F7-GA-1C01

F-7C AND F-7E PROCESSOR GROUP FUNCTIONAL CHARACTERISTICS

HITACHI

Statement on EN55022 Compliance

WARNING: This is a Class A product. In a domestic environment this product may cause radio interference in which case the user may be required to take adequate measures.

Statement on Federal Communications Commission (FCC) Compliance

NOTE: This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference in which case the user will be required to correct the interference at his/her own expense. The user is cautioned that changes or modifications not expressly approved by the manufacturer could void the user's authority to operate the equipment.

Trademarks

F7-GA-1C01 Third Edition (March 1999)

This edition obsoletes F7-GA-1C01 second edition. Although this second edition provides the latest information available, such information is subject to change without prior notice. The reader is responsible for maintaining up-to-date editions.

Change Record

Revision No.	*Date*	*Description*	*Affected Pages*
0	September 1998	Original Edition	—
1	December 1998	Second Edition • Additions and changes associated with introduction of F-7C Turbo Models. • Changes associated with introduction of 31CD, 52CD, and B2C Base Models, and 0XC (B2C) CF Model. • Changes associated with new CHPID assignments. • Changes in increments of MS and ES. • Addition of note on available channel numbers. • Changes in external views of F-7C models. • Changes in F-7C basic logical components. • Editorial changes.	All.
2	March 1999	Third Edition • Additions and changes associated with introduction of F-7E Turbo Models. • Addition of information relating to Integrated Disk. • Addition of information on Virtual Server Phase-1.	All.

PREFACE

This document serves as a functional characteristics manual draft for the F-7C and F-7E Processor Group.

For the purpose of generating a user manual, Hitachi, Ltd. allows the OEM customers to reproduce, rewrite, modify and edit all or part of this document for their use for the assumed OEM marketing/sales of the F-7C and F-7E Processor Group while retaining copyright and all rights associated with copyright.

This document describes primarily the functional characteristics and configuration of the F-7C and F-7E Processor Group hardware, without particular regard to how it is utilized by software. It should be understood, therefore, that not all the hardware capabilities may be supported by the target systems. It should be also noted that the breakdown models given here do not necessarily match the models defined for ordering in the OEM Sales Agreement; for example, models transparent to the end users are omitted, renamed or treated as included in appropriate host features.

Information contained in this document is subject to change without notice.

Reference manuals

- *F-7C and F-7E Processor Group Installation Planning Guide* (F7-GA-1C02)
- *Hitachi MLPF 3.3.3 User Manual* (HT-GA-1012)
- *Hitachi MLPF 3.5.0 User Manual* (HT-GA-1015)
- *Hitachi HCCFA 1.0–3.1 User Manual* (HT-GA-1020)
- *Hitachi HCCFA 4.0 User Manual* (HT-GA-1025)
- *IBM System/360 and System/370 I/O Interface Channel to Control Unit Original Equipment Manufacturers' Information* (GA22-6974)
- *IBM System/370 Extended Architecture Principles of Operation* (SA22-7085)
- *IBM System/370 Extended Architecture Interpretive Execution* (SA22-7095)
- *IBM Enterprise Systems Architecture/390 Principles of Operation* (SA22-7201)
- *IBM Enterprise Systems Architecture/390 ESCON I/O Interface* (SA22-7202)
- *IBM Enterprise Systems Architecture/390 ESCON Channel to Channel Adapter* (SA22-7203)
- *IBM Enterprise Systems Architecture/390 Common I/O Device Commands* (SA22-7204)
- *IBM ES/9000 Input/Output Configuration Program User's Guide and ESCON Channel-to-Channel Reference* (GC38-0401)

TABLE OF CONTENTS

CHAPTER 4 CHANNEL SUBSYSTEM

CHAPTER 5 SERVICE PROCESSOR AND CONSOLE DEVICE

CHAPTER 6 SUPPORTED FEATURES AND FACILITIES

CHAPTER 7 COOLANT DISTRIBUTION UNIT

LIST OF FIGURES

LIST OF TABLES

LIST OF ABBREVIATIONS

AA	Access Administrator (user level)
ADMF	Asynchronous Data Mover Facility
AFT	ASN first table
AIX	Advanced Interactive Executive
ALB	ART Lookaside Buffer
AO	Advanced Operator (user level)
AP	Alternate Processor
ART	access register translation
ASN	address-space number
AST	ASN second table
ATM	asynchronous transfer mode
BC-mode	Basic Control mode
BL	block-multiplexer (channel)
BPU	Basic Processing Unit
BS	buffer storage
BTA	1,2,3-Benzotriazole
BY	byte-multiplexer (channel)
CCW	channel command word
CD	Console Device
CDU	Coolant Distribution Unit
CE	Customer Engineer (user level)
CF	coupling facility
CFC	chlorofluorocarbon
CFR	coupling facility receiver (channel)
CFS	coupling facility sender (channel)
CH	channel
CHP	Channel Processor
CHPID	channel path ID
CH-PK	channel package
CICS	Customer Information Control System
CM	concurrent maintenance
CMOS	complementary metal oxide semiconductor
CNC	connection channel
CP	Central Processor; Control Program
CPN	Central Processor Node
CPU	Central Processing Unit
CTC	channel-to-channel
CU	control unit
CVC	converter channel
DASD	direct access storage device
DAT	dynamic address translation
DB2	Data Base 2
DCS	dynamic channel subsystem
DFSORT	Data Facility Sort
DL/I	Data Language I
dpi	dot per inch
DRM	Dynamic I/O Reconfiguration Management
DSP	Disk System Processor
DSR	Dynamic Storage Reconfiguration
EC-mode	Extended Control mode
EHSA	Extended Hardware System Area
EMIF	ESCON Multiple Image Facility
EN	European Norm
ES	expanded storage
ESA	Enterprise Systems Architecture
ESCH	Extended Serial Channel
ESCON	Enterprise Systems Connection
ETAF	External Timer Attachment Feature
Ex-1	Expansion Frame
FAR	floating address register
FBCH	Fibre Channel
FCC	Federal Communications Commission
FCP	Fibre Channel Protocol
FDDI	fiber distributed data interface
FRR	Functional Recovery Routine
FRU	field-replaceable unit
GB	gigabyte(s)
GUI	graphical user interface
HCCF	Highspeed Coupling Control Feature
HCCFA	Highspeed Coupling Control Feature Assist
HCD	Hardware Configuration Definition
HDD	hard disk drive
HDM	High-Density Module
HFC	hydrofluorocarbon
HMC	Hardware Management Console
HSA	Hardware System Area
IBM	International Business Machines Corporation
ICMF	Integrated Coupling Migration Facility
ID	identifier
IDAW	Indirect-Data-Address Word
IDK	Integrated Disk
IEF	Interpretive Execution Facility
IML	initial microcode loading
IMS	Information Management System
IntCF	Integrated Coupling Facility
I/O	input/output
IOCDS	I/O Configuration Data Set
IOCP	I/O Configuration Program
IOP	Input Output Processor
IP	Instruction Processor
IPL	initial program load
ISCH	Inter-System Coupling Channel
ISCH2	Inter-System Coupling Channel 2
KB	kilobyte(s)
LAN	local area network
LAP	low address protection
LCU	logical control unit

LCUW	logical control unit word
LED	light-emitting diode
LPAR	Logically-Partitioned (mode); logical partition
LRU	least recently used
LSI	large-scale integration
MB	megabyte(s)
MCD	Management Console Device
MHPGS	Multiple High-Performance Guest Support
MLPF	Multiple Logical Processor Feature
MO	magneto-optical
MS	main storage
MVPG	Move Page (Facility 1)
MVPG-2	Move Page Facility 2
MVS	Multiple Virtual Storage
NEMA	National Electrical Manufacturers Association
OCF	Open Systems Connection Feature
OEM	original equipment manufacturer; original equipment manufacturing
Op	Operator (user level)
OS	operating system(s)
OSA	Open Systems Adapter
OSCH	Open Systems Connection Channel
OS/390	Operating System/390
PAF	Processor Availability Facility
PC	personal computer
PCI	power control interface
PCU	physical control unit
PE	processor element
PER	Program Event Recording
PK	package
PMC	Processor Management Console
PMU	Power Management Unit
ppm	parts per million
PSA	Prefixed Storage Area
PSTR	Processor Storage
PSW	program status word
RAM	random access memory
RAS	reliability, availability and serviceability
RCP	real control program
RFU	Refrigeration Unit
RMF	Resource Measurement Facility
RPC	Remote PCI Controller
SACF	Set Address Space Control Fast
SAP	System Assist Processor
SC	System Controller
SCD	Service Console Device
SCP	System Control Program
SDC	self-describing component
SIE	start interpretive execution
SP	System Processor; System Programmer (user level)
SPOC	single-point-of-control
SSSP	Subsystem Storage Protection
SVC	Supervisor Call
SVGA	Super Video Graphics Array
SVP	Service Processor
SYSIML	system initial microprogram loading
Sysplex	System Complex
S/370	System/370
S/390	System/390
TCP/IP	Transmission Control Protocol/Internet Protocol
TIC	transfer-in-channel (command)
TLB	Translation Lookaside Buffer
TOD	time-of-day
TSCF	Target System Control Facility
UA	unit address
UCW	unit control word
VLSI	very large-scale integration
VM	Virtual Machine (operating system)
VMA	Virtual Machine Assist
VSE	Virtual Storage Extended
VSF	Virtual Server Facility
WAN	wide area network
WS	work storage
XA	Extended Architecture

Note: Common abbreviations and instruction mnemonics are excluded.

CHAPTER 1 F-7C AND F-7E PROCESSOR GROUP OVERVIEW

1.1 Introduction

The F-7C and F-7E Processor Group is a group of processors based on complementary metal oxide semiconductor (CMOS) technology featured with cost-efficient capacity, flexibility, and processing power in a small footprint. The F-7C and F-7E Processor Group consists of the F-7C Base Models, the F-7C Turbo Models (performance-boosted models of F-7C), the F-7C CF Models (coupling facility models), and the F-7E Turbo Models (performance-boosted models beyond F-7C). The F-7C Base Models, the F-7C Turbo Models, and the F-7E Turbo Models support ESA/390 and S/390 industry standards and an extensive range of system software scenarios. All these models function in nonparallel or Parallel Sysplex environments, operating in either Basic or LPAR (logically-partitioned) mode. Each F-7C Base Model, each F-7C Turbo Model, and each F-7E Turbo Model also support various types of industry-standard channels—ESCON, parallel, coupling link, and OCF— in separate, intermixed, or shared configurations. By contrast, the F-7C CF Models act as coupling facilities in Parallel Sysplex environments, operating in LPAR mode. Each F-7C CF Model utilizes the same hardware technologies as the F-7C Base Models, and supports coupling link channels in separate, intermixed, or shared configurations.

The F-7C and F-7E processor lines are progressively scaled so that each succeeding processor model provides greater processing power while retaining many configuration possibilities. This scalability assures that upgrades can be easily sized for a particular situation and designed for minimal disruption.

1.2 Model Grouping

Depending on the purpose and the hardware implementation represented by the number of System Controllers (SCs), the F-7C Processor Group provides 1SC Base Models, 2SC Base Models, 2SC Turbo Models, 1SC CF Models, and 2SC CF Models.

The F-7E Processor Group further extends the said scalability with performance-boosted 2SC Turbo Models.

Each F-7C Base Model uses a three- or four-character model ID, where the first character shows the number of [regular] Instruction Processors (IPs), the second character shows the number of SCs, the third character identifies the processor series “C” (fixed), and the fourth character indicates degraded or full performance: D degraded, blank full. Similarly, each F-7C Turbo Model and each F-7E Turbo Model use three- or four-character model ID, where the first character shows the number of [regular] IPs, the second character “2” (fixed) shows the number of SCs, the third character “C” or “E” identifies the processor series, and, where applicable, the fourth character “H” (fixed) indicates the processor module hardware type of boosted performance. For the F-7C CF Models, “0XC” is fixed to identify the CF Model, followed by three characters in parenthesis: the first character shows the number of Integrated Coupling Facility (IntCF) IPs, the second character shows the number of SCs, and the third character identifies the processor series “C” (fixed).

Figure 1-1 shows the F-7C and F-7E model lineup. Tables 1-1 through 1-4 show determination of specific processor models in the F-7C and F-7E lines.

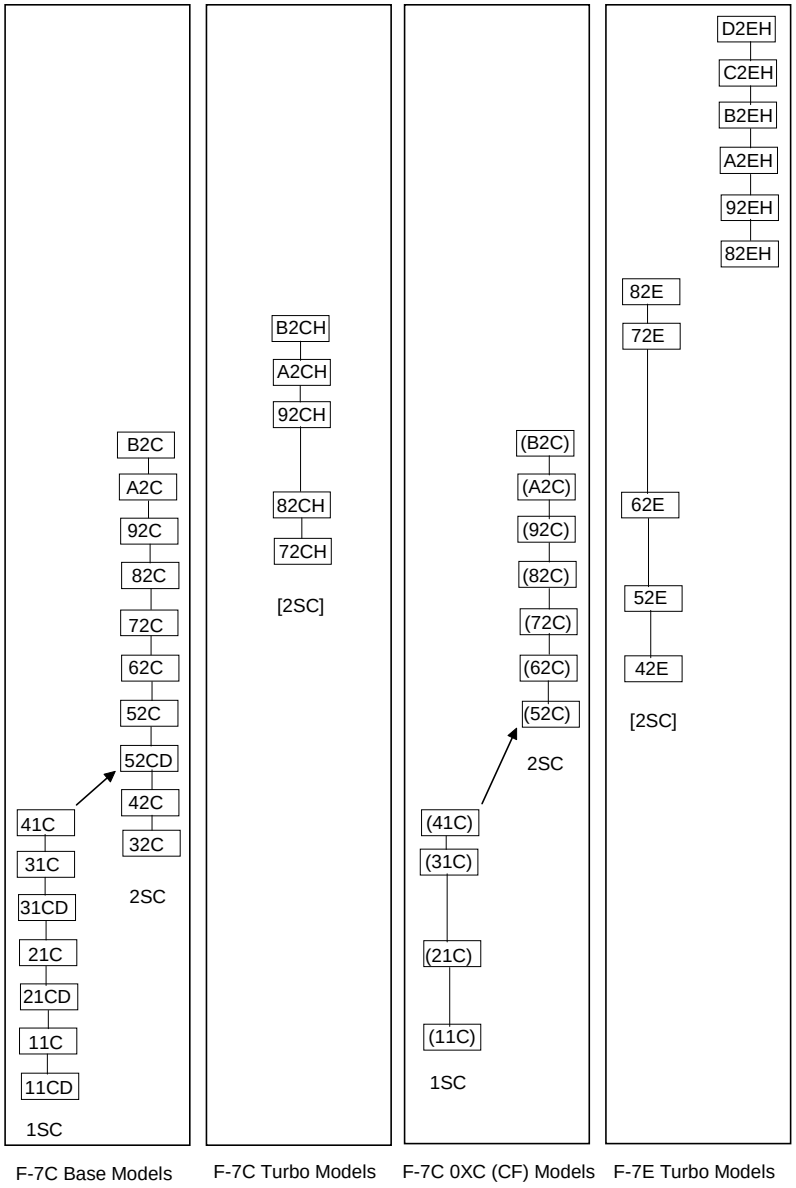

Figure 1-1 F-7C and F-7E Model Lineup

Table 1-1 Determination of F-7C Base Models

Model		*No. of [Regular] IPs*	*No. of SCs*	*No. of APs*				*9672, R6 Ver. Code*
				Total	*IntCF IPs*	*DSPs*	*Standby APs*	
1SC Base	11CD	1	1	1	0–1	0–1	Δ	8F
				4*	0–4*	0–2*	Δ*	8F
	11C	1	1	1	0–1	0–1	Δ	81
				4*	0–4*	0–2*	Δ*	81
	21CD	2	1	3	0–3	0–2	Δ	8E
	21C	2	1	3	0–3	0–2	Δ	82
	31CD	3	1	2	0–2	0–2	Δ	03
	31C	3	1	2	0–2	0–2	Δ	8D
	41C	4	1	1	0–1	0–1	Δ	8C
2SC Base	32C	3	2	2	0–2	0–2	Δ	83
				5*	0–5*	0–2*	Δ*	83
				8*	0–8*	0–2*	Δ*	83
	42C	4	2	1	0–1	0–1	Δ	84
				4*	0–4*	0–2*	Δ*	84
				7*	0–7*	0–2*	Δ*	84
	52CD	5	2	3	0–3	0–2	Δ	05
				6*	0–6*	0–2*	Δ*	05
	52C	5	2	3	0–3	0–2	Δ	85
				6*	0–6*	0–2*	Δ*	85
	62C	6	2	2	0–2	0–2	Δ	86
				5*	0–5*	0–2*	Δ*	86
	72C	7	2	1	0–1	0–1	Δ	87
				4*	0–4*	0–2*	Δ*	87
	82C	8	2	3	0–3	0–2	Δ	88
	92C	9	2	2	0–2	0–2	Δ	89
	A2C	10	2	1	0–1	0–1	Δ	8A
	B2C	11	2	0	0	NA	0	0B

Δ Total minus IntCF IPs and DSPs. * Optional configuration.

Table 1-2 Determination of F-7C Turbo Models

Model		*No. of [Regular] IPs*	*No. of SCs*	*No. of APs*			*9672, R6 Ver. Code*
				Total	*IntCF IPs*	*Standby APs*	
[2SC] Turbo	72CH	7	2	4	0–4	Δ	4D
	82CH	8	2	3	0–3	Δ	4C
	92CH	9	2	2	0–2	Δ	4B
	A2CH	10	2	1	0–1	Δ	8B
	B2CH	11	2	0	0	0	0F

Δ Total minus IntCF IPs.

Table 1-3 Determination of F-7C CF Models

Model		*No. of [Regular] IPs*	*No. of SCs*	*No. of APs*			*9672, R6 Ver. Code*
				Total	*IntCF IPs*	*Standby APs*	
1SC CF	0XC (11C)	0	1	2	1	1	80
	0XC (21C)	0	1	5	2	3	80
	0XC (31C)	0	1	5	3	2	80
	0XC (41C)	0	1	5	4	1	80
2SC CF	0XC (52C)	0	2	8	5	3	80
	0XC (62C)	0	2	8	6	2	80
	0XC (72C)	0	2	8	7	1	80
	0XC (82C)	0	2	11	8	3	80
	0XC (92C)	0	2	11	9	2	80
	0XC (A2C)	0	2	11	10	1	80
	0XC (B2C)	0	2	11	11	0	80

Table 1-4 Determination of F-7E Turbo Models

Model		*No. of [Regular] IPs*	*No. of SCs*	*No. of APs*				*9672, R6 Ver. Code*
				Total	*IntCF IPs*	*DSPs*	*Standby APs*	
[2SC] Turbo	42E	4	2	2	0–2	0–2	Δ	TBD
	52E	5	2	1	0–1	0–1	Δ	TBD
	62E	6	2	3	0–3	0–2	Δ	TBD
	72E	7	2	2	0–2	0–2	Δ	TBD
	82E	8	2	1	0–1	NA	Δ	TBD
	82EH	8	2	3	0–3	0–2	Δ	TBD
	92EH	9	2	2	0–2	0–2	Δ	TBD
	A2EH	10	2	1	0–1	0–1	Δ	TBD
	B2EH	11	2	2	0–2	0–2	Δ	TBD
	C2EH	12	2	1	0–1	0–1	Δ	TBD
	D2EH	13	2	0	0	NA	0	TBD

Δ Total minus IntCF IPs and DSPs.

1.3 Benefits

The F-7C and F-7E Processor Group provides an extensive range of functionality and processing scale. The general benefits are explained in the following subsections.

1.3.1 Leading-Edge Technologies

The F-7C and F-7E processor models utilize high-density modules, advanced logic design based on CMOS, and other technologies to enhance the overall performance.

