GA22-7139-02

Enterprise System/9000
Models 520, 640, 660, 740, 820, 860, and 900

Functional Characteristics and Configuration Guide

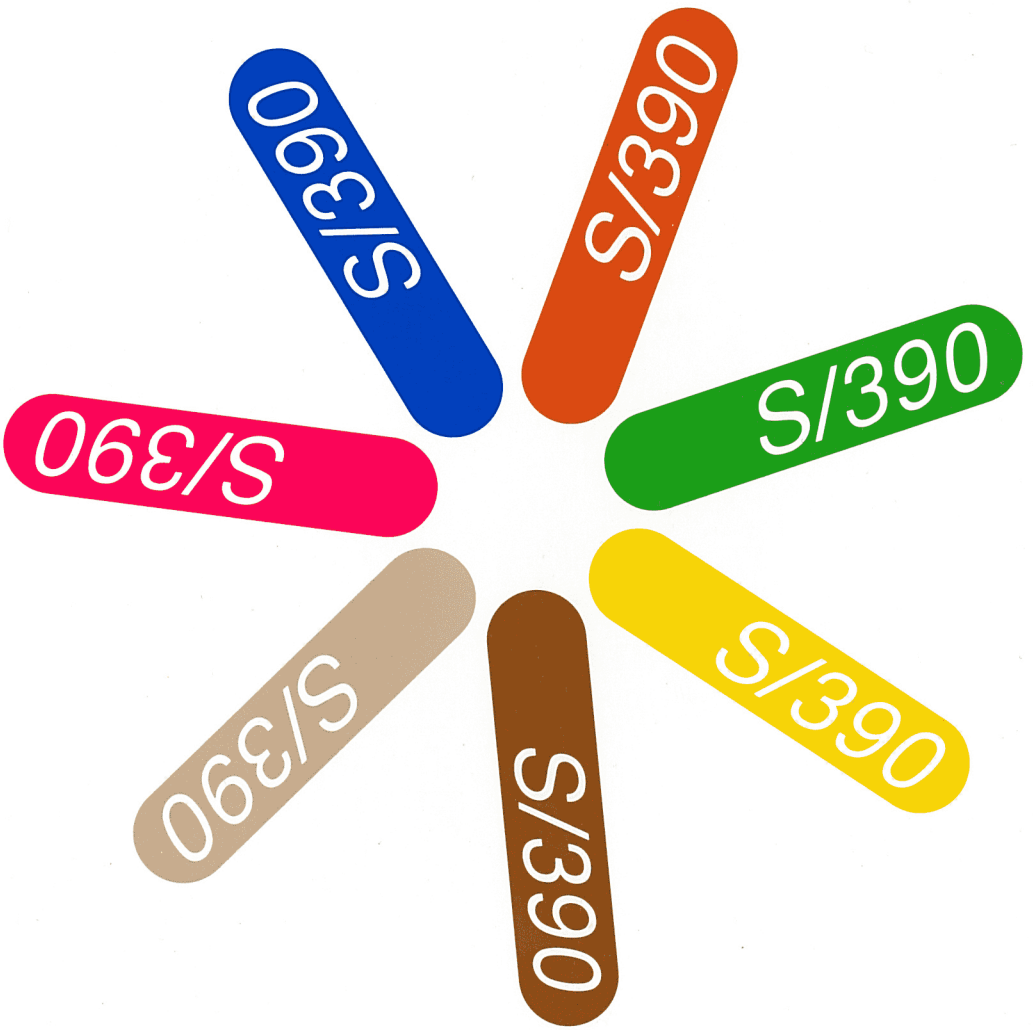

GA22-7139-02

Enterprise System/9000
Models 520, 640, 660, 740, 820, 860, and 900

Functional Characteristics and Configuration Guide

For questions or problems regarding this publication:

In the **U.S.A., Canada, and Puerto Rico,** call

1-800-426-9021

In **all countries,** either mail the Readers' Comment Form or FAX the Readers' Comment Form to

1-914-432-9405 **(FAX number)**

Note: In countries other than the U.S.A., Canada, and Puerto Rico, dial the appropriate international access code before the above FAX number.

Note

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

Third Edition (September 1991)

This edition obsoletes and replaces GA22-7139-1. Changes or additons 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, Department B97
P.O. Box 950
Poughkeepsie, NY 12602
U.S.A.

Contents

Notices xi
Trademarks and Service Marks xi

About This Publication xiii
What is Included in This Publication xiii
Related Publications xiv

Summary of Changes xv

Chapter 1. Introduction 1-1
Overview 1-2
Upgrade Progression 1-2
Hardware Configuration 1-3
Model 520 1-3
Model 640 1-3
Model 660 1-4
Model 740 1-4
Model 820 1-5
Model 860 1-5
Model 900 1-6
Design Highlights 1-6
ESA/390 Architectural Mode of Operation 1-7
Logically Partitioned (LPAR) Operating Mode 1-7
ESCON Architecture and Technology 1-8
Expanded Storage (Optional) 1-10
Vector Facility Feature (Optional) 1-10
Integrated Cryptographic Feature (Optional) 1-10
IBM 9037 Sysplex Timer Attachment (Optional) 1-11
High Levels of Reliability, Availability, and Serviceability 1-11
Softcopy Publications 1-11
Programming Compatibility 1-11
Programming Support 1-12
Control Programs for Basic Modes 1-12
Control Programs for Guest 1-13
Control Programs for Logically Partitioned Operation 1-13

Chapter 2. Characteristics of the ES/9000 2-1
Overview 2-2
Machine Requirements 2-2
Standard Requirements 2-3
Corequisite Requirements 2-3
Optional Machines and Features 2-4
Console and Display Configuration 2-4
Power and Cooling Functions 2-4
Power 2-5
Cooling 2-5
Input/Output Operations 2-5
Storage Operations 2-6
Data Representation 2-7
System Security 2-7
Data Integrity 2-7

System Access Control 2-7
IBM 9022 Processor Controller 2-8
Reliability, Availability, and Serviceability 2-8
Reliability 2-9
Availability 2-9
Serviceability 2-10

Chapter 3. Processor Components 3-1
Overview 3-2
Central Processors 3-3
Control Storage Element 3-3
Buffer Control Element 3-4
Instruction Element 3-4
Execution Element 3-5
CPU ID 3-5
Vector Facility Feature 3-6
Integrated Cryptographic Feature 3-7
Storage 3-8
Common Storage Buffer Element 3-9
Central Storage Element 3-9
Expanded Storage 3-11
Channel Subsystem 3-11
Interconnect Communication Element 3-12
Channels 3-12

Chapter 4. Consoles and Displays 4-1
Overview 4-2
System Console Interactive Facilities 4-3
System Definition 4-3
LPAR Mode Definition 4-3
System Activity Display 4-4
I/O Configuration Data Set Content 4-4
I/O Problem Determination Information 4-4

Chapter 5. Understanding the 9022 Processor Controller 5-1
Overview 5-2
Corequisites 5-2
System Power Panel 5-3
Dual Support Processors 5-3
Operation Monitoring and Control 5-4
Error Recovery 5-4
Configuration 5-4
Remote Support Facility 5-4
RSF Configuration 5-5
Authorization 5-5
Call Details 5-5

Chapter 6. Handling Errors 6-1
Overview 6-2
Automatic Error Recovery 6-2
Error Checking and Correction 6-2
Machine-Check Handling 6-2
I/O Operations 6-4
Problem Analysis 6-5
Power Malfunction 6-5

Other Malfunctions 6-5
Remote Support Facility 6-5

Chapter 7. Standard and Optional Features 7-1
Overview 7-4
Features Dependent on Architectural Mode 7-5
Features Not Dependent on Architectural Mode 7-8
Programming Assists Dependent on Architectural Mode 7-9
Describing the Features 7-10
Access-List-Controlled Protection 7-10
Advanced Address Space 7-10
Basic Control Mode 7-11
Bimodal Addressing 7-11
Branch and Save 7-11
Broadcasted Purging 7-11
Byte-Oriented Operand 7-11
Cancel I/O 7-11
Channels 7-12
Channel Indirect Data Addressing 7-12
Channel-Set Switching 7-12
Channel Subsystem 7-12
Channel-Subsystem Call 7-12
Clear I/O 7-12
Command Retry 7-12
Concurrent Sense 7-13
Conditional Swapping 7-13
Concurrent CP TCM Maintenance 7-13
Concurrent Channel Maintenance 7-13
Concurrent Power Subsystem 7-13
Console Integration 7-13
Control-Switch Assist 7-13
CPU Retry 7-14
CPU Timer and Clock Comparator 7-14
Data Streaming 7-14
DB2 Sort Facility 7-14
Dynamic Reconfiguration Management 7-15
Enhanced Move Page for VM 7-15
Enhanced Power Subsystem 7-15
Error Checking and Correction 7-15
ESCON Channels 7-16
ESCON Channel-to-Channel Adapter 7-16
Expanded Storage 7-16
Extended-Precision Divide 7-16
Extended-Precision Floating Point 7-16
Extended Real Addressing 7-16
Extended Sorting 7-17
Extensions for Virtual Machines 7-17
Fast Release 7-17
Fault-Tolerant Dynamic Memory Arrays 7-17
Floating Point 7-17
Halt Device 7-17
High-Speed Buffer Storage 7-17
Incorrect-Length Indication-Suppression Facility 7-18
Integrated Cryptographic Feature 7-18
Interpretive Execution 7-18

Interval Timer 7-18
I/O-Device Self-Description 7-19
I/O Error Alert 7-19
I/O Interface Reset 7-19
I/O Power Sequence Control 7-19
Key-Controlled Storage Protection 7-19
Limited Channel Logout 7-20
Logging Volume Reduction 7-20
Monitoring 7-20
Move Page 7-20
Multiprocessing 7-20
MVS Dynamic Physical Partitioning 7-21
Page Protection 7-21
PER Extensions 7-21
Preferred Machine Assist 7-21
Private Space 7-21
Processor Availability Facility 7-21
Processor Resource/Systems Manager 7-22
PSW-Key Handling 7-24
Recovery Extensions 7-24
SCP-Initiated Reset 7-24
Service Signal 7-24
Sorting Instructions 7-24
Storage Background Scrubbing 7-24
Storage-Key Instruction Extensions 7-25
Storage-Key Instructions 7-25
Storage-Key 4KB Block 7-25
Storage-Reconfiguration 7-25
Subsystem Storage Protection 7-25
IBM 9037 Sysplex Timer 7-25
System/370 Extended Facility 7-26
Test Block 7-26
Time-of-Day Clock 7-26
Tracing 7-27
Translation 7-27
Vector Facility Feature 7-28
Virtual-Machine Assist 7-28
VM Assists for the CPU Timer 7-28
3033 Extension 7-29
31-Bit Indirect Data Address Word 7-29
31-Bit Real Addressing 7-29

Chapter 8. Characteristics of Channel Operation 8-1
Overview 8-2
Subchannels 8-2
Channel Control 8-3
Channel Commands 8-3
ESA/390 Mode 8-3
Chaining 8-5
I/O Interruptions 8-5
Channel Implementation 8-6
Channel Time-Out Functions 8-6
Device Priority on an I/O Interface 8-7
Dynamic Reconnection 8-8

Chapter 9. Channel Subsystem Performance 9-1
Overview 9-2
Performance Characteristics 9-2
Description 9-2
I/O Device Characteristics 9-5
ESCON Channels 9-7
Configuring ESCON Channels 9-7
Elapsed Time 9-7
Critical Time 9-7
Deferred Access 9-8
Parallel Channels Only 9-9
Device Class 9-9
Block Multiplexer Mode of Operation 9-9
ESCON Channels Using Link-Level and Device-Level Protocols 9-9
ESCON Channels Attached to a 9034 and Parallel Channels 9-10
Byte Multiplexer Mode of Operation (Parallel Channels Only) 9-10
Byte Multiplexer Mode and Burst Mode 9-10

Chapter 10. Guidelines for Channel Subsystem Configuration 10-1
Overview 10-2
ESCON Channel Subsystem Configurations 10-2
Recommendations for ESCON Channel Subsystem Configurations . . . 10-2
Guidelines for Parallel Channel Subsystem Configurations 10-3
Device Sequence on Block Multiplexer Channels 10-3
Device Intermix 10-3
Data Chaining 10-4
CCW Prefetching for ESCON Channels Attached to a 9034 10-5
Byte Multiplexer Channel Analysis 10-5
Balancing Device Load across Channels 10-5
Sequencing Devices Within a Channel 10-6
General Device Considerations 10-6
All ESCON and Parallel Channels 10-6
ESCON Channels Attached to a 9034 and Parallel Channels 10-7

Appendix A. Architectural Deviations A-1
Concurrent Indication of PER Events with Operand-Access Exception A-1
Protection Violation Instead of Delayed Access Exception A-1

Appendix B. Summary of the Configurations B-1

Glossary of Terms and Abbreviations X-1

Index X-19

Figures

1-1. IBM ES/9000 Model 900 1-2
2-1. Unit Identification for the ES/9000 Model 900 (not to scale) 2-3
3-1. Components for Model 900 3-2
3-2. Elements of a Central Processor 3-3
3-3. CPU ID - Doubleword Content in Basic Mode 3-5
3-4. CPU ID - Doubleword Content in LPAR Mode 3-5
3-5. Channel Subsystem Diagram 3-14
4-1. IBM 3206 Display Station Model 100 4-3
5-1. IBM 9022 Processor Controller Model 1A 5-2
5-2. 9022 Model 1A System Power Panel 5-3
8-1. Device Priority on Parallel Channels 8-7

Tables

1-1. Upgrade Progression 1-2
1-2. Connectivity Distances Using ESCON XDF 1-9
3-1. HSA Sizing Table in LPAR Mode 3-11
6-1. Machine-Check Interruption Code (MCIC) 6-4
7-1. Features Dependent on Architectural Mode 7-5
7-2. Features Not Dependent on Architectural Mode 7-8
7-3. Programming Assists Dependent on Architectural Mode 7-9
9-1. Channel Data Rate as a Function of Distance for Models 520 — 900 9-5
B-1. Summary of the Configurations B-1

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's product, program, or service may be used. Any functionally equivalent program that does not infringe any of IBM's 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 the IBM Director of Commercial Relations, IBM Corporation, Purchase, NY 10577.

Trademarks and Service Marks

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

AIX
AT
BookManager
CICS/ESA
DB2
ECKD
Enterprise System/3090
Enterprise System/9000
Enterprise Systems Architecture/370
Enterprise Systems Architecture/390
Enterprise Systems Connection Architecture
ESCON
ESCON XDF
ES/3090
ES/9000
ESA/370
ESA/390
Hardware Configuration Definition
Hiperspace
IBM
ICRF
IMS/ESA
MVS/DFP
MVS/ESA
MVS/SP
Personal Security
Personal System/2
PR/SM
Processor Resource/Systems Manager
PS/2
S/390
System/370
System/390
VM/ESA
VM/XA
VSE/ESA
3090

About This Publication

This publication describes the IBM* Enterprise System/9000* Models 520, 640, 660, 820, 860 and 900, their design, components, functions, features and their capability. It also provides basic guidelines and recommendations for configuring channel loads.

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, "Characteristics of the ES/9000" describes the characteristics of each model.
- Chapter 3, "Processor Components" describes the logical components of the processor unit.
- Chapter 4, "Consoles and Displays" describes the interactive consoles and displays, and the facilities provided by the system console.
- Chapter 5, "Understanding the 9022 Processor Controller" describes the 9022 processor controller.
- Chapter 6, "Handling Errors" describes error recovery procedures the processor controller performes automatically, customer problem analysis procedures, and remote support facility procedures.
- Chapter 7, "Standard and Optional Features" describes the standard and optional features.
- Chapter 8, "Characteristics of Channel Operation" describes the characteristics, function, and structure of the channels.
- Chapter 9, "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 10, "Guidelines for Channel Subsystem Configuration" provides recommendations for ESCON channel configurations and describes how to configure I/O devices for operation on block multiplexer or byte multiplexer channels.
- Appendix A, "Architectural Deviations" outlines the exceptions to the *IBM Enterprise Systems Architecture/390 Principles of Operation*.

- Appendix B, "Summary of the Configurations" contains a table that shows a summary outline of the configurations.

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.
- *ES/9000 Processor Complex Models 330, 340, 500, 580, 620, and 720: Functional Characteristics and Configuration Guide*, GA22-7138.
- *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 Device Summary*, GA32-0039.
- *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 Models 330—900 Installation Manual—Physical Planning*, GC22-7083.
- *ES/9000 Models 520, 640, 660, 740, 820, 860, and 900 Operator Guide*, GC38-0086.
- *ES/9000 Models 520, 640, 660, 740, 820, 860, and 900 Recovery Guide*, GC38-0090.
- *ES/9000 Models 520, 640, 660, 740, 820, 860, and 900 Operator Messages*, GC38-0094.
- *ES/9000 ES/3090 Input/Output Configuration Program User Guide*, GC38-0097.
- *ES/9000 Models 330—900 Customer Library CD-ROM*, GK2T-6685.
- *Introducing Enterprise Systems Connection*, GA23-0383.
- *Introduction to Nonsynchronous Direct Access Storage Subsystems*, GC26-4519.
- *Enterprise Systems Connection: Planning for Migration*, GC66-3181.

Summary of Changes

This edition contains editorial updates and the following product updates:

- Three new processors Models 520, 640, 660, and 740 to complement Models 820, 860, and 900
- Enhanced system reliability, availability, and serviceability
  - Subsystem storage protection
  - Storage background scrubbing
  - Concurrent channel maintenance
  - Concurrent CP TCM maintenance
  - Concurrent power maintenance
  - Concurrent sense
  - Fault tolerant dynamic memory array
- Enhanced system performance functions
  - Enhanced move page for VM
  - Extended sorting
  - Nine new vector instructions
  - Two new scalar square root instructions
  - Access-list-controlled protection
  - Cancel I/O
  - Channel subsystems
  - Console integration
  - Extended sorting
  - I/O interface reset
  - MVS dynamic physical partitioning
- Enhanced IBM Sysplex* Timer functions
  - External time source
  - Automatic propagation delay adjustment
- New ESCON and channel functions
  - ESCON Extended Distance Feature (ESCON XDF*)
  - ESCON data rate increase to 17 MB/sec
- New Processor Resource/Systems Manager* (PR/SM*) functions
  - LPAR high-performance parallel interface (HIPPI)
  - LPAR support of ICRF
  - LPAR preferred path
  - LPAR auto-reconfiguration
  - LPAR management time report
- New Integrated Cryptographic Feature (ICRF) function
  - ICRF Personal Security* card

Chapter 1. Introduction

Overview 1-2
Upgrade Progression 1-2
Hardware Configuration 1-3
Model 520 1-3
Model 640 1-3
Model 660 1-4
Model 740 1-4
Model 820 1-5
Model 860 1-5
Model 900 1-6
Design Highlights 1-6
ESA/390 Architectural Mode of Operation 1-7
Logically Partitioned (LPAR) Operating Mode 1-7
ESCON Architecture and Technology 1-8
ESCON Architecture 1-8
ESCON Channels 1-9
Fiber Optic Cabling, Transmission, and Reception 1-9
Dynamic Connectivity (Switched Point-to-Point Topology) 1-10
Interconnectivity with Other Networks 1-10
Expanded Storage (Optional) 1-10
Vector Facility Feature (Optional) 1-10
Integrated Cryptographic Feature (Optional) 1-10
IBM 9037 Sysplex Timer Attachment (Optional) 1-11
High Levels of Reliability, Availability, and Serviceability 1-11
Softcopy Publications 1-11
Programming Compatibility 1-11
Programming Support 1-12
Control Programs for Basic Modes 1-12
Control Programs for Guest 1-13
Control Programs for Logically Partitioned Operation 1-13

Overview

The IBM ES/9000 9021 family of processors Models 520, 640, 660, 740, 820, 860, and 900 are general-purpose data processing systems that operate in ESA/390 mode (see "ESA/390 Architectural Mode of Operation" on page 1-7). The design of these models combined with their new functions contributes to improved reliability, availability, serviceability and performance. Figure 1-1 shows a standard Model 900.

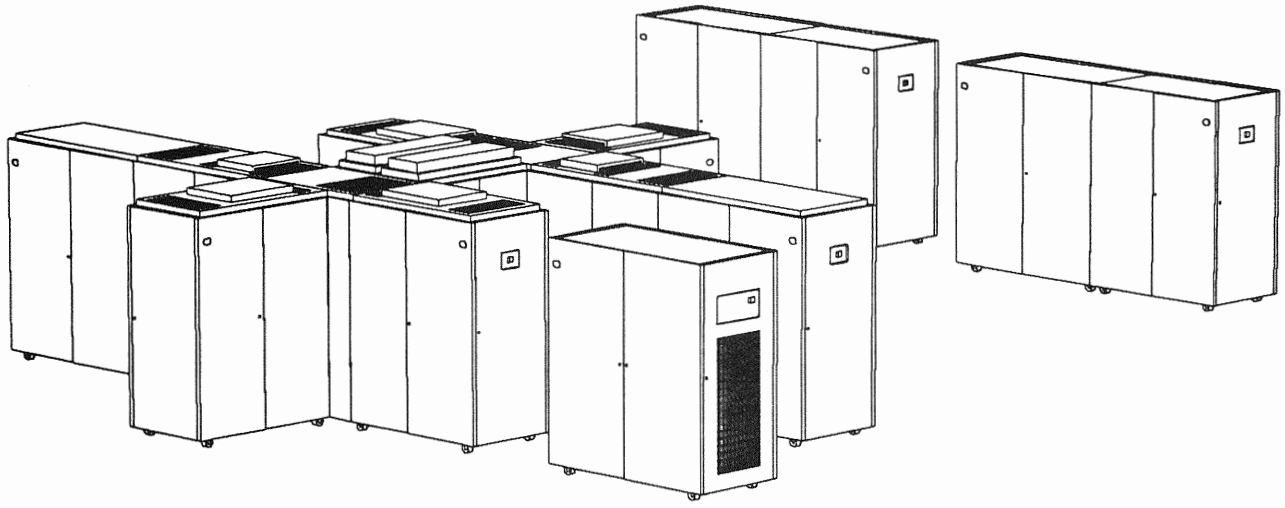

Figure 1-1. IBM ES/9000 Model 900

Upgrade Progression

Models 520, 640, 660, 740, 820, 860, and 900 provide compatible growth from Models 330, 340, 500, 580, 620, and 720. Models 330, 340, 500, 580, 620, and 720 provide compatible growth from IBM Enterprise System/3090* processor complex. For a list of the upgrade progression from Models 330, 340, 500, 580, 620, and 720 to Models 520, 640, 660, 740, 820, 860, and 900 refer to Table 1-1. For a list of the upgrade progression to Models 330, 340, 500, 580, 620, and 720 refer to *ES/9000 Models 330, 340, 500, 580, 620, and 720 Functional Characteristics and Configuration Guide*.

Table 1-1. Upgrade Progression

Model	Upgrade Progression
520	640 or 660
620	720 or 820 (see note)
640	740 or 820
660	820
720	820, 860, or 900 (see note)
740	820, 860, or 900
820	860 or 900
860	900

Note: Upgrades from 620 to 820 and from 720 to 820, 860, or 900 involve substantial changes to the existing processor.

Hardware Configuration

The following describes the hardware configuration for Models 520, 640, 660, 740, 820, 860, and 900.

Model 520

Model 520 is a uniprocessor containing one integrated central processor (CP). The CP has access to a common central storage, expanded storage, and channels. Model 520 provides:

- One integrated CP
- 128MB (MB equals 1 048 576 bytes) minimum of central storage, with a maximum of 256MB
- Channels: 32 minimum, 64 maximum in increments of 16. They can be installed in any valid combination of ESCON and parallel channels: 0, 16, 32, 48, or 64 ESCON channels or 0, 16, 32, or 48 can be parallel channels.
- Ability to operate with as many as seven logical partitions
- An optional shared expanded storage: 256MB, 512MB, 1GB, 1536MB, or 2GB (GB equals 1 073 741 824 bytes)
- An optional Integrated Cryptographic Feature (ICRF) that is mutually exclusive with the vector facility on the CP
- An optional vector facility feature on each CP that is mutually exclusive with ICRF.

Model 640

Model 640 is a two-way processor containing two integrated central processors (CPs). Each CP has access to a common central storage, expanded storage, and channels. Model 640 provides:

- Two integrated CPs
- 128MB (MB equals 1 048 576 bytes) minimum of central storage, with a maximum of 512MB, installed in increments of 128MB or 256MB
- Channels: 64 minimum, 128 maximum in increments of 16. They can be installed in any valid combination of ESCON and parallel channels: 16, 32, 48, 64, 80, 96, 112, or 128 ESCON channels or 0, 16, 32, or 48 can be parallel channels.
- Ability to operate with as many as seven logical partitions
- An optional shared expanded storage: 256MB, 512MB, 1GB, 1536MB, 2GB, 3GB, or 4GB (GB equals 1 073 741 824 bytes)
- An optional Integrated Cryptographic Feature (ICRF) that is mutually exclusive with the vector facility on a CP
- An optional vector facility feature on each CP that is mutually exclusive with ICRF.

Model 660

Model 660 is a two-way multiprocessor containing two integrated central processors (CPs). Each CP has access to a common central storage, expanded storage, and channels. Model 660 provides:

- Two integrated CPs
- 256MB (MB equals 1 048 576 bytes) minimum of central storage, with a maximum of 512MB, installed asymmetrically in increments of 128MB
- Channels: 64 minimum (32 per side), 128 maximum (64 per side) in increments of 16. They can be installed asymmetrically in any valid combination of ESCON and parallel channels: 0, 16, 32, 48 or 64 ESCON channels per side or 0, 16, 32, or 48 can be parallel channels per side.
- Ability to operate in a physically partitioned configuration as two uniprocessors
- Ability to operate with as many as seven logical partitions in a single-image configuration and with as many as 14 logical partitions in a physically partitioned configuration
- An optional shared expanded storage: 256MB, 512MB, 1GB, 1536MB, or 2GB (GB equals 1 073 741 824 bytes) per side
- An optional Integrated Cryptographic Feature (ICRF) for each side that is mutually exclusive with the vector facility on a CP
- An optional vector facility feature on each CP that is mutually exclusive with ICRF.

Model 740

Model 740 is a three-way (triadic) processor containing three integrated central processors (CPs). Each CP has access to a common central storage, expanded storage, and channels. Model 740 provides:

- Three integrated CPs
- 256MB (MB equals 1 048 576 bytes) minimum of central storage, with a maximum of 512MB, installed in increments of 256MB
- Channels: 64 minimum, 128 maximum in increments of 16. They can be installed in any valid combination of ESCON and parallel channels: 16, 32, 48, 64, 80, 96, 112, or 128 ESCON channels or 0, 16, 32, or 48 can be parallel channels.
- Ability to operate with as many as seven logical partitions
- An optional shared expanded storage: 256MB, 512MB, 1GB, 1536MB, 2GB, 3GB, or 4GB (GB equals 1 073 741 824 bytes)
- An optional Integrated Cryptographic Feature (ICRF) that is mutually exclusive with the vector facility on a CP
- An optional vector facility feature on each CP that is mutually exclusive with ICRF.

Model 820

Model 820 is a four-way multiprocessor containing four integrated central processors (CPs). Each CP has access to a common central storage, expanded storage, and channels for single-image configuration or for physically partitioned configuration. Model 820 provides:

- Four integrated CPs
- 128MB (MB equals 1 048 576 bytes), 256MB, or 512MB of shared central storage on each side, installed asymmetrically in increments of 128MB and 256MB
- Channels: 64 minimum, 128 maximum per side in increments of 16. They can be installed asymmetrically in any valid combination of parallel and ESCON channels (16 ESCON channels per side are standard); 0, 16, 32, or 48 can be parallel channels and 16, 32, 48, 64, 80, 96, 112, 128, can be ESCON channels.
- Ability to operate in a physically partitioned configuration as two dyadic processors
- Ability to operate with as many as seven logical partitions in a single-image configuration and with as many as 14 logical partitions in a physically partitioned configuration
- An optional shared expanded storage: 256MB, 512MB, 1GB, 1536MB, 2GB, 3GB, or 4GB (GB equals 1 073 741 824 bytes) on each side installed asymmetrically in increments of 256MB and 512MB.
- An optional Integrated Cryptographic Feature (ICRF) for each side that is mutually exclusive with the vector facility on a CP
- An optional vector facility feature for each CP that is mutually exclusive with ICRF.

