GA22-7148-06

S/390® 9672 Parallel Transaction Server
S/390® 9672 Parallel Enterprise Server
S/390® 9674 Coupling Facility

System Overview
For R1/R2/R3 Based models

Note

Before using this information and the product it supports, be sure to read the general information under “Notices” on page xv.

Seventh Edition (April 1996)

This edition applies to the 9672 Parallel Transaction Server, 9672 Parallel Enterprise Server, and 9674 Coupling Facility R1/R2/R3 based machines and obsoletes and replaces GA22-7148-05.

Changes or additions to the text and illustrations are indicated by a vertical line to the left of the change. Make sure you are using the correct edition for the level of the product.

Order publications through your IBM representative or the IBM branch office serving your locality. Publications are not stocked at the address given below.

A form for readers' comments appears at the back of this publication. If the form has been removed, address your comments to:

IBM Corporation
Information Development
Department C9MG
1701 North Street
Endicott, NY 13760
U.S.A.

You may also send comments to PUBFDBK@VNET.IBM.COM or PUBFDBK@GDLVME.

Contents

Figures

Tables

Notices

References in this publication to IBM products, programs, or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM product, program, or service is not intended to state or imply that only IBM products, programs, or services may be used. Any functionally equivalent program that does not infringe any IBM intellectual property rights may be used instead of the IBM product, program, or service. Evaluation and verification of operation in conjunction with other products, except those expressly designated by IBM, is the user's responsibility.

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
International Business Machines Corporation
500 Columbus Avenue
Thornwood, NY 10594, USA

Electronic Emission Notices

Federal Communications Commission (FCC) Statement

The following statement applies to this IBM product. The statement for other IBM products intended for use with this product will appear in their accompanying manuals.

Federal Communications Commission (FCC) Statement

Note: This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instruction 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 or by 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 equipment does not exceed Class A limits per radio noise emissions for digital apparatus, set out in the Radio Interference Regulation of the Canadian Department of Communications. Operation in a residential area may cause unacceptable interference to radio and TV reception requiring the owner or operator to take whatever steps are necessary to correct the interference.

Avis de conformlté aux normes du ministère des Communications du Canada

Cet équipement ne dépasse pas les limites de Classe A d'émission de bruits radioélectriques pour les appareils numériques, telles que prescrites par le Règlement sur le brouillage radioélectrique établi par le ministère des Communications du Canada. L'exploitation faite en millieu résidentiel peut entrainer le brouillage des réceptions radio et télé, ce qui obligerait le propriétaire ou l'opérateur à prendre les dispositions nécessaires pour en éliminer les causes.

The following paragraphs and Caution statements apply to products with the IBM FL200 Module (laser transducer) option installed. If your system has an ESCON Channel with the Extended Distance Feature, then the following paragraphs and Caution statements apply to your system.)

The IBM FL200 is certified in the U.S. to conform to the requirements of DHHS 21 CFR Subchapter J for Class 1 laser products. Elsewhere, it is certified to be in compliance with IEC 825 (1st edition 1984) and CENELEC HD 482 S1 as a class 1 laser product.The FL200 has been tested and approved in Sweden as a Class 1 laser product by Statens Provningsanstalt (Swedish National Testing Institute), and assigned the approval number SP LA 1989:118.

Class 1 laser products are not considered to be hazardous. Internally, the IBM FL200 Module contains a Class 3b laser that is nominally a 5.0 milliwatt Indium Gallium Arsenide Phosphide laser operating in the wavelength region of 1270-1340 nanometers. The FL200 module is designed such that there is never any human access to laser radiation above a Class 1 level during normal operation, user maintenance, or prescribed service conditions.

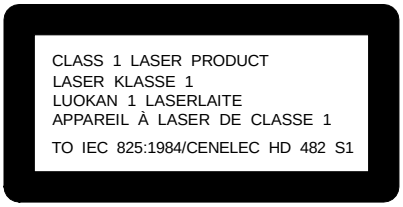

Figure 0-1. IBM FL200 Module label

The label in Figure 0-1 is required by IEC 825 to be affixed to the system, and is located inside the central processor complex (CPC).

CAUTION:
The FL200 laser ports are designed and certified for use only with single mode optical fiber and connectors having characteristics specified by IBM. The use of any other connectors/fiber may result in emission of laser power levels capable of producing injury to the eye if directly viewed. Use of non-specified connectors/fiber could violate the Class 1 Laser classification.

Coupling Facility Channel Laser Safety Information

CAUTION:
The coupling facility channel laser ports are designed and certified for use only with optical fiber and connectors having characteristics specified by IBM. The use of any other connectors or fiber may result in emission of laser power levels capable of producing injury to the eye if directly viewed. Use of non-specified connectors or fiber could violate the class 1 certification.

CAUTION:
Data processing environments can contain equipment transmitting on system links with laser modules operating at greater than class 1 power levels. For this reason, it is advised never to look into the end of an optical fiber cable or open receptable. The inspection or repair of optical fiber cable assemblies and receptacles should be performed by trained service personnel only.

Laser Compliance

The coupling facility channel modules are certified in the U.S. to conform to the requirements of DHHS 21 CFR Subchapter J for Class 1 laser products. Elsewhere, they are certified to be in compliance with IEC 825 (first edition 1984) and CENELEC HD 482 S1 as a Class 1 laser product. The coulping facility channel modules have also been tested and approved to comply with international Class 1 laser product certifications. Consult the label on each part for laser certification numbers and approval information.

Class 1 laser products are not considered to be hazardous. Internally, the coupling facility channel modules operating on multimode fiber contain a Class 3b laser that is nominally a 5.0 milliwatt source operating in the wavelength region of 780-850 nm. Internally, the coupling facility channel modules operating on single-mode fiber contain a class 3b laser that is nominally a 5.0 milliwatt source operating in the wavelength region of 1270-1340 nanometers. All coupling facility channel modules are designed so that there is never any human access to laser radiation above a Class 1 level during normal operation, user maintenance, or prescribed service conditions.

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 publications. 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.

Authorized Use of IBM Online Books

For online versions of this book, we authorize you to:

- Copy, modify, and print the documentation contained on the media, for use within your enterprise, provided you reproduce the copyright notice, all warning statements, and other required statements on each copy or partial copy.
- Transfer the original unaltered copy of the documentation when you transfer the related IBM product (which may be either machines you own, or programs, if the program's license terms permit a transfer). You must, at the same time, destroy all other copies of the documentation.

You are responsible for payment of any taxes, including personal property taxes, resulting from this authorization.

THERE ARE NO WARRANTIES, EXPRESS OR IMPLIED, INCLUDING THE WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE.

Some jurisdictions do not allow the exclusion of implied warranties, so the above exclusion may not apply to you.

Your failure to comply with the terms above terminates this authorization. Upon termination, you must destroy your machine readable documentation.

Trademarks and Service Marks

The following terms, denoted by an asterisk (*) in this publication, are trademarks or service marks of the International Business Machines Corporation in the United States or other countries:

AD/Cycle
AIX
AIX/ESA
AT
BookManager
CICS/ESA
DB2
ECKD
ESCON
Enterprise System/9000
Enterprise Systems Architecture/370
Enterprise Systems Architecture/390
Enterprise Systems Connection Architecture
Enterprise System/3090
Hardware Configuration Definition
ES/3090
ES/9000
ESA/370
ESCON XDF
ESA/390
Hiperspace
IBM
PS/2
SAA
S/390
Site Alert Client
Site Connect Server
Sysplex Timer
System/360
System/370
System/390
VM/ESA
VM/XA
VSE/ESA
Thinkpad
IMS/ESA
MVS/DFP
MVS/ESA
MVS/SP
Personal Security
Personal System/2
PR/SM
Processor Resource/Systems Manager

The following terms, denoted by double asterisks (**) in this publication, are trademarks or registered trademarks of other organizations:

ADMOD (Trademark of EXIDE Electronics)
EXIDE Electronics (Trademark of EXIDE Electronics)
NetWare (Trademark of Novell, Inc.)
NFS (Trademark of Sun Microsystems, Inc.)

About This Publication

This publication describes the design, components, functions, features, and capabilities of the:

- IBM* S/390* R1 Based 9672 Parallel Transaction Server

  The R1 Based 9672 Parallel Transaction Server refers to the 9672 E-Models and P-Models.
- IBM S/390 R1/R2/R3 Based 9672 Parallel Enterprise Server R1/R2/R3 Models

  The R1/R2/R3 Based 9672 Parallel Enterprise Server refers to the 9672 R1, R2, and R3 Models.
- IBM S/390 R1/R2/R3 Based 9674 Coupling Facility C01/C02/C03 Models

  The R1/R2/R3 Based 9674 Coupling Facility refers to the 9674 C01, C02, and C03 Models.

This publication is intended for executives, data processing managers, and data processing technical staff.

Before reading this publication you should be familiar with IBM System/390* and IBM Enterprise Systems Architecture/390* (ESA/390*) as described in *IBM Enterprise Systems Architecture/390 Principles of Operation*, SA22-7201.

You should also be familiar with the IBM Enterprise Systems Connection Architecture* (ESCON* Architecture) as described in *Introducing Enterprise Systems Connection*, GA23-0383.

What is Included in This Publication

This publication contains the following chapters and appendixes:

- Chapter 1, “Introduction” outlines each model's hardware configuration, design highlights, and control programs.
- Chapter 2, “General Product Characteristics” describes the characteristics of each model.
- Chapter 3, “Hardware Description” describes the logical components of the Central Processor Complex (CPC).
- Chapter 4, “Hardware Management Console” describes the Hardware Management Console.
- Chapter 5, “Standard and Optional Features” describes the standard and optional features.
- Chapter 6, “Coupling Facility” describes the operation of the coupling facility.
- Chapter 7, “Characteristics of ESCON and Parallel Channel Operation” describes the characteristics, function, and structure of the channels.
- Chapter 8, “ESCON Multiple Image Facility (EMIF)” describes the operation of the ESCON Multiple Image Facility (EMIF).
- Chapter 9, “Coupling Facility Channels” describes the characteristics, function, and structure of the coupling facility channels.
- Chapter 10, “ESCON and Parallel Channel Subsystem Performance” describes ESCON and parallel channel performance concepts and characteristics, and the criteria for determining the sequence of attachment of the input/output (I/O) devices for parallel channels.
- Chapter 11, “Supported I/O Devices” contains a table that lists the I/O devices capable of attaching to a 9672 Central Processor Complex (CPC).
- Chapter 12, “Upgrade Progression” contains a table that lists the model upgrade progression within and between models.

Related Publications

Other IBM publications that you will find helpful and that you should use along with this publication include:

- *System/360 and System/370 I/O Interface Channel to Control Unit Original Equipment Manufacturers' Information*, GA22-6974
- *System/370: Principles of Operation*, GA22-7000
- *System/370: Input/Output Configurator*, GA22-7002
- *System/370 Extended Architecture Interpretive Execution*, SA22-7095
- *Enterprise Systems Architecture/390 Vector Operations*, SA22-7125
- *Enterprise Systems Architecture/370: Principles of Operation*, SA22-7200
- *Enterprise Systems Architecture/390: Principles of Operation*, SA22-7201
- *Enterprise Systems Architecture/390: ESCON I/O Interface*, SA22-7202
- *Enterprise Systems Architecture/390: Common I/O Device Commands*, SA22-7204
- *Input/Output Equipment Installation Manual–Physical Planning for System/360, System/370, and 4300 Processors*, GC22-7064
- *ES/9000, ES/3090* Processor Resource/Systems Manager Planning Guide*, GA22-7123
- *ES/9000 Input/Output Configuration Program User's Guide*
  ESCON Channel-to-Channel Reference, GC38-0401
- *Introducing Enterprise Systems Connection*, GA23-0383
- *Introduction to Nonsynchronous Direct Access Storage Subsystems*, GC26-4519

9672 Parallel Transaction Server, 9672 Parallel Enterprise Server R1/R2/R3 Models, and 9674 Coupling Facility C01/C02/C03 Models

- *Hardware Management Console Guide*, GC38-0453
- *Operations Guide*, GC38-0454
- *PR/SM Planning Guide*, GA22-7123
- *IOCP User's Guide*, GC38-0401
- *Stand-Alone IOCP User's Guide*, GC38-0455

System/390 (S/390) Parallel Sysplex: A parallel sysplex uses a coupling facility to improve data sharing among the systems participating in the sysplex. The following publications provide additional information to help you understand and prepare for a sysplex that uses a coupling facility for data sharing purposes.

- *Sysplex Overview*, GC28-1208, provides an overview of the sysplex environment and coupling facility data sharing.
- *Sysplex Hardware and Software Migration*, GC28-1210, provides hardware and software planning information to help you migrate to a sysplex with a coupling facility.
- *Sysplex Systems Management*, GC28-1209, provides planning considerations for sysplex systems management.
- *Sysplex Application Migration*, GC28-1211, provides planning considerations for defining CICS/ESA and IMS/ESA applications to run in an parallel sysplex environmemt.
- *Setting Up a Sysplex*, GC28-1449, provides information to help you define MVS/ESA policies for managing a sysplex that uses a coupling facility for data sharing.

S/390 Opens Systems Adapter (OSA) Feature: The following publications provide additional information for planning and using the OSA feature:

- *Planning for the S/390 Open Systems Adapter Feature*, GC23-3870.
- *Using the S/390 Open Systems Adapter Support Facility for MVS/ESA*, SC23-3872.

Summary of Changes

Summary of Changes for GA22-7148-06

S/390 9672 Parallel Transaction Server E/P Models
S/390 9672 Parallel Enterprise Server R1/R2/R3 Models
S/390 9674 Coupling Facility C01/C02/C03 Models

System Overview
For R1/R2/R3 Based Models

This revision contains editorial changes and the following technical changes.

General: References to 9672 R-Models and 9674 C-Models in the previous version of this document have been changed to 9672 Parallel Enterprise Server R1/R2/R3 Models and 9674 Coupling Facility C01/C02/C03 Models respectively. All general references to 9672's include 9672 Parallel Transaction Server E and P Models and 9672 Parallel Enterprise Server R1/R2/R3 Models.

New Information

- **Enhancements**
  - OSA-2 Asynchronous Transfer Mode (ATM)

    A new connectivity option is available on 9672 R2 and R3 Models that allows direct attachment of Asynchronous Transfer Mode (ATM)-based networks to the central processor complex (CPC). ATM technology allows the integration of voice, video, traditional data, and other traffic types on a single network.

Chapter 1. Introduction

1.1 Overview

The 9672 Parallel Transaction Server, 9672 Parallel Enterprise Server R1/R2/R3 Models, and 9674 Coupling Facility C01/C02/C03 Models represent a new generation of IBM Central Processor Complexs (CPCs) that feature an enhanced price-to-performance ratio, support for the S/390 Parallel Sysplex, enhanced hardware management controls, and enhanced CPC packaging to address business processing needs.

1.2 The Benefits

1.2.1 Enhanced Price-to-Performance Ratio

The 9672 Parallel Transaction Server, 9672 Parallel Enterprise Server R1/R2/R3 Models, and 9674 Coupling Facility C01/C02/C03 Models, built using complementary metal oxide semiconductor (CMOS) technology, take advantage of the reduced manufacturing cost associated with this less expensive, high performance technology to produce a more competitively priced product.

Additionally, these CMOS-based CPCs offer the added advantage of reduced operational costs in terms of reduced power, cooling, and floor space requirements.

1.2.2 Support for the S/390 Parallel Sysplex

The 9672 Parallel Transaction Server, the 9672 Parallel Enterprise Server R1/R2/R3 Models, and the 9674 Coupling Facility C01/C02/C03 Models all support the S/390 Parallel Sysplex. S/390 Parallel Sysplex support consists of having the capability to do one or more of the following:

- Install coupling facility channels
- Define, as a logical partition, a portion or all of the Central Processor Complex (CPC) hardware resources (central processors, storage, and channels) for use as a coupling facility that connects to MVS images which share data in the coupling facility
- Connect to a coupling facility to share data
- Support the integrated coupling migration facility (ICMF) which allows you to define a test or migration coupling facility without coupling facility channel hardware on a single CPC.

The 9672 Parallel Transaction Server and 9672 Parallel Enterprise Server R1/R2/R3 Models support all of these capabilities, except for the Model R11 which does ***not*** support coupling facility channels. The 9674 Coupling Facility C01/C02/C03 Models do ***not*** support ICMF or connecting to a coupling facility. See Chapter 6, “Coupling Facility” on page 6-1.

1.2.3 Enhanced Hardware Management Controls

The Hardware Management Console provides a hardware management platform that can control and monitor status for multiple 9672 and 9674 Central Processor Complexs (CPCs) providing a single point of control and single system image for those CPCs defined to it.

The Hardware Management Console focuses multi-CPC operations by using exception processing to monitor status according to criteria you set. Panel colors indicate different system status conditions and color changes notify you when a status change meaningful to you requires your attention.

Additionally, the Hardware Management Console provides the following capabilities:

- An easy-to-use, object-oriented interface running as an OS/2 application (Hardware Management Console Application)
- Customizable groups of hardware objects you can perform operator actions on (for example, CPCs, CPC images, or coupling facility CPC images)
- Customizable Hardware Management Console Application presentation characteristics
- Consolidated hardware status reporting
- Consolidated hardware message presentation
- Consolidated operating system and Coupling Facility Control Code message presentation of messages received through the console integration interface
- Consolidated problem analysis and reporting
- Consolidated Licensed Internal Code (LIC) control and distribution
- Remote I/O configuration definition capability through use of the Hardware Configuration Definition* (HCD)
- Scheduled operations capability
- A variety of remote operations capabilities
- Automation support
- Remote support facility

See Chapter 4, “Hardware Management Console” on page 4-1 for more information.

1.2.4 Enhanced CPC Packaging to Better Address Business Processing Needs

The 9672 Parallel Transaction Server, the 9672 Parallel Enterprise Server R1/R2/R3 Models, and the 9674 Coupling Facility C01/C02/C03 Models are packaged to economically address business processing needs. In addition to providing an enhanced price-to-performance ratio and reduced operational costs, use of CMOS technology enhances hardware cost effectiveness in support of the S/390 Parallel Sysplex.

1.2.4.1 9672 Parallel Transaction Server E-Models and P-Models

9672 Parallel Transaction Server E-Models and P-Models are strings of individual CPCs. 9672 Parallel Transaction Server E-Models and P-Models offer coupling facility, ESCON, and parallel channel attachment. Parallel channel attachment is available through:

- 78-pin D-shell connectors (E-Models and P-Models)
- Bus and tag with tailgates (P-Models only)

Cost Effective Data Sharing Capability, Capacity, and Scaleability: 9672 Parallel Transaction Server E-Models and P-Models are strings of CPCs (from one to eight CPCs, each with their own power, central processors, storage, and channel resources) physically grouped together in close proximity to share network and Sysplex Timer* connectivity.

9672 Parallel Transaction Server E-Models and P-Models provide hardware data sharing capability and capacity in support of the S/390 Parallel Sysplex. 9672 E-Models and P-Models allow you to maximize S/390 parallel processing capability and capacity by:

- Enhancing the number of CPCs and logical partitions (LPs) that can take part in S/390 Parallel Sysplex data sharing which means more parallelizing of the workload across more distinct CPC and CPC image boundaries.

- Enhancing the data sharing capability of each CPC by supporting an increased number of optional coupling facility channels that can be shared among logical partitions that attach to a coupling facility.
- Allowing local area network (LAN) and Sysplex Timer connectivity to be shared among a string of CPCs, maximizing the connectivity and availability of these necessary resources and conserving ports on the Sysplex Timer.

9672 E-Models and P-Models allow you to incrementally add CPCs to existing 9672 strings to enhance S/390 parallel processing capability and capacity. You can dynamically install and configure additional CPCs to the data sharing configuration without disrupting workload or system availability.

Continuous Workload and System Availability: By implementing flexible S/390 parallel processing across a number of CPCs and CPC LPs, 9672 Parallel Transaction Server E-Models and P-Models provide continuous workload and system availability by:

- Insulating the workload and cumulative hardware system from individual hardware system outages, scheduled or unscheduled. You can better control the available capacity of the workload in the face of individual system outages by spreading the workload across more CPCs and LPs.
- Non-disruptive addition of data sharing processing capability during peak processing occasions by allowing the dynamic addition of CPCs, LPs, or both to the operating sysplex.
- Scheduled, non-disruptive upgrade of the 9672 Parallel Transaction Server Licensed Internal Code (LIC) (patches and engineering changes). You can enhance the quality of the system LIC by rippling the LIC changes on a CPC-by-CPC basis by non-disruptively removing CPCs one at a time from the ongoing sysplex and receiving, testing, and applying the LIC changes before non-disruptively returning the CPC to the data sharing configuration.

For more detail on the content of 9672 E and P models, see 3.10, “Hardware Configuration for the 9672 Parallel Transaction Server” on page 3-20.

1.2.4.2 9672 Parallel Enterprise Server R1/R2/R3 Models

| The 9672 Parallel Enterprise Server R1/R2/R3 Models include:

- R1 Models: R11, R21, R31, R41, R51, and R61.
- R2 Models: RA2, R12, R22, R32, R42, R52, and R72. (RA2 has 1 integrated CP.)
- R3 Models: R53, R63, R73, R83, and RX3. (RX3 model has 10 integrated CP's.)

9672 Parallel Enterprise Server R1/R2/R3 Models are Central Processor Complexs (CPCs) that provide a cost effective way to take advantage of S/390 processing capability while protecting your current investment in information assets, S/390 applications, and systems management. The 9672 Parallel Enterprise Server R1/R2/R3 Models offer ESCON, parallel, coupling facility, and OSA channel attachment, except for the Model R11 which supports ESCON, parallel, and OSA channels. Parallel channel attachment is available through:

- For R1: standard bus and tag with a tailgate or 78 pin D-shell connectors
- For R2 and R3: 78 pin D-shell connectors

Additionally, the 9672 Parallel Enterprise Server R1/R2/R3 Models (except for the Model R11) can participate fully in a S/390 parallel sysplex. Prior to adding coupling facility channels, you can test data sharing applications by using the integrated coupling migration facility (ICMF) on the 9672 Parallel Enterprise Server R1/R2/R3 Models. Following installation of coupling facility channels, ICMF is available for further testing purposes.

The 9672 R2 and R3 models provide not only large increases in processing capacity over the 9672 R1 models, but also can have up to twice as much memory and up to four times the channel capacity. Customers can upgrade current 9672 Parallel Enterprise Server R1 models to the new models, extending the performance capability. Customers upgrading from 9672 Parallel Transaction Servers have increased flexibility in their system choices.

For more detail on the content of 9672 R models, see 3.11, “Hardware Configuration for the 9672 Parallel Enterprise Server R1/R2/R3 Models” on page 3-24.

1.2.4.3 9674 Coupling Facility C01/C02/C03 Models

The 9674 Coupling Facility C01/C02/C03 Models CPCs provide cost effective performance only as coupling facilities. In addition to providing an enhanced price-to-performance ratio and reduced operational costs, the 9674 Coupling Facility C01/C02/C03 Models are recommended as the coupling facility to be used in production data sharing configurations.

The 9674 Coupling Facility C01/C02/C03 Models are designed to run as a single dedicated coupling facility logical partition to reduce the risk of long outages, achieve acceptable performance, provide large capacity shared storage, and maximize connectivity to the coupling facility (up to 32 coupling facility channels.).

For high availability, IBM recommends that you install a second, similarly configured, 9674 to reduce the possibility of a single point of failure. A second 9674 improves application software availability by allowing fast recovery from one coupling facility to the other in the event of a coupling facility outage.

The 9674 C02 and C03 models provide not only large increases in coupling efficiencies over the 9674 C01 models, but also can have up to twice as much memory. Customers can upgrade current 9674 Coupling Facility models to the new models, maximizing the coupling efficiency.

For more detail on the content of 9674 C models, see 3.12, “Hardware Configuration for the 9674 Coupling Facility C01/C02/C03 Models” on page 3-30.

1.3 Design Highlights

1.3.1 9672 Parallel Transaction Server and 9672 Parallel Enterprise Server R1/R2/R3 Models

9672 models provide high performance and flexibility due to an improved design and use of technology advances. The design of these models supports:

- ESA/390 architectural mode of operation (ESA/390 or ESA/390 TPF)
- Logically partitioned (LPAR) operating mode. LPAR mode allows you to define ESA/390, ESA/390 TPF, System/370*, and coupling facility logical partitions. You can define up to ten logical partitions for each CPC.
- IBM Enterprise Systems Connection Architecture (ESCON Architecture) and technology for the ESCON channels
- ESCON Multiple Image Facility (EMIF)
- S/390 parallel sysplex
- S/390 Opens Systems Adapter (OSA) feature
- Data compression
- Dynamic I/O configuration
- IBM Sysplex Timer attachment
- Subspace group facility
- Asynchronous data mover facility (ADMF)
- Logical string assist
- Asynchronous pageout facility
- High levels of reliability, availability, and serviceability
- Online information
- Softcopy publications

1.3.2 9674 Coupling Facility C01/C02/C03 Models

The 9674 Coupling Facility C01/C02/C03 Models provide high performance and flexibility due to an improved design and use of technology advances. The design of these models supports:

- Logically partitioned (LPAR) operating mode that allows you to define coupling facility logical partitions. You can define up to ten logical partitions for a CPC.
- S/390 parallel sysplex
- High levels of reliability, availability, and serviceability
- Online information
- Softcopy publications

1.3.3 ESA/390 Architectural Mode of Operation

9672 CPCs provide the Enterprise Systems Architecture/390 (ESA/390 or ESA/390 TPF) mode in three ways:

- Basic mode
- Mode of a logical partition in LPAR mode
- Mode of a guest virtual machine

In ESA/390 mode (which includes the functions of ESA/370* mode), these models have problem-program compatibility with System/360*, System/370, and 4300 processors. They can access virtual storage in multiple address spaces and data spaces. This extends addressability for system, subsystem, and application functions that use ESA/390.

ESA/390 mode provides:

- 31-bit addressing with a virtual address range of 2GB
- ESCON, parallel, coupling facility, and Open Systems Adapter (OSA) channels
- Channel path selection and I/O-busy-condition management as hardware functions (rather than control program functions) that provide:
  - As many as eight channel paths available to each I/O device.
  - Increased I/O device accessibility by allowing each central processor to initiate operations with any of the I/O devices and to handle any I/O interruption conditions.
- A significantly extended addressability through access to multiple address spaces and data spaces while maintaining compatibility with existing 24-bit and 31-bit subsystems and user applications. Each address space can contain as many as 2GB of programs and data. Each data space can contain as many as 2GB of data.

1.3.4 Logically Partitioned (LPAR) Operating Mode

You select the operating mode during system initialization (power-on reset). For 9672 CPCs, the operating mode can be basic mode (ESA/390 or ESA/390 TPF) or LPAR mode. For the 9674 Coupling Facility C01/C02/C03 Models, the operating mode is LPAR mode.

When you select LPAR mode, you must specify the:

- Architectural mode (ESA/390, ESA/390 TPF, System/370, or coupling facility mode) for each logical partition. The 9674 Coupling Facility C01/C02/C03 Models only support coupling facility mode.
- Resources for each logical partition
  - Central processors (CPs)
  - Storage (central storage and optional expanded storage)
  - Channels

You can define up to ten logical partitions.

1.3.4.1 Central Processors (CPs)

CPs can be dedicated to a single logical partition or shared among multiple logical partitions. The allocation of CPs to a logical partition is made when the logical partition is activated. The use of CP resources shared between logical partitions can be limited and modified by operator commands while the logical partitions are active. CPs that are dedicated to a logical partition are available only to that logical partition.

1.3.4.2 Storage

Before you can activate logical partitions, you must define central storage and optional expanded storage to the logical partitions. When a logical partition is activated, the storage resources are allocated in contiguous blocks. 9672 Parallel Enterprise Server R1/R2/R3 Models and 9674 Coupling Facility C01/C02/C03 Models support a 1MB storage granularity for both central and expanded storage. You can dynamically reallocate storage resources for ESA/390 mode logical partitions using dynamic storage reconfiguration. After you define and activate an ESA/390, ESA/390 TPF, or S/370 logical partition, you can load a supporting operating system into that logical partition.

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

1.3.4.3 Channels

You can allocate channels to logical partitions as follows:

- Dedicated channels
- Reconfigurable channels
- Shared channels. See Chapter 8, “ESCON Multiple Image Facility (EMIF)” on page 8-1.

Dedicated Channels: For each logical partition that you define, you can dedicate individual:

- Parallel channels (BY or BL channel paths)
- ESCON channels that attach to an ESCON Converter Model 1 (CVC or CBY channel paths)
- ESCON channels (CNC or CTC channel paths)
- Coupling facility channels (CFS or CFR channel paths)
- Open Systems Adapter channels (OSA channel paths)

Dedicated channels are unshared channels and can only be used by those logical partitions to which they are allocated.

Reconfigurable Channels: You can define the following channels as reconfigurable among logical partitions:

- Parallel channels (BY or BL channel paths)
- ESCON channels that attach to an ESCON Converter Model 1 (CVC or CBY channel paths)
- ESCON channels (CNC or CTC channel paths)
- Coupling facility sender channels (CFS channel paths)
- Open Systems Adapter channels (OSA channel paths)

Reconfigurable channels are unshared channels that can be dynamically moved among logical partitions but can only belong to one logical partition at a given time.

Shared Channels: You can define the following channels as shared among multiple logical partitions so that the shared channels can be accessed by more than one logical partition at the same time:

- ESCON channels (CNC or CTC channel paths)
- Coupling facility sender channels (CFS channel paths)
- Open Systems Adapter channels (OSA channel paths)

See 8.3.2, “Channel Path Definition” on page 8-4.

Device Sharing: You can share a device among logical partitions by:

- Using a separate channel for each logical partition
- Using a shared channel. See 8.2.2, “Device (or Data) Sharing” on page 8-3.

For more information, see *ES/9000, ES/3090* Processor Resource/Systems Manager Planning Guide*, GA22-7123 and *ES/9000 Input/Output Configuration Program User's Guide and ESCON Channel-to-Channel Reference*, GC38-0401.

1.3.5 ESCON Architecture and Technology

ESCON combines technology, architecture, and a set of interrelated hardware and software products and services that provide:

- ESCON Architecture
- ESCON channels, including the ESCON extended distance feature (ESCON XDF*)
- Fiber optic cabling, transmission, and reception
- Dynamic connectivity through switched point-to-point topology and data flow
- Interconnectivity with other networks

For detailed information about the basic concepts of ESCON technology and architecture, see *Introducing Enterprise Systems Connection,* and *ESA/390 Principles of Operation*.

1.3.5.1 ESCON Architecture

ESCON Architecture, a channel architecture, is designed to support the fiber optic environment and dynamic connectivity.

ESCON Architecture uses link-level and device-level protocols for the transfer of information. Link-level protocols establish and maintain the physical and logical path for the transmission and reception of information. Device-level protocols establish and maintain the transfer of information between a channel and control units that use device-level protocols.

For detailed information about the ESCON interface protocol, see *IBM Enterprise Systems Architecture/390: ESCON I/O Interface*.

1.3.5.2 ESCON Channels

ESCON Channels transfer information on a link in a serially transmitted synchronous bit stream (serial transmission) through fiber optic channel cables and:

- Operate using link-level protocols and device-level protocols, or
- Attach to an IBM 9034 ESCON Converter to attach to control units with parallel interfaces, using bus and tag cable.

For detailed information about the I/O commands used to operate the new ESCON interface protocols in the fiber optic environment, see *IBM Enterprise Systems Architecture/390: Common I/O Device Commands*.

1.3.5.3 Fiber Optic Cabling, Transmission, and Reception

ESCON LED uses multimode fiber optic technology to provide a direct channel attachment range of up to 3 kilometers (1.9 miles); and control units can be attached to a range of up to 9 kilometers (5.6 miles) from a channel through two optional ESCON Directors.

ESCON XDF uses Laser/single-mode fiber optic technology to provide a direct channel attachment range of up to 60 kilometers (37.3 miles); and control units can be attached to a range of up 43 kilometers (26.7 miles) from a channel through two optional ESCON Director.s. The following table shows the maximum cable distances that ESCON I/O devices can be located from the CPC using ESCON Directors with ESCON XDF and ESCON I/O.

Table 1-1. Connectivity Distances Using ESCON XDF, ESCON Directors, and ESCON I/O

Device	Maximum Link Distance	Minimum Number of ESCON Director.s Required
ESCON Converter Models 1 and 2	3 km (1.9 miles)	1
3990 Models 2 and 3	15 km (9.3 miles)	1
9343 Model D04 and DC4	43 km (26.7 miles)	2
3490 Models A01, A02, A10, A20, C10, C11, C22, D31, D32, D41, and D42 (see note)	23 km (14.3 miles)	1
3174 Models 12L and 22L	43 km (26.7 miles)	2
3172 Model 1	43 km (26.7 miles)	2
3745	43 km (26.7 miles)	2
RS/6000	43 km (26.7 miles)	2
ES/9000 ESCON Channel-to-Channel (CTC)	60 km (37.3 miles)	2

By utilizing Laser/single-mode fiber optic technology (i.e. 9/125 micron fiber) connectivity distances can significantly exceed those offered by LED/multimode fiber technology. While ESCON LED/multimode technology is the cabling implementation for campus-wide interconnection, ESCON XDF can be used when unrepeated link distances up to a maximum of 20 kilometers, are required. Repeated distances of up to 60 kilometers can be obtained with ESCON XDF.

1.3.5.4 Dynamic Connectivity (Switched Point-to-Point Topology)

Using switched point-to-point topology and ESCON Directors, you can handle many channel and control unit connections simultaneously. The path between each point-to-point interconnection is called a link. With dynamic connectivity:

- Many control units can share a single link to a channel through an ESCON Director.
- Many channels can share a single link to a control unit through an ESCON Director.
- I/O configuration management is simplified through the use of dynamic I/O configuration.
- Alternate paths can be used if a failure occurs.

1.3.5.5 Interconnectivity with Other Networks

Using the IBM 3172 Interconnect Controller and ESCON interface you can send large amounts of data at wideband speeds to remote CPCs, or a parallel (bus and tag) interface to an IBM 9034 ESCON Converter to use existing hardware and I/O interfaces that access commonly used local area networks.

1.3.6 ESCON Multiple Image Facility (EMIF)

The ESCON Multiple Image Facility (EMIF) allows multiple logical partitions to directly share channels and to optionally share any of the control units and associated I/O devices configured to these shared channels. ESCON Multiple Image Facility also provides a way to limit the logical partitions that can access a reconfigurable channel or a shared channel to enhance security.

With ESCON Multiple Image Facility, the CSS provides an independent set of I/O controls for each logical partition called a CSS image. Each logical partition is configured to a separate CSS image in order to allow the I/O activity associated with each logical partition to be processed independently as if each logical partition had a separate CSS.

For more detailed information about EMIF, see Chapter 8, “ESCON Multiple Image Facility (EMIF)” on page 8-1.

1.3.7 System/390 (S/390) Parallel Sysplex Support

The S/390 Parallel Sysplex makes use of a broad range of hardware and software products to process in parallel a transaction processing workload across multiple MVS images running in a sysplex and sharing data in a coupling facility.

The S/390 Parallel Sysplex allows you to manage a transaction processing workload, balanced across multiple MVS images running on multiple Central Processor Complexs (CPCs), as a single data management system. It also offers workload availability and workload growth advantages.

The S/390 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 for the S/390 Parallel Sysplex consists of having the capability to do any or all of the following:

- Install coupling facility channels

- Define, as a logical partition, 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 MVS images for data sharing purposes
- Connect to a coupling facility to share data
- Support the integrated coupling migration facility (ICMF) which allows you to define test or migration coupling facility ***without*** coupling facility channel hardware on a single CPC.

Summary of CPC Support for the S/390 Parallel Sysplex: The 9672 Parallel Transaction Server, the 9672 Parallel Enterprise Server R1/R2/R3 Models, and the 9674 Coupling Facility C01/C02/C03 Models provide the following support for the S/390 Parallel Sysplex:

- Each CPC that is part of a 9672 Parallel Transaction Server provides full support for the S/390 Parallel Sysplex.
- 9672 Parallel Enterprise Server R1/R2/R3 Models (except for Model R11) provide full support for the S/390 Parallel Sysplex. Model R11 ***cannot*** have coupling facility channels installed.
- The 9674 Coupling Facility C01/C02/C03 Models can only define CFR channel paths and run as a coupling facility; it ***cannot*** run MVS or any other operating system and it does ***not*** support ICMF.

1.3.7.1 Coupling Facility

The coupling facility allows applications, running on multiple MVS 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.

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 MVS images in the sysplex define the shared structures used in the coupling facility.

Enhancements to PR/SM* LPAR allow 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 permitted 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 hardware operator console. PR/SM LPAR limits the hardware operator controls usually available for LPs to avoid unnecessary operator activity. However, Coupling Facility Control Code does provide a limited set of new hardware operator controls unique to the coupling facility LP. See 6.10, “Coupling Facility Control Code Commands” on page 6-8.

For more detailed information about Coupling Facilities, see Chapter 6, “Coupling Facility” on page 6-1.

New 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 Chapter 9, “Coupling Facility Channels” on page 9-1.

1.3.8 S/390 Open Systems Adapter (OSA) Feature Support

9672 models E, P, and R1 at engineering change (EC) D79759 include support for the optional S/390 Opens Systems Adapter (OSA-1) feature. Support for the optional S/390 Opens Systems Adapter (OSA-1 and OSA-2) features on 9672 R2 and R3 models is available at EC E45748 or later.

The OSA feature provides:

- OSA-1 and OSA-2 connectivity options that allow direct attachment of Ethernet, Token Ring, and Fiber Distributed Data Interface (FDDI) local area networks (LANs) to the central processor complex (CPC)
- A new OSA-2 connectivity option on 9672 R2 and R3 models at EC E45748B (with CPC EC E45568 + MCL, and Hardware Management Console EC E45552 + MCL) or later that allows for attachment of Asynchronous Transfer Mode (ATM)-based networks to the CEC.

The OSA feature supports the CPC acting as a enterprise-wide server for multiple networks. LANs, LAN backbones, high speed workstations, hubs, bridges, and routers that conform to these LAN interface standards can connect directly to the S/390 Opens Systems Adapter, benefiting from the performance, connectivity, security, availability, and growth opportunities associated with the S/390 architecture.

New OSA channel hardware, available as an orderable feature, provides the connectivity required for direct LAN attachment to the CPC. See 1.3.8.3, “S/390 Opens Systems Adapter (OSA) Feature” on page 1-16.

1.3.8.1 OSA Support Facility for MVS/ESA

9672 E/P/R1 models at EC D79759 and 9672 R2/R3 models at EC E45748 provide support for the OSA Support Facility (OSA/SF) for MVS/ESA. The OSA Support Facility for MVS/ESA is an application that operates as a started task under MVS/ESA. It provides configuration and operations management capabilities for the S/390 Opens Systems Adapters in the configuration and allows you to install supported application function onto S/390 Opens Systems Adapters to support client server applications integrated at the CPC. S/390 OSA Support Facility Version 1 Release 2 is required to download the ATM function onto the 155 ATM S/390 Open Systems Adapter 2 feature.

Software Corequisites: The OSA Support Facility for MVS/ESA requires the following minimum levels of software support:

- MVS/ESA Version 4 Release 3 with the Hardware Configuration Definition 5.1 Usability Feature ***or*** MVS/ESA Version 5 Release 1 or higher
- Time Sharing Options Extensions (TSO/E) Version 2 Release 4
- C/370 Library Version 2 Release 2 For MVS
- Interactive System Productivity Facility (ISPF) Version 3 Release 5, or equivalent, to edit MVS/ESA datasets at the host

1.3.8.2 Support for the TCP/IP, SNA, and IPX Network Protocols

The OSA feature supports data transfer between S/390 host server programs and their network clients using the following network, or transport, protocols.

Transmission Control Protocol/Internet Protocol (TCP/IP) through the OSA TCP/IP Passthru mode of operation. Both an OSA-1 and an OSA-2 feature can run in this mode to transfer data between TCP/IP host programs running in an OS/390, MVS/ESA, or VM/ESA environment and their network clients.

| ***SNA/APPN (SNA) network protocol*** through the OSA SNA mode of operation. Both an OSA-1 and an OSA-2 feature can run in this mode to transfer data between ACF/VTAM applications running in an OS/390, MVS/ESA, or VM/ESA environment and their network-based clients.

| ***Internet Packet Exchange (IPX) protocol in the OSA LANRES/MVS mode*** through two OSA modes that are available on an OSA-1 feature in either an OS/390 or MVS/ESA environment:

- In the LANRES/MVS mode, an OSA-1 feature transfers data between NetWare** LAN-based clients and the LAN Resource Extension and Services/MVS (LANRES/MVS) host program.
- In the LANRES/MVS SAA mode, an OSA-1 feature transfers data between Netware for SAA** LAN-based clients and ACF/VTAM in the host.

| ***The OSA Support Facility (OSA/SF) is also available.***

Unique among the OSA modes of operation, the TCP/IP Passthru mode does not require the use of the OSA Support Facility (OSA/SF) on an OSA-1, a FDDI OSA-2, or an ENTR OSA-2 feature as long as the OSA ports are not defined to be shared among the logical partitions to which the OSA is defined in this mode. When the ports are to be shared, the OSA must recognize the shared home IP addresses for which it requires input that the user supplies to OSA/SF. With or without portsharing, OSA/SF is always required for the TCP/IP Passthru mode on an ATM OSA-2. Also, it is always required for the SNA, LANRES/MVS, and LANRES/MVS SAA modes on any OSA feature.

On an MVS/ESA system image, OSA/SF runs as a licensed program (program number 5655-104). In an OS/390 system, OSA/SF is a base, non-exclusive element. Although OSA/SF does not run directly on a VM/ESA system, it can provide indirect management to OSAs running in the TCP/IP Passthru and SNA modes. These OSAs are defined to a partition in which VM/ESA is running. OSA/SF is on MVS/ESA, which is running either as a guest in the VM partition or in another partition to which the OSA is also defined. OSA/SF has an Operating System/2* (OS/2*) user interface that runs on a programmable workstation with the appropriate communications protocol to the S/390 host.

| ***An OSA must be defined in the system hardware I/O configuration.***

Not only must an OSA be defined as a S/390 channel with a logical, though not physical, control unit, each OSA mode has its own requirements on the I/O device numbers that need to be associated with the OSA channel path. As a unique type of S/390 channel, the OSA correlates these device numbers with the appropriate OSA port or ports to ensure that data is transferred to correct network client. The user must define these devices plus a device number for OSA/SF for each OSA channel path in the hardware I/O configuration using the appropriate host I/O configuration program (IOCP, HCD, dynamic I/O reconfiguration, etc.).

Given the variety of OSA modes of operation, OSA/SF user interfaces, and the host operating systems that are supported, documenting the programming requisites for OSA in useful detail falls outside the scope of this book. For that information, refer to:

- *Planning for the System/390 Open Systems Adapter Feature,* GC23-3870
- *Using the System/390 Open Systems Adapter Feature Support Facility for MVS/ESA,* SC23-3872
- *OS/390 Open Systems Adapter Feature Support Facility User's Guide,* SC23-3872
- The OSA/SF program directory

1.3.8.3 S/390 Opens Systems Adapter (OSA) Feature

The Open Systems Adapter (OSA) feature is a S/390 channel. Depending on the OSA, it provides one or more OSA ports that allow attachment for a FDDI LAN, Ethernet LANs, Token Ring LANs, or ATM-based networks (on 9672 R2 and R3 models at EC E45748B with CPC EC E45568 + MCL, and Hardware Management Console EC E45552 + MCL or later only) to the CPC.

You must use the OSA Support Facility for MVS/ESA to take advantage of the OSA feature's processing capability for applications other than TCP/IP passthru. The S/390 Opens Systems Adapter Support Facility for MVS/ESA is an MVS application that allows you to download supported applications to an S/390 Opens Systems Adapter, including code updates.

For additional detail on OSA channel options, see 3.5, “Optional Open Systems Adapter (OSA)” on page 3-10.

Chpid Numbering Considerations: OSA channel chpids are dependent on the model, expansion cage feature, and the number of IBB's in the system. Installing OSA channels can reduce the total number of ESCON or parallel channels installed on your CPC.

See 3.2.3.10, “Open Systems Adapter 1 (OSA-1) Channel Extender Cards” on page 3-6 and 3.2.3.12, “Open Systems Adapter 2 (OSA-2) Channel Cards” on page 3-7.

OSA Channels (TYPE=OSA): You can configure an OSA channel path as:

- A channel path in a basic mode I/O configuration
- An unshared dedicated channel path to a single logical partition (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 concurrently used by the LPs to which it is configured. The OSA Support Facility for MVS/ESA is required for LPs to share LAN ports associated with shared OSA channels.

1.3.9 Data Compression

9672 CPCs support hardware-assisted data compression for DB2*, IMS, DFSMS/MVS, and VTAM running under MVS/ESA*

New hardware facilities allow the compression and expansion of data which can provide significant reductions in the total cost of computing by significantly reducing:

- DASD and tape storage requirements
- The number of communication lines or required line speeds for network data transmission
- Central processor (CP) cycles needed to perform data compression and expansion (when compared to software-only data compression)

Data compression, along with the increased data rate of ESCON channels and sequential data striping, may provide further enhancements to the effective bandwidth of 9672 systems. When combined, these attributes may improve channel utilization and enhance system level performance in situations where I/O bandwidth has been a limiting factor.

1.3.9.1 Software Corequisites

Hardware-assisted compression requires the following minimum levels of software support:

- **For DB2:**
  - DB2 Version 3
  - MVS/ESA SP Version 4 Release 2.2 (4.2.2)
- **For IMS:**
  - IMS Version 4 Release 1 (4.1)
  - MVS/ESA SP Version 4 Release 3 (4.3)
- **For DFSMS/MVS:**
  - DFSMS/MVS Version 1 Release 2 (1.2)
  - MVS/ESA SP Version 4 Release 3 (4.3) with APAR support
- **For VTAM:**
  - VTAM Version 4 Release 2 (4.2)
  - MVS/ESA SP Version 4 Release 3 (4.3) with APAR support

1.3.10 Dynamic I/O Configuration

Dynamic I/O configuration allows you to modify the I/O configuration without performing:

- Power-on reset (POR) of the CPC
- Initial program load (IPL) of the system control program (SCP)

Dynamic I/O configuration allows you to add, delete, or modify the I/O definitions of channel paths, control units, and I/O devices in the CPC. It also allows you to save the changes you made to the I/O configuration definitions and apply them to the active I/O configuration data set (IOCDS).