1.3.2 Excellent Performance Efficiencies

The use of CMOS with other enhancing technologies, streamlined component design, and advanced diagnostics provide the following performance benefits:

- Lower operating costs
- Greatly-reduced footprints
- Noticeably improved resource usage
- Enhanced reliability, availability, and serviceability (RAS)
- Improved results in better price-to-performance ratios

1.3.3 Full Compatibility

All the F-7C Base Models, F-7C Turbo Models, and F-7E Turbo Models support full S/390 industry standards and ESA/390 architecture in parallel and non-parallel environments. All the F-7C CF Models support S/390-compatible coupling facility (CF) functions in parallel environments. Multiple Logical Processor Feature supports up to 15 LPARs on all models of the F-7C and F-7E Processor Group. Each LPAR in turn supports ESA/390 architecture or S/390-compatible CF functions on all the F-7C Base Models, F-7C Turbo Models, and F-7E Turbo Models, and S/390-compatible CF functions on all the F-7C CF Models.

1.3.4 Balanced Scalability

Scaling provides the ability to create subtle and not-so-subtle configuration and performance differences within the same model. When extended among the wide range of the F-7C and F-7E lines, scaling simplifies upgrading from one model to another. In both instances, scaling provides configuration flexibility and compatibility which simplifies meeting changing business requirements.

1.3.5 Optimized Reliability, Availability, and Serviceability

The F-7C and F-7E processor models use minimal physical components and connection points to simplify the design and reduce the chance of error. The use of CMOS chips and high-density semiconductor packaging creates a tough superstructure that withstands excessive physical stress and helps minimize disruption due to physical damage.

Additionally, the F-7C and F-7E processor models are backed by comprehensive professional services and technical support which includes online diagnostics and other nondisruptive maintenance techniques.

1.4 Physical Design Highlights

The F-7C and F-7E processor models each integrate the Basic Processing Unit (BPU), the Service Processor (SVP), power supplies, and one standard Input Output Processor (IOP) and one optional IOP into a single cabinet called Basic Frame. The BPU contains all the hardware for the Instruction Processor (IP), System Processor (SP), Processor Storage (PSTR), and System Controller (SC).

One Expansion Frame (Ex-1) may be needed to house up to two more units of optional IOP. The necessity of the Ex-1 depends upon the number of channels to be installed and their channel path IDs to be assigned. The Ex-1 incorporates cooling fans and power supplies.

Up to four Extended Frames are optionally installable for support of the Integrated Disk (IDK) on the F-7C Base Models except B2C models and on the F-7E Turbo Models except 82E and D2EH models. One Extended Frame accommodates two IDKs.

One Coolant Distribution Unit (CDU) Frame is mandatory on each F-7C Turbo Model and F-7E Turbo Model to achieve its performance boost through reinforced cooling of processor elements. The CDU Frame incorporates power supplies and dualized Refrigeration Units (RFUs).

One through three Console Devices (CDs) and up to four optional Remote PCI Controllers (RPCs) are each housed in separate cabinets.

Table 1-5 shows the number of frames and components of the F-7C and F-7E processor models. Figures 1-2 through 1-5 illustrate the external views and frame contents of the F-7C processor models.

Table 1-5 Numbers of Frames and Components by Processor Models

Frame or Component	*Q'ty on F-7C*					*Q'ty on F-7E*
	1SC Base	*2SC Base*	*[2SC] Turbo*	*1SC CF*	*2SC CF*	*[2SC] Turbo*
Basic Frame	1	1	1	1	1	1
BPU	1	1	1	1	1	1
IOP	1–2	1–2	1–2	1–2	1–2	1–2
SVP	1	1	1	1	1	1
Ex-1 (option)	0–1	0–1	0–1	0–1	0–1	0–1
IOP (option)	0–2	0–2	0–2	0–2	0–2	0–2
Extended Frame (option)	0–4	0–4	NA	NA	NA	0–4
IDK (option)	0–8	0–8	NA	NA	NA	0–8
CDU Frame	NA	NA	1	NA	NA	1
CD	1–3	1–3	1–3	1–3	1–3	1–3
RPC (option)	0–4	0–4	0–4	NA	NA	0–4

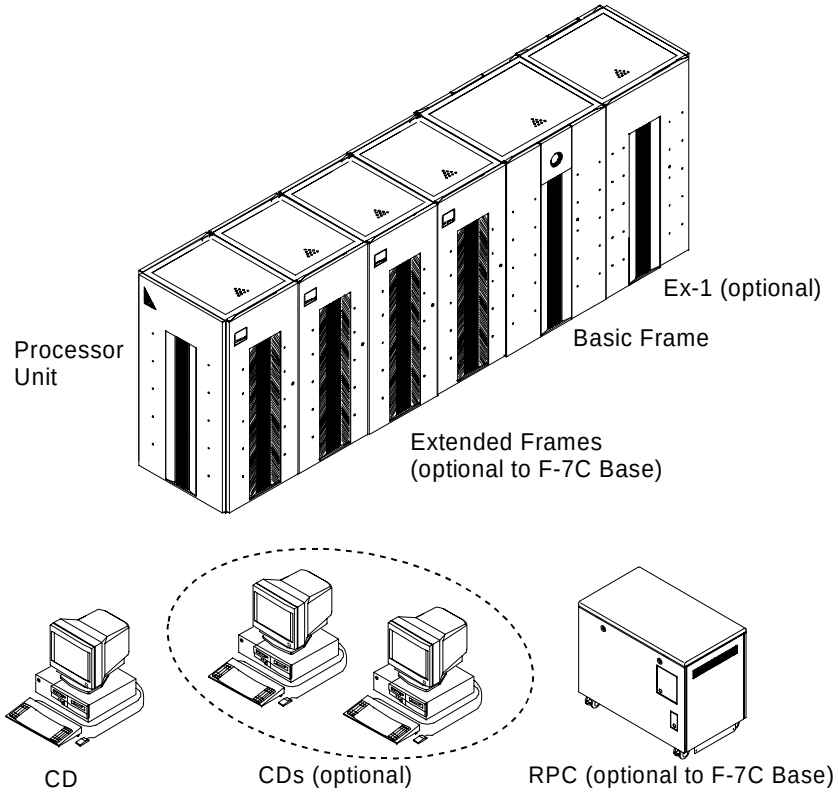

Figure 1-2 External View of F-7C Base and F-7C CF

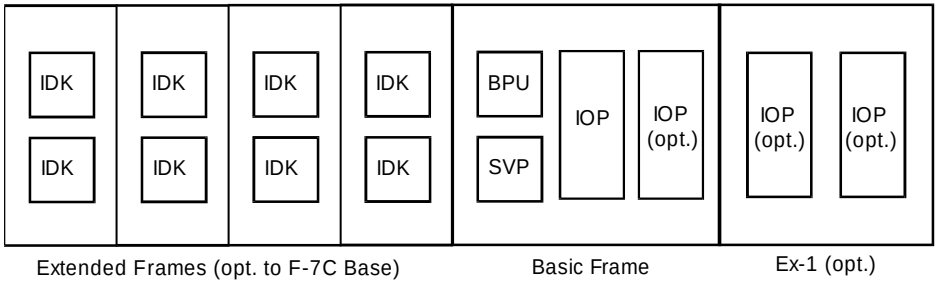

Legend:
BPU = Basic Processing Unit
Ex-1 = Expansion Frame
IDK = Integrated Disk
IOP = Input Output Processor
SVP = Service Processor

Figure 1-3 Processor Unit Frame Contents of F-7C Base and F-7C CF

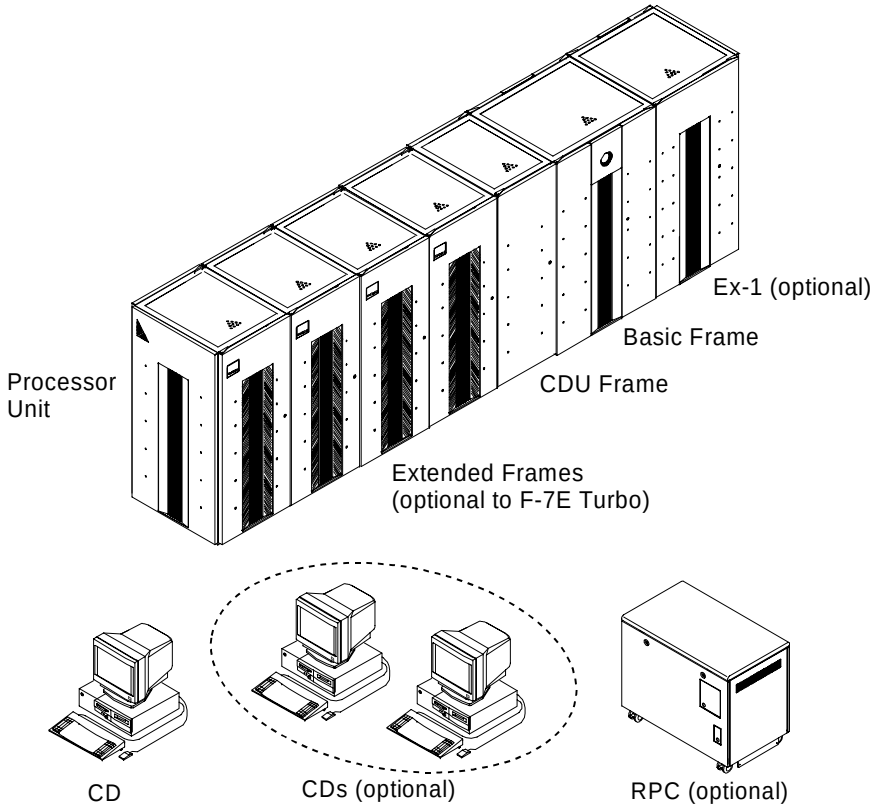

Figure 1-4 External View of F-7C Turbo and F-7E Turbo

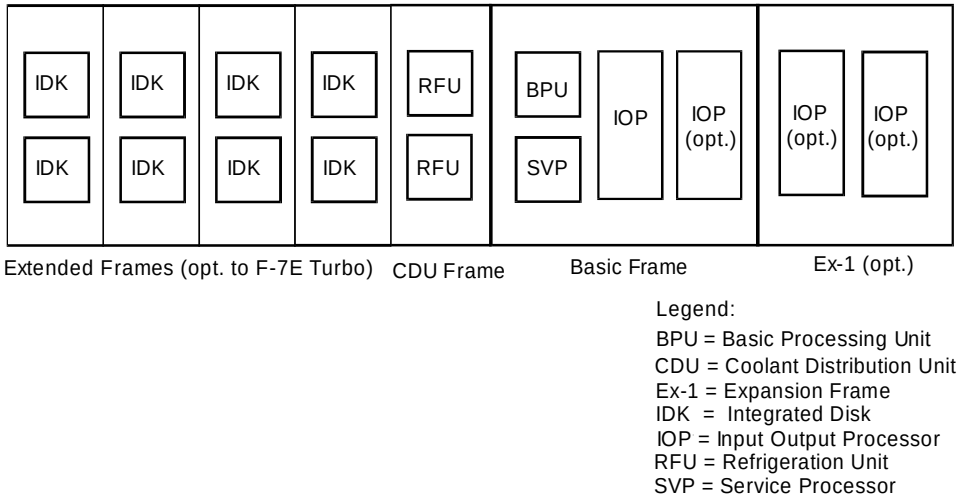

Figure 1-5 Processor Unit Frame Contents of F-7C Turbo and F-7E Turbo

1.5 Configuration

Figure 1-6 conceptually illustrates how the basic elements of the F-7C and F-7E processor models are configured in relationship with each other. Note that each box does not represent a single distinct physical component, but rather a logical component that performs a general function using an aggregate of physical components.

Figures 1-7 through 1-10 illustrate possible configurations for the F-7C and F-7E processor models. Following these figures are subsections 1.5.1 through 1.5.6 that provide description of each component and summary of the processor configurations.

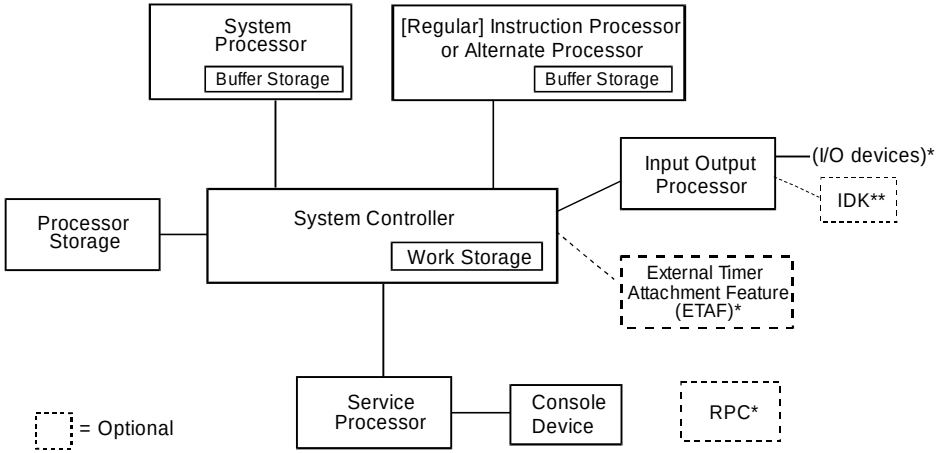

* Not applicable to F-7C CF Models.
** Only applicable to F-7C Base Models other than B2C and F-7E Turbo Models other than 82E and D2EH.