Model 860

Model 860 is a five-way multiprocessor containing five integrated central processors (CPs). Each CP has access to a common central storage, expanded storage, and channels for single-image configuration or for physically partitioned configuration. Model 860 provides:

- Five integrated CPs
- 256MB (MB equals 1 048 576 bytes) or 512MB of shared central storage on A-side and 128MB, 256MB, or 512MB of shared central storage on B-side, installed asymmetrically in increments of 128MB and 256MB
- Channels: 64 minimum, 128 maximum per side in increments of 16. They can be installed asymmetrically in any valid combination of parallel and ESCON channels (16 ESCON channels per side are standard); 0, 16, 32, or 48 can be parallel channels and 16, 32, 48, 64, 80, 96, 112, 128, can be ESCON channels.
- Ability to operate in a physically partitioned configuration as one dyadic and one triadic processor
- Ability to operate with as many as seven logical partitions in a single-image configuration and with as many as 14 logical partitions in a physically partitioned configuration

- An optional shared expanded storage: 256MB, 512MB, 1GB, 1536MB, 2GB, 3GB, or 4GB (GB equals 1 073 741 824 bytes) on each side, installed asymmetrically in increments of 256MB, 512MB, and 1GB
- An optional Integrated Cryptographic Feature (ICRF) for each side that is mutually exclusive with the vector facility on a CP
- An optional vector facility feature for each CP that is mutually exclusive with ICRF.

Model 900

Model 900 is a six-way multiprocessor containing six integrated central processors (CPs). Each CP has access to a common central storage, expanded storage, and channels for single-image configuration or for physically partitioned configuration. Model 900 provides:

- Six integrated CPs
- 256MB (MB equals 1 048 576 bytes) or 512MB of shared central storage on each side, installed asymmetrically in increments of 256MB
- Channels: 64 minimum, 128 maximum per side in increments of 16. They can be installed asymmetrically in any valid combination of parallel and ESCON channels (16 ESCON channels per side are standard); 0, 16, 32, or 48 can be parallel channels and 16, 32, 48, 64, 80, 96, 112, 128, can be ESCON channels.
- Ability to operate in a physically partitioned configuration as two triadic processors
- Ability to operate with as many as seven logical partitions in a single-image configuration and with as many as 14 logical partitions in a physically partitioned configuration
- An optional shared expanded storage: 256MB, 512MB, 1GB, 1536MB, 2GB, 3GB, 4GB (GB equals 1 073 741 824 bytes) on each side, installed asymmetrically in increments of 256MB, 512MB and 1GB
- An optional Integrated Cryptographic Feature for each side that is mutually exclusive with the vector facility on a CP.
- An optional vector facility feature on each CP that is mutually exclusive with ICRF.

Design Highlights

Models 520, 640, 660, 740, 820, 860, and 900 provide high performance and flexibility due to its improved design and use of technology advances. The design of these models incorporates:

- ESA/390 architectural mode of operation
- Logically partitioned LPAR operating mode. LPAR mode provides both ESA/390 and System/370 logical partitions.
- IBM Enterprise Systems Connection Architecture (ESCON Architecture) and technology for the ESCON Channels
- Optional expanded storage
- Optional vector facility feature

- Optional Integrated Cryptographic Feature (ICRF)
- Optional IBM Sysplex Timer attachment
- High levels of reliability, availability, and serviceability (including concurrent CP TCM maintenance, concurrent power maintenance, and concurrent channel maintenance)
- Softcopy publications

ESA/390 Architectural Mode of Operation

Models 520, 640, 660, 740, 820, 860, and 900 provide the Enterprise Systems Architecture/390 (ESA/390) mode in three ways:

- Basic mode
- Mode of a logical partition in the 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 (2 147 483 648 bytes).
- ESCON and parallel 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.
- Support for the Start Interpretive Execution (SIE) instruction, allowing support of guest ESA/390 and System/370 virtual machines.
- Support for the VM data space facility which makes the ESA/390 access-register architecture more useful in virtual machine applications. For more information, see *VM/ESA CP Programming Services*, SC24-5520.

Logically Partitioned (LPAR) Operating Mode

You select the operating mode during system initialization (power-on reset). This occurs for a single-image configuration or for each side of a physically partitioned configuration. The operating mode can be basic mode (ESA/390) or LPAR mode. When you select LPAR mode, you must specify the:

- Architectural mode (ESA/390 or System/370) of each partition
- Resources of each partition

You can reconfigure most resources without having to reinitialize the system. Before you can activate 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 1MB contiguous blocks. These allocations are dynamically reallocated. After you define and activate a System/370 or ESA/390 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.

You can allocate individual channel paths to each logical partition. But a channel path can be allocated only to one logical partition at a time. Also, channel paths can be dynamically reconfigured between logical partitions. To share a device between logical partitions, you must use a separate channel path for each logical partition.

You can now 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 refer to the *ES/9000 ES/3090 Input/Output Configuration Program User Guide*, GC38-0097.

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 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.

If the optional vector facility feature or the optional ICRF feature is installed on a CP, it is available for use by all the partitions that share that CP. CPs that are dedicated to a logical partition (including associated vector or ICRF facilities) are available only to that logical partition.

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.*

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*.

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*.

Fiber Optic Cabling, Transmission, and Reception

ESCON LED multimode fiber optic technology providing 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 single-mode fiber optic technology providing a direct channel attachment range of up to 60 kilometers (37.3 miles); and control unit attachment range of up 43 kilometers (26.7 miles) from a channel through two optional ESCON directors. The following table shows the maximum cable distances that ESCON I/O devices can be located from the processor using ESCON Directors with ESCON XDF.

Table 1-2. Connectivity Distances Using ESCON XDF

Device	Maximum Link Distance	Minimum Number of ESCON Directors Required
ESCON Converter Models 1 and 2	3 km (1.9 miles)	1
9343 Model D04	9 km (5.6 miles)	1
3490 Models A01, A02, A10, A20,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
ES/9000 Channel-to-Channel (CTC)	60 km (37.3 miles)	2

Note: Special features must be installed on the IBM 3490 Magnetic Tape Subsystem when they are attached at a distance greater than 9 km from the host processor with ESCON XDF.

Combining fiber optic cable with ESCON Architecture, moves data faster than with copper wire cable (bus and tag cable for parallel channels). Using ESCON

channels, data can move at a rate of up to 17 million bytes per second (MB/s) compared to a maximum of 4.5 MB/s when using parallel channels.

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 reconfiguration management.
- Alternate paths can be used if a failure occurs.

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 processors, 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.

Expanded Storage (Optional)

Optional high-speed, high-capacity expanded storage is available as an integrated part of the ES/9000 Processor Complex. Expanded storage improves system response and system performance balancing. Models 820, 860, and 900 offer as many as 4GB (4 294 967 296 bytes) of optional expanded storage per side, for a total of 8GB.

See Appendix B, "Summary of the Configurations" for a list of the optional expanded storage for each side.

Vector Facility Feature (Optional)

The optional vector facility feature improves the performance of applications that use a great deal of arithmetic or logical instructions. You can replace the scalar instruction loops with the vector instructions and the application will take advantage of the order inherent in vector data; thereby improving performance.

Integrated Cryptographic Feature (Optional)

The optional Integrated Cryptographic Feature (ICRF) improves the security of data. ICRF adheres to international encryption standards and resists tampering. ICRF is available for a central processor on Models 520, 640, and 740 and for one central processor on each side on Models 660, 820, 860, and 900. ICRF is mutually exclusive with the vector facility on a CP and uses the IBM Personal Security* card to enter cryptographic key information.

For more information refer to "Integrated Cryptographic Feature" on page 3-7.

IBM 9037 Sysplex Timer Attachment (Optional)

The optional IBM 9037 Sysplex Timer, supported by MVS/ESA* SP Version 4.1 and subsequent releases, synchronizes the processor's time-of-day clock in multiple systems.

High Levels of Reliability, Availability, and Serviceability

Models 520, 640, 660, 740, 820, 860, and 900 offer high levels of reliability, availability, and serviceability (RAS) by reducing down time and using the following standard features:

- Online error detection and fault isolation techniques
- Concurrent CP TCM maintenance
- Concurrent channel maintenance
- Concurrent power maintenance
- Concurrent sense
- Concurrent power subsystem
- Enhanced power subsystem
- Dynamic modification of the I/O configuration
- Deferred maintenance capability
- A remote support strategy
- Subsystem storage protection
- Storage background scrubbing
- Fault-tolerant dynamic memory array
- Automatic error detection

For more information, see "Reliability, Availability, and Serviceability" on page 2-8.

Softcopy Publications

All ES/9000 customer publications are available on a CD-ROM disc. To view the CD-ROM disc, you need to use IBM BookManager* READ licensed programs. One CD-ROM disc is included with the hardcopy publications shipped with all ES/9000 models. You can order, from IBM, additional copies of the CD-ROM disc (GK2T-6685).

Programming Compatibility

This information applies to Models 520, 640, 660, 740, 820, 860, and 900 having a single-image configuration, independently to each side of an ES/9000 in the physically partitioned configuration, and to operation in a logical partition.

Any program written for ESA/390 mode can operate on Models 520, 640, 660, 740, 820, 860, and 900 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*.

In general, any program, including its programming support, that complies with the programming compatibility statements described above for Models 520, 640, 660, 740, 820, 860, and 900 will operate in a logical partition with the following exceptions:

- VM/XA* Migration Aid, VM/XA Systems Facility, and VM/XA System Product Release 1 are not supported in a logical partition.
- When operating in a logical partition, the CPU ID presented to that logical partition is not the same as when operating in basic mode. This may impact operation of software products that use the CPU ID field.

For more information, see the *IBM Processor Resource/Systems Manager Planning Guide*.

Programming Support

The control program you use depends on the mode in which the processor complex is operating. For example, an ES/9000 operating in ESA/390 mode (basic mode) requires a control program for ESA/390 mode.

The term *basic mode* refers to the mode the processor complex is operating in when it is not operating in LPAR mode. ESA/390 is an example of a basic mode and guest.

Control Programs for Basic Modes

The following IBM control programs operate in basic mode:

- MVS/ESA System Product Version 4 Release 1 and subsequent releases
- MVS/SP Version 3 Release 1 and subsequent releases
- MVS/SP Version 2 Release 2 and subsequent releases
- VM/ESA System Product Releases 1 (ESA Feature) and 1.1
- VM/XA System Product Release 2.1
- Transaction Processing Facility (TPF) Version 3 Release 1

Control Programs for Guest

VM/XA SP Release 2.1 and VM/ESA System Product Releases 1 and 1.1 provide Start Interpretive Execution (SIE) support for the following guest environments when operating in basic mode:

- MVS/SP Release 3.5
- MVS/SP Version 2 Release 1.3
- MVS/SP Version 3 (requires VM/XA System Product Release 2); Release 2.1 supports the use of the Move Page facility by an MVS/SP Version 3 Release 3.1 guest
- VM/SP Release 4
- VM/SP High Performance Option (HPO) Release 4.2
- OS/VS1 Release 7 with Basic Programming Extensions Release 4
- VSE/Advanced Function Version 2 and Version 4
- VSE/SP Version 2, Version 3, and Version 4
- VM/XA Systems Facility (V=V only)
- VM/XA System Product (V=V only, except that a V=R or V=F VM/XA SP 2.1 guest is supported with a VM/XA SP 2.1 host of Models 520, 640, 660, 740, 820, 860, and 900)
- Advanced Interactive Executive/370 (AIX*/370) (V=V and V=R only).

Control Programs for Logically Partitioned Operation

The following IBM control programs operate in LPAR mode:

- MVS/SP Version 2 Release 1.3 and subsequent releases
- MVS/SP Release 3.5
- VM/SP Release 5 and subsequent releases
- VM/XA SP Release 2.1
- VM/ESA 1 Release 1 (370 Feature)
- VM/SP HPO Release 5.0 and subsequent releases
- VM/SP Release 5 and subsequent releases
- VSE/ESA Release 1
- VSE/SP Version 4 with APAR DY39560
- Transaction Processing Facility (TPF) Version 3 Release 1
- MUSIC/SP Version 2 Release 2 or Version 2 Release 3
- AIX/ESA

The following ESA/390 control programs operate in LPAR mode:

- MVS/ESA Version 4 Release 1 and subsequent releases
- MVS/SP Version 3 Release 1 and subsequent releases
- VM/ESA System Product Release 1 (ESA Feature) and 1.1

For information about the characteristics of those control programs operating in LPAR mode, see the *ES/9000, ES/3090 Processor Resource/Systems Manager Planning Guide*, GA22-7123. Complex, copyrighted and licensed by IBM.

Chapter 2. Characteristics of the ES/9000

Overview 2-2
Machine Requirements 2-2
  Standard Requirements 2-3
  Corequisite Requirements 2-3
Optional Machines and Features 2-4
Console and Display Configuration 2-4
Power and Cooling Functions 2-4
  Power 2-5
  Cooling 2-5
Input/Output Operations 2-5
    Channel-to-Channel Connection 2-6
      ESCON Channels (Attached to a 9034) and Parallel Channels 2-6
      ESCON Channels (Using Link-Level and Device-Level Protocols) . . 2-6
Storage Operations 2-6
Data Representation 2-7
System Security 2-7
  Data Integrity 2-7
  System Access Control 2-7
IBM 9022 Processor Controller 2-8
Reliability, Availability, and Serviceability 2-8
  Reliability 2-9
  Availability 2-9
  Serviceability 2-10

Overview

This chapter describes the following characteristics for the Models 520, 640, 660, 740, 820, 860, and 900:

- Machine requirements
- Optional machines and features
- Console and display configuration
- Power and cooling
- Input/output operations
- Storage operations
- Data representation
- System security
- Technology
- Processor controller
- Reliability, availability and serviceability.

Machine Requirements

Machine requirements fall into two categories:

- Standard requirements
- Corequisite requirements for the operation and maintenance of processors complex ordered separately.

For a plan view of the ES/9000 Model 900, see Figure 2-1 on page 2-3

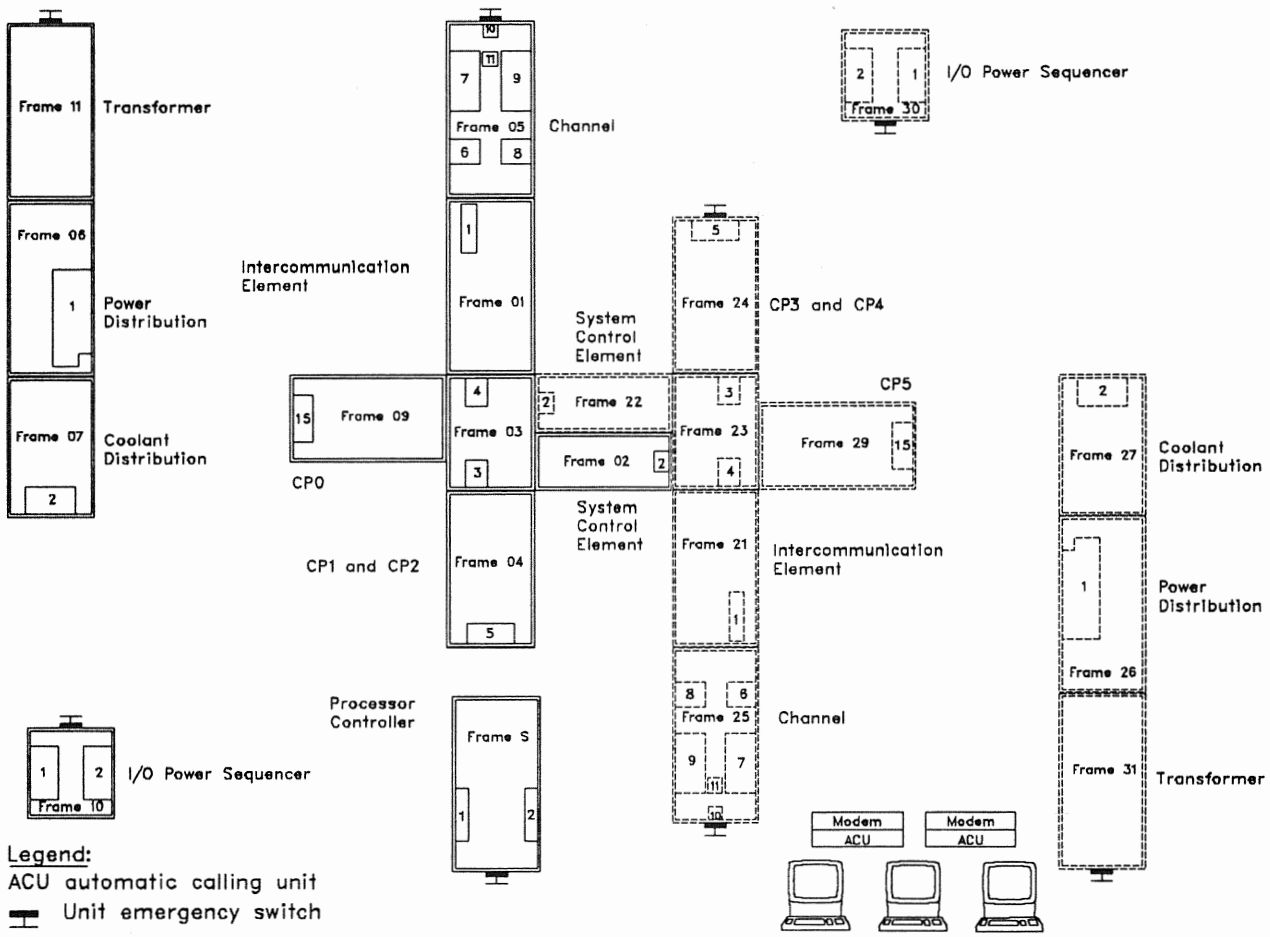

Figure 2-1. Unit Identification for the ES/9000 Model 900 (not to scale)

Standard Requirements

The standard machine requirements are:

- One IBM Model 520, 640, 740, 820, 860, or 900.
- One IBM 9022 Processor Controller Model 1A. Models 660, 820, 860, and 900 require two IBM 9022 Processor Controller Model 1A with specify code 5050.
- Two or three IBM 3206 Display Stations Model 100 (or equivalent). Three display stations are required for Models 660, 820, 860, and 900; two display stations are required for Models 520, 640, and 740. For more information, see Chapter 4, "Consoles and Displays."

Corequisite Requirements

The corequisite machine requirements are:

- One or two IBM 5853 Modems Model 1 (or equivalent). Models 660, 820, 860, and 900 require two modems of the same type; Models 520, 640, and 740 require one modem.
- Two or three channel-attached operator display stations for communication with the control program. Models 660, 820, 860, and 900 in physically partitioned configuration require three display stations; Models 520, 640, and 740 require two. For more information see Chapter 4, "Consoles and Displays."

Note: IBM 3180 Display Station Model 140 and the 3864 Modem Model 2 can be used with Models 520, 640, 660, 740, 820, 860, and 900.

Optional Machines and Features

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

Optional machines and features you can order are shown in Appendix B, "Summary of the Configurations." The following optional features are available for Models 520, 640, 660, 740, 820, 860, and 900:

- Vector facility
- Central storage increments
- Expanded storage
- ESCON channels
- Parallel channels
- Integrated Cryptographic Feature (ICRF)
- Sysplex Timer attachment.

To augment the standard machine requirement of three display stations, you can attach as many as three additional 3206 Display Stations Model 100 (or equivalent).

Console and Display Configuration

A console is a device that performs logical functions. A display is a hardware device and is attached to particular ports of the processor controller. See "Standard Requirements" on page 2-3 and "Optional Machines and Features" for standard and optional consoles available. See Chapter 4, "Consoles and Displays" on page 4-1 for more information.

The system and service support displays are controlled by the processor controller and the operator display is channel-attached to Models 520, 640, 660, 740, 820, 860, and 900.

Optional displays are used for one or more of the following console functions:

- System console
- Service console
- Program mode console
- System monitor console
- Service monitor console.

Power and Cooling Functions

Power and cooling functions are combined in one unit, the IBM 9027 ES/9000 Power and Coolant Distribution Unit Model 1.

Power

The 9027 Model 1 distributes the 50/60 Hz power to the ES/9000 Processor Complex.

The 9027 does not contain an I/O power sequence control. To power on and off as many as 192 control units, you can attach an optional IBM 9023 ES/9000 I/O Power Sequencing Unit (IOPS Unit) to the 9027.

Cooling

Water cooling is provided by a two-loop design:

- A system loop integral to the processor complex
- A customer loop.

In the system loop, chilled, distilled water is pumped from the coolant distribution unit (CDU) to the water-cooled components. The water picks up the generated heat and the warm water returns to the CDU where it passes through a heat exchanger and releases the heat to the chilled water flowing in the customer loop.

The 9027 contains the controls that maintain the correct temperature within the densely packed circuits. A thermal warning message is sent to the system console when the processor water temperature exceeds a specified threshold.

If a cooling problem occurs because of a malfunction in the operating pump, the alternate pump is automatically turned on for continued operation.

Input/Output Operations

I/O operations are handled by the channel subsystem in the processor complex. You can configure ESCON and parallel channels for block multiplexer mode of operation.

Byte multiplexer mode of operation can be configured only as parallel channels. There can be as many as 48 parallel channels per side, of which 8 per side can be byte multiplexer channels.

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.

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 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 Models 520, 640, 660, 740, 820, 860, and 900 and other IBM processors is accomplished by using the channel-to-channel adapter (CTCA) feature on those processors 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 processor and the 3088 are supported in both data-streaming and interlock modes.

The 3088 Model A1 provides two-processor 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 processor channels and can provide the equivalent function of as many as 126 CTCAs. The 3088 Model 2 can interconnect as many as eight processor channels.

ESCON Channels (Using Link-Level and Device-Level Protocols): You can achieve ESCON channel-to-channel connections between Models 520, 640, 660, 740, 820, 860, and 900 and other IBM processors with ESCON channels if one of the ESCON channels is set up in channel-to-channel mode.

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

Storage Operations

A hierarchical storage structure increases performance and contributes to system reliability. Each central processor contains a 256KB high-speed buffer (cache), divided into two 128KB halves:

- One used exclusively for data
- One used exclusively for instructions required by the CP

Central storage provides storage capacity for the Models 520, 640, 660, 740, 820, 860, and 900 as shown in "Hardware Configuration" on page 1-3. Central storage is shared by all central processors. Expanded storage (an optional feature) is available for ES/9000 models as shown in the table in Appendix B, "Summary of the Configurations."

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, which represents an address range of 2GB (2 147 483 648 bytes). 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.

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. Models 520, 640, 660, 740, 820, 860, and 900 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.

System Security

Data integrity features, a two-level system access control, and the Integrated Cryptographic Feature (ICRF) contribute to a high level of system security. Customer planning and management are responsible for the implementation and adequacy of the following controls and for the use of the privileged operator controls such as display and alter storage.

Data Integrity

Data integrity is maintained through:

- Key-controlled storage protection (store and fetch)
- Low-address storage protection
- Storage error checking and correction
- Parity and other internal error checking
- Page protection
- Block multiplexer channel command retry
- Remote support authorization
- Clear reset of registers and main storage.

System Access Control

System access control protects against inadvertent system damage by restricting commands and the use of display frames only to persons at specified authorization levels. System access control is implemented through a hierarchical structure such that a user will have access to functions at a specified level as well as all levels below the specified level. Access levels can be defined for the system console and the service console.

Also, the 3206 Display Station provides:

- A security keylock on the display that allows authorized access and that prevents unauthorized access.
- User authorization for remote access (remote support facility) necessitates the matching of a user-assigned access code, as well as enablement of automatic dialing.
- The Integrated Cryptographic Feature (ICRF) is an optional integrated encryption feature, coupled with the central processor, that is capable of high performance in protecting both high-speed transactions and large amounts of data being stored or transmitted. ICRF adheres to international encryption standards and is tamper-resistant.

  In addition, ICRF employs the Personal Security card to enter crytographic key information.

  For more information. refer to "Integrated Cryptographic Feature" on page 3-7.

IBM 9022 Processor Controller

The IBM 9022 Processor Controller is a stand-alone support unit consisting of dual processors. When using a single-image configuration, one processor is active and the other acts as the backup. The backup processor monitors the active processor thereby providing a high level of availability. In a physically partitioned configuration, both processors are active. The 9022:

- Continuously monitors the ES/9000 operation through direct communication with each component in the processor complex
- Initializes the system
- Distributes Licensed Internal Code (LIC) to control storage
- Monitors voltage levels and coolant temperature
- Provides the control unit function for the attached display stations
- Provides extensive error recording, recovery, and diagnostic support for the processor complex.

The ES/9000 LIC is a fundamental component of the ES/9000 Processor. Each ES/9000 model is delivered with a set of LIC licenses and customized to the machine ordered. The LIC enables the models to operate in accordance with its published specifications.

Changes, such as model upgrades, feature additions, and system engineering, may require an update to the LIC. The updated LIC replaces the existing LIC. The existing LIC must be returned to IBM or erased. LIC is provided on optical disks.

Reliability, Availability, and Serviceability

Models 520, 640, 660, 740, 820, 860, and 900 reduce downtime by using standard features that provide high levels of reliability, availability, and serviceability (RAS).

Reliability

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

- Use of emitter-coupled logic in the thermal conduction modules (ECL/TCM technology) that provides a low intrinsic failure rate
- A dual processor controller that incorporates switchover and initialization of the functional side in the single-image configuration
- Internal direct access storage devices (DASDs) that support switchover in the single-image configuration
- Multiple consoles for monitoring functional console activity and for backup
- Central storage that now attaches to the same board as the TCM, eliminating cables and extra gates
- Subsystem storage protection, for use with CICS/ESA* Storage Protection, prevents application software from overwriting CICS system software and control blocks
- Storage background scrubbing provides continuous monitoring of storage for the correction of detected faults before the storage is used.

Availability

The standard features 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 most situations where a power supply fails.
- Four or more central processors available on certain models.
- Concurrent CP TCM maintenance, which allows one of the central processors to be varied offline and CP TCMs to be replaced without having to take down active partitions. Following CP maintenance you can restore the CP online as a shared processor.
- 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 the one parallel or two ESCON channels on the card without having to take down the system or the channel group.
- Fault-tolerant dynamic memory array, which enables the hardware to detect a failure in central or expanded storage chips and dynamically select a backup chip located on the same memory card.
- Automatic error detection and correction in both central storage, common storage buffer element (CSBE), and expanded storage
  - Single-bit error correction and double-bit error detection in central storage and CSBE
  - Single-bit and double-bit error correction and triple-bit error detection in expanded storage
- Storage deallocation in 4MB increments under system program control
- Dynamic reconfiguration management 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 with MVS/ESA SP Version 4.2.0

- Processor availability facility which enhances system availability and offers additional protection from problems that could impact end users or critical programs. This facility transparently moves programs in process from a failing CP to another operational CP and is available in basic mode and LPAR mode.
- Ability to vary offline parallel channels in single channel increments
- Customer problem analysis effects recovery without a service call.

Serviceability

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

- The location of many functional elements on power boundaries
- 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. Processor Components

Overview 3-2
Central Processors 3-3
  Control Storage Element 3-3
  Buffer Control Element 3-4
    High-Speed Buffer 3-4
    Buffer Directory 3-4
    Translation Lookaside Buffer 3-4
    Dynamic Address Translation 3-4
  Instruction Element 3-4
  Execution Element 3-5
  CPU ID 3-5
Vector Facility Feature 3-6
Integrated Cryptographic Feature 3-7
Storage 3-8
  Common Storage Buffer Element 3-9
  Central Storage Element 3-9
    Error Checking and Correction for Central Storage 3-9
    Dynamic Page Deallocation 3-9
    Key-Controlled Storage Protection 3-10
    Hardware System Area 3-10
      Sizing HSA in Basic Mode 3-10
      Sizing HSA in LPAR Mode 3-10
  Expanded Storage 3-11
    Error Checking and Correction for Expanded Storage 3-11
Channel Subsystem 3-11
  Interconnect Communication Element 3-12
  Channels 3-12
    Block Multiplexer Mode of Operation 3-12
    Byte Multiplexer Mode of Operation 3-13

Overview

Models 520, 640, 660, 740, 820, 860, and 900 consist of components (see Figure 3-1) that execute instructions and commands and perform storage and channel operations. The components of the processor unit are:

- One or more central processors
- Central storage
- Optional expanded storage
- One or two channel subsystems
- One or two system control elements.

The vector facility feature is an optional feature that is available for each central processor. The Integrated Cryptographic Feature (ICRF) is an optional feature that is available for one central processor on each side of the processor unit. ICRF is mutually exclusive with the vector facility on a central processor.

The Models 520, 640, and 740 have one processor unit. The central processors, processor storage, and the channel subsystem are contained in the processor unit.

Models 660, 820, 860, and 900 have one processor unit with two sides (side 0 and side 1). The central processors, processor storage, and the channel subsystems are each associated with either side 0 or side 1. In a single-image configuration, the two sides function as one. In a physically partitioned configuration, the two sides are logically and physically distinct, but are located within the same processor unit.