Dynamic I/O configuration is operational in basic and LPAR modes. Dynamic I/O configuration does not support CFR channel paths.

Software Corequisites: For 9672 models, dynamic I/O configuration requires the following minimum software corequisites:

- MVS/ESA* SP Version 4 Release 2 (MVS/ESA SP 4.2) with APAR support or MVS/ESA SP 4.2.2 with APAR support provide the capability to make dynamic I/O configuration changes to configurations without shared channels.
- MVS/ESA SP 4.3 or later adds support for dynamic I/O configuration and ESCON Multiple Image Facility for ESCON channels, allowing dynamic I/O configuration changes for I/O configurations that have both shared and unshared channels.
- MVS/ESA SP 5.1 or later adds support for dynamic I/O configuration for ESCON byte multiplexer (CBY), coupling facility sender (CFS), and OSA channel paths, allowing dynamic I/O configuration changes.
- VM/ESA* Version 2 Release 1 or later provides support for dynamic I/O configuration for ESCON, parallel, coupling facility sender (CFS), and OSA channels, allowing dynamic I/O configuration changes for I/O configurations that have both shared and unshared channels.

1.3.11 IBM Sysplex Timer Attachment

The IBM Sysplex Timer, supported by MVS/ESA SP Version 4.1 and subsequent releases, synchronizes the processor's time-of-day clock in multiple systems. It is required by the 9672 Parallel Transaction Server and 9672 R models with coupling facility channels, and is optional for the 9672 Parallel Enterprise Server R1/R2/R3 Models without coupling facility channels. For more information, see 3.2.3.15, “Sysplex Timer Adapter Card” on page 3-7

1.3.12 Subspace Group Facility Support

9672 CPCs provide support for the subspace group facility which can enhance the data integrity and reliability of application server subsystems like CICS/ESA*, reducing application failures, service outages, and incorrect modification of critical business data.

Application server subsystems like CICS/ESA use a single application server program to provide common function to multiple application programs all running within a single address space. The subspace group facility allows the application server program to assign address space private storage to each application program running in the address space. This prevents one application program from accidentally overwriting the code and data of another application program running in the same address space. Preventing accidental overwrites reduces application failures, service outages, and incorrect modification of critical business data.

Additionally, there should be little or no modification required of existing application programs because you implement storage isolation and protection through the subsystem application server program.

1.3.12.1 CICS/ESA Use of the Subspace Group Facility

CICS/ESA uses the subspace group facility to implement transaction isolation to prevent one application program from accidentally overwriting the storage of another application program running in the same address space. Transaction isolation extends the subsystem storage protection capability which protects CICS system storage from being accidentally overwritten by user-written application programs.

Additionally, transaction isolation use of the subspace group facility enhances application development and maintenance for the subsystem. When an application attempts to access the address space private storage of another application, CICS/ESA abends the program making it easy to identify the failing application that could cause a data integrity problem.

Software Corequisites: Use of the subspace group facility and CICS/ESA transaction isolation requires the following minimum level of software support:

- MVS/ESA SP Version 5 Release 1 (5.1)
- CICS/ESA Version 4 Release 1 (4.1)

1.3.13 Asynchronous Data Mover Facility (ADMF)

9672 CPCs include support for the asynchronous data mover facility (ADMF). ADMF allows DB2 users in an MVS/ESA environment to more efficiently move large amounts of data between central and expanded storage. DB2 users may see a reduction in elapsed time for selected query applications by using Central Processor Complex (CPC) storage rather than DASD. ADMF can be used with data stored in a compressed format.

ADMF offloads the movement of data between central and expanded storage to the CPC I/O subsystem allowing CPs to accomplish additional work.

1.3.13.1 PR/SM LPAR Mode Considerations

In PR/SM LPAR mode, logical partitions (LPs) enabled to use ADMF share the facility. LP requests to use the facility are serviced on a first come, first served basis. You must deactivate an LP to enable or disable LP use of ADMF. Coupling facility LPs and LPs running in System/370 mode ***cannot*** use ADMF.

1.3.14 Logical String Assist

The logical string assist provides instructions that more efficiently process string data types. These instructions are of particular value to large, complex C-programming language applications that rely heavily on null-character-terminated string manipulation functions.

The logical string assist provides the following architected instructions:

COMPARE LOGICAL STRING
: Compares bytes between two strings until it either finds an inequality or a specified ending character

MOVE STRING
: Moves a string until it finds a specified ending character

SEARCH STRING
: Searches a string of a specified length for a specified character

The first two instructions–COMPARE LOGICAL STRING and MOVE STRING–are particularly useful to the C-programming language that normally delimits strings by a null ending character.

The IBM SAA* AD/Cycle* C/370 compiler will provide a compile-time option to generate these new instructions for new or re-compiled code.

For more information, see:

- *ESA/390 Principles of Operation*, SA22-7201-01

Software Corequisites: The logical string assist requires the following minimum level of software support:

- MVS/ESA SP Version 4 Release 3
- AD/Cycle C/370 Compiler Version 1 Release 1.1 with APAR support

1.3.15 Asynchronous Pageout Facility

The asynchronous pageout facility allows the CPC to more efficiently move pages of data from central storage to expanded storage. This can be valuable to processing environments where a large volume of block paging takes place (for example, numerically intensive computing [NIC]).

With the asynchronous pageout facility, a CPC can move a page of data from central storage to expanded storage without delaying the processing of subsequent instructions by the central processor (CP). The CP can continue to process instructions as soon as data movement begins.

This facility contrasts with the synchronous movement of data from central storage to expanded storage which caused the CP to delay processing of instructions until data movement completed.

Software Corequisites: The asynchronous pageout facility requires the following minimum level of software support:

- MVS/ESA SP Version 4 Release 3

1.3.16 Input/output Intensive Environments

The maximum number of central processors (CPs) available in a central processor complex (CPC) depends on the number of physical processors installed and the CPC configuration. The actual number available is determined automatically when the CPC is activated. An operating environment and power on reset mode should be chosen to yield optimal performance of the applications loaded on the CPC or its logical partitions.

When a CPC is activated in the processor intensive operating environment, the actual number of available CPs is the maximum number.

Activating a CPC in the input/output (I/O) intensive operating environment may reduce the actual number of CPs available. Depending on the 9672 model, one or two available CPs may be assigned as System Assist Processors (SAPs) instead of CPs. This increases the capacity of the channel subsystem to perform I/O operations.

No additional action is necessary if you intend to activate an I/O intensive CPC in a basic operating mode.

However, if you intend to activate an I/O intensive CPC in logically partitioned (LPAR) mode, a reduction in the number of CPs available will reduce the number of logical processors you can activate. Activation of a logical partition will fail if the number of logical processors you attempt to activate exceeds the number of CPs available. To avoid a logical partition activation failure, you must verify the number of logical processors assigned to a logical partition does not exceed the number of CPs available.

For each combination of operating environment and power-on reset mode, the following table shows examples of applications that require the environment and mode for their optimal performance.

Note: One or more applications shown as examples may not be available for the machine model of your CPC.

	Processor Intensive Environment	**Input/Output Intensive Environment**	**Coupling Facility Environment**
Logically Partitioned (LPAR) mode	Applications: - non-TPF system control program	Applications: - TPF Note: number of CPs may change	Applications: - CFCC
ESA/390 mode	Applications: - VM	Applications: - MVS, VSE, AIX, VSE under VM	Applications: - none
ESA/390 TPF mode	Applications: - none	Applications: - TPF Note: number of CPs may change	Applications: - none

For those application environments with high I/O rates the following models provide the capability to convert CPs into additional System Assist Processors (SAPs).

Models	**I/O Intensive Mode**
R41, R51, R61, R32, R42, R52, R72, R53	Converts a single CP into a second SAP
R63, R73, R83, RX3	Converts two additional CPs into a third and fourth SAP

This capability better balances the resources of the 9672 for some TPF environments.

1.3.16.1 Central Processor Identification

| The physical processors available in the 9672 R1, R2, and R3 models are assigned as Central Processors
| (CPs) and System Assist Processors (SAPs). The number of CPs and SAPs assigned for a particular model depends on the configuration. The configuration can be either ***Standard*** or ***I/O Intensive***. Figure 1-1 shows how the CPs are assigned for the 9672 R-models.

Model	Processor Intensive Configuration	I/O Intensive Configuration
R11	0 s	0 s
R21	0 1 s	0 1 s
R31	0 1 2 s	0 1 2 s
R41	0 1 2 3 s	0 1 2 s s
R51	0 1 2 3 4 s	0 1 2 3 s s
R61	0 1 2 3 4 5 s	0 1 2 3 4 s s
RA2	0 s	0 s
R12	0 s	0 s
R22	0 1 s	0 1 s
R32	0 1 2 s	0 1 s s (See Note 4)
R42	0 1 2 3 s	0 1 2 s s
R52	0 1 2 3 4 s	0 1 2 3 s s
R72	0 1 2 3 4 5 6 s	0 1 2 3 4 5 s s (See Note 4)
R53	0 1 2 3 4 s	0 1 2 3 s s
R63	0 1 2 3 6 7 s s	0 1 2 3 s s s s (See Note 4)
R73	0 1 2 3 4 6 7 s s	0 1 2 3 4 s s s s
R83	0 1 2 3 4 5 6 7 s s	0 1 2 3 4 5 s s s s
RX3	0 1 2 3 4 5 6 7 8 9 s s	0 1 2 3 4 5 6 7 s s s s

Figure 1-1. CP and SAP assignments for 9672 R-models.

Notes:

1. The CPs are identified in the Processor column of the Monitor Processors window. (Example: 2 on the figure indicates CP2).
2. The CP numbers are not continuous for some models.
3. The SAPs are identified as “s”
4. This configuration is not recommended for I/O Intensive operation.

Features of the 9674 C-models (C01, C02 and C03) allow the same selections of CPs as the 9672 R-models. (Example: C02 models can have 1, 2, 3, 4, 5, or 7 CPs). To determine the CP configurations for the C01, C02 and C03 models:

1. determine the number of CPs
2. find the line on Figure 1-1 for the corresponding 9672 R-model with the same number of CPs
3. look under Processor Intensive Configurations to find the CP configuration.

| **Note:** The 9674 C02 sub-uni model corresponds to the 9672 RA2 model.

1.3.17 High Levels of Reliability, Availability, and Serviceability

9672 and 9674 CPCs offer high levels of reliability, availability, and serviceability (RAS) with the following standard capabilities:

- Automated approach to problem analysis and isolation
- Channel recovery
- Error checking and correction for processor storage
- The ability to vary channels offline in single channel increments
- Concurrent channel maintenance
- Enhanced power subsystem with concurrent power maintenance for some power supplies

9672 CPCs add the following RAS advantages:

- Dynamic I/O configuration capability which allows you to dynamically add, delete or modify the I/O configuration definitions for channel paths, control units, and I/O devices without a power-on reset (POR) or initial program load (IPL)
- Subsystem storage protection and subspace group facility support for enhanced CICS availability

9672 R2 and R3 models and 9674 C02 and C03 models add the following new additional availability advantages:

- Dual Power Feeds
- Local Uninterruptible Power Supply (also available for all earlier 9672 and 9674 models)
- Redundant Power Thermal Subsystem
- CP Restart
- SAP Reassignment
- Enhanced Error Correction Code (ECC)
- Concurrent Licensed Internal Code (LIC) patch for: Hardware Management Console, ESCON channels, parallel channels, Coupling Facility channels (CFC), support element, power control, and to some extent OSA-1 and OSA-2

For more information, see 2.6, “Reliability, Availability, and Serviceability” on page 2-5.

1.3.18 Online Information

Online information is available on the system to define tasks and to aid in completing tasks.

The following information is available under the Books icon on the Hardware Management Console:

- Application Programming Interface (API) Book
- Coupling Facility (CF) books
- Hardware Management Console (HMC) Guide

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

1.3.19 Softcopy Publications

All customer publications are available on a CD-ROM disc. To view the CD-ROM disc, you need to use IBM BookManager* READ licensed programs.

1.4 Programming Compatibility

This information applies to 9672 CPCs running in basic mode and to operation in a logical partition.

Any program written for ESA/390 mode can operate on 9672 CPCs operating in ESA/390 mode, provided that the program:

- Is not time-dependent.
- Does not depend on the presence of system facilities (such as storage capacity, I/O equipment, or optional features) when the facilities are not included in the configuration.
- Does not depend on the absence of system facilities when the facilities are included in the configuration.
- Does not depend on results or functions that are defined as unpredictable or model dependent in the *IBM Enterprise Systems Architecture/390: Principles of Operation*.
- Does not depend on results or functions that are defined in this publication (or, for logically partitioned operation, in the *ES/9000, ES/3090 Processor Resource/Systems Manager Planning Guide*, GA22-7123) as being differences or deviations from the appropriate *Principles of Operation* publication.
- Does not depend on the contents of instruction parameter fields B and C on interception of the SIE instruction. See *IBM System/370 Extended Architecture Interpretive Execution* for additional information.

Any problem state program written for System/370 operates in ESA/390 mode and any problem state or control program written for 370-XA or ESA/370 operates in ESA/390 mode, provided that in each case the program:

- Observes the limitations in the preceding statements.
- Does not depend on any programming support facilities that are not provided or that have been modified.
- Takes into account other changes made that affect compatibility between modes. These changes are described in the *IBM Enterprise Systems Architecture/390: Principles of Operation*.

1.5 Engineering Change (EC) Level Considerations

| Some enhancements available for 9672 E/P/R1/R2/R3 model CPCs and 9674 C01/C02/C03 model CPCs
| are dependent on the EC change level of the CPC and/or Hardware Management Console. Additionally,
| some enhancements are further dependent on the Microcode Load (MCL) level of the EC on the CPC
| and/or Hardware Management Console. The required MCL level is available to the IBM field represen-
| tative.

1.5.1 9672 Models E, P, and R1 EC Level Considerations

Following is a list of ECs for 9672 E/P/R1 model CPCs.

1.5.1.1 EC D79759

| This EC adds support for the S/390 OSA Support Facility for MVS/ESA and support for OSA-1.

EC D79759 includes the following central processor complex (CPC) and Hardware Management Console Licensed Internal Code engineering change (LIC EC) levels:

- CPC EC level - D79756
- Hardware Management Console EC level - D79524

1.5.1.2 EC D57262

| This EC adds power save state support, coupling facility enhancements, and coupling facility support for 9672 R1 model Parallel Enterprise Servers.

EC D57262 includes the following central processor complex (CPC) and Hardware Management Console Licensed Internal Code engineering change levels:

- CPC EC level - D79533
- Hardware Management Console EC level - D79524

1.5.2 9674 Model C01 EC Level Considerations

Following is a list of ECs for 9674 C01 model CPCs.

1.5.2.1 EC D57264

| This EC adds power save state support, and coupling facility enhancements.

EC D57264 includes the following central processor complex (CPC) and Hardware Management Console Licensed Internal Code engineering change (LIC EC) levels:

- CPC EC level - D79533
- Hardware Management Console EC level - D79524

1.5.3 9672 Model R3 EC Level Considerations

Following is a list of ECs for 9672 R3 model CPCs.

1.5.3.1 EC E12923

| This EC is the base level of the 9672 R3 models.

EC E12923 includes the following central processor complex (CPC) and Hardware Management Console Licensed Internal Code engineering change (LIC EC) levels:

- CPC EC level - E12981
- Hardware Management Console EC level - E12933

| 1.5.3.2 EC E45748

| This EC adds support for OSA-1 and OSA-2, availability enhancements, and channel enhancements on the R3 models.

| EC E45748 includes the following central processor complex (CPC) and Hardware Management Console Licensed Internal Code engineering change (LIC EC) levels:

- CPC EC level - E45568
- Hardware Management Console EC level - E45552

| 1.5.3.3 EC E45748B
| (with CPC EC E45568 + MCL, and Hardware Management Console EC E45552 + MCL)

| This EC adds support for OSA-2 Asynchronous Transfer Mode (ATM) on the R3 models.

| EC E45748B includes the following central processor complex (CPC) and Hardware Management Console
| Licensed Internal Code engineering change (LIC EC) and MCL levels:

| - CPC level - EC E45568 + MCL
| - Hardware Management Console level - EC E45552 + MCL

1.5.4 9674 Model C03 EC Level Considerations

Following is a list of ECs for 9674 C03 model CPCs.

1.5.4.1 EC E12971

| This EC is the base level of the 9674 C03 models.

EC E12971 includes the following central processor complex (CPC) and Hardware Management Console Licensed Internal Code engineering change (LIC EC) levels:

- CPC EC level - E12981
- Hardware Management Console EC level - E12933

1.5.4.2 EC E45748

This EC adds availability enhancements for the C03 models.

EC E45748 includes the following central processor complex (CPC) and Hardware Management Console Licensed Internal Code engineering change (LIC EC) levels:

- CPC EC level - E45568
- Hardware Management Console EC level - E45552

1.5.5 9672 Model R2 EC Level Considerations

Following is a list of ECs for 9672 R2 model CPCs.

1.5.5.1 EC E45748

This EC is the base level of the 9672 R2 models.

EC E45748 includes the following central processor complex (CPC) and Hardware Management Console Licensed Internal Code engineering change (LIC EC) levels:

- CPC EC level - E45568
- Hardware Management Console EC level - E45552

1.5.5.2 EC E45748B
(with CPC EC E45568 + MCL, and Hardware Management Console EC E45552 + MCL)

This EC adds support for OSA-2 Asynchronous Transfer Mode (ATM) on the R2 models.

EC E45748B includes the following central processor complex (CPC) and Hardware Management Console Licensed Internal Code engineering change (LIC EC) and MCL levels:

- CPC level - EC E45568 + MCL
- Hardware Management Console level - EC E45552 + MCL

1.5.6 9674 Model C02 EC Level Considerations

Following is a list of ECs for 9674 C02 model CPCs.

1.5.6.1 EC E45748

This EC is the base level of the 9674 C02 models.

EC E45748 includes the following central processor complex (CPC) and Hardware Management Console Licensed Internal Code engineering change (LIC EC) levels:

- CPC EC level - E45568
- Hardware Management Console EC level - E45552

1.5.7 Verifying EC and MCL Levels

To verify that these ECs and MCLs are on your system, display the LIC EC levels and MCL levels currently running on the CPC and the Hardware Management Console.

You can display the LIC EC and MCL level of a 9672 or 9674 CPC using the procedure in 1.5.8, “Displaying the CPC EC and MCL Level” on page 1-28.

You can display the LIC EC and MCL level of a 9672 or 9674 Hardware Management Console using the procedure in 1.5.9, “Displaying the Hardware Management Console EC and MCL Level.”

1.5.8 Displaying the CPC EC and MCL Level

You can display the LIC EC and MCL level of a CPC as follows:

1. Select the CPC (or CPCs)
2. Drag and drop the CPC (or CPCs) onto the System Information object in the Change Management Tasks list.

   The CPC EC level and (CPC) MCL level (Activated Level) displays.

1.5.9 Displaying the Hardware Management Console EC and MCL Level

You can display the LIC EC and MCL level of a Hardware Management Console as follows:

1. Open (double-click on) the Console Actions object
2. Open (double-click on) the View Console Information object

   The Hardware Management Console EC level and (Hardware Management Console) MCL level (Activated Level) displays.

1.6 Programming Support

1.6.1 9672 Parallel Transaction Server

The 9672 Parallel Transaction Server (E0n models and P0n models) supports control programs as follows.

Note: E0n models, where n equals the number of CPCs in the system, includes models E01, E02, E03, E04, E05, and E06. P0n models, where n equals the number of CPCs in the system, includes models P01, P02, and P03.

1.6.1.1 Control Programs for Basic Modes

- MVS/ESA SP Version 5 Release 2.2
- MVS/ESA SP Version 5 Release 2
- MVS/ESA SP Version 5 Release 1

1.6.1.2 Control Programs for Logically Partitioned (LPAR) Mode

- MVS/ESA SP Version 5 Release 2.2
- MVS/ESA SP Version 5 Release 2
- MVS/ESA SP Version 5 Release 1

1.6.1.3 Control Programs for CPCs Using CPC Level Pricing

CPC level software pricing is available for CPCs that are part of a 9672 Parallel Transaction Server and provides the same control program support available for the 9672 Parallel Enterprise Server R1/R2/R3 Models. See your IBM representative for more information.

1.6.2 9672 Parallel Enterprise Server R1/R2/R3 Models

The 9672 Parallel Enterprise Server R1/R2/R3 Models support control programs as follows.

1.6.2.1 For 9672 R1 models

| **Note:** R1 models include R11, R21, R31, R41, R51, and R61.

Control Programs for Basic Modes

- **MVS**
  - MVS/ESA SP Version 5 Release 2.2
  - MVS/ESA SP Version 5 Release 2
  - MVS/ESA SP Version 5 Release 1
  - MVS/ESA SP Version 4 Release 3
  - MVS/ESA SP Version 4 Release 2.2
  - MVS/ESA SP Version 4 Release 2
  - MVS/SP* Version 3 Release 1.3
  - MVS/SP Version 2 Release 2.0
- **VM**
  - VM/ESA Version 2 Release 1.0
  - VM/ESA Version 1 Release 2.2
  - VM/ESA Version 1 Release 2.1

- **VSE**
  - VSE/ESA* Version 1 Release 3
  - VSE/ESA Version 1 Release 2

  Note: VSE/ESA can run on the R11 in basic mode.
- **Transaction Processing Facility (TPF)**
  - TPF Version 4 Release 1
  - TPF Version 3 Release 1

Control Programs for Logically Partitioned Operation

- **MVS**
  - MVS/ESA SP Version 5 Release 2.2
  - MVS/ESA SP Version 5 Release 2
  - MVS/ESA SP Version 5 Release 1
  - MVS/ESA SP Version 4 Release 3
  - MVS/ESA SP Version 4 Release 2.2
  - MVS/ESA SP Version 4 Release 2
  - MVS/SP Version 3 Release 1.3
  - MVS/SP Version 2 Release 2.0
- **VM**
  - VM/ESA Version 2 Release 1.0
  - VM/ESA Version 1 Release 2.2
  - VM/ESA Version 1 Release 2.1
  - VM/ESA Version 1 Release 1.5 (S/370 mode only)
- **VSE**
  - VSE/ESA Version 1 Release 3
  - VSE/ESA Version 1 Release 2

  Note: VSE/ESA can run on all 9672 Parallel Enterprise Server R1/R2/R3 Models in LPAR mode.
- **Transaction Processing Facility (TPF)**
  - TPF Version 4 Release 1
  - TPF Version 3 Release 1

1.6.2.2 For 9672 R2 and R3 models

| **Note:** R2 models include RA2, R12, R22, R32, R42, R52, and R72. (RA2 has 1 CP.) R3 models
| include R53, R63, R73, R83, and RX3. (RX3 has 10 integrated CP's.)

Control Programs for Basic Modes

- **MVS**
  - OS/390 Version 1 Release 1
  - MVS/ESA SP Version 5 Release 2.2
  - MVS/ESA SP Version 5 Release 2
  - MVS/ESA SP Version 5 Release 1
  - MVS/ESA SP Version 4 Release 3
  - MVS/ESA SP Version 4 Release 2.2
  - MVS/ESA SP Version 4 Release 2
  - MVS/SP Version 3 Release 1.3

  - MVS/SP Version 2 Release 2.0

- **VM**
  - VM/ESA Version 2 Release 1.0
  - VM/ESA Version 1 Release 2.2
  - VM/ESA Version 1 Release 2.1
- **VSE**
  - VSE/ESA Version 2 Release 1
  - VSE/ESA Version 1 Releases 2 and 3
- **Transaction Processing Facility (TPF)**
  - TPF Version 4 Release 1
  - TPF Version 3 Release 1

Control Programs for Logically Partitioned Operation

- **MVS**
  - OS/390 Version 1 Release 1
  - MVS/ESA SP Version 5 Release 2.2
  - MVS/ESA SP Version 5 Release 2
  - MVS/ESA SP Version 5 Release 1
  - MVS/ESA SP Version 4 Release 3
  - MVS/ESA SP Version 4 Release 2.2
  - MVS/ESA SP Version 4 Release 2
  - MVS/SP Version 3 Release 1.3
  - MVS/SP Version 2 Release 2.0

  Note: MVS/ESA Version 5 Release 1 or higher is required on the 9672 R2 and R3 Models attached to the Coupling Facility, together with the proper levels of subsystems that exploit the Coupling Facility.

- **VM**
  - VM/ESA Version 2 Release 1.0
  - VM/ESA Version 1 Release 2.2
  - VM/ESA Version 1 Release 2.1
  - VM/ESA Version 1 Release 1.5 (S/370 mode only)
- **VSE**
  - VSE/ESA Version 2 Release 1
  - VSE/ESA Version 1 Releases 2 and 3
- **Transaction Processing Facility (TPF)**
  - TPF Version 4 Release 1
  - TPF Version 3 Release 1

1.6.3 9674 Coupling Facility C01/C02/C03 Models

The 9674 Coupling Facility C01/C02/C03 Models can only run in LPAR mode with coupling facility control code. No control programs are supported.

Chapter 2. General Product Characteristics

2.1 Overview

This chapter describes the following characteristics for the 9672 Parallel Transaction Server, 9672 Parallel Enterprise Server R1/R2/R3 Models, and 9674 Coupling Facility C01/C02/C03 Models:

- Machine Requirements
  - Standard requirements
  - Optional features
- I/O operations
- Storage operations
- Data representation
- Security
- Reliability, availability, and serviceability (RAS) considerations

2.2 Machine Requirements

Machine requirements fall into two categories:

- Standard requirements
- Optional features

2.2.1 Standard Requirements

The standard machine requirements are:

- One 9672 or 9674 model
- One Hardware Management Console
- One support element (SE) per CPC
- Two 8228 Token Ring Multistation Access Units
- 128MB, 256MB, or 512MB storage minimum, depending on the model
- 9672 R3 models and 9674 C03 models require an I/O expansion cage
- ESCON or parallel channels (9672 models only)
- Coupling facility channels
- Sysplex Timer attachment per CPC (9672 Parallel Transaction Server models and 9672 Parallel Enterprise Server R1/R2/R3 Models with coupling facility channels installed)

2.2.2 Optional Features

Optional features for expanding the size, function, or performance of the system are ordered separately and are added to a CPC.

The following optional features are available for the 9672 models and the 9674 models:

- Additional central processors (9672 Parallel Transaction Server and 9674 Coupling Facility C01/C02/C03 Models only)
- Additional storage increments
- ESCON channels (9672 models only)
- Parallel channels (9672 models only)
- Additional coupling facility channels
- Additional Hardware Management Consoles
- Battery backup feature (available on 9672 E, P, and R1 models and on 9674 model C01 only)
- Local Uninterruptible Power Supply (LUPS)

- Sysplex Timer attachment (9672 Parallel Enterprise Server R1/R2/R3 Models without coupling facility channels)
- Open Systems Adapter (OSA) channels (9672 models only)

2.3 Input/Output Operations

I/O operations are handled by the channel subsystem in the Central Processor Complex (CPC). You can configure ESCON and parallel channels for block and byte multiplexer mode of operation.

You can remove failing channels from the operating configuration. As many as eight control units can be physically attached to a parallel channel, and each channel can address as many as 256 I/O devices.

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

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

2.3.1.1 Channel-to-Channel Connection

ESCON Channels (Attached to a 9034) and Parallel Channels: Parallel channel-to-channel connection between multiple systems is accomplished by using the IBM 3088 Multisystem Channel Communication Unit (MCCU) Model A1, 1, or 2. For channel-to-channel communication between an ESCON channel and parallel channels, the ESCON channel can be extended by fiber optic cable to a 9034K. that is connected to a 3088.

Parallel channel-to-channel connection between a 9672 CPC and other IBM Central Processor Complexs is accomplished by using the channel-to-channel adapter (CTCA) feature on those CPCs that offer it, or by using the 3088.

Both data-streaming and interlock modes are standard on the 3088. Data-streaming mode provides for data transfers of as many as 4.5 million bytes per second, independent of cable length. Cable distances of 122 meters (400 feet) between the CPC and the 3088 are supported in both data-streaming and interlock modes.

The 3088 Model A1 provides two-CPC connectivity and as many as 63 logical CTCA links. The 3088 Model A1 can be field upgraded to a 3088 Model 1 or Model 2. The 3088 Model 1 can interconnect as many as four channels and can provide the equivalent function of as many as 126 CTCAs. The 3088 Model 2 can interconnect as many as eight channels.

ESCON CTC Support: Parallel channel-to-channel (CTC) communication makes use of the parallel I/O CTC architecture. The parallel I/O CTC architecture defines two operating modes for CTC communication: basic mode and extended mode. ESCON CTC support for both of these modes is available. By replacing parallel CTC connections with ESCON CTC connections, you can remove the 3088 hardware requirement. See 7.6.1.2, “ESCON Channels (Using Link-Level and Device-Level Protocols).”

ESCON Channels (Using Link-Level and Device-Level Protocols): You can achieve ESCON channel-to-channel connections between a 9672 CPC and other IBM CPCs with ESCON channels if one of the ESCON channels is defined to operate in channel-to-channel (CTC) mode.

9672 CPCs support ESCON CTC connections that support CTC applications operating in parallel CTC basic mode. This can assist in the migration of parallel CTC applications to an ESCON environment and remove the 3088 hardware requirement.

ESCON channels that operate in CTC mode (extended mode or basic mode) can be defined as shared ESCON channels. For more information, see Chapter 8, “ESCON Multiple Image Facility (EMIF)” on page 8-1.

For detailed information about the ESCON channel-to-channel adapter, see *IBM Enterprise Systems Architecture/390: ESCON Channel-to-Channel Adapter* and *ES/9000 Input/Output Configuration Program User's Guide ESCON Channel-to-Channel Reference*, GC38-0401.

2.4 Storage Operations

Central storage provides storage capacity for the 9672 CPCs and the 9674 Coupling Facility C01/C02/C03 Models. Central storage is shared by all central processors. Expanded storage is physically a section of central storage and includes any installed storage not defined as central storage.

Expanded storage is controlled by the control program and transfers 4KB pages to and from central storage. The control program can use expanded storage to reduce the paging and swapping load to channel-attached paging devices in a storage-constrained environment and a heavy-paging environment.

In ESA/390 mode, storage addressing is extended from 24 bits to 31 bits, and allows for an addressing range up to 2GB. The different models, however, provide a processor storage of 128MB to 4GB. A maximum of 2GB can be used for central storage. Since the processor storage can be configured as central and expanded storage (with increments of 16MB), processor storage above 2GB may be used as expanded storage. In addition, this mode permits the use of either 24-bit or 31-bit addressing, under program control, and permits existing application programs to run with existing control programs.

In ESA/390 mode, an additional channel command word (CCW) format is provided to permit direct addressing of storage of more than 16MB for I/O operations. With this format, channel programs can also reside in storage of more than 16MB.

2.5 Data Representation

The basic addressable data unit is an 8-bit byte that can be used as one character, 2 decimal digits, or 8 binary bits. 9672 CPCs and the 9674 Coupling Facility C01/C02/C03 Models provide the following data representation features:

- Efficient use of storage and effective I/O rates for decimal data
- Variable-length fields
- Broad and flexible code conversion
- Decimal arithmetic
- Fixed-point and floating-point arithmetic
- 32-bit words, 64-bit doublewords, and 128-bit extended words (for floating-point arithmetic)
- Instructions for functions such as translate, edit, convert, move, and compare

2.6 Reliability, Availability, and Serviceability

9672 and 9674 CPCs reduce downtime by using standard features that provide high levels of reliability, availability, and serviceability (RAS).

2.6.1 Reliability

The standard features that provide a high level of reliability include:

- Subsystem storage protection and subspace group facility support, for use with CICS/ESA*, prevents application software from overwriting CICS system software, control blocks, and address spaces
- Storage background scrubbing provides continuous monitoring of storage for the correction of detected faults before the storage is used.

2.6.2 Availability

The standard features on all models that provide a high level of availability include:

- An enhanced power subsystem, which supports concurrent or deferred maintenance on the power supplies for your system and provides an additional backup power supply in situations where a power supply fails.
- Concurrent power subsystem, which allows concurrent replacement of most power supplies without turning off system power.
- Concurrent channel maintenance, which allows concurrent replacement of a channel card by varying offline a channel on a channel card without having to take down the system or the channel group.
- Fault-tolerant memory array, which enables the hardware to detect and correct most failures in central or expanded storage chips and, with power on reset, select a backup chip located on the same memory card.
- Dynamic I/O configuration which enhances system availability by supporting the dynamic addition, removal, or modification of channel paths, control units, I/O devices, and I/O configuration definitions to both hardware and software without requiring a planned outage.
- Capability to vary offline channels in single channel increments
- Customer problem analysis effects recovery without a service call.

Additional standard features on 9672 R2 and R3 models and on 9674 C02 and C03 models that provide a high level of availability include:

- Dual Power Feeds

  The power supply of the S/390 Parallel Enterprise Server R2 and R3 models has been extensively redesigned and offers dual primary power feeds. Each feed is electrically isolated and enables redundant power paths to each server. Customers may elect to provide a dual electrical service to the server, further minimizing any outage due to a single path power interruption.

- Local Uninterruptible Power Supply

  The Local Uninterruptible Power Supply, Machine Type 9910, can be installed as a supplemental or as an alternative to a central UPS to secure system availability.

- Redundant Power Thermal Subsystem

  The 9672 R2 and R3 Models' AC and DC power is designed with N+1 redundancy. Failure of a power thermal component does not cause a system outage. Concurrent replacement of the failed component results in an avoidance of a planned outage.

- CP Restart

  In LPAR mode, if a CP is checkstopped due to an error, self test will be automatically invoked. If the test is successful, the CP will be dynamically varied back on-line to the pool of shared CPs, potentially eliminating an outage to replace the failed part.

- System Assist Processor (SAP) Reassignment

  If the System Assist Processor (SAP) fails, a functioning CP is stopped and dynamically reassigned as a SAP to maintain continuity of I/O processing. This capability eliminates an unplanned outage and permits the repair to be deferred to a more convenient time.

- Concurrent Channel Upgrade

  It is possible to concurrently add ESCON, Parallel, Coupling Facility (CF), and OSA channels to an I/O cage provided an IBB is installed and there are unused channel positions and a channel driver card or (CF) channel adapter card that controls the unused channel positions in the I/O cage. This capability may help eliminate an outage to upgrade the channel configuration as is currently available on 9672 E, P, and R1 models.

- Enhanced Error Correction Code (ECC)

  Memory error checking and correction code is enhanced to detect and correct not only single bit errors, but also four bit errors in a single memory chip. This enhancement further reduces the possibility of an outage due to a memory chip failure.

- Spare Memory

  Memory cards are equipped with spare memory chips. During normal operations, the system monitors and records accumulation of failing bits in memory chips that are corrected by Error Correction Code (ECC). Before a failure threshold is reached that could result in an uncorrectable error, during a planned power-on reset, the system invokes a spare memory chip in place of the one with accumulated failing bits. This action may prevent an unscheduled outage on a need for replacement of the memory card.

- Concurrent Licensed Internal Code (LIC) Patch

  In addition to Licensed Internal Code (LIC) patches for the Hardware Management Console, patches for ESCON channels, parallel channels, coupling links, support element, power control, and to some extent OSA-1 and OSA-2, can now be applied concurrently. This significantly reduces the planned outage necessary for LIC maintenance.

2.6.3 Serviceability

The standard features that provide a high level of serviceability include:

- Automatic fault isolation (analysis routines) concurrent with operation
- Automatic remote support capability
- On-site problem isolation
  - Field-replaceable unit (FRU) isolation
  - Trace tables
  - Error logout recording

Chapter 3. Hardware Description

3.1 9672 and 9674 Frame Configuration

A minimum configuration 9672 model (9672 Model E01, 9672 Model P01, or a 9672 Parallel Enterprise Server R1/R2/R3 Models) and the 9674 Coupling Facility C01/C02/C03 Models consists of a single frame containing a single Central Processor Complex (CPC) and its associated central processor (CP), storage, channel, and power components. The frame is a rack enclosure built to Electronic Industry Association (EIA) standards. It is divided into two cages with the CPC occupying the top cage. The bottom cage can accommodate the following depending on the model:

- An optional expansion cage for additional channel attachment capability for the CPC installed in the top cage (9672 Model E01, 9672 R-models, or 9674 Model C01)
- An expansion cage for coupling facility channel attachment 9674 C03 models.
- A tailgate for standard bus and tag parallel channel attachment for the CPC installed in the top cage (9672 Model P01 or 9672 Model R11)
- An additional CPC (9672 Model E01 upgrade only)
- An OSA-1 module (9672 model E01, 9672 R1 and R2 models)

The 9674 Coupling Facility C01/C02/C03 Models are only available as a single CPC in a single frame configuration. Each 9672 R-model is only available as a single CPC in a one, two, three, or four frame configuration. However, the larger 9672 Parallel Transaction Server E-Models and P-Models are multi-frame, multi-CPC configurations called strings that can have up to six frames attached together. Each CPC in the string has its own distinct CP, storage, channel, and power supply resources but shares through internal cabling:

- Local area network (LAN) attachment to the Hardware Management Console
- Sysplex Timer attachment
- An emergency power off (EPO) switch

LAN attachment (through one of two 8228 Multistation Access Units) and the EPO switch are located in the first frame of the string. See 3.2.4, “8228 Token Ring Network Multistation Access Unit” on page 3-8.

Sysplex Timer attachment is distributed through the first two CPCs in the string. See 3.2.3.15, “Sysplex Timer Adapter Card” on page 3-7.

3.2 Central Processor Complex (CPC) Cage

The physical components of an individual Central Processor Complex (CPC) cage in a 9672 Parallel Transaction Server string, in 9672 Parallel Enterprise Server R1/R2/R3 Models, or in 9674 Coupling Facility C01/C02/C03 Models consist of the following:

- Support element (SE)
- Base CPC board
- Two 8228 Multi-Station Access Units (9674 CPC or first CPC in a 9672 string only)
- Optional battery backup feature (available on 9672 E, P, and R1 models and on 9674 model C01)

3.2.1 Support Element (SE)

The support element (SE) is directly attached to the CPC and provides the interface between the Hardware Management Console and the CPC. This distributes hardware operations and management controls to the Hardware Management Console for the CPC and allows for independent and parallel operational control of a CPC from the Hardware Management Console. Netview applications (for example, Target System Control Facility [TSCF] and Automated Operations Control [AOC]) can also send commands to the SE.

During normal operation, the SE monitors CPC operation, collecting and sending CPC status, hardware messages, and operating system messages to the Hardware Management Console for consolidation and exception processing. The SE can also send SNA alerts and operating system messages to Netview applications.

The SE contains the:

- Licensed Internal Code for the CPC
- Hardware system definitions for the CPC (contained in the reset, image, and load profiles for the CPC)
- Battery-powered clock used to set the CPC time-of-day (TOD) clock at power-on reset. CPC attachment to a Sysplex Timer causes the CPC TOD clock to take its time setting from the Sysplex Timer.
- A LAN port to connect the CPC to the 8228 Multistation Access Unit
- On the 9672 E, P, and R1 models, and on the 9674 model C01, connections for the service cart display, keyboard, and mouse so that local service for the CPC can be carried out without use of the Hardware Management Console.
- 2.88MB diskette drive (for use during service operations)

3.2.2 Base Central Processor Complex (CPC) Board

The base Central Processor Complex (CPC) board consists of specialized cards, or modules and cards, that provide central processor, storage, channel, voltage and power control resources, and Sysplex Timer attachment capability. You can add additional channel capability to the CPC by installing an optional I/O expansion cage in the bottom of the same frame in 9672 E, R1 and R2 models, and in 9674 model C01. An I/O expansion cage is required for 9672 R3 models and 9674 C03 models. See 3.4, “I/O Expansion Cages” on page 3-9.

3.2.3 System Cards

3.2.3.1 Central Processor Card (for 9672 E/P/R1/R2 and 9674 C01/C02)

The central processor card contains the CPs for the CPC. Each CP has its own execution unit, cache storage, control storage, and floating point processor. The CP card also has the processor oscillator and the memory bus adapter (MBA).

3.2.3.2 Central Processor Module (for 9672 R3 and 9674 C03)

The central processor module contains the CPs for the CPC. Each CP has its own execution unit, cache storage, control storage, and floating point processor. The processor oscillator is located on the oscillator card. The memory bus adapter (MBA) is located on the CPC board.

3.2.3.3 Storage Cards

Processor storage, even though physically the same, can be configured as both central and expanded storage for each CPC. A maximum of 2GB storage can be configured as central storage. The 9672 model RA2 requires only one 128MB storage card. 9672 models E, P, R1, R2, and 9674 models C01 and C02 require two equal capacity storage cards. 9672 models R3 and 9674 model C03 require four equal capacity storage cards.

Processor Storage

- 9672 E, P, R1, and 9674 C01 Models
  - Minimum - 128MB; Maximum - 2GB
  - Increments - 128MB to 256MB, 256MB to 512MB, 512MB to 1GB, 1GB to 2GB
- 9672 RA2 Model (1 memory card)
  - Minimum - 128MB; Maximum - 1GB
  - Increments - 128MB to 256MB, 256MB to 512MB, 512MB to 1GB
- 9672 R2 Models (except RA2) and 9674 C02 Model (2 memory cards)
  - Minimum - 256MB; Maximum - 2GB
  - Increments - 256MB to 512MB, 512MB to 1GB, 1GB to 2GB
- 9672 R3 Models and 9674 C03 Model (4 memory cards)
  - Minimum - 512MB; Maximum - 4GB
  - Increments - 512MB to 1GB, 1GB to 2GB, 2GB to 4GB

All storage upgrades require the removal of existing storage cards and installation of new storage cards. All parts removed or replaced are the property of IBM and must be returned to IBM.

3.2.3.4 Voltage Regulator Cards (9672 E, P, R1, and 9674 C01 models)

Voltage regulator cards consist of a basic type that receives 24V DC bulk voltage as input and generates +6.0V, +5.1V, -5.1V, and +3.6V DC (at +/-3%) as output.

3.2.3.5 Distributed Converter Assembly (DCA) Cards (9672 R2, R3, and 9674 C02, C03 models)

DC to DC converter cards in the CPC or expansion cage (FC 2020 only) that converts –350 Volts DC to logic voltages.

3.2.3.6 Internal Bus Buffer Card

The IBB card provides the bus from the channels (ESCON, parallel, and coupling facility) to the memory bus adapter (MBA). An IBB cable is used to connect the IBB card to the processor card.

IBB domains are groups of I/O slots within the CPC (9672 E, P, R1, and R2 models, and 9674 C01 and C02 models only) and the expansion cage(s) that are associated with specific IBB cards and cables.

The 9672 E, P, R1, RA2, R53 models and the 9674 model C01 support two IBBs. The 9672 R2 (except RA2) and R3 (except R53) models support 4 IBB's. The 9674 C02 and C03 models use 2 IBB's.

There are two types of IBB cards:

- A single wide card is available for all models which can drive Coupling Facility adapters and links as well as any of the other adapter cards. This configuration supports 8 I/O slots.
- A double wide card is available for 9672 R3 models **ONLY**. The double wide card does **NOT** support Coupling Facility adapters and links, but does support the other adapters. This configuration supports 16 I/O slots.

The only difference in the plugging of the double is that there is only one IBB cable and the card is plugged into two adjacent card slots.

3.2.3.7 Channel Driver Cards

The number of channel driver cards depend on the model and the number of channels installed. The number of channel driver cards are:

- For 9672 E, P, and R1 models: 1 to 4
- For 9672 R2 and R53 models: 1 to 8
- For 9672 R3 models (except R53): 1 to 12

Channel driver cards can be plugged in the base CPC board or I/O expansion cage for 9672 E, P, R1, and R2 models. Channel driver cards must be plugged in an I/O expansion cage for 9672 R3 models.

Channel driver cards are required for all channel cards except coupling facility channel adapter cards.

3.2.3.8 ESCON Channel Cards

Channel driver cards support up to four ESCON channel cards. An ESCON channel card supports:

- 1 to 3 ESCON channels on 9672 E, P, and R1 models
- 1 to 3 ESCON channels on 9672 R2 and R3 models in an I/O cage upgraded from a 9672 E, P, or R1 model (base CPC board or expansion cage feature code 1000)
- 1 to 3 ESCON channels on 9672 R3 models prior to EC E45748
- 1 to 4 ESCON channels on 9672 R2 and R3 models at EC E45748 or later

3.2.3.9 Parallel Channel Cards

Channel driver cards support up to four parallel channel cards. A parallel channel card supports up to three parallel channels on all 9672 models.

3.2.3.10 Open Systems Adapter 1 (OSA-1) Channel Extender Cards

All 9672 models support up to 3 OSA-1 channel extender cards. All OSA-1 channel extender cards must be installed in the same cage (base CPC board or I/O expansion cage). Each OSA-1 channel extender card installed reduces by 3 or 4 the total number of ESCON and parallel channels available. OSA-1 channels are available on 9672 E/P/R1 models with EC D79759. OSA-1 channels are available on 9672 R2 and R3 models at EC E45748 or later.

3.2.3.11 Open Systems Adapter 1 (OSA-1) Channel Card

Each OSA-1 channel extender card can accommodate up to three OSA-1 channel cards, which are plugged in the OSA-1 module. Each OSA-1 channel card supports one OSA channel. OSA-1 channels are available on 9672 E/P/R1 models with EC D79759. OSA-1 channels are available on 9672 R2 and R3 models at EC E45748 or later.

3.2.3.12 Open Systems Adapter 2 (OSA-2) Channel Cards

9672 R2 and R3 models can accommodate up to 12 Open Systems Adapter 2 (OSA-2) channel cards. An OSA-2 channel card supports one OSA channel. Each OSA-2 channel card installed reduces by 3 or 4 the total number of ESCON and parallel channels available. OSA-2 channels are available only on 9672 R2 and R3 models at EC E45748 or later.

