorough review of the actual physical connections/links of the I/O configuration must be performed to establish that the physical configuration is identical to that specified in the IOCDS source file. Specific attention should be given to devices with multiple device path capability, to help ensure that one device (or control unit) does not (accidentally) connect to more than one partition's channel paths.
- All IOCDSs should be write-protected except for the few minutes during which they are actually updated.
- The time stamps of the production-level IOCDSs should be recorded. By dragging the CPC Icon over to the I/O Configuration task under the CPC

Operations task list, a display of the available IOCDSs will be generated. Periodic audits should be made to assure that the IOCDSs have remained unchanged.

Activation

A reset profile includes information for activating a central processor complex (CPC) and its images (logical partitions).

- In the reset profile, after selecting an LPAR IOCDS (A0-A3) deemed valid for secure operation via the Input/Output (I/O) Configuration task, the operating mode selected must be Logically partitioned.
- Dynamic I/O changes can be disabled on the Dynamic Page of the Power-on Reset Notebook displayed during Power on Reset of the CPC. Ensuring the 'Allow dynamic changes to the channel subsystem input/output (I/O) definition' is not selected, disables dynamic I/O for the CPC. Globally disabling dynamic I/O configuration narrows the control of the I/O configuration control parameter to only controlling a logical partition's reading and writing of IOCDS. Dynamic I/O changes should be disabled as described above, thereby globally disabling dynamic I/O configuration.
- Enabling Workload Manager (found on the Security page of the Image Profile), will permit the sharing of I/O resources among clustered partitions using the Intelligent Resource Director (IRD) clustering technology. Therefore, Workload Manager (IRD) should not be enabled in a secure environment without a thorough review of the security implications. See Chapter 3, "Determining the Characteristics of Logical Partitions" for more information.

Control Authority

- A logical partition's initial security settings are set in the image profile used to activate it. Afterwards, the change LPAR security task can be used to view or change the settings. Changes must be saved in the profile in order to have them available for subsequent use. Security settings are saved by the system across activations for the current configuration. Therefore, if the same configuration is used, Security settings need not be reentered (but should be checked).

Important Note

The default values for the LPAR Security parameters are not appropriate for secure operation, and must not be specified.

- Partition control authority must NOT be given to any of the logical partitions in a secure mode of operation. Abuse of control authority by a program executing in a logical partition can disrupt the processing in other logical partitions.
- The following LPAR Security parameter settings are required for a secure mode of operation:
 - **ISOLATION should be enabled.** This option binds the partition's allocated I/O configuration to it, even when a Channel Path (CHPID) is in an offline state. An overt, auditable operator action is required to unbind an item of the I/O configuration and move it to another partition.
 - **I/O CONFIGURATION CONTROL should be disabled for every partition.** By negating this option, the partitions are prevented from accessing (read or write) the existing IOCDS data sets, or dynamically altering the current I/O configuration. IOCDSs can be a means to surreptitiously pass data between partitions. In addition, dynamic alteration of the current I/O configuration can result in a partition having access to data that it is not authorized to access.

Dynamic I/O Configuration is supported by the Hardware Configuration Definition (HCD) product for the z/OS or z/VM operating system.

Note: I/O Configuration control should be enabled for a single, specific logical partition only during the short period of time when it is permitted to write a new IOCDS. Only the IOCDS to be written should have its write-protection temporarily reset. All other IOCDSs should remain write-protected during an IOCDS update operation. The Security Administrator should remain logged onto the console until the IOCDS update is complete, and the IOCDS update authority is disabled.

Note: Neither Isolation nor I/O Configuration Control option has any effect on the sharing of CHPIDS or I/O Devices. Sharing is enabled by parameters of the CHPID statement used in the definition of the IOCDS.

- **GLOBAL PERFORMANCE DATA AUTHORITY should be disabled.** This recommendation is based on a desire to block any possibility of a partition extracting meaning from another partition's performance data.
- **CROSS-PARTITION CONTROL should be disabled.** Enabling cross-partition control permits one partition to disrupt processing in other partitions, resulting in the threat of denial of service to those partition's users. When cross-partition control is disabled, Automatic Reconfiguration Facility (ARF) is disabled. ARF uses the cross-partition control capability of PR/SM. ARF is not generally appropriate in a tightly managed, secure system.
- **WORKLOAD MANAGER should be disabled.** Enabling Workload Manager (WLM) enables Dynamic CHPID Management (DCM) to optimize I/O throughput across an LPAR cluster by sharing CHPIDs among partitions who have joined the cluster. Any logical partition that is WLM enabled may join the cluster and therefore share CHPIDs reserved for use by members of the specified cluster. Furthermore, partitions within a cluster may issue a special DIAGNOSE instruction to obtain information about other partitions within the same cluster even when Global Performance Data Authority is not enabled.

Reconfiguring the System

Deconfiguration

The recommended way to deconfigure objects owned by a logical partition is to first deconfigure the object from the operating system's point of view, and when necessary (z/OS interacts with PR/SM to complete the reconfiguration process, other operating systems may not), use the hardware system console's commands to request PR/SM to deconfigure the identical object. The z/OS operating system expects operations personnel to use the HMC/SE based configuration tasks to request deconfiguration of a logical partition object.

Reconfiguration

If the objects are not presently part of the logical partition's configuration, they must be made available to the partition with the use of facilities at the hardware systems console. Channel Paths (CHPIDs) may be made available to the target logical partition using the CHPID Operations tasks; reserved storage may be available, but if it isn't, it can be made available by a Security Administrator action by updating the Image profile's Storage page of the Customize/Delete Activation Profiles task. There are many operational considerations relating to reconfiguration that are covered in greater detail in the *z/OS MVS Planning: Operations* document and the *z/OS MVS Recovery and Reconfiguration Guide*.