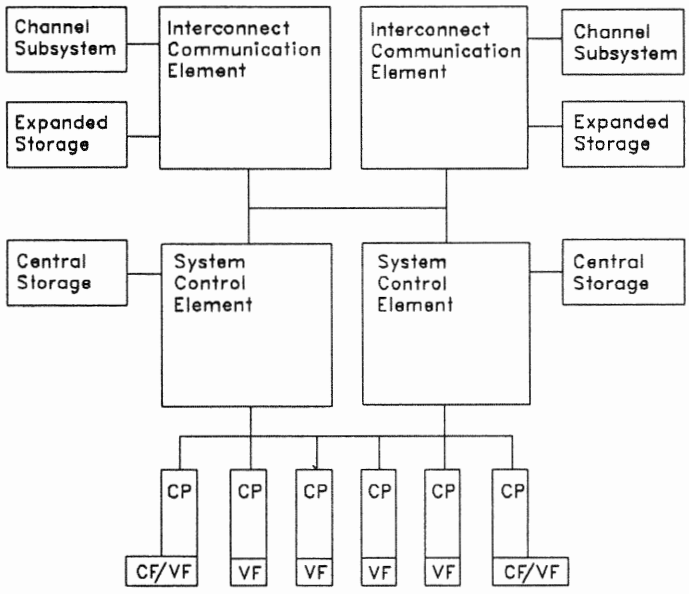

Figure 3-1. Components for Model 900

Central Processors

Each central processor contains:

- A control storage element (CSE)
- A buffer control element (BCE)
- An instruction element (IE)
- An execution element (EE).

The control storage element fetches microinstructions that control instruction execution in the execution element and instruction element. The buffer control element controls the transfer of data between central storage and the central processor containing that BCE. Dynamic address translation is an automatic function of the BCE. Figure 3-2 shows the elements of the central processor.

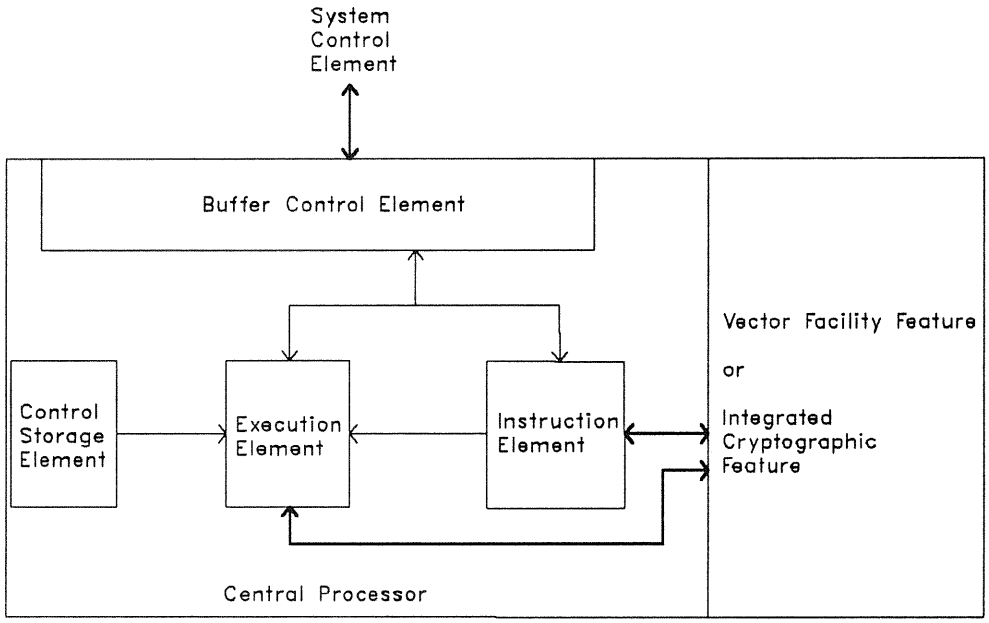

Figure 3-2. Elements of a Central Processor

The system control element (SCE) accepts and processes storage requests from the central processors and the channel subsystem. The SCE analyzes each request and performs the following actions:

- Establishes request priority
- Processes requests
- Performs error checking
- Performs error reporting
- Handles request responses
- Performs cross-interrogation (to ensure that the requester receives the most recent copy of data that is shared).

Control Storage Element

The control storage element (CSE) controls the LIC in the central processor and contains the registers that are used by the central processors. The LIC controls the execution element.

Buffer Control Element

The buffer control element (BCE) handles all central processor requests to and from central storage, performs dynamic address translation, and controls the high-speed buffer. The function of the BCE is to provide the CPUs with very quick access to data.

The buffer control element consists of:

- A 256KB high-speed buffer
- A buffer directory
- A translation lookaside buffer (TLB)
- Dynamic address translation (DAT) hardware.

High-Speed Buffer

The high-speed buffer provides much faster access to data than if the data were in central storage. The high-speed buffer is transparent to programs that are running. When data is referred to during instruction execution, the high-speed buffer, the buffer directory, and the TLB are accessed for address comparison.

Buffer Directory

The buffer directory contains the absolute addresses of central storage for data in the high-speed buffer.

Translation Lookaside Buffer

The translation lookaside buffer stores the real address of the referenced page for a translated virtual address in central storage. Therefore, subsequent translations for the same virtual address are not required because the real address is immediately available in the TLB.

Dynamic Address Translation

Dynamic address translation performs high-speed translation from virtual to real addresses for loading the TLB.

Instruction Element

The instruction element (IE) controls the sequencing of all instructions. It performs the following operations:

- Decodes instructions
- Calculates addresses
- Sends fetch requests to the BCE (for instructions and data) in central storage
- Determines fetch priority
- Controls storage requests
- Provides the execution element with:
  - Operation codes
  - Operands
  - Operand addresses.

The IE can process multiple instructions at the same time by handling the instructions in steps. As one instruction is fetched, decoded, and sent to a queue, the IE begins processing another instruction.

Execution Element

The execution element (EE) executes instructions set up by the IE and operates in parallel with the IE. The EE performs the following operations:

- Processes instructions
- Processes interruptions
- Overlaps operations with the IE
- Initiates control functions.

The EE performs the logical decisions, arithmetic functions, and many control functions of System/370 and ESA/390 architecture instructions. Arithmetic results provided by the EE include:

- Fixed point
- Fixed-point multiply
- Convert to binary
- Convert to decimal
- Floating point
- Extended-precision floating point.

CPU ID

A doubleword designated by the second operand address of the Store CPU ID (STIDP) instruction provides information about the ES/9000 Processor Complex. For detailed information refer to the *IBM Enterprise Systems Architecture/390 Principles of Operation*. A summary of the contents of the doubleword follows:

Basic Mode:

VV	A	SSSSS	9021	0000

Figure 3-3. CPU ID - Doubleword Content in Basic Mode

LPAR Mode:

VV	L	PSSSS	9021	0000

Figure 3-4. CPU ID - Doubleword Content in LPAR Mode

Where:

VV is the version code (two hexadecimal digits):

- A1 (Model 520)
- A2 (Model 640)
- A7 (Model 660)
- A3 (Model 740)
- A4 (Model 820)
- B5 (Model 860)
- A6 (Model 900).

ASSSSS is the central processor identification number (six hexadecimal digits). LPSSSS 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 "CPU Address Identification" on page 7-20):

Digit	Model	
1	520	
1 or 2	640	
1	660	(side 0)
3	660	(side 1)
0, 1, or 2	740	
1 or 2	820	(side 0)
3 or 4	820	(side 1)
0, 1, or 2	860	(side 0)
3 or 4	860	(side 1)
0, 1, or 2	900	(side 0)
3, 4, or 5	900	(side 1)

- The next five digits (SSSSS) are, for the basic modes, selected from the serial number of the ES/9000 Processor Unit. In LPAR mode, the logical partition identifier replaces the first of the five digits (PSSSS).

9021 is the processor unit machine type.

0000 is the reserved hexadecimal digit.

Vector Facility Feature

Central processors with the optional vector facility feature provide significantly increased levels of performance for many numerically-intensive engineering and scientific applications. The vector facility feature is an extension of the instruction and execution elements of a central processor. Some of the vector facility feature characteristics follow:

- The vector facility feature performs vector arithmetic and logical operations on as many as 256 sets of operands with a single instruction.
- Arithmetic and logical units can produce a 32-bit or 64-bit sum, difference, or product during each cycle.
- Compound operations can produce both a product and a sum during each cycle.

The vector architecture provides:

- Storage vector addressing
- Contiguous, noncontiguous, and indirect element selection
- Compound multiply-and-add instructions
- Vector results placed in vector registers
- Scalar results placed in scalar registers
- Logical, binary, short floating-point, and long floating-point operands
- Sixty-three instructions with 180 operation codes (106 operation codes are for floating point).

The vector facility feature provides the following enhanced instruction set:

- Vector square root

  - Two vector-register-to-vector register operands (short and long operands)
  - Two storage-to-vector-register operands (short and long operands)
- Multiply-then-add/subtract
  - Two multiply-then-add instructions (short and long operands)
  - Two multiply-then-subtract instructions (short and long operands)
- Load vector interruption index (VIX) from address

  This instruction loads the vector interruption index independently of the rest of the vector status register. This is useful for operations when only some of the elements of a vector are needed.

In addition, an enhanced scalar execution element design provides new square root instructions and a faster divide instruction.

The following support the vector facility:

- MVS/SP Versions 2, 3, and 4, VM/ESA SP, VM/XA SF, VM/XA SP, and VM/SP HPO including automatic support for asymmetric configurations
- Multitasking facility (MTF) under MVS/SP Versions 2 and 3 for assignment of multiple processors to a job
- Advanced Interactive Executive/370 (AIX/370) as a guest of VM/SP, VM/SP HPO, or VM/XA SP
- VS FORTRAN Version 2 with auto-vectoring capabilities
- APL2 Release 3; no modifications of existing APL2 application programs are required
- Engineering and Scientific Subroutine Library (ESSL): a set of high-performance mathematical routines compatible with the vector facility feature of the ES/9000 Processor Complex.

For additional information on the vector facility feature refer to *IBM Enterprise Systems Architecture/370 and System/370 Vector Operations.*

Integrated Cryptographic Feature

The Integrated Cryptographic Feature (ICRF) is an optional feature for a central processor on Models 520, 640, and 740 and for one central processor on each side on Models 660, 820, 860, and 900. ICRF is mutually exclusive with the vector facility feature on a central processor.

Central processors with ICRF experience significant increased levels of system security. ICRF is an extension of the instruction and execution elements of a central processor. Some of the ICRF characteristics follow:

- ICRF provides total system security when using IBM Resource Access Control Facility (RACF), PR/SM, and IBM Integrated Cryptographic Service Facility/MVS (ICSF/MVS).
- ICRF provides hardware instructions that, when running ICSF/MVS, provide cryptographic functions for data privacy, data integrity, cryptographic key installation and generation, electronic cryptographic key distribution, and personal identification number (PIN) processing.

- ICRF requires physical metal keys to operate key locks that select the secure modes on the key storage unit.
- ICRF includes hardware implementation of the Data Encryption Standard (DEA implementation) cryptographic algorithm.
- ICRF provides two options for entering cryptographic key information into the KSU:
  - A Personal Security card reader integrated on the front panel of the KSU
  - A keypad on the front panel of the KSU

  Note: The hand held KEU is no longer required.

The IBM Personal Security card provides ease of use for master key entry and should:

- Shorten the time required for entry of the system master keys
- Eliminate the possibility of manual errors during master key entry
- Improve overall system security by eliminating data transfer from the KEU to the KSU

In addition, the display screen on the KSU now displays contextual help and error messages for ease of use.

You can write master key information to the Personal Security card offline on an IBM 4754 Security Interface Unit Model 1. When you enter master keys into the KSU, the Personal Security cards containing the requisite key information are sequentially entered into the KSU and read directly into the secure boundary. In addition to master keys, you can enter other key types in the same way.

The Personal Security card requires the following hardware:

- One IBM 4754 Security Interface Unit Model 1
- One IBM Personal System/2* (PS/2*) computer or IBM Personal Computer AT* (or equivalent) with a serial port

The integrated keypad on the front panel of the KSU allows manual entry of cryptographic key information. This provides:

- Key entry consistency during the transition from KEU manual key entry to Personal Security card automatic key entry
- Additional time to plan for installation of the Personal Security card

ICRF is supported by ICSF/MVS and is compatable with:

- MVS/SP Version 3 Release 1.3 and subsequent releases
- TSO/E Version 2 Release 1 and subsequent releases
- ISPF Version 2 Release 3 and subsequent releases
- RACF Release 8 and subsequent releases, or an equivalent access control product.

Storage

Storage is implemented in monolithic and large scale integration technologies. The storage subsystem hierarchy for Models 520, 640, 660, 740, 820, 860, and 900 consists of four elements or levels:

L1 Buffer control element
L2 Common storage buffer element

L3 Central storage element
L4 Expanded storage element

The buffer control element is described under "Buffer Control Element" on page 3-4.

Common Storage Buffer Element

The common storage buffer element (CSBE) is the intermediate level of storage that helps reduce the gap between processing and storage access time. The CSBE consists of the common storage buffer (CSB) and the common storage control (CSC).

The common storage buffer element provides 2MB of storage over four 512KB arrays called interleaves. The CSC controls access to the CSB and manages other system functions such as storage protection, interlocks, and broadcasts. When data is current in the CSB, minimum fetch time by a CPU is only 11 cycles. The difference between access time in the buffer control element and the access time in the common storage element is primarily a result of the data transmission time between boards as well as data management overhead.

Central Storage Element

Central storage contains the logic for:

- Data storage and retrieval for the processor complex
- Communication with and control of the optional expanded storage
- Error checking and correction (ECC)
- Communication with the processor complex (via the system control element).

The central storage available for each model is shown in Appendix B, "Summary of the Configurations."

The addressable portion of central storage is synonymous with main storage, as described in the *IBM Enterprise Systems Architecture/390 Principles of Operation*. A hardware system area (HSA) is reserved within central storage to contain specific system information and cannot be addressed by user programs.

Error Checking and Correction for Central Storage

Error checking and correction (ECC) code bits are stored with data in central storage. Single-bit errors detected during data transfer are corrected. Certain multiple-bit 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.

Dynamic Page Deallocation

A dynamic page deallocation technique, under system program control, temporarily recovers some double-bit failures and allows the operating system to deallocate the failing page frame if the page is not fixed to that page frame. In these instances, the job does not end and processing continues normally. Central storage can be deallocated in 4KB increments under system program control. Maintenance service can be deferred until a predetermined amount of central storage has been deallocated.

Key-Controlled Storage Protection

Key-controlled storage protection provides both store and fetch protection. It prevents unauthorized access to or modification 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 currently active program status word (PSW) are compared with bits 0—3 from the storage key associated with the 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.

Hardware System Area

When you initialize the configuration the hardware system area (HSA) reserves a section of central storage to contain system information. The HSA section is not available for program use. Several factors affected the size of the HSA:

- Mode of operation: basic mode or LPAR mode
- Characteristics of the I/O configuration as defined in the I/O configuration data set (IOCDS)
  - Number of ESCON channels
  - Number of subchannels

Sizing HSA in Basic Mode: To determine the size of the HSA in basic mode, keep in mind the following:

- A minimum of 7MB is automatically allocated to support:
  - The first 4K (K equals 1024) subchannels
  - All parallel channel paths
  - The first 64 ESCON channel paths defined in the IOCDS (whether or not they are physically installed)
- Additional HSA is allocated in increments of:
  - 128KB for each additional 16 ESCON channel paths (or portion)
  - 128KB for each additional 1K subchannels (or portion)
- A maximum of 11MB can be allocated.

Sizing HSA in LPAR Mode: To determine the size of the HSA in LPAR mode, keep in mind the following:

- A minimum of 12MB is automatically allocated.
- A maximum of 20MB can be allocated.

The actual amount allocated in LPAR mode depends on the number of logical partitions and the number of subchannels and ESCON channel paths defined in the IOCDS.

Table 3-1 on page 3-11 shows the size of the HSA based on the system configuration.

Table 3-1. HSA Sizing Table in LPAR Mode				
Number of Logical Partitions	**Number of Subchannels**		**Number of ES Connection CHPs**	**HSA (MB)**
1 to 3	up to 12K	and	up to 64	12
1 to 3	more than 12K	or	more than 64	16
4 to 5	up to 24K	and	up to 256	16
6 to 7	up to 12K	and	up to 64	16
6 to 7	more than 12K	or	more than 64	20

For additional information on how the size of HSA can affect the amount of central storage available to logical partitions, see the *ES/9000 ES/3090 Processor Resource/Systems Manager Planning Guide*, GA22-7123.

Expanded Storage

Expanded storage is an optional high-speed, high-volume, electronic extension of central storage that is accessed in synchronous 4KB increments. Expanded storage reduces the paging and swapping load to channels. It is located on the ICE board. To see a list of expanded storage available for each model refer to Appendix B, "Summary of the Configurations."

Data movement between central storage and expanded storage is initiated by the control program. No data can be transferred to expanded storage without passing through central storage.

Error Checking and Correction for Expanded Storage

Error checking and correction (ECC) code bits within expanded storage are used to permit:

- Single-bit and double-bit error detection and correction
- Triple-bit error detection
- Some multiple-bit error detection.

Unrecoverable errors are flagged.

Channel Subsystem

Models 520, 640, and 740 contain one channel subsystem (CSS). Models 660, 820, 860, and 900 each contain two channel subsystems (CSSs), one on each side. In a single-image configuration, the two channel subsystems appear to the control program as a single dynamic channel subsystem.

The following information applies to:

- A channel subsystem for a processor that is not operating in a physically partitioned configuration
- Either side 0 or side 1 channel subsystem of Models 660, 820, 860, and 900 in a physically partitioned configuration
- Each logical partition of a logically partitioned processor.

A channel subsystem consists of one interconnect communication element (ICE) and as many as 128 channels. A channel controls each channel interface and a channel element controls each group of eight channels. Figure 3-5 on page 3-14

shows a representative channel subsystem for ES/9000 Models 520, 640, 660, 740, 820, 860, and 900.

The channel subsystem handles all I/O operations for the central processor. The CSS controls the communication between a configured channel and the control unit and device. The I/O configuration data set (IOCDS) defines:

- The channel paths on the processor complex
- The mode of operation of the channel paths
- The control units attached to the channel paths, switches, and links
- The I/O devices assigned to the control units.

The IOCDS is created by the I/O configuration program (IOCP) and is stored on DASD internal to the processor controller. At system initialization, the IOCDS information is used to build the necessary control blocks in the hardware system area of central storage.

Interconnect Communication Element

The interconnect communication element (ICE) interacts with central storage, the central processors, and the channels to provide the following services:

- Initialize and end all channel operations
- Centralize storage access control
- Prioritize I/O operations
- Detect and decode central processor (CP) requests for I/O activity
- Fetch, update, and restore I/O control blocks
- Detect and decode channel requests for service
- Present I/O interruptions to the CPs
- Establish communication with expanded storage.

Note: The ESA/390 control blocks represent logical channel queues and sub-channels (devices).

Channels

The channels control all data flow between the interconnect control element and the attached control units. Each channel operates independently to handle all interface sequences, CCW fetches, and data transfers.

Each channel is initialized at system initialization (or when assigned to a logical partition) for one of the following modes of operation:

- All ESCON and parallel channels can be initialized for ESA/390 block multiplexer mode of operation.
- Only parallel channels can be initialized for ESA/390 byte multiplexer mode of operation.

Block Multiplexer Mode of Operation

All ESCON and parallel channels can be configured for block multiplexer mode of operation.

In block multiplexer mode of operation, ESCON channels attached to a 9034 and parallel channels can operate in either:

- Interlocked (high-speed transfer) mode
- Data-streaming mode.

They can be attached to control units operating in:

- High-speed transfer mode
- Data-streaming mode.

ESCON channels using serial transmission of information can be attached to control units that operate using link-level and device-level protocols. Data rates can be as high as 4.5 million bytes per second for data-streaming mode and as high as 17 million bytes per second for serial transmission over fiber optic cables.

On Models 520, 640, 660, 740, 820, 860, and 900, you can connect parallel and ESCON channels to the same control unit provided they are attached to a 9034.

Byte Multiplexer Mode of Operation

When operating in byte multiplexer mode, you can configure only parallel channels but you can configure as many as 16 parallel channels (8 per side). If you have any parallel channels that you do not need for byte multiplexer mode you can configure them for block multiplexer mode.

In byte multiplexer mode of operation you can use parallel channels in byte multiplexer mode or in burst mode. Byte multiplexer mode permits the concurrent operation of several relatively slow-speed devices.

When logical partitioning is in effect, individual channel paths can be allocated to each logical partition. A channel path can be allocated only to one logical partition at a time. A device can be shared between logical partitions by using a separate channel path from each logical partition. Channel paths can be dynamically reconfigured between logical partitions without requiring a power-on reset of the processor complex.

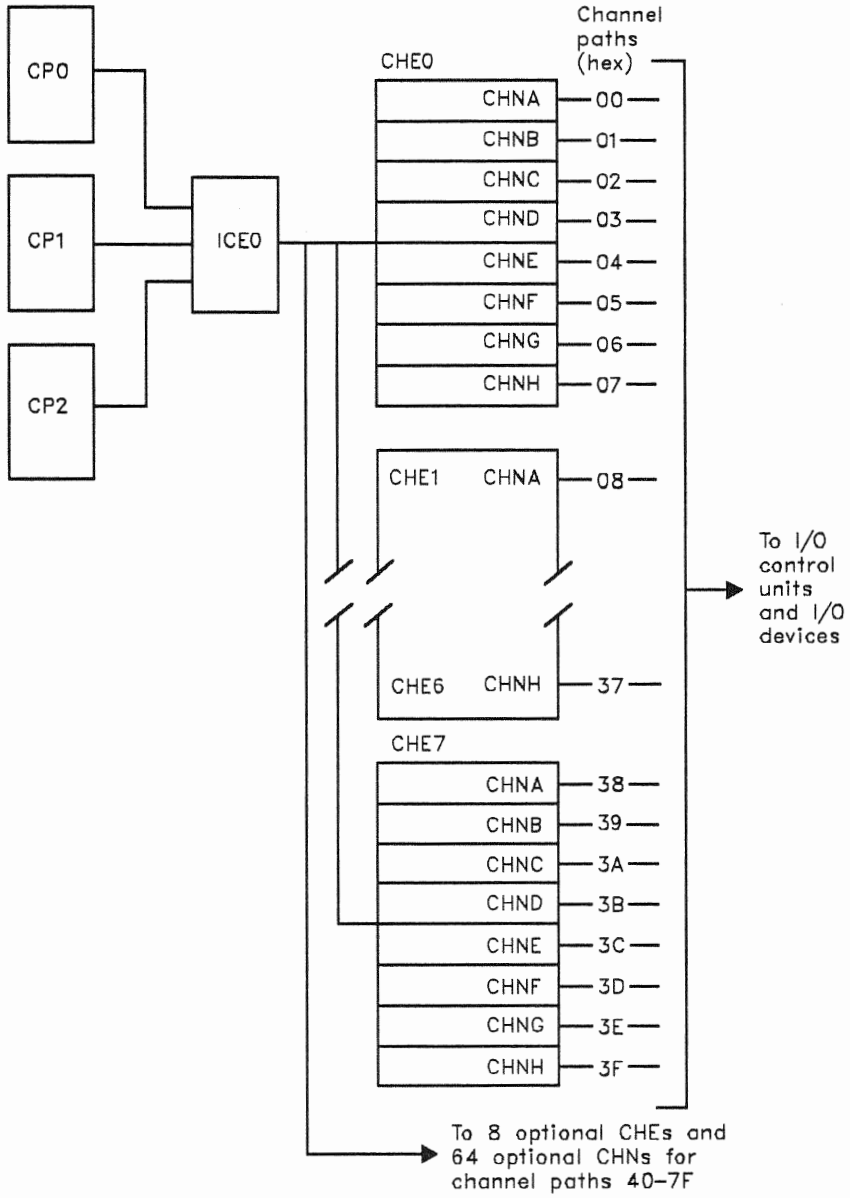

Figure 3-5. Channel Subsystem Diagram

Legend:

ICE Interconnect Communication Element
CHE Channel element
CHN Channel
CP Central processor

Chapter 4. Consoles and Displays

Overview 4-2
System Console Interactive Facilities 4-3
System Definition 4-3
LPAR Mode Definition 4-3
System Activity Display 4-4
I/O Configuration Data Set Content 4-4
I/O Problem Determination Information 4-4

Overview

This chapter describes the required and optional consoles that support Models 520, 640, 660, 740, 820, 860, and 900 and describes the system console interactive facilities.

A console is a device that performs logical functions by means of a display station. A display station is an input/output device.

Note: You can use the IBM 3180 Display Station Model 140 and the IBM 3864 Modem Model 2 with automatic calling unit feature.

Models 520, 640, and 740 require two 3206 Display Stations Model 100 (or equivalent):

- One as a system display (system console)
- One as a service support display (service console).

The operating system requires an additional channel-attached display for an operator console.

Models 660, 820, 860, and 900 require three 3206 Display Stations Model 100 (or equivalent):

- Two as system displays (system consoles)
- One as a service support display (service console).

The operating system requires two additional channel-attached displays for operator consoles for operation in a physically partitioned configuration.

The 9022 Processor Controller Model 1A supports as many as two additional (optional) 3206 Display Stations Model 100 (or equivalent). Each of these displays can be used as a:

- System console
- Service console
- Program mode console
- System monitor console
- Service monitor console.

A Console Assignment frame is provided for assigning selected logical consoles to specified displays. The Swap Cons function key on the display keyboard is used to reassign the logical consoles to a different display. The 3206 Model 100 (Figure 4-1 on page 4-3) and Model 110 are interchangeable.

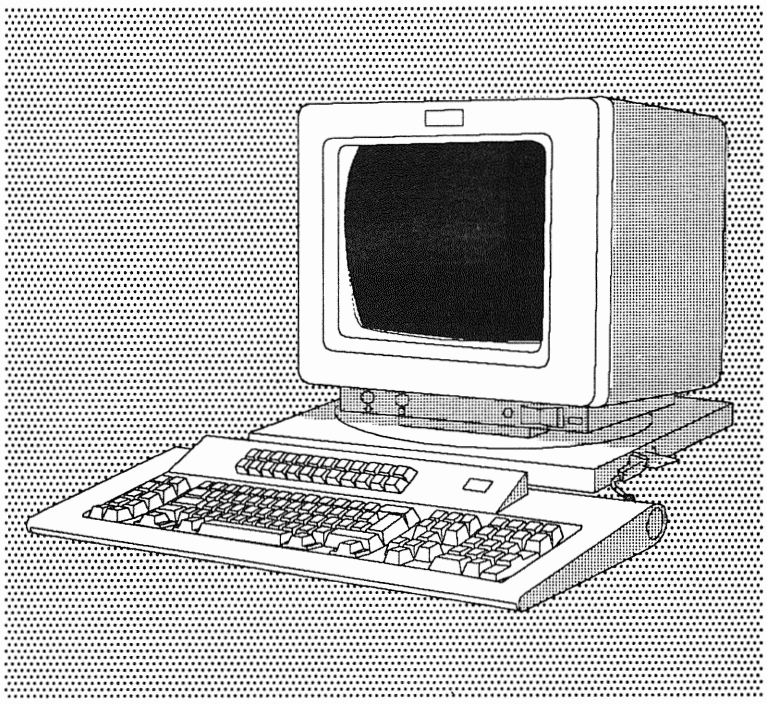

Figure 4-1. IBM 3206 Display Station Model 100

System Console Interactive Facilities

The following functions provide additional user control and flexibility in system operation:

- System definition
- System activity display (SAD)
- I/O configuration data set (IOCDS) content
- I/O problem determination (IOPD) information.

System Definition

The system definition frame provides selections to:

- Supply a unique system name that is displayed on line 24 (system status line) of the screen
- Identify the frame that is to be displayed automatically after a power-on reset
- Enter initial program load (IPL) information (device identification and target central processor) for ESA/390 mode of operation for automatic IPL following a power-on reset.

LPAR Mode Definition

The LPAR mode definition frames provide selections to define:

- The number of CPs
- Central and expanded storage allocation
- Resource allocation parameter.

System Activity Display

A maximum of 24 System Activity Display (SAD) frames (48 in a physically partitioned configuration) can be defined to provide extensive flexibility in the different types of system activity. Frames can be defined to display central processor and channel activity for certain work shifts or types of jobs. SAD frame definition allows you to request that as many as 17 high-usage or low-usage channels or channel paths be dynamically identified and displayed.

A SAD index lists by number all defined SAD frames with their unique names (if specified when defined).