3.2.3.13 Coupling Facility Channel Adapter Cards

Coupling facility channel adapter cards can be installed as follows:

- 1 to 4 per base CPC board on 9672 E/P/R1 models and 9674 model C01.
- 1 to 4 in I/O expansion cage (feature code 1000) on 9672 E/P/R1 models and 9674 model C01.
- 1 to 8 in I/O expansion cage (feature code 2000) on 9672 E and P models.
- 1 to 16 in I/O expansion cage (feature code 2000) on 9674 model C01.
- 1 to 2 per I/O expansion cage (feature code 1000) on 9672 R2 (except RA2) and R3 models in a cage upgraded from a base CPC board or I/O expansion cage from a 9672 E, P, or R1 model.
- 1 to 6 in base CPC on 9672 RA2 models with a maximum of 6 coupling facility channel adapter cards for the CPC.
- 1 to 6 per cage (base CPC board or I/O expansion cage feature code 2020) on 9672 R2 (except RA2) models with a maximum of 6 coupling facility channel adapter cards for the CPC.
- 1 to 6 per I/O expansion cage (feature code 2020) on 9672 R3 models, with a maximum of 6 coupling facility channel adapter cards for the CPC.
- 1 to 16 in base CPC board on 9674 C02 models with a maximum of 16 coupling facility channel adapter cards for the CPC.
- 1 to 16 in I/O expansion cage (feature code 2020) on 9674 C03 models with a maximum of 16 coupling facility channel adapter cards for the CPC.

For 9672 E, P, and R1 models, the third and fourth coupling facility channel adapter cards installed each reduce by 3 the total number of ESCON and parallel channels available. For 9672 R2 and R3 models, each coupling facility channel adapter card installed reduces by 3 or 4 the total number of ESCON and parallel channels available.

3.2.3.14 Coupling Facility Channel Cards

The coupling facility channel adapter cards can accommodate up to 2 coupling facility channel cards. Each coupling facility channel card supports one coupling facility channel.

3.2.3.15 Sysplex Timer Adapter Card

A Sysplex Timer adapter card is required for all CPCs that are part of a 9672 Parallel Transaction Server E–model and 9672 Parallel Enterprise Server R1/R2/R3 Models with coupling facility channels.

For 9672 Model E01, 9672 Model P01, or 9672 Parallel Enterprise Server R1/R2/R3 Models, a dual port Sysplex Timer adapter card feature code #6152 provides two ports for high availability fiber attachment to an external 9037 Sysplex Timer.

For 9672 Parallel Transaction Server E–models with strings of 2 or more CPCs, a standard single port Sysplex Timer adapter card feature code #6150 provides one fiber attachment to an external 9037 Sysplex Timer for each of the first 2 CPCs in the string. This card distributes the Sysplex Timer attachment to the other CPCs in the string. The 9037 Sysplex Timer can synchronize the time-of-day clocks for systems attached to it. Installation of two Sysplex Timer adapter cards feature code #6150, one in each of

the first 2 CPCs in the string, provides two ports for high availability fiber attachment to an external 9037 Sysplex Timer.

For the remaining CPCs in the 9672 Parallel Transaction Server E–model string (CPCs 3 through 8), a Sysplex Timer adapter card feature code #6151 provides attachment, through internal copper cabling, to the primary Sysplex Timer adapter feature code #6150 cards in the first two CPCs. The Sysplex Timer adapter card on CPCs 3 through 8 is cross-wired to each primary Sysplex Timer adapter card in the first two CPCs providing a redundant path to the 9037 Sysplex Timer in the event of a scheduled or unscheduled outage of either of the first two CPCs in the 9672 Parallel Transaction Server string.

Important: If both of the first two CPCs in a 9672 Parallel Transaction Server string are powered-off, Sysplex Timer attachment is lost for the entire string.

| For 9672 R2 and R3 models at EC E45748 or later, ETR External cable feature code #6153 can be used in conjunction with feature code #6150, to enable the sharing of sysplex timer ports, providing high availability.

3.2.3.16 Universal Processor Controller (UPC) Card

The universal power controller (UPC) card consists of a small processor with control store, EEPROM, and oscillator for driving the power control logic (power-on and power-off sequencing) and the blower and thermal supervision logic. Additionally, the UPC card contains interfaces to the memory bus adapter, the frame control panel, and the support element (SE).

3.2.4 8228 Token Ring Network Multistation Access Unit

The 8228 is a Token Ring Network Multistation Access Unit that allows the attachment of up to eight devices into a Token Ring local area network (LAN). It provides the required LAN attachment between the Hardware Management Console and all support elements (SEs) contained in a 9672 or 9674. There are two 8228s installed above the top cage of the first frame of each 9672 model. You attach an SE to an 8228 through cables plugged directly into the 8228 from the SE. You can also attach an 8228 to other CPCs outside a 9672 string to form larger networks.

3.2.5 Optional Battery Backup Feature

3.2.5.1 For 9672 E, P, and R1 models, and the 9674 model C01

The optional battery backup feature #2011 uses three internal batteries to provide a backup power source for short term utility power outages. It can be used to provide power as follows:

- 3.5 minutes at full CPC operation
- 80 minutes of storage-only preservation if a coupling facility logical partition running on the CPC is operating in power save mode. See 6.4, “Coupling Facility Nonvolatility” on page 6-3.

Battery backup support for these support durations is dependent on the batteries being fully charged. You can extend the duration of support by attaching an external 24-volt battery. EXIDE Electronics** Auxiliary DC Power MODule (ADMOD**) can extend the battery backup support duration to up to 48 hours.

Concurrent maintenance of a battery backup battery is possible if two of the three batteries are still operational.

3.3 Local Uninterruptible Power Supply (LUPS)

3.3.1.1 For all 9672 and 9674 models

To reduce any power outages, a Local Uninterruptible Power Supply (LUPS), machine type 9910, is available to be ordered with the server. This unit provides 9 to 18 minutes of full power hold-up in the event of extended customer power line disturbances.

3.4 I/O Expansion Cages

| You can install optional I/O expansion cage(s) to provide additional channel attachment for a Central
| Processor Complex (CPC), depending on the model. An optional expansion cage allows you to add more channels up to the amount supported by a particular I/O expansion cage and the CPC.

An I/O expansion cage is required for 9672 R3 models and for 9674 C03 models.

There are three types of I/O expansion cages:

| • **Feature code 1000** – For 9672 E, P, and R1 models at EC D57262 (except Model R11), 9674 model
| C01, 9672 R2 models (except Model RA2) and R3 models **ONLY** when upgrading from R1 models up to a maximum of two.
  - Allows you to add up to 24 more ESCON or parallel channels.
  - Allows you to add up to 8 more coupling facility channels on 9672 E, P, and R1 models and 9674 model C01. Allows you to add up to 4 more coupling facility channels on 9672 R2 and R3 models. For 9672 E, P, and R1 models, the third and fourth coupling facility channel adapter cards installed each reduce by 3 the total number of ESCON and parallel channels available. For 9672 R2 and R3 models, each coupling facility channel adapter card installed reduces by 3 the total number of ESCON and parallel channels available.
  - Allows you to add up to 9 OSA-1 channels on all 9672 models. (3 OSA-1 channel extender cards in the expansion cage each supporting up to 3 OSA-1 channel cards in the OSA-1 module). Each OSA-1 channel extender card installed reduces by 3 the total number of ESCON and parallel channels available. If OSA-1 channel extenders are already installed in the CPC, you ***cannot*** install additional OSA-1 channel extenders in the I/O expansion cage. OSA-1 channels require that you also install the OSA-1 module that provides the tailgates necessary for direct LAN attachment to the CPC.
- **Feature code 2000** – For 9672 E and P models only, and 9674 model C01 only
  - For 9672 CPCs (E-Models and P-Models only), allows you to add up to 24 more ESCON or parallel channels.
  - For 9672 CPCs (E-Models and P-Models only), allows you to add 16 additional coupling facility channels. For 9674 CPCs (Model C01 only), allows you to add 24 additional coupling facility channels.
| • **Feature code 2020** – For 9672 R2 models (except model RA2) and R3 models only and 9674 model C03 only

  The following bullets refer to channel cards. To determine the number of channels available with each type of channel card, see 3.2.3, “System Cards” on page 3-4.
  - For 9672 R2 models, allows you to add up to 16 channel cards.
  - | For 9672 R3 models, allows you to add up to 16 channel cards prior to EC E45748. With EC
    | E45748 or later, allows you to add up to 22 channel cards.
  - For 9672 R2 and R3 models only, allows you to install up to 6 coupling facility channel adapter cards. Each coupling facility channel adapter card installed reduces by 3 or 4 the total number of

ESCON and parallel channels available. For 9674 model C03 only, allows you to install up to 16 coupling facility channel adapter cards.

- For 9672 R2 and R3 models only, allows you to install up to 9 OSA-1 channels. (3 OSA-1 channel extender cards in the expansion cage each supporting up to 3 OSA-1 channel cards in the OSA-1 module). Each OSA-1 channel extender card installed reduces by 3 or 4 the total number of ESCON and parallel channels available. If OSA-1 channel extenders are already installed in the CPC, you ***cannot*** install additional OSA-1 channels extenders in the I/O expansion cage.
- For 9672 R2 and R3 models only, allows you to install up to 12 OSA-2 channel cards. Each OSA-2 channel card installed reduces by 3 or 4 the total number of ESCON and parallel channels available.

| **Note:** Until EC E45748, the feature code 2020 expansion cage supports two single IBB domains with 8 I/O slots each, or a double IBB domain on R3 models. In the double configuration, the IBB cable is connected to a double wide IBB card. The double wide card supports 16 I/O slots. The double wide card does **not** support coupling facility adapters, but does support the other channel
| cards. With EC E45748 or later, an additional single IBB domain supporting 6 more I/O slots is available. The 2020 expansion cage can support a mix of this single and double IBB domain.

For a listing of valid channel configurations, see Table 3-3 on page 3-21 (9672 CPCs–E-Models and P-Models), Table 3-4 on page 3-25 (9672 CPCs–9672 R-Models), and Table 3-6 on page 3-30 (9674 CPCs).

3.5 Optional Open Systems Adapter (OSA)

The optional Open Systems Adapter (OSA) function provides for direct LAN attachment to the CPC.

| The OSA function appears to the application software as a unique typr of S/390 channel with associated
| I/O devices and OSA ports that allow the attachment of FDDI, Ethernet, and Token Ring LANs, plus
| ATM-based networks.

3.5.1.1 Open Systems Adapter 1 – OSA-1

| OSA-1 features #5000 (available 6/95 for 9672 E, P, and R1 models) and #5100 (available with EC
| E45748 or later for 9672 R2 and R3 models) are modules which consist of an OSA-1 cage, a power cage located behind the OSA-1 cage, and a tailgate that is located below the OSA-1 cage.

OSA-1 channel extender cards, located on the base CPC board or in the I/O expansion cage, connect to the OSA-1 module through internal cables.

Note: An I/O expansion cage is required for 9672 R3 models.

OSA-1 channel feature code #5003 for FDDI LANs is a single LAN port offering. OSA-1 channel feature codes #5002 and #5005, for Token Ring and Ethernet respectively, each provide 5 LAN ports.

You can install a maximum of one optional OSA-1 module for each 9672 CPC. For 9672 Parallel Transaction Server models E02 through E08 and P02 through P03, you can install a maximum of 2 optional OSA-1 modules.

OSA-1 module feature #5000 is available for 9672 E, P, and R1 models. It is available on 9672 R2 and
| R3 models **only** when carried over by an upgrade. OSA-1 module feature #5100 is available with EC
| E45748 or later for 9672 R2 and R3 models. Both features allow a maximum of 9 OSA-1 channel cards per cage.

Note: OSA-1 module feature #5100 is **not** compatible with I/O expansion cage feature #1000.

3.5.1.2 Open Systems Adapter 2 – OSA-2

| This feature is available only on 9672 R2 and R3 models with EC E45748 or later.

The OSA-2 features are I/O cards which are plugged directly into any channel card slot except into a feature code 1000 (from an upgrade) becoming integral components of the multiframe system, enabling convenient LAN attachment. In 9672 R2 models, the cards can be plugged in the CPC or in an optional I/O expansion cage. In 9672 R3 models, the cards can only be plugged in the feature code #2020 I/O expansion cage.

| There are three OSA-2 features:

- One is a single ENTR (Ethernet/Token Ring) feature card (#5201) that can be defined as Ethernet or Token Ring when the connector is attached. It contains two ports which can be configured as Ethernet or Token Ring (4 or 16 Mbits/second).
- The second is a FDDI feature card (#5202) that supports both dual ring and single ring attachments.
- The third, available only on 9672 R2 and R3 models at EC E45748B (with CPC EC E45568 + MCL, and Hardware Management Console EC E45552 + MCL) or later, are ATM feature cards #5204 (155 MB Multimode) and #5205 (155 MB Single Mode) that support Asynchronous Transfer Mode (ATM)-based networks.

| OSA-2 feature codes #5201 and #5202 are available for 9672 R2 and R3 models only. OSA-2 feature codes #5204 and #5205 are available only on 9672 R2 and R3 models at EC E45748B (with CPC EC E45568 + MCL, and Hardware Management Console EC E45552 + MCL) or later. A maximum of 12 OSA-2 channel cards are allowed per system.

OSA-2 feature cards can not be plugged in the OSA-1 module feature codes #5000 or #5100, or in I/O expansion cage feature code #1000.

3.5.1.3 OSA-1 and OSA-2 Differentiation

| The size of the OSA-2 package is smaller than the OSA-1 package, allowing the OSA-2 feature to be plugged into any I/O slot in a CEC or I/O expansion cage, except feature code #1000. This smaller package supports the networking protocols, TCP/IP for MVS and VM and SNA/APPN for MVS (ACF/VTAM for MVS/ESA).

| OSA-1 supports FDDI, Ethernet, and Token Ring LANs. OSA-2 supports FDDI, Ethernet, and Token Ring LANs, and ATM-based networks.

OSA-1 and OSA-2 FDDI features support one dual-ring or single-ring attachment, as well as attachment to an optical bypass switch.

OSA-2 has an ENTR (Ethernet/Token Ring) feature that has two independent ports that can be configured as two Ethernets, two Token Rings, or one Ethernet and one Token Ring, providing maximum flexibility. OSA-1 has additional port capacity. The Ethernet and Token Rings features are two-card features with five ports each.

| The OSA-2 ATM feature, available only on 9672 R2 and R3 models at EC E45748B (with CPC EC E45568 + MCL, and Hardware Management Console EC E45552 + MCL) or later, supports ATM-based networks.

Note: Only OSA-1 is available on 9672 E, P, and R1 models. OSA-1 and OSA-2 can co-exist on 9672 R2 and R3 models.

3.5.1.4 Software Requirements for OSA-1 and OSA-2 Features

For software requirements for OSA-1 and OSA-2, see 1.3.8.2, “Support for the TCP/IP, SNA, and IPX Network Protocols” on page 1-14.

3.6 Hardware Management Console

The Hardware Management Console is attached by local area network (LAN) to the support element (SE) in a 9672 or 9674 Central Processor Complex (CPC) through the 8228 Multistation Access Unit. By using the SE connection, the Hardware Management Console can control and monitor status for multiple CPCs providing a single point of control and single system image.

The Hardware Management Console is a PC running OS/2, Communications Manager/2, the distributed console access facility (DCAF), and the Hardware Management Console Application (HMCA). It also contains configuration information about its own configuration and about support elements (SEs) defined to it.

The Hardware Management Console hardware configuration also includes:

- Choice of a 17" or 21" display monitor
- Keyboard
- Mouse
- Parallel printer port
- 2.88MB disk drive
- 300+MB fixed disk
- 128MB R/W optical drive
- SDLC communication adapter and associated modem
- 3270 interface adapter (coaxial cable)
- Serial communication port

3.7 CPU ID

CPU ID information is system-generated for each Central Processor Complex (CPC) and consists of a version code for the model number field, a CPU identification number that is unique for each logical processor, a model number for the machine type, and four hexadecimal digits of zeros.

The Store CPU ID (STIDP) instruction stores the CPU ID for each logical processor in storage in the following format (Figure 3-1):

Figure 3-1. CPU ID Format

The **version code** is associated with the number of central processors in the CPC and is not affected by the operating mode (basic or LPAR) of the CPC.

For the basic modes, the **CPU identification number** consists of the CPU address of the central processor and digits derived from the serial number of the CPC.

For LPAR mode, the **CPU identification number** for each logical processor consists of a one-digit logical CPU address, a one-digit logical partition identifier, and digits derived from the serial number of the CPC.

Figure 3-2 on page 3-13 and Figure 3-3 on page 3-13 show the format of the CPU identification number (bits 8 through 31 of the CPU ID format) for the basic modes and LPAR mode.

Figure 3-2. CPU ID Format for LPAR Mode and Basic Mode 9672 CPC

Figure 3-3. CPU ID Format for LPAR Mode 9674 CPC

Where:

VV is the version code (2 hexadecimal digits).

For **9672 CPCs (R1 models at EC D57262, E-Models and P-Models)**, VV indicates the following:

- 32 (1-way CPC in a 9672 Model R11)
- 33 (2-way 9672 CPC)
- 34 (3-way 9672 CPC)
- 35 (4-way 9672 CPC)
- 36 (5-way 9672 CPC)
- 37 (6-way 9672 CPC)

For **9672 CPCs (R1 models)**,

- 42 (1-way 9672 CPC)
- 43 (2-way 9672 CPC)
- 44 (3-way 9672 CPC)
- 45 (4-way 9672 CPC)
- 46 (5-way 9672 CPC)
- 47 (6-way 9672 CPC)

For **9674 CPCs (model C01)**,

- 22 (1-way 9674 CPC)
- 23 (2-way 9674 CPC)
- 24 (3-way 9674 CPC)
- 25 (4-way 9674 CPC)
- 26 (5-way 9674 CPC)
- 27 (6-way 9674 CPC)

For **9672 CPCs (R2 models) and 9674 CPCs (model C02)**,

- 1F (1-way CPC in 9672 Model RA2 and 9674 Model C02 sub-uni)
- 12 (1-way CPC in 9672 Model R12 and 9674 Model C02)
- 13 (2-way CPC in 9672 Model R22 and 9674 Model C02)
- 14 (3-way CPC in 9672 Model R32 and 9674 Model C02)
- 15 (4-way CPC in 9672 Model R42 and 9674 Model C02)
- 16 (5-way CPC in 9672 Model R52 and 9674 Model C02)
- 18 (7-way CPC in 9672 Model R72 and 9674 Model C02)

For **9672 CPCs (R3 models) and 9674 CPCs (model C03)**,

- 06 (5-way CPC in 9672 Model R53 and 9674 Model C03)
- 08 (6-way CPC in 9672 Model R63 and 9674 Model C03)
- 09 (7-way CPC in 9672 Model R73 and 9674 Model C03)
- 0A (8-way CPC in 9672 Model R83 and 9674 Model C03)
- 0C (10-way CPC in 9672 Model RX3 and 9674 Model C03)

Annnnn is the (six hexadecimal digits) central processor identification number

LPnnnn is the (six hexadecimal digits) central processor identification number for LPAR mode

- The first digit (A in basic mode, L in LPAR mode) is the CPU address, as stored by the Store CPU Address (STAP) instruction (see 5.4.54.1, “CPU Address Identification” on page 5-20).
- The next 5 digits (nnnnn) are, for the basic modes, selected from the serial number of the CPC. In LPAR mode, the logical partition identifier replaces the first of the 5 digits (Pnnnn).

9672 *or* **9674** is the CPC machine type.

0000 is the reserved hexadecimal digit.

3.8 Storage

Each 9672 Central Processor Complex (CPC) and each 9674 CPC has its own storage resources. CPC storage can consist of both central and expanded storage.

3.8.1 Central Storage

Central storage provides:

- Data storage and retrieval for the central processors (CPs)
- Communication with the CPs
- Communication with and control of optional expanded storage
- Error checking and correction (ECC)

The amount of central storage for a CPC can range from 64MB to 2GB. Part of central storage is used as the hardware system area.

3.8.2 Hardware System Area (HSA)

The hardware system area (HSA) contains the CPC Licensed Internal Code (LIC) and configuration information used by the CPC during the current power-on reset of the machine. HSA is not available for program use.

The HSA size varies according to the:

- Power-on reset (POR) mode of the CPC
- Size of the system I/O configuration
- Whether or not dynamic I/O configuration is enabled

3.8.2.1 Basic Mode HSA Allocation for 9672 CPCs

For 9672 CPCs, the following basic mode system I/O configurations require the given HSA size:

- **11.2MB** is allocated to support the following small system I/O configuration:
  - 4096 HSA subchannels
  - 24 ESCON/parallel channels defined in the IOCDS
  - ***Without*** dynamic I/O configuration enabled
- **15.0MB** is allocated to support the following large system I/O configuration:
  - 18 432 HSA subchannels
  - 192 ESCON/parallel channels defined in the IOCDS
  - With dynamic I/O configuration enabled
- Additional HSA is allocated for OSA channels as follows:
  - **1.8MB** if OSA channels are supported on the model.

Intermediate HSA sizes between the given small and large system I/O configurations do exist.

3.8.2.2 LPAR Mode HSA Allocation for 9672 CPCs

For 9672 CPCs, the following LPAR mode system I/O configurations require the given HSA size:

- **18.6MB** is allocated to support the following small system I/O configuration:
  - 4 Logical Partitions
  - 4096 HSA subchannels (2048 shared)
  - 24 ESCON/parallel channels defined in the IOCDS
  - ***Without*** dynamic I/O configuration enabled
- **29.2MB** is allocated to support the following large system I/O configuration:
  - 10 Logical Partitions
  - 18 432 HSA subchannels
  - 192 ESCON/parallel channels defined in the IOCDS
  - With dynamic I/O configuration enabled
- Additional HSA is allocated for OSA channels as follows:
  - **1.8MB** if OSA channels are supported on the model.

Intermediate HSA sizes between the given small and large system I/O configurations do exist.

Note: Engineering change (EC) level D79759 on 9672 E/P/R1 models provides support for OSA-1 channels. Support for OSA-1 and OSA-2 channels on 9672 R2 and R3 models is available with EC E45748 or later. See 1.5, “Engineering Change (EC) Level Considerations” on page 1-25.

HSA Size Notes:

1. When counting ESCON and parallel channels without dynamic I/O enabled, count the number of ***defined*** channels, not the number of installed channels. With dynamic I/O enabled, HSA is allocated assuming maximum channels.
2. To determine the number of HSA subchannels, refer to the IZP IOCP IODEVICE report under the **HSA Total** column for Subchannels. If you use HCD, you can find the total number of HSA subchannels in the Device Detail report.

   If dynamic I/O configuration is enabled, apply the expansion factors you specified on the Dynamic I/O Options panel in the current Reset profile.

Some amount of installed central storage is reserved for the Hardware System Area (HSA) to hold processor licensed internal code. When you initially defined the amount of central storage to allocate to your partitions (in the Image profiles), it was recommended that you assume that the maximum amount (42MB) was required for the HSA.

3.8.3 Error Checking and Correction

Error checking and correction (ECC) code bits are stored with data in central storage.

For 9672 E, P, and R1 models, and 9674 model C01, single-bit and double-bit errors are detected and logged; single-bit errors are corrected.

For 9672 R2 and R3 models and 9674 C02 and C03 models, memory error checking and correction code is enhanced to detect and correct not only single bit errors, but also four bit errors in a single memory chip.

Unrecoverable errors are flagged for follow-on action.

Data paths from the central processors and the channels are checked for parity. Parity bits are included in each command or data word.

3.8.4 Key-Controlled Storage Protection

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

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

3.8.5 Expanded Storage

For 9672 and 9674 Central Processor Complexs (CPCs), you have the option of defining a portion of CPC as expanded storage. Expanded storage is high-speed storage that can reduce the paging and swapping load of I/O channels.

Expanded storage is configured at power-on reset. Storage above 64MB can be allocated as expanded storage in 16MB increments. The maximum expanded storage is equal to the installed processor storage minus the configured central storage. At most, for the 9672 R3 models, this is 4GB minus 64 MB.

The control program initiates the movement of data between central storage and expanded storage. No data can be transferred to expanded storage without passing through central storage.

Coupling Facility Logical Partition Considerations: You should ***not*** define expanded storage for coupling facility logical partitions (LPs) defined on 9674 CPCs and 9672 CPCs.

3.8.5.1 Error Checking and Correction

For 9672 E, P, and R1 models, and 9674 model C01, error checking and correction (ECC) code bits in expanded storage permit the following:

- Single-bit and double-bit error detection
- Single-bit error correction

For 9672 R2 and R3 models and 9674 C02 and C03 models, memory error checking and correction detects and corrects not only single bit errors, but also four bit errors in a single memory chip.

Unrecoverable errors are flagged.

3.9 Channel Subsystem (CSS)

Each 9672 Central Processor Complex (CPC) and each 9674 CPC has its own channel subsystem. The channel subsystem (CSS) operates in ESA/390 mode only and supports ESCON, parallel, coupling facility, and OSA channels. (9674 Coupling Facility C01/C02/C03 Models CPCs only support coupling facility channels.)

The CSS handles all I/O operations for the central processors (CPs). The CSS controls communication between a configured channel and the control unit and device. The I/O configuration data set (IOCDS) that is selected when the system is initialized defines the channel paths on the CPC, the control units attached to the channel paths, and the I/O devices assigned to the control units. The IOCDS is created using the input/output configuration program (IOCP) and it is stored on the support element hard disk associated with a CPC. At system initialization, the IOCDS information is used to build the necessary control blocks in the hardware system area (HSA) of central storage. HSA is not available for program use.

ESCON and Parallel channel considerations:

- The CSS for 9672 E, P, and R1 models contain one or two internal bus buffers (IBBs). Each IBB supports one or two channel driver cards which, in turn, support from one to four channel cards. Each channel card supports either three ESCON or three parallel channels.
- The CSS for 9672 model RA2 contains one or two internal bus buffers (IBBs). For ESCON and parallel channels, each IBB supports one or two channel driver cards which, in turn, support from one to four channel cards. Each channel card supports either three parallel channels, or four ESCON channels.
- The CSS for 9672 model R53 contains one or two internal bus buffers (IBBs). Each IBB supports one to four channel driver cards which, in turn, support from one to four channel cards. Each channel card
| supports either three parallel channels, or three ESCON channels prior to EC E45748, or four ESCON
| channels with EC E45748 or later.
- The CSS for 9672 R2 models contain one to four internal bus buffers (IBBs). Each IBB supports one or two channel driver cards which, in turn, support from one to four channel cards. Each channel card supports either three parallel channels, or four ESCON channels.
- The CSS for 9672 R3 models (except R53) contain one to four internal bus buffers (IBBs). Up to two of the IBB's can support one to four channel driver cards which, in turn, support from one to four channel cards; up to four of the IBB's can support one to two channel driver cards which, in turn, support from one to four channel cards. Each channel card supports either three parallel channels, or
| three ESCON channels prior to EC E45748, or four ESCON channels with EC E45748 or later.

Table 3-1. Max No. of ESCON and Parallel Channels -- New Builds				
Model	**ESCON LED**	**Parallel**	**Total**	**Notes**
9672 E, P, R1	48	48	48	
9672 RA2	64	48	64	
9672 R53	128	96	128	1
9672 R2 (except RA2)	128	96	128	
9672 R3 (except R53)	192	96	192	2

Notes:

| 1. Prior to EC E45748, max. no. of channels is 96. With EC E45748 or later, max. no. of channels is 128.

| 2. Prior to EC E45748, max. no. of channels is 144. With EC E45748 or later, max. no. of channels is 192.

Coupling Facility channel considerations:

- For 9672 E and P models, a maximum of 24 coupling facility channels are supported.
- For 9672 R1 models, a maximum of 16 coupling facility channels are supported.
- For 9672 R2 and R3 models, a maximum of 12 coupling facility channels are supported
- For all 9674 models, a maximum of 32 coupling facility channels are supported.

Note: When plugging Coupling Facility channel adapter cards into the CPC or I/O expansion cage (feature code 2020) in 9672 R3 models, they cannot be plugged in domains controlled by a double IBB.

OSA channel considerations:

- For 9672 E, P, and R1 models, a maximum of 9 OSA-1 channels are supported.
- For 9672 R2 and R3 models, a maximum of 9 OSA-1 channels and 12 OSA-2 channels are supported.

9672 CPCs can support a mix of ESCON, parallel, OSA, and coupling facility channels. See 3.10, “Hardware Configuration for the 9672 Parallel Transaction Server” on page 3-20 for a summary of the valid channel configurations available for a 9672 E-Models and P-Models. See 3.11, “Hardware Configuration for the 9672 Parallel Enterprise Server R1/R2/R3 Models” on page 3-24 for a summary of the valid channels configurations available for a 9672 R1 model.

Coupling facility adapter cards can be pre-installed allowing the concurrent plugging of coupling facility channel cards. See 4.6.3, “Coupling Facility Continuous Availability” on page 4-16.

Table 3-2. Max No. of OSA-1, OSA-2, Coupling Facility Channels -- New Builds			
Model	**OSA-1 Channels**	**OSA-2 Channels**	**Coupling Facility Channels**
9672 E/P	9	N/A	24
9672 R1	9	N/A	16
9672 R2/R3	9	12	12
9674 All	N/A	N/A	32

3.9.1.1 Block Multiplexer Mode of Operation

ESCON and parallel channels can be configured for block multiplexer mode of operation. In block multiplexer mode of operation, parallel channels and ESCON channels configured as CVC channel paths can operate in either interlock (high-speed transfer) mode or in data-streaming mode. CVC channels require a 9034 ESCON converter. They can also be attached to control units that operate in high-speed transfer or in data-streaming mode. Data rates can be as high 4.5MB per second for parallel channels and for ESCON CVC channel paths.

3.9.1.2 Byte Multiplexer Mode of Operation

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

3.10 Hardware Configuration for the 9672 Parallel Transaction Server

The following describes the hardware configuration for the 9672 E-Models and P-Models.

3.10.1 Common Hardware Elements

Each 9672 E-Model and 9672 P-Model provides the following hardware elements:

- One or more Central Processor Complexs (CPCs)
- Two to six integrated CPs per CPC
- 128MB minimum of storage, with a maximum of 2 GB per CPC.

  Valid storage configurations are 128MB, 256MB, 512MB, 1GB, and 2GB per CPC.
- Channels per CPC:
  - ESCON channels, parallel channels, or a mixture of both:
    three minimum, 48 maximum available in increments of 3.
  - Coupling Facility Channels:
    One minimum, 16 maximum available in increments of 1
    (24 maximum available at EC D57262)
  - Open systems adapter 1 (OSA-1) channels:
    9 maximum available in increments of 1.

  See Table 3-3 on page 3-21.
- Up to ten logical partitions (LPs) per CPC are supported.
- One Hardware Management Console (up to three additional optional Hardware Management Consoles are available).
- One support element (SE) per CPC.
- One service cart for portable attachment to any SE for service activity. The service cart includes a cart, display, keyboard, and cabling.
- Two 8228 Token-Ring Network Multistation Access Units (MAUs)
- Optional battery backup feature per CPC.
- Optional Local Uninterruptible Power Supply (LUPS)
- Sysplex Timer attachment per CPC. Models E01 and P01 optionally support two Sysplex Timer attachments.

Table 3-3 (Page 1 of 2). 9672 Parallel Transaction Server Valid Channel Configurations per CPC

CPC Channel Features	ESCON or Parallel Channels (or both)	OSA Channels	Coupling Facility Channels
Without optional expansion cage	3–24 3–21 3–18 3–15	0 1–3 4–6 7–9	1–4 1–4 1–4 1–4
	3–21 3–18 3–15 3–12	0 1–3 4–6 7–9	5–6 5–6 5–6 5–6
	3–18 3–15 3–12 3–9	0 1–3 4–6 7–9	7–8 7–8 7–8 7–8
With optional expansion cage (Feature code 1000. See note 3.)	21–48 18–45 18–42 18–39	0 1–3 4–6 7–9	1–8 1–8 1–8 1–8
	3–45 3–42 3–39 3–36	0 1–3 4–6 7–9	9–10 9–10 9–10 9–10
	3–42 3–39 3–36 3–33	0 1–3 4–6 7–9	11–12 11–12 11–12 11–12
	3–39 3–36 3–33 3–30	0 1–3 4–6 7–9	13–14 13–14 13–14 13–14
	3–36 3–33 3–30 3–27	0 1–3 4–6 7–9	15–16 15–16 15–16 15–16
With optional expansion cage (Feature code 2000. See note 3.)	27–48 24–45 21–42 3–39 3–36 3–33 3–30 3–27 3–24 3–21 3–18	NA NA NA NA NA NA NA NA NA NA NA	1–4 5–6 7–8 9–10 11–12 13–14 15–16 17–18 19–20 21–22 23–24

Table 3-3 (Page 2 of 2). 9672 Parallel Transaction Server Valid Channel Configurations per CPC

CPC Channel Features	ESCON or Parallel Channels (or both)	OSA Channels	Coupling Facility Channels

Notes:

1. For 9672 E-Models, parallel channel attachment is through a 78 pin D-shell connector cable.
2. For 9672 P-Models, parallel channel attachment can be either through standard bus and tag cable or through a 78 pin D-shell connector cable. Standard bus and tag cables on the 9672 P models require tailgate hardware.
3. For more information, see 3.4, "I/O Expansion Cages" on page 3-9.
4. For OSA-1 channels, feature #5000 an OSA-1 expansion module is required. See 3.5, "Optional Open Systems Adapter (OSA)" on page 3-10.

3.10.2 9672 Model E01

The 9672 Model E01 consists of a single Central Processor Complex (CPC) housed in one or two frames; the CPC has from two to six integrated central processors (CPs). The CPs have access to a common central storage, expanded storage, and channels. See 3.10.1, "Common Hardware Elements" on page 3-20 for a more complete description of the hardware configuration.

3.10.3 9672 Model E02

The 9672 Model E02 consists of two Central Processor Complexs (CPCs), housed in one, two, or three frames; each CPC has from two to six integrated central processors (CPs). The CPs for each CPC have access to a common central storage, expanded storage, and channels. See 3.10.1, "Common Hardware Elements" on page 3-20 for a more complete description of the hardware configuration.

3.10.4 9672 Model E03

The 9672 Model E03 consists of three Central Processor Complexs (CPCs), housed in two, three, or four frames; each CPC has from two to six integrated central processors (CPs). The CPs for each CPC have access to a common central storage, expanded storage, and channels. See 3.10.1, "Common Hardware Elements" on page 3-20 for a more complete description of the hardware configuration.

3.10.5 9672 Model E04

The 9672 Model E04 consists of four Central Processor Complexs (CPCs), housed in two, three, four, or five frames; each CPC has from two to six integrated central processors (CPs). The CPs for each CPC have access to a common central storage, expanded storage, and channels. See 3.10.1, "Common Hardware Elements" on page 3-20 for a more complete description of the hardware configuration.

3.10.6 9672 Model E05

The 9672 Model E05 consists of five Central Processor Complexs (CPCs), housed in three, four, five, or six frames; each CPC has from two to six integrated central processors (CPs). The CPs for each CPC have access to a common central storage, expanded storage, and channels. See 3.10.1, "Common Hardware Elements" on page 3-20 for a more complete description of the hardware configuration.

3.10.7 9672 Model E06

The 9672 Model E06 consists of six Central Processor Complexs (CPCs), housed in three, four, five, six, or seven frames; each CPC has from two to six integrated central processors (CPs). The CPs for each CPC have access to a common central storage, expanded storage, and channels. See 3.10.1, “Common Hardware Elements” on page 3-20 for a more complete description of the hardware configuration.

3.10.8 9672 Model E07

The 9672 Model E07 consists of seven Central Processor Complexs (CPCs), housed in four, five, six, or seven frames; each CPC has from two to six integrated central processors (CPs). The CPs for each CPC have access to a common central storage, expanded storage, and channels. Only five CPCs can add an optional expansion cage. See 3.10.1, “Common Hardware Elements” on page 3-20 for a more complete description of the hardware configuration.

3.10.9 9672 Model E08

The 9672 Model E08 consists of eight Central Processor Complexs (CPCs), housed in four, five, six, or seven frames; each CPC has from two to six integrated central processors (CPs). The CPs for each CPC have access to a common central storage, expanded storage, and channels. Only four CPCs can add an optional expansion cage. See 3.10.1, “Common Hardware Elements” on page 3-20 for a more complete description of the hardware configuration.

3.10.10 9672 Model P01

The 9672 Model P01 consists of single Central Processor Complex (CPC) housed in one, two, or three frames; the CPC has from two to six integrated central processors (CPs). The CPs have access to a common central storage, expanded storage, and channels. See 3.10.1, “Common Hardware Elements” on page 3-20 for a more complete description of the hardware configuration.

3.10.11 9672 Model P02

9672 Model P02 consists of two Central Processor Complexs (CPCs) housed in two, three, four, five, or six frames; each CPC has from two to six integrated central processors (CPs). The CPs for each CPC have access to a common central storage, expanded storage, and channels. See 3.10.1, “Common Hardware Elements” on page 3-20 for a more complete description of the hardware configuration.

3.10.12 9672 Model P03

9672 Model P03 consists of three Central Processor Complexs (CPCs) housed in three, four, five, six, or seven frames; each CPC has from two to six integrated central processors (CPs). The CPs for each CPC have access to a common central storage, expanded storage, and channels. See 3.10.1, “Common Hardware Elements” on page 3-20 for a more complete description of the hardware configuration.

3.11 Hardware Configuration for the 9672 Parallel Enterprise Server R1/R2/R3 Models

The following describes the hardware configuration for the 9672 9672 R-Models:

3.11.1 Common Hardware Elements

Each 9672 R-model provides the following hardware elements:

- One Central Processor Complex (CPC)
- Integrated CPs
  - R1: 1 to 6 CPs
  - RA2: 1 CP
  - R2: 1 to 5, 7 CPs
  - R3: 5 to 8
  - RX3: 10 CPs
- Storage:
  - R1: 128MB minimum, 2GB maximum
  - RA2: 128MB minimum, 1GB maximum
  - R2 (except RA2): 256MB minimum, 2GB maximum
  - R3: 512MB minimum, 4GB maximum

  Valid storage configurations are 128MB, 256MB, 512MB, 1GB, 2GB and 4GB per CPC.
- Up to ten logical partitions (LPs) are supported.
- One Hardware Management Console (up to three additional optional Hardware Management Consoles are available).
- One support element (SE).
- For 9672 R1 models: one service cart for portable attachment to any SE for service activity. (The service cart includes a cart, display, keyboard, and cabling.)
- Two 8228 Token-Ring Network Multistation Access Units (MAUs)
- Optional battery backup feature (available on R1 models only).
- Optional Local Uninterruptible Power Supply (LUPS)
- Sysplex Timer attachment (required on models with coupling facility channels, optional on models without coupling facility channels).
- Channels:
  - ESCON channels, parallel channels, or a mixture of both: three or four minimum; available in increments of 3 or 4; 144 maximum prior to EC E45748, 192 maximum with EC E45748 or later depending on model. For 9672 R1 models, parallel channel attachment is through standard bus and tag cable or 78 pin D-shell connectors. Standard bus and tag cables on the 9672 R1 models require tailgate hardware. For 9672 R2 and R3 models, parallel channel attachment is through 78 pin D-shell connectors.
  - Optional Coupling facility channels on 9672 models R2, R3, and R1 (at EC D57262): 12 or 16 maximum available in increments of 1, depending on model.

Note: (Model R11 does ***not*** support coupling facility channels.

– Optional OSA channels:
  - 9 OSA-1 maximum available in increments of 1
  - 12 OSA-2 maximum available in increments of 1 on 9672 R2 and R3 models.

– For valid 9672 R1 models channel configurations, see Table 3-4.

– The Channel Subsystem (CSS) for 9672 R2 and R3 models may consist of numerous combinations of ESCON, parallel, OSA-1, OSA-2, and coupling facility channels, up to the maximums described in 3.9, “Channel Subsystem (CSS)” on page 3-17. Replacement of any ESCON or parallel channel card by any OSA-1 channel extender card, OSA-2 channel card, or coupling facility channel adapter card ***reduces*** by three or four (depending on the specific ESCON channel card) the total number of ESCON and parallel channels available.

– For minimum and maximum I/O expansion cages for 9672 R2 and R3 models without and with upgrade from R1 models, see Table 3-5 on page 3-26. The minimum and maximum I/O expansion cages for R1 models are zero and one respectively.

Table 3-4 (Page 1 of 2). 9672 R1 model valid Channel Configurations per CPC

CPC Channel Features	ESCON or Parallel Channels (or both)	OSA Channels	Coupling Facility Channels
R1 Models ***Without*** optional expansion cage	3–24	0	1–4
	3–21	1–3	1–4
	3–18	4–6	1–4
	3–15	7–9	1–4
	3–21	0	5–6
	3–18	1–3	5–6
	3–15	4–6	5–6
	3–12	7–9	5–6
	3–18	0	7–8
	3–15	1–3	7–8
	3–12	4–6	7–8
	3–9	7–9	7–8
R1 Models With optional expansion cage (Feature code 1000. See note 2.)	21–48	0	1–8
	18–45	1–3	1–8
	18–42	4–6	1–8
	18–39	7–9	1–8
	3–45	0	9–10
	3–42	1–3	9–10
	3–39	4–6	9–10
	3–36	7–9	9–10
	3–42	0	11–12
	3–39	1–3	11–12
	3–36	4–6	11–12
	3–33	7–9	11–12
	3–39	0	13–14
	3–36	1–3	13–14
	3–33	4–6	13–14
	3–30	7–9	13–14
	3–36	0	15–16
	3–33	1–3	15–16
	3–30	4–6	15–16
	3–27	7–9	15–16

Table 3-4 (Page 2 of 2). 9672 R1 model valid Channel Configurations per CPC			
CPC Channel Features	**ESCON or Parallel Channels (or both)**	**OSA Channels**	**Coupling Facility Channels**

Notes:

1. For 9672 R1 Models, parallel channel attachment is through standard bus and tag cable. Standard bus and tag cables on the 9672 R1 models require tailgate hardware.
2. For more information, see 3.4, “I/O Expansion Cages” on page 3-9.
3. For OSA-1 channels, feature code #5000 an OSA-1 expansion module is required. See 3.5, “Optional Open Systems Adapter (OSA)” on page 3-10.
4. Model R11 does ***not*** support coupling facility channels.

Table 3-5. Minimum and Maximum I/O Expansion Cages without and with upgrade from R1 models.

9672 R2 & R3 Models	Without Upgrade from R1	With Upgrade from R1
RA2	0	0
R12	0 – 1	1 – 2
R22	0 – 1	1 – 2
R32	0 – 1	1 – 2
R42	0 – 1	1 – 2
R52	0 – 1	1 – 2
R72	0 – 1	1 – 2
R53	1 – 2	1 – 2
R63	1 – 3	1 – 4
R73	1 – 3	1 – 4
R83	1 – 3	1 – 4
RX3	1 – 3	1 – 4

Note: When an OSA-1 module is added to an R3 model, the maximum number of feature code 2020 expansion cages is reduced to 2.

3.11.2 9672 Model R11

The 9672 Model R11 consists of one Central Processor Complex (CPC) housed in one, two, or three frames; the CPC has one integrated central processor (CP) which has access to a common central storage, expanded storage, and channels. See 3.11.1, “Common Hardware Elements” on page 3-24 for a more complete description of the hardware configuration.

3.11.3 9672 Model R21

The 9672 Model R21 consists of one Central Processor Complex (CPC) housed in one, two, or three frames; the CPC has two integrated central processors (CPs). The CPs have access to a common central storage, expanded storage, and channels. See 3.11.1, “Common Hardware Elements” on page 3-24 for a more complete description of the hardware configuration.

3.11.4 9672 Model R31

The 9672 Model R31 consists of one Central Processor Complex (CPC) housed in one, two, or three frames; the CPC has three integrated central processors (CPs). The CPs have access to a common central storage, expanded storage, and channels. See 3.11.1, “Common Hardware Elements” on page 3-24 for a more complete description of the hardware configuration.

3.11.5 9672 Model R41

The 9672 Model R41 consists of one Central Processor Complex (CPC) housed in one, two, or three frames; the CPC has four integrated central processors (CPs). The CPs have access to a common central storage, expanded storage, and channels. See 3.11.1, “Common Hardware Elements” on page 3-24 for a more complete description of the hardware configuration.

3.11.6 9672 Model R51

The 9672 Model R51 consists of one Central Processor Complex (CPC) housed in one, two, or three frames; the CPC has five integrated central processors (CPs). The CPs have access to a common central storage, expanded storage, and channels. See 3.11.1, “Common Hardware Elements” on page 3-24 for a more complete description of the hardware configuration.

3.11.7 9672 Model R61

The 9672 Model R61 consists of one Central Processor Complex (CPC) housed in one, two, or three frames; the CPC has six integrated central processors (CPs). The CPs have access to a common central storage, expanded storage, and channels. See 3.11.1, “Common Hardware Elements” on page 3-24 for a more complete description of the hardware configuration.

3.11.8 9672 Model RA2

The RA2 consists of one Central Processor Complex (CPC) housed in one frame; the CPC has one integrated central processor (CP) which has access to a common central storage, expanded storage, and channels. The RA2 has 128MB minimum storage and 1GB maximum. See 3.11.1, “Common Hardware Elements” on page 3-24 for a more complete description of the hardware configuration.

3.11.9 9672 Model R12

The R12 consists of one Central Processor Complex (CPC) housed in one, two, or three frames; the CPC has one integrated central processor (CP) which has access to a common central storage, expanded storage, and channels. The R12 has 256MB minimum storage and 2GB maximum. See 3.11.1, “Common Hardware Elements” on page 3-24 for a more complete description of the hardware configuration.

3.11.10 9672 Model R22

The R22 consists of one Central Processor Complex (CPC) housed in one, two, or three frames; the CPC has two integrated central processors (CPs) which has access to a common central storage, expanded storage, and channels. The R22 has 256MB minimum storage and 2GB maximum. See 3.11.1, “Common Hardware Elements” on page 3-24 for a more complete description of the hardware configuration.

3.11.11 9672 Model R32

The R32 consists of one Central Processor Complex (CPC) housed in one, two, or three frames; the CPC has three integrated central processors (CPs) which has access to a common central storage, expanded storage, and channels. The R32 has 256MB minimum storage and 2GB maximum. See 3.11.1, “Common Hardware Elements” on page 3-24 for a more complete description of the hardware configuration.

3.11.12 9672 Model R42

The R42 consists of one Central Processor Complex (CPC) housed in one, two, or three frames; the CPC has four integrated central processors (CPs) which has access to a common central storage, expanded storage, and channels. The R42 has 256MB minimum storage and 2GB maximum. See 3.11.1, “Common Hardware Elements” on page 3-24 for a more complete description of the hardware configuration.