The following elements can be reconfigured from the z/OS Operator's console using a CONFIG command. Such a reconfiguration is limited to the objects owned by the logical partition:

- Logical Processors
- Central Storage
- Expanded Storage
- Channel Paths

See *z/OS MVS System Commands* for further detail on the CONFIG command.

MVS is aware of the logical partition objects it owns, and interacts with PR/SM to reconfigure them using the service call instruction. This Execution of this instruction results in a mandatory interception (by the System z9 processor hardware) which causes every use thereof to be mediated by PR/SM. PR/SM mediates the instruction to limit the scope of such requests to the objects that the security administrator defined for the specific logical partition.

Audit Trail

All security-relevant events initiated from the HMC/SE by the System Administrator will be written to the security log. When these logs become full, they are *pruned*. This means that the oldest one-third of the entries are deleted and the log is reduced to two-thirds full. The log has the capability to store many weeks worth of security relevant events under normal system operation.

To insure the no security relevant information is lost, the security log should be offloaded periodically to removable media provided with the processor. When the security log is successfully off loaded to removeable media, the active log is pruned so that it does not exceed 20% of it's capacity. If the active security log is below 20%, then no entries are removed. If it is above 20%, then enough active security log entries are removed (from oldest to newest) to reduce the size of the active security log to 20%. The oldest entries are still in the offloaded log.

Refer to the *Support Element Operations Guide* for additional information on how to look at the contents of the security log.

Recovery Planning

You should read "Recovery Strategy" on page 5-10, and then adapt it to your configuration's requirements for security and processing priorities.

Installation-specific recovery procedures must be developed and documented in advance, always giving consideration to where the sensitive data will be after each recovery scenario has completed.

Service and Maintenance

Many secure accounts are hesitant about enabling remote service. Consideration should be given to enabling outbound RSF calls that contain the data necessary to automatically dispatch an IBM service representative. Since there is considerable customizing capability provided, RSF can probably be tailored to match your installation's security policy and practices.

This product has support for the concurrent service and maintenance of hardware. The following can be serviced concurrently while normal customer operations continue:

- Power supplies
- Channel cards

- Licensed Internal Code (LIC)
- Processor Book

When service is performed on the above-listed elements of the processor, the physical elements are logically and electrically isolated from the remaining portions of the system still in use. This is begun by first logging on the HMC with a SERVICE ID and then performing the desired maintenance or service task. Refer to the *Service Guide* for information on how to perform the desired task.

Note: Before placing a reconfigurable or shared channel path into a service configuration, record the logical partition name(s) that it's currently assigned to. This will assure that after service is complete, the channel path will be returned to the logical partition(s) to which it was previously allocated, even if different operations personnel are now in charge.

When a partial or complete configuration is surrendered for service or maintenance the following recommendations should be followed:

- The IOCDSs should remain write-protected.
- All installation configuration data should be, or has been previously, saved. The installation configuration data should be restored, and the initial activation must be fully manual. When activation completes, use the following procedure to check the active I/O configuration:
 - go to the **CPC Configuration** panel
 - select **Input/Output (I/O) Configuration**
 - select the IOCDS which is marked **Active**
 - select the **View** pulldown
 - in the **View** pulldown, select **Channel Path Configuration**
- Prior to giving the system to the service representative for maintenance, it is advisable to idle the partitions (perform an orderly shutdown of the applications and control programs occupying the partitions, followed by stopping each partition) rather than deactivating them. Doing this allows the system to perform automatic (re)activation on the subsequent activation. Automatic activation offers fewer opportunities for human error to affect the controlling parameters of the system, and hence is more secure.

After completion of a service operation, By dragging the CPC Icon over to the I/O Configuration task under the CPC Operations task list, a display of the available IOCDSs will be generated. Use this list to check the IOCDS time stamps against the values recorded the last time the IOCDSs were updated. This is to help ensure that the IOCDSs remain unchanged.

Logical Processors

A logical CP may be taken offline as the result of an MVS operator entering an MVS CONFIG command to take one (or more) CPs offline. When this is done, MVS performs the work necessary to no longer dispatch work on the CP(s), and then executes a service call instruction to request that PR/SM take the logical CP(s) offline. See *z/OS MVS System Commands* for further detail on the CONFIG command. Lastly, a logical CP may be taken off line at the next activation of the partition by reducing the number of CPs defined for a logical partition in the image profile for the logical partition.

The maximum number of logical processors for each logical partition is defined at logical partition activation, and remains fixed for the duration of the activation. Each of these logical CPs is represented by a data structure that is associated only with

its specific logical partition. There are no circumstances where a logical CP can be "transferred" to another logical partition, nor is there a capability within the system to accomplish this.

When a logical CP is taken offline, the data structure that represents it is marked as "offline", and continues to be maintained in PR/SM-accessible storage, remaining absolutely bound to its logical partition for the duration of that partition's activation. An offline logical CP presents a checkstopped status when interrogated by the other logical CPs in the partition. An offline logical CP can be restored to the online status by issuing an MVS CONFIG command. MVS uses the service call instruction to request PR/SM bring an offline logical CP back on line. If successful, MVS prepares its control structures to add the CP to its pool of available resources.

Initial Program Load

An Initial Program Load (IPL) resets a logical partition to prepare it for loading an operating system, and then loads the operating system using the specified IPL address and IPL parameter. The IPL address and IPL parameter are normally entered manually in the image activation profile or the load profile. However, a partition with I/O Configuration authority has the capability of dynamically changing an IPL address and IPL parameter. This could potentially cause an unintended operating system to be loaded into the partition. To prevent this, the 'Use dynamically changed address' and 'Use dynamically changed parameter' checkboxes must not be selected in the activation or load profile.

Cryptographic Usage and Control Domains

Trusted Key Entry (TKE) is a workstation application supported by ICSF to allow an alternative method of securely loading cryptographic keys (DES and PKA master keys and operational keys). A unique set of cryptographic keys is maintained for each domain index within the cryptographic facility. Only one partition can perform TKE Workstation functions at a time. The partition with this control is referred to as the TKE host. The other partitions that receive key updates from the TKE host are referred to as the TKE targets.