Figure 1-6 F-7C and F-7E Basic Logical Components

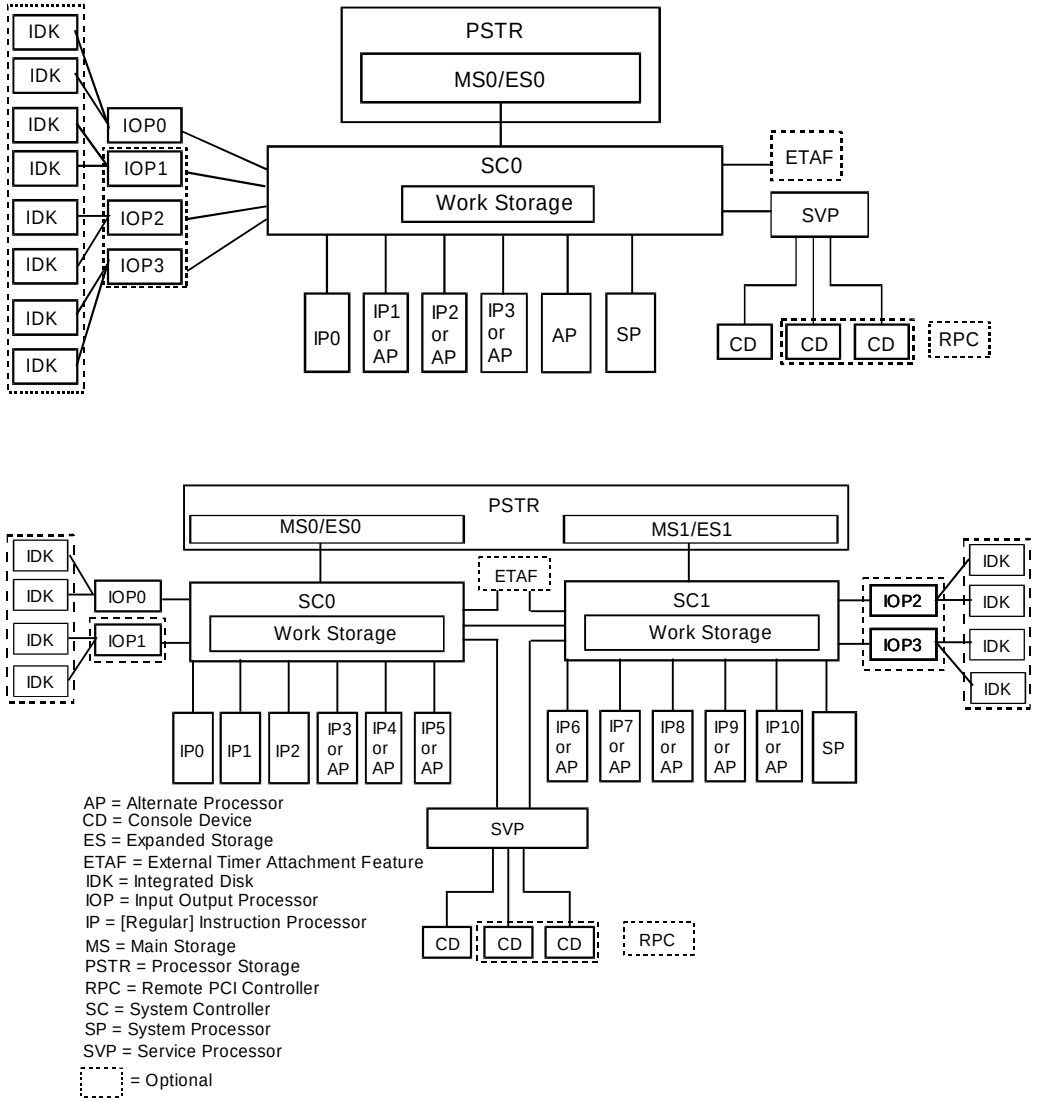

Figure 1-7 Configuration of F-7C 1SC Base Models and 2SC Base Models

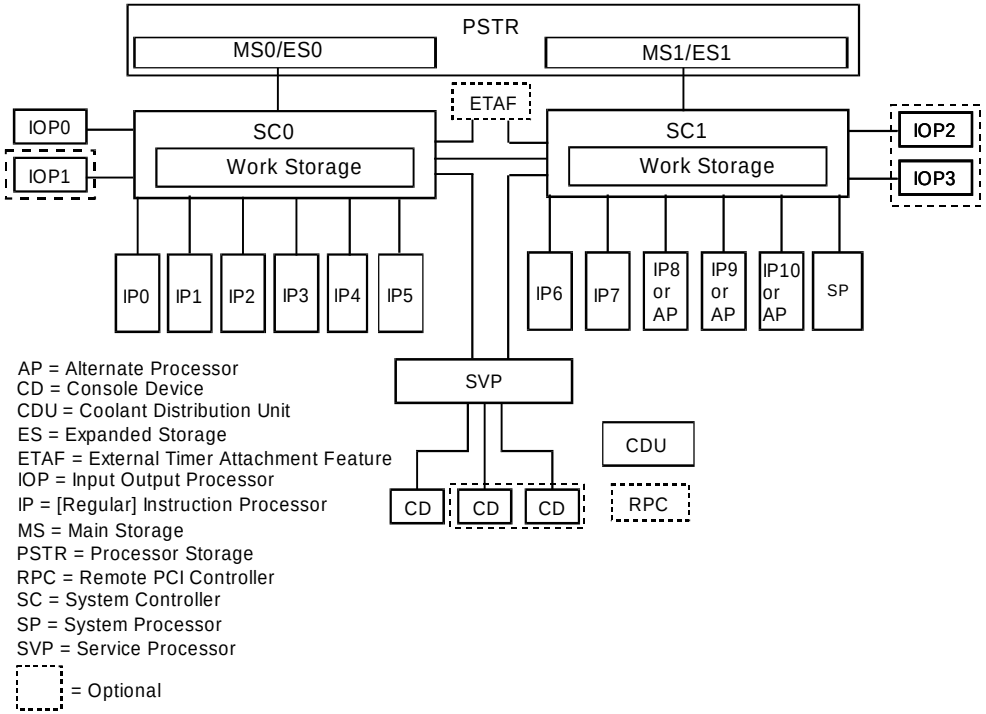

Figure 1-8 Configuration of F-7C 2SC Turbo Models

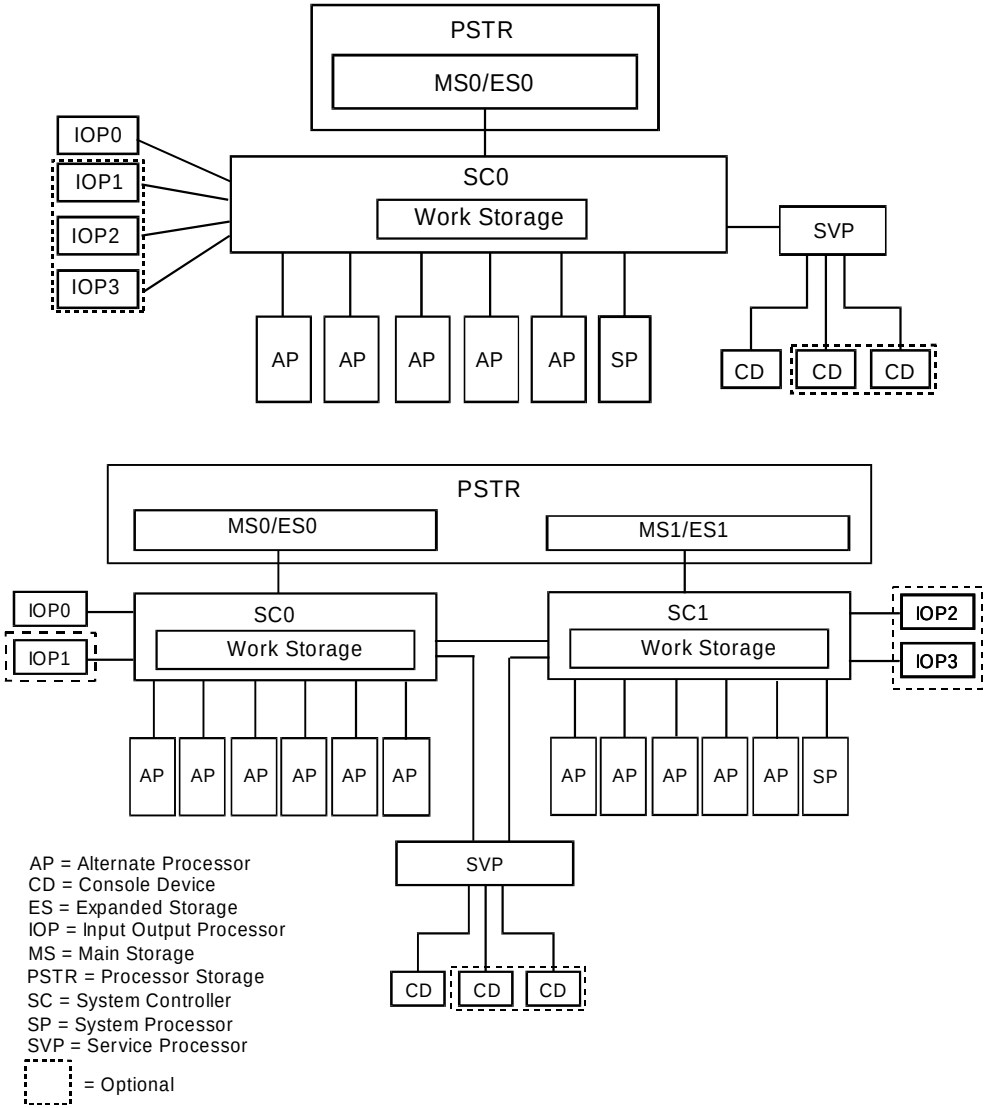

Figure 1-9 Configuration of F-7C 1SC CF Models and 2SC CF Models

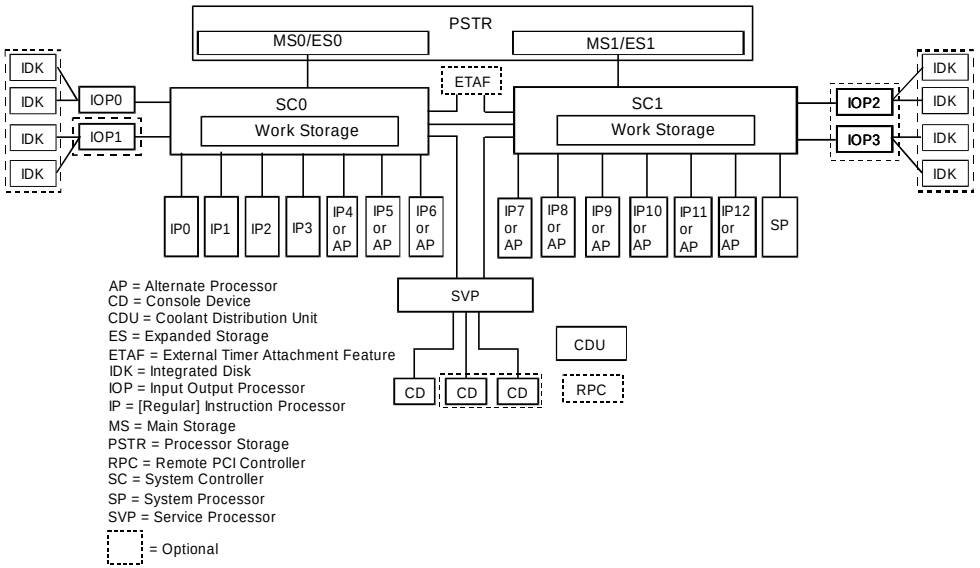

Figure 1-10 Configuration of F-7E 2SC Turbo Models

1.5.1 Instruction Processor and System Processor

The Instruction Processors (IPs) and System Processor (SP) manage the execution of program instructions and the architectural facilities for their support. The F-7C Base Models, F-7C Turbo Models, and F-7E Turbo Models support all instruction formats and instructions defined in IBM's *Enterprise Systems Architecture/390 Principles of Operation*.

Each processor is an aggregate of physically-associated processing circuitry and cache memory arrays. Each F-7C and F-7E processor model has all of its processors and associated cache in a single High-Density Module (HDM).

Two levels of cache (high speed buffer storage) reside between the IPs (and SP) and the main storage: Buffer Storage (BS) and Work Storage (WS). Cache improves processor performance by reducing the effects of delays encountered when a process fetches data from main storage.

1.5.2 System Controller and Processor Storage

The System Controller (SC) manages access to all of the processor's physical storage called processor storage. Processor storage provides all the storage capacity for each processor complex. The F-7C and F-7E models partition processor storage into main storage and optional expanded storage.

The basic addressable storage unit in main storage is the 8-bit byte. This byte can be used as a single hexadecimal character, as 2 decimal digits, or as 8 binary digits. All architecturally-defined data formats are supported.

On the F-7C 1SC Base Models and F-7C 1SC CF Models, the maximum amount of installable processor storage is 16 GB. On the F-7C 2SC Base Models, F-7C Turbo Models, F-7C 2SC CF Models, and F-7E Turbo Models, the maximum amount of installable processor storage is 32 GB.

In Basic (non-LPAR) mode running under ESA/390, main storage is limited to 2 GB with all the remaining processor storage assignable to expanded storage.

In LPAR mode running under Multiple Logical Processor Feature, the F-7C and F-7E models (both 1SC and 2SC) can partition up to 16 GB of processor storage as main storage to be sub-allocated to LPARs. Then, each LPAR is limited to 2 GB of main storage.

1.5.3 Input/Output Device Management

The Input Output Processor (IOP) manages channels to which are attached all I/O devices except the Console Device (CD). The Service Processor (SVP) manages all CDs.

The channel subsystem contains IOPs, the channels they support, and the IOP-associated logic residing in the IPs. This subsystem manages all I/O activity independent of the IPs to allow the channels and processors to work simultaneously with minimal wait for I/O operations to complete.

The channel subsystem supports the following combinations of channels depending on the processor models:

- F-7C Base: ESCH (ESCON), parallel, ISCH2 (coupling link), Open Systems Connection Feature (OCF) channels, and Fibre Channel (FBCH) of Fibre Channel Protocol (FCP) channel type*
- F-7C Turbo: ESCH (ESCON), parallel, ISCH2 (coupling link), and Open Systems Connection Feature (OCF) channels
- F-7C CF: ISCH2 channels
- F-7E Turbo: ESCH (ESCON), parallel, ISCH2 (coupling link), Open Systems Connection Feature (OCF) channels, and Fibre Channel (FBCH) of Fibre Channel Protocol (FCP) channel type*

* Dedicated to Integrated Disk (IDK).

1.5.4 SVP/CD Subsystem

The SVP/CD subsystem contains the Service Processor (SVP) which supports one to three CDs and the optional interfaces to the Processor Management Console (PMC) and Hardware Management Console (HMC). These consoles can also back up a regular operator console when needed.

1.5.5 Optional Components

The following are optional, where "*" indicates "NA on F-7C CF Models" and "**" indicates "only applicable to F-7C Base Models and F-7E Turbo Models":

- Processor storage
- Channels: ESCON-compatible ESCH (multimode duplex)*, Parallel*, OCF*, FBCH (FCP)**, and ISCH2 coupling link (single-mode duplex)
- Remote PCI Controller*
- Additional Console Devices
- External Timer Attachment Feature (ETAF)*
- Integrated Disk (IDK)**

1.5.6 Processor Configuration Summary

Table 1-6 provides processor configuration summary of the F-7C and F-7E processor models.

Table 1-6 Processor Configuration Summary of F-7C and F-7E Processor Models

Component			*Q'ty on F-7C*					*Q'ty on F-7E*
			1SC Base	*2SC Base*	*[2SC] Turbo*	*1SC CF*	*2SC CF*	*[2SC] Turbo*
Processor elements	[Regular] IP		1–4	3–11	7–11	0	0	4–13
	SP		1	1	1	1	1	1
	SC		1	2	2	1	2	2
	AP	Total	1–4	0–8	0–4	2–5	8–11	0–3
		IntCF IP	0–3 (0–4*)	0–3 (0–8*)	0–4	1–4	5–11	0–3
		DSP	0–2	0–2	–	–	–	0–2
		Standby AP	Δ	Δ	Δ	Δ	Δ	Δ
Max. PSTR capacity (GB)			16	32	32	16	32	32
ETAF			0–1	0–1	0–1	–	–	0–1
IOP			1–4	1–4	1–4	1–4	1–4	1–4
Channel	Parallel		0–96	0–96	0–96	–	–	0–96
	ESCH		0–256	0–256	0–256	–	–	0–256
	ISCH2		0–32	0–32	0–32	0–32	0–32	0–32
	OCF		0–12	0–12	0–12	–	–	0–12
	FBCH (FCP)		0–16	0–16	–	–	–	0–16
	Total		8–256	8–256	8–256	2–32	2–32	8–256
CDU			–	–	1	–	–	1
SVP and CD	SVP		1	1	1	1	1	1
	CD		1–3	1–3	1–3	1–3	1–3	1–3
RPC			0–4	0–4	0–4	–	–	0–4

* Optional configuration. Δ Total minus IntCF IPs and DSPs.

1.6 System Design Highlights

1.6.1 Operating System Support

The following are supported by the F-7C Base Models, F-7C Turbo Models, and F-7E Turbo Models:

- OS/390
- MVS/ESA 4.3, 5.1, 5.2 and subsequent releases
- VM/ESA 1.2.2, 2.1, 2.2, 2.3
- VSE/ESA 1.4, 2.1, 2.2, 2.3

1.6.2 Advanced Logic Design

The F-7C and F-7E processor models enhance overall performance using a design that optimizes combined and individual aspects when integrating components. Techniques and technologies used include:

- High-speed system cycles and execution
- High-speed/high-density IOPs and channels built on CMOS LSI
- Execution pipeline controls
- Multilevel storage hierarchy

1.6.3 CPU ID

The F-7C and F-7E processor models generate CPU ID information which is presented to software through the use of STORE CPU ID (STIDP) instruction. The 8-byte data area returned by STIDP contains a version code which identifies the processor model, a CPU identification number which is unique to each IP, a model number (which identifies the entire F-7C and F-7E processor models as "9672"), and 2 bytes of hexadecimal zero padding on the right. STIDP returns this data in the following format:

Byte	0	1	4	6
	Version Code	CPU Identification Number	Model Number	0000

The format of the CPU Identification Number field is different when the system runs in Basic mode and when it runs in LPAR mode.

In Basic mode, the CPU Identification Number identifies the physical IP by the following format:

Byte	1
	ISSSSS

In LPAR mode, the CPU Identification Number identifies the logical IP by the following format:

Byte	1
	LPSSSS

Each field of the CPU ID information contains the following:

- Version Code: Two hexadecimal digits identifying the processor model. See section 1.2, tables 1-1 through 1-4.
- CPU Identification Number
  - I: Physical CPU address as 1 hexadecimal digit.
  - L: Logical CPU address as 1 hexadecimal digit.
  - P: LPAR number as 1 hexadecimal digit.
  - SSSSS or SSSS: Serial number as 5 or 4 hexadecimal digits derived from the low order digits of the system serial number.
- Model Number: "9672" as 4 hexadecimal digits.
- 0000: "0000" as 4 hexadecimal digits.

1.6.4 Packaging Technology

The F-7C and F-7E processor components use the design features for optimum packaging and performance. Table 1-7 summarizes the characteristics of the primary technologies used by the F-7C and F-7E processor models. Following this table are explanations about such technologies and how they affect functionality.

Table 1-7 F-7C and F-7E Processor Group Technology Summary

Component	*Topic*	*Characteristics*
Logic LSIs in IPs and SP	Technology	CMOS
	No. of transistors	25 million
Memory LSIs in Processor Storage	Technology	CMOS
	Capacity per chip	64, 128, or 256 Mbits
Logic LSIs in IOP	Technology	CMOS
High-Density Module (HDM)	Material	Glass ceramic
	Size (mm)	136 × 136
	No. of layers	55
	No. of pins	5042
	No. of LSIs (Max.)	36
	No. of IPs (Max.)	11 on F-7C, 13 on F-7E
Processor Package	Size (mm)	364 × 534
	No. of layers	30
	No. of HDMs	1
Cooling	Method	• F-7C Base or F-7C CF: air • F-7C Turbo or F-7E Turbo: air and water

Chip technologies

Explained below are chip technologies used in the F-7C and F-7E processor models.

- Logic LSIs in IPs and SP: The Instruction Processors (IPs) and System Processor (SP) use high-speed, high-density, CMOS-based LSIs. CMOS technology increases the number of logic gates per chip and reduces the distance that signals must travel between gates. This high gate density (25 million transistors per chip) enables functions previously using multiple chips to be consolidated onto a single chip. Similarly, functions previously requiring multiple cards or modules are consolidated into a single multi-chip package.

- Memory LSIs in Processor Storage: Processor Storage uses 64, 128, or 256 M-bit CMOS-based memory LSIs. These chips sustain high memory throughput and reduce the physical frame size and footprint used by large-capacity processor storage.

- Logic LSIs in IOPs: The Input Output Processors (IOPs) use logic-in-RAM CMOS LSIs. Each is made up of a 364-kilobit RAM and a 1,150-kilogate logic area. This technology, in conjunction with circuitry innovations, allows concurrent channel maintenance and packaging efficiencies.

High-Density Module (HDM)

On each system, a single High-Density Module (HDM) comprises Instruction Processors (IPs), System Processor (SP), cache, and System Controller (SC). The HDM is a very small package providing excellent environmental characteristics.

Processor Package

Processor Package comprises HDM, Processor Storages, and Key Storages.

Cooling

The F-7C Base Models and CF Models use air-cooling technology only. The F-7C Turbo Models and F-7E Turbo Models use water for cooling on the HDM and air for cooling on the other components. For cooling specifications, see *F-7C and F-7E Processor Group Installation Planning Guide* (F7-GA-1C02).

1.7 Functional Design Highlights

The following lists the main facilities and features supported by the design of the F-7C and F-7E processor models, where "*" indicates "Not available or not effective on F-7C CF Models." For further information, refer to the subsections that follow.

- ESA/390 architecture support*
- S/390 architectural features and functions*
- SIE compatibility*
- Parallel Sysplex computing
- LPAR support
- Virtual Server Facility Phase-1*
- ESCON Channel compatibility*
- Parallel channel support*
- Open Systems Connection Feature (OCF)*
- Console support including GUI and security
- Reliability, Availability, and Serviceability approach to non-stop computing

1.7.1 ESA/390 Architecture Support

The F-7C Base Models, the F-7C Turbo Models, and the F-7E Turbo Models support ESA/390 architecture operationally in Basic mode, in LPAR mode, or as a guest virtual machine.

The F-7C Base Models, the F-7C Turbo Models, and the F-7E Turbo Models provide the necessary hardware characteristics to fully support ESA/390 architecture, including:

- Large storage capacity
- High speed data transfer using:
  - Parallel channels
  - Extended Serial Channels (ESCHs)
  - Inter-System Coupling Channels 2 (ISCH2s)

1.7.2 S/390 Architectural Features and Functions

The F-7C Base Models, the F-7C Turbo Models, and the F-7E Turbo Models provide complete compatibility for all S/390 features and functions. The following is a partial list of features supported. Refer to chapter 6 for further details and for a complete listing of which features are supported in ESA/390 mode.

- Asynchronous Data Mover Facility (ADMF)
- Asynchronous Page-out Facility
- Called Space Identifiers
- Cancel I/O
- Checksum
- Compare-and-Move Extend
- Data Compression
- Dynamic I/O Reconfiguration Management (DRM)
- Enhanced Move-Page (MVPG-2)
- Immediate-and-Relative Instructions
- Logical String Assist
- Parallel Sysplex participation
- Scalar Square Root Instructions
- Set Address Space Control Fast (SACF)
- Subspace Group Facility
- Subsystem Storage Protection (SSSP)
- Suppression on Protection

1.7.3 SIE Compatibility

All characteristics and functions needed to support SIE compatibility are provided on the F-7C Base Models, the F-7C Turbo Models, and the F-7E Turbo Models including:

- Expedited SIE
- Interpreted SIE
- Multiple High Performance Guest Support (MHPGS)
- Region Relocation Facility for VM
- Region Relocation Facility for MLPF

1.7.4 Parallel Sysplex Computing

The F-7C and F-7E processor models fully participate in a Parallel Sysplex environment, or implement Parallel Sysplex by providing connectivity to a coupling facility function.

As a participant in a Parallel Sysplex environment, the F-7C and F-7E processor models provide:

- Coupling code called Highspeed Coupling Control Feature Assist (HCCFA) effective under MLPF
- Connectivity to coupling facility via Inter-System Coupling Channels 2 (ISCH2s)
- Connectivity to the Sysplex Timer via External Timer Attachment Feature (ETAF)
- ESCH channels
- Support for MVS/ESA, Ver. 5 and OS/390

To implements a Parallel Sysplex environment, the F-7C and F-7E processor models provide:

- ISCH environment with standalone HCCF or F-7C CF Model
- ISCH environment with logically-partitioned HCCF
- ICMF environment

1.7.5 LPAR Support

Multiple Logical Processor Feature (MLPF) supports up to 15 logical partitions (LPARs) which enables the use of a logically-partitioned operating mode. Other support for LPARs includes:

- Dynamic storage reconfiguration
- ESA/390 architecture support (F-7C Base Models, F-7C Turbo Models, and F-7E Turbo Models only)
- Support for up to a 16 GB pool of main storage for LPARs (a single LPAR cannot exceed 2 GB)

1.7.6 Virtual Server Facility Phase-1

Virtual Server Facility (VSF) Phase-1 provides the user with virtually fenced, fixed groups of physical IPs called "Virtual Servers" under MLPF on the F-7C Base Models, the F-7C Turbo Models, and the F-7E Turbo Models. Each virtual server has a unique set of version code and serial number to be reasonably determined by the capacity of the server, thus contributing to possible reduction of the user's software cost. Changes to the physical IP grouping are only allowed to the OEM customer's authorized service representative.

1.7.7 ESCON Channel Compatibility

The F-7C Base Models, the F-7C Turbo Models, and the F-7E Turbo Models provide the ESCH channels compatible with IBM's ESCON:
- ESCON architecture and protocols
- Fiber optic cabling capability
- Dynamic connectivity
- Connectivity to ESCON Directors
- Connectivity to ESCON Converters (9034, Model 1 and Model 2)
- 17 MB/s data rate
- Connectivity to multimode fiber optic cable
- ESCON Multiple Image Facility (EMIF)

1.7.8 Parallel Channel Support

Parallel channels are fully supported on the F-7C Base Models, the F-7C Turbo Models, and the F-7E Turbo Models to allow attachment of pre-ESCON devices via bus and tag cables. Parallel channels support data rates up to 4.5 MB/s on the standard (IBM-defined) I/O interface or 6 MB/s on the Hitachi-specified I/O interface.