3.11.13 9672 Model R52

The R52 consists of one Central Processor Complex (CPC) housed in one, two, or three frames; the CPC has five integrated central processors (CPs) which has access to a common central storage, expanded storage, and channels. The R52 has 256MB minimum storage and 2GB maximum. See 3.11.1, “Common Hardware Elements” on page 3-24 for a more complete description of the hardware configuration.

3.11.14 9672 Model R72

The R72 consists of one Central Processor Complex (CPC) housed in one, two, or three frames; the CPC has seven integrated central processors (CPs) which has access to a common central storage, expanded storage, and channels. The R72 has 256MB minimum storage and 2GB maximum. See 3.11.1, “Common Hardware Elements” on page 3-24 for a more complete description of the hardware configuration.

3.11.15 9672 Model R53

The R53 consists of one Central Processor Complex (CPC) housed in one, two, or three frames; the CPC has five integrated central processors (CPs) which has access to a common central storage, expanded storage, and channels. The R53 has 512MB minimum storage and 4GB maximum. See 3.11.1, “Common Hardware Elements” on page 3-24 for a more complete description of the hardware configuration.

3.11.16 9672 Model R63

The R63 consists of one Central Processor Complex (CPC) housed in one, two, three, or four frames; the CPC has six integrated central processors (CPs) which has access to a common central storage, expanded storage, and channels. The R63 has 512MB minimum storage and 4GB maximum. See 3.11.1, “Common Hardware Elements” on page 3-24 for a more complete description of the hardware configuration.

3.11.17 9672 Model R73

The R73 consists of one Central Processor Complex (CPC) housed in one, two, three, or four frames; the CPC has seven integrated central processors (CPs) which has access to a common central storage, expanded storage, and channels. The R73 has 512MB minimum storage and 4GB maximum. See 3.11.1, “Common Hardware Elements” on page 3-24 for a more complete description of the hardware configuration.

3.11.18 9672 Model R83

The R83 consists of one Central Processor Complex (CPC) housed in one, two, three, or four frames; the CPC has eight integrated central processors (CPs) which has access to a common central storage, expanded storage, and channels. The R83 has 512MB minimum storage and 4GB maximum. See 3.11.1, “Common Hardware Elements” on page 3-24 for a more complete description of the hardware configuration.

3.11.19 9672 Model RX3

The RX3 consists of one Central Processor Complex (CPC) housed in one, two, three, or four frames; the CPC has ten integrated central processors (CPs) which has access to a common central storage, expanded storage, and channels. The R3 has 512MB minimum storage and 4GB maximum. See 3.11.1, “Common Hardware Elements” on page 3-24 for a more complete description of the hardware configuration.

3.12 Hardware Configuration for the 9674 Coupling Facility C01/C02/C03 Models

The 9674 models are stand-alone coupling facilities and are the models recommended for use as production coupling facilities.

The following describes the hardware configuration for the 9674 Coupling Facility C01/C02/C03 Models.

3.12.1 Common Hardware Elements

Each 9674 model provides the following hardware elements:

- One Central Processor Complex (CPC)
- Integrated CPs
  - C01: 1 to 6 CPs
  - C02: 1 to 5, 7 CPs
  - C03: 5 to 8, 10 CPs
- Storage:
  - C01: 128MB minimum, 2GB maximum
  - C02: 256MB minimum, 2GB maximum
  - C03: 512MB minimum, 4GB maximum

  Valid storage configurations are 128MB, 256MB, 512MB, 1GB, 2GB, and 4GB.
- Up to ten logical partitions (LPs) are supported.
- One Hardware Management Console (up to three additional optional Hardware Management Consoles are available).
- One support element (SE).
- For 9674 C01 models: One service cart for portable attachment to any SE for service activity. (The service cart includes a cart, display, keyboard, and cabling).
- Two 8228 Token-Ring Network Multistation Access Units (MAUs)
- Optional battery backup feature (available on model C01 only).
- Optional Local Uninterruptible Power Supply (LUPS)
- Channels:
  - Coupling facility channels:
    two minimum, 32 maximum available in increments of 1

Table 3-6 (Page 1 of 2). Valid Channel Configurations per CPC

CPC Channel Features	Coupling Facility Channels
Model C01 ***Without*** optional expansion cage	2–8
Model C01 With optional expansion cage (Feature code 1000. See note.)	9–16

Table 3-6 (Page 2 of 2). Valid Channel Configurations per CPC

CPC Channel Features	Coupling Facility Channels
Model C01 With optional expansion cage (Feature code 2000. See note.)	9–32
Model C02 ***No*** expansion cage	2–32
Model C03 With required expansion cage (Feature code 2020. See note.)	2–32

Note: For more information, see 3.4, “I/O Expansion Cages” on page 3-9.

3.12.2 C01

The C01 consists of one Central Processor Complex (CPC) housed in a single frame; the CPC has from one to six integrated central processors (CPs). The CPs have access to a common central storage, expanded storage, and channels. The C01 has 128MB minimum storage and 2GB maximum. See 3.12.1, “Common Hardware Elements” on page 3-30 for a more complete description of the hardware configuration.

3.12.3 9674 Model C02

The 9674 C02 consists of one Central Processor Complex (CPC) housed in a single frame; the CPC has one, two, three, four, five, or seven integrated central processors (CPs). The CPs have access to a common central storage, expanded storage, and channels. The C02 has 256MB minimum storage and 2GB maximum. See 3.12.1, “Common Hardware Elements” on page 3-30 for a more complete description of the hardware configuration.

3.12.4 9674 Model C03

The 9674 C03 consists of one Central Processor Complex (CPC) housed in a single frame; the CPC has five, six, seven, eight, or ten integrated central processors (CPs). The CPs have access to a common central storage, expanded storage, and channels. The C03 has 512MB minimum storage and 4GB maximum. See 3.12.1, “Common Hardware Elements” on page 3-30 for a more complete description of the hardware configuration.

Chapter 4. Hardware Management Console

4.1 Capabilities

A Hardware Management Console provides a single point of control and single system image for those Central Processor Complexs (CPCs) configured to it. CPCs configured to a Hardware Management Console are those CPCs whose support elements (SEs) are:

- Attached by local area network (LAN) to the Hardware Management Console
- Defined to have the same domain name as the Hardware Management Console
- Defined in the Defined CPCs group at the Hardware Management Console

This includes any 9672 CPC or any 9674 CPC.

Figure 4-1. Hardware Management Consoles. Connected to a 9672 Parallel Transaction Server, 9672 Parallel Enterprise Server R1/R2/R3 Models, and a 9674 Coupling Facility C01/C02/C03 Models over a Private Token-Ring Local Area Network (LAN).

The SE for each CPC allows the Hardware Management Console to monitor the CPC by providing status information. Each SE provides the Hardware Management Console with operator controls for its associated CPC so that you can target operations:

- In parallel to multiple or all CPCs
- To a single CPC

The Hardware Management Console provides the following:

- Hardware Management Console Application object-oriented interface
- Customizable groups of hardware objects.
- Customizable Hardware Management Console Application Settings
- Consolidation of:
  - Operator controls
  - Hardware status reporting
  - Hardware message presentation
  - Operating system messages
  - Problem analysis and reporting
  - Licensed Internal Code (LIC) control and distribution.
- Remote I/O configuration and IOCDS management
- Scheduled operations
- Remote operations
- Automation and API support.

Note: There is an interactive tutorial (form SK2T-1196) available for the Hardware Management Console.

4.2 Hardware Management Console Application

The Hardware Management Console Application (HMCA), running on the Hardware Management Console, is an OS/2 based LIC application that provides an easy-to-use, object-oriented interface. You can directly manipulate objects displayed in the Hardware Management Console Workplace using a mouse or key combinations.

Figure 4-2. Hardware Management Console Workplace. The CPC Console task is renamed to the Single Object Operations task with engineering change (EC) D79535.

The Hardware Management Console Workplace is divided into three areas:

- Views Area

  The Views Area contains icons that represent views or collections of objects that make up your system.

- Work Area

  The Work Area contains the objects of your system, based on the view you selected.

- Tasks Area

  The Tasks Area contains objects that represent tasks you can perform on other objects.

  The Hardware Management Console Application supports the following tasks:

  For the 9672 E, P, and R1 models, and the 9674 C01 model:

  – **Daily Tasks**

    - Activate
    - Reset Normal
    - Deactivate
    - Single Object Operations

      The CPC Console task is renamed to the Single Object Operations task with engineering change (EC) D57262 for 9672 CPCs and EC D57264 for 9674 CPCs. See 1.5, “Engineering Change (EC) Level Considerations” on page 1-25.
    - Grouping
    - Activity

      Available with EC D57262 for 9672 CPCs and EC D57264 for 9674 CPCs. See 1.5, “Engineering Change (EC) Level Considerations” on page 1-25.

  – **CPC Recovery Tasks**

    - Start
    - Stop
    - Reset Normal
    - PSW Restart
    - Reset Clear

      Available with EC D57262 for 9672 CPCs and EC D57264 for 9674 CPCs. See 1.5, “Engineering Change (EC) Level Considerations” on page 1-25.

  – **Service Tasks**

    - Service Mode
    - Service Status

Available with EC D57262 for 9672 CPCs and EC D57264 for 9674 CPCs. See 1.5, “Engineering Change (EC) Level Considerations” on page 1-25.

- Perform Problem Analysis
- View Service History
- Backup Critical Data
- Hard Disk Restore
- Checkout Tests
- Report a Problem
- Transmit Service Data

– **Change Management Tasks**

- Engineering Changes (ECs)
- CPC EC Details
- Retrieve Internal Code
- Change Internal Code
- System Information

– **CPC Remote Customization Tasks**

- Remote Service
- Problem Management
- Operations Management
- Account Information

– **CPC Operational Customization Tasks**

- Activation Profiles
- Automatic Activation
- Scheduled Operations
- Customize Date/Time

Available with EC D57262 for 9672 CPCs and EC D57264 for 9674 CPCs. See 1.5, “Engineering Change (EC) Level Considerations” on page 1-25.

– **Object Definition Tasks**

- Change Object Definition
- Add Object Definition
- Remove Object Definition

– **CPC Configuration Tasks**

- Perform Model Conversion
- Transmit Vital Product Data

For the 9672 R2 and R3 models, and the 9674 C02 and C03 models:

– **Daily Tasks**

- Activate
- Reset Normal
- Deactivate
- Grouping
- Activity

– **CPC Recovery Tasks**

- Single Object Operations
- Start
- Stop

- Reset Normal
- PSW Restart
- Reset Clear
- Load

– **Service Tasks**

- Service Status
- Perform Problem Analysis
- View Service History
- Backup Critical Data
- Hard Disk Restore
- Checkout Tests
- Report a Problem
- Transmit Service Data

– **Change Management Tasks**

- Engineering Changes (ECs)
- CPC EC Details
- Retrieve Internal Code
- Change Internal Code
- Product Engineering Directed Changes
- System Information

– **CPC Remote Customization Tasks**

- Remote Service
- Problem Management
- Operations Management
- Account Information

– **CPC Operational Customization Tasks**

- Customize Activation Profiles
- View/Delete Activation Profiles
- View Activation Profiles
- Automatic Activation
- Scheduled Operations
- Customize Date/Time

– **Object Definition Tasks**

- Change Object Definition
- Add Object Definition
- Remove Object Definition

– **CPC Configuration Tasks**

- Perform Model Conversion
- Transmit Vital Product Data
- View Frame Layout
- Edit Frame Layout

Before you can use the Hardware Management Console, you must logon in one of the following user modes:

- Access Administrator
- Operator
- Advanced Operator
- System Programmer

After logging on, the initial view of the Hardware Management Console Workplace displays the objects associated with the Groups View in the Work Area.

The Hardware Management Console allows you to customize the groups of hardware objects. You use the Grouping task to define groups of hardware objects (CPCs or CPC Images, and Coupling Facility Logical Partition CPC Images). The Grouping task allows you to:

- Create and name a new group
- Add to an existing group
- Delete from an existing group

You can perform operator actions on:

- All defined system hardware objects (CPCs, CPC Images, or Coupling Facility CPC Images)
- Groups of hardware objects you define to be meaningful
- Individual hardware objects.

You can perform the following operator actions in parallel on groups of objects:

- Activate
- Activity (available with EC D79535)
- Deactivate
- Reset clear (available with EC D79535)
- System reset normal
- Start central processors
- Stop central processors
- Program status word (PSW) restart
- Retrieve internal code
- Change internal code
- Problem management
- Operations management
- Account information
- Automatic activation
- Scheduled operations
- Add object definition
- Remove object definition
- Perform model conversion

You can also use the Hardware Management Console to access the CPC console of an individual SE to perform more detailed operator actions on a single CPC (or on LPs on the CPC) when required.

4.2.1 Customization

You can use the Hardware Management Console Settings task to customize the presentation characteristics of the Hardware Management Console. These customized settings can be saved to a diskette and used on other Hardware Management Consoles if desired. The Hardware Management Console Settings task allows you to:

- Modify the default colors or use grey patterns instead of colors.
- Associate a color or pattern with any unacceptable status value you define to distinguish between types of exceptions
- Change the background color of the Views Area used to indicate exception and non-exception situations
- Modify the default color associated with pending hardware or operating system messages

- Enter the Hardware Management Console SNA and TCP/IP addresses and domain name

4.2.2 Status Reporting

Each SE monitors the operation of its associated CPC and any CPC images running on the CPC and sends status to the Hardware Management Console for consolidation and exception processing.

Exception processing surfaces only those hardware status items you define as unacceptable to the running of your systems. You can define the conditions for exception processing at the Hardware Management Console using the CPC Details panel associated with each CPC or CPC image. The Hardware Management Console uses customizable panel colors to indicate hardware status requiring your attention.

The Hardware Management Console displays hardware status by using color (or grey patterns) to indicate acceptable or unacceptable status for CPCs or CPC Images. Unacceptable status results in an exception situation that causes the color to change for the:

- Entire Views Area background
- CPC or CPC Image object background in the Work Area for the CPC or CPC Image with the unacceptable status
- Group object background in the Work Area for any group that the CPC or CPC Image with the unacceptable status is part of.

The default color change is from green (acceptable status) to red (unacceptable status). You can customize these colors (or patterns) using the Hardware Management Console Settings task.

4.2.3 Hardware Messages

The Hardware Management Console allows you to monitor the hardware messages from any CPC, CPC Image, or any group of CPCs or CPC Images configured to it.

The Hardware Management Console uses color (or grey patterns) and flashing icons to indicate when hardware messages are pending for CPCs or CPC Images.

Hardware messages present information about problems that are detected, suggest actions where possible, and aid in requesting service when appropriate. When a message has been handled, it is deleted from all the Hardware Management Consoles and SEs.

When hardware messages are pending for a hardware object or group of hardware objects:

- The background of the object and its associated group turns blue (the default)
- The Hardware Messages icon turns blue and flashes

4.2.4 Operating System Messages

Local operating systems and coupling facility control code running in a coupling facility logical partition can use the console integration facility of the hardware to send operator messages to be displayed by the Hardware Management Console. The Hardware Management Console allows you to monitor and respond to the operating system messages from any CPC Image, coupling facility, or any group of CPC Images configured to it.

The Hardware Management Console uses color (or patterns) and flashing icons to indicate when operating system messages are pending for CPCs or CPC Images.

When operating system messages are pending for a hardware object or group of hardware objects:

- The background of the object and its associated group turns cyan (the default)
- The Operating System Messages icon turns cyan and flashes

For a coupling facility logical partition, coupling facility control code uses the console integration facility to display coupling facility messages and to accept coupling facility control code commands. The console integration facility, through the Operating System Messages task, provides the only interface for entering commands to an operating coupling facility LP.

4.2.5 Hardware Management Console Pager (For 9672 R2 and R3, and 9674 C02 and C03 Models)

The Message and State Monitor facilities of the 9672 R2 and R3 model and 9674 C02 and C03 model Hardware Management Console can be enabled to send text notification to a personal paging device whenever a specific message is received from either the hardware subsystem or an operating system, or when a CPC (hardware object) or a CPC Image (Operating system object) changes from one “state” to another “state”. The “state” of an object represents its current condition and functional capability, so a state change represents a change in functional capability that may require attention. Hardware and operating system messages are intended to keep the operator informed of conditions that may require attention. However, not all messages and not all state changes are important; only specific ones require attention and notification of a responsible person.

By enabling the Message and State Monitor facilities and customizing the Site Alert Client* message filters, a user can cause an alphanumeric page message to be sent whenever some message or state change occurs. For example, if the state change notification is enabled and the message filters are customized to detect a state change from “operating” state to “not operating” state, an alphanumeric page message can be sent to an operator, a system programmer, a service representative, or a group of responsible people whenever the processor goes from operating to not operating. This allows notification of critical events even when the system is unattended. Up to five different people can be identified to receive the page notifications.

The Message and State Monitor facilities consist of:

- A Hardware Management Console “console action” to allow the Message and State Monitor facilities to be enabled for any combination of:
  - Operating system message monitoring
  - Hardware message monitoring
  - State change monitoring
- Hardware Management Console LIC to monitor the enabled message and state notifications and to create the messages for the Site Alert Client
- Site Alert Client facilities for customizing the message filter characteristics
- Site Alert Client monitoring facilities
- Site Connect Server* facilities for customizing the paging service parameters and for identifying the recipients of the alphnumeric pager messages
- Site Connect Server facilities for sending the resulting pager messages to the pager service

To make the Message and State Monitor facilities functional, the user must:

- Enable the Message and State Monitor facilities desired on a Hardware Management Console that will always be active
- Customize the Site Alert Client message filters in the Hardware Management Console

- Customize the Site Connect Server users and paging services in the Hardware Management Console
- Customize the operating system to send messages to the console integration facilities of the Hardware Management Console (if operating system message monitoring is enabled)
- Add an asynchronous modem and telephone connection to the Hardware Management Console
- Subscribe to a paging service
- Provide appropriate alphanumeric pagers to the people to be notified

4.2.6 Problem Analysis and Reporting

Each SE monitors and analyzes problems detected on the associated CPC. For problems that are isolated to a single CPC, the results are reported to the Hardware Management Console as a hardware message. For those problems that potentially involve multiple CPCs, that problem data is sent to the problem management Hardware Management Console where data from multiple CPCs is analyzed and reported. The Hardware Management Console configured as a problem management focal point can perform:

- Problem analysis for ESCON, coupling facility, and Sysplex Timer link faults encountered by the CPCs configured to it.
- Service calls for all CPCs configured to it. Enabling the Hardware Management Console as a phone server identifies the Hardware Management Console as having a modem that all CPCs configured to it can use for placing service calls.

When properly configured, the problem analysis function generates SNA generic alerts to report all problems to the host problem management focal point.

4.2.7 Licensed Internal Code

Each SE and each Hardware Management Console has Licensed Internal Code (LIC) and is subject to periodic updates from IBM. One of the Hardware Management Consoles should be configured as a LIC change management focal point to:

- Receive Licensed Internal Code updates remotely from IBM
- Retrieve and distribute LIC updates to all the support elements of all the CPCs configured to the Hardware Management Console.

4.2.8 Remote I/O Configuration Definition and IOCDS Management

Each CPC requires a definition of the I/O attached to it. The Hardware Configuration Definition (HCD) is an MVS application that aids in the definition of all the I/O and aids in the distribution of the appropriate I/O definitions to the appropriate CPCs. The Hardware Management Console configured as a change management focal point assists HCD in finding the names of all defined CPCs. A single HCD then distributes the appropriate IOCDS and IPL parameters to the various SEs of the CPCs defined to the same Hardware Management Console with change management capability.

Remote I/O configuration definition and IOCDS management requires MVS/ESA 5.1 or higher.

4.2.9 Scheduled Operations

The Hardware Management Console provides support for scheduling the times and dates for automatic Licensed Internal Code (LIC) updates and backup of critical hard disk data for any CPCs configured to the Hardware Management Console. You can accomplish this by using the Scheduled Operations task available from the CPC Operational Customization Tasks list.

The Scheduled Operations task allows you to schedule the following LIC-related operations:

- Activate LIC changes
- Deactivate the CPC
- Disruptive install and activate LIC changes (R1 models)
- Concurrent install and activate LIC changes (R2 and higher models)
- Accept LIC changes
- Remove LIC changes
- Concurrent remove and activate LIC changes (R2 and higher models)
- Backup critical hard disk data

4.2.10 Remote Operations

The 9672 hardware systems management products support remote operations in a variety of ways over a variety of communications connections. In each case, the objective is to enable a human or programmed operator to monitor or control a remote system in essentially the same manner as if the operator were at the same site as the remote system. Remote operations include:

- Hardware Management Console operation of S/390 9672s and 9674s.

  The Hardware Management Console is attached by a token-ring LAN to the 9672 or 9674 processors it controls. The Hardware Management Console(s) may be connected to the CPC(s) using bridged LANs. The Hardware Management Console response time will be dependent of the connection speeds. A minimum of 64K bps connection is recommended.

- Remote manual operation of a Hardware Management Console

  A remote operator console capability will be supported via SDLC-attached PS/2 consoles. The SDLC-attached remote operator console capability will use a modem on the Hardware Management Console. The Distributed Console Access Facility (DCAF) product will be used to provide the remote operator console for the Hardware Management Console. To use the remote operator console facilities a PS/2 is required with OS/2, Communication Manager/2 (CM/2) and DCAF products. The Hardware Management Console will have the necessary target portions of DCAF already installed.

- Remote operation of S/390 9672s and 9674s from other platforms

  The facilities and APIs of the Hardware Management Console are available to other management applications by means of the SNMP protocol supported by the HMCA. This protocol supports remote management applications which may be provided by IBM, other software vendors, or the enterprise.

4.2.10.1 Remote Support Facility (RSF)

The Hardware Management Console provides Remote Support Facility (RSF) to aid in the service and maintenance of your system. RSF provides:

- Automatic or customer initiated call for service
- Automatic or customer down-loading of the latest LIC change levels.

RETAIN support of HMC-attached processors will be via SDLC links on the Hardware Management Console to an IBM RETAIN system.

4.2.11 Automation and API Support

NetView provides facilities that enable automation support for SNA networks. You can exploit NetView's automation capabilities by using automation products like TSCF and AOC/MVS that provide automation routines that can manage 9672 or you can write your own automation routines using NetView command lists (CLISTS) or REXX. Problem reports can be generated and sent to NetView using SNA generic alerts.

In addition to providing an end-user with the ability to view and manipulate managed objects, the Hardware Management Console application programming interfaces (APIs) provide support of the SNMP protocol for development and use by remote management applications. The Application Programming Interfaces are:

1. **User interface transitioning** APIs
2. **Management** APIs

The user interface transitioning API provides a local application with the ability to transfer into the HMC user interface in context. This will be exploited by IBM's ESCON manager Release 3 product. The management APIs contain the ability to get/set a HMC managed object's attributes, issue commands to be performed on a managed object from a local or remote application, and receive asynchronous event notifications. The management APIs provide a mechanism to IBM, independent system management vendors, and an enterprise, to integrate with the Hardware Management console application.

4.2.12 Activation

Activating a CPC is an intelligent (LIC controlled) process that takes the CPC from its current state to a fully operational state. The activation may involve a power-on, power-on reset, and IPL, as necessary.

To activate a CPC, you must specify system activation information for each CPC configured to the Hardware Management Console. You specify the CPC system activation information using activation profiles. Activation profiles allow you to define multiple power-on reset (POR) configurations for the same CPC. These profiles allow you to tailor the CPC resources (central processors, storage, and channels) to meet varying business and workload needs.

You use activation profiles to define basic mode and PR/SM LPAR mode configurations. Activation profiles are stored on the SE hard disk so that they are available for future PORs.

You can modify default activation profiles that are provided with the CPC and use them as templates to produce customized profiles to fit your processing needs.

There are three types of activation profiles you can produce:

- Reset - Used during power-on processing
- Load - Used to load an operating system
- Image - Used to define an LPAR partition

You must define a Reset profile and a Load profile for a basic mode system. For PR/SM LPAR mode, you must define a Reset profile and one Image profile for each logical partition.

4.3 Availability

One Hardware Management Console is required for system monitoring and operation of the CPCs configured to it. For high availability applications, IBM recommends that you have at least two Hardware Management Consoles for your configuration to guarantee that the Hardware Management Console functions are available when needed.

You can repair a Hardware Management Console concurrent with the operation of the CPCs configured to it. Having more than one Hardware Management Console provides continuous availability of Hardware Management Console functions, including the following:

- Hardware operator controls, hardware status monitoring, and hardware and operating system messages for all configured CPCs
- Capability to call for service
- Remote operations control
- Problem analysis

4.4 Security Considerations

Since multiple Hardware Management Consoles and SEs are connected via a LAN and remote connections are possible, it is important to understand the use and capabilities enabled for each Hardware Management Console.

Hardware Management Consoles operate as peers with equal access to the CPCs configured to them. The support elements (SEs) for each CPC serialize command requests from Hardware Management Console Applications on a first come, first served basis. There is no guarantee of exclusive control across a sequence of commands sent from a single Hardware Management Console. You should consider these security recommendations:

- Following installation of the CPCs and Hardware Management Consoles in your configuration, the access administrator should change the default logon passwords at the Hardware Management Console.
- Define all Hardware Management Consoles, all CPCs, and all attaching control units to a private LAN. Using a private LAN for your configuration offers several security, availability, and performance advantages as follows:
  - Direct access to the LAN is limited to those Hardware Management Consoles and CPCs attached to it. Outsiders cannot connect to it.
  - Traffic disruption due to temporary outages on the LAN is reduced– including disruptions caused by plugging in and powering on new devices on the LAN (minor) to LAN adapters being run at the wrong speed (catastrophic).
  - LAN traffic is minimized reducing the possibility of delays at the Hardware Management Console user interface.

  Additionally, Hardware Management Console and SE activity, if included on a non-private LAN, could potentially disrupt other LAN activity.
- If using a private LAN is not practical, isolate the LAN used by the Hardware Management Consoles, SEs, and control units by providing a LAN bridge between the isolated LAN and the backbone LAN to provide an intermediate level of security.

  Important: If the LAN bridge exists between a Hardware Management Console and its SEs, the LAN bridge must allow SNA, TCP/IP, and NetBios traffic to flow across the bridge.

- Assign a unique domain name that includes all the CPCs controlled from one or more Hardware Management Consoles.
- Install one or more Hardware Management Consoles that have all of the CPCs you want to control defined to it. Place at least one of these Hardware Management Consoles in the machine room near the CPCs that form its domain.

  The following "enable/disable" controls are provided to help you control access and provide focal point capabilities:

  – Remote Licensed Internal Code (LIC) update (change management focal point)
  – Remote service support
  – Remote customer access
  – Remote service access
  – Auto-answer of the modem

  Change Management Considerations: If you want to limit the number of Hardware Management Consoles that have change management capability (for example, LIC update control and I/O definition and remote IOCDS management capability using HCD), only enable the remote LIC update Hardware Management Console capability for those Hardware Management Consoles to be used as change management consoles. All Hardware Management Consoles are shipped with change management enabled. A least one Hardware Management Console in the domain must have change management enabled.

  Remote Access Considerations: If you want to limit remote access to Hardware Management Consoles, only enable auto-answer of the modem for the Hardware Management Consoles you want to be able to control from a remote console. Remote access is then funneled through the Hardware Management Consoles you designate. Secure the Hardware Management Console (keep it in a locked room).
- If a remote console is used for remote operations access, assign a secure DCAF logon password. You may want to consider the DCAF high security feature to enhance security for the remote connection.
- Logoff each Hardware Management Console when it is not in use. The Hardware Management Console provides a status bar capable of displaying status colors (or grey patterns) to alert you when operator activity is needed, even when no one is logged on.
- Use the OS/2 lockup function to lock the Hardware Management Console keyboard after a predetermined period of keyboard inactivity has expired (if this does not disrupt your site's operating habits).

4.5 Using Scheduled Operations

You can use the Hardware Management Console to schedule operations that allow you to automate and control the frequency of central CPC (CPC) and Hardware Management Console Licensed Internal Code (LIC) updates. This can positively affect the availability by minimizing the frequency and duration of recovery and service activity. You can schedule operations for the Hardware Management Console and for any defined CPC.

CPC configuration data is stored on the hard disk of the SE associated with each CPC and includes IOCDSs, activation profiles (including any logical partition definitions), and System Activity Display (SAD) profiles.

Hardware Management Console configuration data is stored on the Hardware Management Console hard disk and includes information about things configured to the Hardware Management Console, such as CPCs, CPC images, coupling facilities, and user access codes.

Maintaining current backup copies of this configuration information is critical to restoring the production configuration in recovery and service situations.

The configuration data is backed up from the SE or Hardware Management Console hard disk to rewriteable optical disk at the Hardware Management Console. To use scheduled backup, the correct optical disk must be mounted in the Hardware Management Console at the time of the backup.

4.5.1 Licensed Internal Code (LIC) Maintenance

The Hardware Management Console also allows you to schedule the automatic retrieval, distribution, and installation and activation of Licensed Internal Code (LIC) changes for each SE and for the Hardware Management Console. You can also schedule the automatic removal or acceptance of LIC changes to the base LIC used for system activation. Any or all of these operations can also be performed manually.

LIC changes can represent fixes to known LIC problems or delivery of new function (for example, an RPQ). LIC changes can be retrieved from IBM service support system through the Remote Support Facility (RSF) or through diskette or optical disk cartridge. Once retrieved, SE LIC changes are loaded onto the SE hard disk and Hardware Management Console LIC changes are loaded onto the Hardware Management Console hard disk.

LIC maintenance can either be a service representative task or a customer task. LIC maintenance options are available to both the Advanced Operator and System Programmer user modes.

4.5.2 Backup Recommendations

If you prefer to control the backup of critical configuration data and LIC maintenance of the S/390 Parallel Transaction Server solution. IBM recommends that you schedule automatic operations for the Hardware Management Console and any CPC, group of CPCs, or all CPCs configured to the Hardware Management Console and do the following:

1. Backup critical data
2. Accept LIC changes
3. Retrieve LIC changes
4. Install and activate LIC changes

4.6 I/O Configuration

When installing 9672 and 9674 CPCs, you must define the I/O configuration to the channel subsystem (CSS) of each CPC. The CSS for a CPC controls the communication between the CPC (or the logical partitions on the CPC if the CPC is running in LPAR mode), its channels, and any control units and devices attached to the channels.

To do this, you must create the source necessary to build an input/output configuration data set (IOCDS) that contains the appropriate I/O configuration information for each CPC.

4.6.1 Recommended Approach

If you do not have the MVS Hardware Configuration Definition (HCD) tool, you must use stand-alone IOCP to define the initial IOCDS for your CPC. MVS, VM, and VSE users can then use the batch version of IZP IOCP available on their operating system to build I/O configurations.

If you have the MVS Hardware Configuration Definition tool, IBM recommends using both stand-alone

IOCP (for the initial IOCDS) and HCD. This can simplify creating and distributing the necessary IOCDSs across CPCs. This approach takes advantage of the capability to remotely write IOCDSs featured by HCD running on MVS/ESA SP 5.1 and supported on 9672 and 9674 CPCs that are configured to the same Hardware Management Console with change management capability.

Note: The 9674 Coupling Facility C01/C02/C03 Models support another CPC remotely writing an IOCDS to it; however, neither a 9674 C01, C02, or C03 Model nor a coupling facility LP running on a 9672 CPC can remotely write an IOCDS using HCD because HCD ***cannot*** run in a coupling facility.

4.6.1.1 HCD Advantages

This approach:

- Eliminates the need to run stand-alone IOCP at the installation of each CPC
- Allows you to enable dynamic I/O configuration for all IOCDSs so that dynamic I/O configuration is available at the initial activation of each CPC assuming HCD was used to write the IOCDSs.

  Note: Dynamic I/O configuration does ***not*** support CFR channel paths.

4.6.2 Planning Considerations

IBM recommends that you consider the following when planning the definition of your I/O configuration:

1. Prior to installation of 9672 and 9674 CPCs, run HCD on another CPC to define, in advance, a single production input/output definition file (IODF) for all CPCs configured to the same Hardware Management Console. Then, use HCD to create an error-free production IOCP input file from this IODF for one of the CPCs and save it to tape.

   Note: If you do not have HCD running on a CPC at your site, use a batch version of IOCP (MVS, VM, or VSE) to manually create an IOCP input file for each CPC. If you do not have access to a batch version of IOCP, consult your IBM representative to help you create the IOCP input files. Without HCD, you must use stand-alone IOCP on each CPC to define the initial I/O configuration for each CPC.
2. Install the 9672(s) and 9674(s), ensuring that the support element (SE) for each CPC configured to the Hardware Management Console is LAN-attached to the Hardware Management Console. Additionally, at the Hardware Management Console, you ***must***:
   - Define each CPC to the Hardware Management Console
   - Enable remote licensed internal code (LIC) changes

     Enabling remote LIC changes gives the Hardware Management Console change management capability and allows the Hardware Management Console to support the HCD capability to remotely write IOCDSs.
3. Use the Hardware Management Console to run stand-alone IOCP on the CPC you created the IOCP input file for in Step 1. Use stand-alone IOCP to define a minimum configuration IOCDS that identifies the channel-attached tape drive with the tape containing the production IOCP input file.
4. Power-on reset (POR) the CPC using the minimum configuration IOCDS.
5. Use stand-alone IOCP to import the production IOCP input file from tape. On 9672 and 9674 CPCs, the Import from Tape function creates an IOCDS and writes it to the CPC SE providing the IOCP input file contains no errors.
6. POR the CPC using the IOCDS created in Step 5. IPL MVS and start HCD.
7. Use the HCD Update IOCDS function to build and write production IOCDSs to the SEs of the rest of the CPCs configured to the Hardware Management Console with change management capability.
8. POR the CPCs using the IOCDSs created in Step 7.

4.6.3 Coupling Facility Continuous Availability

You can use HCD running on a 9672 CPC to remotely write an IOCDS to a 9674 CPC or to another 9672 CPC that is running a coupling facility logical partition (LP) as long as both CPCs are configured to the same Hardware Management Console with change management capability.

The IOCDS can then be used during activation of the CPC serving as a coupling facility. However, you ***cannot*** use HCD to dynamically add or delete channel paths defined in the coupling facility channel configuration.

This can be a problem if you want to add CPCs or LPs to a parallel sysplex configuration to increase its processing capacity. In this case, you need to define and, possibly install, additional coupling facility channel hardware to connect from the coupling facility to the newly added CPCs or LPs.

Defining additional coupling facility channels to the coupling facility I/O configuration requires you to update or create a new IOCDS to include the additional coupling facility channel definitions, deactivate the coupling facility, then re-POR and reactivate the coupling facility using the updated or new IOCDS. If you also need to install coupling facility channel hardware, you may have to power-off the CPC that the coupling facility is running on.

In either case, you could cause a partial or complete disruption of data sharing activity depending on the number and capacity characteristics of the coupling facilities in the parallel sysplex configuration. To avoid this situation and to enhance the continuous availability of the parallel sysplex configuration, you should consider doing the following:

- Pre-define a maximum number of coupling facility channels in the coupling facility IOCDS whether they are installed or not.

  This will eliminate the need to create or modify an IOCDS and re-POR the coupling facility to add the channels to the coupling facility channel configuration when additional CPCs are added to the parallel sysplex configuration.

  Non-Dedicated Coupling Facility Configurations: If the coupling facility is not the only logical partition (LP) running on the CPC (***not*** recommended for production use), you will have to weigh the channel requirements of the non-coupling facility LPs against the availability requirements you have for the coupling facility when determining how many additional coupling facility channels you should predefine. Pre-defining coupling facility channels in this case could deny use of installed ESCON or parallel channels by the non-coupling facility LP or LPs.

- Pre-install additional coupling facility channel adapter cards so you can install, at a later date, additional coupling facility channel cards concurrent with coupling facility operation.

  The coupling facility channel hardware implementation consists of a coupling facility channel adapter card that contains slots for the plugging of two coupling facility channel cards. One coupling facility channel plugs into each coupling facility channel card.

  Coupling facility channel adapter cards cannot be installed without turning off CPC power. However, coupling facility channel cards (and coupling facility channels) can be installed concurrent with CPC (and coupling facility) operation as long as the coupling facility channel adapter card has been preinstalled.

  To maintain continuous availability, you could pre-install additional coupling facility channel adapter cards at CPC installation or during a scheduled outage.

Chapter 5. Standard and Optional Features

5.1 Overview

The column headings in the tables under “Features Dependent on Architectural Mode” and “Programming Assists Dependent on Architectural Mode” identify the architectural modes of the 9672 CPC.

Additional information about the mode-dependent features can be found in the following publications:

- *IBM Enterprise Systems Architecture/390: Principles of Operation*
- *IBM System/370 Principles of Operation*
- *IBM Enterprise Systems Architecture/370* and System/370 Vector Operations*
- *IBM System/370 Extended Architecture Interpretive Execution*
- *ES/9000 ES/3090 Processor Resource/Systems Manager Planning Guide*, GA22-7123

5.2 Features Dependent on Architectural Mode

The following table identifies the features that are available in ESA/390 and System/370 architectural mode. System/370 is available only in LPAR mode and for a virtual machine.

Table 5-1 (Page 1 of 4). Features Dependent on Architectural Mode

Feature	S/370	ESA/390
Advanced address space	N/D	Standard
Asynchronous data mover facility	N/D	Standard
Asynchronous pageout facility	N/D	Standard
Basic control mode	Standard	N/D (Note 1)
Bimodal addressing	N/D	Standard
Branch and save	Standard	Standard
Broadcasted purging	N/D	Standard
Byte-oriented operand	Standard	Standard
Cancel I/O	N/D	Standard
Channel indirect data addressing	Standard	Standard
Channel subsystem	N/D	Standard
Channel-subsystem call	N/D	Standard
Clear I/O	Standard	N/D (Note 2)
Command retry	Standard	Standard
Concurrent Channel Maintenance	Standard	Standard
Concurrent Sense	Standard	Standard
Conditional swapping	Standard	Standard
Console Integration	N/D	Standard
Coupling Facility	N/D	Standard
CPU timer and clock comparator	Standard	Standard
Data compression (hardware assisted)	N/D	Standard
Data Streaming	Standard	Standard
DB2* sort facility	Standard	Standard
Dynamic I/O configuration	N/D	Standard
Enhanced move page	N/D	Standard
Error Checking and Correction	Standard	Standard
ESCON Channels	N/D	Optional
ESCON Channel-to-Channel Adapter	N/D	Optional
ESCON Multiple Image Facility	N/D	Standard
Expanded storage	Optional	Optional
Extended-precision divide	N/D	Standard
Extended-precision floating point	Standard	Standard
Extended real addressing (26-bit)	Standard	N/D (Note 3)
Extended sorting	N/D	Standard

Table 5-1 (Page 2 of 4). Features Dependent on Architectural Mode		
Feature	**S/370**	**ESA/390**
Extensions for virtual machines	N/D	Standard
Fast release	Standard	N/D (Note 2)
Floating point	Standard	Standard
Halt device	Standard	N/D (Note 2)
High-Speed Buffer Storage	Standard	Standard
Incorrect-length-indication suppression	N/D	Standard
Interpretive execution (SIE)	N/D	Standard
Interval timer	Standard	N/D
I/O-device self-description	N/D	Standard
I/O Error Alert	Standard	Standard
I/O interface reset	N/D	Standard
I/O Power Sequence Control	Standard	Standard
Key-controlled storage protection	Standard	Standard (Note 4)
Limited channel logout	Standard	N/D (Note 2)
Logging volume reduction	N/D	Standard
Logical String Assist	N/D	Standard
Monitoring	Standard	Standard
Move page	N/D	Standard
Multiprocessing: CPU address identification CPU signaling and response Prefixing Shared main storage TOD clock synchronization	 Standard Standard Standard Standard Standard	 Standard Standard Standard Standard Standard
Open Systems Adapter (OSA)	N/D	Standard
Page protection	N/D	Standard
PER extensions	N/D	Standard
Private space	N/D	Standard
Processor Resource/Systems Manager*	N/D	Standard
PSW-key handling	Standard	Standard
Recovery extensions	Standard	N/D
Scalar square root	Standard	Standard
SCP-Initiated Reset	Standard	Standard
Segment protection	Standard	N/D (Note 5)
Service signal	Standard	Standard
Sorting instructions	N/D	Standard
Storage-key instruction extensions	Standard	Standard
Storage-key instructions (ISK, SSK)	Standard	N/D (Note 4)

Table 5-1 (Page 3 of 4). Features Dependent on Architectural Mode

Feature	S/370	ESA/390
Storage-key 4KB block:		
Single-key 4KB blocks	Standard (Note 6)	Standard
Storage-key exception control	Standard	N/D (Note 4)
Subspace Group Facility Support	N/D	Standard
Subsystem storage protection	N/D	Standard
Suppression on Protection	Standard	Standard
Sysplex Timer	N/D	Optional
Automatic propagation delay adjustment	N/D	Optional
External time source	N/D	Optional
System/370 extended facility: Non-MVS-dependent portion	Standard	Standard
System/370 I/O instructions	Standard	N/D (Note 2)
Test block	Standard	Standard
Time-of-day (TOD) clock	Standard	Standard
Tracing (ASN, branch, and explicit)	N/D	Standard
Translation:		
Dynamic address translation:		
2KB page size	N/I	N/D
4KB page size	Standard	Standard
64KB segment size	Standard	N/D
1MB segment size	Standard	Standard
Extended mode control	Standard	N/D (Note 1)
Program-event recording 2 (PER 2)	N/D	Standard
Set-system-mask suppression	Standard	Standard
Store status	Standard	Standard
3033 extension:		
Dual-address space (DAS)	Standard	Standard (Note 7)
SIOF queuing	Standard	N/D (Note 2)
Suspend and resume	Standard	N/D (Note 2)
31-bit IDAWs	Standard	Standard
31-bit real addressing	N/D	Standard

Legend:

N/D This function is not implemented because it is not defined in the principles of operation publication for this architectural mode.

N/I This function is not implemented but is defined in the principles of operation publication for this architectural mode.

Table 5-1 (Page 4 of 4). Features Dependent on Architectural Mode

Feature	S/370	ESA/390

Notes:

a. Operation in ESA/390 mode is comparable to operation in EC mode of System/370.

b. Replaced by standard functions of the channel subsystem operating in ESA/390 mode; channel program compatibility with System/370 is maintained.

c. Replaced by 31-bit real addressing.

d. The storage-key instruction extensions provide the required function to manage the storage keys in ESA/390 mode. System/370 instructions ISK, RRB, and SSK are not in ESA/390 architecture.

e. Replaced by page protection.

f. Double-key 4KB blocks are not implemented.

g. Does not include dual-address space (DAS) tracing. Address space number (ASN) tracing provides a comparable function.

5.3 Programming Assists Dependent on Architectural Mode

Table 5-2. Programming Assists Dependent on Architectural Mode

Programming Assists	S/370	ESA/390
System/370 extended facility:		
MVS-dependent portion		
Four lock-handling instructions	Standard	Standard
Six tracing instructions	Standard	N/D
Fix page instruction	Standard	N/D
SVC assist instruction	Standard	Standard
Add FRR instruction	Standard	Standard
VM assist for MVS/370 assists	Standard	N/D
Virtual-machine assists under SIE	Standard	Standard

Legend:

N/D This function is not implemented because it is not defined in the principles of operation publication for this architectural mode.

5.4 Describing the Features

This section describes the features and programming assists for the 9672 Parallel Transaction Server (E-Models and P-Models) and the 9672 Parallel Enterprise Server R1/R2/R3 Models. The feature descriptions should be used with the preceding tables.

5.4.1 Advanced Address Space

Advanced address-space functions of 9672 CPCs provide capabilities not available in the 370-XA architecture; these ESA/390 functions allow programs and their data to reside in different address spaces. Data can be accessed in multiple address spaces concurrently, and address spaces can be selected for processing without control program intervention, if authority has been previously granted by the control program. The five components of the advanced address space facilities are:

- Access-register mode
- Linkage stack mechanism
- Home-space mode
- Real-storage access
- Data movement

5.4.1.1 Access-Register Mode

Sixteen access registers allow the program immediate access to as many as sixteen 2GB address spaces including the address space in which the program resides. In the access-register translation mode, the B-fields and/or R-fields of many instructions can designate both a general register and an access register. The contents of the access register, along with the contents of protected tables in storage, specify the operand address space to be accessed.

By changing the contents of the access registers, the program (under the control of an authorization mechanism) can have fast access to hundreds of different operand address spaces. Instructions are provided to load and store the contents of the access registers and to change between access-register mode and other translation modes.

5.4.1.2 Linkage Stack Mechanism

A linkage stack can be used to pass control between programs residing in the same or different address spaces. The called and calling programs can have degrees of privilege and authority that are arbitrarily different. The state of the calling program is saved on the linkage stack and is restored during the return linkage.

5.4.1.3 Home-Space Mode

The home-space mode provides an efficient means for the control program to obtain control in the home address space where principal control blocks for a dispatchable unit (a task or a process) are kept.

5.4.1.4 Real-Storage Access

The Load Using Real Address and Store Using Real Address instructions allow the control program to access data in real storage more efficiently.

5.4.1.5 Data Movement

The Move with Destination Key instruction and the Move with Source Key instruction can improve performance when data is to be moved alternately (in both directions) between two storage areas that are fetch-protected by means of unequal keys.

5.4.2 Asynchronous Data Mover Facility (ADMF)

ADMF allows DB2 users in an MVS/ESA environment to more efficiently move large amounts of data between central and expanded storage. DB2 users may see a reduction in elapsed time for selected query applications by using central processor complex (CPC) storage rather than DASD. ADMF can be used with data stored in a compressed format.

ADMF offloads the movement of data between central and expanded storage to the CPC I/O subsystem allowing CPs to accomplish additional work.

5.4.2.1 PR/SM LPAR Mode Considerations

In PR/SM LPAR mode, logical partitions (LPs) enabled to use ADMF share the facility. LP requests to use the facility are serviced on a first come, first served basis. Coupling facility LPs and LPs running in System/370 mode ***cannot*** use ADMF. You must deactivate an LP to enable or disable LP use of ADMF.

5.4.3 Asynchronous Pageout Facility

The asynchronous pageout facility allows the CPC to more efficiently move pages of data from central storage to expanded storage. This can be valuable to processing environments where a large volume of block paging takes place (for example, numerically intensive computing [NIC]).