I/O Configuration Data Set Content

I/O configuration data sets (IOCDSs) provide the flexibility to change I/O configurations. Models 520, 640, and 740 have six IOCDSs. Models 660, 820, 860, and 900 have six IOCDSs on each side in a physically partitioned configuration and a total of 12 IOCDSs in a single-image configuration.

Information about specified control units and devices in the I/O configuration is available on the IOCDS frames; they provide information about all:

- Control units (CUs) that are associated with a specified channel. The displayed information includes CU machine type, attached channel paths, protocol (data-streaming or interlock), and port addresses.
- Devices that are accessible from a specified channel path.

I/O Problem Determination Information

I/O Problem Determination (IOPD) frames provide extensive I/O information and can be useful in identifying I/O problems.

The IOPD frames display status information about all installed channel paths, specified subchannel content, shared control units, and device configuration.

Chapter 5. Understanding the 9022 Processor Controller

Overview 5-2
Corequisites 5-2
System Power Panel 5-3
Dual Support Processors 5-3
Operation Monitoring and Control 5-4
Error Recovery 5-4
Configuration 5-4
Remote Support Facility 5-4
  RSF Configuration 5-5
  Authorization 5-5
  Call Details 5-5

Overview

The 9022 Processor Controller Model 1A (see Figure 5-1) monitors and controls the status of all physical units within Models 520, 640, 660, 740, 820, 860, and 900.

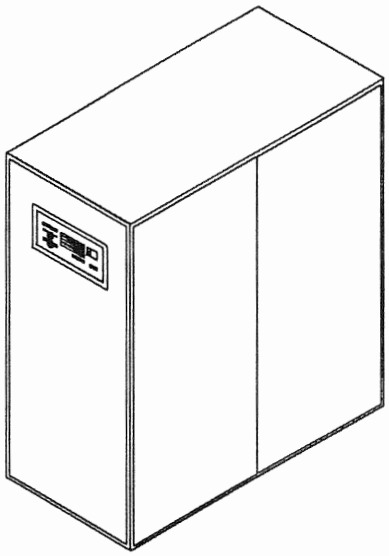

Figure 5-1. IBM 9022 Processor Controller Model 1A

The 9022 Model 1A supports:

- Power on and power off (including I/O units)
- Selection of basic mode or LPAR mode
- System initialization
- Control of the configuration of hardware elements
- Control unit function for required and optional consoles and optional printers
- Monitoring of power supplies, temperature, and water flow
- Control and assistance for error recovery
- Automatic analysis of data (analysis routines) for field-replaceable unit (FRU) isolation
- Collection and storage of error data (logout data) for later analysis
- Full processor complex remote support capability
- Problem analysis procedures
- Collection of information for System Activity Display (SAD) frames
- Collection of information for I/O Problem Determination (IOPD) frames
- Collection of status information for Problem Analysis (PA) frames.

Corequisites

The 9022 requires the following corequisites for full processor complex support:

- Internal DASD
- Internal optical disk storage
- An IBM 5853 Modem Model 1 (or equivalent).

The 5853 Modem Model E is interchangeable with the 5853 Modem Model 1.

System Power Panel

The 9022 system power control panel contains the:

- Power on switch
- Power off switch
- Emergency power off switch
- System configuration switches (if applicable)
- Power status and service mode indicators.

Figure 5-2 shows a typical system power control panel. For more information, see *ES/9000 Models 820 and 900 Operator Guide*, GC38-0086.

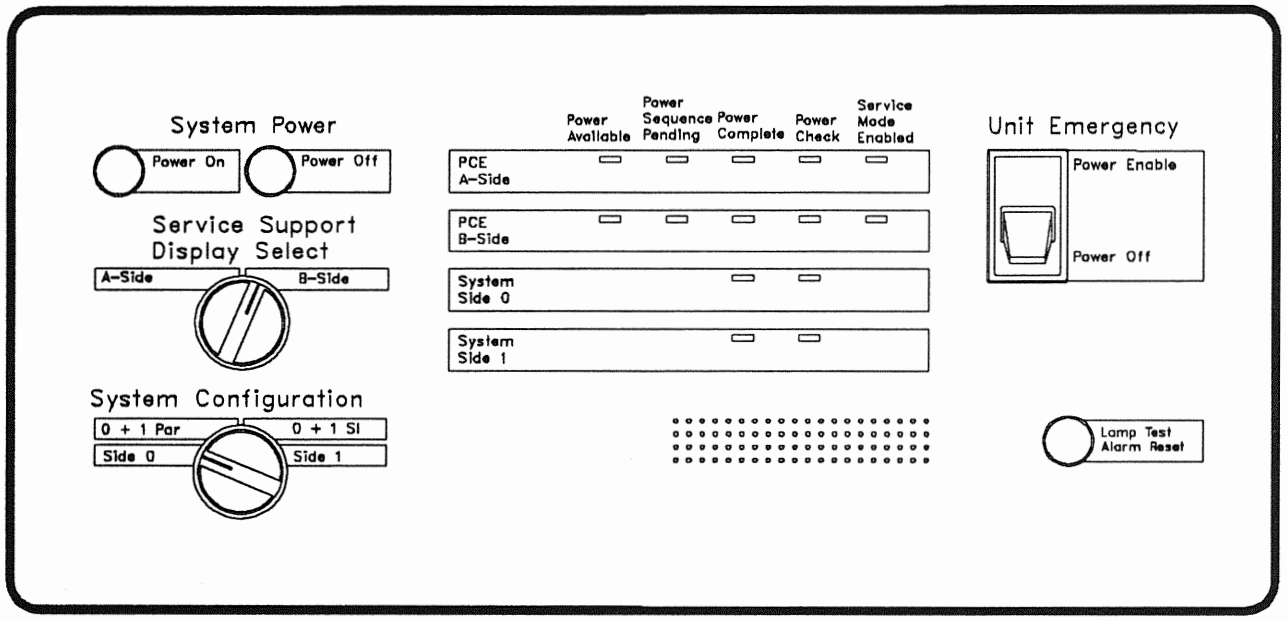

Figure 5-2. 9022 Model 1A System Power Panel

Dual Support Processors

The 9022 Model 1A contains dual processors (A-side and B-side). When used with Models 520, 640, and 740, the backup processor monitors the active processor to help maintain availability. If the active processor fails, a switchover to the backup processor occurs.

When you press the **Power On** pushbutton, both A-side and B-side of the 9022 automatically power on and an initial machine load (IML) and IPL occur. The battery-operated clock (BOC) is read and the processor controller time-of-day (TOD) clock is set. Each side communicates its state to the other side to set the active and backup sides. If both sides are functional, the A-side is made active and the B-side is the backup. If both sides are functional but the states of the DASD differ, the side with the most current DASD is made active. If one side is not functional, the functioning side is made active.

The 9022 Model 1A with specify code 5050 operates as described for Models 660, 820, 860, and 900 in single-image configuration. For a physically partitioned configuration, both sides are active. The A-side of the 9022 controls side 0 of an Models 660, 820, 860, and 900 in a physically partitioned configuration, and the

B-side of the 9022 controls side 1 of an Models 660, 820, 860, and 900 in a physically partitioned configuration.

Operation Monitoring and Control

The 9022 monitors and controls the operations of the ES/9000 Processor Complex. The 9022 initializes and distributes power to all ES/9000 Processor Complex components and to all interconnected I/O control units that are under power-sequence control. During initialization of the ES/9000 Processor Complex, the 9022 validates areas of central storage as error-free data locations, records failing storage locations, and assigns the hardware system area in central storage based on contiguous error-free storage locations. After power sequencing is complete, the processor controller performs an IML.

During processing, the 9022 monitors voltage levels and coolant flow. If the coolant flow rate significantly decreases or stops, the second pump in the 9027 is switched into the coolant circuit to avoid thermal shutdown.

Error Recovery

The 9022 logs errors as they occur and then analyzes them for service personnel. Failure symptoms that are saved at the time of a malfunction are analyzed on a time-sharing basis with other processor controller functions; this operation is concurrent with operation of the processor complex.

The 9022 saves the symptoms of these errors, correlates multiple symptoms, performs error analysis, and isolates the failure to the failing FRU or group of FRUs. When automatic error-recovery attempts fail or the error occurs frequently, failure information is displayed on the system console, and an audible alarm is sounded to alert the operator of a problem requiring action. During IML, similar notification to the operator occurs when loss of storage exceeds a threshold that may degrade system performance.

Configuration

You can use the 9022 to initiate configuration changes in:

- Central storage
- Central processors
- Online channels.

However, operator-initiated action through the system control program is the preferred method of reconfiguration because it makes both hardware and software changes as a single, integrated action.

Remote Support Facility

The Remote Support Facility (RSF) consists of three parts:

- Configuration
- Authorization
- Call details.

RSF Configuration

At each installation, the service representative tailors RSF to customer requirements. RSF comes with two configuration frames to specify RSF information unique to each site. The configuration outlines:

- The number of telephones (up to four) having access to RSF
- The mode of operation of the equipment used (for example, manual dialing, automatic dialing, modem type)
- The agreed-on service update schedule
- The allowance or disallowance of incoming and outgoing calls
- The unique RSF customer access code
- Additional customer and IBM information:
  - The names of the responsible personnel
  - The location of the system at the site
  - The customer's business, system console, and modem telephone numbers
  - The type of modem
  - The IBM branch office number and telephone number
  - The prime shift and off-hour shift dispatcher's telephone numbers.

You can change the contents of the configuration by invoking the frames.

Authorization

No incoming or outgoing calls are allowed without customer authorization. Customer authorization includes:

- Reason for call
- Type of call: outgoing, incoming, automatic calling, manual dialing
- Name of the person at the customer location to be contacted
- Whether a service representative is on-site
- Whether the system is immediately available
- Whether RSF can be enabled, can be deferred, or is not authorized.

Call Details

Details of all RSF calls (both incoming and outgoing) are recorded and each call is listed on an RSF log index. The RSF log index can be displayed and any call that is listed can be selected from the index. The RSF log index shows the date, time, status, and reason for the call. When a call is selected, the first of a set of frames is displayed. The frames contain detailed information about the selected call.

For information about customer problem analysis and RSF procedures refer to Chapter 6, "Handling Errors."

Chapter 6. Handling Errors

Overview 6-2
Automatic Error Recovery 6-2
  Error Checking and Correction 6-2
  Machine-Check Handling 6-2
  I/O Operations 6-4
Problem Analysis 6-5
  Power Malfunction 6-5
  Other Malfunctions 6-5
Remote Support Facility 6-5

Overview

The 9022 Processor Controller automatically performs certain error recovery procedures. If automatic error recovery is not successful, a set of Problem Analysis frames and procedures is available to facilitate recovery by the user for certain types of failures, and remote support facility (RSF) procedures are available for all failures.

Automatic Error Recovery

Error handling by the processor controller provides both automatic recovery from many hardware malfunctions and reporting by machine-check or channel-check interruption.

Error recovery functions are provided for errors in central storage and for channel errors. When an error is detected, the 9022 automatically performs error analysis procedures to isolate the malfunctioning area directly and to identify (if applicable) the field-replaceable unit (FRU) or group of FRUs. These procedures include problem recognition, recording, and diagnosis.

Error Checking and Correction

Error checking and correction in central storage provides automatic single-bit error detection and correction of all data read from central storage. Error checking also detects all double-bit errors and some multiple-bit errors of the data read from central storage, but does not correct the errors.

Some double-bit errors are temporarily recovered to allow the system control program an opportunity to deallocate the failing page frame.

Error checking and correction in expanded storage provides automatic single-bit and double-bit error detection and correction of all data read from expanded storage. Error checking also detects all triple-bit errors and some multiple-bit errors of the data read from expanded storage, but does not correct the errors.

Parity checking is used to verify other data in the processor complex that is not contained in central or expanded storage.

Machine-Check Handling

If the 9022 is operating in basic mode and a machine-check occurs, the 9022 collects the error information and records the information in a log. The central processor then presents a machine-check interruption to the system control program.

All machine-check interruptions store a doubleword that contains a machine-check interruption code. Bits that are not assigned or are not implemented are stored as 0's. Certain bits in control register 14 are associated with machine-check handling. Models 520, 640, 660, 740, 820, 860, and 900 use bits 3—7 in ESA/390 mode.

When a malfunction makes it undesirable or impossible to continue processing, the central processor enters the check-stop state. The central processor always enters the check-stop state when any of the following occurs:

- PSW bit 13 is 0 and an exigent machine-check condition occurs.
- Another exigent machine-check condition is detected during the execution of an interruption that was caused by an exigent machine-check condition.
- The machine-check interruption code cannot be stored or the new PSW cannot be fetched during the execution of an interruption.
- An invalid ECC is detected in the prefix register.

The central processor is removed from the check-stop state by a CPU reset.

Note: When multiple central processors are part of the configuration, the processor controller generates a malfunction external interruption to all the configured central processors. If the processor controller cannot identify which central processor should be put into the check-stop state, all central processors are put into the check-stop state.

Table 6-1 on page 6-4 shows the machine-check interruption code (MCIC).

Table 6-1. Machine-Check Interruption Code (MCIC)	
MCIC Bit	**Meaning**
0	System damage
1	Instruction processing damage
2	System recovery
4	Timing facility damage
5	External damage
6	Vector facility failure
9	Channel report pending
10	Service-processor damage
11	Channel subsystem damage
13	Vector facility source
14	Backed up
16	Storage error uncorrected
17	Storage error corrected
18	Storage key error uncorrected
19	Storage degradation
20	PSW-EMWP validity
21	PSW mask and key validity
22	PSW program mask, condition code validity
23	PSW instruction address validity
24	Failing storage address validity
27	Floating-point register validity
28	General register validity
29	Control register validity
31	Storage logical validity
32	Indirect storage error
33	Access register validity
34	Delayed access exception
46	CPU timer validity
47	Clock comparator validity

I/O Operations

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

- Interface control check (IFCC)
- Channel control check (CCC)
- Channel data check (CDC).

Machine-check interruptions include the following:

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

Problem Analysis

To attempt recovery before initiating the remote support facility, the operator can invoke problem analysis from the System Console Index frame. If the problem was caused by a power malfunction, the first of a set of Power Status Problem Analysis frames display. If the problem was caused for some other reason, the first of a second set of Problem Analysis frames display.

Power Malfunction

To assist recovery from a power malfunction, the first set of Problem Analysis frames display:

- The power boundary errors with reason codes
- A list of suggested recovery actions
- A defined service action.

For more information, see the *ES/9000 Models 520, 640, 660, 740, 820, 860, and 900 Recovery Guide*, GC38-0090.

Other Malfunctions

Status information is displayed for the central processors, hardware, interface control checks, and channels or channel paths. Based on the status information, the operator can select any of a variety of problem analysis categories that include:

- Non-I/O hardware errors
- Unsuccessful IPL
- Enabled or disabled wait state
- Interface control checks (IFCCs)
- I/O device errors
- Operator console lockout.

Each procedure provided by the Problem Analysis frames gives current status information for that type of malfunction, and then lists possible recovery actions. The IOPD frames can also be used for troubleshooting IFCC and I/O device problems. If the attempt at recovery fails and if remote support is required, a selection from the frame invokes the RSF Authorization frame.

Remote Support Facility

The operator can initiate the remote support facility from the problem analysis procedures or invoke the RSF Authorization frame and establish the remote connection from there.

After the service request is authorized, a telephone number is automatically dialed over the public switched network to establish a connection with a remote modem. The remote modem acknowledges the connection and the remote support facility is enabled.

When the remote support facility is connected by the data link, RSF has access to and control of the ES/9000 Processor Complex.

Chapter 7. Standard and Optional Features

Overview 7-4
Features Dependent on Architectural Mode 7-5
Features Not Dependent on Architectural Mode 7-8
Programming Assists Dependent on Architectural Mode 7-9
Describing the Features 7-10
Access-List-Controlled Protection 7-10
Advanced Address Space 7-10
Access-Register Mode 7-10
Linkage Stack Mechanism 7-10
Home-Space Mode 7-11
Real-Storage Access 7-11
Data Movement 7-11
Basic Control Mode 7-11
Bimodal Addressing 7-11
Branch and Save 7-11
Broadcasted Purging 7-11
Byte-Oriented Operand 7-11
Cancel I/O 7-11
Channels 7-12
Channel Indirect Data Addressing 7-12
Channel-Set Switching 7-12
Channel Subsystem 7-12
Channel-Subsystem Call 7-12
Clear I/O 7-12
Command Retry 7-12
Concurrent Sense 7-13
Conditional Swapping 7-13
Concurrent CP TCM Maintenance 7-13
Concurrent Channel Maintenance 7-13
Concurrent Power Subsystem 7-13
Console Integration 7-13
Control-Switch Assist 7-13
CPU Retry 7-14
CPU Timer and Clock Comparator 7-14
Data Streaming 7-14
DB2 Sort Facility 7-14
Dynamic Reconfiguration Management 7-15
Enhanced Move Page for VM 7-15
Enhanced Power Subsystem 7-15
Error Checking and Correction 7-15
ESCON Channels 7-16
ESCON Extended Distance Feature (ESCON XDF) 7-16
ESCON Channel-to-Channel Adapter 7-16
Expanded Storage 7-16
Extended-Precision Divide 7-16
Extended-Precision Floating Point 7-16
Extended Real Addressing 7-16
Extended Sorting 7-17
Extensions for Virtual Machines 7-17
Fast Release 7-17
Fault-Tolerant Dynamic Memory Arrays 7-17

Floating Point 7-17
Halt Device 7-17
High-Speed Buffer Storage 7-17
Incorrect-Length Indication-Suppression Facility 7-18
Integrated Cryptographic Feature 7-18
Interpretive Execution 7-18
Interval Timer 7-18
I/O-Device Self-Description 7-19
I/O Error Alert 7-19
I/O Interface Reset 7-19
I/O Power Sequence Control 7-19
Key-Controlled Storage Protection 7-19
Limited Channel Logout 7-20
Logging Volume Reduction 7-20
Monitoring 7-20
Move Page 7-20
Multiprocessing 7-20
CPU Address Identification 7-20
CPU Signaling and Response 7-20
Prefixing 7-20
Shared Main Storage 7-21
TOD Clock Synchronization 7-21
MVS Dynamic Physical Partitioning 7-21
Page Protection 7-21
PER Extensions 7-21
Preferred Machine Assist 7-21
Private Space 7-21
Processor Availability Facility 7-21
Processor Resource/Systems Manager 7-22
Logical Partitioning 7-22
Multiple High-Performance Guests 7-22
PR/SM LPAR Auto-Reconfiguration 7-23
PR/SM LPAR Management Time Reporting 7-23
Retention of Logical Partition Processing Weight Definitions 7-23
PR/SM LPAR Preferred Path 7-23
PR/SM LPAR High-Performance Parallel Interface (HIPPI) 7-24
PR/SM LPAR Support of ICRF 7-24
PSW-Key Handling 7-24
Recovery Extensions 7-24
SCP-Initiated Reset 7-24
Service Signal 7-24
Sorting Instructions 7-24
Storage Background Scrubbing 7-24
Storage-Key Instruction Extensions 7-25
Storage-Key Instructions 7-25
Storage-Key 4KB Block 7-25
Storage-Reconfiguration 7-25
Subsystem Storage Protection 7-25
IBM 9037 Sysplex Timer 7-25
External Time Source 7-26
Automatic Propagation Delay Adjustment 7-26
System/370 Extended Facility 7-26
Test Block 7-26
Time-of-Day Clock 7-26
Tracing 7-27

Translation 7-27
Dynamic Address Translation 7-27
Extended Control Mode 7-27
Program-Event Recording 2 (PER2) 7-27
Set-System-Mask Suppression 7-28
Store Status 7-28
Vector Facility Feature 7-28
Virtual-Machine Assist 7-28
VM Assists for the CPU Timer 7-28
3033 Extension 7-29
Dual-Address Space 7-29
Start I/O Fast Queuing 7-29
Suspend and Resume 7-29
31-Bit Indirect Data Address Word 7-29
31-Bit Real Addressing 7-29

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 ES/9000 Processor Complex.

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

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

Features Dependent on Architectural Mode

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

Table 7-1 (Page 1 of 3). Features Dependent on Architectural Mode

Feature	S/370	ESA/390
Access-list-controlled protection	N/D	Standard
Advanced address space	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-set switching	Standard	N/D (Note 2)
Channel subsystem	N/D	Standard
Channel-subsystem call	N/D	Standard
Clear I/O	Standard	N/D (Note 2)
Command retry	Standard	Standard
Concurrent sense	N/D	Standard
Conditional swapping	Standard	Standard
Console Integration	N/D	Standard
CPU timer and clock comparator	Standard	Standard
DB2 sort facility	Standard	Standard
Dynamic reconfiguration management	N/D	Standard
Enhanced move page	N/D	Standard
ESCON Channels	N/D	Optional
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
Extension for virtual machines	N/D	Standard
Fast release	Standard	N/D (Note 2)
Fault-tolerant dynamic memory arrays	N/D	Standard
Floating point	Standard	Standard
Halt device	Standard	N/D (Note 2)

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 7-1 (Page 2 of 3). Features Dependent on Architectural Mode

Feature	S/370	ESA/390
Incorrect-length-indication suppression	N/D	Standard
Integrated Cryptographic Feature	N/D	Optional
Interpretive execution (SIE)	N/D	Standard
Interval timer	Standard	N/D
I/O-device self-description	N/D	Standard
I/O interface reset	N/D	Standard
Key-controlled storage protection	Standard	Standard (Note 4)
Limited channel logout	Standard	N/D (Note 2)
Logging volume reduction	N/D	Standard
Monitoring	Standard	Standard
Move page	N/D	Standard
Multiprocessing:		
CPU address identification	Standard	Standard
CPU signaling and response	Standard	Standard
Prefixing	Standard	Standard
Shared main storage	Standard	Standard
TOD clock synchronization	Standard	Standard
MVS dynamic physical partitioning	N/D	Standard
Page protection	N/D	Standard
PER extensions	N/D	Standard
Private space	N/D	Standard
Processor availability facility	Standard	Standard
PSW-key handling	Standard	Standard
Recovery extensions	Standard	N/D
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)
Storage-key 4KB block:		
Single-key 4KB blocks	Standard (Note 6)	Standard
Storage-key exception control	Standard	N/D (Note 4)
Storage-reconfiguration	N/D	Standard
Subsystem storage protection	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 7-1 (Page 3 of 3). Features Dependent on Architectural Mode

Feature	S/370	ESA/390
Sysplex timer	N/D	Optional
Automatic propagation relay	N/D	Optional
External time source	Standard	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
Vector facility	Optional	Optional
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.

Notes:

1. Operation in ESA/390 mode is comparable to operation in EC mode of System/370.
2. Replaced by standard functions of the channel subsystem operating in ESA/390 mode; channel program compatibility with System/370 is maintained.
3. Replaced by 31-bit real addressing.
4. 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.
5. Replaced by page protection.
6. Double-key 4KB blocks are not implemented.
7. Does not include dual-address space (DAS) tracing. Address space number (ASN) tracing provides a comparable function.

Features Not Dependent on Architectural Mode

Table 7-2 (Page 1 of 2). Features Not Dependent on Architectural Mode

Feature	Model 520	Model 640	Model 660	Model 740	Model 820	Model 860	Model 900
Parallel Channels							
Parallel Channel group 1st Additional (16, side 0)	Opt	Opt	Opt	Opt	Opt	Opt	Opt
Parallel Channel group 1st Additional (16, side 1)	NA	NA	NA	NA	Opt	Opt	Opt
Parallel Channel group 2nd Additional (16, side 0)	Opt	Opt	Opt	Opt	Opt	Opt	Opt
Parallel Channel group 2nd Additional (16, side 1)	NA	NA	Opt	NA	Opt	Opt	Opt
Parallel Channel group 3rd Additional (16, side 0)	Opt	Opt	Opt	Opt	Opt	Opt	Opt
Parallel Channel group 3rd Additional (16, side 1)	NA	NA	Opt	NA	Opt	Opt	Opt
ESCON Channels							
ESCON Channel group 1st Additional (16, side 0)	Opt	Opt	Opt	Opt	Opt	Opt	Opt
ESCON Channel group 1st Additional (16, side 1)	NA	NA	Opt	NA	Opt	Opt	Opt
ESCON Channel group 2nd Additional (16, side 0)	Opt	Opt	Opt	Opt	Opt	Opt	Opt
ESCON Channel group 2nd Additional (16, side 1)	NA	NA	Opt	NA	Opt	Opt	Opt
ESCON Channel group 3rd Additional (16, side 0)	Opt	Opt	Opt	Opt	Opt	Opt	Opt
ESCON Channel group 3rd Additional (16, side 1)	NA	NA	Opt	NA	Opt	Opt	Opt
ESCON Channel group 4th Additional (16, side 0)	Opt	Opt	Opt	Opt	Opt	Opt	Opt
ESCON Channel group 4th Additional (16, side 1)	NA	NA	Opt	NA	Opt	Opt	Opt
ESCON Channel group 5th Additional (16, side 0)	NA	Opt	Opt	Opt	Opt	Opt	Opt
ESCON Channel group 5th Additional (16, side 1)	NA	NA	Opt	NA	Opt	Opt	Opt
ESCON Channel group 6th Additional (16, side 0)	NA	Opt	Opt	Opt	Opt	Opt	Opt
ESCON Channel group 6th Additional (16, side 1)	NA	NA	Opt	NA	Opt	Opt	Opt
ESCON Channel group 7th Additional (16, side 0)	NA	Opt	Opt	Opt	Opt	Opt	Opt
ESCON Channel group 7th Additional (16, side 1)	NA	NA	Opt	NA	Opt	Opt	Opt
ESCON Channel (16 or 16 each side for MPs)	NA	NA	NA	NA	Std	Std	Std

Table 7-2 (Page 2 of 2). Features Not Dependent on Architectural Mode

Feature	Model 520	Model 640	Model 660	Model 740	Model 820	Model 860	Model 900
Other Features							
Concurrent power subsystem	Std	Std	Std	Std	Std	Std	Std
CPU retry	Std	Std	Std	Std	Std	Std	Std
Data streaming	Std	Std	Std	Std	Std	Std	Std
Enhanced power subsystem	Std	Std	Std	Std	Std	Std	Std
Error checking and correction	Std	Std	Std	Std	Std	Std	Std
High-speed buffer storage	Std	Std	Std	Std	Std	Std	Std
I/O error alert	Std	Std	Std	Std	Std	Std	Std
9023 I/O power sequence control (for 1st to 64th control unit)	Opt	Opt	Opt	Opt	Opt	Opt	Opt
9023 I/O power sequence control (for 65th to 128th control unit)	NA	Opt	Opt	Opt	Opt	Opt	Opt
I/O power sequence control (for 129th to 192nd control unit)	NA	Opt	Opt	Opt	Opt	Opt	Opt
Processor Resource/Systems Manager	Std	Std	Std	Std	Std	Std	Std
Storage background scrubbing	Std	Std	Std	Std	Std	Std	Std
Treshold setting	Std	Std	Std	Std	Std	Std	Std

Programming Assists Dependent on Architectural Mode

Table 7-3. 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 instructions	Standard	N/D
SVC assist instructions	Standard	Standard
Add FRR instruction	Standard	Standard
VM assist for MVS/370 assists	Standard	N/D
Virtual-machine assist	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.

Describing the Features

This section describes the features and programming assists for Models 520, 640, 660, 740, 820, 860, and 900. The feature descriptions should be used with the preceding tables.

Access-List-Controlled Protection

Access-list-controlled protection provides the ability to prohibit store-type storage references to an address space by means of a bit in the access-list entry used to access the space. This allows different users with different access lists to have different capabilities to store in the same address space.

Advanced Address Space

Advanced address-space functions of Models 520, 640, 660, 740, 820, 860, and 900 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.

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.

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.

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.

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.

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.

Basic Control Mode

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

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.

Branch and Save

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

Broadcasted Purging

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

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).

Cancel I/O

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.

Channels

ESCON and parallel channels can be assigned for block multiplexer mode of operation.

Only parallel channels can be assigned for byte multiplexer mode of operation. See Appendix B, "Summary of the Configurations" for standard and optional channels for each model.

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.

Channel-Set Switching

Channel-set switching permits program-controlled switching of channel sets between two central processors so that if one central processor fails, either channel set may be assigned to the other central processor.

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.

Models 660, 820, 860, and 900 have a channel subsystem on each side. However, when these models are in a single-image configuration, the two channel subsystems operate as one dynamic channel subsystem.

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.

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.

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.

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.

Conditional Swapping

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

Concurrent CP TCM Maintenance

Concurrent CP TCM maintenance provides the ability to deconfigure a single CP to perform maintenance on a TCM without having to take down active partitions. Following CP TCM maintenance, you can restore the CP online as a shared processor.