1.7.9 Open Systems Connection Feature

Open Systems Connection Feature (OCF) provides connectivity and interoperability to a variety of defacto-standard local area network (LAN) such as Ethernet LAN, fiber distributed data interface (FDDI) LAN, or asynchronous transfer mode (ATM) LAN, and wide area network (WAN) for the F-7C Base Models, the F-7C Turbo Models, and the F-7E Turbo Models. OCF aids in deriving the benefits of client/server computing while achieving improved connectivity throughout the enterprise.

1.7.10 Console Support Including GUI and Security

Console support includes the following features:
- Console Integration (F-7C Base, F-7C Turbo, and F-7E Turbo only)
- Graphical user interface (GUI)
- Console Device (CD)
- Hardware Management Console (HMC) connectivity (optional)
- Remote Console
- Customer-administered security

1.7.11 RAS Approach to Non-Stop Computing

Reliability

Reliability is assured with minimal physical components and connection points.

Availability

Characteristics supporting system availability include:
- Concurrent channel upgrades
- Concurrent maintenance of channels, power supplies, Refrigeration Units (F-7C Turbo and F-7E Turbo only), fans, microcode, and MLPF
- IP/SP standby
- Console Device redundancy
- Fan redundancy
- Chip alternation
- N + 1 power supplies

Serviceability

Serviceability support functions include:
- Field-replaceable unit (FRU) identification
- Online problem diagnostics
- Auto Call
- Remote operations and downloading of firmware through support centers
- Remote operations through support centers
- Remote firmware placeholders
- Deferred maintenance for redundant components

CHAPTER 2 INSTRUCTION PROCESSOR AND SYSTEM PROCESSOR

2.1 Processor Overview

The F-7C and F-7E processor models contain one to thirteen Instruction Processors (IPs), depending upon the model, and one System Processor (SP). Both the IP and the SP use the same architectural and hardware implementations, while the logic managing these processors and their interfaces with other system components differ. Specifically, each IP or SP is implemented on a single CMOS VLSI chip, whereas the IP and the SP provide different functions as explained below.

The IP executes operating systems and application programs. It interprets processor instructions and executes those instructions using the appropriate hardware and firmware.

The SP manages external interfaces needed by the IPs to the IOPs and the SVP. It performs the back-end portion of I/O processing by routing I/O interrupts to the IPs for processing. It also serializes IP requests to the SVP and performs asynchronous execution for special IP instructions.

To summarize the foregoing, the SP performs the following general functions:

- Instruction execution for the back-end of I/O, asynchronous Expanded Storage access, and DIAGNOSE and SERVICE CALL instructions
- I/O interrupt routing
- Reset operations and SVP frame operations
- System check handling for failure detection, logout, and failure recovery

2.2 System Processor Interface Design and Function

When an IP needs to interface with an IOP or the SVP during instruction processing, it interfaces directly with the SP instead. The SP in turn handles the SVP and IOP interfaces. Note that the SP handles only back-end interrupts and is not involved with starting an I/O. The SP uses the following types of interfaces.

2.2.1 Interface between SP and IP

- SP to IP(s): This interface uses a request table with locks in the Hardware System Area (HSA). An SP issues requests to an individual IP or to all IPs (as in failure situations) and updates the HSA request table. When the IP processes the signal using the request table, it either handles the request or has the SP execute the required action for the IP.

- IP to SP: This interface uses a request table with locks in the HSA. An IP issues requests to the SP via the request table. When the SP polls the table, it executes the requested action and removes the request from the table.

2.2.2 SVP Interface

This interface uses a four byte read/write register and service word protocol. Both the SVP and SP pass requests back and forth using the register.

The SP uses this interface for tasks such as DIAGNOSE, SERVICE CALL, and HYPERVISOR CALL instructions and for performing status report and logout requests.

The SVP uses this interface for tasks such as maintenance functions and reset, and to indicate the end of a required instruction such as END OF DIAGNOSE or SERVICE SIGNAL.

2.2.3 IOP Interface

- IOP to SP: The IOP uses an attention request mechanism to report interruptions of I/O completions or errors. Or, a control interface is used to set/reset hardware registers which enable/disable IOP attention.

- IOP to SP to IP: The SP is not used to start I/O activity but instead routes back-end I/O interruptions to the appropriate IPs. The SP generates interruptions by sending I/O interrupt requests to the IP.

2.3 Address Translation Process with Lookaside Buffers

Accessing instructions and data in main storage requires translating virtual addresses into real addresses, which is accomplished by the dynamic address translation (DAT) mechanism.

Address translation tables architecturally reside in main storage, however, and main storage accesses can degrade performance even with sophisticated cache mechanisms. The F-7C and F-7E processor models therefore contain two types of lookaside buffers to help prevent main storage accesses for address translation. When an address to be translated is found in these buffers, access to the main storage translation tables is bypassed.

2.3.1 Virtual Address Formats

The F-7C and F-7E processor models support 31-bit and 24-bit addressing of main storage under program control. Formats are shown in figure 2-1.

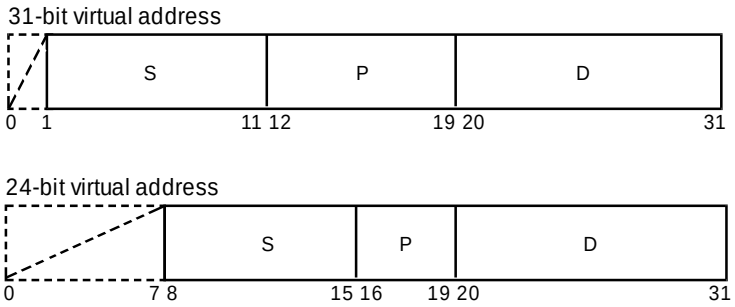

S: Segment field. Specifies the entry in the segment table.
P: Page field. Specifies the entry in the page table.
D: Displacement field. Specifies the address within the page.

Figure 2-1 Examples of Virtual Address Formats

2.3.2 Address Translation Using Translation Lookaside Buffer

Each IP and SP has a Translation Lookaside Buffer (TLB) to accelerate translating virtual addresses into real addresses. The TLBs hold virtual-to-real address pairs translated by the dynamic address translation (DAT) facility. The reuse of entries held in the TLB allows the full DAT process (storage access of segment and page tables) to be bypassed, thus improving overall system performance.

The TLB uses high-speed array chips to provide high performance and can hold up to 256 virtual-real address pairs. The TLB employs a set-associative addressing structure. Lines in the TLB are replaced by the least recently used (LRU) algorithm.

If the TLB does not contain the sought-after valid virtual-real address pair, the DAT process, which makes use of control register contents, segment tables, and page tables, translates the virtual address into the real address. The following description explains these address translation steps, including the registration of the translated virtual-real address pairs in the TLB.

1. Adds the value in the segment field of the virtual address to the segment table origin address stored in either control register 1 or 7, depending upon whether the address space is primary or secondary.
2. References the segment table using the address derived in step 1.
3. Adds the value in the page field of the virtual address to the page-table-origin address found in the segment table entry.
4. References the page table using the address derived in step 3.
5. Obtains the real page frame address.
6. Converts the real page frame address into the absolute page address by prefixing and by concatenating the value in the displacement field of the virtual address to the right of the absolute page frame address.
7. Registers into the TLB the pair of the segment and page fields of the virtual address and the absolute page address.

As part of the address translation process, the lengths of the segment and page tables and the entries in these tables are checked for validity. If any errors are found, an address translation exception is recognized and a program interruption is generated.

2.3.3 Address Translation Using Access Register

Each IP and SP has an ART Lookaside Buffer (ALB) to accelerate translating virtual addresses into real addresses. This translation process is called access register translation (ART).

In ESA mode, data in any space can be referenced by an access register with the base register field of the instruction. For example, in a sequence where a MOVE CHARACTER instruction (MVC) in the address space selects target data using an access register (access register translation), the following takes place:

1. The base register field of the instruction designates an access register.
2. The value in the selected access register is added to the access list origin address, and the result selects an entry from the access list.
3. The access list entry refers to a control table.
4. The segment table origin designated by the control table selects a target data space.
5. The target data address in the data space is obtained as a virtual address.
6. The virtual address of the target data address is translated to a real address by conventional dynamic address translation using the segment table origin in the virtual address.

In this translation process, the segment table origin is obtained through two levels of conversion: access list and control table, which are both maintained in main storage. If such reference to main storage occurs at every translation, the translation speed remains relatively low. To improve this, the ALB is implemented in the IP and the SP as a dedicated lookaside buffer for access register translation allowing the translation process to bypass two main storage references.

2.4 Pipeline Design and Features

The F-7C and F-7E processor models use pipelines for instruction execution. This pipeline operates in overlap mode so it is able to have a different instruction at each stage of the pipeline and complete an instruction once every system cycle.

Each instruction uses only the pipeline stages needed. Complex instructions can loop in one or more stages or through the entire pipeline several times before completing. Simpler instructions may wait in the pipeline while any preceding, more complex instructions complete. The pipeline's design includes features that enhance its performance while managing the differences that exist among processor instructions.

The F-7C and F-7E models use the following seven-stage pipeline for instruction execution:

- Instruction decode cycle: The instruction is decoded and prepared for execution.
- Address generation cycle (C0 cycle): Storage addresses required by the instruction are generated and resolved.
- Access cache cycle (C1 cycle): Cache is accessed for fetching of necessary data.
- Data return to execution unit cycle (C2 cycle): The data is returned to the execution unit.
- Instruction execution cycle: The instruction is executed according to its particular architectural rules.
- Put away cycle: Instruction execution result operands are put away for future use. If a result operand resides in main storage, the result is placed in Buffer Storage (BS) for eventual movement to main storage.
- Checkpoint data cycle: The checkpoint data for the instruction execution is saved.

2.5 Alternate Processors

In addition to the IPs and the SP, most F-7C and F-7E processor models contain Alternate Processors (APs) as aforementioned in chapter 1. Physically, APs are identical to the IPs and the SP. By contrast, the IPs to run general purpose work are sometimes notated as "[regular] IPs."

APs, depending upon the processor model, serve as any of the following processors:

- an Integrated Coupling Facility (IntCF) IP
- a Disk System Processor (DSP)
- a standby processor to replace a failed SP
- a standby processor to replace a failed [regular] IP
- a standby processor to replace a failed IntCF IP
- a standby processor to replace a failed DSP

When an AP acting as a standby processor (Standby AP) replaces a failed SP, IP, or DSP, the failed processor is automatically placed in a state where it is unavailable to process any work.

APs do not change the model number of any F-7C or F-7E processor model. For example, the F-7 Model 92C with nine [regular] IPs and two APs remains a 9-way 92C model. The APs do not affect software licensing charges based on aggregate system capacity.

CHAPTER 3 STORAGE SUBSYSTEM

3.1 Storage Subsystem Overview

Controlled by processor activity, storage is used as:
- input areas for instructions and data waiting to be processed
- queueing areas for prefetched information
- a scratch pad during instruction processing
- areas for generated data being prepared for output
- buffer areas and workspace for calculating and accumulating information

The purpose of the storage subsystem is to:
- provide storage area types tailored to optimize different processor needs
- ensure that the data is stored in formats recognized by all supported System Control Programs (SCPs) and architectures
- ensure that the data contents are protected from inappropriate access or use
- monitor and adjust for hardware and system error conditions that affect storage integrity
- provide instructions and data to the IPs with minimal delay

In the F-7C and F-7E processor models, the processor storage subsystem includes the physical processor storage implemented on CMOS random access memory (RAM) chips, the internal resource management logic, and the architecture imposed by the system. The processor storage cache hardware is integrated into the IP and SP structures. Processor storage, using a Storage Controller (SC), stores all data and instructions, provides access control and protection, and assures that addressing and representation conventions are met. All processor storage and storage-related logic reside in the Basic Processing Unit (BPU).

In addition to the processor storage, most of the F-7C Base Models and F-7E Turbo Models can optionally configure the Integrated Disk.

Sections 3.2 through 3.7 describe the general physical design, architecture, and supported features of the F-7C and F-7E Processor Group processor storage. Section 3.8 describes the Integrated Disk.

3.2 Hardware Organization

The F-7C and F-7E processor systems implement a hierarchy of areas in which instructions and data are stored.

The most fundamental area is processor storage. Processor storage consists of two sub-types: main and expanded. Main storage is byte-addressable and contains executable instructions and data accessed by those instructions. This storage is ultimately manipulated by the programs executing in the IP. Instructions and data in main storage exist for the life of the program or as specifically controlled by the program. Expanded storage is block-addressable and is used for specialized program functions such as buffering and paging.

Although it is the most fundamental area in the hierarchy, processor storage is also the slowest. Between the IP and processor storage are two intermediary hierarchical storage components that exist for performance reasons only: Buffer Storage (BS) and Work Storage (WS). They may sometimes be called cache or buffer storage.

Figure 3-1 illustrates the storage hierarchy and shows general relationships between the storage components.

As illustrated, storage access time slows but storage capacity increases the further the storage component is from the IP. Consequently, storage access becomes faster but storage capacity decreases the closer the storage is to the IP. This design allows the IP to access data as quickly as possible while simultaneously providing very large amounts of storage capacity.

Main storage is cached. Instructions and data exist in the cache hierarchy transparently to the program. Expanded storage is not cached. Data is transferred to and from main storage without flowing through the caches.

The cache components perform no architectural function. The caches exist solely to allow the IP to execute instructions with minimal delays due to having to access storage. The caches are not generally known by or visible to programs running in the IP, although there is a limited set of system control instructions exploited by operating systems to optimize cache usage.

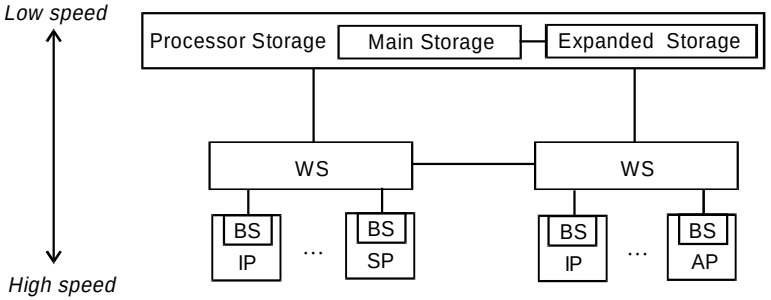

Figure 3-1 Storage Hierarchy for F-7C and F-7E Models

3.3 Cache Buffers

Buffer storage areas contain frequently-accessed instructions and data. They also function as temporary scratch pads for work in progress. The access time for obtaining data from these buffers is much faster than retrieval from main storage. The result is that the use of buffers increases throughput as well as optimizing processor use. Data is moved in fixed-length blocks called “lines.”

3.3.1 Buffer Storage (BS)

This type of cache buffer is physically located in each IP or SP chip. This buffer is the fastest of the processor storage areas and the smallest. Each BS is used as a holding area and scratch pad by the IP or SP. Data is moved in 256-byte lines.

3.3.2 Work Storage (WS)

This type of cache buffer resides on a chip next to each IP or SP. This buffer is slower than BS but has much greater capacity. Each WS contains copies of the instructions and data from main storage. Data is moved in 256-byte lines.

3.3.3 Cache Algorithms and Data Movement

The fetching of data or instructions by an IP logically begins in main storage and ends with the instructions and data positioned in BS for execution. As during the execution process, incremental results are generated and stored in BS and simultaneously in WS via store-through algorithms. The data stays in WS until internal conditions trigger a move to main storage simultaneously.

Table 3-1 shows the characteristics of the two types of cache where PE (Processor Element) means IP, SP, or AP.

Table 3-1 Cache Characteristics

Cache	*Attribute*	*Data*
Buffer Storage (BS)	Capacity	256 KB/PE
	Algorithm	Store-through
	Line Size	256 bytes
Work Storage (WS)	Capacity	4 MB/6 PEs on F-7C, 8 MB/7 PEs on F-7E
	Algorithm	Store-in
	Line Size	256 bytes

3.4 System Controller

The System Controller (SC) controls access to main storage, expanded storage, and channel subsystem. Each 1SC model of the F-7C processor system has one SC. Each 2SC model of the F-7C and F-7E processor systems has two SCs.

The SC is a distributed design comprised of multiple VLSI chips, all of which reside in the BPU. Some of the SC logic is housed in the High Density Module (HDM) with the IPs, the SP, and cache. This logic provides access not only to processor storage but also to the channel subsystem and to the optional External Timer Attachment Feature (ETAF). Other portions of the SC logic are physically packaged with the RAM comprising processor storage.

3.5 Processor Storage

The contiguous processor storage area is composed of CMOS dynamic random access memory (DRAM) chips. The area is available in maximum sizes of 16 GB on each F-7C 1SC Model, and 32 GB on each F-7C and F-7E 2SC Model. See figure 3-2.

Main storage is partitioned from processor storage in 64 MB increments. The size of main storage is set during initial microcode loading (IML) in accordance with the system installation information previously entered from the console.

The remaining processor storage is configured as expanded storage by default. Whether it is desirable or not to divide processor storage into main and expanded depends on the needs and capabilities of the operating system environment.

Architecturally, an S/390 operating system may use up to a maximum of 2 GB of main storage. However, when an F-7C or F-7E processor system (either 1SC or 2SC) is running in LPAR mode via MLPF, up to 16 GB of main storage can be made available for use by all LPARs combined. The remaining processor storage is usable as expanded storage if so partitioned at IML. Each LPAR can be allocated as all or as a portion of the expanded storage.

Main and expanded storages allocated to an LPAR are dedicated to that LPAR, and cannot be concurrently shared or accessed by any other LPAR. With proper configuration planning, however, main and expanded storages can be non-disruptively moved to another LPAR subject to configuration rules using the Dynamic Storage Reconfiguration facility.

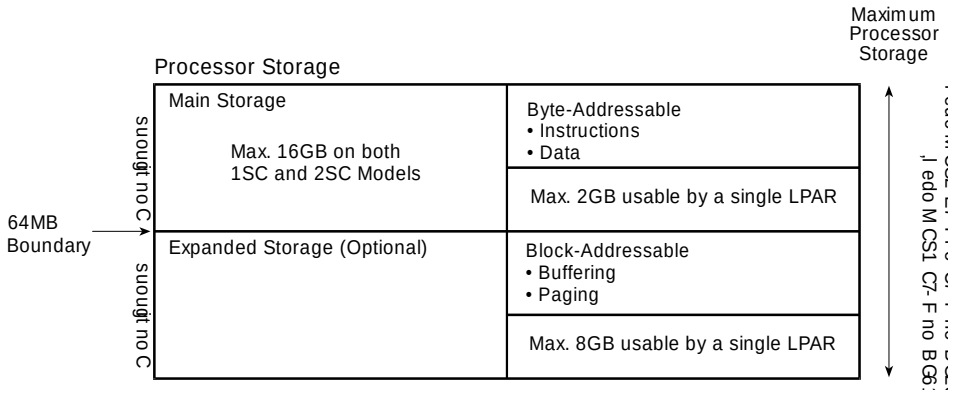

Figure 3-2 Processor Storage

3.5.1 Main Storage

Main storage is byte-addressable and contains all the instructions and data actively in use by the programs running in the IPs. All instructions and data must initially reside in main storage in order to be fetched and processed by an IP. The instructions and data are automatically prefetched up to a cached buffer for optimal processor access, but the starting point is always main storage.

Main storage uses floating address registers (FARs) to map physical blocks of storage to a program's absolute addresses.

An FAR is assigned for each block of processor storage. The size of the block depends on the total size of physical processor storage. Mapping of absolute-to-physical addresses is established by placing the block's absolute address in its FAR. The FAR is used to:

- alter address mapping correlations
- implement Dynamic Storage Reconfiguration
- remove a block with a fault from use (remove the block from the system configuration)

3.5.2 Expanded Storage

Expanded storage is block addressable in 4 KB blocks and is used by some operating systems for buffering and paging. As mentioned previously, the size of main storage is set at IML; the remaining processor storage is configured as expanded storage by default. Any data in expanded storage must be transferred to main storage before it is used by a program. A program can also transfer data to expanded storage.