TKE host and target partitions are defined with control and usage domain indexes. The usage domain is the area a logical partition will use for actual cryptographic functions. Usage domains are unique for logical partitions within a single System z9. The TKE host partition's control domain index list must include the usage domain index for all the TKE target partitions for which it has control. This TKE control allows one logical partition to set-up the domains for usage by other logical partitions. Domains specified at activation remain fixed for the duration of the activation.

Refer to *z/OS ICSF Trusted Key Entry PCIX Workstation User's Guide*, in addition to this appendix, for further detail setting up partitions for TKE.

CP Assist for Cryptographic Functions

The CPACF is a cryptographic feature for the fast processing of symmetric clear key functions (DES, TDES, AES) and hashes (SHA-1, SHA-256). Each PU has a CPACF attached to it. To use CPACF with logical partitions, image profiles do not need to be changed. The CPACF will require enablement via the loading of configuration data (FC 3863) using the Support Element or Single-Object operations on the Hardware Management Console. If CPACF crypto work runs on any image on a machine, that workload can be moved to another image, without requiring a change to the image, or an outage to the LPAR. Access to the CPACF by applications can be made using either the Integrated Cryptographic Services Facility

(ICSF) callable services designated for execution on the CPACF or by issuing any one of five problem state instructions (KMAC, KM, KMC, KIMD, KLMD) documented in the *z/OS Principals of Operation*.

Crypto Express2

Crypto Express2 is an optional feature that combines and replaces the functionality provided by PCI Crypto Accelerator (PCICA) and the PCI X Crypto Coprocessor (PCIXCC). Crypto Express2 is immediately available for use if the LPAR has been activated with the Cryptographic number assigned to it and the image has been placed in the Cryptographic Online List.

On the z9 EC, the maximum number of Crypto Express2 features (0863) is eight (8), (The z9 BC can have a combination of Crypto Express2 and Crypto Express2-1P, but the max of both would be 8.)

There is no CHPID association to Crypto Express2 in the System z9. However, a Cryptographic number is required to be associated with each Crypto Express2 feature. Crypto Express2 is associated with a Cryptographic number during its installation. This Cryptographic number is to be specified in the Cryptographic Candidate and Online Lists. (See “Cryptographic Characteristics” on page 3-77).

Trusted Facility Library

Use the manuals listed in this section as needed for background or supplemental information.

Check the edition notices in the beginning of each manual for correspondence to the appropriate driver level. For manuals that do not specify driver level, the highest (most recent) suffix (last two digits of the form number) is required.

Table B-1. Trusted Facility Library for PR/SM

Title	Order Number
<i>System z9 Enterprise Class System Overview</i>	SA22-6833
<i>System z9 Business Class System Overview</i>	SA22-1083
<i>Hardware Management Console Operations Guide</i>	SC28-6859
<i>Support Element Operations Guide</i>	SC28-6860
<i>Service Guide for Hardware Management Consoles and Support Elements</i>	GC28-6861
<i>Service Guide for Trusted Key Entry Workstations</i>	GC28-6862
<i>System z9 Enterprise Class Service Guide</i>	GC28-6841
<i>System z9 Business Class Service Guide</i>	GC28-6853
<i>System z9 Enterprise Class Installation Manual</i>	GC28-6840
<i>System z9 Business Class Installation Manual</i>	GC28-6852
<i>System z9 Enterprise Class Installation Manual for Physical Planning (IMPP)</i>	GC28-6844
<i>System z9 Business Class Installation Manual for Physical Planning (IMPP)</i>	GC28-6855
<i>Stand-alone IOCP User's Guide</i>	SB10-7152
<i>Input/Output Configuration Program User's Guide</i>	SB10-7037
<i>z/OS Hardware Configuration Definition: User's Guide</i>	SC33-7988
<i>System z9 Business Class Safety Inspection</i>	GC28-6854

Table B-1. Trusted Facility Library for PR/SM (continued)

Title	Order Number
<i>System z9 Enterprise Class Safety Inspection</i>	GC28-6842
<i>Application Programming Interfaces</i>	SB10-7030
<i>Security Architecture: Securing the Open Client/Server Distributed Enterprise</i>	SC28-8135
<i>Enterprise Systems Architecture/390 ESCON I/O Interface</i>	SA22-7202
<i>Introducing Enterprise Systems Connection</i>	GA23-0383
<i>z/OS MVS System Commands</i>	SA22-7627
<i>z/OS MVS Planning: Operations</i>	SA22-7601
<i>MVS Planning: Security</i>	GC28-1439
<i>z/OS MVS Recovery and Reconfiguration Guide</i>	SA22-7623
<i>z/OS Cryptographic Services ICSF TKE Workstation User's Guide</i>	SA22-7524

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

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

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 and Service Marks

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

AIX/ESA
BatchPipes
CICS
DB2
Enterprise Systems Architecture/390
ESCON
eServer
ES/9000
FICON
Geographically Dispersed Parallel Sysplex
HiperSockets
IBM
ibm.com
IMS
MVS
MVS/ESA
OS/390
Parallel Sysplex
Processor Resource/Systems Manager
PR/SM
RACF
Resource Link
RMF
Sysplex Timer
System z
'System z9
S/390

VSE/ESA
zSeries
z/Architecture
z/OS
z/VM
z/VSE
z9

INFINIBAND, InfiniBand Trade Association, and the INFINIBAND design marks are trademarks and/or service marks of the INFINIBAND Trade Association.

Java and all Java-based trademarks and logos are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

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

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

Electronic Emission Notices

The following statements apply 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 Paragraf 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 product is a Class A Information Technology Equipment and conforms to the standards set by the Voluntary Control Council for Interference by Information Technology Equipment (VCCI). In a domestic environment, this product may cause radio interference in which case the user may be required to take adequate measures.