The function is available in basic and PR/SM LPAR mode.

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.

Concurrent Power Subsystem

Concurrent power subsystem provides the ability to replace, repair, and verify most power supplies without turning off system power.

Console Integration

The Console Integration facility, supported by MVS/ESA SP Version 4 Release 2, provides the ability to use the hardware system console during MVS system initialization and during console recovery to recover the MVS operator console. This facility provides an availability enhancement when the MVS master and alternate operator consoles become unavailable. In this case, you can temporarily use the hardware system console as an MVS operator console to reestablish an operational MVS console.

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

Control-Switch Assist

Control-switch assist enhances the functions of the preferred-machine assist by increasing the speed with which interruptions on central processor (CP) owned channels are presented to a preferred virtual machine, and by allowing a preferred virtual machine to access certain control program DIAGNOSE codes. VM/SP HPO Release 3.2 and subsequent releases support control-switch assist.

CPU Retry

CPU retry automatically examines any instruction when an error occurs during the execution of the instruction. CPU retry usually attempts to execute the instruction again.

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.

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.

DB2 Sort Facility

See "Extended Sorting" on page 7-17.

Dynamic Reconfiguration Management

With dynamic reconfiguration management you can modify the I/O configuration without performing:

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

With dynamic reconfiguration management you can add, delete, or modify the I/O definitions of channel paths, control units, and I/O devices in the processor. You can save the changes made to the I/O configuration definitions and apply them to the active I/O configuration data set (IOCDS). To initiate these changes use MVS/ESA SP Version 4 Release 2 and the Hardware Configuration Definition* (HCD is a component of MVS/ESA).

Dynamic reconfiguration management operates in basic and LPAR modes.

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

Enhanced Power Subsystem

The enhanced power subsystem provides the ability to perform concurrent or deferred maintenance on the system power supplies.

A new design of the power supplies, featuring an additional backup power supply, provides continued system operation in most situations where a power supply fails. The design allows the failing power supply to take itself out of the configuration and be replaced by the backup power supply.

Additionally, there is new function that searches for failures among power supplies for deferred service calls.

Error Checking and Correction

Data paths between expanded storage (if installed) 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 the central storage and in the expanded storage data arrays. ECC codes apply to data stored in and fetched from central storage and expanded storage; single-bit and multiple-bit error detection are performed in both central storage and expanded storage. Single-bit error correction takes place in both central storage

and expanded storage. Double-bit error correction takes place in expanded storage.

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.

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.

For more information, see "Fiber Optic Cabling, Transmission, and Reception" on page 1-9.

ESCON Channel-to-Channel Adapter

The ESCON channel-to-channel adapter (ESCON CTCA) provides the same type of function for ESCON channel paths that is available for the parallel channels.

Expanded Storage

Expanded storage is an optional high-speed, high-capacity storage that transfers 4KB blocks to and from central storage. See "Optional Machines and Features" on page 2-4 for the expanded storage available for Models 520, 640, 660, 740, 820, 860, and 900.

Extended-Precision Divide

Extended-precision divide provides the Divide (DXR) instruction for extended-precision floating-point operands.

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).

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.

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.

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.

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.

Fault-Tolerant Dynamic Memory Arrays

Fault-tolerant dynamic memory arrays enable the hardware to determine when a chip has exceeded its error threshold so that it can move the data from that chip to a backup chip located on the same memory card, preventing an unrecoverable error.

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.

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.

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. For more information, see "Buffer Control Element" on page 3-4.

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).

Integrated Cryptographic Feature

The IBM Integrated Cryptographic Feature (ICRF) together with IBM Integrated Cryptographic Service Facility/MVS (ICSF/MVS), provides a system for secure, tamper-detecting, high-speed cryptographic services. It enables you to encrypt, decrypt, and authenticate data, and to generate and manage cryptographic keys.

ICRF includes:

- A thermal conduction module (TCM)
- A key-storage unit (KSU)

If you choose the Personal Security* card reader method of entering encryption keys, you also need:

- An IBM 4754 Security Interface Unit Model 1
- An IBM PS/2 or an IBM Personal Computer AT (or equivalent) with a serial port
- IBM Personal Security cards

ICRF uses the American National Standards Institute (ANSI) Data Encryption Algorithm (DEA), also known as the U.S. National Institute of Science and Technology Data Encryption Standard (DES) algorithm, to encrypt and decrypt data.

ICRF can be installed on Models 520, 640, 660, 740, 820, 860, and 900. On a dual-sided processor complex, one ICRF can be installed on each side of the processor complex. On a single-sided processor complex, only one ICRF can be installed.

ICRF and the vector facility feature are mutually exclusive on a central processor. ICRF should be installed on the highest-numbered central processor for processors that can install the vector facility feature. No changes in installation planning are required for ICRF.

Interpretive Execution

The interpretive execution facility is used by VM/ESA and VM/XA SP and 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.

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.

I/O-Device Self-Description

I/O-device self-description allows a device to describe itself and its position in the I/O configuration.

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.

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 check 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. Use the F-selection field on the System Definition (SYSDEF) frame to enable the function.

For more information, see the ES/9000 Models 520, 640, 660, 740, 820, 860, and 900 Operator Guide, GC38-0086.

I/O Power Sequence Control

The 9023 provides I/O power sequence control. One I/O power sequence control on the 9023 directs the sequence of power on or off for 64 control units. An option is available to extend the power sequence to 2 x 64 increments. This extends the total capability to 192 control units. Models 660, 820, 860, and 900 require two 9023s for the first 64 control units on each side. Two optional additional features can be added to extend the power sequence to 192 control units on each side.

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. For more information refer to "Central Storage Element" on page 3-9.

Limited Channel Logout

Limited channel logout provides 4 bytes of channel status information for model-independent recovery from channel errors.

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.

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.

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*. For VM, see "Enhanced Move Page for VM" on page 7-15.

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. See Appendix B, "Summary of the Configurations" for the number of CPs for each model.

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. See "CPU ID" on page 3-5.

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.

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.

Shared Main Storage

Shared main storage permits all central processors to have access to common main storage locations.

TOD Clock Synchronization

TOD clock synchronization provides a uniform appearance to a clock synchronization program in all ES/9000 Processor Complexes, 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.

MVS Dynamic Physical Partitioning

Improvements to MVS dynamic physical partitioning have significantly reduced the amount of time required to reconfigure central storage.

A new storage-reconfiguration command allows multiple requests for reconfiguration to be made by means of a single communication with the service processor.

In a laboratory test environment, reconfiguring 256MB of central storage was reduced from 16 minutes to 1 minute. This decrease could enhance your ability to move quickly from single-image to physically partitioned operation.

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.

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.

Preferred Machine Assist

Preferred machine assist permits a single MVS/SP virtual-equals-real (V=R) virtual machine operating under VM/SP HPO to operate with a minimum of simulated instruction execution. This allows the MVS/SP virtual machine to achieve near basic performance. With preferred machine assist, any MVS/SP release that supports more than 16MB of real storage can use real storage above 16MB when MVS/SP is operating as a V=R virtual machine.

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.

Processor Availability Facility

The processor availability facility enhances system availability and offers additional protection from problems that could impact end users or critical programs. Currently, if a central processor (CP) encounters a failure, the program in process at the time of the failure abends and data is lost. The Processor Availability Facility reduces the impact of many formerly unrecoverable CP failures.

Systems with two or more CPs move the programs in process from the failing CP to another operational CP.

The processor availability facility is an extension of the hardware logic that supports CPU retry. CPU retry requires that status information about the program in process be available to the CP. When an error occurs in the execution of an instruction, CPU retry uses the status information to rerun the instruction, resuming the instruction flow at the precise point where the error occurred.

The processor availability facility makes the status information required by CPU retry available to the operating system responsible for recovering the program in process. Using the hardware status, the operating system or LPAR determines whether or not Processor Availability Facility recovery can be performed. If so, Processor Availability Facility moves the program in process to another operational CP. This is performed transparently to the application as if a normal interrupt occurred and the program rescheduled on a different CP.

In a PR/SM LPAR mode, the processor availability facility offers additional benefits. A system running in basic mode with a single physical CP cannot use Processor Availability Facility. However, in a PR/SM LPAR partition where a single logical CP shares a pool of multiple physical CPs, a CP failure results in the rescheduling of the partition on another CP from the pool. This provides greater resilience in basic mode.

In addition, systems that do not support multiple CPs but still require high availability can benefit from Processor Availability Facility. For example, a VSE system running in a partition with shared access to multiple physical CPs can have improved protection and system availability by using Processor Availability Facility.

Processor Resource/Systems Manager

Processor Resource/Systems Manager (PR/SM) enables logical partitioning in the processor complex, and supports the VM/XA System Product enhancement for multiple preferred guests.

Logical Partitioning

PR/SM enables the ES/9000 Processor Complex (or a side of a physically partitioned ES/9000 Processor Complex) to be initialized for logically partitioned operation. Models 520, 640, 660, 740, 820, 860, and 900 support as many as seven logical partitions (14 when operating in a physically partitioned configuration). See "Logically Partitioned (LPAR) Operating Mode" on page 1-7 for more information.

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.

PR/SM LPAR Auto-Reconfiguration

This function in PR/SM is an automatic reconfiguration facility 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.

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 make it easier for you to make better capacity planning decisions due to more precise information in the processor resource utilization report.

Retention of Logical Partition Processing Weight Definitions

PR/SM LPAR supports the retention of logical partition processing weight definitions across power-on reset.

This improves system operation by providing a new method for retaining the logical partition processing weights defined on the Logical Partition Control (LPCTL) frame.

Previously, the system associated the processing weight values (displayed in the Weight Value and Weight Capped fields) with a particular I/O configuration data set (IOCDS) slot (for example, A0, A1, A2). If you rewrote an IOCDS slot, the system retained the values only if:

- The number of logical partitions defined in the new IOCDS was equal to the number of logical partitions in the old version of the IOCDS.
- The names of the logical partitions defined in the new IOCDS were the same as the names of the logical partitions in the old version of the IOCDS.

If you did not meet these two conditions when you rewrote an IOCDS slot, the weights reverted to the default values after the system was power-on reset with the new IOCDS.

The system now associates the LPCTL frame processing weight values with the individual logical partition names and retains them as long as the logical partition name exists in any IOCDS.

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 the *ES/9000, ES/3090 Input/Output Configuration Program User Guide*, GC38-0097.

PR/SM LPAR High-Performance Parallel Interface (HIPPI)

PR/SM LPAR supports the assignment of a HIPPI interface to a logical partition providing the same HIPPI capabilities available in basic mode. Support of HIPPI in LPAR mode includes:

- Enhanced large system resource utilization
- Ability to mix traditional computing and numerically intensive computing (NIC) on the same processor
- Ability of one logical partition to access two HIPPIs for Models 660, 820, 860, and 900.

PR/SM LPAR Support of ICRF

PR/SM LPAR mode now supports the use of ICRF by an MVS/ESA guest running under VM/ESA that is operating in a logical partition.

PSW-Key Handling

PSW-key handling provides the Set PSW Key from Address (SPKA) and Insert PSW Key (IPK) instructions.

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.

SCP-Initiated Reset

See "I/O Interface Reset" on page 7-19.

Service Signal

Service signal provides an external interruption that is used by the 9022 Processor Controller to signal information to the control program.

Sorting Instructions

Sorting instructions are used by IBM Data Facility Sort (DFSORT) Release 7 and subsequent releases, running under MVS an VM systems control program. The sorting instructions are used by DFSORT when sorting fixed-length records using the block-set sorting technique.

Storage Background Scrubbing

Storage background scrubbing reduces system or application failure due to accumulation of multiple storage errors.

Storage background scrubbing acts as a constant reading and rewriting of the central and expanded storage arrays not in active use. The system continually, on a cyclic time basis, scrubs the arrays for soft errors, makes correction to

these errors, and records information about hard errors that cannot be corrected.

The recorded information collected during storage background scrubbing enables the processors to detect a failure in central or expanded storage chips and to initiate fault tolerant dynamic memory arrays. For more information, see "Fault-Tolerant Dynamic Memory Arrays" on page 7-17.

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.

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.

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.

Storage-Reconfiguration

See "MVS Dynamic Physical Partitioning" on page 7-21.

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.

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.

IBM 9037 Sysplex Timer

The optional Sysplex Timer, supported by MVS/ESA SP Version 4.1 and subsequent releases, synchronizes the processor time-of-day clocks of the systems in a complex. The Sysplex Timer provides an external time source for attached processors when sharing data or workloads.

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.

Automatic Propagation Delay Adjustment

The automatic propagation delay adjustment function of Sysplex Timer removes the restriction that cables attached to the central processor complexes must be equal in length. This function automatically adjusts the signals to compensate for propagation delay; thereby making the differences in cable length appear nearly transparent to the processing time-of-day (TOD) clock synchronization.

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.

Test Block

Test block provides the Test Block (TB) instruction for testing the usability of a 4KB block of central storage.

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.

Tracing

Tracing provides three aids for problem-program analysis:

- Address-space-number (ASN) tracing
- Branch tracing
- Explicit tracing.

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).

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 ES/9000 model.

Models 520, 640, 660, 740, 820, 860, and 900 use 4KB pages and either 64KB segments or 1MB segments. ESA/390 mode uses only the 1MB segment size. A System/370 guest virtual machine under control of interpretive execution can use both segment sizes.

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.

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.

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.

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.

Vector Facility Feature

The vector facility feature is optional for each central processor in the ES/9000 Processor Complex and is mutually exclusive on the same CP with ICRF. Central processors with the vector facility feature experience significant increased levels of performance for many computer-intensive engineering and scientific applications.

The vector facility feature is an extension of the central processor's instruction and execution elements that allows the CP to perform vector arithmetic and logical operations on vectors of any length (from one element up to 256 sets of operands) with a single instruction. Simple instruction loops handle longer vectors. Vector instructions are precisely interruptible for easy exception handling.

For Models 520, 640, 660, 740, 820, 860, and 900, the section size is 256 and the partial-sum number is 4.

The vector facility feature includes an enhanced instruction set. For more information, see "Vector Facility Feature" on page 3-6 and *IBM Enterprise Systems Architecture/370 and System/370 Vector Operations*.

Virtual-Machine Assist

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.

VM Assists for the CPU Timer

VM assists for the CPU timer permit a central processor to execute directly the Set CPU Timer (SPT) and Store CPU Timer (STPT) instructions for a virtual machine operating under VM/SP.

3033 Extension

The 3033 extension provides the following facilities:

- Dual-address space
- Start-I/O-fast queuing
- Suspend and resume.

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 protection by problem programs.

In ESA/390 mode the tracing facility provides an alternative set of aids.

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.

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.

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.

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 8. Characteristics of Channel Operation

Overview 8-2
Subchannels 8-2
Channel Control 8-3
  Channel Commands 8-3
  ESA/390 Mode 8-3
  Chaining 8-5
    Data Chaining 8-5
    Command Chaining 8-5
  I/O Interruptions 8-5
    Resets 8-6
Channel Implementation 8-6
  Channel Time-Out Functions 8-6
  Device Priority on an I/O Interface 8-7
  Dynamic Reconnection 8-8

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:

- ESA/390 Architectural mode of operation
- Data mode of operation
  - Block multiplexer (all ESCON and parallel channels)
  - Byte multiplexer (parallel channels only).
- 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.

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 four channel paths, which may be configured to an I/O device.

In ESA/390 mode, one UCW is assigned for each subchannel (device).

Models 520, 640, 660, 740, 820, 860, and 900 provide one set of I/O display frames that display the contents of a specified subchannel or UCW.

The data mode of operation is determined by the multiplexer mode (byte or block). This is selected:

- For specific channels, when either the processor complex or a logical partition is initialized.
- For one side of a processor complex operating in a physically partitioned configuration.

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

Channel Control

I/O operations over the interface are controlled by channel commands and mode-dependent I/O instructions, chaining operations, and I/O interruptions.

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.

ESA/390 Mode

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. VM/ESA uses this function to improve guest I/O recovery in its Missing Interrupt Handler.

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.

Chaining

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.

Data Chaining

When the data transfer specified by the current CCW is finished, data chaining causes the operation to continue by fetching a new CCW and using the storage area defined by the new CCW. Execution of the operation at the I/O device is not affected.

Command Chaining

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

I/O Interruptions

I/O interruptions report the completion of I/O operations to the CPs 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.

For a listing of the machine-check interruption codes, see Table 6-1 on page 6-4.

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.

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.

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.

Device Priority on an I/O Interface

Device priority on an I/O interface applies only to ESCON channels attached to a 9034 ES Connection Converter and parallel channels.

Device priority on the I/O interface of a channel depends on the order in which they were attached. If the devices are connected to the 'select out' line, the first device has the highest priority. If the devices are attached to the 'select in' line, the priority sequence is reversed. Devices attached to the 'select out' line have priority over devices attached to the 'select in' line.

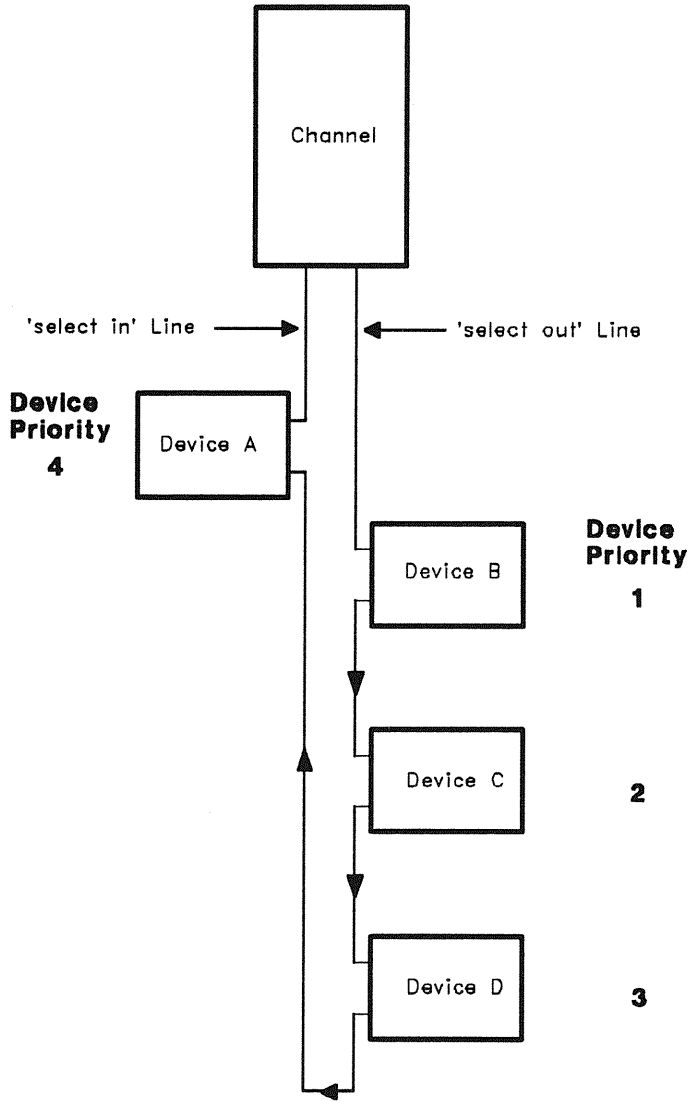

Figure 8-1. Device Priority on Parallel Channels

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 four possible channel paths defined in an IOCP parameter). The selected path is not necessarily the one used initially in the I/O operation.

Chapter 9. Channel Subsystem Performance

Overview 9-2
Performance Characteristics 9-2
  Description 9-2
    All ESCON and Parallel Channels 9-2
    Synchronous and Nonsynchronous I/O Operation 9-3
      Synchronous Operation 9-3
      Nonsynchronous Operation 9-3
    Data Transfer Rate 9-4
  I/O Device Characteristics 9-5
    Channel Subsystem Workload Characteristics 9-5
    Channel Program Characteristics 9-5
    Cache-Hit Rates 9-6
  ESCON Channels 9-7
    Distance 9-7
  Configuring ESCON Channels 9-7
  Elapsed Time 9-7
  Critical Time 9-7
  Deferred Access 9-8
Parallel Channels Only 9-9
  Device Class 9-9
    Device Class 1 9-9
    Device Class 2 9-9
    Device Class 3 9-9
Block Multiplexer Mode of Operation 9-9
  ESCON Channels Using Link-Level and Device-Level Protocols 9-9
  ESCON Channels Attached to a 9034 and Parallel Channels 9-10
    Interlocked 9-10
    Data Streaming 9-10
Byte Multiplexer Mode of Operation (Parallel Channels Only) 9-10
  Byte Multiplexer Mode and Burst Mode 9-10

Overview

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 an ES/9000 Processor Complex. 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*.

Performance Characteristics

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 vs synchronous)
- Data transfer rates
- Distance
- Device characteristics
- Workload characteristics.

Description

This section provides the description of the components of channel performance and the factors that affect channel performance characteristics:

- All ESCON and parallel channels
- ESCON channels (using link-level and device-level protocols)
- ESCON channels (attached to a 9034) and parallel channels
- Parallel channels only.

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.

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.

Synchronous and Nonsynchronous I/O Operation

For detailed information about concepts described in this section, see *Introduction to Nonsynchronous Direct Access Storage Subsystems.*

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

All work involved in ending an operation and advancing to the next operation must be completed before the DASD head reaches the next record (commonly referred to as the inter-record gap). If this does not occur, a rotational positional sensing/sensor (RPS) miss or an overrun is generated and the operation must wait for one DASD revolution before continuing.

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

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

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

Data Transfer Rate

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

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

- Channel data rate
- 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 17 MB/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 9-1 on page 9-5 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 4.5 MB/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 9-1. Channel Data Rate as a Function of Distance for Models 520 – 900

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

I/O Device Characteristics

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

Channel Subsystem Workload Characteristics

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

Channel Program Characteristics

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

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

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

Chained data commands increase response time due to an additional interlocked exchange between the channel and control unit. See "Distance" on page 9-7 for more information.

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

Cache-Hit Rates

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

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

ESCON Channels

This section describes the performance characteristics of ESCON channels.

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.

Configuring ESCON 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 parallel and ESCON channels include:

- Elapsed time
- Critical time
- Deferred access.

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.

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 "Deferred Access" and "Device Class" on page 9-9.

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'. The probability of deferred accesses occurring during data chaining can be reduced by following the programming recommendations stated in "Data Chaining" on page 10-4.

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 "I/O Interruptions" on page 8-5 for additional information.

Parallel Channels Only

This section describes the characteristics of devices that operate in byte multiplexer mode.

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.

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.

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.

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).

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 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 17 MB/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.

ESCON Channels Attached to a 9034 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.

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.

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.

Byte Multiplexer Mode of Operation (Parallel Channels Only)

Only parallel channels can operate in byte multiplexer mode of operation.

Byte multiplexer 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.

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. For example, an IBM 3211 Printer can specify either burst mode or byte multiplexer mode with 1-byte or 6-byte transfers.

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 (for example, by a factor of six for a 3211 Printer operating in byte multiplexer mode with 6-byte transfers). 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 (see Chapter 10, "Guidelines for Channel Subsystem Configuration" on page 10-1) are followed for byte multiplexer channels of an ES/3090 Processor Complex, 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.

Rules governing the placement of devices that can operate in byte multiplexer or burst mode on one or more byte multiplexer channels are discussed in Chapter 10, "Guidelines for Channel Subsystem Configuration" on page 10-1.

Chapter 10. Guidelines for Channel Subsystem Configuration

Overview 10-2
ESCON Channel Subsystem Configurations 10-2
  Recommendations for ESCON Channel Subsystem Configurations . . . 10-2
Guidelines for Parallel Channel Subsystem Configurations 10-3
Device Sequence on Block Multiplexer Channels 10-3
  Device Intermix 10-3
  Data Chaining 10-4
  CCW Prefetching for ESCON Channels Attached to a 9034 10-5
Byte Multiplexer Channel Analysis 10-5
  Balancing Device Load across Channels 10-5
  Sequencing Devices Within a Channel 10-6
General Device Considerations 10-6
  All ESCON and Parallel Channels 10-6
    I/O Devices Using Caches 10-6
    I/O Devices Using Buffers 10-6
  ESCON Channels Attached to a 9034 and Parallel Channels 10-7

Overview

This chapter provides recommendations for the installation and operation of an ESCON channel subsystem and provides guidelines for I/O devices using caches and buffers to maximize data transfer rates in a channel subsystem.

Information in this chapter that applies to a parallel environment is intended to assist in the physical placement of the control units connected to the I/O interface of a channel subsystem. The Input/Output Configuration Program establishes the physical and logical addressing relationships of the I/O devices in the ES/9000 Processor Complex configuration.

ESCON Channel Subsystem Configurations

This section provides publication references for those who need to understand basic ESCON concepts or to plan for the installation and operation of an ESCON channel subsystem.

The ESCON environment supports Models 520, 640, 660, 740, 820, 860, and 900. In a channnel subsystem, all channels have equal priority.

ESCON combines ESCON Architecture, fiber optic technology, dynamic connectivity, and interconnectivity with other networks to provide new function and capabilities including enhanced connectivity, improved systems availability, better security, and better data center management.

ESCON LED multimode fiber optic technology provides a direct channel attachment range of up to 3 kilometers (1.864 miles); and control unit attachment range of up to 9 km (5.592 miles) from a channel through two optional ESCON Directors.

ESCON XDF single-mode fiber optic technology provides a direct channel attachment range of up to 60 kilometers (37.3 miles); and control unit attachment range of up 43 kilometers (26.7 miles) from a channel through two optional ESCON directors. These distances are significantly greater than the currently supported 400 feet. See "Fiber Optic Cabling, Transmission, and Reception" on page 1-9 for more information.

The ESCON environment is also designed to provide a platform for growth into a new generation of powerful and flexible products.

Recommendations for ESCON Channel Subsystem Configurations

The following publication references are for those who require additional information about ESCON channels that use link-level and device-level protocols or that attach to a 9034:

- To gain an understanding of basic ESCON concepts, see *Introducing Enterprise Systems Connection*.

  This publication provides detailed information about the benefits and components of an ESCON environment, and includes a comprehensive list of the publications required for planning an ESCON channel subsystem.

- Detailed guidelines for interconnecting ESCON channels with existing control units and parallel interfaces are provided in *Enterprise Systems Connection: Planning for Migration*.

Note: IBM 3490 Magnetic Tape Subsystem and 3990 Storage Control support both parallel and ESCON interfaces (the same control unit can be accessed in either environment by different systems) to provide the capability of having coexisting parallel and ESCON environments while migrating in stages toward a full ESCON environment.

- For guidelines that apply to the planning and operation of a specific channel subsystem, contact your IBM representative.

Guidelines for Parallel Channel Subsystem Configurations

This section applies to configurations for a parallel environment. It applies to both ESCON channels that attach to a 9034 and to parallel channels. In a channel subsystem, all channels have equal priority.

- Device sequence on block multiplexer channels
  - Device intermix
  - Data chaining
- Byte multiplexer channel analysis
  - Balancing device load across channels
  - Sequencing devices within a channel.

For positional placement of control units on a particular I/O interface, consider the service priority that is obtained from attaching devices to the 'select out' line or 'select in' line. Devices attached to the 'select out' line have priority over devices attached to the 'select in' line. (See "Device Priority on an I/O Interface" on page 8-7 and Figure 8-1 on page 8-7 for more information.)

Device Sequence on Block Multiplexer Channels

To establish the sequence of devices on any channel, consider the following:

Device Class. Class-2 devices should have higher priority for service than class-3 devices.

Total System Priorities. Attach the devices (control units) that are most important to system performance first (in the highest priority position).

Channel Service Time. Within a system priority, attach the devices with the shortest channel service time first.

Device Intermix

Although different devices (such as an IBM 3380 Direct Access Storage and an IBM 3800 Printing Subsystem) can share channels, the decision to intermix devices depends on the load imposed on the channel by each device.

Usually, different types of devices should be attached to different channels because of their design and use.

Reasons for separate channel requirements include:

- Type of channel
- Paging and response time
- Random versus sequential records (such as disk versus tape)

- Chained records (such as tape or disk)
- Logging
- Different critical times
- Number of devices
- Channel utilization.

For example, to improve system performance, the following types of devices should be on separate block multiplexer channels:

- Direct-access storage devices (such as the IBM 3380 Direct Access Storage and 3390 Direct Access Storage)
- Tape devices (such as the IBM 3480 Magnetic Tape Subsystem and 3490 Magnetic Tape Subsystem).

Data Chaining

Data chaining can put a very heavy load on the channel. Factors that increase channel loading include:

- Data address alignment
- Low CCW byte counts
- High device-data rates
- Other channel activity within the channel subsystem
- Additional central storage accesses for CCWs during device data transfer.