With the asynchronous pageout facility, a CPC can move a page of data from central storage to expanded storage without delaying the processing of subsequent instructions by the central processor (CP). The CP can continue to process instructions as soon as data movement begins.

This facility contrasts with the synchronous movement of data from central storage to expanded storage which caused the CP to delay processing of instructions until data movement completed.

Software Corequisites: The asynchronous pageout facility requires the following software corequisites:

- MVS/ESA SP Version 4 Release 3

- VM/ESA Version 1 Release 2

  VM/ESA 1.2 guest support is for MVS/ESA SP 4.3 preferred guests only. If you run MVS as a V=V guest under VM/ESA 1.2, MVS will use the synchronous expanded storage access facility to avoid performance degradation.

5.4.4 Basic Control Mode

Basic control (BC) mode provides a PSW format that is compatible with the PSW format of System/360.

5.4.5 Bimodal Addressing

Bimodal addressing permits 31-bit logical addressing, yet allows users to continue running System/370 problem programs that use 24-bit logical addresses.

5.4.6 Branch and Save

Branch and save provides the Branch and Save instruction (BAS and BASR).

5.4.7 Broadcasted Purging

Broadcasted purging provides for conditionally updating tables associated with address translation and clearing address-translation lookaside buffers in multiple CPUs.

5.4.8 Byte-Oriented Operand

Byte-oriented operand allows storage operands of most unprivileged instructions to appear on any byte boundary without causing a specification exception and a program interruption. This feature applies to fixed-point, floating-point, and logical operands. It does not apply to instruction addresses, privileged instructions, or channel command words (CCWs).

5.4.9 Cancel I/O or Cancel Subchannel

The Cancel I/O facility, or Cancel Subchannel (XSCH) instruction, allows the program to withdraw a pending start function from the designated subchannel without signaling the device.

Use Cancel I/O to simulate a START I/O instruction for a 370-mode guest and in certain error recovery situations.

5.4.10 Channels

ESCON and parallel channels can be assigned for block or byte multiplexer mode of operation. Coupling facility channels can be assigned for either coupling facility sender or coupling facility receiver mode of operation. Open Systems Adapter channels can be assigned for OSA mode of operation.

5.4.11 Channel Indirect Data Addressing

The addresses contained in channel command words (CCWs) in virtual storage must be translated by the system control program before execution. Channel indirect data addressing allows immediately adjacent areas of virtual storage to be mapped into nonadjacent areas of absolute storage.

5.4.12 Channel Subsystem

The ESA/390 dynamic channel subsystems queue I/O requests, selects from as many as eight channel paths to any I/O device, and handle I/O busy conditions. Fourteen ESA/390 I/O instructions are associated with the channel subsystem.

9672 and 9674 CPCs have a single channel subsystem.

5.4.13 Channel-Subsystem Call

Channel-subsystem call provides various functions for use in the management of the I/O configuration. Some of the functions acquire information about the configuration from the accessible elements of the configuration, while others dynamically change the configuration.

5.4.14 Clear I/O

Clear I/O provides the clear I/O function in a channel when the privileged Clear I/O (CLRIO) instruction is executed. The clear I/O function causes a channel to discontinue its current I/O operation with an addressed I/O device by storing the status of the operation in the channel status word (CSW) and by making the associated subchannel available.

5.4.15 Command Retry

Command retry allows a subchannel to retry a command without causing an I/O interruption. The retry is initiated by a control unit.

5.4.16 Concurrent Sense

Concurrent sense improves performance by allowing sense information to be presented at the time of an interruption due to a unit-check condition. This avoids the need for a separate I/O operation to obtain the sense information.

5.4.17 Conditional Swapping

Conditional swapping makes available the Compare and Swap (CS) and Compare Double and Swap (CDS) instructions.

5.4.18 Concurrent Channel Maintenance

Concurrent channel maintenance allows the repair and replacement of a channel card without having to take down the system or the channel group that is being serviced. By extending the function provided by single channel service mode, the service representative can perform maintenance on a selected channel or channels without impact to the system.

5.4.19 Console Integration

The Console Integration facility, supported by MVS/ESA SP Version 4 Release 2 or higher and Coupling Facility Control Code, allow you to use the Hardware Management Console to initialize the MVS system and to recover the MVS operator console. The facility provides an availability enhancement when the MVS master and alternate operator consoles become unavailable. In this case, you can temporarily use the Hardware Management Console as an MVS operator console to reestablish an operational MVS console.

Note: The console integration facility is not intended to replace the channel-attached MVS operator consoles which are still required for normal system operation.

5.4.20 Coupling Facility

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 MVS images in the sysplex define the shared structures used in the coupling facility.

Enhancements to PR/SM LPAR allow 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).

New 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 connection.

5.4.21 CPU Timer and Clock Comparator

The CPU timer of each central processor is a high-resolution timer that causes an interruption whenever its value is negative. The interruption request is allowed by setting bit 21 in control register 0 and the external mask bit in the PSW.

The CPU timer measures central processor elapsed time and causes an interruption at the end of the period that is specified by the program. The timer is decremented when the central processor is executing instructions and during the wait state, but is not decremented when the central processor is in the stopped state. The program can initiate inspection of the CPU timer by using the Store CPU Timer (STPT) instruction and can set the timer to a specific value by using the Set CPU Timer (SPT) instruction. The contents of the CPU timer are reset to 0 by initial CPU reset.

Note: When the time-of-day (TOD) clock is in the stopped state or in the error state, the CPU timer is not decremented.

The clock comparator of each central processor provides for an interruption when the TOD clock reaches a value specified by the program. The interruption is allowed when the central processor sets bit 20 in control register 0 and the external mask bit in the PSW.

The format of the clock comparator is the same as that of the TOD clock. A clock-comparator interruption is an external interruption. The program can initiate inspection of the clock comparator by using the Store Clock Comparator (STCKC) instruction and can set it by using the Set Clock Comparator (SCKC) instruction. The contents of the clock comparator are reset to 0 by an initial CPU reset.

Note: When the TOD clock is in the stopped state or in the error state, the clock comparator is not operating.

5.4.22 Data Compression

9672 CPCs support hardware-assisted data compression for DB2, IMS, DFSMS/MVS, and VTAM running under MVS/ESA.

New hardware facilities allow the compression and expansion of data which can provide significant reductions in the total cost of computing by significantly reducing:

- DASD and tape storage requirements
- The number of communication lines or required line speeds for network data transmission

- Central processor (CP) cycles needed to perform data compression and expansion (when compared to software-only data compression)

Data compression, along with the increased data rate of ESCON channels and sequential data striping, may provide further enhancements to the effective bandwidth of 9672 systems. When combined, these attributes may improve channel utilization and enhance system level performance in situations where I/O bandwidth has been a limiting factor.

5.4.22.1 Software Corequisites

Hardware-assisted compression requires the following minimum levels of software support:

- **For DB2:**
  - DB2 Version 3
  - MVS/ESA SP Version 4 Release 2.2 (4.2.2)
- **For IMS:**
  - IMS Version 4 Release 1 (4.1)
  - MVS/ESA SP Version 4 Release 3 (4.3)
- **For DFSMS/MVS:**
  - DFSMS/MVS Version 1 Release 2 (1.2)
  - MVS/ESA SP Version 4 Release 3 (4.3) with APAR support
- **For VTAM:**
  - VTAM Version 4 Release 2 (4.2)
  - MVS/ESA SP Version 4 Release 3 (4.3) with APAR support

5.4.23 Data Streaming

Data streaming is available on all block multiplexer channels. It permits higher data rates (up to 4.5MB per second, depending on the control units attached) and longer cable lengths. Data streaming is initiated by the control unit. The channel subsystem permits the intermixed attachment of data-streaming and non-data-streaming control units on the same channel.

5.4.24 DB2 Sort Facility

See 5.4.35, “Extended Sorting” on page 5-17.

5.4.25 Dynamic I/O Configuration

Dynamic I/O configuration allows you to modify the I/O configuration without performing:

- Power-on reset (POR) of the CPC
- Initial program load (IPL) of the system control program (SCP)

Dynamic I/O configuration allows you to add, delete, or modify the I/O definitions of channel paths, control units, and I/O devices in the CPC. It also allows you to save the changes you made to the I/O configuration definitions and apply them to the active I/O configuration data set (IOCDS).

Initiate these changes by using MVS/ESA SP Version 4 Release 2 and the Hardware Configuration Definition (HCD) or by using VM/ESA Version 2 Release 1. It is operational in basic and LPAR modes.

Software Corequisites: For 9672 models, dynamic I/O configuration requires the following minimum software corequisites:

- MVS/ESA* SP Version 4 Release 2 (MVS/ESA SP 4.2) with APAR support or MVS/ESA SP 4.2.2 with APAR support provide the capability to make dynamic I/O configuration changes to configurations without shared channels.
- MVS/ESA SP 4.3 or later adds support for dynamic I/O configuration and ESCON Multiple Image Facility for ESCON channels, allowing dynamic I/O configuration changes for I/O configurations that have both shared and unshared channels.
- MVS/ESA SP 5.1 or later adds support for dynamic I/O configuration for ESCON byte multiplexer (CBY), coupling facility sender (CFS), and OSA channel paths, allowing dynamic I/O configuration changes.
- VM/ESA Version 2 Release 1 or later provides support for dynamic I/O configuration for ESCON, parallel, coupling facility sender (CFS), and OSA channel paths, allowing dynamic I/O configuration changes for I/O configurations that have both shared and unshared channels.

5.4.26 Enhanced Move Page for VM

Enhanced move page for VM adds problem-program options that allow application programs to make more effective use of Move Page, increasing throughput and reducing processor usage. Enhanced move page provides the following:

- Movement between storage with different keys.
- Control over whether a translation exception results in an interruption or sets a condition code, providing greater flexibility in handling these situations.
- Destination reference-intention indication, which can lead to reduced page handling.
- Additional information provided with translation exceptions makes exception handling by the control program more efficient.

5.4.27 Error Checking and Correction

Data paths between expanded storage (if configured) and central storage, and between central storage and the channels and central processors, are checked using either parity or error checking and correction.

Error checking and correction (ECC) code bits are stored with the data in central storage and in expanded storage data arrays. ECC codes apply to data stored in and fetched from central storage and expanded storage.

For 9672 E, P, and R1 models, and 9674 model C01, double-bit error detection is performed in both central and expanded storage. Single-bit error correction takes place in both central storage and expanded storage.

For 9672 R2 and R3 models, and 9674 C02 and C03 models, Error Checking and Correction detects and corrects not only single bit errors, but also four bit errors in a single memory chip.

5.4.28 ESCON Channels

ESCON Channels transfer information on a link in a serially transmitted synchronous bit stream (serial transmission) through fiber optic channel cables and:

- Operate using link-level protocols and device-level protocols, or
- Attach to an IBM 9034 ESCON Converter for using bus and tag cables attachment to control units with parallel interfaces.

5.4.28.1 ESCON Extended Distance Feature (ESCON XDF)

ESCON XDF significantly lengthens the connectivity distances possible using the ESCON Architecture.

By using ESCON XDF single-mode fiber optic technology, connectivity distances can significantly exceed those offered by the LED multimode fiber technology.

5.4.29 ESCON Channel-to-Channel Adapter

The ESCON channel-to-channel adapter (ESCON CTCA) architecture uses the ESCON I/O interface, which has one operating mode. This operating mode is similar to parallel CTC extended mode as defined by the parallel I/O interface channel-to-channel adapter architecture. See *ESA/390 ESCON Channel-to-Channel Adapter*, SA22-7203.

5.4.29.1 ESCON CTC Support for CTC Applications Running in Parallel CTC Extended Mode

CTC applications running in parallel CTC extended mode require the IBM 3088 Multisystem Channel Communication Unit (MCCU). However, you can remove the 3088 hardware requirement by migrating CTC applications running in parallel CTC extended mode to ESCON CTC channel paths.

5.4.29.2 ESCON CTC Support for CTC Applications Running in Parallel CTC Basic Mode

Beginning with MVS/ESA SP 4.3, MVS can provide the operating system support for ESCON CTC support for CTC applications running in parallel CTC basic mode. By migrating CTC applications running in parallel CTC basic mode, to ESCON CTC channel paths, you can remove the 3088 hardware requirement for parallel CTC basic mode operation.

5.4.30 ESCON Multiple Image Facility (EMIF)

The ESCON Multiple Image Facility (EMIF) allows multiple logical partitions to directly share channels and to optionally share any of the control units and associated I/O devices configured to these shared channels. Enhancements to the channel subsystem (CSS) allow the CSS to be operated as multiple logical CSSs when the CPC is operating in LPAR mode so that one logical CSS is automatically configured to each logical partition. ESCON Multiple Image Facility also provides a way to limit the logical partitions that can access a reconfigurable channel or a shared channel to enhance security.

See Chapter 8, “ESCON Multiple Image Facility (EMIF)” on page 8-1.

5.4.31 Expanded Storage

Expanded storage is available on 9672 and 9674 CPCs as a portion of central storage that can be used as high-speed, high-capacity storage for transferring 4KB blocks to and from central storage.

5.4.32 Extended-Precision Divide

Extended-precision divide provides the Divide (DXR) instruction for extended-precision floating-point operands.

5.4.33 Extended-Precision Floating Point

Extended-precision floating point provides seven floating-point instructions that use the extended-precision format (a signed 7-bit characteristic and a 28-digit fraction).

5.4.34 Extended Real Addressing

Extended real addressing permits the addressing of real storage in excess of 16MB. The system control program uses extended real addressing for locating user programs and portions of the system control program in central storage at real addresses to 64MB. Extended real addressing does not affect virtual addressability, which cannot exceed 16MB in each address space.

5.4.35 Extended Sorting

Extended sorting provides instructions that improve the performance of the DB2 sorting function. These improvements significantly reduce the amount of elapsed time and CPU time required for DB2 queries using sort when used with DB2 Version 2 Release 3.

5.4.36 Extensions for Virtual Machines

The extensions for virtual machines consist of:

- VM-data-space facility, which makes ESA/390 access-register architecture more useful in virtual machine applications. The facility improves the ability to address a larger amount of data and to share data.

  This facility is available when used with VM/ESA Version 1 Release 1
- A new storage-key function that improves performance by removing the need for the previously used RCP area.
- Improvement of interpreted SIE (available with region relocation) to permit preferred guests under VM when VM itself is operating as a high performance guest.
- An optional special-purpose lookaside for some of the guest state information and greater freedom in certain implementation choices.

5.4.37 Fast Release

Fast release provides the start I/O fast release function on a channel when the Start I/O Fast Release (SIOF) instruction is executed. This function provides for early release of the central processor that executes the instruction. Fast release occurs before the device-selection procedure is completed, thereby reducing the central processor delay associated with the operation.

5.4.38 Floating Point

Floating point provides the floating-point instructions and the floating-point registers. In System/370, floating point combined with the commercial instruction set is sometimes referred to as the System/370 universal instruction set.

5.4.39 Halt Device

Using the privileged Halt Device (HDV) instruction, the halt-device function signals the addressed I/O device to terminate its current I/O operation.

5.4.40 High-Speed Buffer Storage

High-speed buffer storage in each central processor satisfies many storage fetch requests, making the effective storage access time much shorter than the actual central storage cycle time.

5.4.41 Incorrect-Length Indication-Suppression Facility

For format-1 channel programs, the Incorrect-Length Indication-Suppression facility provides the incorrect-length indication mode. The Incorrect-Length Indication-Suppression facility includes the incorrect-length suppression-mode control, bit 24 of word 1 of the operation request block (ORB).

5.4.42 Interpretive Execution

The interpretive execution facility provides hardware support for several areas of virtual machine operation such as interval timer operation, prefixing, address translation, and privileged instruction handling. This facility provides the Start Interpretive Execution (SIE) instruction (with interception format 2 installed), which is used to dispatch all virtual machines.

5.4.43 Interval Timer

The interval timer of each central processor provides external interruptions on a program-controlled basis. The value stored at a specified storage location is automatically decremented by 1 in bit position 23 every 3.33 milliseconds. The program receives an external interruption request when the interval timer decrements from 0 to a negative value. (Bit 7 of the PSW and bit 24 of control register 0 must be on.) The range of the interval timer is approximately 15.5 hours.

Note: When the TOD clock of a central processor is in the stopped or error state, the interval timer is not operating.

5.4.44 I/O-Device Self-Description

I/O-device self-description allows a device to describe itself and its position in the I/O configuration.

5.4.45 I/O Error Alert

I/O error alert permits a channel to be alerted when a malfunction affects the ability of a control unit to continue operating.

5.4.46 I/O Interface Reset

I/O interface reset, or SCP-initiated reset, allows a system control program (SCP) to reset its I/O configuration prior to entering the disabled wait state following certain checkstop conditions. This minimizes the occurrence of shared DASD Sysplex hang conditions.

When the SCP places the system in some disabled wait states, the I/O interface reset facility enables the system to issue an I/O interface reset to the channel interfaces connected to the operating configuration. The I/O interface reset releases any shared devices designated as reserved by the SCP upon entering disabled wait state. This allows access to the shared devices by other SCPs.

This function is available in basic mode and LPAR mode. At the Hardware Management Console, use the Activation Profiles task (Options page) to enable the function. The Activation Profiles task is available from the CPC Operational Customization Task Area. See 5.4.59.3, “PR/SM LPAR Auto-Reconfiguration” on page 5-22 for usage restrictions.

For more information, see *CPC Console Operations*, GC38-0453.

5.4.47 I/O Power Sequence Control

9672 CPCs (P-Models and 9672 R-Models with parallel channels) offer I/O power sequence control for 16 control units. You can extend this by two additional 16-unit increments.

5.4.48 Key-Controlled Storage Protection

Key-controlled storage protection prevents unauthorized access to information in central storage. Key-controlled storage protection includes both store protection and fetch protection. If storage protection is violated, data is not stored into the protected area; if fetch protection is violated, data is not retrieved from the protected area. When a violation is recognized, a program interruption occurs.

5.4.49 Limited Channel Logout

Limited channel logout provides 4 bytes of channel status information for model-independent recovery from channel errors.

5.4.50 Logging Volume Reduction

Logging volume reduction together with IMS/ESA* Version 3 should result in reduced logging volume for DL/I databases accessed from IMS/DC or CICS.

5.4.51 Logical String Assist

The logical string assist provides instructions that more efficiently process string data types. These instructions are of particular value to large, complex C-programming language applications that rely heavily on null-character-terminated string manipulation functions.

The logical string assist provides the following architected instructions:

COMPARE LOGICAL STRING
: Compares bytes between two strings until it either finds an inequality or a specified ending character

MOVE STRING
: Moves a string until it finds a specified ending character

SEARCH STRING
: Searches a string of a specified length for a specified character

The first two instructions–COMPARE LOGICAL STRING and MOVE STRING–are particularly useful to the C-programming language that normally delimits strings by a null ending character.

The IBM SAA AD/Cycle C/370 compiler will provide a compile-time option to generate these new instructions for new or re-compiled code.

For more information, see:

- *ESA/390 Principles of Operation*, SA22-7201-01

Software Corequisites: The logical string assist requires the following minimum software corequisites:

- MVS/ESA SP Version 4 Release 3 or higher
- VM/ESA Version 1 Release 2

- AD/Cycle C/370 Compiler Version 1 Release 1.1 with APAR support

5.4.52 Monitoring

Monitoring provides a means of selectively recording designated events in the execution of a program. This facility is implemented by the Monitor Call (MC) instruction.

5.4.53 Move Page

Move page allows an application program to move a page (4KB) of data efficiently between main and expanded storage, provided that the source and destination pages are both valid. This instruction is used with the MVS/ESA HSPSERV routine, which handles the situation when one or both of the pages are invalid. The Move Page instruction is also used by MVS/DFP* VSAM Hiperspace*

5.4.54 Multiprocessing

The multiprocessing feature permits a multiprocessing configuration. Multiprocessing provides CPU address identification, CPU signaling and response, prefixing, shared main storage, and TOD clock synchronization.

5.4.54.1 CPU Address Identification

CPU address identification provides an address by which each of the central processors can be identified by the Signal Processor (SIGP) instruction. It also provides new external interruption conditions and the Store CPU Address instruction by which the system control program can determine the address of a central processor.

5.4.54.2 CPU Signaling and Response

CPU signaling and response provides for communication among the central processors. This feature provides the Signal Processor (SIGP) instruction and the mechanism to interpret and to act on several order codes, such as sense, stop, and restart.

5.4.54.3 Prefixing

For each central processor, prefixing provides a means of assigning real addresses 0 through 4095 to different 4KB blocks of central storage. One area of central storage (represented by a single contiguous range of absolute addresses) is assigned to each central processor.

5.4.54.4 Shared Main Storage

Shared main storage permits all central processors to have access to common main storage locations.

5.4.54.5 TOD Clock Synchronization

TOD clock synchronization provides a uniform appearance to a clock synchronization program in 9672 CPCs allowing the program to be independent of the actual number of TOD clocks and central processors in a configuration. It includes a TOD clock synchronization control bit in control register 0.

5.4.55 Open Systems Adapter (OSA) Feature

The S/390 Opens Systems Adapter feature provides:

- Connectivity options that allow direct attachment of Ethernet, Token-Ring, and Fiber Distributed Data Interface (FDDI) LANs, and Asynchronous Transfer Mode (ATM)-based networks to the central processor complex (CPC)

The S/390 Opens Systems Adapter feature supports the CPC acting as a enterprise-wide server for multiple networks. LANs, LAN backbones, high speed workstations, hubs, bridges, and routers that conform to these LAN interface standards can connect directly to the S/390 Opens Systems Adapter, benefiting from the performance, connectivity, security, availability, and growth opportunities associated with the S/390 architecture.

New S/390 Opens Systems Adapter channel hardware, available as an orderable feature, provides the connectivity required for direct LAN attachment to the CPC. See 1.3.8, “S/390 Open Systems Adapter (OSA) Feature Support” on page 1-14.

5.4.55.1 OSA Support Facility for MVS/ESA

The OSA Support Facility for MVS/ESA is an application that operates as a started task under MVS/ESA. It provides configuration and operations management capabilities for the S/390 Opens Systems Adapters in the configuration and allows you to install supported application function onto S/390 Opens Systems Adapters to support client server applications integrated at the CPC. For additional information on OSA software requirements, see 1.3.8.2, “Support for the TCP/IP, SNA, and IPX Network Protocols” on page 1-14.

5.4.56 Page Protection

Page protection provides protection against improper storing. It controls access to virtual storage by using the page protection bit in each page table entry.

5.4.57 PER Extensions

The PER extensions are an augmentation of PER 2 that provide additional information about PER events and also, in the interpretive-execution mode, additional PER function.

5.4.58 Private Space

Private space provides a bit in the segment-table designation to prevent the use of TLB entries for common segments, and to prevent the application of low-address protection and fetch-protection override to the specified address space.

5.4.59 Processor Resource/Systems Manager

Processor Resource/Systems Manager (PR/SM) enables logical partitioning in the CPC and provides additional byte multiplexer channel capability.

5.4.59.1 Logical Partitioning

PR/SM enables 9672 and 9674 CPCs to be initialized for logically partitioned operation. 9672 and 9674 CPCs support up to ten logical partitions.

5.4.59.2 Multiple High-Performance Guests

PR/SM supports the VM/XA* System Product Enhancement for multiple preferred guests. It allows the support of multiple high-performance guests running concurrent with other virtual machines. Six V = F preferred guests are supported. When running a V=R preferred guest virtual machine, as many as five V=F preferred guests are supported. PR/SM supports devices dedicated to V=R and V=F guests. VMA, under SIE, supports V=R and V=F VM/SP and VM/SP HPO guests.

5.4.59.3 PR/SM LPAR Auto-Reconfiguration

This function in PR/SM is an automatic reconfiguration facility (ARF) that improves the operation of “hot standby” or backup logical partitions. When it is used with MVS/ESA, a logical partition can invoke the automatic reconfiguration facility which:

- Resets the primary logical partition
- Automatically reconfigures storage to the secondary logical partition
- Automatically deactivates the primary logical partition

All this occurs without operator intervention.

Do not use ARF with the I/O Interface Reset facility provided as an option on the Customize Options panel in the Reset profile.

5.4.59.4 PR/SM LPAR Dynamic Storage Reconfiguration

Dynamic storage reconfiguration provides logical partition definitions that allow a logical partition to attach storage (both central and expanded storage) to its configuration when another logical partition releases its storage and to detach storage from its configuration while MVS remains active. For example, central and expanded storage detached from a logical partition can be used to activate another logical partition.

Hardware/Software Corequisites: MVS/ESA and MVS/SP 2.2 (and later releases) support dynamic storage reconfiguration in ESA/390 logical partitions.

See the appropriate preventive service planning (PSP) bucket subset ID for hardware and software maintenance information.

5.4.59.5 PR/SM LPAR Management Time Reporting

PR/SM LPAR processor utilization reporting has improved the Resource Measurement Facility (RMF) partition data report. The improved report provides LPAR management time information and should help you understand PR/SM's low utilization effects in LPAR environments. It also makes it easier for you to make better capacity planning decisions due to more precise information in the processor resource utilization report.

Software Corequisites: The Resource Measurement Facility (RMF) requires APAR and PTF upgrades to support PR/SM LPAR support for 10 logical partitions. See the appropriate PSP bucket for the necessary APAR and PTF information.

5.4.59.6 PR/SM LPAR Processor Weight Management

For 9672 and 9674 CPCs, PR/SM LPAR enforces LP processing weights as follows:

- For LPs with processor resource capping, PR/SM LPAR now 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.

- 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.
- 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.

5.4.59.7 PR/SM LPAR Dispatching Enhancement (For 9672 R2/R3 and 9674 C02/C03)

For EMIF-capable processors, PR/SM management time is decreased as a result of improved efficiency in the PR/SM LPAR dispatcher. The system can now accommodate additional logical central processor (CP) dispatching. As a result, slight increase in the MVS Task Control Block (TCB) and Service Request Block (SRB) times may be observed due to increased CP resource contention in pre-MVS/ESA SP Version 4 Release 3 environments for systems with more than one physical CP.

MVS/ESA SP 4.3 Benefits in a Shared LPAR Environment: MVS/ESA 4.3 provides significant benefits when operating in a shared environment for 711-based processors (new 9672 processors; older ones with a PTF to MVS):

- A reduction in PR/SM LPAR low utilization effect (LUE) due to enhancements in MVS CP management.
- Addtional reduction in the PR/SM LPAR management time beyind that provided by the PR/SM LPAR dispatcher enhancement.

The reduction of PR/SM LPAR LUE in support of MVS 4.3 n-way LPs provides a more accurate reading of the capacity required to support MVS workloads. LUE made it difficult to plan for future workload changes because the CP time required by the existing workload could not easily be determined. The reduction in LUE and the reduction in LPAR management time adds capacity by reducing the CP resources used to manage the processor complex.

5.4.59.8 PR/SM LPAR Support for Logical Central Processor (CP) Vary

PR/SM LPAR logical CP vary allows logical partitions (LPs) to reallocate central processors (CPs).

PR/SM LPAR logical CP vary has the following aspects:

- An activated LP no longer needing a particular CP can send the appropriate system control program (SCP) command to release it, making it available for use by other LPs.
- PR/SM LPAR automatically adjusts the processing shares whenever you vary online or vary offline a logical CP in a shared LP. The share of CP resource for each ***online*** logical CP in the LP is calculated by dividing the share for each LP by the number of ***online*** logical CPs in the LP.
- PR/SM LPAR will no longer require LP deactivation for an LP to restore the capacity associated with a dedicated CP that was previously lost due to CP checkstop. By using the appropriate MVS or VM SCP command, the LP can claim access to another CP if that resource is physically available and ***not*** required by other LPs.
- For dedicated LPs, logical CP vary allows you to vary online a CP from the shared CP pool to the dedicated LP to replace a physical CP lost due to CP checkstop. This may be useful when you wish

to restore the CP capacity to the LP while it continues to run. Previously, to restore CP capacity in this situation, you would have had to deactivate the LP, reactivate it, then re-IPL its operating system.

5.4.59.9 PR/SM LPAR Preferred Path

You can specify preferred channel paths for devices in LPAR mode as well as basic mode. Use IOCP to define a preferred path for a device by coding the PATH parameter in the IODEVICE macroinstruction. For more information, see *ES/9000 Input/Output Configuration Program User's Guide and ESCON Channel-to-Channel Reference*, GC38-0401.

5.4.60 PSW-Key Handling

PSW-key handling provides the Set PSW Key from Address (SPKA) and Insert PSW Key (IPK) instructions.

5.4.61 Recovery Extensions

Recovery extensions consist of:

- The clear channel function in a channel, which can be used to perform an I/O system reset in a channel when the Clear Channel (CLRCH) instruction is executed
- Machine-check extensions, which include a machine-check external damage-code validity bit and provide a detailed indication of the cause of external damage
- Limited channel logout extensions, which consist of two additional logout bits, to indicate whether the I/O interface is operative and whether the logout is valid.

5.4.62 Scalar Square Root

An enhanced scalar execution element design provides new square root instructions and a faster divide instruction.

5.4.63 SCP-Initiated Reset

See 5.4.46, “I/O Interface Reset” on page 5-18.

5.4.64 Service Signal

Service signal provides an external interruption that is used by the processor console to signal information to the control program.

5.4.65 Set Address-Space Control Fast (SACF)

This instruction is used to change the address space mode setting the same as had been done with the Set Address-space Control (SAC) instruction. However, SACF does not perform serialization and check-point synchronization functions as the SAC instruction does. This can result in improved performance in situations where there is a need to extensively use the Set Address-space Control function.

5.4.66 Sorting Instructions

Sorting instructions are used by IBM Data Facility Sort (DFSORT) Release 7 and subsequent releases, running under MVS. The sorting instructions are used by DFSORT when sorting fixed-length records using the block-set sorting technique.

5.4.67 Storage-Key Instructions

The storage-key instructions Set Storage Key (SSK) and Insert Storage Key (ISK) allow initialization and inspection of the storage key associated with each block of storage that is available in the configuration.

5.4.68 Storage-Key Instruction Extensions

The storage-key instruction extensions provide the Set Storage Key Extended (SSKE), Insert Storage Key Extended (ISKE), and Reset Reference Bit Extended (RRBE) instructions, which provide 31-bit addresses and operate on the storage-key associated with a 4KB block of storage.

5.4.69 Storage-Key 4KB Block

Storage-key 4KB block allows a single key to be associated with each 4KB block of storage and, in System/370 mode, provides the storage-key exception control bit in control register 0.

5.4.70 Suppression on Protection

Suppression on Protection is a hardware feature available on 9672 CPCs that allows a protection exception condition due to a page protection to result in suppression of instruction execution rather than termination of execution. Information about the protected page is recorded in low storage and is available to the system when executing.

5.4.71 Subspace Group Facility Support

9672 CPCs provide support for the subspace group facility which can enhance the data integrity and reliability of application server subsystems like CICS/ESA, reducing application failures, service outages, and incorrect modification of critical business data.

Application server subsystems like CICS/ESA use a single application server program to provide common function to multiple application programs all running within a single address space. The subspace group facility allows the application server program to assign address space private storage to each application program running in the address space. This prevents one application program from accidentally overwriting the code and data of another application program running in the same address space. Preventing accidental overwrites reduces application failures, service outages, and incorrect modification of critical business data.

Additionally, there should be little or no modification required of existing application programs because you implement storage isolation and protection through the subsystem application server program.

5.4.71.1 CICS/ESA Use of the Subspace Group Facility

CICS/ESA uses the subspace group facility to implement transaction isolation to prevent one application program from accidentally overwriting the storage of another application program running in the same address space. Transaction isolation extends the subsystem storage protection capability which protects CICS system storage from being accidentally overwritten by user-written application programs.

Additionally, transaction isolation use of the subspace group facility enhances application development and maintenance for the subsystem. When an application attempts to access the address space private storage of another application, CICS/ESA abends the program making it easy to identify the failing application that could cause a data integrity problem.

Software Corequisites: Use of the subspace group facility and CICS/ESA transaction isolation requires the following minimum level of software support:

- MVS/ESA SP Version 5 Release 1 (5.1)
- CICS/ESA Version 4 Release 1 (4.1)

5.4.72 Subsystem Storage Protection

Subsystem storage protection provides a mechanism for subsystems to isolate their code, data, and control blocks from application code running in the same address space while still enabling access to the application code.

Subsystem storage protection requires MVS/ESA SP Version 4 Release 2.2 or later.

CICS/ESA* Version 3 with the CICS/ESA storage protection function utilizes subsystem storage protection to:

- Enhance user productivity by preventing outages caused by application software overwriting CICS/ESA system software
- Reduce the system management time required to diagnose errant programs that could potentially overwrite CICS system storage
- Allow most current programs to make use of this facility without software changes, protecting customer investment.

For the system programmer: Examine the SCHED*xx* member of SYS1.PARMLIB that your site uses to ensure that you do not attempt to add any programs to the program properties table (PPT) with program protect key 9. If you attempt to assign key 9, the system will assign the default value of key 8 to that program, and the results will be unpredictable.

5.4.73 IBM Sysplex Timer

The Sysplex Timer (required by the 9672 Parallel Transaction Server and 9672 Parallel Enterprise Server R1/R2/R3 Models with coupling facility channels: optional for the 9672 Parallel Enterprise Server R1/R2/R3 Models without coupling facility channels and the 9674 Coupling Facility C01/C02/C03 Models) synchronizes the processor time-of-day clocks of the systems in a complex and is supported by MVS/ESA SP Version 4.1 and subsequent releases. The Sysplex Timer provides an external time source for attached processors when sharing data or workloads.

5.4.73.1 External Time Source

The external time source function of the Sysplex Timer can improve the precision of the Sysplex Timer and compensate for the effects of long-term TOD clock drift.

This function allows the Sysplex Timer to:

- Receive timing signal broadcasts from a stable external time source
- Adjust periodically to the time provided by the external time source

5.4.73.2 Automatic Propagation Delay Adjustment

The automatic propagation delay adjustment function makes cable length effects nearly transparent to processor time-of-day (TOD) clock synchronization.

This function removes the restriction that cables to the attached CPCs be equal in length by automatically adjusting signals to compensate for propagation delay.

5.4.74 System/370 Extended Facility

The non-MVS-dependent portion of the System/370 extended facility consists of:

- Low-address protection, which improves system integrity by providing special protection for storage (at fixed storage addresses 0 through 511) that is vital to the system control program
- Invalidate Page Table Entry (IPTE) instruction and the common-segment bit, which increase the efficiency of dynamic address translation
- Test Protection (TPROT) instruction, which performs tests for potential protection violations without causing program interruptions for protection exceptions.

The MVS-dependent portion of the System/370 extended facility consists of:

- Supervisor Call (SVC) Assist instruction, which improves central processor performance by reducing the time needed to enter MVS supervisory services
- Fix Page instruction, Add FRR (Add Functional Recovery Routine) instruction, six tracing instructions, and four lock-handling instructions, which improve central processor performance.

5.4.75 Test Block

Test block provides the Test Block (TB) instruction for testing the usability of a 4KB block of central storage.

5.4.76 Time-of-Day Clock

The time-of-day (TOD) clock for each central processor provides a consistent measurement of elapsed time that can be used for indicating the time of day. The TOD clock for each central processor is initialized by the Set Clock (SCK) instruction by a central processor.

- Bit 51 increments at 1-microsecond intervals.
- Bits 52–55 are monotonic to ensure 1-microsecond counting in bit 51.
- Bits 61–63 contain the central processor address.

5.4.77 Tracing

Tracing provides three aids for problem-program analysis:

- Address-space-number (ASN) tracing
- Branch tracing
- Explicit tracing

5.4.78 Translation

Translation includes the following features:

- Dynamic address translation
- Extended control (EC) mode
- Program-event recording 2 (PER 2)
- Set-system-mask suppression
- Store status

As part of these features, translation also provides the following instructions:

- Load Read Address (LRA)
- Purge Translation Lookaside Buffer (PTLB)

- Reset Reference Bit (RRB)
- Store Then AND System Mask (STNSM)
- Store Then OR System Mask (STOSM)

5.4.78.1 Dynamic Address Translation

Dynamic address translation (DAT) provides hardware translation of virtual addresses to real addresses during program execution. DAT supports real storage sizes according to the amount of storage available for each 9672 CPC.

9672 CPCs use 4KB pages and either 64KB segments or 1-MB segments. ESA/390 mode uses only the 1-MB segment size. A System/370 guest virtual machine under control of interpretive execution can use both segment sizes.

5.4.78.2 Extended Control Mode

System/370 includes extended control (EC) mode, in which virtual storage and high-speed DAT are available. Operation in ESA/390 mode is comparable to operation in EC mode.

5.4.78.3 Program-Event Recording 2 (PER2)

Program-event recording 2 (PER 2) aids in debugging programs. During program execution, PER 2 can monitor the following actions:

- Successful branches
- Instruction fetches from a specified storage area
- Alteration of a specified storage area
- Branches to a specified storage area
- Alteration of storage within specified address spaces

Other extensions provide additional information about PER 2 events and, in the interpretive-execution mode, additional PER 2 function.

5.4.78.4 Set-System-Mask Suppression

Set-system-mask suppression permits suppression of execution of the Set System Mask (SSM) instruction and provides the special-operation program interruption code.

5.4.78.5 Store Status

Store status is an operator-initiated function that places the contents of the current PSW and the program-addressable registers in permanently assigned locations within the first 512 bytes of absolute storage. Store status also includes a noninitializing manual reset function.

5.4.79 Virtual-Machine Assists under SIE

Virtual-machine assist (VMA), which is an assist for VM/SP, directly executes certain privileged virtual-machine instructions and validates page-table entries in the shadow tables. VMA improves performance on virtual-storage system operation under VM/SP by reducing the amount of time VM/SP spends in the real supervisor state. The reduction is achieved by emulation (instead of software simulation) of certain privileged operation codes used by the virtual-storage (guest) control program. An interpretively-executed System/370-mode guest virtual machine also can benefit from the advantages offered by VMA.

5.4.80 3033 Extension

The 3033 extension provides the following facilities:

- Dual-address space
- Start-I/O-fast queuing
- Suspend and resume

5.4.80.1 Dual-Address Space

Dual-address space aids communication between virtual address spaces and provides:

- Twelve additional instructions
- Two address spaces for immediate use by a program
- A means of changing to other virtual address spaces
- A table-based subroutine linkage
- The use of multiple access keys for key-controlled by problem programs

5.4.80.2 Start I/O Fast Queuing

Start I/O fast queuing allows a Start I/O Fast Release (SIOF) instruction to complete execution independently of device selection or a channel-busy condition. Control unit or device busy conditions encountered before execution of a SIOF instruction cause the I/O operation to remain pending until facilities are available to initiate the operation at the device.

5.4.80.3 Suspend and Resume

Suspend and resume provides:

- The suspend flag in the channel command word (CCW), which indicates execution of a channel program is to be suspended
- A channel address word (CAW) bit, which controls whether the suspend flag of the CCW should cause suspension of execution of a channel program
- A channel status word (CSW) bit, which indicates execution of a channel program has been suspended
- The Resume I/O (RIO) instruction, which causes the execution of a suspended channel program to be resumed

5.4.81 31-Bit Indirect Data Address Word

The 31-bit indirect data address word (IDAW) extends the size of the address field in IDAWs to 31 bits.

5.4.82 31-Bit Real Addressing

The 31-bit real addressing feature ensures that certain fields contain 31-bit real addresses regardless of the setting of the addressing-mode control bit in the PSW.

Chapter 6. Coupling Facility

6.1 Overview

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 MVS images in the sysplex define the shared structures used in the coupling facility.

Enhancements to PR/SM LPAR allow 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 permitted 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 hardware operator console. PR/SM LPAR limits the hardware operator controls usually available for LPs to avoid unnecessary operator activity. However, Coupling Facility Control Code does provide a limited set of new hardware operator controls unique to the coupling facility LP. See 6.10, “Coupling Facility Control Code Commands” on page 6-8.

New 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 Chapter 9, “Coupling Facility Channels” on page 9-1.

6.2 Test or Migration Coupling Facility 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 MVS images connected to the coupling facility
- Running non-coupled production work

A single CPC configuration has the following considerations:

- Simultaneous loss of the coupling facility and any MVS images coupled to it (a more likely possibility in a single CPC configuration) can potentially cause extended recovery times
- The 9674 Coupling Facility C01/C02/C03 Models can provide better connectivity (up to 32 coupling facility channels). See 6.3, “Production Coupling Facility Configuration.”

You can define a test or migration coupling facility with or ***without*** coupling facility channel hardware. See 9.2, “Integrated Coupling Migration Facility (ICMF)” on page 9-3 for information on how to define a test or migration facility ***without*** coupling facility channel hardware.

6.3 Production Coupling Facility Configuration

IBM recommends that you run your shared data 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.

IBM recommends a production coupling facility configuration that consists of:

- One 9674 Coupling Facility C01/C02/C03 Models running as a single dedicated coupling facility LP to provide large capacity shared storage and maximum coupling facility channel connectivity (up to 32 coupling facility channels).
- A second 9674 Coupling Facility, similarly configured, to reduce the possibility of a single point of failure. A second 9674 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.

Additionally, the 9674 Coupling Facility C01/C02/C03 Models provide availability advantages that can enhance the nonvolatility of the shared structures stored in the coupling facility. See 6.4.1, “Nonvolatility Advantages of a 9674 Coupling Facility C01/C02/C03 Models.”

6.4 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 ensure this, you ***must*** provide an ***optional*** backup power supply to make coupling facility storage contents nonvolatile across utility power failures.

6.4.1 Nonvolatility Advantages of a 9674 Coupling Facility C01/C02/C03 Models

The 9674 Coupling Facility C01/C02/C03 Models offer advantages for nonvolatility of coupling facility storage.

6.4.1.1 For a 9674 model C01, you can install:

- An optional uninterruptible power supply (UPS)

  Optional UPS provides a secondary power source for use during extended utility power outages allowing continuous coupling facility operation.

- An optional battery backup feature

  The optional battery backup feature installed on 9674 C01 will preserve coupling facility storage across a utility power failure for short term outages. Additionally, the battery backup feature provides connection to an external 24 volt power supply that can extend the battery backup feature support. See 3.2.5, “Optional Battery Backup Feature” on page 3-8.

- Both an optional UPS and an optional battery backup feature

  In this case, the optional battery backup feature can provide an override capability during switchover from primary power to the optional UPS.

  Note: The optional battery backup feature can also be used for a coupling facility LP defined on a CPC that is part of a 9672 Parallel Transaction Server.

6.4.1.2 For all 9674 Coupling Facilities, you can install:

- An optional Local Uninterruptible Power Supply (LUPS)

  Optional LUPS provides 9 to 18 minutes of full power hold-up in the event of extended customer power line disturbances. See 3.3, “Local Uninterruptible Power Supply (LUPS)” on page 3-8.

  Note: The LUPS can also be used for a coupling facility LP defined on a CPC that is part of a 9672 Parallel Transaction Server.

6.4.2 Setting the Conditions for Monitoring Coupling Facility Nonvolatility Status

In addition to installing an optional backup power supply to 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 POWERSAVE (the default) determines the volatility status of the coupling facility based on whether or not:
  - A battery backup feature, UPS, or LUPS is installed (depending on the model).
  - The battery backup feature, UPS, or LUPS has batteries that are online and charged

  If battery backup, or a UPS, or a LUPS is installed (depending on the model) and its batteries are online and charged, Coupling Facility Control Code sets its status to nonvolatile. If either condition is ***not*** met, coupling facility volatility status is volatile.

  The battery backup feature, UPS, or LUPS will preserve coupling facility storage across a utility power outage following the expiration of a rideout interval when power save state is enabled for the CPC (see 6.5, “Power Save State for Coupling Facility LPs”). The default rideout interval is 10 seconds. You can change the rideout interval using the Coupling Facility Control Code RIDEOUT command.

  MODE POWERSAVE is useful only on CPCs that support power save state and have installed an operating optional battery backup feature, UPS, or LUPS (depending on the model).

  Note: Engineering change (EC) D57264 for 9674 model C01 CPCs and EC D57262 for 9672 E/P model CPCs provide support for power save state.

- MODE NONVOLATILE sets coupling facility volatility status to nonvolatile and should be used if a 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.

6.5 Power Save State for Coupling Facility LPs

For 9674 CPCs and 9672 CPCs (except Model R11), power save state is available to preserve coupling facility storage across a utility power failure. Power save state allows coupling facility structures to be saved intact. This simplifies the subsystem recovery and restart activity that would have been required had the structures been lost.

Enabling power save state for the CPC (the default setting) provides the following capability:

- Coupling facility logical partition (LP) operation shuts down and storage contents for coupling facility LPs are preserved across a utility power failure regardless of the volatility MODE setting of the cou-

pling facility LP. When utility power is restored, coupling facility LPs return to operation with their structures intact.

A CPC draws less power in power save state allowing coupling facility LP storage contents to be preserved for a longer period of time than if the CPC continued operation.

- Any non-coupling facility LP running on the same CPC as the coupling facility LP is automatically system reset when the CPC loses utility power. When utility power is restored, you must re-IPL the non-coupling facility LPs.

Power save state requires:

- An installed and charged battery backup feature on the CPC or an operational uninterruptible power source (UPS or LUPS), depending on the model.
- Enabling prior to power-on reset of the CPC (enabled is the default power save state setting).

If you ***disable*** power save state for a CPC that has coupling facility LPs activated or the CPC has no coupling facility LPs activated, the battery backup feature or UPS maintains CPC operation through the utility power failure as long as the backup power source remains operational. Continuous CPC operation is effective to the extent that the DASD used by the CPC survives the power outage.

You can enable and disable power save state for a CPC using the Power Save option on the Customize pulldown of the System Monitor panel. The default setting for the power save option is enabled.

6.6 Parallel Sysplex Recovery with Coupling Facilities that Use Power Save State

There are two possible scenarios for parallel sysplex recovery in configurations where one or more coupling facilities has power save state enabled when a utility power failure occurs.

6.6.1 Parallel Sysplex Recovery When No MVS Systems Survive Utility Power Failure

When a utility power failure causes all MVS systems in the configuration to fail and one or more coupling facilities to enter power save state, parallel sysplex recovery proceeds as follows when utility power is restored:

1. Coupling facilities that entered power save state resume operation with their structures intact.
2. MVS systems that are re-IPLed allow MVS to gain ownership of the operating coupling facilities with the new instance of the sysplex.
3. As MVS gains ownership of the operating coupling facilities, it cleans up non-persistent structures and connectors and places persistent connectors into the failed state.
4. As subsystems that use the coupling facility return to operation, they reconnect as new instances of the failed persistent connection and perform the appropriate recovery processing using the data, locks, or both as available in the coupling facilities.