この装置は、情報処理装置等電波障害自主規制協議会（VCCI）の基準に基づくクラス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。

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

absolute address. (1) A direct address that identifies a location without reference to a base address. An absolute address may itself be a base address. (T) (2)

In ESA/390 modes, an address that exists after translation and prefixing, but before configuration occurs.

active logical partition. An operator-initiated procedure that performs a system reset to a logical partition 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 configuration. In an ESCON environment, the ESCON Director configuration determined by the status of the current set of connectivity attributes.

adapter. (1) Hardware that provides some transitional functions between two or more devices. (2) A mechanism for attaching parts, for example, parts having different diameters. (3) In an ESCON or FICON environment, link hardware used to join different optical fiber connector types.

address. (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) 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.

allocate. To assign a resource, such as a disk or a diskette file, to perform a task. Contrast with *deallocate*.

AP. Adjunct processor.

APAR. Authorized program analysis report.

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 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.

ARF. Automatic Reconfiguration Facility. A hardware/LPAR function that is part of the cross-partition authority control setting. ARF functions are used by the Sysplex Failure Manager (SFM) policy

functions within z/OS, when RESETTIME, DEACTTIME, or the RECONFIG statement is coded in the SFM policy.

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*.

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.

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*.

bit. Either of the digits 0 or 1 when used in the binary numeration system. (T) See also *byte*.

block. A string of data elements recorded or transmitted as a unit. The element may be characters, words, or physical records. (T)

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.

C

CBP. Integrated cluster bus coupling facility peer channel.

central processor. A processor that contains the sequencing and processing facilities for instruction execution, interruption action, timing functions, initial program loading, and other machine-related functions.

central processor complex (CPC). A physical collection of hardware that consists of central storage, one or more central processors, timers, and channels.

central storage. Storage that is an integral part of the processor and includes both main storage and the hardware system area.

CF. Coupling Facility.

CFCC. Coupling Facility control code.

CFP. Coupling facility peer channel.

channel. The system element that controls one channel path, whose mode of operation depends on the type of hardware to which it is attached.

channel control check. A category of I/O errors affecting channel controls and sensed by the channel to which a device is attached.

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.

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 or FICON 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 *point-to-point channel path configuration*, and *switched point-to-point channel path configuration*.

channel path identifier (CHPID). In a channel subsystem, a value assigned to each installed channel path of the system that uniquely identifies that path to the system.

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-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.

CHP. Channel path.

CHPID. Channel path identifier.

CICS. Customer Information Control System.

CMC. Communications management configuration.

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.

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.

console. A logical device used for communication between the user and the system.

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.

controller. A unit that controls input/output operations for one or more devices.

conversion. (1) In programming languages, the transformation between values that represent the same data item but belong to different data types. Information can be lost through conversion because accuracy of data representation varies among different data types. (I) (2) The process of changing from one method of data processing to another or from one data processing system to another. (3) The process of changing from one form of representation to another; for example, to change from decimal representation to binary representation.

coupling facility. A special logical partition that provides high-speed caching, list processing, and locking functions in a 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 (or logical partitions) directly attached to it.

coupling facility control code. The Licensed Internal Code (LIC) that runs in a coupling facility logical partition to provide shared storage management functions for a sysplex.

CP. Central processor.

CPC. Central processor complex.

CPU. Central processing unit.

CS. (1) Central storage. (2) Channel set. (3) Control storage.

CTC. (1) Channel-to-channel. (2) Mnemonic for an ESCON channel attached to another ESCON channel.

CUADD. Control unit logical address.

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.

D

DASD. Direct access storage device.

DCM. Dynamic CHPID Management.

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.

deactivate logical partition. An operator-initiated procedure that releases the hardware assigned to a logical partition, 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*.

deconfigure. To remove a system resource from the currently active configuration, usually through the system control program (SCP) or through the Configuration (CONFIG) frame on the system console.

default. Pertaining to an attribute, value, or option that is assumed when none is explicitly specified. (I)

degraded. Pertaining to a mode of operation in which the system operates with some resources not available.

device. A mechanical, electrical, or electronic contrivance with a specific purpose.

device address. The 8 rightmost bits of an I/O address that identify a particular I/O device and a control unit on the designated channel.

device number. In a channel subsystem, four hexadecimal digits that uniquely identify an I/O device.

direct access storage device (DASD). A device in which access time is effectively independent of the location of the data.

domain. (1) That part of a computer network in which the data processing resources are under common control. (T) (2) In SNA, a system services control point (SSCP) and the physical units (PUs), logical units (LUs), links, link stations, and all associated resources that the SSCP has the ability to control by means of activation requests and deactivation requests.

dynamic connection. In an ESCON Director, a connection between two ports, established or removed by the ESCD and that, when active, appears as one continuous link. The duration of the connection depends on the protocol defined for the frames transmitted through the ports and on the state of the ports.

dynamic reconfiguration. Pertaining to a processor reconfiguration between a single-image (SI) configuration and a physically partitioned (PP) configuration when the system control program is active.

dynamic storage reconfiguration. A PR/SM LPAR function that allows central or expanded storage to be added or removed from a logical partition without disrupting the system control program operating in the logical partition.

E

element. A major part of a component (for example, the buffer control element) or a major part of a system (for example, the system control element).

environmental error record editing and printing program (EREP). The program that makes the data contained in the system recorder file available for further analysis.

EREP. Environmental error record editing and printing program.

error. A discrepancy between a computed, observed, or measured value or condition and the true, specified, or theoretically correct value or condition. (I) (A)

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.

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.

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.

extent. (1) Continuous space on a disk or diskette that is occupied by or reserved for a particular data set, data space, or file. (2) A set of consecutively addressed tracks that a channel program can access. The limits of an extent are defined by specifying the addresses of the first and last tracks in the extent.

F

failure. An uncorrected hardware error.