Expanded storage is architecturally and physically connected to the system directly; so, data in expanded storage can be accessed much faster than data in the peripheral devices which are accessed through relatively slow I/O operations. However, expanded storage does not employ storage protection features such as key protection, and conventional instructions cannot directly access expanded storage.

Some operating systems and subsystems significantly improve performance by using expanded storage for paging, swapping, program controlled caching, or as a file device. How expanded storage is functionally used is defined by software, not by the architecture or the processor hardware.

Two methods of transferring blocks between main storage and expanded storage are used. Neither of these transfer modes requires IOP initiation or a closing procedure (interruption).

- Synchronous transfer (for example, MVPG and MVPG2, or Move Page): The system uses special CPU instructions to directly transfer data between main storage and expanded storage. Instructions are processor synchronous and hence instruction execution does not complete until data transfer is complete or an error condition is detected.
- Asynchronous transfer (for example, Asynchronous Data Mover Facility): There is asynchronous data transfer between main storage and expanded storage. Instruction execution continues after request for data movement is accepted and continues under program control concurrently with data movement. ADMF offloads data movement between main storage and expanded storage from the IPs. On the F-7C and F-7E processors, asynchronous transfer is implemented by the System Processor (SP).

3.5.3 Processor Storage Sizes

Table 3-2 lists the sizes of processor storage which can be installed on and configured to the various F-7C and F-7E processor models.

Table 3-2 Available Processor Storage Sizes

Model	*Processor Storage Capacities (B)*				
	Total			*MS**	*ES*
	Min.	*Increment [Range]*	*Max.*	*Min–Max [Increment]*	*Min–Max [Increment]*
1SC Model	256M	256M [256–3072M], 512M [3072–8192M], 2G [2–16G]	16G**	256M–16G [64M]	0–Total minus MS [64M]
2SC Model	256M	256M [256–3072M], 512M [3072–8192M], 2G [2–16G], 8G [16–32G]	32G***	256M–16G [64M]	0–Total minus MS [64M]

* Maximum size of main storage per LPAR is 2GB.
** 8GB is the maximum processor storage capacity on F-7C 1SC Models in current shipment. The maximum processor storage capacity of 16GB will be available from 10/99.
*** 16GB is the maximum processor storage capacity on F-7C 2SC Models in current shipment. The maximum processor storage capacity of 32GB will be available from 10/99.

3.6 Storage Protection

Five protection mechanisms are provided in the F-7C and F-7E processor models to protect the information in virtual/real storage: key-controlled protection, low-address protection, page protection, Subsystem Storage Protection, and Subspace Group Facility.

3.6.1 Key-Controlled Protection

A storage key is provided for each 4KB block of main storage. Every time a storage access is made, the key of the associated 4KB block is checked against the access key of the requesting IP or channel to verify the validity of the store or fetch request.

Change and reference bits, used by the operating system for virtual/real storage management, are also maintained in each protection key. Whenever a Translation Lookaside Buffer (TLB) entry is created, a copy of the protection key is placed in the TLB entry.

3.6.2 Low-Address Protection

Low-address protection protects the first 512 bytes of the prefixed storage area (which contains data critical to the system's operation) from being destroyed. Real addresses 0 to 511 of each active prefixed storage area are protected whenever bit 3 of control register 0 is set to 1.

3.6.3 Page Protection

Page protection provides a means to stop data from being stored anywhere within a page of virtual storage. Whenever bit 22 of a page table entry is set to 1, store operations into that page are prohibited. Whenever a TLB entry is created, the status of the page protection bit is placed in the TLB entry. This permits page protection checking to be part of the dynamic address translation.

3.6.4 Subsystem Storage Protection

Subsystem Storage Protection (SSSP) allows subsystem storage to be isolated from application programs, preventing any application programs from overwriting (or overlaying) subsystem storage. SSSP is exploited by CICS/ESA, Version 3 and above.

3.6.5 Subspace Group Facility

Subspace Group Facility provides a more robust protection mechanism than SSSP for use by CICS. CICS runs both CICS subsystem programs and one or more application programs in a single address space. With Subspace Group Facility and the CICS storage isolation facility, each application program is assigned storage private to itself; this private storage cannot be accidentally or maliciously modified by any other application program. This protection mechanism is exploited by CICS/ESA, Version 4, Release 1 and above.

3.7 Hardware System Area

The Hardware System Area (HSA) is a special area of main storage used only by the system hardware. The HSA portion of main storage is set aside for use by the system hardware and is not directly accessible by any operating system, subsystem, or application software. The HSA provides storage areas for:

- Communication among different system components such as the SP, the IPs, and the IOPs.
- Storing data related to internal hardware functions and locks.
- Storing unit control words (UCWs) representing individual I/O devices.
- Storing management tables and queues used by the dynamic channel subsystem.
- Common area for system directories.
- Holding a special additional subarea for the I/O trace data when an I/O maintenance trace has been selected at initial microcode loading (IML). In this case the HSA is expanded in its area size.

In some MLPF environments, an additional main storage area called the Extended Hardware System Area (EHSA) is reserved.

HSA is allocated as a single contiguous area of storage. The size used depends upon the number of I/O devices, the complexity of the I/O device pathing, the mode of system operation, and whether or not certain resource management features such as DRM and EMIF are enabled or disabled.

Figures 3-3 through 3-5 illustrate the logical layouts of HSA in Basic mode, in LPAR mode with up to 2 GB main storage for use by LPARs, and in LPAR mode with more than 2 GB and up to 16 GB main storage for use by LPARs.

Since HSA size varies with operating environments, three calculation models (standard Basic, standard LPAR, and maximum LPAR) are assumed as shown in table 3-3. Table 3-4 provides the calculation results under the above-mentioned conditions.

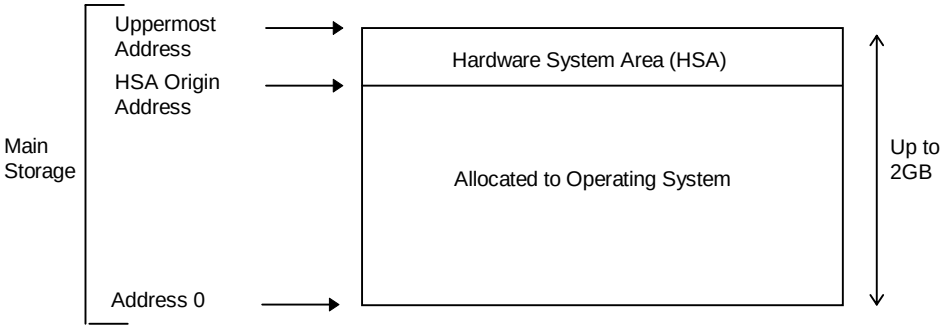

Figure 3-3 HSA in Basic Mode

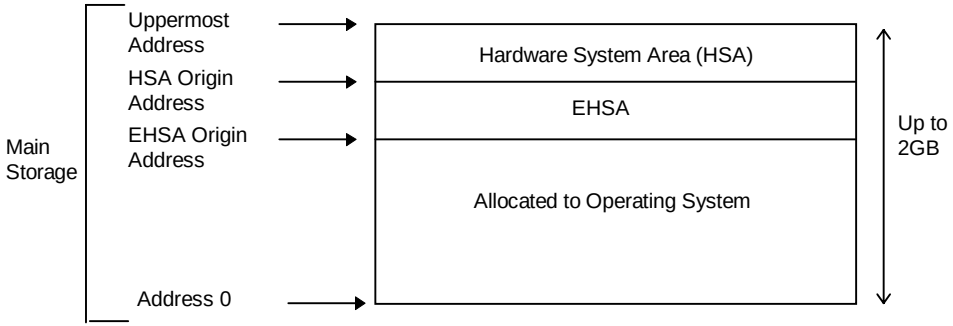

Figure 3-4 HSA in LPAR Mode with up to 2 GB MS

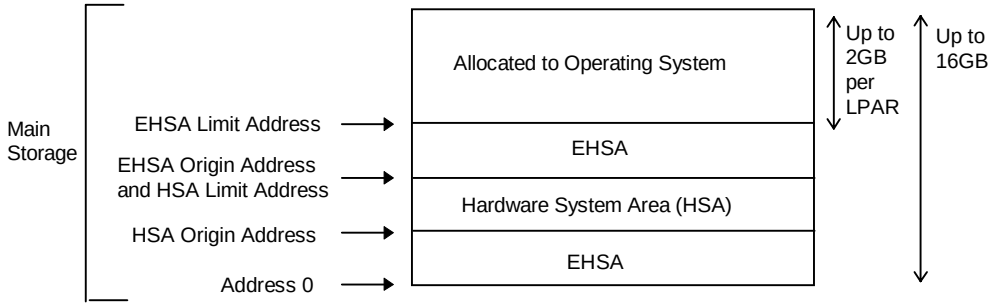

Figure 3-5 HSA in LPAR Mode with More than 2 GB up to 16 GB MS

Table 3-3 Calculation Models (Assumptions)

Calculation Factor		*Assumptions by Calculation Model*		
		Standard Basic	*Standard LPAR*	*Maximum LPAR*
Operating mode		Basic	LPAR	LPAR
DRM setting		NA	50 % expansion ratio	50 % expansion ratio
EMIF setting		NA	50 % shared ratio	50 % shared ratio
No. of CHs	Metal	64	64 (Reconfigurable)	256
	ESCH	32	32 (Shared)	
I/O configu-ration	No. of devices	64/CH	64/CH	40 K
	No. of logical CUs	4/CH	4/CH	5 K
	No. of physical CUs	4/CH	4/CH	6 K
	No. of paths	2/CU	2/CU	2/CU

Table 3-4 Calculation Results (Typical HSA Sizes)

Item		*Sizes in MB by Calculation Model*					
		F-7C			*F-7E*		
		Standard Basic	*Standard LPAR*	*Maximum LPAR*	*Standard Basic*	*Standard LPAR*	*Maximum LPAR*
I/O area	Conditional	0.5	3.2	6.8	0.5	3.2	6.8
	Fixed	11	11	11	11	11	11
Other than I/O area		6	6	6	7	7	7
Total (rounded up)		18	21	24	19	22	25

3.8 Integrated Disk

The Integrated Disk (IDK) and its associated features, optionally available on the F-7C Base Models except B2C model and on the F-7E Turbo Models except 82E and D2EH models, provide faster data access by PSTR-Cache, IDK, and Fibre Channel interface. An IDK and its associated features comprise the following dedicated components.

- Level 1 Disk Cache: The Level 1 Disk Cache is the storage area added in PSTR and utilizes write-through management. The cache has the capacity of 512 MB per IDK (subsystem) and 4 GB per system.

- IDK: The IDK is a disk subsystem to be accommodated in the Extended Frame. Up to eight IDKs are installable per system. Each disk subsystem offers 60 to 130 GB user area depending on the emulation mode. The IDK stores duplicated data by mirroring.

- Level 2 Disk Cache: The Level 2 Disk Cache resides in the IDK. The cache utilizes write-after management and mirroring. The available size of this cache on each IDK is 512 MB.

- Disk System Processor (DSP): The DSP is dedicated to control the IDK, and one DSP manages up to four IDKs. Addition of a DSP reduces the number of available Standby APs by one.

- Fibre Channel (FBCH): The FBCH allows connection to/from the IDK on the special protocol called Fibre Channel Protocol (FCP). One FBCH PK has two channels, and two channels on different PKs are required to connect one IDK. Each FBCH transfers data at 100 megabytes (MB) per second.

- Extended Frame: The Extended Frame incorporates up to two IDKs. Up to four Extended Frames are installable per system.

Figure 3-6 illustrates the processor configuration with the IDK. Table 3-5 shows the limits of the components for the IDK and its associated features.

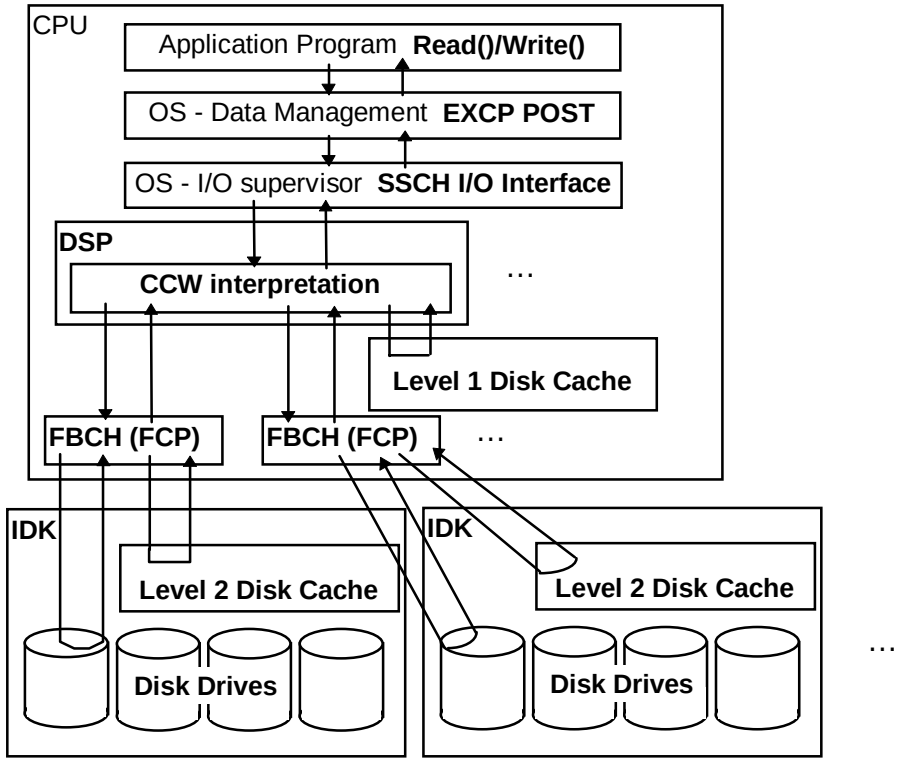

Figure 3-6 Processor Configuration with Integrated Disk

Table 3-5 Limits for Integrated Disk and Its Associated Features

Category	*Contents*	
Max. Level 1 Disk Cache capacity	4 GB/system (512 MB/subsystem)	
Max. Level 2 Disk Cache capacity	512 MB/subsystem	
Max. user area capacity on disk	3380-K	968 GB/system (121 GB/subsystem)
	3390-1	482 GB/system (60 GB/subsystem)
	3390-2	968 GB/system (121 GB/subsystem)
	3390-3	1086 GB/system (136 GB/subsystem)
Max. No. of disk subsystems	8/system	
Max. No. of Disk System Processors	2/system	
No. of Fibre Channels	2/subsystem	

CHAPTER 4 CHANNEL SUBSYSTEM

4.1 Channel Subsystem Overview

In support of ESA/390 architecture, the channel subsystem in the F-7C Base Models, the F-7C Turbo Models, and the F-7E Turbo Models manages all details of an I/O request. The subsystem performs all translation between logical and physical I/O device addresses, all device path selection and queuing, and handles I/O queue requests and restarts resulting from the device being shared with another system.

In ESA/390 mode, the operating system reduces its overhead by queuing I/O requests only at the logical I/O device level. This approach:

- Eliminates the need to fix the physical I/O device and channel affinity to a specific Instruction Processor (IP).
- Allows the initiation of I/O process from any IP regardless of the physical I/O configuration.
- Enables the channel subsystem to send I/O interrupts to the operating system using any IP available which is enabled for I/O interrupts.

The channel subsystem in the F-7C CF Models manages all details of inter-system coupling interfaces in Parallel Sysplex environments.

4.2 System Architecture Supported by Channel Subsystem

The channel subsystem of the F-7C and F-7E Processor Groups comprises the Input Output Processor (IOP) which accommodates the Channel Processor (CHP) and various types of channels.

The IOP of the F-7C Base Models, the F-7C Turbo Models, and the F-7E Turbo Models works as part of the dynamic channel subsystem (DCS) functional in ESA/390 mode. In ESA/390 mode, the operating system queues requests at the logical I/O device level only, since all device path selection and queuing operations are carried out by the channel subsystem. In addition, I/O request queuing and restarts resulting from the device being shared with another system are handled by the channel subsystem. These operations, performed by the channel subsystem, reduce complexity and overhead in the operating system.

Because logical devices are used, there is no need for a fixed affinity between a physical I/O device and its channel with a specific Instruction Processor (IP). In systems with more than one IP, any I/O operation can be initiated from any IP without considering the physical I/O configuration. The channel subsystem can send I/O interrupts back to be dealt with by the operating system without regard to which IP initiated the associated I/O operation. A separate, dedicated processor called the System Processor (SP) polls for I/O interrupts from the IOPs and selects an IP to complete the I/O process.

The IOP of the F-7C CF Models works in a similar manner as that of the F-7C Base Models, the F-7C Turbo Models, and the F-7E Turbo Models. The only difference is that the channel types accommodated by the F-7C CF Models are limited to those of the Inter-System Coupling channels.

The F-7C and F-7E channel subsystems realize various I/O controls by the following hardware facilities:

- I/O control logic inside the IP
- I/O control logic inside the SP
- Hardware system area, which accommodates control information for all I/O related activities
- Unit control words and logical control unit words, which control all I/O related activities
- Channel command words, which control devices

Roles of these facilities in providing such I/O controls are described in the following subsections.

4.2.1 I/O Control Logic inside Instruction Processor

Logic built into each Instruction Processor (IP) takes care of front-end I/O processing. An IP initiates an I/O request to the channel subsystem using data and communication areas in the HSA. When, for example, an IP executes a START SUBCHANNEL (SSCH) instruction, it gathers the data necessary to identify the device and channel program areas involved in the resultant I/O operation and posts a request in the HSA to start the operation.

While other components specialized for I/O processing cause the I/O to be scheduled, executed, and terminated, the IP asynchronously processes other instructions.

The IP provides the architectural interface between programs executing on the F-7C and F-7E processor systems and the systems' attached I/O devices. The IP does so by executing the instructions that control channel operations. The implementation of these instructions uses HSA to initiate I/O operations and to receive status information from the IOPs.

The following ESA/390 mode I/O instructions are implemented on the F-7C Base Models, the F-7C Turbo Models, and the F-7E Turbo Models by the IP's I/O control logic:

- CLEAR SUBCHANNEL (CSCH)
- HALT SUBCHANNEL (HSCH)
- MODIFY SUBCHANNEL (MSCH)
- RESET CHANNEL PATH (RCHP)
- RESUME SUBCHANNEL (RSCH)
- SET ADDRESS LIMIT (SAL)
- SET CHANNEL MONITOR (SCHM)
- START SUBCHANNEL (SSCH)
- STORE CHANNEL PATH STATUS (STCPS)
- STORE CHANNEL REPORT WORD (STCRW)
- TEST PENDING INTERRUPTION (TPI)
- TEST SUBCHANNEL (TSCH)
- CANCEL SUBCHANNEL (XSCH)

4.2.2 I/O Control Logic inside System Processor

The System Processor (SP) is designed to take care of the back-end I/O processing. The SP polls I/O communication areas in the HSA, and, when an I/O interrupt occurs, the SP assembles the data needed and selects an eligible IP to complete the I/O process.

4.2.3 Hardware System Area

The HSA stores all the I/O communication areas, IOCDS information, and tables for the IPs, the IOPs, and the SP. All these hardware units use storage locking to maintain system integrity while accessing HSA data. There is one HSA per system.

The HSA is built during power-on reset or during a SYSIML CLEAR. It contains information from the system Input/Output Configuration Data Set (IOCDS) as defined by the system's user. Up to six user-selectable IOCDS files are provided by the console subsystem file system. Information from the selected IOCDS is used to build, in HSA, the control blocks required for I/O subsystem control. These control blocks are used to map physical I/O devices to unit control words (UCWs) and to pass I/O initiations and interruptions between the IPs and the channel subsystem.

The IP uses the control blocks in the HSA, together with hardware connections to the channel subsystem, to initiate and monitor I/O operations. The instructions MSCH, XSCH, STSCH, TPI, TSCH and STCRW are executed by the IP using information from the HSA without signaling the channel subsystem. The remaining I/O instructions cause the channel subsystem to be signaled after any changes to the HSA data so the channel subsystem can begin to process the request independently of the IP. The IP and the channel subsystem use storage locking to assure system integrity while accessing HSA data.