6.6.2 Parallel Sysplex Recovery When One or More MVS Systems Survive Utility Power Failure

When one or more MVS systems survive a utility power failure, parallel sysplex recovery proceeds as follows:

1. Coupling facilities that enter power save state perform an I/O system reset. All attached MVS systems, if any, that are still operational following the utility power failure view the I/O system reset as a loss of connectivity to the coupling facility.
2. If the sysplex meets all the requirements for automatic initiation of structure rebuild due to loss of connectivity, MVS automatically triggers structure rebuild for all structures that resided in the coupling facilities that entered power save state.
   - If there are coupling facilities that remain operational following the power failure, MVS attempts to rebuild the structures in the operational coupling facilities. If the structures are successfully rebuilt, operational MVS systems continue to use the rebuilt structures in their new location. The structure in the coupling facility that entered power save state is discarded when that coupling facility resumes operation.
   - If there are no coupling facilities that remain operational following the power failure, structure rebuild is not possible. Connectors to structures disconnect and can become failed-persistent.
3. When utility power returns, MVS gains ownership of coupling facilities that return from power save state. Structures in these coupling facilities that have been successfully rebuilt in new coupling facility locations are discarded; the structure survived the power outage but is no longer needed. Structures in these coupling facilities that have ***not*** been successfully rebuilt in new coupling facility locations are kept, if persistent, or discarded, if non-persistent.

6.7 Coupling Facility LP Definition Considerations

You define a coupling facility logical partition (LP) by specifying new Reset and Image profile parameters. For 9674 CPCs and 9672 CPCs (except Model R11), you can define coupling facility mode for an LP at the Hardware Management Console using the Activation Profiles task available from the CPC Operational Customization Tasks Area.

You can define coupling facility LPs with dedicated or shared CPs. 9674 Coupling Facility C01/C02/C03 Models and 9672 Parallel Transaction Server support for defining coupling facility LPs as shared LPs is available with engineering change (EC) D79262.

Coupling facility LPs must be defined with at least 8MB of central storage and at least one CFR channel path targeted to be brought online at LP activation. Additional storage above 8MB must be defined to allow applications to define shared structures within the coupling facility LP. For more information, see *PR/SM Planning Guide*, GA22-7123.

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, reserved expanded storage, or a reserved storage origin (coupling facility LPs do ***not*** support dynamic storage reconfiguration)
- Asynchronous data mover facilities
- Automatic IPL
- Automatic IPL address
- Automatic IPL parameters

6.8 Coupling Facility LP Activation Considerations

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.

Parallel, ESCON, CFS, and OSA channel paths are prevented from being brought online to a coupling facility LP at LP activation.

Note: Any parallel, ESCON, CFS, or OSA channel path targeted to be brought online will automatically be configured online if the coupling facility LP is redefined as an ESA mode LP.

6.9 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 hardware console. PR/SM LPAR limits the hardware operator controls usually available for LPs to avoid unnecessary operator activity.

Coupling facility LPs do ***not*** support the following tasks typically available to the Hardware Management Console operator:

- Daily Tasks
  - Reset normal
- CPC Recovery Tasks
  - Start
  - Stop
  - Reset normal
  - PSW restart
  - Reset clear
  - Load

Additionally, coupling facility LPs do ***not*** support the following panel selections typically available to the operator at the CPC console:

- Monitor the System panel/Operate menu
  - Start all processors in a system
  - Stop all processors in a system
- Recover the System panel
  - Load normal
  - Load clear
  - Reset system normal
  - Reset system clear
  - Restart a processor
  - Store status
- Monitor Processors panel/Operate menu
  - Start a processor
  - Stop a processor
  - Restart a processor
  - Load normal
  - Load clear

  - Interrupt
- Monitor Processors panel/Analyze menu
  - Store status
  - Display or alter
  - Stop on an address match
  - Trace
- Monitor Channels panel/Operate menu
  - Configure on

    You can use the CONFIGURE Coupling Facility Control Code command instead.
  - Configure off

    Configure off is available but IBM recommends using the CONFIGURE Coupling Facility Control Code command instead.

6.10 Coupling Facility Control Code Commands

Coupling Facility Control Code does provide a limited set of new hardware operator controls unique to the coupling facility LP. These controls are available from the Operating System Messages panel.

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)
- HELP (display coupling facility control code command syntax)
- MODE (define coupling facility volatility mode)
- RIDEOUT (define rideout interval for POWERSAVE mode)
- SHUTDOWN (shutdown coupling facility operation)
- TIMEZONE (sets timezone offset from Greenwich Mean Time for a coupling facility)

Note: Support for the CP and HELP coupling facility control code commands is available with engineering change (EC) D57264 for 9674 CPCs and EC D57262 for 9672 CPCs.

6.11 MVS/ESA SP 5.1 Considerations

MVS images that attach to a coupling facility must run MVS/ESA SP 5.1 or higher. To define how to manage MVS 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.

- Logical partition (LP) information for the coupling facility

  The LP numbers used to identify LPs in IOCP input files must match the numbers in the PARTITION keyword in the CF statement in the CFRM policy. IBM also recommends that the LP names for the 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 information in either the IOCP or HCD reports.

6.12 MVS/ESA SP 5.2 Considerations

6.12.1 Dynamic Reconfiguration of Coupling Facility Structures

Beginning with MVS/ESA SP 5.2 and supported at engineering change (EC) D57264 for 9674 CPCs and EC D57262 for 9672 CPCs, MVS supports dynamic reconfiguration of coupling facility structures.

Dynamic reconfiguration of coupling facility structures, also called structure alteration, provides improved availability by allowing allocated cache, list, and lock structures in a coupling facility logical partition (LP) to:

- Dynamically expand or contract without requiring MVS to deallocate the structure first.
- Dynamically reapportion the ratio of entries to elements for a structure while the structure remains allocated. This allows for dynamic tuning of structures providing more effective use of coupling facility LP storage resources.

Subsystems can now specify optional IXLCONN macro keywords that:

- Allow or disallow structure alteration (ALLOWALTER=NO is the default)
- Allow or disallow structure reapportionment (RATIO=YES is the default) if structure alteration is allowed
- Set minimum percentages of entries (MINENTRY) and elements (MINELEMENT) that should exist following dynamic structure reapportionment, expansion, or contraction if structure alteration is allowed

A structure can expand up to the maximum storage size you specify based on structure usage and coupling facility storage availability. Use the SIZE keyword in the coupling facility resource management policy (CFRM) to specify maximum structure size. Additionally, you can optionally specify an initial structure size (new INITSIZE keyword) that is smaller than the maximum structure size.

6.12.2 List Processing Enhancements

Beginning with MVS/ESA SP 5.2 and supported at EC D57264 for 9674 CPCs and EC D57262 for 9672 CPCs, there are list structure processing enhancements available in a level 1 coupling facility accessible through XES. New IXLLIST parameters provide these new functions. For more information, see the *Sysplex Services Guide*, GC28-1495.

6.12.3 Support for 255 Data Elements Per Entry in a Coupling Facility Structure:

Beginning with MVS/ESA SP 5.2 and supported at EC D57264 for 9674 CPCs EC D57262 for 9672 CPCs, data entries can contain up to 255 data elements. Though the total data size for an entry remains limited to a maximum of 64KB, a subsystem using the coupling facility can now segment the entry data across a greater number of smaller data elements. This allows greater control of the granularity of storage allocation and can provide for more efficient use of storage in the coupling facility structure.

Chapter 7. Characteristics of ESCON and Parallel Channel Operation

7.1 Overview

The channels in the channel subsystem permit transfer of data between central storage and I/O devices under control of a channel program. The channels act independent of other operations being performed by a central processor. A central processor is therefore free to resume other operations after initiating an I/O operation.

The channel facility required to perform an I/O operation is called a subchannel.

The control and implementation of I/O operations depend on:

- Channel Hardware Technology
- Data mode of operation
  - Block multiplexer
  - Byte multiplexer
- I/O Operation Controls
- I/O interface protocol
  - ESCON channels using serial transmission, link-level and device level
  - ESCON channels attached to a 9034 and parallel channels
    - Interlocked
    - Data streaming

7.2 Subchannels

One unit control word (UCW) is assigned to each subchannel. UCWs are stored in the hardware system area of central storage and are moved to a channel subsystem during execution of an I/O operation.

Each UCW contains the following control information:

- Protection key
- Data address
- Identity of the operation specified by the command code
- CCW flags
- Byte count
- Channel status
- Address of the next CCW
- Channel path identifiers for as many as eight channel paths that may be configured to an I/O device

In ESA/390 mode, one UCW is assigned for each subchannel (device).

9672 Central Processor Complexs (CPCs) provide panels that display the contents of a specified sub-channel or UCW.

7.3 Channel Hardware Technology

7.3.1 ESCON Channel

The ESCON Channel provides a high-speed connection between host and control units for I/O devices. It is a channel that uses fiber optic cabling, which is thinner and weighs less than copper cabling and eliminates the need for "raised floor". Fiber optic cabling also reduces electromagnetic emissions, which provides for improved security.

The ESCON Channel is available in two versions:

- ESCON Channel, which provides a light-emitting diode (LED) light source for fiber optic cables. It can extend up to three kilometers (1.86 U.S. miles), a range that can be further extended to six or nine kilometers (km) by retransmission through one or two ESCON directors.
- ESCON Channel with the **Extended Distance Feature (XDF)**, available only on 9672 E, P, and R1 models, which provides a laser light source for fiber optic channels allowing for greater distances. It can support unrepeated link distances up to 20 km (12.4 U.S. miles). In combination with ESCON directors, interconnection distances may be up to 60 km, depending on the devices attached.

With the ESCON Channel, "channel-attached" achieves a new meaning. Devices directly attached to the processor no longer need to be in close range. With fiber links, the channel-to- channel and channel-to-control-unit distances can be up to 20 kilometers. With the signal repeated through ESCON directors, the channel-to-channel distances can be up to **60** kilometers. Display users within a **43** kilometer range can be connected through a channel link instead of a telecommunications link. The IBM 3490/3490E Tape range can extend up to **23** kilometers, and IBM 3990 or 9343 DASD Storage Controllers can be located up to **9** kilometers away.

The use of the ESCON Channel in combination with ESCON directors saves space and relocation costs - less space is required to house cabling for any distance; the same cabling can be shared by multiple channels and control units interconnected with an ESCON director; the total length of cabling will be considerably reduced; lighter and more easily installed fiber optic cabling is easier to relocate if necessary.

You can also use ESCON Channel hardware and cabling, in combination with an ESCON converter, to extend the range of a Parallel Channel. This capability allows you to mix your current data processing environment with the new ESCON environment. You can continue to use existing control units and copper cables while you evolve to fiber optic cables and ESCON control units as your requirements for enterprise growth increase.

7.3.2 Parallel Channel

The parallel channel can be used in ESA/390 and in S/370 modes of operation. It provides byte as well as block multiplex capability with a data rate of up to 4.5 megabytes per second in block multiplex mode. The parallel channel is an electrical channel that uses copper cabling to transmit data and control information between host channel and control units for I/O devices. It extends up to 122 meters for the attachment of control units, a range that can be extended by parallel channel extenders, subject to restrictions imposed by the types of I/O devices attached.

A major advantage of the parallel channel is that it allows customers who are migrating/upgrading from installed 4300, 9370, or 9221 systems to continue to use most of their installed channel-attached I/O devices, including devices attached to byte multiplexer channels.

7.4 Data Mode of Operation

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

7.4.1 Block Multiplexer Mode of Operation

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

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

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

7.4.2 Byte Multiplexer Mode of Operation

Byte interleave mode of operation allows the execution of multiple I/O operations concurrently. Each addressed device is selected, one at a time, for transfer of a byte or a group of bytes to or from central storage. Bytes from multiple devices are interleaved on the channel and routed to or from the desired locations in central storage.

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

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

Note: 9672 CPCs with engineering change (EC) D57262 will be able to define and use CBY channel paths to connect to 3720, 3725, and 3745 control units requiring byte multiplexer channel path attachment. Connection is through a 9034 ESCON Converter Model 1.

See the 9672IO subset id in the 9672DEVICE upgrade id of the preventive service planning (PSP) bucket for prerequisite 9034 EC level information.

7.4.3 Byte Multiplexer Mode and Burst Mode

A byte multiplexer channel can be monopolized by one I/O device (burst mode) or shared by many I/O devices (byte multiplexer mode). The number of bytes transferred at a time in byte multiplexer mode can be one (single byte transfers) or more than one (multibyte transfers). Most control units that operate in byte multiplexer mode can also operate in burst mode. A manually set switch at the control unit determines whether the control unit operates in burst mode or byte multiplexer mode.

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

Because most of the time spent in a data-transfer control sequence is for control, increasing the burst size (the number of bytes transferred per sequence) results in a relatively small increase in the total channel busy time for the sequence.

Also, increasing the burst size reduces the number of data transfer sequences required. The net effect is a significant improvement in channel efficiency and a higher allowable data rate.

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

If the device configuration guidelines are followed for byte multiplexer channels on an 9672 CPC, deferred accesses are minimized and data transfer sequences exceeding 32 microseconds are acceptable when large burst sizes are specified.

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

7.5 I/O Operations Control

I/O operations over the interface are controlled by channel commands and mode-dependent I/O instructions, chaining operations, and I/O interruptions.

7.5.1 Channel Commands

The six basic channel commands are:

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

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

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

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

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

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

7.5.2 ESA/390 Mode I/O Instructions

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

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

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

Start Subchannel (SSCH)
Starts execution of a channel program with the I/O device associated with the specified subchannel.

Test Subchannel (TSCH)
Checks subchannel status and clears the subchannel control bits.

Clear Subchannel (CSCH)
Clears the subchannel and signals the channel subsystem to perform the clear function at the associated I/O device.

Halt Subchannel (HSCH)
Terminates the current operation at the specified subchannel and signals the channel subsystem to perform the halt function at the I/O device.

Resume Subchannel (RSCH)
Signals the channel subsystem to resume execution of a suspended channel program with the I/O device associated with the specified subchannel.

Store Subchannel (STSCH)
Stores control and status information about the specified subchannel.

Modify Subchannel (MSCH)
Allows the control program to influence the execution of path management functions and some basic I/O functions.

Test Pending Interruption (TPI)
Stores the interruption code for a pending I/O interruption and clears the interruption request.

Set Address Limit (SAL)
Sets the address limit used in a comparison with the absolute storage address each time central storage is accessed for I/O data.

Reset Channel Path (RCHP)
Initiates a reset of the specified channel path.

Set Channel Monitor (SCHM)
Activates and deactivates the channel monitoring modes (measurement-block update and device-connect-time measurement).

Store Channel Path Status (STCPS)
Identifies what channel paths are being used when STCPS is executed.

Store Channel Report Word (STCRW)
Stores error-related information about a malfunction affecting the channel subsystem.

Cancel I/O (XSCH)
Withdraws a pending start function from the designated subchannel without signaling the device.

In ESA/390 mode, all I/O instructions except the SAL, SCHM, and STCPS instructions set the PSW condition code. This mode uses interruption response blocks rather than channel status words for returning interruption status to the program.

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

7.5.3 Chaining Operations

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

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

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

7.5.3.1 Data Chaining

When the data transfer specified by the current CCW is finished, data chaining causes the operation to continue by fetching a new CCW and using the storage area defined by the new CCW. Execution of the operation at the I/O device is not affected.

7.5.3.2 Command Chaining

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

7.5.4 I/O Interruptions

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

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

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

- Interface control check (IFCC); for example, interface tag errors and time-outs
- Channel control check (CCC); for example, parity, decode, or control errors
- Channel data check (CDC); for example, a parity error detected in central storage

Machine-check interruptions include the following:

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

7.5.4.1 Resets

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

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

An I/O system reset occurs as part of:

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

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

7.6 I/O Interface Protocol

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

7.6.1 Channel-to-Channel Connection

7.6.1.1 ESCON Channels (Attached to a 9034) and Parallel Channels

Parallel channel-to-channel connection between multiple systems is accomplished by using the IBM 3088 Multisystem Channel Communication Unit (MCCU) Model A1, 1, or 2. For channel-to-channel communication between an ESCON channel and parallel channels, the ESCON channel can be extended by fiber optic cable to a 9034 ESCON Converter Model 1 that is connected to a 3088.

Parallel channel-to-channel connection between a 9672 CPC and other IBM Central Processor Complexs is accomplished by using the channel-to-channel adapter (CTCA) feature on those CPCs that offer it, or by using the 3088.

Both data-streaming and interlock modes are standard on the 3088. Data-streaming mode provides for data transfers of as many as 4.5 million bytes per second, independent of cable length. Cable distances of 122 meters (400 feet) between the CPC and the 3088 are supported in both data-streaming and interlock modes.

The 3088 Model A1 provides two-CPC connectivity and as many as 63 logical CTCA links. The 3088 Model A1 can be field upgraded to a 3088 Model 1 or Model 2. The 3088 Model 1 can interconnect as many as four channels and can provide the equivalent function of as many as 126 CTCAs. The 3088 Model 2 can interconnect as many as eight channels.

ESCON CTC Support: Parallel channel-to-channel (CTC) communication makes use of the parallel I/O CTC architecture. The parallel I/O CTC architecture defines two operating modes for CTC communication:

basic mode and extended mode. ESCON CTC support for both of these modes is available. By replacing parallel CTC connections with ESCON CTC connections, you can remove the 3088 hardware requirement. See 7.6.1.2, “ESCON Channels (Using Link-Level and Device-Level Protocols).”

7.6.1.2 ESCON Channels (Using Link-Level and Device-Level Protocols)

You can achieve ESCON channel-to-channel connections between a 9672 CPC and other IBM CPCs with ESCON channels if one of the ESCON channels is defined to operate in channel-to-channel (CTC) mode.

9672 CPCs support ESCON CTC connections that support CTC applications operating in parallel CTC basic mode. This can assist in the migration of parallel CTC applications to an ESCON environment and remove the 3088 hardware requirement.

ESCON channels that operate in CTC mode (extended mode or basic mode) can be defined as shared ESCON channels. For more information, see Chapter 8, “ESCON Multiple Image Facility (EMIF)” on page 8-1.

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

7.6.2 Parallel Channel Implementation

Each I/O interface can attach as many as eight control units and address as many as 256 I/O devices. As many as 16 control units can be attached to an I/O interface using a switching unit (such as an IBM 3814 Switching Management System).

Multiplexing refers to the capability a channel and device have to disconnect and reconnect during an operation. Block multiplexing takes place between blocks of data, and byte multiplexing takes place between either bytes of data or groups of bytes of data.

Channel time-out functions and device priority described in the following sections apply only to ESCON channels attached to a 9034 ES Connection Converter and parallel channels.

7.6.3 Channel Time-Out Functions

The optional time-out function described here applies only to ESCON channels that attach to a 9034 ES Connection Converter and parallel channels.

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

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

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

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

7.6.4 Control Unit (CU) Priority on an I/O Interface

CU priority on an I/O interface applies only to ESCON channels attached to a 9034 ES Connection Converter and parallel channels.

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

Figure 7-1. Control Unit (CU) Priority on Parallel Channels

7.6.5 Dynamic Reconnection

In ESA/390 mode, the channel subsystem permits dynamic reconnection of I/O devices that have the dynamic-reconnection feature installed and that are set up to operate in a multipath mode, such as the IBM 3390 Direct Access Storage Model A14 or A22. Dynamic reconnection allows the device to reconnect and continue a chain of I/O operations using the first available channel path (one of as many as eight possible channel paths defined in an IOCP parameter). The selected path is not necessarily the one used initially in the I/O operation.

Chapter 8. ESCON Multiple Image Facility (EMIF)

8.1 Overview

With ESCON Multiple Image Facility (EMIF), multiple logical partitions can directly share channels and optionally share any of the control units and associated I/O devices configured to these shared channels. Enhancements to the channel subsystem (CSS) allow the CSS to be operated as multiple logical CSSs when the CPC is operating in LPAR mode so that one logical CSS is automatically configured to each logical partition. ESCON Multiple Image Facility also provides a way to limit the logical partitions that can access a reconfigurable channel or a shared channel to enhance security.

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

With ESCON Multiple Image Facility, you can configure channels as follows:

ESCON (TYPE=CNC or TYPE=CTC), Coupling Facility Sender (TYPE=CFS), Open Systems Adapter (TYPE=OSA): You can configure a channel path as:

- An unshared dedicated channel path to a single logical partition
- An unshared reconfigurable channel path that can be configured to only one logical partition at a time but which can be dynamically moved to another logical partition by channel path reconfiguration commands
- A shared channel path that can be concurrently used by the logical partitions to which it is configured.

Coupling Facility Receiver (TYPE=CFR): You can configure a channel path as:

- An unshared dedicated channel path to a single logical partition

ESCON (TYPE=CVC and TYPE=CBY) and Parallel Channels: You can configure a channel path as:

- An unshared dedicated channel path to a single logical partition
- An unshared reconfigurable channel path that can be configured to only one logical partition at a time but which can be dynamically moved to another logical partition by channel path reconfiguration commands

Neither ESCON (TYPE=CVC or TYPE=CBY), coupling facility receiver (TYPE=CFR), nor parallel channels (TYPE=BL or TYPE=BY) can be configured as shared channels.

With ESCON Multiple Image Facility, shared channel paths can provide extensive control unit and I/O device sharing. ESCON Multiple Image Facility allows all, some, or none of the control units attached to shared channels to be shared by multiple logical partitions limited only by the sharing limitations of the attached control units.

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

ESCON Multiple Image Facility enhances the following controls:

- Input/output configuration program (IOCP)
- Hardware Configuration Definition* (HCD)
- Logical path management
- Hardware system operations

8.1.1 ESCON Multiple Image Facility Software Corequisites

Any operating system that ***fully*** supports the ESCON architecture can use ESCON Multiple Image Facility. Additionally, non-ESCON operating systems can still operate on ESCON Multiple Image Facility-capable CPCs.

Note: VSE Version 1 Release 3 (1.3) is the first VSE release that ***fully*** supports the ESCON architecture.

Additionally, there is a new stand-alone input/output configuration program (IOCP) for EMIF-capable CPCs, IZP IOCP. Operating system program temporary fixes (PTFs) will introduce the corresponding batch versions of IZP IOCP for MVS, VM, and VSE/ESA*. For operating system releases that support dynamic I/O configuration, see 5.4.25, “Dynamic I/O Configuration” on page 5-14.

Note: EMIF is available on 9672 CPCs. On 9674s, you can only define CFR channel paths which cannot be shared.

8.2 ESCON Multiple Image Facility Functional Summary

ESCON Multiple Image Facility provides the capability to share channels connected to control units among multiple logical partitions for CPCs operating in LPAR mode and has implications for PR/SM LPAR resource sharing and device (or data) sharing.

8.2.1 PR/SM LPAR Resource Sharing

Prior to ESCON Multiple Image Facility, logical partitions required either unshared dedicated or unshared reconfigurable channel resources. With ESCON Multiple Image Facility, logical partitions can share channels. Optionally, multiple logical partitions can share the control units and I/O devices attached to shared channels.

8.2.2 Device (or Data) Sharing

Though device sharing was available prior to ESCON Multiple Image Facility, the only way to accomplish this was to define a separate dedicated or reconfigurable channel to each control unit from each logical partition wishing to share the associated devices. With ESCON Multiple Image Facility, logical partitions can share the same device through a single shared channel or set of shared channels.

8.3 IOCP Considerations

IZP IOCP supports 9672 and 9674 CPCs and ESCON Multiple Image Facility.

Note: The 9674 runs the same IZP IOCP used by CPCs capable of sharing channels. However, 9674s can only define CFR channel paths which cannot be shared.

Enhancements to IOCP allow you to define controls for ESCON Multiple Image Facility-capable CPCs. These changes include changes to the way you define logical partitions, channel paths, and I/O devices.

8.3.1 Logical Partition Definition

For ESCON Multiple Image Facility LPAR mode IOCDSs, you can specify an optional RESOURCE macroinstruction. You can use the RESOURCE macroinstruction to define all logical partitions in the configuration and assign a partition number to each logical partition. If you do not specify logical partition numbers using the RESOURCE macroinstruction, IZP IOCP assigns them.

The RESOURCE macroinstruction is required when you define shared ESCON channel paths for ESCON CTC communications and when you define coupling facility resource management (CFRM) policies for parallel sysplex configurations. See 6.11, “MVS/ESA SP 5.1 Considerations” on page 6-8.

8.3.2 Channel Path Definition

For ESCON Multiple Image Facility LPAR mode IOCDSs, you can define shared channel paths in addition to dedicated and reconfigurable channel paths. The CHPID macroinstruction has a new SHARED keyword to accomplish this. You can define:

- All channel paths as dedicated
- All channel paths ***except*** CFR channel paths as reconfigurable
- Only CNC, CTC, CFS, and OSA channel paths as shared

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

The access list parameter specifies the logical partition or logical partitions that will have the channel path configured online at logical partition activation following the intial power-on reset of an LPAR IOCDS. For exceptions, see *ES/9000 ES/3090 Processor Resource/Systems Manager Planning Guide*, GA22-7123.

The candidate list parameter specifies the logical partitions that can configure the channel path on-line.

Additionally, the candidate list provides security control by limiting the logical partitions that can access shared or reconfigurable channel paths.

Note: PR/SM LPAR manages the channel path configuration across POR. See *ES/9000 ES/3090 Processor Resource/Systems Manager Planning Guide*.

8.3.3 I/O Device Definition

For ESCON Multiple Image Facility LPAR mode IOCDSs, you can specify either the optional PART|PARTITION keyword or the optional NOTPART keyword on the IODEVICE macroinstruction to limit device access by logical partitions for devices assigned to shared ESCON or OSA channels.

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

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

  The system establishes logical paths between a logical partition and an ESCON channel path only when you define one or more of the devices associated with the ESCON control unit as accessible to the logical partition by using the IODEVICE PART|PARTITION or NOTPART keyword.

8.4 Hardware Configuration Definition (HCD) Considerations

Hardware Configuration Definition (HCD) support for ESCON Multiple Image Facility is available beginning with MVS/ESA SP 4.3.

For MVS, HCD provides the capability to make both dynamic hardware and software I/O configuration changes.

HCD also provides:

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

Enhancements to HCD allow you to define controls for ESCON Multiple Image Facility-capable CPCs. These include changes to the way you define logical partitions, channel paths, and I/O devices.

Changes include several new panels and new fields on previously defined panels. Also, there are changes to the way the FORCE parameter is implemented.

8.4.1 Enhanced HCD Panels to Support EMIF

The following HCD panels support ESCON Multiple Image Facility:

Add Partition
Allows explicit definition of logical partitions and associated logical partition numbers.

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

Define Candidate List (for channel paths)
Allows definition of candidate list for channel path access control of shared and reconfigurable channel paths.

Define Candidate List (for devices)
Allows definition of candidate list for device access control for devices assigned to shared channels.

Add Processor
Allows you to determine whether a CPC is ESCON Multiple Image Facility-capable.

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

Define Device / Processor
New field to specify candidate list.

8.5 Logical Path Establishment Considerations

For ESCON channels (both shared and unshared TYPE=CNC and TYPE=CTC), ESCON Multiple Image Facility introduces several enhancements to both the way logical paths are established and how you can manage their establishment. These are as follows:

- Logical paths are established between a logical partition and control units configured to the logical partition at logical partition activation instead of at POR. Non-ESCON Multiple Image Facility-capable CPCs establish logical paths at POR.

  At logical partition deactivation, logical paths associated with the logical partition are removed at both the channel subsystem and at the control units configured to the logical partition. Non-ESCON Multiple Image Facility-capable CPCs remove logical paths only at power off or when a channel path was deconfigured.

- Configure/deconfigure channel path commands for shared ESCON channel paths apply on a channel image basis. Configuring (or deconfiguring) a channel path for a logical partition will establish (or remove) the logical paths associated with that channel image.

- Channel path access and candidate lists can control the establishment of logical paths by controlling which logical partitions have a channel path configured online
- IODEVICE device candidate lists can control the establishment of logical paths.

8.6 Hardware System Area (HSA) Considerations

ESCON Multiple Image Facility and shared ESCON, coupling facility, and OSA channels, introduce new HSA considerations. See 3.8.2, “Hardware System Area (HSA)” on page 3-14.

8.7 ESCON Multiple Image Facility Channel Consolidation Benefits

Depending on how you take advantage of ESCON Multiple Image Facility, you can typically realize the following improved functionality and potential cost efficiencies for your system:

- **Reduced channel requirements and better resource management with comparable system performance**

  You can reduce channel requirements to better manage your channel resources by consolidating unshared dedicated and unshared reconfigurable channels and redefining them as shared channels.

  ESCON CTC connections can be redefined so that a single CTC/CNC pair can serve all logical partitions on a CPC. A second pair can be configured for performance, reconfigurability, and redundancy purposes.

  You can accomplish all of this and enhance system performance in comparison to a comparable unshared channel configuration.

- **Increased channel utilization**

  By having several logical partitions share a channel, channel utilization and efficiency can increase for a better return on investment.

- **I/O configuration flexibility**

  Use of shared channels can improve I/O connectivity, minimizing or eliminating channel reconfiguration among logical partitions.

- **Potential lower ESCON entry cost/Free ESCON Director (ESCD) ports and ESCON control unit adapters**

  By reducing ESCON channel requirements, you can migrate from parallel channels to shared ESCON channels at a lower cost than possible without ESCON Multiple Image Facility.

  By consolidating ESCON channels through ESCON channel sharing, you can reduce the number of ESCD ports you use without decreasing the I/O connectivity you want. Freed ESCD ports can increase connectivity by allowing attachment to previously unattached control units.

  Also, ESCON control unit adapters can be freed as a result of sharing ESCON channels. This can enhance I/O connectivity by creating the opportunity for additional control unit attachment.

- **Remote site flexibility**

  ESCON Multiple Image Facility can reduce fiber cable costs by reducing the number of ESCON channels used for data backup and communication with a recovery site.

- **Additional logical partitions**

  With ESCON Multiple Image Facility, you can add new logical partitions without necessarily adding more channel resources because new logical partitions can share the same installed channels and I/O devices.

Additionally, you can define logical partitions without channel resources.

- **Additional security**

  By defining a candidate list of logical partitions, you can enhance security by limiting the logical partitions that can access a reconfigurable channel.

Chapter 9. Coupling Facility Channels

9.1 Overview

Coupling facility channels are System/390 channels that use fiber optic cables 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.

Coupling facility channels:

- Require a point-to-point connection (direct channel attach between a CPC or LP and a coupling facility)
- 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
- 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
- Require CFR channel path definition (coupling facility receiver channel capability) at the coupling facility end of a coupling facility channel connection
- Require CFS channel path definition (coupling facility sender channel capability) at the CPC or MVS end of a coupling facility channel connection. You can define CFS channel paths for basic or LPAR mode channel configurations.
- Support the following specifications:
  - 50/125 micron multimode fiber with a maximum distance of one kilometer
  - 9/125 micron single mode fiber with a maximum distance of three kilometers

    9/125 Micron Single Mode Fiber: Support for the 9/125 micron single mode fiber is available with engineering change (EC) D57264 for 9674 CPCs and EC D57262 for 9672 CPCs.

9.1.1 Coupling Facility Sender Channels (TYPE=CFS)

You can configure a CFS 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 concurrently used by the LPs to which it is configured.

9.1.2 Coupling Facility Receiver Channels (TYPE=CFR)

You can configure a CFR channel path as:

- An unshared dedicated channel path to a single coupling facility LP

Changing the CFR channel path configuration requires a power-on reset. Dynamic I/O configuration does ***not*** support CFR channel paths.

9.1.3 I/O Configuration Considerations

IZP IOCP Version 1 Release 3 (1.3) and above support coupling facility channel path definition on 9674 CPCs and 9672 CPCs.

With MVS/ESA SP 5.1, HCD provides the following enhancements to support coupling facility definition:

- Controls for defining coupling facility channels. HCD also automatically generates the control unit and device definitions associated with CFS channel paths.
- Controls for defining a logical partition (LP) as either a coupling facility or an operating system LP. Additionally, you must use these controls when defining a new LP in HCD.

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

Additionally, because MVS/ESA SP 5.1 does ***not*** support the MVS configuration program (MVSCP), you must use HCD to define the channel configuration to MVS.

9.2 Integrated Coupling Migration Facility (ICMF)

The Integrated Coupling Migration Facility (ICMF) provides a test or migration alternative to purchasing Coupling Facility channel hardware. By emulating Coupling Facility channel hardware, ICMF allows you to connect coupling facility logical partitions (LPs) to LPs running MVS on the same central processor complex (CPC) for data sharing purposes.

You can define coupling facility LPs that use ICMF to connect to other LPs running MVS that use ICMF to assist in the test and development of data sharing applications.

At LP activation, a coupling facility LP using ICMF automatically loads Coupling Facility Control Code. Coupling Facility Control Code provides shared storage management functions for the coupling facility LP using ICMF. No initial program load (IPL) of an operating system is necessary or supported in a coupling facility LP that uses ICMF.

You can enable ICMF prior to installing Coupling Facility channel hardware hardware and continue to use it after you have added Coupling Facility channel hardware hardware to the CPC.

ICMF provides you with a great deal of flexibility when defining PR/SM LPAR configurations. Coupling facility LPs using ICMF and MVS LPs using ICMF can coexist on the same CPC with:

- Coupling facility LPs that use Coupling Facility channel hardware hardware
- LPs running MVS that use Coupling Facility channel hardware hardware
- LPs running workloads not associated with a coupling facility

However, you ***cannot*** use coupling facility channels and emulated coupling facility channels concurrently in the same LP.

Use of ICMF for production purposes has the following considerations:

- Additional overhead is associated with the use of ICMF
- A sysplex defined using ICMF can only involve LPs on the same CPC
- The potential for extended recovery exists with running a coupling facility LP on the same CPC as the MVS LPs connected to it

If your environment allows you to accept these conditions, you can use ICMF for production use.

9.2.1 Emulated Coupling Facility Channels

PR/SM LPAR emulates coupling facility channel paths to provide the connectivity for data sharing between coupling facility LPs using ICMF and MVS LPs using ICMF. PR/SM LPAR emulates coupling facility channel paths as follows:

- Coupling facility LPs using ICMF and MVS LPs using ICMF must be defined in the IOCDS used for power-on reset. If channel resources are defined for LPs using ICMF, they are removed as follows:
  - For coupling facility LPs, all channel definitions, if defined, are removed from the coupling facility LP channel configuration.
  - For MVS LPs, CFS and CFR channel paths, if defined, are removed from the MVS LP channel configuration. ESCON, OSA, and parallel channel paths, if defined to the MVS LP channel configuration, are configured and can be brought online.
- For coupling facility using ICMF, emulated CFR channel paths allow connection to MVS LPs using ICMF.
- For MVS LPs, emulated CFS channel paths allow connection to coupling facility LPs using ICMF.

Additionally, PR/SM LPAR emulates CFS control units, CFS devices, and subchannels associated with CFS channel paths. Because PR/SM LPAR emulates the coupling facility channel configuration for those LPs using ICMF, there is no information on these emulated channel resources contained in any IODF or IOCDS. However, PR/SM LPAR does assign numbers for the emulated channel resources.

Dynamic I/O configuration does ***not*** support emulated CFR or emulated CFS channel paths but does support the ESCON (CNC and CTC channel paths), OSA, and parallel channel paths defined to an MVS LP using ICMF.

9.2.2 Emulated Coupling Facility Channel Configuration

For the emulated coupling facility channel configuration, PR/SM LPAR assigns numbers to the following emulated resources:

- CFR and CFS channel paths
- CFS control units
- CFS devices
- Subchannels associated with CFS channel paths

To avoid confusing emulated coupling facility channel resources with actual installed and configured channel resources, you should avoid using the numbers PR/SM LPAR assigns to the emulated coupling facility channel configuration, if possible.

For more information, see *PR/SM Planning Guide*, GA22-7123.

9.2.3 LP Definition Considerations for LPs Using ICMF

You enable ICMF for coupling facility LPs and MVS LPs by enabling ICMF from the Activation Profiles task available from the CPC Operational Customization Tasks Area. Target the LP to use ICMF. Select the General Information pushbutton on the CPC Operational Customization panel, then select the ICMF field that displays. You must do this for each coupling facility LP and each MVS LP to use ICMF.

You can define as many LPs as your CPC supports. However, you can only activate two coupling facility LPs using ICMF; one having an even partition number and one having an odd partition number.

Coupling facility LPs using ICMF can run with either dedicated CPs or shared CPs. Coupling facility LPs using ICMF and MVS LPs using ICMF are defined like coupling facility LPs and MVS LPs that do ***not*** use ICMF.

9.2.4 LP Activation Considerations for LPs Using ICMF

Coupling facility LPs using ICMF and MVS LPs using ICMF activate like LPs that do not use ICMF. However, at LP activation:

- All channel paths are removed from the channel configuration of a coupling facility LP using ICMF
- All CFS and CFR channel paths are removed from the channel configuration of an MVS LP using ICMF

Note: Any parallel, ESCON, CFS, or OSA channel paths targeted to be brought online will automatically be configured online if the coupling facility LP is redefined as an ESA mode LP.

9.2.5 LP Operation Considerations for LPs Using ICMF

PR/SM LPAR limits the same hardware operator controls for coupling facility LPs using ICMF as it does for coupling facility LPs not using ICMF. Operator activity is confined to the hardware console. Additionally, there are further considerations as follows:

- Emulated CFR channel paths ***cannot*** be displayed using frames or panels associated with the hardware console. However, the Coupling Facility Control Code DISPLAY command will display emulated CFR channel paths.
- Emulated CFR channel paths:
  - Do ***not*** appear in any HCD or IOCP input or output
  - ***Cannot*** be dynamically reconfigured using dynamic I/O configuration

MVS LPs using ICMF operate like other ESA mode LPs with the following exceptions:

- Emulated CFS channel paths ***cannot*** be displayed using frames or panels associated with the hardware console. However, the MVS DISPLAY CF command will display emulated CFS channel paths.
- Emulated CFS channel paths:
  - Do ***not*** appear in any HCD or IOCP input or output
  - ***Cannot*** be configured online or offline
  - ***Cannot*** be dynamically reconfigured using dynamic I/O configuration

9.2.6 MVS Operational Considerations for MVS LPs Using ICMF

For MVS LPs using ICMF, these MVS commands will operate as follows:

- DISPLAY CF will display emulated CFS channel paths and any associated device numbers and subchannel numbers.
- DISPLAY MATRIX will ***not*** display the online or offline status of emulated CFS channel paths; emulated CFS channel paths will always display as offline.
- CF CHP(nn),ONLINE/OFFLINE will ***not*** function for emulated CFS channel paths.

Chapter 10. ESCON and Parallel Channel Subsystem Performance

10.1 Introduction

This chapter describes the characteristics of ESCON channels and parallel channels that affect the performance of the channel subsystem.

Maximizing channel subsystem performance is an important consideration in configuring I/O devices to a 9672 CPC. Channel subsystem performance depends on the factors described in this chapter.

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

10.2 Overview

Channel subsystem performance can be examined by observing two measurements:

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

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

- Queuing and setup time
- Control unit and device time
- Data transfer time
- Completion time

These major components of channel performance are affected by:

- Type of operation (nonsynchronous versus synchronous)
- Data transfer rates
  - Distance
- Attached device characteristics
- Workload characteristics

10.3 Components of Channel Response Time and Throughput

This section provides the description of the components of channel performance for all ESCON and parallel channels.

- Queueing and Setup Time
  - The time taken for a channel path, control unit, and device to become available
  - The time taken for a channel to send the I/O operation commands to the control unit
- Control Unit and Device Time

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

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

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

10.4 Factors Affecting the Components of Channel Performance

This section provides the description of the factors that affect channel performance characteristics.

Factors that affect the various components of performance include:

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

10.4.1 Synchronous and Nonsynchronous I/O Operation

For detailed information about concepts described in this section, see *Introduction to Nonsynchronous Direct Access Storage Subsystems*.

10.4.1.1 Synchronous Operation

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

All work involved in ending an operation and advancing to the next operation must be completed before the DASD head reaches the next record (commonly referred to as the inter-record gap). If this does not occur, a rotational positional sensing/sensor (RPS) miss or an overrun is generated and the operation must wait for one DASD revolution before continuing.

10.4.1.2 Nonsynchronous Operation

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

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

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

10.4.2 Data Transfer Rate

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

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

- Channel data rate
- Control unit data rate
- Device data rate

The channel data rate is the rate that the channel transfers data between the transmission link and processor storage during an I/O operation.

For ESCON channels, the channel data rate is 17MB/s for block sizes less than or equal to 2048 bytes. However, it is a function of distance for block sizes larger than 2048 bytes. Table 10-1 shows channel data rate as a function of distance for block sizes larger than 2048 bytes.

For parallel channels and ESCON channels attached to a 9034, the channel data rate is 4.5MB/s.

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

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

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

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

Table 10-1. ESCON Channel Data Rate as a Function of Distance for 9672 CPCs for block sizes > 2048 bytes.

Distance (km)	**Distance (miles)**	**Channel Data Rate (MB/s)**
0.122 (122 meters)	0.075 (133 yards)	17
3	1.86	17
6	3.72	17
9	5.59	17
12	7.46	14.2
15	9.32	11.9
20	12.43	9.4
30	18.64	6.3
43	26.72	4.7
60	37.28	3.4

10.4.3 I/O Device Characteristics

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

10.4.4 Channel Subsystem Workload Characteristics

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

10.4.4.1 Channel Program Characteristics

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

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

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

Chained data commands increase response time due to an additional interlocked exchange between the channel and control unit. See 10.5.1.1, “Distance” on page 10-6 for more information.

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

10.4.4.2 Cache-Hit Rates

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

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

10.5 Channel Performance Characteristics

10.5.1 All ESCON Channels

This section describes the performance characteristics of ESCON channels.

10.5.1.1 Distance

The distance between the channel and control unit affects the setup and completion times of an I/O operation; as the distance between the channel and the control unit increases, the response time increases and the throughput decreases. Channel and control unit utilization also increase as distance between the channel and control unit increases.

The speed of data transfer through fiber optic cable is subject to the propagation delay time. Propagation delay time is determined by two factors: the speed of light through the optical fiber (which is fixed), and the length of the fiber optic link. Propagation delay time increases as the distance between elements in a fiber optic environment increase.

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

10.5.1.2 Shared ESCON Channels

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

10.5.2 ESCON Channels (Attached to a 9034) and Parallel Channels

This section describes channel performance concepts for ESCON channels (attached to a 9034) and parallel channels.

Performance-related factors that should be considered in configuring ESCON channels (attached to a 9034) and parallel channels include:

- Elapsed time
- Critical time
- Deferred access

10.5.2.1 Elapsed Time

The elapsed time to complete service for an I/O device connected to a channel includes the device tag time, the cable propagation time, the channel busy time to service the requested control sequence, and the wait time for channel resources.

10.5.2.2 Critical Time

The time that an I/O device can wait for channel service without a negative impact on the performance of the device is called critical time. Each I/O device has limits on the elapsed times of various sequences between the device and the channel. When these time limits are not satisfied, device performance may be reduced.

For devices operating on block multiplexer channels, the control sequences that have the most significant critical-time constraints are those related to chaining operations. For devices operating on byte multiplexer channels, the time for connecting and disconnecting of the device to the channel for each byte or burst of data sent is the most critical sequence.

In some cases, the time limit is related to synchronized motion (for example, the time between columns on a card moving through an IBM 2501 Card Reader, or the time between the end of a Search ID Equal and

the beginning of a Read or Write at an IBM 3380 Direct Access Storage). In each case, a control sequence must be completed within a time limit that relates directly to the physical motion of the device to sustain maximum I/O performance.

The critical time of a device can be exceeded because of other traffic on the channel or other traffic in the channel subsystem. Central storage loading can also contribute to the elapsed time, but not significantly. If the elapsed time is greater than the critical time, some performance degradation occurs. The result of exceeding critical time is described under 10.5.2.3, “Deferred Access” and 10.6.1, “Device Class” on page 10-8.

10.5.2.3 Deferred Access

A data deferred access is caused by the inability of the channel to transmit or accept data at the rate requested or transmitted by an I/O device.

A data deferred access is much less likely to occur on buffered devices than on unbuffered devices because buffered devices can wait for channel service. Unbuffered devices (such as start-stop terminals) may have data deferred accesses when the time required for an error-recovery system logout exceeds the critical time.

Data-chaining operations involving the transfer of one or more blocks of data increase the probability of data deferred accesses with devices that do not respond to ‘suppress out’.

A chaining check is an error detected in a channel when a channel accepts more data (in an input operation) than was specified by the count in the CCW. The check occurs when an I/O data rate is too high to be handled by the channel and storage.

A command deferred access is the inability of the channel to present the next command within the critical command-chaining time of a control unit.

Degradation (a loss in performance) can result from a deferred access. A deferred access that requires operator intervention can create significant degradation. In most cases, a deferred access that is handled automatically by retry does not significantly affect throughput.

Depending on the device and the type of deferred access, the operation may be halted when the need for a deferred access occurs, or it may continue transferring data until the end of the block is reached. A deferred access may cause a unit check to be presented to the channel. Any chaining is suppressed and an I/O interruption request is generated at the end of the operation. Certain control units, however, may initiate a command retry sequence without generating an I/O interruption request. See 7.5.4, “I/O Interruptions” on page 7-8 for additional information.

10.6 ESCON Channels (CBY Channel Paths) and Parallel Channels

ESCON (CBY channel paths) and parallel channels can operate in byte multiplexer mode.

Note: 9672 CPCs with engineering change (EC) D57262 will be able to define and use CBY channel paths to connect to 3720, 3725, and 3745 control units requiring byte multiplexer channel path attachment. Connection is through a 9034 ESCON Converter Model 1.

See the 9672IO subset id in the 9672DEVICE upgrade id of the preventive service planning (PSP) bucket for prerequisite 9034 EC level information.

10.6.1 Device Class

Devices that can operate in byte multiplexer mode of operation are classified by what happens when the device is not serviced within the critical time for the requested control sequence for that device. Depending on how overall channel performance is impacted by the critical time for the device being exceeded, a device falls into one of three classes: 1, 2, or 3.