Note: Failures are either recoverable or not recoverable by the software or the operator. The operator is always notified when failures occur. Usually, system recovery occurs through a hardware reconfiguration. If this is not possible, recovery requires a repair of the failed hardware.

FCC. Federal Communications Commission.

feature. A part of an IBM product that may be ordered separately by the customer. A feature is designated as either special or specify and may be designated also as diskette-only.

Federal Communications Commission (FCC). A board of commissioners appointed by the President under the Communications Act of 1934, having the power to regulate all interstate and foreign communications by wire and radio originating in the United States.

G

GB. Gibabyte.

GMT. Greenwich mean time.

guest. In interpretive execution mode, the interpreted or virtual machine as opposed to the real machine (the host).

H

hardware system area (HSA). A logical area of central storage, not addressable by application programs, used to store Licensed Internal Code and control information.

hex. Hexadecimal.

hexadecimal. (1) Pertaining to a selection, choice, or condition that has 16 possible values or states. (I) (2) Pertaining to a fixed-radix numeration system, with radix of 16. (I) (3) Pertaining to a numbering system with base of 16; valid numbers use the digits 0–9 and characters A–F, where A represents 10 and F represents 15.

HSA. Hardware system area.

I

I/O. Input/output.

ICF. Internal coupling facility.

ICP. Internal coupling facility peer channel.

ID. Identifier.

identifier (ID). (1) One or more characters used to identify or name a data element and possibly to indicate certain properties of that data element. (T) (2) A sequence of bits or characters that identifies a program, device, or system to another program, device, or system. (3) In an ESCON Director, a user-defined symbolic name of 24 characters or fewer that identifies a particular ESCD.

IFL. Integrated Facility for Linux.

IFA. Integrated Facility for Applications. Also referred to as zAAPs, or System z9 Application Assist Processors.

IML. Initial machine load.

IMS. Information Management System.

initial program 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.

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 attaches to the processor.

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.

interactive. Pertaining to a program or system that alternately accepts input and then responds. An interactive system is conversational; that is, a continuous dialog exists between the user and the system. Contrast with *batch*.

interface. (1) A shared boundary between two functional units, defined by functional characteristics, signal characteristics, or other characteristics as appropriate. The concept includes the specification of the connection of two devices having different functions. (T) (2) Hardware, software, or both, that links systems, programs, or devices.

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.

interruption. Synonym for *interrupt*.

IOCDS. I/O configuration data set.

IOCP. I/O configuration program.

IOP. Input/Output processor.

IPL. Initial program load.

IRD. Intelligent Resource Director.

ITR. Internal throughput rate.

L

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. Microcode can be Licensed Internal Code and licensed as such.

link. (1) In an ESCON or FICON environment, the physical connection and transmission medium used between an optical transmitter and an optical receiver. A link consists of two conductors, one used for sending and the other for receiving, providing a duplex communication path. (2) In an ESCON or FICON I/O interface, the physical connection and transmission medium used between a channel and a control unit, a channel and an ESCD, a control unit and an ESCD, or, at times, between two ESCDs.

local. Synonym for *channel-attached*.

log. (1) To record, for example, to log all messages on the system printer. (2) A record of events that have occurred.

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.

logical partition (LP). A subset of the processor hardware that is defined to support the operation of a system control program (SCP). See also *logically partitioned (LPAR) mode*.

logical processor. In LPAR mode, a central processor in a logical partition.

logically partitioned (LPAR) mode. A central processor mode, available on the Configuration frame when using the PR/SM feature, that allows an operator to allocate processor unit hardware resources among logical partitions. Contrast with *basic mode*.

logout. Log data that has been collected, formatted, and recorded.

LP. Logical partition.

LPAR. Logically partitioned.

M

machine check. An error condition that is caused by an equipment malfunction.

MB. Megabyte.

megabyte (MB). (1) A unit of measure for storage size. One megabyte equals 1 048 576 bytes. (2) Loosely, one million bytes.

monitor. A device that observes and records selected activities within a data processing system for analysis.

multidrop. A network configuration in which there are one or more intermediate nodes on the path between a central node and an endpoint node. (T)

multidrop topology. A network topology that allows multiple control units to share a common channel path, reducing the number of paths between channels and control units. Contrast with *switched point-to-point topology*. See also *point-to-point topology*.

multiple image facility (MIF). A facility that allows channels to be shared among PR/SM logical partitions in an ESCON environment.

MVS/ESA. Multiple Virtual Storage/Enterprise Systems Architecture.

N

network. (1) An arrangement of nodes and connecting branches. (T) (2) A configuration of data processing devices and software connected for information exchange.

nybble. Equal to 4 bits.

O

offline. (1) Pertaining to the operation of a functional unit that takes place either independently of, or in parallel with, the main operation of a computer. (T) (2) Neither controlled by, nor communicating with, a computer. Contrast with *online*.

online. (1) Pertaining to the operation of a functional unit when under the direct control of a computer. (T) (2) Pertaining to a user's ability to interact with a computer. (A) (3) Pertaining to a user's access to a computer via a terminal. (A) (4) Controlled by, or communicating with, a computer. Contrast with *offline*.

operating system. 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.

option. (1) A specification in a statement that may be used to influence the execution of the statement. (2) A hardware or software function that can be selected or enabled as part of a configuration process. (3) Hardware (such as a network adapter) that can be installed in a device to change or enhance device functions.

P

page. In a virtual storage system, a fixed-length block that has a virtual address and is transferred as a unit between real storage and auxiliary storage. (I) (A)

parallel. (1) Pertaining to a process in which all events occur within the same interval of time, each handled by a separate but similar functional unit; for example, the parallel transmission of the bits of a computer word along the lines of an internal bus. (T) (2) Pertaining to concurrent or simultaneous operation of two or more devices or to concurrent performance of two or more activities in a single device. (A) (3) Pertaining to concurrent or simultaneous occurrence of two or more related activities in multiple devices or channels. (A) (4) Pertaining to the simultaneity of two or more processes. (A) (5) Pertaining to the simultaneous processing of the individual parts of a whole, such as the bits of a character and the characters of a word, using separate facilities for the various parts. (A)

parameter. (1) A variable that is given a constant value for a specified application and that may denote the application. (I) (A) (2) An item in a menu for which the user specifies a value or for which the system provides a value when the menu is interpreted. (3) Data passed between programs or procedures.

partition. See *logical partition*.

path. In a network, any route between any two nodes. (T)

PCI. Peripheral component interconnect.