4.2.4 Unit Control Words and Logical Control Unit Words

Unit control words (UCWs) are commonly called subchannels and are used by system components to record the status of I/O devices and I/O operations. Logical control unit words (LCUWs) are used to map logical control units and physical channel connections to the UCWs. LCUWs are also used for I/O queuing and channel path selection within the channel subsystem.

The data contained in the UCWs and the LCUWs enable the channel subsystem to correctly handle I/O operations and to map physical devices, I/O controllers, and channels for I/O operations.

The data in the UCWs and LCUWs are loaded from the IOCDS into HSA during system power on or during SYSIML CLEAR. IOCDSs are generated by the HCD or IOCP software processes and are stored by the console subsystem.

CAUTION: Failure to correctly specify the type of any control unit and/or the correct range of device addresses supported by that control unit may result in malfunction of an I/O device, control unit, or channel. In particular, I/O interrupts may be lost.

CAUTION: The device address range specified for a control unit must be what the control unit actually supports. This range may be greater than the actual number of devices attached to that control unit. Failure to specify the actual supported range may cause missing I/O interrupts.

CAUTION: Some I/O devices, such as tape units, require shared subchannels. “Shared” in this context refers to the control unit’s ability to process multiple I/O requests concurrently. If a device requires a shared subchannel, specifying a non-shared subchannel may cause data integrity problems. Possible errors include missing records, incorrect records, tape mispositioning, and application program failures. Shared subchannels are specified by use of the SHARED= operand of the IOCP macroinstruction CNTLUNIT. Do not confuse shared subchannels with shared channels. The latter are specified by the SHARED= operand of the CHPID IOCP macroinstruction and are used in ESCON Multiple Image Facility (EMIF) environments.

4.2.5 I/O Device Control by Channel Command Words

A channel program consisting of one or more channel command words (CCWs) controls the operation of each device attached to a processor. The transfer-in-channel (TIC) command is a special command used to chain together CCWs residing in non-contiguous storage areas. The TIC CCW, used in conjunction with the status modifier bit during I/O device status presentation, allows branching between the CCWs.

CCWs are divided into the following basic types:
- Write
- Read
- Read backward
- Control
- Sense
- Sense ID
- Transfer in channel

Channel interface error conditions and status presented to the channel by the I/O device are initially handled by the SP. The action taken varies according to the type of error, the current status of the UCW, and the type of channel in use.

The channel subsystem supports the command retry facility where an I/O control unit signals the channel to reissue a CCW command which has failed. Command retry facility reduces overhead associated with retrying failed I/O operations at the operating system level.

4.3 Hardware Organizaion

Each Input Output Processor (IOP) in the F-7C and F-7E Processor Group consists of two Channel Processors (CHPs).

The CHP handles I/O initiation, path management, queuing, de-queuing, and interruptions for the channel packages (CH-PKs) connected to it. The CHP also handles data transfer functions for 1 through 16 channel packages plugged into the 16 slots as illustrated in figure 4-1.

The number of I/O channels in a channel package depends on the type of channel being attached. Parallel channels and Extended Serial Channels (ESCHs) have four to a package respectively, Inter-System Coupling Channels 2 (ISCH2s) and Fibre Channel (FBCH) have two, and Open Systems Connection Feature (OCF) channels have one.

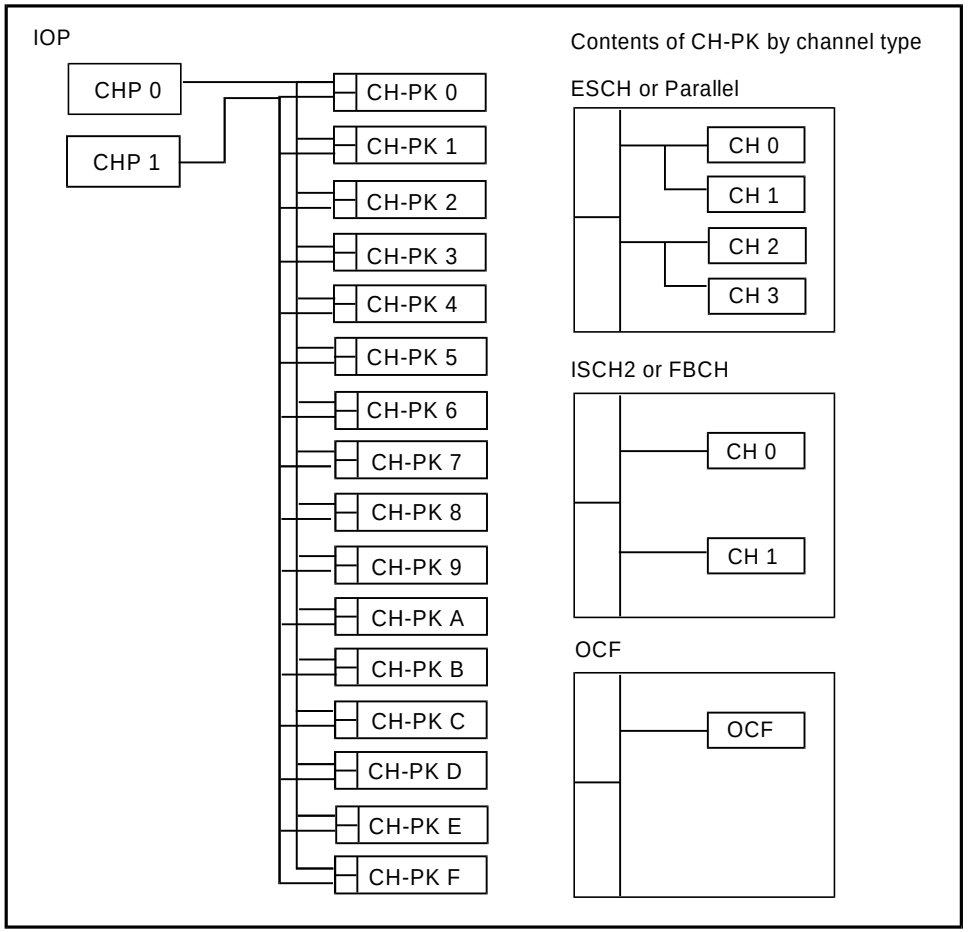

Figure 4-1 Hardware Organization of Each IOP

4.4 Channel Subsystem Functions

Table 4-1 lists what channel subsystem function is supported by what channel type.

Table 4-1 Channel Subsystem Functions by Channel Type

Functions Supported	*Channel Type*
Incorrect length suppression	BL/BY/CVC/CNC/CTC
Concurrent sense	CNC
Channel path Time-out Indication	BL/BY/CVC/CNC
Channel Subsystem Call Instruction	All
Event information report	CNC/CTC
XSCH (Cancel Subchannel)	All
17MB/s transfer rate	CNC/CTC
DRM	BL/BY/CVC/CNC/CTC/CFS
CFS channel (ISCH2)	CFS
CFR channel (ISCH2)	CFR
Open Systems Connection Feature (OCF)	OSCH
Fibre Channel for IDK	FBCH*
CHP microcode online patch	All
Channel microcode online patch	All
Channel package concurrent service	All
Channel package concurrent install/deinstall	All

* Strictly, defined as FCP (Fibre Channel Protocol channel) when referred to with the subject microcode included, and viewed as CNC from the operating system.

4.5 Channel Subsystem Configuration

4.5.1 Channel Subsystem Capacities

Table 4-2 shows channel and I/O configuration limits for the F-7C and F-7E processor models.

Table 4-2 Channel Subsystem Capacities

Category	*Contents*
Max. No. of subchannels	36,864/IOCDS 81,920/HSA
Max. No. of Logical CUs (LCUs)	4,096/IOCDS 10,240/HSA
No. of channels/LCU	8
Max. No. of physical CUs (PCUs)	8,192/system 48/parallel (BL, BY) 120/CNC 120/CTC 48/CVC 1/CFS 1/CFR 1/OCF (OSCH) 8/device
Max. No. of channels/PCU	8
Max. No. of devices	256/parallel (BL, BY) 1,024/CNC 512/CTC 256/CVC 256/CFS 1/CFR 254/OCF (OSCH)*
Max. No. of links	120/CNC 120/CTC

* Ethernet PK uses 4 UAs; FDDI PK uses 2 UAs.

4.5.2 CHPID Assignments

Table 4-3 summarizes the number of channels available depending on the physical frame configuration.

Tables 4-4 and 4-5 summarize channel configuration limits by physical frame and identify which CHPIDs are configurable for each type of channel. IZPIOCP 1.7 or higher is required.

Table 4-3 Number of Channels Available

Processor Model	*Channel Type*	*No. of Channels Min–Max [Increment]*			
		Basic Frame, 1 IOP	*Basic Frame, 2 IOPs*	*Basic Frame + Ex-1, 3 IOPs*	*Basic Frame + Ex-1, 4 IOPs*
Base Model and Turbo Model	Parallel	0–32* [4]	0–64* [4]	0–96* [4]	0–96 [4]
	Recommended	0–24 [4]	0–48 [4]	0–72 [4]	0–96 [4]
	ESCH	0–64 [4]	0–128 [4]	0–192 [4]	0–256 [4]
	ISCH2	0–12 [2]	0–24 [2]	0–32 [2]	0–32 [2]
	OSCH-E/FD/FS	0–12 [1]	0–12 [1]	0–12 [1]	0–12 [1]
	FBCH (FCP)**	0–16 [2]	0–16 [2]	0–16 [2]	0–16 [2]
	Total	8–64	8–128	8–192	8–256
CF Model	ISCH2	2–16 [2]	2–32 [2]	2–32 [2]	2–32 [2]

* When this maximum number of parallel channels (8 CH-PKs per IOP) are installed, the resulting I/O connector layout blocks the field engineer's concurrent maintenance services. To avoid such situations, it is recommended not to let more than two parallel CH-PKs neighbor each other.

** Only applicable to F-7C Base Models other than B2C models and F-7E Turbo Models other than 82E and D2EH models.

Note: Every increment denotes one CH-PK. Each IOP accommodates 16 CH-PKs.

Table 4-4 CHPID Assignments on F-7C Base, F-7C Turbo, and F-7E Turbo

Frame	*IOP#*	*CH-PK Slot #*	*CHPID*				
			Parallel	*ESCH*	*ISCH2*	*OCF*	*FBCH**
Basic	IOP 0	0	00–03	00–03		00	00, 02
		1	04–07	04–07	04, 05	04	04, 06
		2	08–0B	08–0B	08, 09	08	08, 0A
		3	0C–0F	0C–0F	0C, 0D	0C	0C, 0E
		4	10–13	10–13		10	10, 12
		5	14–17	14–17		14	14, 16
		6	18–1B	18–1B		18	18, 1A
		7	1C–1F	1C–1F		1C	1C, 1E
		8		20–23		20	20, 22
		9		24–27	24, 25	24	24, 26
		A		28–2B	28, 29	28	28, 2A
		B		2C–2F	2C, 2D	2C	2C, 2E
		C		30–33		30	30, 32
		D		34–37		34	34, 36
		E		38–3B		38	38, 3A
		F		3C–3F		3C	3C, 3E
	IOP 1	0	40–43	40–43		40	40, 42
		1	44–47	44–47	44, 45	44	44, 46
		2	48–4B	48–4B	48, 49	48	48, 4A
		3	4C–4F	4C–4F	4C, 4D	4C	4C, 4E
		4	50–53	50–53		50	50, 52
		5	54–57	54–57		54	54, 56
		6	58–5B	58–5B		58	58, 5A
		7	5C–5F	5C–5F		5C	5C, 5E
		8		60–63		60	60, 62
		9		64–67	64, 65	64	64, 66
		A		68–6B	68, 69	68	68, 6A
		B		6C–6F	6C, 6D	6C	6C, 6E
		C		70–73		70	70, 72
		D		74–77		74	74, 76
		E		78–7B		78	78, 7A
		F		7C–7F		7C	7C, 7E

(Continued on next column)

* FCP channel type dedicated to optional Integrated Disk available on F-7C Base Models except B2C and F-7E Turbo Models except 82E and D2EH. FBCH (FCP) should be installed to every two CH-PK slots of (0, 1), (2, 3), … (E, F). When FBCH (FCP) is installed to one slot of any such pair, the other slot is only usable for FBCH (FCP).

Table 4-4 CHPID Assignments on F-7C Base, F-7C Turbo, and F-7E Turbo (Cont.)

Frame	*IOP#*	*CH-PK Slot #*	*CHPID*				
			Parallel	*ESCH*	*ISCH2*	*OCF*	*FBCH**
Ex-1	IOP 2	0	80–83	80–83		80	80, 82
		1	84–87	84–87	84, 85	84	84, 86
		2	88–8B	88–8B	88, 89	88	88, 8A
		3	8C–8F	8C–8F	8C, 8D	8C	8C, 8E
		4	90–93	90–93		90	90, 92
		5	94–97	94–97		94	94, 96
		6	98–9B	98–9B		98	98, 9A
		7	9C–9F	9C–9F		9C	9C, 9E
		8		A0–A3		A0	A0, A2
		9		A4–A7	A4, A5	A4	A4, A6
		A		A8–AB	A8, A9	A8	A8, AA
		B		AC–AF	AC, AD	AC	AC, AE
		C		B0–B3		B0	B0, B2
		D		B4–B7		B4	B4, B6
		E		B8–BB		B8	B8, BA
		F		BC–BF		BC	BC, BE
	IOP 3	0	C0–C3	C0–C3		C0	C0, C2
		1	C4–C7	C4–C7	C4, C5	C4	C4, C6
		2	C8–CB	C8–CB	C8, C9	C8	C8, CA
		3	CC–CF	CC–CF	CC, CD	CC	CC, CE
		4	D0–D3	D0–D3		D0	D0, D2
		5	D4–D7	D4–D7		D4	D4, D6
		6	D8–DB	D8–DB		D8	D8, DA
		7	DC–DF	DC–DF	DC, DD	DC	DC, DE
		8		E0–E3	E0, E1	E0	E0, E2
		9		E4–E7	E4, E5	E4	E4, E6
		A		E8–EB		E8	E8, EA
		B		EC–EF		EC	EC, EE
		C		F0–F3		F0	F0, F2
		D		F4–F7	F4, F5	F4	F4, F6
		E		F8–FB	F8, F9	F8	F8, FA
		F		FC–FF	FC, FD	FC	FC, FE

* FCP channel type dedicated to optional Integrated Disk available on F-7C Base Models except B2C and F-7E Turbo Models except 82E and D2EH. FBCH (FCP) should be installed to every two CH-PK slots of (0, 1), (2, 3), … (E, F). When FBCH (FCP) is installed to one slot of any such pair, the other slot is only usable for FBCH (FCP).

Table 4-5 CHPID Assignments on F-7C CF

Frame	*IOP#*	*CH-PK Slot #*	*CHPID for ISCH2*	*Frame*	*IOP#*	*CH-PK Slot #*	*CHPID for ISCH2*
Basic	IOP 0	0		Ex-1	IOP 2	0	
		1	04, 05			1	84, 85
		2	08, 09			2	88, 89
		3	0C, 0D			3	8C, 8D
		4				4	
		5	E0*, E1*			5	
		6	E4*, E5*			6	
		7				7	
		8				8	
		9	24, 25			9	A4, A5
		A	28, 29			A	A8, A9
		B	2C, 2D			B	AC, AD
		C				C	
		D				D	
		E				E	
		F				F	
	IOP 1	0			IOP 3	0	
		1	44, 45			1	C4, C5
		2	48, 49			2	C8, C9
		3	4C, 4D			3	CC, CD
		4				4	
		5	F4*, F5*			5	
		6	F8*, F9*			6	
		7				7	DC, DD
		8				8	E0, E1
		9	64, 65			9	E4, E5
		A	68, 69			A	
		B	6C, 6D			B	
		C				C	
		D				D	F4, F5
		E				E	F8, F9
		F				F	FC, FD

* Can be assigned when all the other CHPIDs in Basic frame are used up and still more ISCH2s need to be installed in Basic frame.

CHAPTER 5 SERVICE PROCESSOR AND CONSOLE DEVICE

5.1 Overview of SVP/CD Subsystem

The Service Processor (SVP) and the Console Device (CD) constitute a subsystem that provides the means for operating, configuring, and servicing the F-7C and F-7E processor systems. This SVP/CD subsystem allows operators to communicate with the F-7C and F-7E hardware and firmware through standard display, keyboard, and mouse devices.

The SVP/CD subsystem manages all hardware consoles and all interfaces to hardware alarms and status displays. The subsystem also controls power on/off activity, configuration changes, and service-related activities such as system diagnostics, error recovery logging, and maintenance.

The SVP/CD subsystem does not manage channel-attached (3270-type) consoles used for communication with system control programs (SCPs). However, the subsystem supports console integration, allowing the hardware to receive and to reply to operating system messages. This is primarily intended for situations where channel-attached SCP consoles are unavailable.

The primary components of the SVP/CD subsystem are:

- SVP: The SVP manages, monitors, and controls all the processor subsystems. It stores system configuration information and provides communication between the processor subsystems and the CDs. One SVP is standard with each system of the F-7C or the F-7E.

- CDs: These are the PC-based consoles which provide operator-to-system interface. Operators control the F-7C and F-7E systems using CDs. CDs provide a graphical user interface (GUI) environment for dialogs with the operator. They also provide the support for remote service. CDs support logon security protection to limit an operator's scope of actions to those appropriate to their job. The customer may define access levels for each individual person authorized to operate the F-7C or F-7E processor system. The access levels define that person's scope of activity. One CD is standard with each processor system of the F-7C or the F-7E. An additional one or two CDs may optionally be added.

Figure 5-1 shows the physical positioning of the console subsystem elements. The SVP resides in the Processor Unit frame(s), specifically the Basic Frame. The CDs are external to the Processor Unit frame(s). The one standard CD must be located in the machine room within 30 meters of the Basic Frame. The optional CDs can be located wherever convenient to the customer, but must be within 100 meters of the Basic Frame.

Through the optional HMC/PMC interface, multiple systems may be operated from a single console. Such single console may be IBM's Hardware Management Console (HMC) modified with Hitachi CPN Driver Code that can operate multiple F-7, F-9 and HCF7 systems as well as IBM 9672 and 9674 systems. Or, such single console may be Hitachi's Processor Management Console (PMC) that can operate multiple F-7, F-9 and HCF7 systems.

This single-point-of-control (SPOC) operation requires a dedicated CD for service functions on the F-7, F-9, or HCF7 system subject to the SPOC operation. Up to four dedicated CDs each providing service functions for the F-7, F-9, or HCF7 system subject to the SPOC operation may be located in an optional machine room console frame.

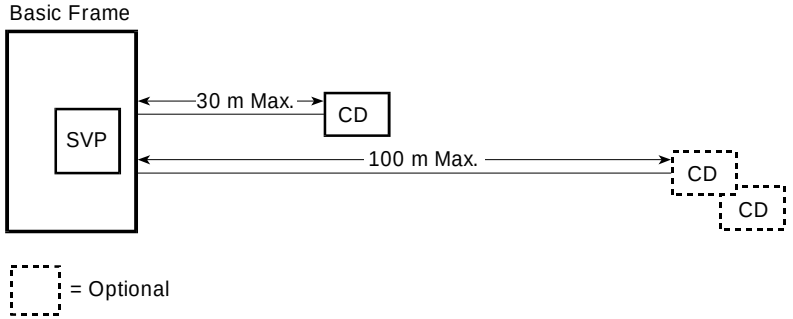

Figure 5-1 Console Subsystem Physical Positioning

5.2 SVP Hardware Design and Features

5.2.1 SVP Hardware Design and Feature Summary

The SVP is located in the Basic Frame and contains a microprocessor, hard disk drives, interface paths to other subsystems, and Ethernet connectivity to the CDs. The SVP supports system configuration and monitoring, power control, internal logging, error recovery, diagnostic, and maintenance functions of the F-7C and F-7E processor systems.