10.6.1.1 Device Class 1

When the critical time is exceeded, a deferred access occurs and the data is not transferred successfully. The consequent error indication causes an I/O interruption request, and program recovery action is required.

10.6.1.2 Device Class 2

When the critical time is exceeded, the device must be resynchronized. The additional delay results in performance degradation. The device performance is degraded by the combined delay of waiting for the channel and resynchronization.

10.6.1.3 Device Class 3

When the critical time is exceeded, the device waits for channel service and causes performance degradation (the delay of waiting for the channel service).

10.7 ESCON Channels Using Link-Level and Device-Level Protocols

The ESCON Architecture provides two protocols for block multiplexer mode of operation on the I/O interface for the serial transmission of data:

- Link-level protocols
- Device-level protocols

Block multiplexer mode of operation using link-level and device-level protocols can sustain a maximum data rate of 17MB/s.

I/O operations for the serial transmission and reception of data require that link-level and device-level protocols be present in both the channel and the control unit.

10.8 ESCON Channels (CVC Channel Paths) and Parallel Channels

The channel subsystem provides two modes for block multiplexer mode of operation on the I/O interface in a parallel environment:

- Interlocked
- Data streaming

10.8.1.1 Interlocked

Operation performance using the interlocked mode depends on overall tag timings (including channel subsystem service), cable length, and control unit service. Block multiplexer mode of operation using the interlocked protocol can sustain a maximum data of 1.5MB per second.

10.8.1.2 Data Streaming

The data-streaming protocol does not require interlocking of data transfer signals between the channel and the control unit; once data transfer is established over the interface, it continues at a rate governed by the control unit. Block multiplexer mode of operation using the data-streaming protocol can sustain a maximum data rate of 4.5MB per second.

10.9 Planning for Channel Subsystem Availability

The following figures show the major components of the channel subsystem (CSS) and suggest ways to configure the CSS for maximum availability.

10.9.1 9672 E/P/R1 Models and 9674 C01 Models

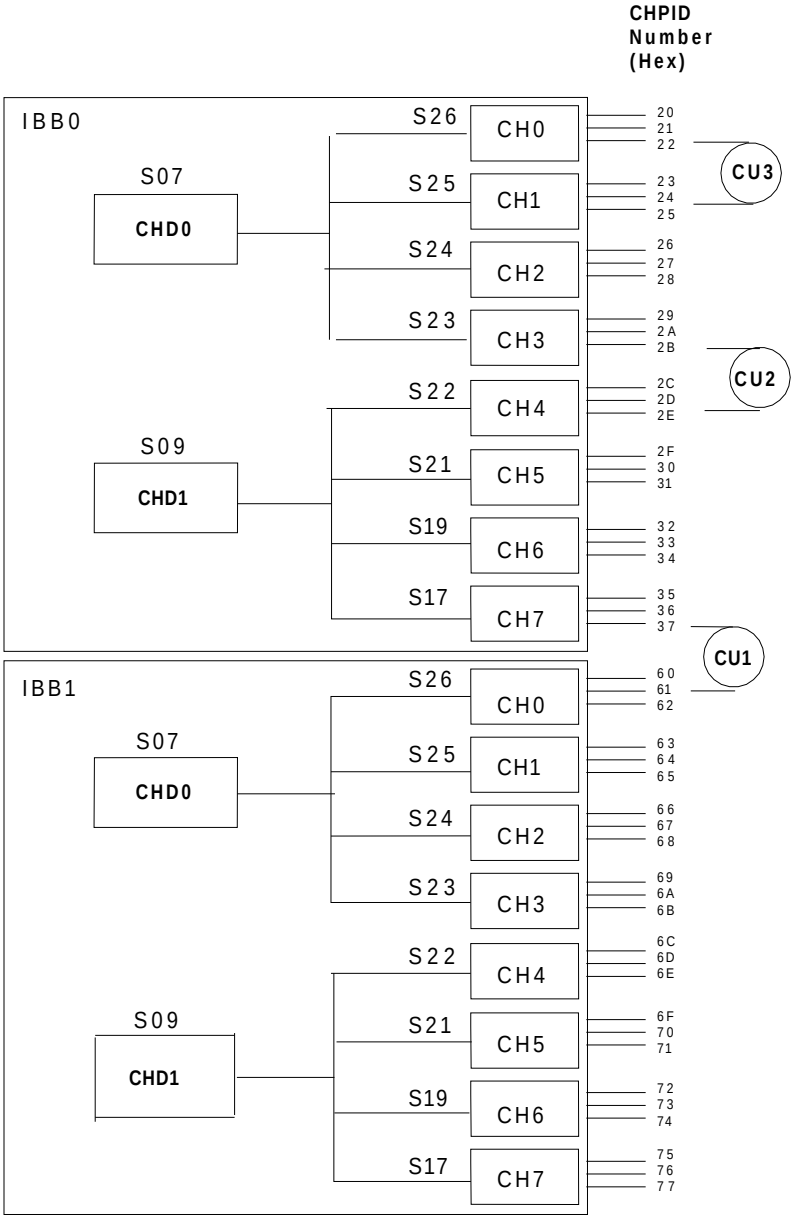

Figure 10-1. I/O channel path configuration. For a representative 9672 E, P, or R1 model. CPCs with an expansion board have two IBBs.

Legend:

IBB	Internal bus buffer
CH	Channel card
CHD	Channel driver card
S	Slot

In Figure 10-1 on page 10-10, each IBB can connect to 24 channels.

When configuring devices with multiple paths to the same processor complex, select any of the channel paths shown in Figure 10-1 on page 10-10 that:

- Are available on the CPC you are defining
- Are the correct type (parallel or ESCON) to meet the control unit attachment requirements
- Satisfy the rules regarding the mixing of channel types

However, for maximum availability of the device, you should consider the following guidelines:

- For systems with two IBBs, distribute paths across both IBBs (see CU 1).
- For systems with one IBB, or if you define multiple paths on the same IBB, distribute paths across both CHDs (see CU 2).
- For single CHD systems or if you define multiple paths on the same CHD, distribute paths across different CHs (see CU 3).

Notes:

1. Examples CU 2 and CU 3 also apply to IBB1.
2. Each example in Figure 10-1 on page 10-10 (CU 1, CU 2, and CU 3) is only an attempt to illustrate an availability guideline and does not represent an overall optimal configuration.

Figure 10-2. Coupling facility channel path configuration. For a representative 9674 or 9672 CPC. This figure applies to 9674 Model C01 and 9672 E, P, and R1 model CPCs. CPCs with an expansion board have two IBBs. Only a 9674 Model C01 CPC with two IBBs can have a maximum of 32 coupling facility channels. A 9672 E or P model CPC with two IBBs can have a maximum of 24 coupling facility channels. A 9672 R1 model CPC with two IBBs can have a maximum of 16 coupling facility channels.

Legend:

IBB	Internal bus buffer
CF	Coupling facility channel adapter card

In Figure 10-2 on page 10-11, each IBB can connect to 16 channels.

When using multiple paths from a coupling facility to the same processor complex or logical partition, select any of the channel paths shown in Figure 10-2 on page 10-11 that are available on the CPC you are defining.

However, for maximum availability of the coupling facility, you should consider the following guidelines:

- For systems with two IBBs, distribute paths across both IBBs (see CPC1).
- If you define multiple paths on the same IBB, distribute across CF0 through CF7 (see CPC2).

Notes:

1. Example CPC2 also applies to IBB1.
2. Each example in Figure 10-2 on page 10-11 (CPC1 and CPC2) is only an attempt to illustrate an availability guideline and does not represent an overall optimal configuration.

10.9.2 9672 R2 and R3 Models and 9674 C02 and C03 Models

The following figures show the major components of the channel subsystem (CSS) and suggest ways to configure the CSS for maximum availability.

The I/O configuration for a 9672 Model R3 with a feature code 2020 I/O cage can be either of the following:

- A double IBB card (see Figure 10-3 on page 10-14) combined with single IBB card 5 (see Figure 10-4 on page 10-15)
- Any combination of single IBB cards 12, 13, and 5 (see Figure 10-4 on page 10-15)

The I/O configuration for a 9672 Model R2 with a feature code 2020 I/O cage can be any combination of single IBB cards 12, 13, and 5 (see Figure 10-4 on page 10-15).

The I/O configuration for a 9672 Model R2 with a feature code 0020 CPC cage consists of single IBB cards 12 and 13 (see Figure 10-4 on page 10-15).

The I/O configuration for a 9672 Model R3 or a 9672 Model R2 with a feature code 1000 I/O cage consists of single IBB card 8 (see Figure 10-5 on page 10-16).

The I/O configuration for a 9674 Model C03 with a feature code 2020 I/O cage or a 9674 Model C02 with CPC cage(0020) consists of single IBB cards 12 and 13 (see Figure 10-6 on page 10-17).

Notes:

1. You will need to know the feature codes of the I/O cages used in your configuration and where each IBB is plugged in each I/O cage to use this section of the publication.
2. The Channel Path Identifier (CHPID) is used by the IOCDS to define each physical path to a device. The CHPID numbers used in Figure 10-3 on page 10-14, Figure 10-4 on page 10-15, Figure 10-5 on page 10-16, and Figure 10-6 on page 10-17 are relative CHPID numbers. The actual CHPID numbers are dependent on the IBB cable number. For the R3 models, each IBB bus has a CHPID range of 64. For the R2 models, each IBB bus has a CHPID range of 32.
   - For 9672 R3 models, refer to Table 10-2 on page 10-19, Table 10-3 on page 10-20, and Table 10-5 on page 10-22.
   - For 9672 R2 models, refer to Table 10-4 on page 10-21 and Table 10-6 on page 10-22
   - For 9674 Model C03, refer to Table 10-3 on page 10-20.
   - For 9674 Model C02, refer to Table 10-4 on page 10-21.

   The available CHPIDs for a specific machine are dependent on the model type and I/O subsystem configuration for that machine.

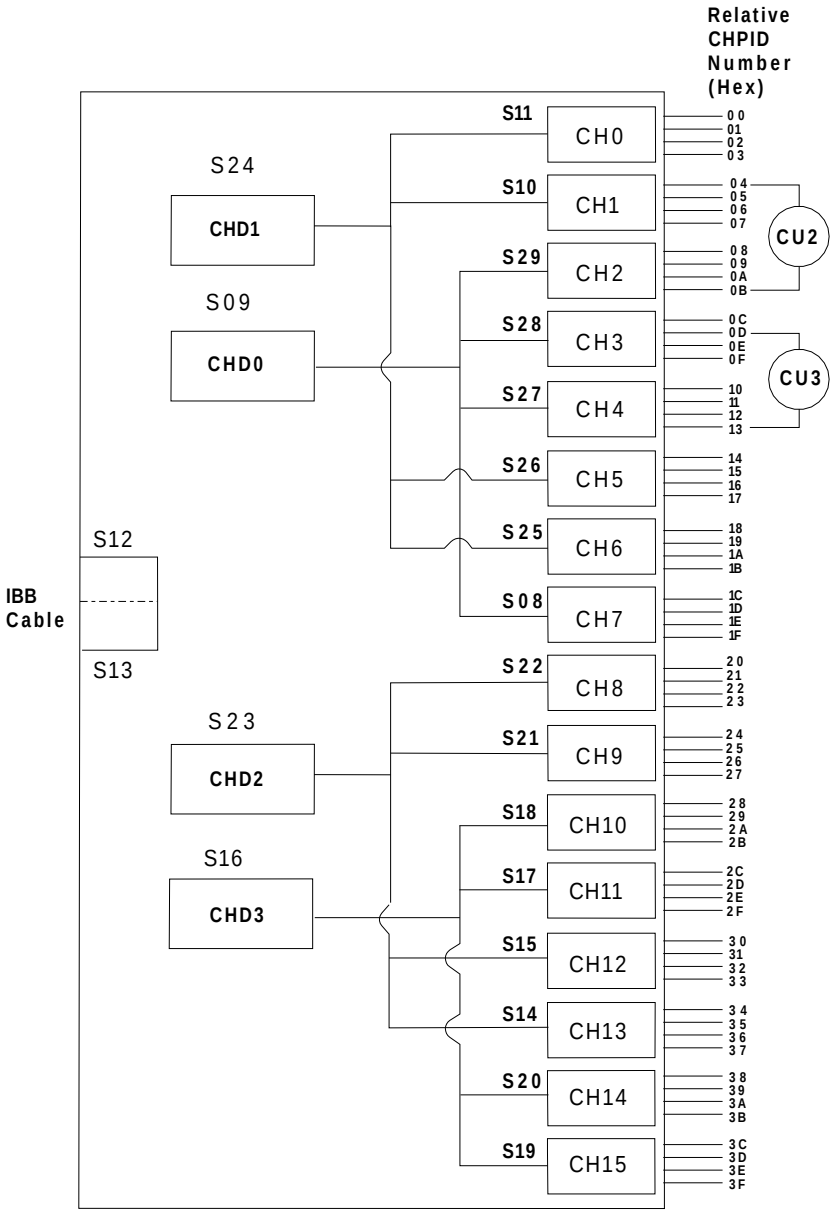

Figure 10-3. I/O channel path configuration. For a representative 9672 R3 model with feature code 2020 I/O cage with double IBB card.

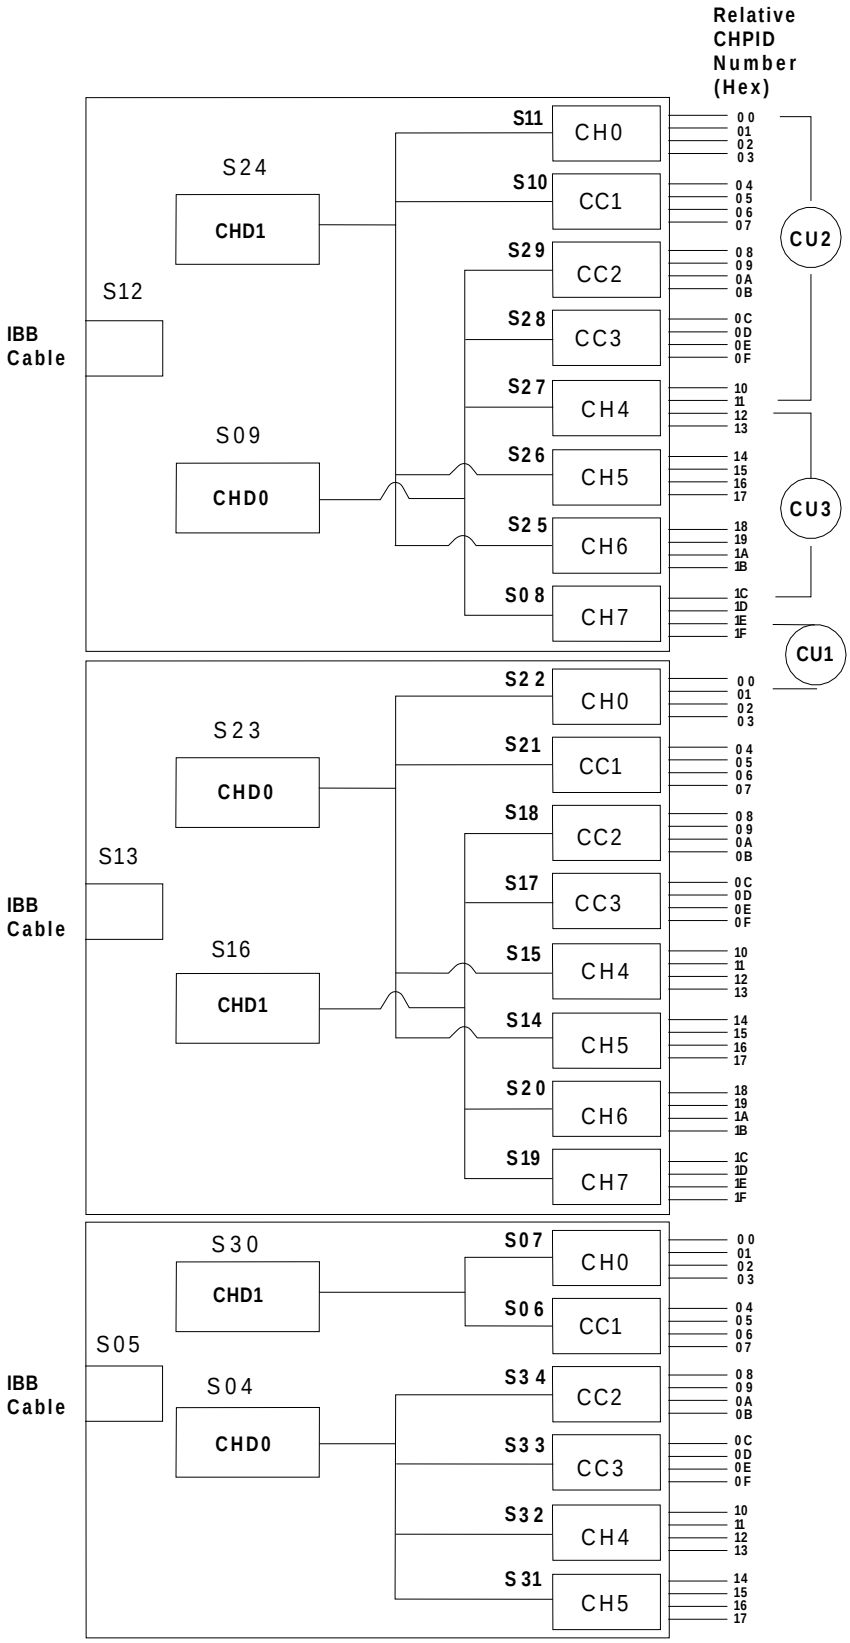

Figure 10-4. I/O channel path configuration. For a representative 9672 R3 model or 9672 R2 model. For 9672 R3 model with feature code 2020 I/O cage and single IBB cards. For 9672 R2 model with feature code 2020 I/O cage and single IBB cards **(only uses cards 12 and 13)**. *For 9672 R2 model feature code 0020 CPC cage with single IBB cards* **(only uses cards 12 and 13)**.

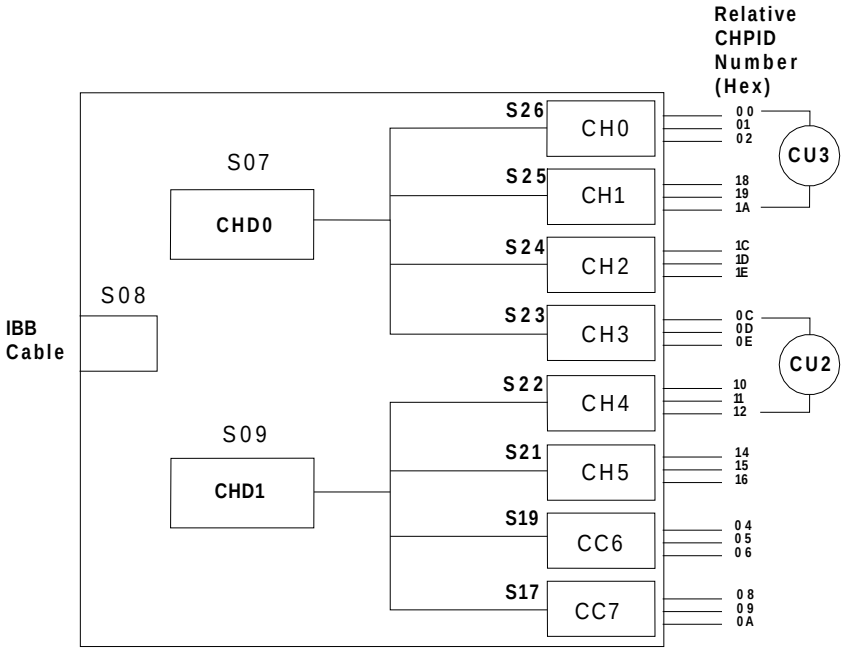

Figure 10-5. I/O channel path configuration. For a representative 9672 R3 model or 9672 R2 model with feature code 1000 I/O cage.

Legend:

IBB	Internal bus buffer
CC	Channel card or Coupling facility channel adapter card (Note: The coupling facility channel adapter card is not driven by a channel driver card.)
CHD	Channel driver card
CH	Channel card
S	Slot

When configuring devices with multiple paths to the same processor complex, select any of the channel paths shown in Figure 10-3 on page 10-14, Figure 10-4 on page 10-15, or Figure 10-5 that:

- Are available on the CPC you are defining
- Are the correct type (parallel or ESCON) to meet the control unit attachment requirements
- Satisfy the rules regarding the mixing of channel types

However, for maximum availability of the device or coupling facility on a 9672 R2 or R3 model, you should consider the following guidelines:

- For systems with multiple I/O cages, distribute paths across the I/O cages (not shown)
- For systems with multiple IBBs, distribute paths across the IBBs (see CU 1).

  Note: This is also recommended for optimum performance of your most heavily used I/O devices. For R63 through RX3, it is also recommended you distribute paths between CHPIDs 00-7F and 80-FF.
- If you define multiple paths on the same IBB, distribute paths across different CHDs (see CU 2). This does not apply to the coupling faciltiy channel adapter cards because they do not require channel driver cards.

  Note: For I/O intensive mode users: Distribution across CHDs is important for optimum performance in I/O intensive mode. In I/O intensive mode, CHPIDs in the range of 00-7F are assigned to two System Assist Processors (SAPs), and CHPIDs in the range 80-FF (only models R63 through RX3) are assigned to two other SAPs. For a given range, one SAP services channels on all even num-

bered CHDs, and the other SAP services channels on all odd numbered CHDs. Within a given CHPID range, you should distribute paths across even and odd numbered CHDs.

- If you define multiple paths on the same CHD, distribute paths across different CHs (see CU 3). Also, if you define multiple coupling facility channels to the same couling facility, processor complex, or logical partition, distribute paths across different coupling facility channel adapter cards.

Note: Each example in Figure 10-3 on page 10-14, Figure 10-4 on page 10-15, and Figure 10-5 on page 10-16 (CU1, CU2, and CU3) is only an attempt to illustrate an availability guideline and does not represent an overall optimal configuration.

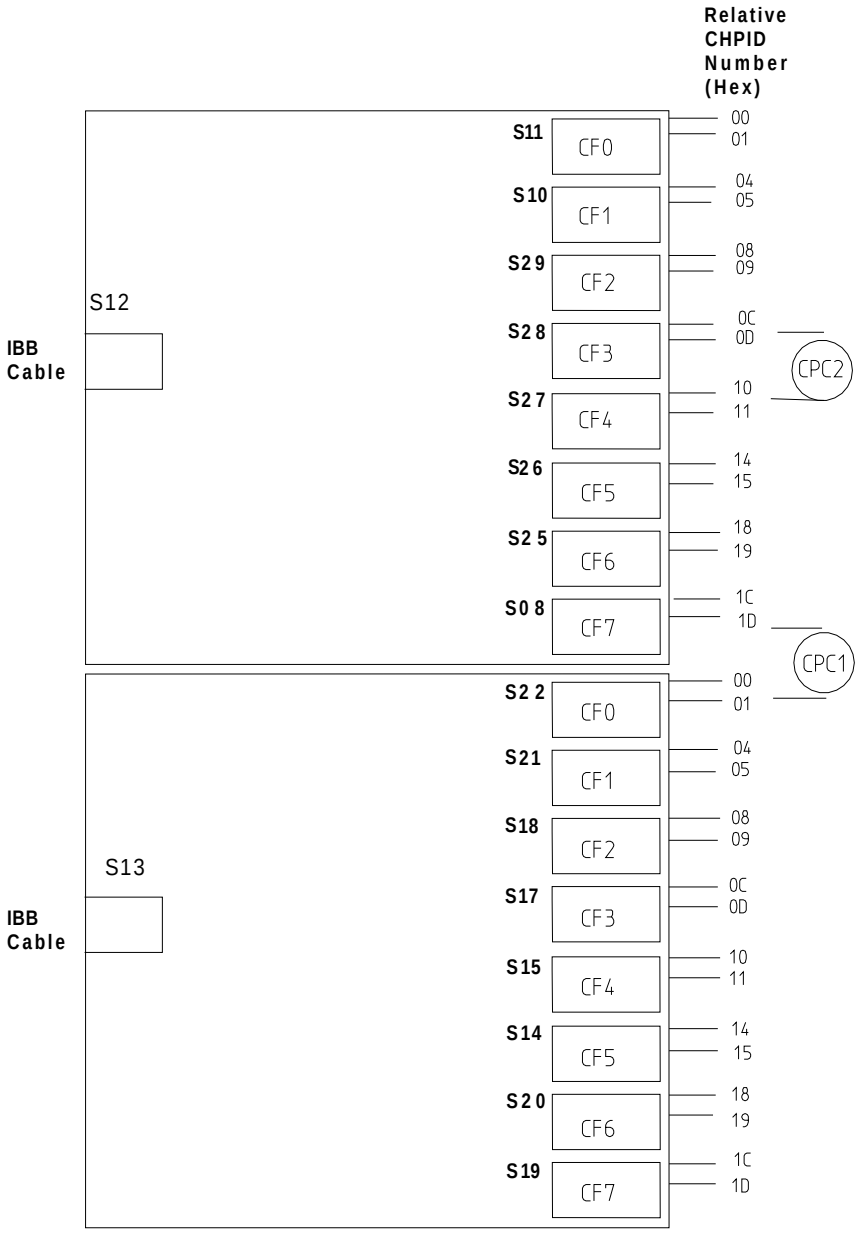

Figure 10-6. Coupling facility channel path configuration. For a representative 9674 Model C03 or 9674 Model C02. For 9674 Model C03 with feature code 2020 I/O cage and single IBB cards. For 9674 Model C02 with feature code 0020 CPC cage and single IBB cards.

Legend:

IBB	Internal bus buffer
CF	Coupling facility channel adapter card
S	Slot

When using multiple paths from a coupling facility to the same CPC or logical partition, select any of the channel paths shown in Figure 10-6 on page 10-17 that are available on the CPC you are defining.

However, for maximum availability of a 9674 Coupling Facility (Models C02 or C03), you should consider the following guidelines:

- For systems with multiple IBBs, distribute paths across the IBBs (see CPC1).
- If you define multiple paths on the same IBB, distribute paths across different CFs (see CPC2).

Note: Each example in Figure 10-6 on page 10-17 (CPC1 and CPC2) is only an attempt to illustrate an availability guideline and does not represent an overall optimal configuration.

Table 10-2. CHPID numbers for 9672 R3 models with double and single IBB card–feature 2020 I/O cage

IBB Card Slot	I/O Card Slot	IBB Cable Number			
		3	2	1	0
12 (Double IBB with 13)	11		80–83	40–43	00–03
	10		84–87	44–47	04–07
	29		88–8B	48–4B	08–0B
	28		8C–8F	4C–4F	0C–0F
	27		90–93	50–53	10–13
	26		94–97	54–57	14–17
	25		98–9B	58–5B	18–1B
	8		9C–9F	5C–5F	1C–1F
13 (Double IBB with 12)	22		A0–A3	60–63	20–23
	21		A4–A7	64–67	24–27
	18		A8–AB	68–6B	28–2B
	17		AC–AF	6C–6F	2C–2F
	15		B0–B3	70–73	30–33
	14		B4–B7	74–77	34–37
	20		B8–BB	78–7B	38–3B
	19		BC–BF	7C–7F	3C–3F
5 (Single IBB)	7	C0–C3	80–83	40–43	00–03
	6	C4–C7	84–87	44–47	04–07
	34	C8–CB	88–8B	48–4B	08–0B
	33	CC–CF	8C–8F	4C–4F	0C–0F
	32	D0–D3	90–93	50–53	10–13
	31	D4–D7	94–97	54–57	14–17

Note: IBB cables 0 and 2 can be plugged into slot 5 **OR** into slots 12/13. The CHPID range is as follows:

- With IBB0 in slot 5, the available IBB0 CHPID range is 00–17
- With IBB0 in slots 12/13, the available IBB0 CHPID range is 00–3F
- With IBB2 in slot 5, the available IBB2 CHPID range is 80–97
- With IBB2 in slots 12/13, the available IBB2 CHPID range is 80–BF

Table 10-3. CHPID numbers for 9672 R3 models and 9674 model C03 with single IBB cards–feature code 2020 I/O cage

IBB Card Slot	I/O Card Slot	IBB Cable Number			
		3	2	1	0
12 Single IBB	11	C0–C3	80–83	40–43	00–03
	10	C4–C7	84–87	44–47	04–07
	29	C8–CB	88–8B	48–4B	08–0B
	28	CC–CF	8C–8F	4C–4F	0C–0F
	27	D0–D3	90–93	50–53	10–13
	26	D4–D7	94–97	54–57	14–17
	25	D8–DB	98–9B	58–5B	18–1B
	8	DC–DF	9C–9F	5C–5F	1C–1F
13 Single IBB	22	C0–C3	80–83	40–43	00–03
	21	C4–C7	84–87	44–47	04–07
	18	C8–CB	88–8B	48–4B	08–0B
	17	CC–CF	8C–8F	4C–4F	0C–0F
	15	D0–D3	90–93	50–53	10–13
	14	D4–D7	94–97	54–57	14–17
	20	D8–DB	98–9B	58–5B	18–1B
	19	DC–DF	9C–9F	5C–5F	1C–1F
5 Single IBB	7	C0–C3	80–83	40–43	00–03
	6	C4–C7	84–87	44–47	04–07
	34	C8–CB	88–8B	48–4B	08–0B
	33	CC–CF	8C–8F	4C–4F	0C–0F
	32	D0–D3	90–93	50–53	10–13
	31	D4–D7	94–97	54–57	14–17

Table 10-4. CHPID numbers for 9672 R2 models and 9674 model C02–feature code 2020 I/O cage or feature code 0020 CPC cage

IBB Card Slot	I/O Card Slot	IBB Cable Number			
		3	2	1	0
12 Single IBB	11		40–43		00–03
	10		44–47		04–07
	29		48–4B		08–0B
	28		4C–4F		0C–0F
	27		50–53		10–13
	26		54–57		14–17
	25		58–5B		18–1B
	8		5C–5F		1C–1F
13 Single IBB	22	60–63		20–23	
	21	64–67		24–27	
	18	68–6B		28–2B	
	17	6C–6F		2C–2F	
	15	70–73		30–33	
	14	74–77		34–37	
	20	78–7B		38–3B	
	19	7C–7F		3C–3F	

Table 10-5. CHPID numbers for 9672 R3 models–feature code 1000 I/O cage

IBB Card Slot	I/O Card Slot	IBB Cable Number			
		3	**2**	**1**	**0**
8 Single IBB	26	C0–C2		40–42	00–02
	19	C4–C6		44–46	04–06
	17	C8–CA		48–4A	08–0A
	23	CC–CE		4C–4E	0C–0E
	22	D0–D2		50–52	10–12
	21	D4–D6		54–56	14–16
	25	D8–DA		58–5A	18–1A
	24	DC–DE		5C–5E	1C–1E

Table 10-6. CHPID numbers for 9672 R2 models–feature code 1000 I/O cage

IBB Card Slot	I/O Card Slot	IBB Cable Number	
		3	**2**
8 Single IBB	26	60–62	40–42
	19	64–66	44–46
	17	68–6A	48–4A
	23	6C–6E	4C–4E
	22	70–72	50–52
	21	74–76	54–56
	25	78–7A	58–5A
	24	7C–7E	5C–5E

Chapter 11. Supported I/O Devices

11.1 9672 Supported I/O Devices

9672 Central Processor Complexs (CPCs) support the following I/O devices:

- **Direct Storage Access Devices**
  - 3880 Storage Controller
  - 3990 Storage Controller
  - 3995 Optical Library Dataserver
  - 9341 Storage Controller
    (supports 78 pin D-shell connector cable, see note)
  - 9343 Storage Controller
    (supports 78 pin D-shell connector cable, see note)
  - 9394 RAMAC Array Subsystem
- **Printers**
  - 3800 Printing Subsystem
  - 3820 Page Printer
  - 3825 Page Printer
  - 3827 Page Printer
  - 3828 Page Printer
  - 3835 Page Printer
  - 3900 Duplex Printing System
- **Tape Devices**
  - 2440 Magnetic Tape Subsystem
  - 3480 Magnetic Tape Subsystem
  - 3490 Magnetic Tape Subsystem
    (supports 78 pin D-shell connector cable, see note)
  - 3494 Tape Library Dataserver
- **Communication Controllers**
  - 3720 Communications Controller
  - 3725 Communications Controller
  - 3745 Communications Controller
- **ESCON Directors**
  - 9032 ESCON Director
  - 9033 ESCON Director
- **ESCON Converters**
  - 9034 ESCON Converter
  - 9035 ESCON Converter
- **Miscellaneous Other**
  - 3088 Multisystem Channel Communication Unit
  - 3814 Switching Management System
  - 7013 RS/6000
  - 7015 RS/6000
  - 9036 ESCON Remote Channel Extender
  - 9076 Power Parallel Systems
- **Miscellaneous Controllers**
  - 3172 Interconnect Controller
  - 3174 Establishment Controller

Note: Supports direct channel attach using a 78 pin D-shell connector channel cable. All other parallel devices can attach to a 78 pin D-shell connector channel cable through an adapter cable that allows you to convert from a 78D connector to bus and tag connectors at the I/O device.

See Chapter 7, “Characteristics of ESCON and Parallel Channel Operation” on page 7-1 for additional detail on I/O devices.

Chapter 12. Upgrade Progression

12.1 Overview

You can upgrade 9672 Parallel Transaction Server E-Models and P-Models, 9672 Parallel Enterprise Server R1/R2/R3 Models, and the 9674 Coupling Facility C01/C02/C03 Models as follows:

12.1.1 9672 Parallel Transaction Server E-Models and P-Models Upgrade Progression

The following table provides the 9672 Parallel Transaction Server E-Models and P-Models model upgrade paths.

Example: Model E01 can be upgraded to model E02, E03, or E04. Model P01 can be upgraded to model P02 or P03.

Table 12-1. 9672 Parallel Transaction Server E-Models and P-Models Upgrade Progression

Model	Upgrade Model
E01	E0(2-4); R(1-5,7)2; R(7-8,X)3 (See notes 2 and 3)
E02	E0(3-4)
E03	E0(4-6)
E04	E0(5-6)
E05	E0(6-8)
E06	E0(7-8)
E07	E08
P01	P0(2-3); R(1-5,7)2: R(7-8,X)3 (See notes 2 and 3)
P02	P03

Notes:

1. 9672 Parallel Transaction Server E-Models and P-Models model upgrades enhance the processing capability and availability of the 9672 Parallel Transaction Server by adding additional Central Processor Complexs (CPCs) to the 9672 Parallel Transaction Server string.

Additionally, you can upgrade the processing capability of individual CPCs in the 9672 Parallel Transaction Server string by adding central processors (CPs). Additional CPs are available as optional features on a CPC basis within the string. Each CPC can have from two to six CPs.

See Chapter 3, “Hardware Description” on page 3-1 for a description of the hardware configuration for each 9672 Parallel Transaction Server model number

12.1.2 9672 Parallel Enterprise Server R1/R2/R3 Models Upgrade Progression

The following table provides the 9672 Parallel Enterprise Server R1/R2/R3 Models model upgrade paths.

Example: Model R11 can be upgraded to model R21, R31, R41, R51, R61, R12, R22, R32, R42, R52, R72, R73, R83, RX3, E01, E02, P01, or P02.

Table 12-2. 9672 Parallel Enterprise Server R1/R2/R3 Models Upgrade Progression

Model	Upgrade Model
R11	R(2-6)1; R(1-5,7)2; R(5-8,X)3; E0(1-2); P0(1-2) (See notes 1 and 2)
R21	R(3-6)1; R(2-5,7)2; R(5-8,X)3; E0(1-2); P0(1-2) (See notes 1 and 2)
R31	R(4-6)1; R(3-5,7)2; R(5-8,X)3; E0(1-2); P0(1-2) (See notes 1 and 2)
R41	R(5-6)1; R(3-5,7)2; R(5-8,X)3; E0(1-2); P0(1-2) (See notes 1 and 2)
R51	R61; R(4-5,7)2; R(5-8,X)3; E0(1-2); P0(1-2) (See notes 1 and 2)
R61	R(4-5,7)2; R(5-8,X)3; E0(1-2); P0(1-2) (See notes 1 and 2)
RA2	R(1-5,7)2; R(6-8,X)3 (See notes 3 and 5)
R12	R(2-5,7)2; R(7-8,X)3 (See notes 3 and 5)
R22	R(3-5,7)2; R(6-8,X)3 (See notes 3 and 5)
R32	R(4-5,7)2; R(6-8,X)3 (See notes 3 and 5)
R42	R(5,7)2; R(6-8,X)3 (See notes 3 and 5)
R52	R72; R(6-8,X)3 (See notes 3 and 5)
R72	R(8,X)3 (See notes 3 and 5)
R53	R(6-8,X)3 (See note 4)
R63	R(7-8,X)3 (See note 4)
R73	R(8,X)3 (See note 4)
R83	RX3 (See note 4)

See Chapter 3, “Hardware Description” on page 3-1 for a description of the hardware configuration for each 9672 Parallel Enterprise Server model number.

12.1.3 9674 Model C01 Upgrade Progression

The following table provides the 9674 Coupling Facility C01/C02/C03 Models upgrade paths.

Example: Model C01 can be upgraded to model C02 or C03.

Table 12-3. 9674 Coupling Facility C01/C02/C03 Models Upgrade Progression.

Model	Upgrade Model
C01	C0(2-3) (See notes 1 and 2)
C02	C03 (See note 3)

You can upgrade the processing capability of the 9674 Coupling Facility C01/C02/C03 Models by adding central processors (CPs). Additional CPs are available as optional features. Upgrades within C03 models will be available 4Q95. Upgrades within C02 models will be available 1Q96. 9674 Coupling Facility C01/C02/C03 Models Models can have from one to eight, or ten CPs, depending on the model.

See Chapter 3, “Hardware Description” on page 3-1 for a description of the hardware configuration for each 9674 Coupling Facility model number.

Glossary of Terms and Abbreviations

This glossary includes terms and definitions from:

- The *Dictionary of Computing*, SC20-1699.
- The *American National Standard Dictionary for Information Systems*, ANSI X3.172-1990, copyright 1990 by the American National Standards Institute (ANSI). Copies can be purchased from the American National Standards Institute, 1430 Broadway, New York, New York 10018. Definitions are identified by the symbol (A) after the definition.
- The *ANSI/EIA Standard - 440A: Fiber Optic Terminology*, copyright 1989 by the Electronics Industries Association (EIA). Copies can be purchased from the Electronic Industries Association, 2001 Pennsylvania Avenue N.W., Washington, D.C. 20006. Definitions are identified by the symbol (E) after the definition.
- The *Information Technology Vocabulary*, developed by Subcommittee 1, Joint Technical Committee 1, of the International Organization for Standardization and the International Electrotechnical Commission (ISO/IEC JTC1/SC1). Definitions of published parts of this vocabulary are identified by the symbol (I) after the definition; definitions taken from draft international standards, committee drafts, and working papers being developed by ISO/IEC JTC1/SC1 are identified by the symbol (T) after the definition, indicating that final agreement has not yet been reached among the participating National Bodies of SC1.

The following cross-references are used in this glossary:

Contrast with. This refers to a term that has an opposed or substantively different meaning.

See. This refers the reader to multiple-word terms in which this term appears.

See also. This refers the reader to terms that have a related, but not synonymous, meaning.

Synonym for. This indicates that the term has the same meaning as a preferred term, which is defined in the glossary.

A

abend. Abnormal end of task.

abnormal end of task (abend). Ending a task before its completion because of an error condition that cannot be resolved by recovery facilities while the task is being executed.

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 System/370 and ESA/390 modes, an address that exists after translation and prefixing, but before configuration occurs. See also *logical address*, *physical address*, *real address*, and *virtual address*.

action. One of the defined tasks that an application performs. Actions modify the properties of an object or manipulate the object in some way.

action bar. The area at the top of the primary window that contains keywords that give users access to actions available in that window. After users select a choice in the action bar, a pull-down menu appears from the action bar.

activate 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. Contrast with *saved configuration*.

active subchannel. A subchannel that is locked and either busy or has a pending interrupt, and is indicated by subchannel status word (SCSW) bit 24 equals 1. The control information resides in the channel subsystem because it is necessary for the current operation. Contrast with *inactive subchannel*. See also *busy subchannel*.

Note: An active subchannel can also reside in the local working storage of an IOP or channel.

active window. The window that users are currently interacting with. This is the window that receives keyboard input.

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.

address translation. (1) The process of changing the address of an item of data or the address of an instruction to the address in main storage at which it is to be loaded or relocated. (2) In virtual storage systems, the process of changing the address of an item of data or an instruction from its virtual storage address to its real storage address. See also *dynamic address translation*.

ADMF. Asynchronous Data Mover Facility.

ADMOD**. EXIDE Electronics** Auxiliary DC Power MODule.

alert. (1) A unit of information, usually indicating the loss of a system resource, passed from one machine or program to a host to signal an error. (2) An error message sent to the system services control point (SSCP) at the host system.

allocate. To assign a resource, such as a disk or a diskette file, to perform a task. Contrast with *deallocate*.

American National Standard Code for Information Interchange (ASCII). The standard code, using a coded character set consisting of 7-bit coded characters (8 bits including parity), used for information interchange among data processing systems, data communication systems, and associated equipment. The ASCII set consists of control characters and graphics characters. (A)

Note: IBM has defined an extension to ASCII code (characters 128–255).

ANSI. American National Standards Institute.

AOC/MVS. automated operations control/multiple virtual storage.

APAR. Authorized program analysis report.

API. application programming interface

application. (1) The use to which an information processing system is put, for example, a payroll application, an airline reservation application, a network application. (2) A collection of software components used to perform specific types of work on a computer.

application program. (1) A program that is specific to the solution of an application problem. (T) (2) A program written for or by a user that applies to the user's work, such as a program that does inventory control or payroll. (3) A program used to connect and communicate with stations in a network, enabling users to perform application-oriented activities.

architected. Defined in *IBM System/370 Principles of Operation*, GA22-7000, or *IBM Enterprise Systems Architecture/390 Principles of Operation*, SA22-7201.

ASCII. American National Standard Code for Information Interchange.

asynchronous. (1) Pertaining to two or more processes that do not depend upon the occurrence of specific events such as common timing signals. (T) (2) Without regular time relationship; unexpected or unpredictable with respect to the execution of program instructions. Contrast with *synchronous*.

authorized program analysis report (APAR). A request for correction of a problem caused by a defect in a current release of a program unaltered by the user.

auto-answer. In data communication, the ability of a station to respond automatically to a call that it receives over a switched line.

auto-call. In data communication, the ability of a station to initiate a call automatically over a switched line.

B

basic mode. A central processor mode that does not use logical partitioning. Contrast with *logically partitioned (LPAR) mode*.

batch. (1) An accumulation of data to be processed. (2) A group of records or data processing jobs brought together for processing or transmission. (3) Pertaining to activity involving little or no user action. Contrast with *interactive*.

BL. Parallel block multiplexer channel.

block. A string of data elements recorded or transmitted as a unit. The element may be characters, words, or physical records. (T)

block multiplexer channel. A multiplexer channel that interleaves blocks of data. Contrast with *selector channel*. See also *byte multiplexer channel.*

BOC. Battery-operated clock.

buffer. (1) A routine or storage used to compensate for a difference in rate of flow of data, or time of occurrence of events, when transferring data from one device to another. (A) (2) To allocate and schedule the use of buffers. (A) (3) A portion of storage used to hold input or output data temporarily. See *fiber buffer*.

burst. In data communication, a sequence of signals counted as one unit in accordance with some specific criterion or measure. (A)

bus. (1) A facility for transferring data between several devices located between two end points, only one

device being able to transmit at a given moment. (T) (2) A network configuration in which nodes are interconnected through a bidirectional transmission medium. (3) One or more conductors used for transmitting signals or power. (A)

bus-in data. Two digits of information transmitted from a control unit to a parallel channel. The information is data, an address, or status.

busy subchannel. A subchannel that is either active and locked, inactive with an I/O operation pending or in process, or between commands in a multiplex operation. See also *active subchannel* and *inactive subchannel*.

BY. Parallel byte multiplexer channel.

byte. (1) A string that consists of a number of bits, treated as a unit, and representing a character. (T) (2) A binary character operated upon as a unit and usually shorter than a computer word. (A) (3) A string that consists of a particular number of bits, usually eight, that is treated as a unit, and that represents a character. (4) A group of eight adjacent binary digits that represent one extended binary-coded decimal interchange code (EBCDIC) character.

byte multiplexer channel. A multiplexer channel that interleaves bytes of data. Contrast with *selector channel*. See also *block multiplexer channel.*

C

CA. channel adapter (card).

cache. (1) A special purpose buffer storage, smaller and faster than main storage, used to hold a copy of the instructions and data obtained from main storage and likely to be needed next by the processor. (T) (2) A buffer storage that contains frequently accessed instructions and data; it is used to reduce access time.

cache structure. A coupling facility structure that enables high-performance sharing of cached data by multisystem applications in a sysplex. Applications can use a cache structure to implement several different types of caching systems, including a store-through or a store-in cache.

CAW. channel address word.

CBY. ESCON byte multiplexer channel.

CCC. Channel control check.

CCE. channel control element.

CCH. Channel-check handler.

CCW. Channel command word.

CDC. Channel data check.

central processor (CP). The part of the computer that contains the sequencing and processing facilities for instruction execution, initial program load, and other machine operations.

central processor complex (CPC). The boundaries of a system, exclusive of I/O control units and devices, that can be controlled by a single operating system. A CPC consists of main storage, one or more central processing units, time-of-day clocks, and channels, which are or can be placed in a single configuration. A CPC also includes channel subsystems, service processors, and expanded storage where installed.

central storage. Storage that is an integral part of the processor and includes both main storage and the hardware system area.

CF. coupling facility (CF).

CFS. coupling facility sender channel.

CFR. coupling facility receiver channel.