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 topology. A network topology that provides one communication path between a channel and a control unit and does not include switching facilities. Contrast with *switched point-to-point topology*. See also *multidrop topology*.

POR. Power-on reset.

PR/SM. Processor Resource/Systems Manager.

problem determination. 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 state. A condition during which the processor cannot execute input/output and other privileged instructions.

Processor Resource/Systems Manager (PR/SM)

feature. 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.

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.

program version. A separate IBM licensed program, based on an existing IBM licensed program, that usually has important new code or functions.

PSW. Program status word.

PTF. Program temporary fix.

R

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.

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.

RMF. Resource Measurement Facility.

S

SCP. System control program.

SCSW. Subchannel status word.

SEC. System engineering change.

security administrator. Any user(s) of the HMC who are defined with a user mode of System Programmer or Service Representative.

Server Time Protocol (STP). A time synchronization feature designed to enable multiple servers to maintain time synchronization with each other. This technology uses coupling links to exchange timing signals directly between participating CECs. STP can operate either in a mode with no Sysplex Timer or can operate in a mixed mode in conjunction with a Sysplex Timer.

service representative. A person who performs maintenance services for IBM hardware products or systems.

SFM. Sysplex Failure Manager.

SIE. Start Interpretive Execution (370-XA or ESA/370 instruction).

SNA. Systems Network Architecture.

start interpretive execution (SIE). An instruction that causes the CPU to enter interpretive-execution mode and to begin running the guest program.

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)

storage control. In a DASD subsystem, the component that connects the DASD to the host channels, performs channel commands, and controls the DASD devices. For example, the 3990 Model 2 and Model 3 are storage controls.

storage director. In a logical entity consisting of one or more physical storage paths in the same storage cluster.

STP. See Server Time Protocol (STP).

subchannel. In 370-XA and ESA/390 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.

subsystem. A secondary or subordinate system, or programming support, usually capable of operating independently of or asynchronously with a controlling system. (T)

switched point-to-point channel path configuration. In an ESCON or FICON I/O interface, a configuration that consists of a link between a channel and an ESCON Director and one or more links from the ESCD, each of which attaches to a control unit. This configuration depends on the capabilities of the ESCD

for establishing and removing connections between channels and control units. Contrast with *point-to-point channel path configuration*.

switched point-to-point topology. A network topology that uses switching facilities to provide multiple communication paths between channels and control units. See also *multidrop topology*. Contrast with *point-to-point topology*.

synchronous. (1) Pertaining to two or more processes that depend on the occurrence of specific events, such as common timing signals. (T). (2) Occurring with a regular or predictable time relationship. Contrast with *asynchronous*.

sysplex. A set of systems communicating and cooperating with each other through certain multisystem hardware components and software services to process customer workloads.

Sysplex Timer. An IBM table-top unit that synchronizes the time-of-day (TOD) clocks in as many as 16 processors or processor sides.

system administrator. Any user(s) with access to the Hardware Management Console (HMC).

system control program (SCP). Programming that is fundamental to the operation of the system. SCPs include MVS, VM, and VSE operating systems and any other programming that is used to operate and maintain the system. Synonymous with *operating system*.

system reset. To reinitialize the execution of a program by repeating the load operation.

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.

T

threshold. A level, point, or value above which something is true or will take place and below which it is not true or will not take place.

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.

TOD. Time of day.

topology. The physical or logical arrangement of nodes in a computer network.

TPF. Transaction processing facility.

Transaction Processing Facility. Transaction Processing Facility is a specialized high availability operating system designed to provide quick response times to very high volumes of messages from large networks of terminals and workstations.

TSO. Time sharing option.

V

variable. (1) In programming languages, a language object that may take different values, one at a time. The values of a variable are usually restricted to a certain data type. (I) (2) A quantity that can assume any of a given set of values. (A) (3) A name used to represent a data item whose value can be changed while the program is running.

virtual machine (VM). (1) A virtual data processing system that appears to be at the exclusive disposal of a particular user, but whose functions are accomplished by sharing the resources of a real data processing system. (T) (2) A functional simulation of a computer system and its associated devices, multiples of which can be controlled concurrently by one operating system.

VM. Virtual machine.

VSE. Virtual storage extended.

W

warning message. An indication that a possible error has been detected.

WLM. WorkLoad Manager.

Z

zAAP. System z9 Application Assist Processor.

zIIP. IBM System z9 Integrated Information Processor.

Index

A

- access list
 - definition 2-30, 3-48
- activation, automatic
 - changing definitions 3-85
 - definition of 3-55, 3-59
 - reserved expanded storage definition considerations 3-25
- allocation, control unit 2-7
- application, potential 1-18
 - examples 1-19
- authority, control 3-51
- automatic activation
 - changing definitions 3-85
 - definition of 3-55, 3-59
 - reserved expanded storage definition considerations 3-25

C

- candidate list
 - definition 2-30
 - I/O device 3-48, 3-49
 - isolation 3-51
- Candidate List 3-78
- capacity backup upgrade (CBU) 2-66
- capacity upgrade on demand 2-62
- capping
 - defining 3-5
 - processing weights 3-39
 - single LP 3-42
 - disabling event-driven dispatching 3-42
- CBU (capacity backup upgrade) capability 2-66
- central processors
 - considerations for Linux-Only 3-30
 - dedicated 3-35
 - suitable workload 3-36
 - limitation 3-28
 - maximum number 3-28
 - number defined 1-15
 - overview 1-3
 - processing weights 3-37
 - recommendations for intelligent resource director (IRD) 3-29
 - reserved 2-62
 - running time 3-43
 - SFM, with 3-26
 - shared 3-36
 - suitable workload 3-36
 - wait completion 3-44
 - workload requirement 3-29
- central storage
 - definition 3-8
 - initial 3-8, 3-10
 - origin 3-9, 3-11
- CFCC patch apply 2-62