The SVP physically and logically interfaces to the CDs, optional Remote PCI Controllers (RPCs), and other subsystems. Figure 5-2 summarizes the general physical components and features related to the SVP.

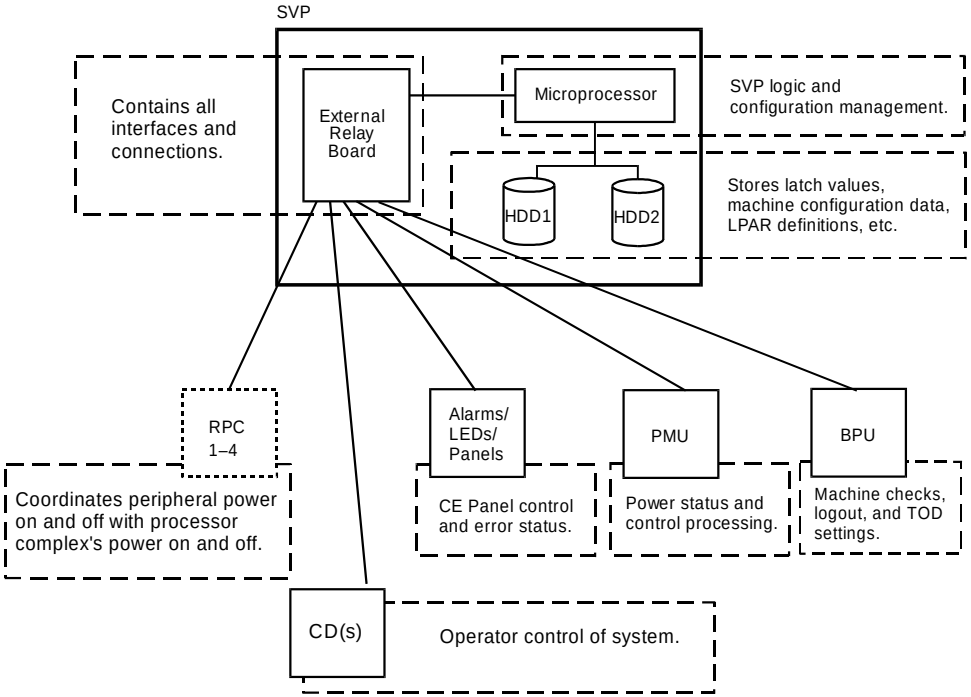

Figure 5-2 SVP Hardware Design and Feature Summary

5.2.2 SVP Interface Design Notes

The SVP contains all the interfaces to support the following hardware units.

- BPU: The SVP interfaces with the Basic Processing Unit (BPU) in support of system control, system configuration, and hardware error recovery such as system checks, logging activities, and time-of-day (TOD) changes and synchronization. The BPU resides with the SVP in the Basic frame.

- Power Management Unit (PMU): The F-7C and F-7E processor systems have multiple PMUs which supply and control power (including power on/off processing) to major system components such as the BPU and IOPs. The SVP interfaces with the PMUs for monitoring, controlling, and powering on and off. Through the PMUs, the SVP supervises power status and control processing.

- Alarms, LEDs, and Panels: The Basic Frame contains several types of alarms. These alarms are LEDs (electronic displays) and panels that communicate codes and status information about system-related errors and features, and also enable service personnel to monitor a system without console support. The SVP updates the information presented by these elements.

- Console Devices (CDs): The SVP is the sole interface between the CDs or system consoles and all other system components. The consoles are used to operate, diagnose, and maintain the F-7C and F-7E processor systems. Message and status traffic between the system and the consoles flows through the SVP.

- Remote PCI Controllers (RPCs): The SVP manages the four optional Remote PCI Controllers when they are installed. Each RPC controls the power flow to as many as 32 peripheral devices. The SVP ensures that each RPC synchronizes its peripherals' power on/off actions with the system's power on/off actions.

5.2.3 SVP Logical Functions

The SVP directs BPU functions such as system and LPAR configuration, IP state (such as started and stopped), LPAR states (such as active and inactive), storage configuration (such as main versus expanded), storage allocation to LPARs, and channel states and allocations to LPARs. The SVP provides communication between the CDs and other subsystems. It continually supervises and monitors each unit's status, stores any error detection data, and reports the results. The SVP provides the file system for storing system firmware, configuration data, and various logs. The following describe major functions supported by the SVP.

- Firmware storage and distribution: The SVP is the repository for all system firmware except firmware used in the CDs. As such, the SVP stores and maintains multiple firmware levels concurrently, distributes firmware to other system functional units during IML or maintenance operations, and records the status and levels of firmware installed in the system. The SVP firmware repository is a staging area for firmware patches and upgrade levels prior to their installation.

- Concurrent maintenance: Performing critical functions related to concurrent maintenance (CM) of the processor system, the SVP directs and controls CM activities, records and monitors information about system parts status and maintenance levels, and distributes firmware patches to concurrently maintainable system units.

- Non-concurrent maintenance: When maintenance cannot be done concurrently, the SVP provides diagnostic and analytical tools for use by service personnel in problem determination and resolution. Similar to concurrent maintenance, the SVP records system parts status and maintenance levels. It also stores and distributes firmware patches and major upgrades to system firmware which cannot be concurrently applied.

- Error recording and monitoring: The SVP monitors and controls the IPs and records system checks and IP logouts. It continually monitors the other system components to ensure proper operating states and logs improper states. Under certain error conditions, the SVP directs the CDs to contact the Support Center to report problems ("phone home").

- Other Functions: The SVP also provides the following functions.
  - Performs file control, troubleshooting, and CE panel control.
  - Sets the time-of-day clocks of the CPU and SVP.
  - Monitors hardware tracing.

5.3 CD Hardware Design and Features

5.3.1 CD Hardware Design

The CD consists of a personal computer (PC) and proprietary system and application software. It is the hardware and software platform used by the operator and service personnel.

Each CD consists of a personal computer with a main processor, a color monitor, a keyboard and a mouse as illustrated in figure 5-3. Depending on whether the CD is powered in synchronism with the processor system through the feature called CD-PCI, the main processor will be a bare PC or a PC housed in an external cabinet. An optional printer may also be connected to each console.

The CD's 17-inch high resolution monitor presents information using graphical user interface (GUI) displays and character displays in visual units called windows. Both the keyboard and the mouse are used for data entry, command entry, icon selection, and dialog interactions. System switches and status displays are represented as a series of graphic icons in each window.

Table 5-1 lists the major hardware components of the F-7C and F-7E Processor Group CDs.

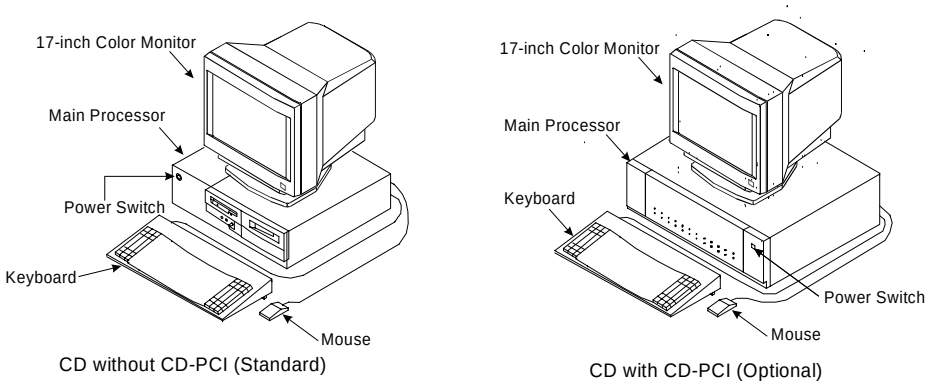

Figure 5-3 CD Physical Components

Table 5-1 CD Hardware Components

Component	*Characteristics*
Main processor	• Hard disk drive • Magneto-optical (MO) disk drive for 3.5 inch MO disks • Floppy disk drive for 3.5 inch floppy disks • Dedicated power on/off capability • Cabinet housing with optional CD-PCI • Power Control Interface (PCI) with optional CD-PCI
Color monitor	• 17-inch screen size • High-contrast, high-quality screen with minimum flicker and reflection • SVGA resolutions of 1024 x 768 (64,000 colors) • Physically adjustable vertically and horizontally
Keyboard	• 101-type keyboard • Movable with coil cord • Adjustable in two angles
Mouse	• 400 dpi resolution
Input power plug cords	• Selectable between 100–120 volts and 200–240 volts input voltages • Equipped with NEMA- or EN60320-complying input power plug

5.3.2 CD Operating Modes

Each CD functions in either of the operating modes defined below, although not at the same time. Which operating mode is activated depends solely on the access level of the userid at logon time.

- Management Console Device (MCD): This mode is meant for customer operation of the processor system. This includes control functions that affect the entire complex. A CD in MCD mode controls functions such as power on/off, LPAR definitions, and configuration control. Multiple MCDs may be active simultaneously.

- Service Console Device (SCD): This mode is meant for service of the processor system by service personnel. A CD in SCD mode supports service console functions for maintenance, diagnostics, and repair. Only one SCD may be active at a time.

5.3.3 CD Features

Each CD provides the following features.

Console functions

For data center personnel, the CD provides the man–machine interface. Console functions provided include:

- Powering the system on and off
- IPLing operating systems (on F-7C Base Models, F-7C Turbo Models, and F-7E Turbo Models)
- Switching the system's operating modes between LPAR mode and Basic mode (on F-7C Base Models, F-7C Turbo Models, and F-7E Turbo Models)
- Reconfiguring LPAR resources
- Activating/deactivating LPARs
- Monitoring system activity
- Online, context-sensitive help
- Administration and control of the CDs (for example, access control for security, online/offline state of each CD)

External interfaces

The CD provides the following interfaces for use by such features as automation software and remote diagnostics.
- Modem interface for connection to a support center
- Modem interface for remote operation for service connection
- Optional token-ring interface for HMC or PMC single-point-of-control operations of multiple systems

Device interfaces

The CD provides other interfaces to other devices, including:
- Integrated 3.5-inch re-writable MO drive—used to transfer bulk system readable data
- Integrated 3.5-inch floppy disk drive—for transferring small amount of data
- External printer—for hardcopy recording

5.4 User Access Levels (Security) and Operating Mode Definitions

The SVP/CD subsystem has five user access levels to dictate the CD's general operating mode (MCD or SCD) and which access-restricted subset of features within that mode are available to a specific user.

Each user is assigned a userid, a password, and associated user access levels.

All CDs can operate in either SCD or MCD mode, but not simultaneously. The CD's current userid is activated during console logon. To change modes, logoff the current userid and re-logon to the CD with a userid that supports the mode desired.

The CD user access levels are explained in the following subsections.

5.4.1 Operator (Op)—MCD Mode

The Operator generally monitors the running of the system during regular workload processing. This person may need help from others when exceptions or problems occur. This access level includes all functions needed to operate the system and the LPARs. The console operates in MCD mode for this user access level.

5.4.2 Advanced Operator (AO)—MCD Mode

The Advanced Operator has more technical skills than the Operator and usually performs preliminary problem identification. This access level includes all functions needed to completely operate the system and LPARs, and to execute basic system and LPAR configuration functions. The console operates in MCD mode for this user access level.

5.4.3 System Programmer (SP)—MCD Mode

The System Programmer is responsible for system level setup, problem determination, changes, and testing. System programmers are involved with all software-related problems and issues, and coordinate with service personnel for hardware support and maintenance. This access level includes all functions necessary to completely configure and operate the system and LPARs. The console operates in MCD mode for this user access level.

5.4.4 Customer Engineer (CE)—SCD Mode

The Customer Engineer is responsible for vendor support of the system's hardware, software, and firmware at the customer site. This access level includes all functions needed to service and configure the system for the support being provided. The console operates in SCD mode for this user access level.

5.4.5 Access Administrator (AA)—MCD Mode

The Access Administrator controls, defines, and assigns userids, passwords, and user access levels. For security, this level's administrative functions are not available from the other user access level. The console operates in MCD mode for this user access level.

CHAPTER 6 SUPPORTED FEATURES AND FACILITIES

The F-7C and F-7E processor models support full ESA/390 functionality, comprehensive RAS capabilities, and other major functions through the implementation and use of specific features and facilities. The following sections present details about the features and facilities available on the F-7C and F-7E processor models.

These features are categorized as follows:
- ESA/390 features dependent on operation mode
- Features not dependent on mode
- Reliability, availability, and serviceability (RAS) features

For additional information on these features and facilities, refer to *IBM Enterprise Systems Architecture/390 Principles of Operation.*

6.1 List of ESA/390 Features Dependent on Operation Mode

Table 6-1 lists ESA/390 features dependent on operation mode applicable to the F-7C Base Models, F-7C Turbo Models, and F-7E Turbo Models. All these models implement ESA/390 architecture in both Basic and LPAR modes and as VM guests.

Legend:
Std = Standard
Opt = Optional
ND = Not defined in the architecture and, therefore, not implemented
NI = Not implemented (but defined in the architecture)
NS = Not supported

Table 6-1 ESA/390 Features Dependent on Operation Mode

Feature	*Basic*	*LPAR*	*Comments*
ESA/390 architecture	Std	Std	
System/370 architecture	ND	ND	
Advanced Address Space	Std	Std	
Asynchronous Data Mover Facility (ADMF)	Std	Std	
Asynchronous Page-out Facility	Std	Std	
Basic Control mode (BC-mode)	ND	ND	
Bimodal Addressing	Std	Std	
Branch and Save	Std	Std	
Branch and Set Authority Facility	Std	Std	
Broadcasted Purging	Std	Std	
Byte-oriented Operand	Std	Std	
Called Space Identification	Std	Std	
Cancel I/O	Std	Std	
Cancel Subchannel	See Cancel I/O.		
Channel Indirect Data Addressing	Std	Std	
Channel Subsystem	Std	Std	
Channel Subsystem Call	Std	Std	
Checksum Facility	Std	Std	
Clear I/O	ND	ND	*
Command Retry	Std	Std	
Compare-and-Move Extended	Std	Std	
Concurrent Sense	Std	Std	
Conditional Swapping	Std	Std	
Console Integration	Std	Std	
Coupling Channel	See Inter-System Coupling Channel (ISCH).		
Coupling Facility	Std	Std	
Coupling Facility Channel	See Inter-System Coupling Channel (ISCH).		
Coupling Links	See Inter-System Coupling Channel (ISCH).		
CPU Timer and Clock Comparator	Std	Std	
Data Compression (hardware assisted)	Std	Std	
Data Streaming	Std	Std	
DB2 sort facility	See Extended Sorting.		
Dynamic Channel Subsystem	See Channel Subsystem.		

(Continued on next page)

* Similar functionality and S/370 CCW compatibility are provided by channel subsystem.

Table 6-1 ESA/390 Features Dependent on Operation Mode (Cont.)

Feature	*Basic*	*LPAR*	*Comments*
Dynamic I/O Configuration	See Dynamic I/O Reconfiguration Management (DRM).		
Dynamic I/O Reconfiguration Management (DRM)	Std	Std	
Dynamic Storage Reconfiguration (DSR)	See Dynamic Storage Reconfiguration (DSR) under MLPF.		
EC-mode	See Extended Control Mode (EC-mode) under Translation.		
Enhanced DSR	See Enhanced DSR under MLPF.		
Enhanced Move-page (MVPG-2)	See Move Page Facility 2.		
ESA/390 Standard Channel Subsystem	See Channel Subsystem.		
ESCON	See Extended Serial Channel (ESCH).		
ESCON Channel-to-Channel Adapter	See Extended Serial Channel (ESCH).		
ESCON Multiple Image Facility (EMIF)	NS	Std	
Expanded Storage	Std	Std	Processor storage is partitionable into main storage and expanded storage.
Extended Control mode (EC-mode)	See Extended Control mode (EC-mode) under Translation.		
Extended Precision Divide	Std	Std	
Extended Precision Floating Point	Std	Std	
Extended Real Addressing (26-bit)	ND	ND	Replaced by 31-bit real addressing.
Extended Serial Channel (ESCH)	Opt	Opt	
Extended Sorting	Std	Std	
Extensions for Virtual Machines	Std	Std	
External Timer Attachment Feature (ETAF)	Opt	Opt	
Fast Release	ND	ND	*
Floating Point	Std	Std	
Halt Device	ND	ND	*
I/O Device Self-Description	Std	Std	
I/O Error Alert	Std	Std	
I/O Interface Reset	Std	Std	
I/O Power Sequence Control	Opt	Opt	
Immediate-and-Relative Instructions	Std	Std	

(Continued on next column)

* Similar functionality and S/370 CCW compatibility are provided by channel subsystem.

Table 6-1 ESA/390 Features Dependent on Operation Mode (Cont.)

Feature	*Basic*	*LPAR*	*Comments*
Incorrect Length Indication Suppression	Std	Std	
Integrated Coupling Facility (IntCF)	NS	Opt	
Integrated Coupling Migration Facility (ICMF)	ND	Std	
Interpreted SIE	See Extensions for Virtual Machines.		
Interpretive Execution Facility (IEF)	Std	Std	
Inter-System Coupling Channel (ISCH)	Opt	Opt	
Interval timer	ND	ND	
Key-controlled Storage Protection	Std	Std	Functionality of S/370's ISK, RRB, and SSK instructions is replaced by ESA/390's storage-key instruction extensions.
Limited Channel Logout	ND	ND	*
Logging Volume Reduction	Std	Std	
Logical String Assist	Std	Std	
Monitoring	Std	Std	
Move Page Facility 1	Std	Std	
Move Page Facility 2	Std	Std	
Multiple High Performance Guests	Std	Std	
Multiple Logical Processor Feature (MLPF)			
15 LPARs	NS	Std	
16 GB Main Storage Support	NS	Std	
Dynamic Storage Reconfiguration (DSR)	NS	Std	
Enhanced DSR	NS	Std	
LPAR Management Time Reporting	NS	Std	
LPAR Preferred Path	NS	Std	
LPAR Processor Weight Management	NS	Std	
LPAR Support for Logical IP Vary Off	NS	Std	
MLPF Policy	NS	Std	
Resource Capping	NS	Std	
Time Machine	NS	Std	
Time Warp	NS	Std	

(Continued on next page)

* Similar functionality and S/370 CCW compatibility are provided by channel subsystem.

Table 6-1 ESA/390 Features Dependent on Operation Mode (Cont.)

Feature	*Basic*	*LPAR*	*Comments*
Multiprocessing			
CPU Address Identification	Std	Std	
CPU Signaling and Response	Std	Std	
Prefixing	Std	Std	
Shared Main Storage	Std	Std	
TOD Clock Synchronization	Std	Std	
Open Systems Connection Feature (OCF)	Opt	Opt	
Open Systems Adapter (OSA)	See Open Systems Connection Feature (OCF).		
Page Protection	Std	Std	
PER Extensions	Std	Std	
Perform Locked Operation Facility	Std	Std	
Private Space	Std	Std	
PSW Key Handling	Std	Std	
Recovery Extensions	ND	ND	
Scalar Square Root	Std	Std	
SCP-initiated Reset	See I/O Interface Reset.		
Segment Protection	ND	ND	Replaced by page protection.
Service Signal	Std	Std	
Set Address Space Control Fast (SACF)	Std	Std	
SIGP Assists	See Extensions for Virtual Machines.		
Sorting instructions	Std	Std	
Storage-key 4 KB block			
Single-key 4 KB block	Std	Std	
Storage-key exception control	ND	ND	ESA/390 implements functions necessary to manage storage keys by the storage-key instruction extensions.
Storage-key instructions ISK, RRB, SSK	ND	ND	
Storage-key instruction extensions	Std	Std	
Subspace Group Facility	Std	Std	
Subsystem Storage Protection (SSSP)	Std	Std	
Suppression on Protection	Std	Std	
Sysplex Timer	See External Timer Attachment Feature (ETAF).		

(Continued on next column)

Table 6-1 ESA/390 Features Dependent on Operation Mode (Cont.)