Because of these factors, data-deferred accesses and chaining checks can occur when data is chained during data transfer sequences. Two programming recommendations help minimize the probability of these two problems:

1. If the device or control unit responds to the suppress-out tag and suppresses data transfers, make the minimum CCW byte count greater than the number of bytes or byte requests in transit on the I/O interface cable. This number varies; depending on the cable length, the data transfer rate, and the protocol.

   With the interlocked protocol, no minimum byte count exists.

   With the data-streaming protocol, specifying a minimum byte count of 16 covers all combinations of local cable lengths and data transfer rates.

   Note: If a channel extender (using the IBM 3044 Fiber-Optic Channel Extender Link) is used, the minimum byte count is four times the cable length (measured in thousands of feet) times the data rate (measured in MB per second).

2. If the device or control unit does not suppress data transfer, specify byte counts to be no less than 54 times the device instantaneous data rate (a maximum data rate of 4.5MB per second) and put data addresses on doubleword or (preferably) quadword boundaries.

No minimum byte count restrictions exist for data chaining within gaps of a direct-access storage device record (that is, gaps between count, key, and data fields).

CCW Prefetching for ESCON Channels Attached to a 9034

Because ESCON channels allow control units to be placed further away from the channel than is allowed for parallel channels, the number of command overruns tends to increase with the longer response times caused by propagation delay. To eliminate this increase in command overruns, command chain prefetching has been implemented for ESCON channels using 9034s. This prefetching is used for all devices that have a device type of 33xx specified in the IOCP input deck.

Command-chain prefetching allows the next CCW in a channel program to be fetched before the data for the current CCW is moved (this applies only if the current CCW has the command chain flag set to 1 and the data chain flag set to 0). If the current CCW is a read, and the next CCW in the channel program is contained in the data read by the current CCW, prefetching cannot be completed (such self-modifying channel programs cannot take advantage of command chain prefetching).

Note: Any channel programs that are dynamically modified by the program during execution may behave differently on ESCON channels attached to a 9034 than stand-alone parallel channels due to a command chain prefetching.

Byte Multiplexer Channel Analysis

Many class-2 and class-3 devices can operate efficiently in burst mode on *block multiplexer* channels. Before attempting to configure any such devices on byte multiplexer channels, place as many of these devices as possible on block multiplexer channels. This reduces the contention of the devices remaining on byte multiplexer channels. Next, ensure that devices with a variable burst size capability are set for large burst size. Finally, balance the device load across the available byte multiplexer channels and sequence the devices within these channels (as described in "Balancing Device Load across Channels").

Balancing Device Load across Channels

To balance the device load across the channels, consider the following factors, in decreasing order of priority:

Burst Size Try to configure I/O devices that are small-burst devices (1, 2, 3, or 6 bytes per transfer) to one or more byte multiplexer channels, and all large-burst devices (16 or 32 bytes per transfer) to other byte multiplexer channels. Class-1 devices should be distributed among the channels that attach small-burst devices.

Critical Time When possible, configure each byte multiplexer channel with both short and long critical time devices.

For example, if two I/O devices have short critical times, put each device on a separate byte multiplexer channel. If you attach multiple devices (such as terminals) with different critical times to the same control unit, configure the control units according to the device with the shortest critical time.

Data Rate Balance data rate across multiple byte multiplexer channels where no conflict exists with class-1 and critical-time considerations.

Availability Follow standard guidelines for availability, which provide for alternate paths.

Sequencing Devices Within a Channel

To sequence the devices (control units that have the most critical I/O devices) within a channel, consider the following factors.

Device Class Attach class-1 devices first, class-2 devices next, and class-3 devices last. For information on class specifications, see "Device Class" on page 9-9.

Increasing Critical Time
Within a class, the device with the shortest critical time has the highest priority. For devices having the same critical time, attach the device with the smallest burst size first. A communication controller with several lines having short critical times should be ahead of a communication controller with one line having a short critical time. Devices operating in burst mode and connected to a byte multiplexer channel should be given lowest attachment priority.

General Device Considerations

In general, ensure that the feature information for each device attached (or to be attached) to the channel subsystem is carefully considered to determine the impact that the device can have on the configuration when attached.

All ESCON and Parallel Channels

This section provides basic guidelines for using I/O devices with caches and buffers to maximize data transfer rates in a channel subsystem.

I/O Devices Using Caches

When the control unit reads or writes to a device, it matches the data transfer rate of the device. Therefore, the control unit will not use a faster transfer rate for the I/O operation.

Having a cache in the control unit improves the data transfers rate and response time. A cache stores the most recently used data that was read or written to the attached device. This eliminates the need for the control unit to transfer information to or from the device during an I/O operation.

For example, if the data to be read is stored in the cache, the control unit can send the data at the higher transfer rate of the control unit instead of the device. In general, the larger the cache, the more likely the needed information is stored in the cache and therefore data transfer will take place at the control unit's maximum transfer rate.

I/O Devices Using Buffers

A buffer is similar to a cache, but does not improve response time directly. A cache normally has a certain amount of data stored in it at any given time, even when there are no I/O operations taking place. This data, if needed, will not be read from the device during an I/O operation. A buffer stores data only during the data transfer portion of an I/O operation (data needed has to be read from the device).

A buffer allows information to be read from the device without the need for the control unit to be connected to a channel requesting that data. The control unit

can disconnect from the requesting channel while the data is being read from the device (at device speed), then reconnect to the channel and send that data at the maximum data transfer rate of the control unit. Though this does not reduce the response time of the I/O operation. It allows the channel to be free for the other I/O operations while the data is being transferred from the device to the control unit, increasing throughput.

ESCON Channels Attached to a 9034 and Parallel Channels

For positional placement of control units on a particular I/O interface, consider the service priority that is obtained from 'select out' line or 'select in' line. Devices attached to the 'select out' line have priority over devices attached to the 'select in' line. (See "Device Priority on an I/O Interface" on page 8-7 and Figure 8-1 on page 8-7 for more information.)

Appendix A. Architectural Deviations

The following information describes architectural deviations from the *IBM Enterprise Systems Architecture/390 Principles of Operation*.

Concurrent Indication of PER Events with Operand-Access Exception

Storage alteration PER events can be indicated by Models 520, 640, 660, 740, 820, 860, and 900 for execution of the instructions Edit, Edit and Mark, and Translate, even though an operand-access exception is encountered that nullifies or suppresses instruction execution.

Protection Violation Instead of Delayed Access Exception

The ESA/390 principles of operation permit considerable extent of unpredictability when a valid and attached dynamic address translation (DAT) table entry is changed and the entry is used for translation before the translation lookaside buffer (TLB) is cleared of copies of that entry. The definition permits changes to all of those result fields that are not protected. Changes can occur, for example, to the condition code, to operands due to be changed in registers, and to those portions of the operands due to be changed in storage for which no access exception exists.

Models 520, 640, 660, 740, 820, 860, and 900 deviate from the architecture in that:

- Protection exceptions are ignored if they are caused by key-controlled protection, page protection, and segment protection that occur after the initial pretest.
- Storing for a store-type reference takes place.
- No interruption occurs.

Appendix B. Summary of the Configurations

Table B-1 (Page 1 of 3). Summary of the Configurations

Feature	Model 520	Model 640	Model 660	Model 740	Model 820	Model 860	Model 900
Processor Controller							
9022 Model 1A	Req	Req	NA	Req	NA	NA	NA
9022 Model 1A with Specify Code 5050	NA	NA	Req	NA	Req	Req	Req
Displays *(IBM 3206 Display Station 100 or equivalent)*							
First display	Req	Req	Req	Req	Req	Req	Req
Second display	Req	Req	Req	Req	Req	Req	Req
Third display	Opt	Opt	Req	Opt	Req	Req	Req
Fourth display	NA	NA	Opt	NA	Opt	Opt	Opt
Fifth display	NA	NA	Opt	NA	Opt	Opt	Opt
Sixth display	NA	NA	Opt	NA	Opt	Opt	Opt
Printers *(IBM 3287 Printer Model 1 or 2, IBM 4224 Printer Model 201 or 202, or equivalent)*							
First printer	Opt	Opt	Opt	Opt	Opt	Opt	Opt
Second printer	NA	NA	Opt	NA	Opt	Opt	Opt
Modems *(IBM 5853 Modem Model 1, or 3864 Modem Model 2 with Automatic Calling Unit #5801, or equivalent)*							
First Modem	Req	Req	Req	Req	Req	Req	Req
Second Modem	NA	NA	Req	NA	Req	Req	Req
Power and Coolant Distribution Units *(IBM 9027 PCDU Model 1)*							
First 9027	Req	Req	Req	Req	Req	Req	Req
Second 9027	NA	NA	Req	NA	Req	Req	Req
I/O Power Sequencing Unit *(IBM 9023 ES/9000 I/O Power Sequencing Unit Model 1)*							
First 9023	Opt	Opt	Opt	Opt	Opt	Opt	Opt
Second 9023	Opt	Opt	Opt	Opt	Opt	Opt	Opt
Central Processor							
CP0	NA	NA	NA	Std	NA	Std	Std
CP1	Std	Std	Std	Std	Std	Std	Std
CP2	NA	Std	NA	Std	Std	Std	Std
CP3	NA	NA	Std	NA	Std	Std	Std
CP4	NA	NA	NA	NA	Std	Std	Std

Table B-1 (Page 2 of 3). Summary of the Configurations

Feature	Model 520	Model 640	Model 660	Model 740	Model 820	Model 860	Model 900
CP5	NA	NA	NA	NA	NA	NA	Std
System Control Element							
SCE0	Std	Std	Std	Std	Std	Std	Std
SCE1	NA	NA	Std	NA	Std	Std	Std
Channel Subsystem							
ICE0	Std	Std	Std	Std	Std	Std	Std
ICE1	NA	NA	Std	NA	Std	Std	Std
Vector Facility							
VE0	NA	NA	NA	Opt	NA	Opt	Opt
VE1	Opt	Opt	Opt	Opt	Opt	Opt	Opt
VE2	NA	Opt	NA	Opt	Opt	Opt	Opt
VE3	NA	NA	Opt	NA	Opt	Opt	Opt
VE4	NA	NA	NA	NA	Opt	Opt	Opt
VE5	NA	NA	NA	NA	NA	NA	Opt
ICRF *(Note 1)*							
CT0	Opt	Opt	Opt	Opt	Opt	Opt	Opt
CT1	NA	NA	Opt	NA	Opt	Opt	Opt
Central Storage *(Note 2)*							
128MB	Std	Std	NA	NA	NA	NA	NA
256MB	Opt	Opt	Std	Std	Std	NA	NA
384MB	NA	NA	Opt *(Note 3)*	NA	Opt *(Note 3)*	Std	NA
512MB	NA	Opt	Opt *(Note 3)*	Opt	Opt *(Note 3)*	Opt *(Note 4)*	Std
640MB	NA	NA	NA	NA	Opt *(Note 3)*	Opt *(Note 4)*	NA
768MB	NA	NA	NA	NA	Opt *(Note 3)*	Opt *(Note 4)*	Opt *(Note 5)*
1024MB	NA	NA	NA	NA	Opt *(Note 3)*	Opt *(Note 4)*	Opt *(Note 5)*
Expanded Storage *(Note 6)*							
Maximum Size	2GB	4GB	4GB *(Note 7)*	4GB	8GB *(Note 7)*	8GB *(Note 7)*	8GB *(Note 7)*
Initial Offering	0MB	0MB	0MB	0MB	0MB	0MB	0MB
PR/SM	Std	Std	Std	Std	Std	Std	Std
IBM 9037 Sysplex Timer Attachment *(Note 8)*							
First Timer	Opt	Opt	Opt	Opt	Opt	Opt	Opt
Second Timer	NA	NA	Opt	NA	Opt	Opt	Opt
Maximum Number of Channels *(ESCON and Parallel)*							
Parallel	48	48	96	48	96	96	96
ESCON	64	128	128	128	256	256	256

Table B-1 (Page 3 of 3). Summary of the Configurations

Feature	Model 520	Model 640	Model 660	Model 740	Model 820	Model 860	Model 900
Total maximum	64	128	128	128	256	256	256
Total minimum	32	64	64	64	64	64	64
Parallel Channels *(Note 9)* **For A-side on Models 520, 640, and 740 and for each side on Models 660, 820, 860, and 900.**							
0	Opt	Opt	Opt	Opt	Opt	Opt	Opt
16	Opt	Opt	Opt	Opt	Opt	Opt	Opt
32	Opt	Opt	Opt	Opt	Opt	Opt	Opt
48	Opt	Opt	Opt	Opt	Opt	Opt	Opt
ESCON Channels *(Note 9)* **For A-side on Models 520, 640, and 740 and for each side on Models 660, 820, 860, and 900.**							
16	Opt	Std	Opt	Std	Std	Std	Std
32	Opt	Opt	Opt	Opt	Opt	Opt	Opt
48	Opt	Opt	Opt	Opt	Opt	Opt	Opt
64	Opt	Opt	Opt	Opt	Opt	Opt	Opt
80	NA	Opt	NA	Opt	Opt	Opt	Opt
96	NA	Opt	NA	Opt	Opt	Opt	Opt
112	NA	Opt	NA	Opt	Opt	Opt	Opt
128	NA	Opt	NA	Opt	Opt	Opt	Opt

Legend:

CP	Central processor
ICE	Interconnect communication element
PCDU	Power and coolant distribution unit
SCE	System control element
VE	Vector
NA	Not applicable

Notes:

1. Mutually exclusive with the vector facility on a central processor.
2. These figures represent the total system configuration.
3. Increments of 128MB and 256MB to 512MB on each side.
4. Increments of 256MB and 512MB for A-side; 128MB, 256MB and 512MB for B-side.
5. Increments of 256MB to 512MB on each side
6. Expanded storage increments range from 256MB to 1GB depending on model and current storage size. Expanded storage configurations for MP models can be asymmetrical. See 9021 Sales Manual Pages for more details.
7. Expanded storage size is for the total MP configuration for Models 660, 820, 860, and 900.
8. Must be symmetrical for Models 660, 820, 860, and 900.
9. Channels can be in combinations of available block-multiplexer channels or ESCON channels.

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*.

access method. (1) A technique to obtain the use of data, storage, or the use of an input/output channel to transfer data; for example, random access method, sequential access method. (T) (2) A technique for moving data between main storage and input/output devices. (3) Positive acknowledgment.

acknowledgment. (1) The transmission, by a receiver, of acknowledge characters as an affirmative response to a sender. (T) (2) An indication that an item sent was received.

action queue. A collection of pending maintenance actions.

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 state indicator. An indicator, viewed from the fixed-format frames, that is used with error indicators (triggers) for isolating failures to a field-replaceable unit (FRU).

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.

adapter. (1) Hardware that provides some transitional functions between two or more devices. (2) A mechanism for attaching parts, for example, parts having different diameters. (3) In an ESCON environment, link hardware used to join different optical fiber connector types.

address. (1) A value that identifies a register, a particular part of storage, a data source, or a data sink. The value is represented by one or more characters. (T) (2) To refer to a device or an item of data by its

address. (I) (A) (3) The location in the storage of a computer where data is stored. (4) In data communication, the unique code assigned to each device or workstation connected to a network. (5) The identifier of a location, source, or destination.

address generation. The method whereby the I-element address adder provides the address of the next instruction.

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*.

allocate. To assign a resource, such as a disk or a diskette file, to perform a task. Contrast with *deallocate*.

analysis routine. A routine that analyzes error records, provided by an error handler, to isolate failures to one or more field-replaceable units (FRUs).

APAR. Authorized program analysis report.

application. (1) The use to which an information processing system is put, for example, a payroll application, an airline reservation application, a network application. (2) A collection of software components used to perform specific types of work on a computer.

application program. (1) A program that is specific to the solution of an application problem. (T) (2) A program written for or by a user that applies to the user's work, such as a program that does inventory control or payroll. (3) A program used to connect and communicate with stations in a network, enabling users to perform application-oriented activities.

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*.

attention identifier (AID). A character in a data stream indicating that the user has pressed a key, such as the Enter key, that requests an action by the system.

authorized program analysis report (APAR). A request for correction of a problem caused by a defect in a current, unaltered release of a program.

automated logic diagram (ALD). A computer-generated diagram that represents functioning circuitry in terms of logic blocks, interconnecting conductor networks, and input/output terminals.

automatic fault isolation manager. Licensed Internal Code that controls the selection and running of analysis routines used to isolate failing field-replaceable units (FRUs).

auto-vectoring. In VS FORTRAN Version 2, an automatic operation that compiles the source code to contain object code having vector instructions.

auxiliary storage. All addressable storage, other than main storage, that can be accessed by means of an input/output channel; for example, storage on magnetic tape or direct access devices. Contrast with *main storage*.

B

backup copy. A copy, usually of information or data, that is kept if the original is changed or destroyed.

balun. A transformer used to connect balanced cables, such as twisted-pair cables, to unbalanced cables, such as coaxial cables, by matching the electrical characteristics of the cables.

basic mode. A central processor mode that does not use logical partitioning. Contrast with *logically partitioned (LPAR) mode*.

basic telecommunications access method (BTAM). An access method that permits read/write communication with remote devices.

BCE. Buffer control element.

bpi. Bits per inch.

Bpi. Bytes per inch.

bps. Bits per second.

Bps. Bytes per second.

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*.

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.

bus-out data. Two digits of information transmitted from a parallel channel to a control unit. 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*.

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

CAC. Common adapter code.

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.

CAI. Channel available interruption.

calculated link loss. In an ESCON environment, the total optical attenuation (loss) calculated for a specific link, the value of which cannot be more than the maximum loss allowed for that link. See also *maximum allowable link loss*.

CAW. Channel address word.

CCA. (1) Channel communication area. (2) Common communication adapter.

CCW. Channel command word.

CDC. Channel data check.

CE. (1) Channel end.

central processor. A processor that contains the sequencing and processing facilities for instruction execution, interruption action, timing functions, initial program loading, and other machine-related functions.

central storage. Storage that is an integral part of the processor and includes both main storage and the hardware system area.

CERF. Customer engineering reporting facility.

channel. The system element that controls one channel path, whose mode of operation depends on the type of hardware to which it is attached.

channel 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 communication area (CCA). An area used to transmit control information between the channels and the IOP.

channel control check. A category of I/O errors affecting channel controls and sensed by the channel to which a device is attached. See also *channel data check*.

channel data check. A category of I/O errors, indicating a machine error in transferring data to or from storage and sensed by the channel to which a device is attached. See also *channel control check*.

channel 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.

check stop. The status when an error makes it undesirable or not possible to continue the operation in progress.

CHN. Channel.

CHP. Channel path.

CHPID. Channel path identifier.

CICS. Customer Information Control System.

cm. Centimeter.

CM. Channel monitor.

CME. Channel monitor expanded.

CMS. Conversational monitor system.

CNC. Mnemonic for an ESCON channel attached to an ESCON-capable device.

coaxial cable. A cable consisting of one conductor, usually a small copper tube or wire, within and insulated from another conductor of larger diameter, usually copper tubing or copper braid.

COB. Card on board.

collision. (1) An unwanted condition that results from concurrent transmissions on a channel. (T) For example, when a frame from a transmitting adapter encounters any other signal in its path (frame, noise, or another type of signal), the adapter stops transmitting and a collision occurs.

command. (1) A character string from a source external to a system that represents a request for system action. (2) A request from a terminal for performance of an operation or execution of a program. (3) A value sent on an I/O interface from a channel to a control unit that specifies the operation to be performed.

command chaining. The fetching of a new channel command word (CCW) immediately following the completion of the previous CCW.

communication controller node. A subarea node that does not contain a system services control point (SSCP).

component. (1) Hardware or software that is part of a functional unit. (2) A functional part of an operating system; for example, the scheduler or supervisor.

computing system RPQ. A customer request for a price quotation on changes or additions to the functional capabilities of a computing system, hardware product, or device. The RPQ can be used with programming RPQs to solve unique data processing problems. See also *programming RPQ (PRPQ)*.

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*.

configuration matrix. In an ESCON environment, an array of connectivity attributes that appear as rows and columns on a display device and can be used to determine or change active and saved configurations.

connectivity attribute. In an ESCON Director, the characteristic that determines a particular element of a port's status.

connectivity capability. (1) The capability that allows attachment of a device to a system without requiring physical reconfiguration of the device or its interconnections. (2) In an ESCON Director, the capability that allows logical manipulation of link connections to provide physical device attachment. See also *configuration matrix, connectivity control,* and *dynamic connection*.

connectivity control. In an ESCON Director, the method used to change a port's connectivity attributes, thereby determining the communication capability of the link attached to that port.

connector. (1) A means of establishing electrical flow. (A) See *optical fiber connector*.

consensus processing. In ESCON Manager, the function that informs other hosts of connectivity changes, thereby allowing general agreement for changes to affected resources.

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*.

controller. A unit that controls input/output operations for one or more devices.

control operation. An action that affects the recording, processing, transmission, or interpretation of data; for example, starting or stopping a process, carriage return, font change, rewind, and end of transmission. (I) (A)

control panel assembly. The field-replaceable unit (FRU) that includes both the System Power panel and the Service panel.

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 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.

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.

conversion. (1) In programming languages, the transformation between values that represent the same data item but belong to different data types. Information can be lost through conversion because accuracy of data representation varies among different data types. (I) (2) The process of changing from one method of data processing to another or from one data processing system to another. (3) The process of changing from one form of representation to another; for example, to change from decimal representation to binary representation.

core. (1) In an optical cable, the central region of an optical fiber through which light is transmitted. (E) (2) In an optical cable, the central region of an optical fiber that has an index of refraction greater than the surrounding cladding material. (E) See also *cladding* and *optical fiber*.

corrective maintenance. Maintenance performed specifically to overcome existing faults. (T) Contrast with *preventive maintenance*.

critical resource. A resource required for system operation. See also *system resource*.

CS. (1) Central storage. (2) Channel set. (3) Control storage.

CSE. Control storage element.

CSS. Channel subsystem.

CSW. Channel status word.

CTC. (1) Channel-to-channel. (2) Mnemonic for an ESCON channel attached to another ESCON channel.

CU. Control unit.

cursor. (1) A movable, visible mark used to indicate the position at which the next operation will occur on a display screen. (A) (2) A visual cue that shows the user where keyboard input will appear on the screen.

Customer Information Control System (CICS). An IBM licensed program that enables transactions entered at remote terminals to be processed concurrently by user-written application programs. It includes facilities for building, using, and maintaining data bases.

CVC. Mnemonic for an ESCON channel attached to a 9034.

D

DASD. Direct access storage device.

DASD subsystem. A storage control and its attached direct access storage devices.

data bus. (1) A bus used to communicate data internally and externally to and from a processing unit, storage, and peripheral devices. (A) See *bus*.

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 station. See *station*.

data stream. (1) All data transmitted through a data channel in a single read or write operation. (2) A continuous stream of data elements being transmitted, or intended for transmission, in character or binary-digit form, using a defined format.

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 terminal equipment (DTE). (1) That part of a data station that serves as a data source, data receiver, or both. (I) (A) (2) Equipment that sends or receives data, or both.

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*.

data transfer phase. The phase of a data call during which data signals can be transferred between data terminal equipment (DTE) connected by the network.

deconfigure. To remove a system resource from the currently active configuration, usually through the system control program (SCP) or through the Configuration (CONFIG) frame on the system console.

dedicated connection. In an ESCON Director, a connection between two ports that is not affected by information contained in the transmission frames. This connection, which restricts those ports from communicating with any other port, can be established or removed only as a result of actions performed by a host control program or at the ESCD console. Contrast with *dynamic connection*.

Note: The two links having a dedicated connection appear as one continuous link.

dedication. Pertaining to the assignment of a system resource; for example, an I/O device, a program, or a whole system, to one application or purpose.

default. Pertaining to an attribute, value, or option that is assumed when none is explicitly specified. (I)

degraded. Pertaining to a mode of operation in which the system operates with some resources not available.

device. A mechanical, electrical, or electronic contrivance with a specific purpose.

device address. 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*.

device identifier. In a channel subsystem, an address, not apparent to the program, that is used to communicate with I/O devices. See also *channel path identifier, device number, input/output address,* and *subchannel number*.

Note: In System/370 mode, the device identifier is called a device address and consists of an 8-bit value.

device-level addressing. In an ESCON I/O interface, one of two levels of addressing, this level pertaining to an I/O device and identifying that device to the channel or control unit once the control unit has been determined through link-level addressing. Contrast with *link-level addressing*.

device number. In a channel subsystem, four hexadecimal digits that uniquely identify an I/O device.

diagnostics. (1) The process of investigating the cause or the nature of a condition or a problem in a product or system. (2) Modules or tests used by computer users and service personnel to diagnose hardware problems.

direct access storage. A storage device that provides direct access to data. (I) (A) See also *random access memory*.

direct access storage device (DASD). A device in which access time is effectively independent of the location of the data.

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*.

directory lookaside table (DLAT). A table in the buffer control element used during dynamic address translation.

display device. (1) An output unit that gives a visual representation of data. Usually the data are displayed temporarily; however, arrangements may be made for making a permanent record. (I) (A) (2) In computer graphics, a device capable of presenting display elements on a display surface; for example, a cathode ray tube, plotter, microfilm viewer, or printer. (3) A device that presents information on a screen. See also *display image* and *display screen*.

display/printer adapter (DPA). An integrated control unit contained on each side of the processor controller that provides attachment capability for displays and printers.

distribution panel. (1) In an ESCON environment, a panel that provides a central location for the attachment of trunk and jumper cables and can be mounted in a rack or wiring closet, or on a wall. (2) In the IBM Token-Ring Network, a wiring board that has a patch panel function and mounts in a rack.

downstream. (1) In the direction of data flow or toward the destination of transmission. (2) From the processor toward an attached unit or end user. Contrast with *upstream*.

downstream load (DSL). The capability of a distributed function terminal to receive its control program from the controller to which it is attached. A disk containing the terminal's control program is loaded into the control unit.

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 connection. In an ESCON Director, a connection between two ports, established or removed by the ESCD and that, when active, appears as one continuous link. The duration of the connection depends on the protocol defined for the frames transmitted through the ports and on the state of the ports. Contrast with *dedicated connection*.

dynamic connectivity. In an ESCON Director, the capability that allows connections to be established and removed at any time.

dynamic reconfiguration. Pertaining to a processor reconfiguration between a single-image (SI) configuration and a physically partitioned (PP) configuration when the system control program is active. Contrast with *static reconfiguration*.

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).

E

EE. Execution element.

emulate. To imitate one system with another, primarily by hardware, so that the imitating system accepts the same data, executes the same programs, and achieves the same results as the imitated system. (A) Contrast with *simulate*.

emulation. (1) The use of programming techniques and special machine features to permit a computing system to execute programs written for another system. (2) Imitation; for example, imitation of a computer or device. Contrast with *simulation*.

end of operation (EOP). A signal controlled by the Licensed Internal Code indicating that an instruction sequence has ended (the end of instruction execution).

error. A discrepancy between a computed, observed, or measured value or condition and the true, specified, or theoretically correct value or condition. (I) (A) Contrast with *failure* and *fault*.

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.

error handler. A component that responds to unsolicited interruptions; it analyzes error conditions and takes recovery steps.

error log. A data set or file in a product or system where error information is stored for later access.

error message. An indication that an error has been detected. (A) See also *information message* and *warning message*.

ESA/390. Enterprise Systems Architecture/390.

ESC. Expanded storage controller.

ESCD. Enterprise Systems Connection (ESCON) Director.

ESCD console. The ESCON Director input/output device used to perform operator and service tasks at the ESCD.

ESCD console adapter. Hardware in the ESCON Director console that provides the attachment capability between the ESCD and the ESCD console.

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.

ESE. Expanded storage element.

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.

exchange. To remove an item and put another in its place; for example, to remove a field-replaceable unit (FRU) and install another of the same type.

execution element. An element in a central processor that performs all floating-point, fixed-point multiply, fixed-point divide, and convert operations.

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.

failure. An uncorrected hardware error. Contrast with *error* and *fault.*

Note: Failures are either recoverable or not recoverable by the software or the operator. The operator is always notified when failures occur. Usually, system recovery occurs through a hardware reconfiguration. If this is not possible, recovery requires a repair of the failed hardware.

feature. A part of an IBM product that may be ordered separately by the customer. A feature is designated as either special or specify and may be designated also as diskette-only.

feature code. A code used by IBM to process hardware and software orders.

fiber. See *optical fiber.*

fiber buffer. Material used to protect an optical fiber from physical damage, thereby providing mechanical isolation or protection, or both.