- changing LP definitions
 - changes available at the next LP activation 3-85
 - changes available dynamically 3-84
 - IOCDS, through 3-85
- channel paths
 - access list 3-48
 - assigning 3-52
 - candidate list 3-48
 - isolation 3-50
- channels
 - coupling facility channels 2-54
- characteristics of LPs 1-15
- clock type assignment 3-60
- communication via ESCON CTC 3-49
- compatibility considerations 1-20
- concurrent memory upgrade 2-66
- concurrent patch 2-61
- configuration, I/O
 - central storage resources 2-5
 - determining size 2-5
 - recovery 2-6
- configuring offline
 - ESCON channel 2-12
 - shared images 2-14
 - unnneeded channels 2-13, 2-14
- connectivity, ESCON 2-11
- control authority
 - cross-partition 3-51
 - global performance 3-51
 - I/O configuration 3-51
- Control Domain Index 3-77
- control program
 - characteristics 3-5
 - device number 2-3
 - EREP 1-21
 - ESA/390 1-5
 - HCD 1-13
 - z/OS 1-20
- control unit
 - allocation 2-7
 - considerations 2-11
- controls, security 3-50
- coupling facility 2-39
- coupling facility channels
 - CBP, CFP, ICP 2-57
 - description 2-54
 - shared channel path recommendations 2-57
- coupling facility level considerations 2-50
- coupling facility nonvolatility 2-45
- coupling facility storage sizes, estimating 2-48
- CPENABLE 5-9
- CPU addresses 1-22
 - machine types 1-22
 - model numbers 1-22
- CPU ID 1-22
 - examples 1-23
 - fields 1-22

- Cryptographic
 - Online List 3-80
- Cryptographic Candidate List 3-78
- cryptographic coprocessor feature 3-56, 3-63, 5-6
- Cryptographic Online List 3-80

D

- DCM 3-50
- deactivating LPs
 - managing logical paths for LPs 2-13
- dedicated
 - channel paths 3-48
 - CPs 3-4, 3-35
 - processing weights 3-37
- dedicated or shared internal coupling facility (ICF)
 - CPs 3-63
- defined capacity, workload charging by
 - soft-capping 3-46
- defining
 - logical partitions 3-55
 - shared channel paths 3-48, 3-49
- definitions, changing LP
 - changes available dynamically 3-84, 3-85
 - IOCDs, through 3-85
- duplexing
 - system-managed coupling facility structure 2-44
- duplicate device numbers
 - different physical devices 2-33
 - examples, number conflicts 2-36
 - migration 1-20
 - number conflicts 2-35
 - resolving conflicts, using 2-37
 - deconfiguring original channel path 2-37
 - I/O configuration 2-37
 - original channel path 2-37
- dynamic CHPID management considerations 3-50
- dynamic coupling facility dispatching 2-42, 2-50
- dynamic I/O configuration 1-14
 - availability 3-52
 - hardware configuration definition (HCD) 3-52
 - managing 3-52
 - planning 3-52
 - publications xxi
 - z/VM support 2-4
- dynamic ICF expansion (dynamic internal coupling facility expansion) 2-42
- dynamic internal coupling facility expansion (dynamic ICF expansion) 2-42
- dynamic storage reconfiguration
 - operation consideration 3-25
 - storage 3-7
 - central storage, origin 3-9
 - central, initial 3-8
 - configuration 3-12
 - expanded, initial 3-10
 - expanded, origin 3-11
 - reserved central 3-9
 - reserved expanded 3-11
- dynamically managed CHPIDs 2-27
- DYNDISP 2-42, 2-50

E

- enhanced book availability 2-67
- ESCON channels
 - configuration rules 2-11
 - example, MIF 2-10
 - overview 2-21
 - recommendations 2-21
 - shared channels
 - control units, infrequently used 2-27
 - CTC configuration 2-25, 2-26
 - ESCD configuration 2-24
 - ESCON configuration 2-23
 - requirements 2-11
 - unshared channels
 - channel utilization 2-27
 - control units, limitations 2-27
 - recommendations 2-27
- ESCON multiple image facility
 - communication via ESCON CTC 3-49
 - concurrent data transfer 2-23
 - defining devices, subset of LPs 2-16
 - IOCP deck example 2-16, 2-18
 - maximum channels 2-23
 - overview 1-3
- expanded storage
 - definition 1-2, 3-10
 - initial 3-10
- extended recovery facility 1-18

F

- FICON Express 2-7
- FICON Express2 2-7

G

- guest coupling simulation 1-5
- guidelines 2-5
 - recovery 2-6

H

- hardware configuration definition (HCD)
 - dynamic I/O configuration 1-13
 - limiting I/O devices 2-16
 - MVS/ESA 2-3
 - z/OS 1-13
- hardware support 1-14
- hardware system area (HSA)
 - allocation 1-24
 - HSA estimation tool 1-24
- HCD
 - dynamic I/O configuration 1-13
 - limiting I/O devices 2-16
 - MVS/ESA 2-3
 - z/OS 1-13
- high-performance coupling facility channels 2-56
- HiPerLinks 2-56
- HSA
 - allocation 1-24