CFRM. coupling facility resource management (policy).

change. An alteration (addition, deletion, or modification) of one or more information system components, of one of the following types: hardware (may include internal code), or software (system or application). The term *change* also refers to an SNA/File Services data object containing internal code, internal code customizing data, software, software customizing data, applications data, procedures, or documentation.

change management. The management discipline that defines the process of planning (for example, scheduling) and controlling (for example, distributing, installing, and tracking) changes in an SNA network.

channel. (1) A path along which signals can be sent, for example, input/output channel. (2) The system element that controls one channel path, whose mode of operation depends on the type of hardware to which it is attached.

channel adapter. (1) A communication controller hardware unit used to attach the controller to a data channel. (2) Hardware that attaches a group of channels to the secondary data stager and prioritizes and stages data between the channels and the channel control element.

channel address. In System/370 mode, the 8 leftmost bits of an input/output address that identify the channel. See also *device address* and *input/output address*.

channel address word (CAW). An area in storage that specifies the location in main storage at which a channel program begins.

channel-attached. (1) Pertaining to attachment of devices directly by data channels (I/O channels) to a computer. (2) Pertaining to devices attached to a controlling unit by cables rather than by telecommunication lines. Contrast with *link-attached.*

channel command word (CCW). A doubleword at the location in main storage specified by the channel address word. One or more CCWs make up the channel program that directs data channel operations.

channel control check. A category of I/O errors affecting channel controls and sensed by the channel to which a device is attached. See also *channel data check*.

channel control element (CCE). In a channel subsystem, hardware that consists of an I/O processor, a primary data stager, and one or two secondary data stagers.

channel data check. A category of I/O errors, indicating a machine error in transferring data to or from storage and sensed by the channel to which a device is attached. See also *channel control check*.

channel data rate. The rate at which a channel can move data between a transmission link and processor storage during the data transfer portion of an I/O operation.

channel Licensed Internal Code. That part of the channel subsystem Licensed Internal Code used to start, maintain, and end all operations on the I/O interface. See also *IOP Licensed Internal Code*.

channel path (CHP). A single interface between a central processor and one or more control units along which signals and data can be sent to perform I/O requests.

channel path configuration. In an ESCON environment, the connection between a channel and a control unit or between a channel, an ESCON Director, and one or more control units. See also *link, point-to-point channel path configuration,* and *switched point-to-point channel path configuration*.

channel path identifier (CHPID). In a channel subsystem, a value assigned to each installed channel path of the system that uniquely identifies that path to the system.

channel set. In System/370 mode, a collection of channels that can be addressed concurrently by a central processor. See also *channel subsystem*.

channel status word (CSW). An area in storage that provides information about the termination of input/output operations.

channel subsystem (CSS). A collection of subchannels that directs the flow of information between I/O devices and main storage, relieves the processor of communication tasks, and performs path management functions.

channel subsystem (CSS) Licensed Internal Code. Code that consists of the IOP Licensed Internal Code and the channel Licensed Internal Code.

channel-to-channel (CTC). Refers to the communication (transfer of data) between programs on opposite sides of a channel-to-channel adapter (CTCA).

channel-to-channel adapter (CTCA). An input/output device that is used by a program in one system to communicate with a program in another system.

check stop. The state that occurs when an error makes it impossible or undesirable to continue the operation in progress.

CHD. channel driver (card).

CHN. Channel.

choice. An item that users may select. Choices appear in the selection fields, action bars, and pull-down menus.

CHP. Channel path.

CHPID. See channel path identifier.

CI. console integration.

CICS. Customer Information Control System.

CICS/ESA. Customer Information Control System/Enterprise Systems Architecture.

CICSplex. A group of connected CICS regions.

CKD. count key data.

click. To press and release a mouse button without moving the mouse pointer off the choice.

CLIST (command list). A data set in which commands and possibly subcommands and data are stored for subsequent execution.

closely coupled. A multisystem structure that requires a high degree of interaction and cooperation between multiple MVS systems to process a workload. See also *loosely coupled, tightly coupled*.

CMOS. Complementary metal-oxide semiconductor.

CMS. Conversational monitor system.

CNC. Mnemonic for an ESCON channel attached to an ESCON-capable device.

command. (1) A character string from a source external to a system that represents a request for system action. (2) A request from a terminal for performance of an operation or execution of a program. (3) A value sent on an I/O interface from a channel to a control unit that specifies the operation to be performed.

command area. The area on a screen that contains the command entry field.

command chaining. The fetching of a new channel command word (CCW) immediately following the completion of the previous CCW.

command entry field. An entry field in which a user types commands.

command list. See *CLIST*.

command retry. A channel and control unit procedure that causes a command to be retried without requiring an I/O interrupt.

communication control unit. A communication device that controls transmission of data over lines in a network.

communication controller. (1) A device that directs the transmission of data over the data links of a network; its operation can be controlled by a program executed in a processor to which the controller is connected or it may be controlled by a program executed within the device. (T) (2) A type of communication control unit whose operations are controlled by one or more programs stored and executed in the unit. It manages the details of line control and the routing of data through a network.

complementary metal-oxide semiconductor (CMOS). A technology that combines the electrical properties of positive and negative voltage requirements to use considerably less power than other types of semiconductors.

concurrent maintenance. Hardware maintenance actions performed by a service representative while normal operations continue without interruption. See also *nondisruptive installation* and *nondisruptive removal*.

configuration. (1) The arrangement of a computer system or network as defined by the nature, number, and the chief characteristics of its functional units. More specifically, the term configuration may refer to a hardware configuration or a software configuration. (I) (A) (2) In an ESCON Director, the physical connection capability determined by a set of attributes. The attribute values specify the connectivity control status and identifiers associated with the ESCD and its ports. See also *active configuration, configuration matrix, connectivity attribute,* and *saved configuration*.

configure. To describe to the system the devices and optional features installed on the system.

connectivity. A term used to describe the physical interconnections of multiple devices/computers/networks employing similar or different technology or architecture together to accomplish effective communication between and among connected members involving data exchange or resource sharing.

console. A logical device used for communication between the user and the system. (A) See *display station, monitor console, operator console, program mode console, programming support console, service console,* and *system console*.

console integration (CI). The hardware and software facilities used to bring operating systems management and hardware systems management under a single control point.

control program. A computer program designed to schedule and to supervise the execution of programs of a computer system. (I) (A)

control unit. A hardware unit that controls the reading, writing, or displaying of data at one or more input/output units.

control unit data rate. The rate at which a control unit can move data between itself and a transmission link during the data transfer portion of an I/O operation.

control-unit end. In I/O operations, a signal from a control unit to the channel indicating that the control unit is no longer needed for the operation.

control unit header (CUH). See *logical control unit (LCU)*.

controller. A unit that controls input/output operations for one or more devices.

conversational monitor system (CMS). A virtual machine operating system that provides general interactive time sharing, problem solving, and program development capabilities, and operates only under the VM control program.

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 directly attached to it.

coupling services. In a sysplex, the functions of XCF that transfer data and status between members of a group residing on one or more MVS systems in the sysplex.

CP. (1) Control program. (2) Central processor.

CPC. Central processor complex.

CPCID. CPC identifier.

CPC Image. The set of CPC resources that support a single control program.

CPU. central processing unit.

CPUID. CPU identifier.

critical resource. A resource required for system operation. See also *system resource*.

cross-system coupling facility (XCF). XCF is a component of MVS that provides functions to support cooperation between authorized programs running within a sysplex.

CRW. Channel report word.

CS. (1) Central storage. (2) Channel set. (3) Control storage.

CSS. Channel subsystem.

CSW. Channel status word.

CTC. (1) Channel-to-channel. (2) Mnemonic for an ESCON channel attached to another ESCON channel.

CTCA. channel-to-channel adapter.

CU. Control unit.

CUA. Control unit address.

CUADD. Control unit logical address.

CUCW. Control unit control word.

CUH. Control unit header.

cursor. (1) A movable, visible mark used to indicate the position at which the next operation will occur on a display screen. (A) (2) A visual cue that shows the user where keyboard input will appear on the screen.

Customer Information Control System (CICS). An IBM licensed program that enables transactions entered at remote terminals to be processed concurrently by user-written application programs. It includes facilities for building, using, and maintaining data bases.

customize. To change a data processing installation or network to meet the needs of particular users.

CVC. Mnemonic for an ESCON channel attached to a 9034.

D

DASD. See direct access storage device.

DASD subsystem. A storage control and its attached direct access storage devices.

DAT. Dynamic address translation.

data processing (DP). The systematic performance of operations upon data; for example, arithmetic or logic operations upon data, merging or sorting of data, assembling or compiling of programs. (T)

data rate. See *channel data rate, control unit data rate, device data rate, effective data rate,* and *input/output subsystem data rate*. See also *link rate*.

data sharing. The ability of concurrent subsystems (such as DB2 or IMS DB) or application programs to directly access and change the same data while maintaining data integrity.

data streaming. In an I/O interface, a mode of operation that provides a method of data transfer at up to 4.5 megabytes per second. Data streaming is not interlocked between the sender and the receiver. Once data transfer begins, the sender does not wait for acknowledgment from the receiver before sending the next byte. The control unit determines the data transfer rate. Contrast with *direct-coupled interlock (DCI)*.

data transfer. (1) The result of the transmission of data signals from any data source to a data receiver. (2) The movement, or copying, of data from one location and the storage of the data at another location.

data transfer mode. The method of information exchange used on an I/O interface. See *direct-coupled interlock (DCI)* and *data streaming*.

DB2. DATABASE 2.

DCAF. distributed console access facility.

DCI. direct-coupled interlock.

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*.

device. A mechanical, electrical, or electronic contrivance with a specific purpose.

device address. In System/370 mode, the 8 rightmost bits of an I/O address that identify a particular I/O device and a control unit on the designated channel. See *channel address, device-level addressing,* and *input/output address*.

DFSMS. Data Facility Storage Management Subsytem.

direct access storage. A storage device that provides direct access to data. (I) (A) See also *random access memory*.

direct access storage device (DASD). (1) A storage device in which the location of each data record can be directly addressed. (2) A device in which the access time is effectively independent of the location of the data. (Restriction: Does not refer to diskette drive.)

direct-coupled interlock (DCI). In an I/O interface, a mode of operation that provides a method of data transfer at up to 1.5 megabytes per second. DCI protocol requires the sender to raise a signal on the interface along with the byte of data being transferred. This signal and the accompanying data must be maintained until the receiver of the data sends a signal acknowledging that the data has been received. Contrast with *data streaming*.

DP. data processing.

dynamic address translation (DAT). In virtual storage systems, the change of a virtual storage address to a real storage address during execution of an instruction. See also *address translation*.

dynamic reconfiguration management. In MVS, the ability to modify the I/O configuration definition without needing to perform a power-on reset (POR) of the hardware or an initial program load (IPL).

dynamic storage reconfiguration. A PR/SM LPAR function that allows central or expanded storage to be added or removed from a logical partition without disrupting the system control program operating in the logical partition.

E

EC. engineering change.

ECC. error checking and correction.

ECKD. extended count key data.

EEPROM. electrically erasable programmable read only memory.

EIA. Electronics Industries Association.

EMIF. ESCON multiple image facility.

Enterprise Systems Architecture/390 (ESA/390). IBM ESA/390 principles of operation.

Enterprise Systems Connection (ESCON). A set of products and services that provides a dynamically connected environment using optical cables as a transmission medium.

EPO. Emergency power off.

EREP. Environmental error record editing and printing program.

error checking and correction (ECC). In a processor, the detection and correction of all single-bit errors, plus the detection of double-bit and some multiple-bit errors.

ESA. (1) Enterprise Systems Architecture. (2) Expanded storage array.

ESA/370. Enterprise Systems Architecture/370.

ESA/390. Enterprise Systems Architecture/390.

ESCD. Enterprise Systems Connection (ESCON) Director.

ESCD console. The ESCON Director input/output device used to perform operator and service tasks at the ESCD.

ESCM. Enterprise Systems Connection Manager.

ESCON. Enterprise Systems Connection.

ESCON channel. A channel having an Enterprise Systems Connection channel-to-control-unit I/O interface that uses optical cables as a transmission medium. Contrast with *parallel channel*.

ESCON Director (ESCD). A device that provides connectivity capability and control for attaching any two links to each other.

ESCON environment. The data processing environment having an Enterprise Systems Connection

channel-to-control-unit I/O interface that uses optical cables as a transmission medium.

ESCON Manager (ESCM). A licensed program that provides host control and intersystem communication capability for ESCON Director connectivity operations.

ESCON multiple image facility (EMIF). A facility that allows channels to be shared among PR/SM logical partitions in an ESCON environment.

Ethernet. A communication network (USA, Xerox 1975).

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 4KB pages to and from central storage.

F

facility. (1) An operational capability, or the means for providing such a capability. (T) (2) A service provided by an operating system for a particular purpose; for example, the checkpoint/restart facility.

FDDI. Fiber Distributed Data Interface (100Mbits/s fiber optic LAN).

FEID. functional element identifier.

fiber. See *optical fiber*.

frame. (1) For a System/390 microprocessor cluster, a frame contains one or two central processor complexes (CPCs), support elements, and AC power distribution. (2) A housing for machine elements. (3) The hardware support structure, covers, and all electrical parts mounted therein that are packaged as one entity for shipping. (4) A formatted display. See *display frame* and *transmission frame*.

FRU. Field-replaceable unit.

ft. Foot.

function key. (1) In computer graphics, a button or switch that may be operated to send a signal to the computer program controlling the display. (T) (2) A key that, when pressed, performs a specified set of operations.

functional element identifier (FEID). An alphanumeric character or characters identifying a component or element in service language or other diagnostic application programs.

G

GB. Gigabyte.

gigabyte. (1) A unit of measure for storage size. One gigabyte equals 1 073 741 824 bytes. (2) Loosely, one billion bytes.

GMT. Greenwich mean time.

graying. The indication that a choice in a pull-down menu is unavailable for selection.

H

Hardware Management Console. A console used to monitor and control hardware such as the System/390 microprocessors.

hardware system area (HSA). A logical area of central storage, not addressable by application programs, used to store Licensed Internal Code and control information.

HCD. hardware configuration definition.

HSA. hardware system area.

HMCA. Hardware Management Console Application.

I

IBB. Internal Bus Buffer

ICMF. Integrated Coupling Migration Facility.

IDAW. indirect data address word.

IFCC. interface control check.

IML. Initial machine load.

IMS. Information Management System.

initial machine load (IML). A procedure that prepares a device for use.

initial power controller (IPC). Logic in the processor controller that controls the base power of both the primary support processor (PSP) and the input/output support processor (IOSP).

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 attach to the processor complex.

input/output configuration data set (IOCDS). The data set that contains an I/O configuration definition built by the I/O configuration program (IOCP).

input/output configuration program (IOCP). A program that defines to a system all the available I/O devices and the channel paths.

input/output configuration source. The file used as input to the IOCP when an IOCDS is created. It is an 80-column card-image, and is initially provided in EBCDIC form on a system tape, or in ASCII form on a 3.5-inch diskette.

interrupt. (1) A suspension of a process, such as execution of a computer program caused by an external event, and performed in such a way that the process can be resumed. (A) (2) To stop a process in such a way that it can be resumed. (3) In data communication, to take an action at a receiving station that causes the sending station to end a transmission. (4) To temporarily stop a process.

I/O. See input/output.

IOCDS. I/O configuration data set.

IOCP. I/O configuration program.

IODF. input/output definition file.

IOP. Integrated offload processor.

IOP Licensed Internal Code. The part of the channel subsystem Licensed Internal Code used to start, maintain, and end (normally or abnormally) all operations with the central processor. See also *channel Licensed Internal Code*.

IOPS. I/O power sequencer.

IPC. initial power controller.

IPL. See initial program load.

K

KB. (1) kilobyte.

kilobyte. (1) A unit of measure for storage size. (2) Loosely, one thousand bytes.

L

LAN. See local area network.

laser. A device that produces optical radiation using a population inversion to provide light amplification by stimulated emission of radiation and (generally) an optical resonant cavity to provide positive feedback. Laser radiation can be highly coherent temporally, or spatially, or both. (E)

LCU. Logical control unit.

LED. Light-emitting diode.

LIC. Licensed Internal Code.

Licensed Internal Code (LIC). Software provided for use on specific IBM machines and licensed to customers under the terms of IBM's Customer Agreement. Microcode can be Licensed Internal Code and licensed as such.

light-emitting diode (LED). A semiconductor chip that gives off visible or infrared light when activated.

local area network (LAN). A computer network located on a user's premises within a limited geographical area. Communication within a local area network is not subject to external regulations; however, communication across the LAN boundary can be subject to some form of regulation.

Note: A LAN does not use store and forward techniques.

logical address. The address found in the instruction address portion of the program status word (PSW). If translation is off, the logical address is the real address. If translation is on, the logical address is the virtual address. See also *absolute address*, *physical address*, *real address*, and *virtual address*.

logical control unit. A group of contiguous words in the hardware system area that provides all of the information necessary to control I/O operations through a group of paths that are defined in the IOCDS. Logical

control units represent to the channel subsystem a set of control units that attach common I/O devices.

logical partition (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, central processor resources defined to operate in a logical partition like a physical central processor.

logical unit (LU). In SNA, a port to the network through which an end user accesses the SNA network and the functions provided by system services control points (SSCPs). An LU can support at least two sessions–one with an SSCP and one with another LU–and may be capable of supporting many sessions with other LUs.

logically partitioned (LPAR) mode. A central processor complex (CPC) power-on reset mode that enables use of the PR/SM feature and allows an operator to allocate CPC hardware resources (including central processors, central storage, expanded storage, and channel paths) among logical partitions. Contrast with *basic mode*.

LP. Logical partition.

LPAR. See logically partitioned mode.

LU. logical unit.

M

main storage. A logical entity that represents the program addressable portion of central storage. See central storage.

main storage. (1) Program-addressable storage from which instructions and other data can be loaded directly into registers for subsequent processing. (I) (A) (2) That part of internal storage into which instructions and other data must be loaded for subsequent execution or processing. (3) The part of a processor unit where programs are run. See *central storage*.

Notes:

1. Main storage refers to the whole program-addressable execution space and can include one or more storage devices.
2. The term *main storage* is generally used in large and intermediate computers. The term *memory* is primarily used in microcomputers, calculators, and some minicomputers.

maintenance change level (MCL). A change to correct a single licensed internal code design defect. Higher quality than a *patch*, and intended for broad distribution. Considered functionally equivalent to a software PTF.

MAU. multistation access unit.

Mb. Megabit.

MB. Megabyte.

MBA. memory bus adapter.

MCCU. multisystem channel communication unit.

MCL. See maintenance change level.

MEC. Machine engineering change.

megabit (Mb). A unit of measure for throughput. One megabit equals 1 000 000 bits.

megabyte (MB). (1) A unit of measure for storage size. One megabyte equals 1 048 576 bytes. (2) Loosely, one million bytes.

modem (modulator/demodulator). A device that converts digital data from a computer to an analog signal that can be transmitted on a telecommunication line, and converts the analog signal received to data for the computer.

mouse. In computer graphics, a pointing device operated by moving it on a surface.

multiplexer channel. A channel designed to operate with a number of I/O devices simultaneously. Several I/O devices can transfer records at the same time by interleaving items of data. See *block multiplexer channel* and *byte multiplexer channel*.

MVS. Multiple Virtual Storage.

MVS/ESA. Multiple Virtual Storage/Enterprise Systems Architecture.

MVS image. A single occurrence of the MVS/ESA operating system that has the ability to process work.

MVS/SP. Multiple Virtual Storage/System Product.

MVS system. An MVS image together with its associated hardware, which collectively are often referred to simply as a system, or MVS system.

N

NetBIOS. local area network basic input/output system.

NetView. An IBM licensed program used to monitor a network, manage it, and diagnose its problems.

network. (1) An arrangement of nodes and connecting branches. (T) (2) A configuration of data processing devices and software connected for information exchange.

NIC. numerically intensive computing.

O

object-action. A process sequence in which users select an object and then select an action to apply to that object.

operate. To do a defined action, such as adding or comparing, performed on one or more data items.

operating system (OS). Software that controls the execution of programs and that may provide services such as resource allocation, scheduling, input/output control, and data management. Although operating systems are predominantly software, partial hardware implementations are possible. (T)

operator console. (1) A functional unit containing devices that are used for communications between a computer operator and a computer. (T) (2) A display used for communication between the operator and the system, used primarily to specify information concerning application programs and I/O operations and to monitor system operation.

optical cable. A fiber, multiple fibers, or a fiber bundle in a structure built to meet optical, mechanical, and environmental specifications. (E) See also *jumper cable, optical cable assembly,* and *trunk cable*.

optical fiber. Any filament made of dielectric materials that guides light, regardless of its ability to send signals. (E) See also *fiber optics* and *optical waveguide*.

OS. Operating system.

OSA. Open Systems Adapter. A new connectivity option that allows direct attachment of Ethernet, Token Ring, and FDDI LANs to the central processor complex.

OS/2. Operating system/2.

P

parallel channel. (1) A channel having a System/360 and System/370 channel-to-control-unit I/O interface that uses bus-and-tag cables as a transmission medium. Contrast with *ESCON channel*. (2) A data path along which a group of signals representing a character or any other entity of data can be sent simultaneously.

PD. problem determination.

point-to-point channel path configuration. In an I/O interface, a configuration that consists of a single link between a channel and one control unit. Contrast with *switched point-to-point channel path configuration*.

point-to-point connection. A connection established between two data stations for data transmission. Contrast with *multipoint connection*.

Note: The connection may include switching facilities.

pop-up window. A movable window, fixed in size, that provides information that is required by an application so it can continue a user request.

POR. Power-on reset.

power-on reset. A function that re-initializes all the hardware in the system and loads the internal code that enables the machine to load and run an operating system. This function is intended as a recovery function.

power-on reset state. The condition after a machine power-on sequence and before an IPL of the control program.

problem analysis. The process of identifying the source of a problem; for example, a program component, a machine failure, telecommunication facilities, user or contractor-installed programs or equipment, an environment failure such as a power loss, or a user error.

problem determination (PD). The process of determining the source of a problem; for example, a program component, machine failure, telecommunication facilities, user or contractor-installed programs or equipment, environmental failure such as a power loss, or user error.

problem management. The management discipline that handles a problem from its detection through its final resolution. Problem management is composed of the following:

Problem determination
Problem diagnosis
Problem bypass and recovery
Problem resolution
Problem tracking and control.

processor. (1) In a computer, a functional unit that interprets and executes instructions. A processor consists of at least an instruction control unit and an arithmetic and logic unit. (T) (2) The functional unit that interprets and executes instructions. (3) The boundaries of a system, exclusive of I/O control units and devices, that can be controlled by a single operating system. A processor consists of main storage, one or more central processors, time-of-day clocks, and chan-

nels, which are, or can be, placed in a single configuration. A processor also includes channel subsystems, and expanded storage where installed.

processor complex. A system configuration that consists of all the machines required for operation; for example, a processor unit, a processor controller, a system display, a service support display, and a power and coolant distribution unit.

processor console. The workstation from which an operator can control and observe system hardware operation. See also system console.

Processor Resource/Systems Manager (PR/SM) facility. The feature that allows the processor to use several system control programs (SCPs) simultaneously, provides logical partitioning capability for the real machine, and provides support for multiple preferred guests.

profile. (1) A description of the characteristics of an entity to which access is controlled. (2) Data that describes the significant characteristics of a user, a group of users, or one or more computer resources.

program. Sequence of instructions for a computer. A program interacts and relies on either the hardware or other programs.

program status word (PSW). An area in storage used to indicate the sequence in which instructions are executed, and to hold and indicate the status of the computer system.

program temporary fix (PTF). A temporary solution or bypass of a problem diagnosed by IBM as resulting from a defect in a current, unaltered release of the program.

PR/SM. Processor Resource/Systems Manager.

PSP. preventive service planning.

PSW. Program status word.

PTF. Program temporary fix.

R

RAS. reliability, availability, serviceability.

reconfiguration. (1) A change made to a given configuration in a computer system; for example, isolating and bypassing a defective functional unit or connecting two functional units by an alternative path. Reconfiguration is effected automatically or manually and can be used to maintain system integrity. (T) (2) The process of placing a processor unit, main storage, and channels offline for maintenance, and adding or removing components. (3) Contrast with *nondisruptive installation* and *nondisruptive removal*.

recovery. To maintain or regain system operation after a failure occurs. Generally, to recover from a failure is to identify the failed hardware, to deconfigure the failed hardware, and to continue or restart processing.

remote service facility (RSF). (1) A control program plus associated communication equipment that allows local personnel to connect to an IBM service center, and allows remote personnel to operate the remote system or send new internal code fixes to it, if properly authorized. (2) A system facility invoked by Licensed Internal Code that provides procedures for problem determination and error detection.

Remote Technical Assistance and Information Network (RETAIN). A database, accessible to service representatives, of information relating to IBM-installed products.

RETAIN. Remote Technical Assistance and Information Network.

return code. (1) A code used to influence the execution of succeeding instructions. (A) (2) A value returned to a program to indicate the results of an operation requested by that program.

REXX. restructured extended executor language.

ring network. A network configuration in which devices are connected by unidirectional transmission links to form a closed path.

Note: A ring of an IBM Token-Ring Network is referred to as a LAN segment or as a Token-Ring Network segment.

RMF. Resource Measurement Facility.

RPQ. request for price quotation.

RPS. rotational positional sensing/sensor.

RSF. Remote support facility.

S

SAD. system activity display.

SCP. system control program.

scroll. To move a display image vertically or horizontally to display data that cannot be observed within the boundaries of the display frame.

scroll bar. A window component associated with a scrollable area that provides users a visual cue that more information is available and that the unseen infor-

mation can be manipulated into view using the mouse. Users scroll the information in the window by interacting with the scroll bar.

SCS. single channel service.

SDLC. See synchronous data link control.

SE. support element.

service representative. A person who performs maintenance services for IBM hardware products or systems. See also *IBM program support representative*.

single channel service (SCS). Testing a single channel while the other channels remain operational. See also *concurrent maintenance*.

single point of control. The characteristic a sysplex displays when you can accomplish a given set of tasks from a single workstation, even if you need multiple IBM and vendor products to accomplish that particular set of tasks.

single system image. The characteristic a product displays when multiple images of the product can be viewed and managed as one image.

SNA. See systems network architecture.

SNA/Management Services (SNA/MS). Functions distributed among network components to operate, manage, and control the network.

SNA/MS. SNA/Management Services.

SNA network. The part of a user-application network that conforms to the formats and protocols of Systems Network Architecture. It enables reliable transfer of data among end-users and provides protocols for controlling the resources of various network configurations. The SNA network consists of network addressable units (NAUs), boundary function components, and the path control network.

SNMP. simple network management protocol.

storage. (1) A functional unit into which data can be placed, in which they can be retained, and from which they can be retrieved. (T) (2) The action of placing data into a storage device. (I) (A)

structure. A construct used by MVS to map and manage storage on a coupling facility. See cache structure, list structure, and lock structure.

subchannel. In 370-XA 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. See also *channel path identifier, device identifier*, and *device number*.

subsystem. A secondary or subordinate system, or programming support, usually capable of operating independently of or asynchronously with a controlling system. (T) See *DASD subsystem* and *storage subsystem*.

subsystem storage. Synonym for *cache*.

support element (SE). (1) an internal control element of a processor that assists in many of the processor operational functions. (2) A hardware unit that provides communications, monitoring, and diagnostic functions to a central processor complex (CPC).

switched line. In data communication, a temporary connection between computers or devices that is established by dialing. Contrast with nonswitched line.

synchronous data link control (SDLC). A form of communication line control that uses commands to control data transfer over a communication line.

sysplex. A set of MVS systems communicating and cooperating with each other through certain multisystem hardware components and software services to process customer workloads. See also *MVS system*.

Sysplex Timer. An IBM unit that synchronizes the time-of-day (TOD) clocks in multiple processors or processor sides. External Time Reference (ETR) is the MVS generic name for the IBM Sysplex Timer (9037).

system. Comprises the processor complex and all attached and configured I/O and communication devices.

system area. A logical area of central storage used to store Licensed Internal Code and control information (not addressable by application programs).

Systems Network Architecture (SNA). The description of the logical structure, formats, protocols, and operational sequences for transmitting information units through, and controlling the configuration and operation of, networks.

S/370. IBM System/370

S/390. IBM System/390.

T

target logical partition. In LPAR mode, the target logical partition is the logical partition that is selected on the System Monitor panel. It is the current or immediate logical partition; the logical partition that the operator is working with or looking at.

target processor. The processor that controls execution during a program restart, instruction trace, standalone dump, or IPL, and whose ID is identified by highlighting on the status line.

TCP/IP. Transmission Control Protocol/Internet Protocol.

TIC. transfer in channel.

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.

token. A sequence of bits passed from one device to another on the token-ring network that signifies permission to transmit over the network. It consists of a starting delimiter, an access control field, and an end delimiter. The access control field contains a bit that indicates to a receiving device that the token is ready to accept information. If a device has data to send along the network, it appends the data to the token. When data is appended, the token then becomes a frame.

token-ring network. (1) A ring network that allows unidirectional data transmission between data stations, by a token passing procedure, such that the transmitted data return to the transmitting station. (T) (2) A network that uses ring topology, in which tokens are passed in a circuit from node to node. A node that is ready to send can capture the token and insert data for transmission.

Note: The IBM Token-Ring Network is a baseband LAN with a star-wired ring topology that passes tokens from network adapter to network adapter.

TPF. Transaction processing facility.

transaction processing. In batch or remote batch processing, the processing of a job or job step. In interactive processing, an exchange between a terminal and another device that does a particular action; for example, the entry of a customer's deposit and the updating of the customer's balance.

TSCF. target system control facility.

U

UCW. unit control word.

UPC. universal power controller.

UPS. uninterruptible power supply.

user interface. Hardware, software, or both that allows a user to interact with and perform operations on a system, program, or device.

V

VR. voltage regulator (card).

VSE. virtual storage extended.

VTAM. virtual telecommunications access method.

W

window. (1) An area of the screen with visible boundaries through which information is displayed. A window can be smaller than or equal in size to the screen. Windows can overlap on the screen and give the appearance of one window being on top of another. (2) A choice in the action bar of some applications. Users select it to arrange the display of several windows or to change the active window. (3) A choice in the action bar of multiple-document interface applications. (4) A choice in an action bar that allows the user to arrange the display of all open windows and to change the active window. (5) A choice in the action bar of multiple-document interface applications that allows a user to arrange the display of all open windows and to change the active window.

workstation. (1) A functional unit at which a user works. A workstation often has some processing capability. (T) (2) A terminal or microcomputer, usually one that is connected to a mainframe or network, at which a user can perform applications. (3) See also *display station* and *terminal*.

X

XDF. extended distance feature.

Numerics

370-XA. IBM System/370 extended architecture

Index

Numerics

A

B

C

D

E

F

H

I

K

L

M

O

P

R

S

T

U

V

Readers' Comments — We'd Like to Hear from You

S/390® 9672 Parallel Transaction Server
S/390® 9672 Parallel Enterprise Server
S/390® 9674 Coupling Facility
System Overview
For R1/R2/R3 Based models

Publication No. GA22-7148-06

Overall, how satisfied are you with the information in this book?

	Very Satisfied	Satisfied	Neutral	Dissatisfied	Very Dissatisfied
Overall satisfaction	□	□	□	□	□

How satisfied are you that the information in this book is:

	Very Satisfied	Satisfied	Neutral	Dissatisfied	Very Dissatisfied
Accurate	□	□	□	□	□
Complete	□	□	□	□	□
Easy to find	□	□	□	□	□
Easy to understand	□	□	□	□	□
Well organized	□	□	□	□	□
Applicable to your tasks	□	□	□	□	□

Please tell us how we can improve this book:

Thank you for your responses. May we contact you? □ Yes □ No

When you send comments to IBM, you grant IBM a nonexclusive right to use or distribute your comments in any way it believes appropriate without incurring any obligation to you.

Name _____ Address _____

Company or Organization _____

Phone No. _____

Readers' Comments — We'd Like to Hear from You
GA22-7148-06

Cut or Fold
Along Line

Fold and Tape | **Please do not staple** | Fold and Tape

NO POSTAGE
NECESSARY
IF MAILED IN THE
UNITED STATES

BUSINESS REPLY MAIL

FIRST-CLASS MAIL PERMIT NO. 40 ARMONK, NEW YORK

POSTAGE WILL BE PAID BY ADDRESSEE

International Business Machines Corporation
Information Development
Department C9MG
1701 North Street
ENDICOTT NY 13760-9987

Fold and Tape | **Please do not staple** | Fold and Tape

Cut or Fold
Along Line

Printed in U.S.A.

GA22-7148-06

```
DSMKPO653E POSTSCRIPT FILE 'BPI201A' NOT FOUND.
DSMMOM395I '.EDFPO' LINE 80: .po BPI201A
DSMMOM397I '.EDFPO' WAS IMBEDDED AT LINE 910 OF '.EDFAWRK'
DSMMOM397I '.EDFAWRK' WAS IMBEDDED AT LINE 1420 OF 'DB2TO003'
DSMMOM397I 'DB2TO003' WAS IMBEDDED AT LINE 49 OF 'DB2TOMST'
DSMKPO653E POSTSCRIPT FILE 'db2td030' NOT FOUND.
DSMMOM395I '.EDFPO' LINE 80: .po db2td030
DSMMOM397I '.EDFPO' WAS IMBEDDED AT LINE 910 OF '.EDFAWRK'
DSMMOM397I '.EDFAWRK' WAS IMBEDDED AT LINE 1449 OF 'DB2TO003'
DSMMOM397I 'DB2TO003' WAS IMBEDDED AT LINE 49 OF 'DB2TOMST'
DSMKPO653E POSTSCRIPT FILE 'db2td031' NOT FOUND.
DSMMOM395I '.EDFPO' LINE 80: .po db2td031
DSMMOM397I '.EDFPO' WAS IMBEDDED AT LINE 910 OF '.EDFAWRK'
DSMMOM397I '.EDFAWRK' WAS IMBEDDED AT LINE 1454 OF 'DB2TO003'
DSMMOM397I 'DB2TO003' WAS IMBEDDED AT LINE 49 OF 'DB2TOMST'
DSMKPO653E POSTSCRIPT FILE 'db2td02' NOT FOUND.
DSMMOM395I '.EDFPO' LINE 80: .po db2td02
DSMMOM397I '.EDFPO' WAS IMBEDDED AT LINE 910 OF '.EDFAWRK'
DSMMOM397I '.EDFAWRK' WAS IMBEDDED AT LINE 50 OF 'DB2TO004'
DSMMOM397I 'DB2TO004' WAS IMBEDDED AT LINE 50 OF 'DB2TOMST'
DSMKPO653E POSTSCRIPT FILE 'db2g2w0' NOT FOUND.
DSMMOM395I '.EDFPO' LINE 80: .po db2g2w0
DSMMOM397I '.EDFPO' WAS IMBEDDED AT LINE 910 OF '.EDFAWRK'
DSMMOM397I '.EDFAWRK' WAS IMBEDDED AT LINE 103 OF 'DB2TO004'
DSMMOM397I 'DB2TO004' WAS IMBEDDED AT LINE 50 OF 'DB2TOMST'
DSMKPO653E POSTSCRIPT FILE 'maa7801' NOT FOUND.
DSMMOM395I '.EDFPO' LINE 80: .po maa7801
DSMMOM397I '.EDFPO' WAS IMBEDDED AT LINE 910 OF '.EDFAWRK'
DSMMOM397I '.EDFAWRK' WAS IMBEDDED AT LINE 789 OF 'DB2TO007'
DSMMOM397I 'DB2TO007' WAS IMBEDDED AT LINE 53 OF 'DB2TOMST'
DSMKPO653E POSTSCRIPT FILE 'io1g25' NOT FOUND.
DSMMOM395I '.EDFPO' LINE 80: .po io1g25
DSMMOM397I '.EDFPO' WAS IMBEDDED AT LINE 910 OF '.EDFAWRK'
DSMMOM397I '.EDFAWRK' WAS IMBEDDED AT LINE 703 OF 'DB2TO010'
DSMMOM397I 'DB2TO010' WAS IMBEDDED AT LINE 56 OF 'DB2TOMST'
DSMBEG323I STARTING PASS 2 OF 4.
DSMKPO653E POSTSCRIPT FILE 'BPI201A' NOT FOUND.
DSMMOM395I '.EDFPO' LINE 80: .po BPI201A
DSMMOM397I '.EDFPO' WAS IMBEDDED AT LINE 910 OF '.EDFAWRK'
DSMMOM397I '.EDFAWRK' WAS IMBEDDED AT LINE 1420 OF 'DB2TO003'
DSMMOM397I 'DB2TO003' WAS IMBEDDED AT LINE 49 OF 'DB2TOMST'
DSMKPO653E POSTSCRIPT FILE 'db2td030' NOT FOUND.
DSMMOM395I '.EDFPO' LINE 80: .po db2td030
DSMMOM397I '.EDFPO' WAS IMBEDDED AT LINE 910 OF '.EDFAWRK'
DSMMOM397I '.EDFAWRK' WAS IMBEDDED AT LINE 1449 OF 'DB2TO003'
DSMMOM397I 'DB2TO003' WAS IMBEDDED AT LINE 49 OF 'DB2TOMST'
DSMKPO653E POSTSCRIPT FILE 'db2td031' NOT FOUND.
DSMMOM395I '.EDFPO' LINE 80: .po db2td031
DSMMOM397I '.EDFPO' WAS IMBEDDED AT LINE 910 OF '.EDFAWRK'
DSMMOM397I '.EDFAWRK' WAS IMBEDDED AT LINE 1454 OF 'DB2TO003'
DSMMOM397I 'DB2TO003' WAS IMBEDDED AT LINE 49 OF 'DB2TOMST'
DSMKPO653E POSTSCRIPT FILE 'db2td02' NOT FOUND.
DSMMOM395I '.EDFPO' LINE 80: .po db2td02
DSMMOM397I '.EDFPO' WAS IMBEDDED AT LINE 910 OF '.EDFAWRK'
DSMMOM397I '.EDFAWRK' WAS IMBEDDED AT LINE 50 OF 'DB2TO004'
DSMMOM397I 'DB2TO004' WAS IMBEDDED AT LINE 50 OF 'DB2TOMST'
DSMKPO653E POSTSCRIPT FILE 'db2g2w0' NOT FOUND.
DSMMOM395I '.EDFPO' LINE 80: .po db2g2w0
DSMMOM397I '.EDFPO' WAS IMBEDDED AT LINE 910 OF '.EDFAWRK'
DSMMOM397I '.EDFAWRK' WAS IMBEDDED AT LINE 103 OF 'DB2TO004'
DSMMOM397I 'DB2TO004' WAS IMBEDDED AT LINE 50 OF 'DB2TOMST'
DSMKPO653E POSTSCRIPT FILE 'maa7801' NOT FOUND.
DSMMOM395I '.EDFPO' LINE 80: .po maa7801
DSMMOM397I '.EDFPO' WAS IMBEDDED AT LINE 910 OF '.EDFAWRK'
DSMMOM397I '.EDFAWRK' WAS IMBEDDED AT LINE 789 OF 'DB2TO007'
```

DSMMOM397I 'DB2TO007' WAS IMBEDDED AT LINE 53 OF 'DB2TOMST'
DSMKPO653E POSTSCRIPT FILE 'io1g25' NOT FOUND.
DSMMOM395I '.EDFPO' LINE 80: .po io1g25
DSMMOM397I '.EDFPO' WAS IMBEDDED AT LINE 910 OF '.EDFAWRK'
DSMMOM397I '.EDFAWRK' WAS IMBEDDED AT LINE 703 OF 'DB2TO010'
DSMMOM397I 'DB2TO010' WAS IMBEDDED AT LINE 56 OF 'DB2TOMST'
DSMBEG323I STARTING PASS 3 OF 4.
DSMKPO653E POSTSCRIPT FILE 'BPI201A' NOT FOUND.
DSMMOM395I '.EDFPO' LINE 80: .po BPI201A
DSMMOM397I '.EDFPO' WAS IMBEDDED AT LINE 910 OF '.EDFAWRK'
DSMMOM397I '.EDFAWRK' WAS IMBEDDED AT LINE 1420 OF 'DB2TO003'
DSMMOM397I 'DB2TO003' WAS IMBEDDED AT LINE 49 OF 'DB2TOMST'
DSMKPO653E POSTSCRIPT FILE 'db2td030' NOT FOUND.
DSMMOM395I '.EDFPO' LINE 80: .po db2td030
DSMMOM397I '.EDFPO' WAS IMBEDDED AT LINE 910 OF '.EDFAWRK'
DSMMOM397I '.EDFAWRK' WAS IMBEDDED AT LINE 1449 OF 'DB2TO003'
DSMMOM397I 'DB2TO003' WAS IMBEDDED AT LINE 49 OF 'DB2TOMST'
DSMKPO653E POSTSCRIPT FILE 'db2td031' NOT FOUND.
DSMMOM395I '.EDFPO' LINE 80: .po db2td031
DSMMOM397I '.EDFPO' WAS IMBEDDED AT LINE 910 OF '.EDFAWRK'
DSMMOM397I '.EDFAWRK' WAS IMBEDDED AT LINE 1454 OF 'DB2TO003'
DSMMOM397I 'DB2TO003' WAS IMBEDDED AT LINE 49 OF 'DB2TOMST'
DSMKPO653E POSTSCRIPT FILE 'db2td02' NOT FOUND.
DSMMOM395I '.EDFPO' LINE 80: .po db2td02
DSMMOM397I '.EDFPO' WAS IMBEDDED AT LINE 910 OF '.EDFAWRK'
DSMMOM397I '.EDFAWRK' WAS IMBEDDED AT LINE 50 OF 'DB2TO004'
DSMMOM397I 'DB2TO004' WAS IMBEDDED AT LINE 50 OF 'DB2TOMST'
DSMKPO653E POSTSCRIPT FILE 'db2g2w0' NOT FOUND.
DSMMOM395I '.EDFPO' LINE 80: .po db2g2w0
DSMMOM397I '.EDFPO' WAS IMBEDDED AT LINE 910 OF '.EDFAWRK'
DSMMOM397I '.EDFAWRK' WAS IMBEDDED AT LINE 103 OF 'DB2TO004'
DSMMOM397I 'DB2TO004' WAS IMBEDDED AT LINE 50 OF 'DB2TOMST'
DSMKPO653E POSTSCRIPT FILE 'maa7801' NOT FOUND.
DSMMOM395I '.EDFPO' LINE 80: .po maa7801
DSMMOM397I '.EDFPO' WAS IMBEDDED AT LINE 910 OF '.EDFAWRK'
DSMMOM397I '.EDFAWRK' WAS IMBEDDED AT LINE 789 OF 'DB2TO007'
DSMMOM397I 'DB2TO007' WAS IMBEDDED AT LINE 53 OF 'DB2TOMST'
DSMKPO653E POSTSCRIPT FILE 'io1g25' NOT FOUND.
DSMMOM395I '.EDFPO' LINE 80: .po io1g25
DSMMOM397I '.EDFPO' WAS IMBEDDED AT LINE 910 OF '.EDFAWRK'
DSMMOM397I '.EDFAWRK' WAS IMBEDDED AT LINE 703 OF 'DB2TO010'
DSMMOM397I 'DB2TO010' WAS IMBEDDED AT LINE 56 OF 'DB2TOMST'
DSMBEG323I STARTING PASS 4 OF 4.
DSMKPO653E POSTSCRIPT FILE 'BPI201A' NOT FOUND.
DSMMOM395I '.EDFPO' LINE 80: .po BPI201A
DSMMOM397I '.EDFPO' WAS IMBEDDED AT LINE 910 OF '.EDFAWRK'
DSMMOM397I '.EDFAWRK' WAS IMBEDDED AT LINE 1420 OF 'DB2TO003'
DSMMOM397I 'DB2TO003' WAS IMBEDDED AT LINE 49 OF 'DB2TOMST'
DSMKPO653E POSTSCRIPT FILE 'db2td030' NOT FOUND.
DSMMOM395I '.EDFPO' LINE 80: .po db2td030
DSMMOM397I '.EDFPO' WAS IMBEDDED AT LINE 910 OF '.EDFAWRK'
DSMMOM397I '.EDFAWRK' WAS IMBEDDED AT LINE 1449 OF 'DB2TO003'
DSMMOM397I 'DB2TO003' WAS IMBEDDED AT LINE 49 OF 'DB2TOMST'
DSMKPO653E POSTSCRIPT FILE 'db2td031' NOT FOUND.
DSMMOM395I '.EDFPO' LINE 80: .po db2td031
DSMMOM397I '.EDFPO' WAS IMBEDDED AT LINE 910 OF '.EDFAWRK'
DSMMOM397I '.EDFAWRK' WAS IMBEDDED AT LINE 1454 OF 'DB2TO003'
DSMMOM397I 'DB2TO003' WAS IMBEDDED AT LINE 49 OF 'DB2TOMST'
DSMKPO653E POSTSCRIPT FILE 'db2td02' NOT FOUND.
DSMMOM395I '.EDFPO' LINE 80: .po db2td02
DSMMOM397I '.EDFPO' WAS IMBEDDED AT LINE 910 OF '.EDFAWRK'
DSMMOM397I '.EDFAWRK' WAS IMBEDDED AT LINE 50 OF 'DB2TO004'
DSMMOM397I 'DB2TO004' WAS IMBEDDED AT LINE 50 OF 'DB2TOMST'
DSMKPO653E POSTSCRIPT FILE 'db2g2w0' NOT FOUND.
DSMMOM395I '.EDFPO' LINE 80: .po db2g2w0

DSMMOM397I '.EDFPO' WAS IMBEDDED AT LINE 910 OF '.EDFAWRK'
DSMMOM397I '.EDFAWRK' WAS IMBEDDED AT LINE 103 OF 'DB2TO004'
DSMMOM397I 'DB2TO004' WAS IMBEDDED AT LINE 50 OF 'DB2TOMST'
DSMKPO653E POSTSCRIPT FILE 'maa7801' NOT FOUND.
DSMMOM395I '.EDFPO' LINE 80: .po maa7801
DSMMOM397I '.EDFPO' WAS IMBEDDED AT LINE 910 OF '.EDFAWRK'
DSMMOM397I '.EDFAWRK' WAS IMBEDDED AT LINE 789 OF 'DB2TO007'
DSMMOM397I 'DB2TO007' WAS IMBEDDED AT LINE 53 OF 'DB2TOMST'
DSMKPO653E POSTSCRIPT FILE 'io1g25' NOT FOUND.
DSMMOM395I '.EDFPO' LINE 80: .po io1g25
DSMMOM397I '.EDFPO' WAS IMBEDDED AT LINE 910 OF '.EDFAWRK'
DSMMOM397I '.EDFAWRK' WAS IMBEDDED AT LINE 703 OF 'DB2TO010'
DSMMOM397I 'DB2TO010' WAS IMBEDDED AT LINE 56 OF 'DB2TOMST'