Note: Cable fabrication techniques vary. Some result in tight contact between fiber and protective buffering (tight buffer), while others result in a loose fit (loose buffer), permitting the fiber to slide in the buffer tube. Multiple buffer layers can be used for added fiber protection.

fiber optic cable. See *optical cable.*

fiber optics. The branch of optical technology concerned with the transmission of radiant power through fibers made of transparent materials such as glass, fused silica, and plastic. (E)

Note: Telecommunication applications of fiber optics use optical fibers. Either a single discrete fiber or a nonspatially aligned fiber bundle can be used for each information channel. Such fibers are often called *optical fibers* to differentiate them from fibers used in noncommunication applications.

field macro diagram (FMD). Documentation used by service representatives to analyze card-on-board logic.

field-replaceable unit (FRU). An assembly that is replaced in its entirety when any one of its components fails. Sometimes, a field-replaceable unit may contain other field-replaceable units; for example, a brush and a brush block that can be exchanged individually or as a single unit.

filter. A device or program that separates data, signals, or material in accordance with specified criteria. (A)

flag. (1) A variable that indicates a certain condition holds. (T) (2) Any of various types of indicators used for identification; for example, a wordmark. (A) (3) A character that signals the occurrence of some condition, such as the end of a word. (A)

floating-point. (1) Pertaining to the representation of a quantity as a decimal number and using a symbol (exponent) to show the position for the decimal point, for example, 199.9 represented by $1.999-2$. (2) Pertaining to a machine feature for such arithmetic.

FP. Floating point.

frame. (1) A housing for machine elements. (2) The hardware support structure, covers, and all electrical parts mounted therein that are packaged as one entity for shipping. (3) A formatted display. See *display frame* and *transmission frame.*

frame check sequence (FCS). (1) A system of error checking performed at both the sending and receiving station after a block check character has been accumulated. (2) A numeric value derived from the bits in a message that is used to check for any bit errors in transmission. (3) A redundancy check in which the check key is generated by a cyclic algorithm.

FRU. Field-replaceable unit.

G

G. Giga.

GB. Gigabyte.

giga (G). Ten to the ninth power; 1 000 000 000 in decimal notation. When referring to storage size, 2 to

the thirtieth power, 1 073 741 824 in decimal notation.

guest. In interpretive execution mode, the interpreted or virtual machine as opposed to the real machine (the host).

H

hardware-assist control (HAC). A function used to control and simplify hardware information transfer between the central processors and the channel subsystem, and includes the hardware system area directory in central storage, the channel subsystem Licensed Internal Code, and all the control blocks used for I/O operations.

hardware system area (HSA). A logical area of central storage, not addressable by application programs, used to store Licensed Internal Code and control information.

hex. Hexadecimal.

hexadecimal. (1) Pertaining to a selection, choice, or condition that has 16 possible values or states. (I) (2) Pertaining to a fixed-radix numeration system, with radix of 16. (I) (3) Pertaining to a numbering system with base of 16; valid numbers use the digits 0—9 and characters A—F, where A represents 10 and F represents 15.

hexadecimal number. The 1-byte hexadecimal equivalent of an EBCDIC character.

high-speed buffer. A cache or a set of logically partitioned blocks that provides significantly faster access to instructions and data than provided by central storage.

HPO. High performance option.

HSA. Hardware system area.

I

IBM program support representative (PSR). An IBM service representative who performs maintenance services for IBM software at a centralized IBM location. Contrast with *IBM systems engineering operations specialist*.

IBM systems engineering operations specialist. An IBM service representative who performs maintenance services for IBM software in the field. Contrast with *IBM program support representative*.

ICE. Interconnect communication element.

ID. Identifier.

identifier (ID). (1) One or more characters used to identify or name a data element and possibly to indicate certain properties of that data element. (T) (2) A sequence of bits or characters that identifies a program, device, or system to another program, device, or system. (3) In an ESCON Director, a user-defined symbolic name of 24 characters or fewer that identifies a particular ESCD. See also *password identifier* and *port address name*.

IE. Instruction element.

IML. Initial machine load.

IMLP. Initial machine load pending.

immediate access storage. A storage device whose access time is negligible in comparison with other operating times. (A)

inactive subchannel. A subchannel that is busy (even if unlocked), not busy, or has a pending interrupt, and is indicated by subchannel status word (SCSW) bit 24 equals 0. The control information resides outside the channel subsystem because it is not necessary for the current operation. Contrast with *active subchannel*. See also *busy subchannel*.

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.

initial machine load (IML). A procedure that prepares a device for use.

initial program load (IPL). (1) The initialization procedure that causes an operating system to commence operation. (2) The process by which a configuration image is loaded into storage at the beginning of a work day or after a system malfunction. (3) The process of loading system programs and preparing a system to run jobs.

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 address. (1) In System/370 mode, a 16-bit address that consists of two parts: the 8 bits of the leftmost part constitute the channel address, and the 8 bits of the rightmost part constitute the device address. (2) In ESA/390 mode, an address provided by the channel subsystem that consists of four parts: channel path identifiers, subchannel numbers, device

numbers, and addresses (called device identifiers). Device identifiers are dependent on the channel path type and not visible to programs. See *channel path identifier, device identifier, device number*, and *subchannel number*.

input/output configuration. The collection of channel paths, control units, and I/O devices that attaches to the processor.

input/output configuration data set (IOCDS). The data set that contains an I/O configuration definition built by the I/O configuration program (IOCP).

input/output configuration program (IOCP). A program that defines to a system all the available I/O devices and the channel paths.

input/output interface. The interface that connects channels and control units for the exchange of signals and data.

input/output unit. A device in a data processing system by means of which data can be entered into the system, received from the system, or both. (I) (A)

input unit. A device in a data processing system by means of which data can be entered into the system. (I) (A) Contrast with *output unit*.

instruction element (IE). A part of a processor that executes some instructions and generates operand addresses and instruction requests. It also controls the sequencing of instructions through the machine and is usually controlled by Licensed Internal Code.

interrupt. (1) A suspension of a process, such as execution of a computer program caused by an external event, and performed in such a way that the process can be resumed. (A) (2) To stop a process in such a way that it can be resumed. (3) In data communication, to take an action at a receiving station that causes the sending station to end a transmission. (4) To temporarily stop a process.

interruption. Synonym for *interrupt*.

invalidation. The process of removing records from cache because of a change in status of a subsystem facility or function, or because of an error while processing the cache image of the set of records. When such a cache image is invalidated, the corresponding records cannot be accessed in cache and the assigned cache space is available for allocation.

I/O. Input/output.

IOCDS. I/O configuration data set.

IOPD. I/O problem determination.

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*.

IPL. Initial program load.

J

jacket. In an optical cable, the outermost layers of protective covering.

K

k. Kilo.

kb. Kilobit.

KB. Kilobyte.

KB/s. Kilobytes per second.

kilo (k). Thousand.

kilobit (kb). One thousand binary digits.

kilobyte (KB). 1024 bytes for storage size; otherwise, 1000 bytes.

kilometer. One thousand meters; 0.62 mile.

L

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)

LIC. Licensed Internal Code.

Licensed Internal Code (LIC). Software provided for use on specific IBM machines and licensed to customers under the terms of IBM's Customer Agreement. Microcode can be Licensed Internal Code and licensed as such.

link. (1) In an ESCON environment, the physical connection and transmission medium used between an optical transmitter and an optical receiver. A link consists of two conductors, one used for sending and the other for receiving, providing a duplex communication path. (2) In an ESCON I/O interface, the physical connection and transmission medium used between a channel and a control unit, a channel and an ESCD, a control unit and an ESCD, or, at times, between two ESCDs.

link address. In an ESCON environment, an address assigned at initialization that identifies a channel or control unit and allows it to send and receive trans-

mission frames and perform I/O operations. See also *port address*.

link-attached. Pertaining to devices that are connected to a controlling unit by a data link. Contrast with *channel-attached*.

link-level addressing. In an ESCON I/O interface, one of two levels of addressing, this level pertaining to link-level functions and identifying the channel path between the channel and a control unit. Contrast with *device-level addressing*.

link segment. In an ESCON environment, any portion of an optical cable between connectors, including passive components.

local. Synonym for *channel-attached*.

local storage. Synonym for *local working storage*.

logically partitioned (LPAR) mode. A central processor mode, available on the Configuration frame when using the PR/SM feature, that allows an operator to allocate processor unit hardware resources among logical partitions. Contrast with *basic mode*.

logical partition (LP). A subset of the processor hardware that is defined to support the operation of a system control program (SCP). See also *logically partitioned (LPAR) mode*.

logical processor. In LPAR mode, a central processor in a logical partition.

logical switch number. A two-digit number used by the I/O configuration program (IOCP) to identify a specific ESCON Director.

logical-to-physical coordination. In ESCON Manager, the coordination of a physical connectivity change in an ESCON Director with an equal logical change in any associated host.

logical vector. In LPAR mode, a vector associated with a logical processor.

logic support station (LSS). Hardware used to attach processor complex elements that respond to commands from the processor controller and includes scan control logic.

LP. Logical partition.

LPAR. Logically partitioned.

M

m. (1) Milli; one thousandth part. (2) Meter.

machine check. An error condition that is caused by an equipment malfunction.

machine-readable information (MRI). Synonym for *soft copy*.

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 analysis procedure (MAP). A step-by-step procedure for tracing a symptom to the cause of a failure.

MAP. Maintenance analysis procedure.

maximum allowable link loss. In an ESCON environment, the maximum amount of link attenuation (loss), expressed in decibels, that can exist without causing a possible failure condition. Contrast with *calculated link loss*.

Mb. Megabit.

MB. Megabyte.

mega (M). 10 to the sixth power; 1 000 000 in decimal notation. When referring to storage size, two to the twentieth power; 1 048 576 in decimal notation.

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.

minimum acceptable receive level. In an ESCON environment, the calculated level, expressed in decibels, received at a specific point in the link. This value is used as a rejection criterion when measuring the actual level received at the same point.

mm. Millimeter.

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.

monitor. A device that observes and records selected activities within a data processing system for analysis.

monitor console. An optional logical display used to monitor the service or system console. Each monitor console can be assigned to any of the physical displays attached to the processor controller.

MP. (1) Multiprocessor. (2) Maintenance procedure.

MRI. Machine-readable information.

ms. Millisecond.

multiple preferred guests. A VM facility that, when used with the Processor Resource/Systems Manager (PR/SM) feature, supports multiple preferred virtual machines. See also *preferred virtual machine.*

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.*

multiplexing. In data transmission, a function that permits two or more data sources to share a common transmission medium so that each data source has its own channel. (I) (A)

multipoint. Pertaining to communication among more than two stations over a single telecommunication line.

multipoint connection. A connection established among more than two data stations. (I) (A) Contrast with *point-to-point connection.*

multipoint line. A telecommunication line or circuit connecting two or more stations. Contrast with *point-to-point line.*

multiprocessing. (1) A mode of operation for parallel processing by two or more processors of a multiprocessor. (I) (A) (2) Pertaining to the simultaneous execution of two or more computer programs or sequences of instructions by a computer. (A) (3) Loosely, parallel processing. (A) (4) Simultaneous execution of two or more sequences of instructions by a multiprocessor.

multiprocessor (MP). A processor complex that can be physically partitioned to form two operating processor complexes.

MVS/ESA. Multiple Virtual Storage/Enterprise Systems Architecture.

MVS/SP. Multiple Virtual Storage/System Product.

MVS/XA. Multiple Virtual Storage/Extended Architecture.

N

N/A. (1) Not applicable. (2) Not available.

non-physical-contact connector. In an ESCON environment, an optical fiber connector type having an air gap between itself and its receptacle, thereby providing a junction point having (usually) more loss compared to a physical-contact connector.

O

offline. (1) Pertaining to the operation of a functional unit that takes place either independently of, or in parallel with, the main operation of a computer. (T) (2) Neither controlled by, nor communicating with, a computer. Contrast with *online.*

online. (1) Pertaining to the operation of a functional unit when under the direct control of a computer. (T) (2) Pertaining to a user's ability to interact with a computer. (A) (3) Pertaining to a user's access to a computer via a terminal. (A) (4) Controlled by, or communicating with, a computer. Contrast with *offline.*

operating system (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 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.*

optical fiber cable. Synonym for *optical cable.*

optical fiber connector. A hardware component that transfers optical power between two optical fibers or bundles and is designed to be repeatedly connected and disconnected.

optical fiber splice. A permanent joint that couples optical power between two fibers. See also *fusion splice* and *mechanical splice.*

optical fiber waveguide. Synonym for *optical fiber*.

optical link. See *link*.

optical link loss. See *maximum allowable link loss*.

optical link segment. See *link segment*.

optical mode conditioner (OMC). A tool that, when inserted between an optical LED source and the ESCON link under test, provides a consistent method for measuring optical attenuation.

option. (1) A specification in a statement that may be used to influence the execution of the statement. (2) A hardware or software function that can be selected or enabled as part of a configuration process. (3) Hardware (such as a network adapter) that can be installed in a device to change or enhance device functions.

output device. Synonym for *output unit*.

output unit. A device in a data processing system by which data can be received from the system. (I) (A) Contrast with *input unit*.

P

PA. (1) Problem analysis. (2) Program access. (3) Program area.

page. In a virtual storage system, a fixed-length block that has a virtual address and is transferred as a unit between real storage and auxiliary storage. (I) (A)

paging. The transfer of pages between real storage and auxiliary storage. (I) (A)

parallel. (1) Pertaining to a process in which all events occur within the same interval of time, each handled by a separate but similar functional unit; for example, the parallel transmission of the bits of a computer word along the lines of an internal bus. (T) (2) Pertaining to concurrent or simultaneous operation of two or more devices or to concurrent performance of two or more activities in a single device. (A) (3) Pertaining to concurrent or simultaneous occurrence of two or more related activities in multiple devices or channels. (A) (4) Pertaining to the simultaneity of two or more processes. (A) (5) Pertaining to the simultaneous processing of the individual parts of a whole, such as the bits of a character and the characters of a word, using separate facilities for the various parts. (A)

parallel channel. 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*.

parameter. (1) A variable that is given a constant value for a specified application and that may denote the application. (I) (A) (2) An item in a menu for which the user specifies a value or for which the system provides a value when the menu is interpreted. (3) Data passed between programs or procedures.

partition. See *logical partition* and *target logical partition*.

passive component. In an ESCON environment, any link component that does not contain electronic logic; for example, a splice, connector, adapter, coupler, distribution panel, or optical cable.

password. (1) A value used in authentication or a value used to establish membership in a set of people having specific privileges. (2) A unique string of characters known to the computer system, and to a user who must specify it to gain full or limited access to a system and to the information stored within it.

password identifier. In an ESCON Director, a user-defined symbolic name of 24 characters or fewer that identifies the password user.

patch panel. Synonym for *distribution panel*.

path. In a network, any route between any two nodes. (T)

PER. Program event recording.

physical address. The absolute address after configuration (the final address). See also *absolute address*, *logical address*, *real address*, and *virtual address*.

physical-contact connector. In an ESCON environment, an optical fiber connector type having a polished end face that aligns precisely with its receptacle, providing an extremely low-loss junction point. Contrast with *non-physical-contact connector*.

physical partitioning. The process of reconfiguring the processor complex from the single-image configuration to the physically partitioned (PP) 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.

point-to-point line. A switched or nonswitched telecommunication line that connects a single remote station to a computer. Contrast with *multipoint line*.

point-to-point topology. A network topology that provides one communication path between a channel and a control unit and does not include switching facilities. Contrast with *switched point-to-point topology*. See also *multidrop topology*.

polling. (1) On a multipoint connection or a point-to-point connection, the process whereby data stations are invited one at a time to transmit. (I) (2) Interrogation of devices for such purposes as to avoid contention, to determine operational status, or to determine readiness to send or receive data. (A)

port address. In an ESCON Director, an address used to specify port connectivity parameters and to assign link addresses for attached channels and control units. See also *link address*.

port address name. In an ESCON Director, a user-defined symbolic name of 24 characters or fewer that identifies a particular port.

port card. In an ESCON environment, a field-replaceable hardware component that provides the optomechanical attachment method for jumper cables and performs specific device-dependent logic functions.

port information block (PIB). In an ESCON Director, a data area that contains information relating to the connectivity of each available port.

port number. In an ESCON Director, a hexadecimal number that identifies a physical link connection point. This number is the same as the port address unless the service representative has reassigned the port associated with that address.

power-on reset state. The condition after a machine power-on sequence and before an IPL of the control program.

power-on self-test (POST). A series of tests that are run each time power is switched on.

PP. Physically partitioned (as in configuration).

PPC. Primary power compartment.

ppm. Parts per million.

PRB. Problem record block.

preferred virtual machine. (1) In VM, a particular virtual machine to which one or more of the performance options have been assigned. (2) In the VM/XA System Product, a virtual machine that runs in the V = R (virtual = real) area. The control program gives this virtual machine preferred treatment in performance, processor assignment, and I/O interrupt handling. See also *multiple preferred guests*.

prefetch. To fetch data from central storage before it is needed and to store data in a buffer.

preventive maintenance. Maintenance performed specifically to prevent faults from occurring. (A) Contrast with *corrective maintenance*.

primary distribution panel. The connection point at a building's service entrance that feeds the distribution panels for one or more computer systems.

primary ESCON Manager. In a system environment containing multiple ESCON Managers, the source of ESCM commands.

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 Resource/Systems Manager (PR/SM) feature. The feature that allows the processor to use several system control programs (SCPs) simultaneously, provides logical partitioning capability for the real machine, and provides support for multiple preferred guests.

processor storage. The storage available to the central processors (central storage and, if installed, expanded storage).

processor unit. Hardware that consists of one or more central processors, central storage, optional expanded storage, channel subsystem, and system control element.

product data. In an ESCON Director, information contained in an electrically-erasable programmable read-only memory (EEPROM) module that defines specific hardware characteristics and can be displayed or modified.

programming RPQ (PRPQ). A customer request for a price quotation on changes or additions to the functional capabilities of system control programming or licensed programs. The RPQ can be used with computing system RPQs to solve unique data processing problems. See also *computing system RPQ*.

program temporary fix (PTF). A temporary solution or bypass of a problem diagnosed by IBM as resulting from a defect in a current, unaltered release of the program.

program version. A separate IBM licensed program, based on an existing IBM licensed program, that usually has important new code or functions.

protective plug. In an ESCON environment, a type of duplex connector that provides environmental and physical protection. Contrast with *wrap plug*.

protocol. (1) A set of semantic and syntactic rules that determines the behavior of functional units in achieving communication. (I) (2) In SNA, the meanings of and the sequencing rules for requests and responses used for managing the network, transferring data, and synchronizing the states of network components. (3) A specification for the format and

relative timing of information exchanged between communicating parties.

PRPQ. Programming RPQ.

PR/SM. Processor Resource/Systems Manager.

PS. Power supply.

PSB. Present status byte.

PSR. Program support representative.

PSW. Program status word.

PU. (1) Processing unit. (2) Physical unit.

R

RAM. Random access memory.

random access memory (RAM). (1) A storage device in which data can be written and read. (2) A storage device into which data is entered and from which data is retrieved in a nonsequential manner.

RAS. Reliability, availability, serviceability.

RCS. Remote control support.

REA. Request for engineering action.

read-only memory (ROM). (1) A storage device in which data, under normal conditions, can only be read. (T) (2) Memory in which stored data cannot be modified by the user except under special conditions.

real address. An address before prefixing, such as found in the instruction address portion of the channel status word (CSW). If translation is off, the logical address is the real address. See also *absolute address*, *logical address*, *physical address*, and *virtual address*.

REC. Request for engineering change.

receiver. In fiber optics, see *optical receiver*.

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. Synonym for *link-attached*.

remote support facility (RSF). A system facility invoked by Licensed Internal Code that provides procedures for problem determination and error detection.

request for engineering action (REA). A document that starts the release of an engineering change.

request for price quotation (RPQ). See *computing system RPQ* and *programming RPQ (PRPQ)*.

RPQ. Request for price quotation.

RSC. Remote support center.

RSF. Remote support facility.

S

SA. (1) System area. (2) Stand-alone.

scalar. (1) A quantity characterized by a single value. (I) (A) (2) A type of program object that contains either string or numeric data. It provides representation and operational characteristics to the byte string to which it is mapped. (3) Pertaining to a single data item as opposed to an array of data items. (4) Contrast with *vector*.

scattering. In fiber optics, the change in direction of light rays or photons after striking a small particle or particles. In can also be regarded as the diffusion of a light beam caused by the inhomogeneity of the transmitting medium. (E)

SCE. System control element.

SCH. Subchannel.

SCP. System control program.

SCS. Single channel service.

SCSW. Subchannel status word.

SDLC. synchronous data link control.

SEC. system engineering change.

secondary ESCON Manager. In a system environment containing multiple ESCON Managers, an ESCM that receives a command from the primary ESCM through intersystem communication.

selector channel. An I/O channel designed to operate with only one I/O device at a time. After the I/O device is selected, a complete record is transferred 1

byte at a time. Contrast with *byte multiplexer channel*. See also *block multiplexer channel*.

service console. A logical device used by service representatives to maintain the processor complex and to isolate failing field-replaceable units (FRUs). The service console can be assigned to any of the physical displays attached to the processor controller.

service representative. A person who performs maintenance services for IBM hardware products or systems. See also *IBM program support representative*.

SI. Single-image configuration.

signal interruption. A logic signal sent from the channel subsystem to the central processor at the completion of an I/O operation that indicates the subsystem needs to transfer information. The signal is sent as a coded series of bits contained in a command to the system control element.

single channel service (SCS). Testing a single channel while the other channels remain operational. See also *concurrent maintenance*.

single-image configuration. A system configuration that allows the processor controller to use one processor complex side while the other side is available as a backup. The A-side or B-side becomes the operational processor complex, the remaining side becomes the backup, and all the functional elements can be activated.

single-mode optical fiber. An optical fiber in which only the lowest-order bound mode (which can consist of a pair of orthogonally polarized fields) can propagate at the wavelength of interest. (E) Contrast with *multimode optical fiber*.

SLC. Service language command.

soft copy. (1) A nonpermanent copy of the contents of storage in the form of a display image; for example, a cathode ray tube display. (T) (2) Information that a machine can read or obtain from a storage medium. (3) One or more files that can be electronically distributed, manipulated, and printed by a user.

splice. See *optical fiber splice*.

splice loss. Synonym for *insertion loss*.

storage. (1) A functional unit into which data can be placed, in which they can be retained, and from which they can be retrieved. (T) (2) The action of placing data into a storage device. (I) (A)

storage control. In a DASD subsystem, the component that connects the DASD to the host channels, performs channel commands, and controls the DASD devices. For example, the 3990 Model 2 and Model 3 are storage controls.

storage subsystem. A storage control and its attached storage devices.

stride. The number of element positions in storage needed to advance from one vector element to the next. For example, contiguous vectors have a stride of one.

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*.

supervisor state. A condition during which a processor can execute input/output and other privileged instructions. Contrast with *problem state*.

SVCH. Service channel.

switched point-to-point channel path configuration. In an ESCON I/O interface, a configuration that consists of a link between a channel and an ESCON Director and one or more links from the ESCD, each of which attaches to a control unit. This configuration depends on the capabilities of the ESCD for establishing and removing connections between channels and control units. Contrast with *point-to-point channel path configuration*.

switched point-to-point topology. A network topology that uses switching facilities to provide multiple communication paths between channels and control units. See also *multidrop topology*. Contrast with *point-to-point topology*.

switchover. Pertaining to the operation that changes the backup side of the processor controller to the active side and places the active side offline.

synchronous. (1) Pertaining to two or more processes that depend on the occurrence of specific events, such as common timing signals. (T). (2) Occurring with a regular or predictable time relationship. Contrast with *asynchronous*.

SYSGEN. System generation.

Sysplex Timer. An IBM table-top unit that synchronizes the time-of-day (TOD) clocks in as many as 16 processors or processor sides.

system configuration. A process that specifies the devices and programs that form a particular data processing system.

system console. (1) A console, usually having a keyboard and a display screen, that is used by an operator to control and communicate with a system. (2) A logical device used for the operation and control of hardware functions (for example, IPL, alter/display, and reconfiguration). The system console can be assigned to any of the physical displays attached to the processor controller.

system control element (SCE). Hardware that handles the transfer of data and control information associated with storage requests between the elements of the processor.

system control programming (SCP). IBM-supplied programming that is fundamental to the operation and maintenance of the system. It serves as an interface with licensed programs and user programs and is available without additional charge.

system display. A display station that attaches to a port of the processor controller and is required for normal system operation. See also *service support display*.

system generation (SYSGEN). The process of selecting optional parts of an operating system and of creating a particular operating system tailored to the requirements of a data processing installation. (I) (A)

system log (SYSLOG). A data set or file in which job-related information, operational data, descriptions of unusual occurrences, commands, and messages to or from the operator are stored.

System Power panel. The part of the control panel assembly on the processor controller used by both customers and service representatives to determine and control the power status of both the processor controller and the processor complex. Contrast with *Service panel*.

system reference code. A code that contains information, such as a failing field-replaceable unit, for a service representative.

system reset (SYSRESET). To reinitialize the execution of a program by repeating the initial program load (IPL) operation.

system resource. Hardware, such as a central processor, I/O devices, channel paths, software programs, or other components that contribute to system operation. See also *critical resource*.

T

target logical partition. In LPAR mode, the current or immediate logical partition being used or displayed. It is identified on the status line and determined by the SETLP service language command.

target processor. The processor that controls execution during a program restart, instruction trace, stand-alone dump, or IPL, and whose ID is identified by highlighting on the status line.

terminal. (1) A functional unit in a system or communication network at which data may enter or leave. (T) (2) A point in a system or communication network at which data can either enter or leave. (A) (3) In data communication, a device, usually having a keyboard and display device, capable of sending and receiving information. See also *display station* and *workstation*.

thermal conduction module (TCM). A field-replaceable unit (FRU) that has multiple logic modules containing level-sensitive scan design logic.

threshold. A level, point, or value above which something is true or will take place and below which it is not true or will not take place.

time-of-day (TOD) clock. A system hardware feature that is incremented once every microsecond, and provides a consistent measure of elapsed time suitable for indicating date and time. The TOD clock runs regardless of whether the processor is in a running, wait, or stopped state.

time-out. (1) An event that occurs at the end of a predetermined period of time that began at the occurrence of another specified event. (I) (2) A time interval allotted for specific operations to occur; for example, response to polling or addressing before system operation is interrupted and must be restarted. (3) A terminal feature that logs off a user if an entry is not made within a specified period of time.

TOD. Time of day.

trace. (1) A record of the execution of a computer program. It exhibits the sequences in which the instructions were executed. (A) (2) A record of the frames and bytes transmitted on a network.

trunk cable. In an ESCON environment, a cable consisting of multiple fiber pairs that do not directly attach to an active device. This cable usually exists between distribution panels and can be located within, or external to, a building. Contrast with *jumper cable*.

U

user identification (user ID). A string of alphanumeric characters that identifies a user to a system.

V

VE. Vector element.

vector. (1) A quantity usually characterized by an ordered set of numbers. (I) (A) (2) A one-dimensional array. Contrast with *scalar*.

vector facility. An optional feature that enables a processor to run programs that issue vector instructions used for scientific calculations.

VF. Vector facility.

virtual address. The address of a location in virtual storage. A virtual address must be translated into a real address to process the data in processor storage. See also *absolute address, logical address, physical address*, and *real address*.

virtual machine (VM). (1) A virtual data processing system that appears to be at the exclusive disposal of a particular user, but whose functions are accomplished by sharing the resources of a real data processing system. (T) (2) A functional simulation of a computer system and its associated devices, multiples of which can be controlled concurrently by one operating system.

Virtual Machine/System Product (VM/SP). An IBM licensed program that manages the resources of a single computer so that multiple computing systems appear to exist. Each virtual machine is the functional equivalent of a "real" machine.

virtual storage (VS). (1) The storage space that may be regarded as addressable main storage by the user of a computer system in which virtual addresses are mapped into real addresses. The size of virtual storage is limited by the addressing scheme of the computer system and by the amount of auxiliary storage available, not by the actual number of main storage locations. (I) (A) (2) Addressable space that is apparent to the user as the processor storage space, from which the instructions and the data are mapped into the processor storage locations.

VM. Virtual machine.

VM/ESA. Virtual Machine/Enterprise Systems Architecture.

VM/SP. Virtual Machine/System Product.

VM/SP HPO. Virtual Machine/System Product High Performance Option.

VM/XA. Virtual Machine/Extended Architecture.

VM/XA SF. Virtual Machine/Extended Architecture Systems Facility.

VM/XA SP. Virtual Machine/Extended Architecture System Product.

W

wrap plug. In an ESCON environment, a type of duplex connector used to wrap the optical output signal of a device directly to the optical input. Contrast with *protective plug*.