HSA (*continued*)
HSA estimation tool 1-24

I

I/O configuration
central storage resources 2-5
control authority 3-51
determining size 2-5
director, using 2-19
block ports 2-19
prohibit dynamic connections 2-19
recovery 2-6
I/O configuration data set (IOCDS)
assigning channel paths 3-53
requirements 3-6
I/O topology
point-to-point 2-21
IBM System z9 Integrated Information Processor (zIIP) 2-61
ICB (integrated cluster bus channels) 2-56
ICF (internal coupling facility feature) 3-63
ICF (internal coupling facility) 2-41, 2-46, 3-31
ICFs, maximum number supported per CPC
model 2-41
identifier
partition 3-6
IFA 2-60
initial expanded storage 3-10
Input/Output Configuration Program
limiting I/O devices
IOCP deck example 2-16, 2-18
Input/Output Configuration Program (IOCP) 2-28
characteristics 2-4
limiting I/O devices 2-16
overview 1-14
integrated cluster bus channels (ICB) 2-56
Integrated Facility for Applications 2-60
intelligent resource director 3-29
internal coupling facility (ICF) 2-41, 2-46, 3-31
IOCDS
assigning channel paths 3-53
requirements 3-6
IOCP 2-28
characteristics 2-4
limiting I/O devices 2-16
overview 1-14
IRD 3-29
isolated, LP 3-50

K

keyword, IOCP
RESOURCE 2-28
SHARED 2-28, 2-30

L

Linux planning considerations 2-59
logical partitioning
logical vary 1-2

logical partitions
automatic IPL 3-55
changing definitions 3-84
characteristics 1-15
defining 3-55
identifier 3-6
isolation 3-50
management time reporting 5-8
maximum number 2-5
modes, supported 3-7
monitoring activity 5-2
overview 1-2
performance 5-8
dedicated CPs 5-9
shared CPs 5-9
start interpretive execution 5-9
storage configuration 3-7
logical paths
configuration rules 2-11
configuring offline 2-14
consideration 2-11
control unit 2-11
ESCON and FICON connectivity 2-11
ESCON channel configuration 2-11
ESCON converter 2-12
control unit allocation 2-7
definition 2-7
establishing 2-11
managing 2-7, 2-15
example 2-9
recommendations 2-12
parallel channels 2-7

M

maintaining relative percentage 3-42
managing
logical paths 2-7
recommendations 2-12
map planning 3-25
maximum number, CPs 3-28
migration considerations 1-18, 1-20
applications 1-18
control program 1-18
modes, supported 3-7
monitoring activity 5-2
multiple image facility
I/O management 2-22
performance considerations 2-22
performance enhancements 2-22
planning consideration 2-22
requirements 2-21
topology 2-22

N

Nondisruptive Hardware Install 3-78

O

- operation prerequisites 1-2
- operator training 1-14
- origin
 - central storage 3-9
 - expanded storage 3-11
- overview
 - capabilities 1-2
 - logical partitioning 1-2
 - potential application 1-18

P

- parallel channel paths
 - MIF 3-48
- parallel channels 2-7
- parallel sysplex support 1-4
- partition
 - identifier 3-6
- PCI cryptographic processor
 - planning for hot plugging 2-71
- performance
 - capped 3-5
 - controlling 3-4
 - ITR 3-4
 - LPAR mode 5-8
 - dedicated CPs 5-9
 - shared CPs 5-9
 - start interpretive execution 5-9
- planning
 - dynamic I/O configuration 3-52
 - storage map 3-25
- pool, single storage 3-7
- potential application 1-18
 - examples 1-19
- prerequisites for operation
 - hardware support 1-14
 - operator training 1-14
- processing weights
 - capping 3-39
 - dedicated CPs 3-37
 - effects 3-38
 - enforcement 3-39
 - use 3-37
- programs, control
 - characteristics 3-5
 - EREP 1-21
 - ESA/390 1-5
 - HCD 1-13
 - support 1-5
 - z/OS 1-20

R

- reconfigurable
 - channel path 3-48
- reconfigurable storage unit 1-21
- recovery
 - considerations 3-5
 - planning 2-6

- recovery (*continued*)
 - strategy 5-10
- related publications xx
- relative percentage, maintaining 3-42
- report management facility 1-21, 2-28, 3-51, 5-8
- reserved
 - central storage 3-9
 - expanded storage 3-11
- reserved CPs 2-62
- RESOURCE keyword 2-28
- RMF
 - See report management facility
- RSU 1-21
- running time, CP 3-43

S

- security controls 3-50
- Server Time Protocol 1-6, 1-7
- SFM (Sysplex Failure Manager)
 - description 3-26
- shared
 - CPs 3-36
 - wait completion 3-44
- shared channel paths
 - capabilities 3-4
 - configuring 3-84
 - deconfiguring 3-84
 - defining 3-48, 3-49
 - modes, possible 3-48
 - overview 1-3
 - removing for service 3-84
- shared devices
 - ESCON channels, shared 2-31
 - parallel channels, using 2-31
 - unshared channels 2-31
- SHARED keyword 2-28, 2-30
- single storage pool 3-7
- statements
 - reserved words 2-29
 - RESOURCE keyword 2-28
 - SHARED keyword 2-30
- storage
 - configurations 3-7
 - map planning 3-25
 - requirements 3-7
 - resources 3-7
- STP 1-6, 1-7
- supported modes 3-7
- Sysplex Failure Manager (SFM)
 - description 3-26
- sysplex timer 1-25
- System z9 Application Assist Processor 2-60
- system-managed coupling facility structure
 - duplexing 2-44

T

- TOD clock 1-25
- trademarks C-2

U

upgrade, capacity on demand 2-62
Usage Domain Index 3-78

W

wait completion 3-44
weights, processing 3-37
 enforcement 3-39
workload
 balancing 2-13
 requirements, CP 3-29
workload charging by soft-capping to a defined
 capacity 3-46
workload manager LPAR CPU management of shared
 CPs 3-45

X

XRF
 See extended recovery facility

Z

z/VM guest coupling simulation 1-5
zAAP 2-60
zIIP 2-61
zIIPs, shared or dedicated 3-4
zSeries parallel sysplex support 1-4



Printed in USA

SB10-7041-03