Index

A

About this Publication xiii
access
  central storage 3-12
  data deferred 9-8
action queue X-1
activate logical partition X-1
active configuration X-1
active state indicator X-1
active subchannel X-1
adapter X-1
address X-1
address generation X-2
address translation X-2
Advanced Interactive Executive/370 1-13, 3-7
AIX/370 1-13, 3-7
allocate X-2
analysis routine X-2
APAR X-2
application X-2
application program X-2
architectural deviations A-1
architectural mode of operation 8-2
assists 7-9
  control-switch assist 7-13
  CPU Timer, VM Assists 7-28
  preferred machine assist 7-21
  virtual-machine 7-28
asynchronous X-2
attention identifier (AID) X-2
authorized program analysis report (APAR) X-2
auto-vectoring X-2
automated logic diagram (ALD) X-2
automatic fault isolation manager X-2
auxiliary storage X-2
availability (RAS) 1-11, 2-8

B

backup copy X-2
balancing device load 10-5
balun X-2
basic control mode feature 7-11
basic mode X-2
basic modes 1-12
basic telecommunications access method (BTAM) X-2
BCE 3-4, X-2
bimodal addressing feature 7-11
block multiplexer mode of operation 9-9
bpi X-2
bps X-2
branch and save feature 7-11
broadcasted purging 7-11
buffer X-2
buffer control element 3-4
buffer directory 3-4
buffered device 9-8
burst mode 9-10
burst size 10-4
bus X-2
bus-in data X-2
bus-out data X-2
busy subchannel X-3
byte X-3
byte multiplexer
  channel 10-4
  operation 9-10
byte multiplexer channel X-3
byte multiplexer mode 9-10
byte-oriented operand 7-11

C

CAC X-3
cache 2-6, 3-4, X-3
CAI X-3
calculated link loss X-3
CAW X-3
CCA X-3
CCC 6-4, 8-5
CCW 2-7, 7-12, 7-29, X-3
CCW (channel command word) 8-5
CDC 6-4, 8-5, X-3
CE X-3
central processor B-2, X-3
  description of 3-3
  elements of 3-3
central storage 3-8, B-2, X-3
CERF X-3
chaining
  check 9-8
  command 8-5
  data 8-5, 9-8, 10-4
  definition of 8-5
  operation 8-5
  type of 8-5
channel 2-5, X-3
  block multiplexer 10-3
  byte multiplexer 10-4
  command word 8-5
  configuration 2-5, 8-6, 10-2
  display frames 8-2
  extension 10-4
  function 3-12
  implementation 8-6
  initialization 3-12
  load balancing 10-4

channel *(continued)*
operation 2-5, 3-12, 8-2, 8-6
performance 9-2
priority 10-3
program 8-2
separate requirements 10-3
service time 10-3
time-out
active 8-6
inactive 8-6
channel address X-3
channel address word (CAW) X-3
channel command word (CCW) X-3
channel communication area (CCA) X-3
channel control check 6-4, 8-5, X-3
channel data check 6-4, X-3
channel element 3-11
channel groups 7-12
channel indirect addressing 7-12
channel Licensed Internal Code X-3
channel operation characteristics 8-2
channel path configuration X-3
channel path identifier 8-2
channel path identifier (CHPID) X-3
channel path (CHP) X-3
channel server 3-11
channel set X-4
channel status word (CSW) X-4
channel subsystem 2-5, 3-11, 7-12, B-2
block multiplexer 3-11
configuration 3-11
configuration guidelines 10-1
operation 3-11
channel subsystem performance characteristics 9-1
channel subsystem performance characteristics, summary 9-2
channel subsystem performance 9-2
channel subsystem (CSS) X-4
channel subsystem (CSS) Licensed Internal Code X-4
channel-attached X-3
channel-set switching 7-12
channel-subsystem call 7-12
channels 7-12
characteristics 2-2
check stop X-4
CHN 3-11, X-4
CHP X-4
CHPID X-4
CICS X-4
class, device 9-11, 10-3
clear I/O 7-12
cm X-4
CME X-4
CMS X-4
CNC X-4
coaxial cable X-4
COB X-4
collision X-4
command X-4
chaining 8-5
deferred access 9-8
command chaining 8-5, X-4
command retry 7-12
communication controller node X-4
compatibility, programming 1-11
component X-4
components, description 2-10
computing system RPQ X-4
concurrent
channel maintenance 2-9
CP TCM maintenance 2-9
power subsystem 2-9
concurrent channel maintenance 7-13
concurrent CP TCM maintenance 7-13
concurrent maintenance X-4
concurrent power subsystem 2-9, 7-9, 7-13
concurrent sense 7-12
conditional swapping 7-12, 7-13
configuration 1-3, 2-2, 5-4, X-4
channel 8-6, 10-2
channel subsystem 2-5, 3-11
console 2-4, 4-2
model 520 1-3
model 640 1-3
model 660 1-3
model 740 1-4
model 820 1-4
model 860 1-5
model 900 1-6
physically partitioned 5-3
single-image 5-3, 7-12
configuration matrix X-4
configurations 2-4, B-2
central processor B-2
central storage B-2
channel subsystem B-2
CP B-2
IBM 9037 Sysplex Timer B-2
ICE B-2
SCE B-2
system controller B-2
vector facility B-2
VE0-5 B-2
connectivity attribute X-4
connectivity capability X-4
connectivity control X-4
connector X-4
consensus processing X-4
console 4-2, 4-3, X-5
configuration 2-4, 4-2
display 4-2
logical 4-2
requirements 2-4
type 4-2

console integration 7-13
consoles
  3206 4-2
consoles and displays 4-1
control blocks, ESA/390 3-12
control check 6-4, 8-5
control operation X-5
control panel 5-3
control panel assembly X-5
Control program X-5
  ESA/390 1-12
  guest control 1-13
  LPAR mode 1-13
control storage element 3-3
control unit X-5
control units, physical placement 10-3
control-switch assist 7-13
control-unit end X-5
controller X-5
conversational monitor system (CMS) X-5
conversion X-5
cooling 2-4
core X-5
corrective maintenance X-5
CP 3-3, B-2
CPU address identification 7-20
CPU ID, contents of 3-5
CPU retry 7-14
CPU signaling and response 7-20
CPU timer and clock comparator 7-14
critical resource X-5
critical time 9-9, 9-11, 10-4, 10-5, 10-6
CS X-5
CSE 3-3, X-5
CSS B-2, X-5
CSW X-5
CTC X-5
CU X-5
cursor X-5
Customer Information Control System (CICS) X-5
CVC X-5

D

DASD X-5
DASD subsystem X-5
DAT 7-27
DAT (dynamic address translation) 3-4
data
  chaining 8-5, 10-4
  deferred access 9-8
  rate 10-4
data bus X-5
data chaining 8-5
data check 6-4, 8-5
data check, channel 8-5
data encryption algorithm 7-18
data encryption standard 7-18
data integrity 2-7
data processing (DP) X-5
data representation 2-7
data station X-5
data stream X-5
data streaming 7-14, X-5
  DB2 sort facility 7-14
  dynamic reconfiguration management 7-15
data streaming protocol 9-9
data terminal equipment (DTE) X-6
data transfer X-6
data transfer mode X-6
data transfer phase X-6
data transfer rate 9-3
deconfigure X-6
dedicated connection X-6
dedication X-6
default X-6
deferred access 9-8
deferred access, command 9-8
definition
  LPAR mode 4-3
  system 4-3
degradation 9-8
degraded X-6
design highlights 1-6
deviations, architectural A-1
device X-6
  block multiplexer
  buffered 9-8
  channel 10-3
  class 10-3, 10-6
  intermix 10-3
  I/O 9-10
  load-balancing across channels 10-3
  performance 9-10
  priority 8-7
  sequence 10-3, 10-6
device address X-6
device class 9-11, 10-3
device end signal 8-5
device identifier X-6
device number X-6
device-level addressing X-6
diagnostics X-6
direct access storage X-6
direct access storage device (DASD) X-6
direct-coupled interlock (DCI) X-6
directory lookaside table (DLAT) X-6
display 4-2
  optional 2-4
  requirements 2-4
display device X-6
display models 2-4
display/printer adapter (DPA) X-6
distribution panel X-7
downstream X-7

downstream load (DSL) X-7
DP X-7
dual processors 5-3
dual-address space 7-29
dynamic address translation 7-27
dynamic address translation (DAT) 3-4, X-7
dynamic connection X-7
dynamic connectivity X-7
dynamic reconfiguration X-7
dynamic reconfiguration management X-7
dynamic reconnection 8-8

E

EE 3-5, X-7
elements, central processor
  buffer control element 3-3
  control storage element 3-3
  execution element 3-3
  instruction element 3-3
emitter-coupled logic (ECL) 2-8
emulate X-7
emulation X-7
end of operation (EOP) X-7
ending status 8-5
enhanced move page 7-15
enhanced power subsystem 2-9, 7-9, 7-15
error X-7
  automatic error recovery 6-2
  checking and correction 2-8, 2-9, 3-9
  error checking and correction 3-3, 3-9, 3-11, 6-2, 7-9, 7-15, 7-19
  error handling 6-4
  error handling, I/O 6-4
  error logout 5-4
  error recovery 5-2, 5-4, 6-2
  error recovery, automatic 6-2
  ESCON CTCA 7-16
  I/O error alert 7-19
  I/O operation 6-4
  logout 2-10
error checking and correction (ECC) X-7
error handler X-7
error handling 6-1, 8-5
error log X-7
error message X-7
ESA/370 2-6
ESA/390 1-7, X-7
ESC X-7
ESCD X-7
ESCD console X-7
ESCD console adapter X-7
ESCM X-7
ESCON X-7
  architecture 1-8
  channels 1-9, 7-16
  technology 1-8
ESCON channel X-7
ESCON Director (ESCD) X-7
ESCON environment X-7
ESCON Manager (ESCM) X-7
ESE X-8
ES/9000 configurations 2-4
  See also configurations
event X-8
exchange X-8
execution element X-8
  description of 3-5
expanded storage 1-10, 3-11, X-8
extended control (EC) mode 7-27
extended real addressing 7-16
extended sorting 7-16
extended-precision divide 7-16
extended-precision floating point 7-16
extensions for virtual machines 7-17

F

facility X-8
failure X-8
fast release 7-17
fault tolerant dynamic memory array 2-9
fault tolerant dynamic memory arrays 7-17
feature X-8
feature code X-8
feature descriptions
  basic control mode 7-11
  bimodal addressing 7-11
  branch and save 7-11
  byte-oriented operand 7-11
  channel indirect addressing 7-12
  channel subsystem 7-12
  channel-set switching 7-12
  channels 7-12
  clear I/O 7-12
  command retry 7-12
  conditional swapping 7-12
  control-switch assist 7-13
  CPU address identification 7-20
  CPU retry 7-14
  CPU signaling and response 7-20
  CPU timer and clock comparator 7-14
  CPU Timer, VM Assists 7-28
  data streaming 7-14
  error checking and correction 7-15
  expanded storage 7-16
  extended real addressing 7-16
  extended-precision divide 7-16
  extended-precision floating point 7-16
  fast release 7-17
  floating point 7-17
  halt device 7-17
  high-speed buffer storage 7-17
  IBM 9037 Sysplex Timer 7-25
  ICRF 7-18
  indirect data address word (IDAW) 7-29
  interpretive execution 7-17

feature descriptions *(continued)*
interval timer 7-17
I/O error alert 7-19
logging volume reduction 7-20
multiprocessing 7-20
page protection 7-21
preferred machine assist 7-21
prefixing 7-20
Processor Resource/Systems Manager 7-22
PSW-key handling 7-24
real addressing 7-29
recovery extensions 7-24
service signal 7-24
shared main storage 7-21
sorting instructions 7-24
storage key instructions 7-25
storage-key instruction extensions 7-25
storage-key 4KB block 7-25
Sysplex Timer 7-25
System/370 extended facility 7-26
test block 7-26
time-of-day clock 7-26
TOD clock synchronization 7-21
tracing 7-26
translation 7-27
vector facility feature 7-28
virtual-machine assist 7-28
3033 extension 7-29
features
architecture dependent 7-4
optional 7-4
standard 7-4
fiber X-8
fiber buffer X-8
fiber optic
cabling 1-9
reception 1-9
transmission 1-9
fiber optic cable X-8
fiber optics X-8
field macro diagram (FMD) X-8
field-replaceable unit (FRU) X-8
filter X-8
flag X-8
floating point 7-17
floating-point X-8
FP X-8
frame X-8
frame check sequence (FCS) X-8
frame deallocation 3-9
FRU X-8
function
of channels 3-12
of interconnect communication element 3-12

G

G X-8
GB X-8
giga (G) X-8
guest X-9
guest control programs 1-13

H

halt device 7-17
hardware configuration 1-3, 1-6
hardware system area 3-9, 3-10
hardware system area (HSA) X-9
hardware-assist control (HAC) X-9
hex X-9
hexadecimal X-9
hexadecimal number X-9
high-speed buffer 3-4, X-9
high-speed buffer storage 7-17
HPO X-9
HSA 3-10, X-9
HSA (hardware system area) 8-2

I

IBM program support representative (PSR) X-9
IBM systems engineering operations specialist X-9
IBM 9037 Sysplex Timer 1-11, B-2
ICE 3-11, 3-12, X-9
ICRF 1-10, 7-18
ID X-9
identifier (ID) X-9
IE 3-4, X-9
IFCC 6-4, 8-5
IML X-9
IMLP X-9
immediate access storage X-9
implementation, channel 8-6
inactive subchannel X-9
indirect data address word (IDAW) 7-29
initial machine load (IML) X-9
initial program load (IPL) X-9
initialization X-9
channel 3-12
system 3-12
input unit X-10
input/output address X-9
input/output configuration X-10
input/output configuration data set (IOCDS) X-10
input/output configuration program (IOCP) X-10
input/output interface X-10
input/output unit X-10
input/output (I/O) X-9
input/output, see I/O 2-5
instruction element (IE) X-10
instruction element, description of 3-4
instructions
BAS 7-11

instructions *(continued)*
BASR 7-11
Clear Subchannel (CSCH) 8-4
CLRCH 7-24
Halt Subchannel (HSCH) 8-4
IPK 7-24
IPTE 7-26
Modify Subchannel (MSCH) 8-4
Reset Channel Path (RCHP) 8-4
Resume Subchannel (RSCH) 8-4
RIO 7-29
RRBE 7-25
SCK 7-26
SCKC 7-14
Set Address Limit (SAL) 8-4
Set Channel Monitor (SCHM) 8-4
SIE 1-7, 1-12
SIGP 7-20
SIOF 7-17
sorting 7-24
SPKA 7-24
SPT 7-14
SSK 7-25
SSM 7-28
STAP 3-6
Start Subchannel (SSCH) 8-4
STCKC 7-14
STIDP 3-5
storage-key instructions 7-25
Store Channel Path Status (STCPS) 8-4
Store Channel Report Word (STCRW) 8-4
Store Subchannel (STSCH) 8-4
STPT 7-28
SVC 7-26
Test Pending Interruption (TPI) 8-4
Test Subchannel (TSCH) 8-4
TPROT 7-26
Integrated Cryptographic Feature 1-10, 3-7
interconnect communication element 3-11, 3-12
function 3-12
operation 3-12
structure 3-12
interface
I/O 8-2
protocol, I/O
data streaming 9-9
interlocked 9-9
interface control check 6-4, 8-5
interlocked protocol 9-9
interpretive execution 7-18
interrupt X-10
interruption X-10
interruptions
control check 8-5
data check 8-5
interface control check 8-5
I/O 8-5
machine check 8-5
interruptions, types of
channel control check 6-4
channel data check 6-4
interface control check 6-4
interval timer 7-18
introduction 1-1
invalidation X-10
IOCDS X-10
IOCP (Input/Output Configuration Program) 8-2
IOP Licensed Internal Code X-10
IOPD X-10
IOPS Unit 2-5
IPL X-10
I/O 8-6, X-10
device 8-2, 9-10
error handling, I/O 6-4
interface 8-2
interface mode
data streaming 9-9
interlocked 9-9
interruptions 8-5
performance 9-10
problem determination 4-4
system reset 8-6
I/O operations 2-5
I/O Power Sequence Unit, 9023 2-5
I/O-device self-description 7-19

J

jacket X-10

K

k X-10
kb X-10
KB/s X-10
key-controlled storage protection 2-7, 3-9, 3-10, 7-19, 7-25, 7-29, A-1
key-controlled storage protection 7-19
limited channel logout 7-19
monitoring 7-19
keylock, security 2-8
kilo (k) X-10
kilobit (kb) X-10
kilobyte (KB) X-10
kilometer X-10

L

laser X-10
LIC X-10
Licensed Internal Code (LIC) X-10
limited channel logout 7-20
link X-10
link address X-10
link segment X-11
link-attached X-11

link-level addressing X-11
load balancing device 10-4
local X-11
local storage X-11
logging volume reduction 7-20
logic support station (LSS) X-11
logical partition 1-7
logical partition (LP) X-11
logical partitioning 7-22
logical processor X-11
logical switch number X-11
logical vector X-11
logical-to-physical coordination X-11
logically partitioned operating mode 1-7
logically partitioned (LPAR) mode X-11
LP X-11
LPAR X-11
LPAR mode 1-7

M

m X-11
machine check X-11
machine cycle time 3-3
machine requirements
  corequisite 2-2
  standard 2-2
machine-check code bits 6-2
machine-check handling 6-2
machine-check interruption codes 6-2
machine-check interruptions 8-6
machine-readable information (MRI) X-11
main storage X-11
maintenance analysis procedure (MAP) X-11
MAP X-11
maximum allowable link loss X-11
Mb X-11
MCIC 6-2
mega (M) X-11
megabit (Mb) X-11
megabyte (MB) X-11
minimum acceptable receive level X-11
mm X-11
mode
  burst 9-10
  byte multiplexer 9-10
  data 8-2
  of operation
    ESA/390 8-3
  of operation, architectural 8-2
model upgrade progression 1-2
Model 520 1-3
Model 620 1-3
Model 640 1-3
  hardware configuration 1-3
Model 660 1-3, 7-12
  hardware configuration 1-3
Model 720 1-3
Model 740 1-3, 1-4
  hardware configuration 1-4
Model 820 1-3, 1-4, 7-12
  hardware configuration 1-4, 1-5
Model 860 1-3, 1-5, 7-12
Model 900 1-3, 1-6, 7-12
modem (modulator/demodulator) X-11
modem, 5853 2-3
monitor X-12
monitor console X-12
monitoring 7-20
move page 7-20
MP X-12
MRI X-12
ms X-12
multiple console 2-4
  console 2-4
  display 2-4
multiple high-performance guests 7-22
multiple preferred guests X-12
multiplexer channel X-12
multiplexing X-12
multiplexing, definition of 8-6
multipoint X-12
multipoint connection X-12
multipoint line X-12
multiprocessing 7-20, X-12
multiprocessor (MP) X-12
MVS dynamic physical partitioning 7-21
MVS/ESA X-12
MVS/SP 7-21, X-12
MVS/XA X-12

N

non-physical-contact connector X-12
N/A X-12

O

offline X-12
online X-12
operating system (OS) X-12
operation
  block multiplexer 9-9
  byte multiplexer 9-9
  chaining 8-5
  channel 3-12, 8-2, 8-6
  error handling, I/O 6-4
  interconnect communication element 3-12
  monitoring of 5-4
operations, retry 8-5
operator console X-12
optical fiber X-12
optical fiber cable X-12
optical fiber connector X-12
optical fiber splice X-12
optical fiber waveguide X-12

optical link X-13
optical link loss X-13
optical link segment X-13
optical mode conditioner (OMC) X-13
option X-13
output device X-13
output unit X-13

P

PA X-13
page X-13
page protection 7-21
paging X-13
parallel X-13
parallel channel X-13
parameter X-13
partition X-13
passive component X-13
password X-13
password identifier X-13
patch panel X-13
path X-13
PER X-13
PER extensions 7-21
performance
  characteristics, channel subsystem 9-2
  device 9-10
  I/O 9-10
  system 10-3
physical address X-13
physical partitioning X-13
physical-contact connector X-13
physically partitioned configuration 1-5, 1-6, 1-11, 2-3, 4-4, 5-3
  single-image 1-11
point-to-point connection X-13
point-to-point line X-13
point-to-point topology X-13
polling X-14
port address X-14
port address name X-14
port card X-14
port information block (PIB) X-14
port number X-14
power 2-4
power panel, system 5-3
power sequence control 7-19
power-on reset state X-14
power-on self-test (POST) X-14
PP X-14
PPC X-14
ppm X-14
PRB X-14
preferred machine assist 7-21
preferred virtual machine X-14
prefetch X-14
prefixing 7-20
preventive maintenance X-14
primary distribution panel X-14
primary ESCON Manager X-14
priorities
  system 10-3
priority
  device 8-7
priority, channel 10-3
private space 7-21
problem analysis
  power malfunction 6-4
  types 6-4
processor availability facility 7-21
processor complex X-14
processor controller 2-8
  corequisites
  dual processors 5-3
  function 5-4
processor resource/Systems manager feature 1-7
Processor Resource/Systems Manager (PR/SM) feature X-14
processor storage X-14
processor unit X-14
  components of 3-2
product data X-14
program temporary fix (PTF) X-14
program version X-14
program-event recording 2 (PER2) 7-27
programming
  compatibility 1-11
  support 1-11
programming RPQ (PRPQ) X-14
programs
  guest control 1-13
progression, upgrade 1-2
protective plug X-14
protocol X-14
PRPQ X-15
PR/SM X-15
PS X-15
PSB X-15
PSR X-15
PSW X-15
PSW-key handling 7-24
PU X-15
publications, related xiv
purging, broadcasted 7-11

R

RAM X-15
random access memory (RAM) X-15
RAS 1-11, 2-8, X-15
RCS X-15
REA X-15
read-only memory (ROM) X-15
real address X-15
real addressing 7-29

REC X-15
receiver X-15
reconfiguration X-15
reconnection, dynamic 8-8
recovery X-15
recovery extensions 7-24
related publications xiv
reliability (RAS) 1-11, 2-8
remote X-15
remote support facility
  authorization 5-5
  calls 5-5
  configuration 5-5
  enabling of 6-5
remote support facility (RSF) X-15
request for engineering action (REA) X-15
request for price quotation (RPQ) X-15
requirements
  cooling 2-4
  corequisite 2-2
  power 2-4
  standard requirements 2-2
reset, system 8-6
retry
  operations 8-5
RPQ X-15
RSC X-15
RSF X-15

S

SA X-15
SAD 4-4
scalar X-15
scattering X-15
SCE 3-3, B-2, X-15
SCH X-15
SCP X-15
SCS X-15
SCSW X-15
SDLC X-15
SEC X-15
secondary ESCON Manager X-15
security
  keylock 2-8
  system security 2-7
selector channel X-15
sequence, device 10-3, 10-6
service console X-16
service representative X-16
service signal 7-24
serviceability (RAS) 1-11, 2-8
set-system-mask suppression 7-28
shared main storage 7-21
SI X-16
signal interruption X-16
single channel service (SCS) X-16
single-image 1-5, 1-6, 1-11, 5-3, 7-12
single-image configuration 2-8, 4-4, X-16
single-mode optical fiber X-16
SLC X-16
soft copy X-16
sorting instructions 7-24
splice X-16
splice loss X-16
Start I/O Fast Release (SIOF) 7-29
start-I/O-fast queuing 7-29
storage X-16
  buffer 2-6, 3-8
  cache 2-6
  central 2-6, 3-8, 3-9
  element, control 3-3
  expanded 2-6, 3-11
  hierarchical 2-6
  protection, key-controlled 3-9
storage background scrubbing 7-24
storage control X-16
storage key instructions 7-25
storage subsystem X-16
storage-key instruction extensions 7-25
storage-key 4KB block 7-25
storage-reconfiguration 7-25
storage, central B-2
storage, expanded 1-10
store status 7-28
stride X-16
structure 2-2
  of interconnect communication element 3-12
subchannel 8-2, X-16
subchannel number X-16
subsystem X-16
subsystem storage X-16
subsystem storage protection 7-25
supervisor state X-16
suspend and resume 7-29
SVCH X-16
switched point-to-point channel path
  configuration X-16
switched point-to-point topology X-16
switchover X-16
synchronous X-16
SYSGEN X-16
Sysplex Timer 7-25, X-16
system
  access control 2-7
  initialization 3-12
  I/O reset 8-6
  performance 10-3
  priorities 10-3
  security 2-7
  system power panel 5-3
system activity display 4-4
system configuration X-17
system console X-17
system control element 3-3, B-2

system control element (SCE) X-17
system control programming (SCP) X-17
system definition 4-3
system display X-17
system generation (SYSGEN) X-17
system log (SYSLOG) X-17
system power panel 5-3, X-17
system reference code X-17
system reset (SYSRESET) X-17
system resource X-17
system security 2-7
System/370 extended facility 7-26

T

target logical partition X-17
target processor X-17
terminal X-17
test block 7-26
thermal conduction module (TCM) 2-8, X-17
threshold X-17
time
  channel service 10-3
  critical 10-4, 10-6
time-of-day clock 7-26
time-of-day (TOD) clock X-17
time-out X-17
time-out functions, channel
  active 8-6
  inactive 8-6
TLB 3-4
TOD X-17
TOD clock synchronization 7-21
trace X-17
tracing 7-27
translation 7-27
translation lookaside buffer 3-4
trunk cable X-17

U

UCW (unit control word) 8-2
unit control word (UCW)
  assignment 8-2
  contents 8-2
  definition 8-2
upgrade progression 1-2
user identification (user ID) X-18

V

VE X-18
VE 0-5 B-2
vector X-18
vector facility B-2, X-18
vector facility feature 1-3, 1-4, 1-5, 1-6, 1-10, 3-2, 3-6, 7-28
version codes 3-6
  Model 520 3-6
version codes *(continued)*
  Model 640 3-6
  Model 660 3-6
  Model 740 3-6
  Model 820 3-6
  Model 860 3-6
  Model 900 3-6
VF X-18
virtual address X-18
virtual machine (VM) X-18
Virtual Machine/System Product (VM/SP) X-18
virtual storage (VS) X-18
virtual-machine assist 7-28
VM X-18
VM data space facility 1-7
VM/ESA X-18
VM/SP X-18
VM/SP HPO 7-21, X-18
VM/SP-HPO 3-7, 7-13
VM/XA X-18
VM/XA SF X-18
VM/XA SP X-18

W

wrap plug X-18

Numerics

2501 9-8
2914 8-6
3033 extension 7-29
3044 10-4
31-bit indirect data address word 7-29
31-bit real addressing 7-29
3180 2-4, 4-2
3206 2-3, 2-4
3211 9-10, 9-11
3380 8-8, 9-8, 10-3, 10-4
3480 10-4
3490 10-4
3800 10-3
3814 8-6
3864 4-2
5853 Modem 2-3
9022 2-3, 2-8, 5-2
  function 5-2
9022 Processor Controller 5-1
9027 2-4
9027 Model 1 7-19
  I/O power sequence control 7-19

Readers' Comments

Enterprise System/9000
Models 520, 640, 660, 740, 820, 860, and 900
Functional Characteristics
and Configuration Guide

Publication No. GA22-7139-02

You may use this form to communicate your comments about this publication, its organization, or subject matter, with the understanding that IBM may use or distribute whatever information you supply in any way it believes appropriate without incurring any obligation to you. Your comments will be sent to the author's department for whatever review and action, if any, are deemed appropriate. We appreciate your comments.

Note: Copies of IBM publications are not stocked at the location to which this form is addressed. Please direct any requests for copies of publications to your IBM representative or to the IBM branch office serving your locality.

Machine types_____ and models_____ to which this applies.

What control programs do you use? _____

What is your occupation? _____

How did you use this publication?

☐As a reference source ☐As a classroom text ☐As a self-study text

How do you rate this publication?

☐Very Poor ☐Poor ☐Average ☐Good ☐Very Good

What sections or figures were particularly useful or understandable to you?

What sections or figures could be improved?

What sections or figures require added information?

Any other comments?

Note: You can mail this form or FAX it to **1-914-432-9405**.

Name _____ Address _____

Company or Organization _____ _____

Phone No. _____ _____

Readers' Comments
GA22-7139-02

IBM®

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 B97
PO BOX 950
POUGHKEEPSIE NY 12602-0950

Fold and Tape | **Please do not staple** | Fold and Tape

Cut or Fold
Along Line

File Number: 9021-02

Printed in U.S.A.

For questions or problems regarding this publication:

In the **U.S.A., Canada, and Puerto Rico**, call

1-800-426-9021

In **all countries**, either mail the Readers' Comment Form or FAX the Readers' Comment Form to

1-914-432-9405 **(FAX number)**

Note: In countries other than the U.S.A., Canada, and Puerto Rico, dial the appropriate international access code before the above FAX number.

GA22-7139-02