Feature	*Basic*	*LPAR*	*Comments*
System/370 Extended Facility			
MVS-dependent			
4 lock-handling instructions	Std	Std	
6 tracing instructions	ND	ND	
Fix page instruction	ND	ND	
SVC assist instruction	Std	Std	
Add FRR instruction	Std	Std	
Non-MVS dependent			
Low address protection	ND	ND	
IPTE instruction	Std	Std	
TPROT instruction	Std	Std	
System/370 I/O instructions	ND	ND	*
Test Block	Std	Std	
Time-of-day (TOD) clock	Std	Std	
Tracing (ASN, branch, and explicit)	Std	Std	
Translation			
Dynamic Address Translation (DAT)			
2 KB page size	ND	ND	
4KB page size	Std	Std	
64KB segment size	ND	ND	
1MB segment size	Std	Std	
Extended Control mode (EC-mode)	ND	ND	ESA/390 operation is comparable to S/370 EC mode.
Program Event Recording 2 (PER 2)			
Set-system-mask suppression	Std	Std	
Store status	Std	Std	
Virtual Machine Assists under SIE	NI	NI	
3033 Extension			
Dual Address Space (DAS)	Std	Std	Except DAS tracing. ASN tracing is used instead.
Start I/O Fast (SIOF) Queuing	ND	ND	*
Suspend and Resume	ND	ND	*
31-bit IDAWs	Std	Std	
31-bit Real Addressing	Std	Std	

* Similar functionality and S/370 CCW compatibility are provided by channel subsystem.

6.2 List of Features Not Dependent on Mode

Table 6-2 lists the F-7C and F-7E features not dependent on architecture/operation mode. (Std = Standard, Opt = Optional)

Table 6-2 Features Not Dependent on Mode

Feature	*F-7C*			*F-7E*
	Base	*Turbo*	*CF*	*Turbo*
Cache	Std	Std	Std	Std
High Speed Buffer Storage	See Cache.			
HMC Connectivity	Opt	Opt	Opt	Opt
Multiple High Performance Guests	Std	Std	Std	Std
PMC Connectivity	Opt	Opt	Opt	Opt
TSCF target support	Opt	Opt	Opt	Opt

6.3 List of Features Contributing to RAS

Table 6-3 lists the F-7C and F-7E features which support and contribute to the high levels of RAS of the F-7C and F-7E processor models. Some of these features may have also been listed in the previous sections.

Legend:
Std = Standard
Opt = Optional
NA = Not applicable
ND = Not defined in the architecture and, therefore, not implemented
NS = Not supported

Table 6-3 RAS Features

Feature	*F-7C*			*F-7E*
	Base	*Turbo*	*CF*	*Turbo*
Application Preservation - Basic mode	Std	Std	ND	Std
Application Preservation - MLPF	Std	Std	Std	Std
Application Preservation - LPARs under MLPF	Std	Std	ND	Std
Command Retry	Std	Std	ND	Std
Concurrent Maintenance of				
Channels	Std	Std	Std	Std
Fans	Std	Std	Std	Std
Firmware patches	Std	Std	Std	Std
MLPF patches	Std	Std	Std	Std
Power supplies	Std	Std	Std	Std
Concurrent Upgrade of Channels	Std	Std	Std	Std
CP/SAP Sparing	See IP Alternation and SP Alternation.			
Dynamic Coupling Facility Dispatching	Std	Std	Std	Std
Dynamic I/O Reconfiguration Management (DRM)	Std	Std	ND	Std
Dynamic Memory Sparing	See Memory Array Alternation.			
Dynamic SAP Sparing/Reassignment	See SP Alternation.			
Dynamic Storage Configuration (DSR)	See MLPF Dynamic Storage Reconfiguration.			
Enhanced DSR	See MLPF Enhanced DSR.			
Error Checking and Correction	Std	Std	Std	Std
I/O Error Alert	Std	Std	ND	Std
I/O Interface Reset	Std	Std	ND	Std
IP Alternation - static or dynamic	Std	Std	Std	Std
Memory Array Alternation	Std	Std	Std	Std
MLPF Dynamic Storage Reconfiguration (DSR)	Std	Std	ND	Std
MLPF Enhanced DSR	Std	Std	ND	Std
Monitoring	Std	Std	ND	Std
PER Extensions	Std	Std	ND	Std
Processor Availability Facility (PAF)	Std	Std	ND	Std
Recovery Extensions	ND	ND	ND	ND
Redundant Components				
Fans	Std	Std	Std	Std
Power Supplies	Std	Std	Std	Std
Refrigeration Units	NA	Std	NA	Std
Remote Support	Std	Std	Std	Std
SP Alternation - dynamic	Std	Std	Std	Std
Subspace Group Facility	Std	Std	ND	Std
Subsystem Storage Protection (SSSP)	Std	Std	ND	Std

6.4 Feature Definitions

- **Advanced Address Space:** Enables a program and its data to reside in separate address spaces. This facility includes the following components:
  - Access register mode: Allows many instructions to use both general and access registers to have fast access to many different address spaces.
  - Linkage stack mechanism: Passes control between programs in different address spaces.
  - Home-space mode: Used by the control program, this component efficiently accesses dispatchable unit control blocks.
  - Real-storage access: Supports two control program instructions (Load and Store Using Real Storage) for efficient real storage access.
  - Data Movement: Improves performance of moves between fetch protected areas through implementation of the MOVE WITH DESTINATION KEY (MVCDK) and MOVE WITH SOURCE KEY (MVCSK) instructions.
- **Application Preservation - Basic mode:** On a system running in Basic (non-LPAR) mode, this feature allows work in progress on a failed IP to be transferred to a running IP non-disruptively.
- **Application Preservation - MLPF:** On a system running in LPAR mode, this feature allows MLPF work in progress (for example, the MLPF Hypervisor) on a failed IP to be transferred to a running IP non-disruptively.
- **Application Preservation - LPARs under MLPF:** On a system running in LPAR mode, this feature allows LPAR work in progress (for example, a job in an LPAR under MLPF) on a failed IP to be transferred to a running IP non-disruptively.
- **Asynchronous Data Mover Facility (ADMF):** Frees the IP from managing data movement between the main and expanded storage areas. With ADMF, the SP controls the page movement.
- **Asynchronous Pageout Facility:** Enables a main-storage-to-expanded-storage move to occur without delaying subsequent processing by releasing the processor when a page transfer starts. In contrast, the synchronous pageout mode locks the processor while a page or block is being transferred.
- **Basic Control mode (BC-mode):** In System/370 mode, BC-mode provides the facilities necessary to execute legacy application programs intended to operate in System/360 environments. Included in such facilities is a Program Status Word (PSW) compatible with System/360. Most significantly, BC-mode does not provide Dynamic Address Translation and, therefore, lacks in virtual addressing and paging capabilities. In BC-mode, instructions and logical addresses are treated as real addresses (i.e., not translated). Contrast with Extended Control mode (EC-mode).
- **Bimodal Addressing:** Processing support for two modes of logical instruction addressing:
  - 24-bit addressing mode for the execution of S/370 programs
  - 31-bit addressing mode for the execution of XA and ESA/390 programs
- **Branch and Save:** Two formats (BAS, BASR) of instruction used for linking to programs known to be in the same addressing mode as the caller.
- **Branch and Set Authority Facility:** Consists of BRANCH AND SET AUTHORITY (BSA) instruction that replaces PROGRAM CALL, PROGRAM TRANSFER, and SET PSW KEY FROM ADDRESS instructions and provides enhanced performance.
- **Broadcasted Purging:** Allows for the following capabilities.
  - Conditional updates of address translation tables
  - Efficiently purging address-translation lookaside buffers (TLBs) in systems with more than one IP
- **Byte-oriented Operand:** Applicable to logical, fixed-point, and floating-point operand data, this feature prevents specification exceptions or program interruptions resulting from storage operand mis-alignments encountered by unprivileged instructions. The operands may be on any byte boundaries rather than the boundaries implied by the operand type (e.g., halfword, fullword).
- **Cache:** Specialized hardware memory units which lie between IPs and main storage. Cache has much faster access times than main storage and serves to improve performance by reducing the effective delays of fetching instructions and data from main storage into IPs. Also known as High Speed Buffer Storage.
- **Called Space Identification:** Adds space (i.e., address space, data space) identification information to linkage-stack state entries created by certain instances of PROGRAM CALL instruction.
- **Cancel I/O:** Withdraws, by means of CANCEL SUBCHANNEL (XSCH) instruction, a pending start request from a subchannel without signaling the actual device. This ability is exploited in certain types of I/O error recovery.
- **Cancel Subchannel:** See Cancel I/O.
- **Channel Indirect Data Addressing:** Provides a channel command word (CCW) with the ability to select 2K storage area blocks from an indirect-data-address list of designated blocks. This complements dynamic address translation (DAT) ability in CPUs.

- **Channel Subsystem:** Channel-related hardware and processor I/O instructions that direct the flow of information between main storage and I/O devices. Allows instruction processing and I/O activity to proceed concurrently as the IP does not deal directly with each I/O device. When an I/O request is made, the IP forwards the logical request to the channel to offload I/O request queuing, channel path selection, and certain other I/O related functions from the IP. Sometimes called the Dynamic Channel Subsystem or ESA/390 Standard Channel Subsystem.
- **Channel Subsystem Call:** Provides retrieval and update functions used in I/O configuration management activities.
- **Checksum Facility:** Provides CHECKSUM (CKSM) instruction to develop checksums for use in checking the validity of data. Exploited by TCP/IP.
- **Clear I/O:** An obsolete command for stopping an I/O request in the 370 environment. This is replaced by the standard functions of the channel subsystem operating in ESA/390 mode, which maintains System/370 channel program compatibility.
- **Command Retry:** This control-unit-initiated retry allows the channel subsystem to reissue a command without causing an I/O interrupt.
- **Compare-and-Move Extended:** Provides two interruptible instructions COMPARE LOGICAL LONG EXTENDED (CLCLE) and MOVE LONG EXTENDED (MVCLE) for the following purposes:
  - Increase maximum operand size to a 32-bit specification.
  - Interrupt instruction execution whenever approximately 4K of data has been processed.

  The latter allows program polling in multiprocessing environments.
- **Concurrent Maintenance:** Allows components to be maintained or serviced concurrently and non-disruptively. The following components can be concurrently maintained:
  - Channels: All types of channels can be maintained concurrently. Channels are serviced in the increment granularity in which they are packaged.
  - Fans
  - Firmware patches: Some but not all firmware patches can be applied concurrently. All types of firmware patches can be loaded into the system concurrently for later application. System components for which firmware patches can be applied concurrently are BPU, SP, IOP, MLPF (see below for details) and HCCFA.
  - MLPF patches: Some but not all MLPF patches can be applied concurrently. All types of MLPF patches can be loaded into the system concurrently for later application.
  - Power supplies: Some but not all power supplies can be maintained concurrently.
  - Refrigeration Units (F-7C Turbo Models and F-7E Turbo Models only)
- **Concurrent Sense:** Enables unit-check interruptions to present sense information at the time of the interrupt. This eliminates the need to issue a separate I/O to retrieve the interrupt's sense data.
- **Concurrent Upgrade of Channels:** Allows all types of channels to be added to a system concurrently and non-disruptively. Channels are added in the increment granularity in which they are packaged. Channel upgrade requiring additional IOPs or frames is not concurrent.
- **Conditional Swapping:** Makes the instructions COMPARE AND SWAP and COMPARE AND DOUBLE SWAP available for use.
- **Console Integration:** Provides the ability to use the hardware system console for MVS system initialization and recovery of a channel-attached MVS operator console when all other MVS consoles are unusable.
- **Coupling Channel:** See Inter-System Coupling Channel (ISCH).
- **Coupling Facility Channel:** See Inter-System Coupling Channel (ISCH).
- **Coupling Link:** See Inter-System Coupling Channel (ISCH).
- **CP/SAP Sparing:** See IP Alternation and SP Alternation.
- **CPU Timer and Clock Comparator:**
  - CPU Timer: Measures central processor elapsed time. It decrements while the CPU is executing or in a wait state; it does not decrement when the CPU is stopped. When the value is less than zero or at a value specified by the control program, the CPU Timer issues an interrupt. The CPU Timer does not operate when the TOD clock is in a stopped or error state.
  - Clock Comparator: Creates an external interruption when the TOD clock reaches a pre-specified value. Two instructions (SCKC and STCKC) are available for using the Clock Comparator. The Clock Comparator does not operate when the TOD clock is in a stopped or error state.
- **Data Compression (hardware assisted):** Compresses data using hardware facilities. Can reduce external storage requirements and may reduce CPU cycles used by software compression facilities.
- **Data Streaming:** On parallel block multiplexer channels, this feature optimizes the amount of data transferred, thus allowing up to 6 MB/second data rates. Also allows cable length of up to 400 feet between the channel and the control unit. Data streaming operation is initiated by the control unit. Data streaming and non-data streaming control units may be intermixed on a channel.
- **DB2 sort facility:** See Extended Sorting.
- **Dynamic Channel Subsystem:** See Channel Subsystem.
- **Dynamic Memory Sparing:** See Memory Array Alternation.

- **Dynamic I/O Reconfiguration Management (DRM):** Allows the modification of I/O configurations without performing a Power-on Reset or a System Control Program (SCP) Initial Program Loading (IPL). With DRM, the I/O definitions for I/O devices, control units, and channel paths may be added, deleted, or modified non-disruptively. In addition to changing the active I/O configuration, the changes may also be saved in the active I/O configuration dataset (IOCDS). This function is exploited by any of the following Operating Systems.
  - MVS/ESA Version 4 Release 2 and above using Hardware Configuration Definition (HCD)
  - VM/ESA Version 2 Release 1 and above
- **Dynamic IP Alternation:** In the event of a failure on an Instruction Processor (IP), this facility replaces the failed IP with one of the standby Alternate Processors (APs). Replacement of the failed IP is dynamic (non-disruptive). APs do not change the model number of the F-7C or F-7E system. For instance, a 9-way 92C with two APs remains the model 92C. The APs do not affect software licensing charges based on aggregate system capacity.
- **Dynamic SAP Sparing/Reassignment:** See Dynamic SP Alternation.
- **Dynamic SP Alternation:** In the event of a failure on the System Processor (SP), this facility replaces the failed SP with one of the standby Alternate Processors (APs). Replacement of the failed SP is dynamic (non-disruptive).
- **Dynamic Storage Reconfiguration (DSR):** See Dynamic Storage Reconfiguration (DSR) under MLPF.
- **EC-mode:** See Extended Control mode (EC-mode) under Translation.
- **Enhanced DSR:** See Enhanced DSR under MLPF.
- **Enhanced Move-page (MVPG-2):** See Move Page Facility 2.
- **Error Checking and Correction:** Checks data paths using parity or error checking and correction code. Data paths checked are between expanded storage and main storage, and among main storage, channels, and IPs.
- **ESA/390 architecture:** The architecture for computer systems and programs defined in *IBM Enterprise Systems Architecture/390 Principles of Operation*.
- **ESA/390 Standard Channel Subsystem:** See Standard Channel Subsystem.
- **ESCON:** See Extended Serial Channel (ESCH).
- **ESCON Channel-to-Channel Adapter:** See Extended Serial Channel (ESCH).
- **ESCON Multiple Image Facility (EMIF):** Allows channels to be shared concurrently by multiple LPARs. Optionally, the control units and devices on the shared channels can also be shared by the LPARs. ESCH TYPE=CNC, ESCH TYPE=CTC, ISCH TYPE=CFS, and OCF channels can be shared via EMIF. The other channels cannot be shared via EMIF.
- **Expanded Storage:** Provides the ability to use high-speed, high-capacity storage as processor storage that can transfer 4K blocks to and from main storage. Expanded storage is block addressable as opposed to main storage, which is byte addressable.
- **Extended Control mode (EC-mode):** See Extended Control mode (EC-mode) under Translation.
- **Extended Precision Divide:** Instruction (mnemonic: DXR) that allows arithmetic division using extended-precision floating-point operands.
- **Extended Precision Floating Point:** Supports data format consisting of 128-bit data fields with a signed 7-bit characteristic and a 28-digit fraction. This feature adds seven instructions to operate on extended-precision data.
- **Extended Real Addressing (26-bit):** Supports the addressing of real storage from 16MB up to 64MB for use by operating systems only. Virtual addressability remains limited to 16 MB.
- **Extended Serial Channel (ESCH):** ESCON protocol serial transmission bit stream channel providing switched point-to-point connection of I/O devices. Supports, depending on the control unit, data rates up to 17 MB per second. Supports native (CNC), converter (CVC), and channel-to-channel (CTC - both basic and extended mode) channel types. Attaches to 50/125-micron multimode fiber optic cables.
- **Extended Sorting:** Supports instructions which improve the performance of the DB2 sorting function. Exploited in DB2 Version 2 Release 3 or later.
- **Extensions for Virtual Machines:** These extensions consist of five facilities:
  - VM-Data-Space facility that improves the ability of VM applications to address and share large amounts of data through exploitation of the access-register addressing architecture.
  - Storage-key function that eliminates the VM RCP area and thereby improves performance.
  - Interpreted SIE improvement that permits second-level preferred guests to run under a VM operating as a high performance guest.
  - Special-purpose, optional lookaside feature for some guest state information and for expanding some implementation choices.
  - SIGP assist that provides for the interpretation of the signal processor (SIGP) external call and the external call interruption in the SIE environment.

- **External Timer Attachment Feature (ETAF):** Provides an attachment interface for the Sysplex Timer. The Sysplex Timer provides a uniform time source for all of the systems to which it is attached. Uniformity is effected by synchronizing the time-of-day clocks in all the systems. ETAF is required for all systems in a Parallel Sysplex. ETAF itself does not directly provide but indirectly provides the following features of the Sysplex Timer:
  - Automatic Propagation Delay Adjustment: Adjusts processor time-of-day (TOD) clock synchronization to minimize the effect of cable length on the resulting TOD values.
  - External Time Source: Improves the precision of the Sysplex Timer by allowing the timer to receive and adjust to the signals from a stable external time source.
- **Fast Release:** Provides for quick release of the IP that executes Start I/O Fast Release (SIOF) instruction. This action reduces the central processing delay associated with the Start I/O operation.
- **Floating Point:** Provides floating-point instructions, registers, and data formats that support floating-point arithmetic. The floating-point feature supports the short (32-bit) and long (128-bit) data formats.
- **Halt Device:** Provides Halt Device (HDV) instruction, which terminates the current I/O operation of the I/O device addressed.
- **High Speed Buffer Storage:** See Cache.
- **HMC Connectivity:** Allows one or more F-7, F-9 and/or HCF7 systems to be operated from a single-point-of-control Hardware Management Console (HMC). Operation of the Hitachi systems occurs concurrently with operation of IBM systems supported by the HMC.
- **Immediate-and-Relative Instructions:** Provides two new instructions formats (RI and RSI) and thirteen new instructions: ADD HALFWORD IMMEDIATE (AHI), BRANCH RELATIVE AND SAVE (BRAS), BRANCH RELATIVE ON CONDITION (BRC), BRANCH RELATIVE ON COUNT (BRCT), BRANCH RELATIVE ON INDEX HIGH (BRXH), BRANCH RELATIVE ON INDEX LOW OR EQUAL (BRXLE), COMPARE HALFWORD IMMEDIATE (CHI), LOAD HALFWORD IMMEDIATE (LHI), MULTIPLY SINGLE (MSR), MULTIPLY SINGLE (MS), MULTIPLY HALFWORD IMMEDIATE (MHI), TEST UNDER MASK HIGH (TMH), and TEST UNDER MASK LOW (TML).
- **Incorrect Length Indication Suppression:** Provides the incorrect-length indication mode for format-1 channel programs.
- **Integrated Coupling Facility (IntCF):** Optionally provides coupling facility functionality in an LPAR. This allows a logical partition to serve as a coupling facility in an MVS/ESA or OS/390 Parallel Sysplex environment. The IntCF IPs are dedicated to Coupling Facility functions and are not available for general purpose work, such as MVS. Addition of an IntCF reduces the number of Alternate Processor available for Dynamic IP Alternation or Dynamic SP Alternation by one. IntCFs do not change the model number of the F-7C or F-7E system. For instance, a 9-way F-7 model 92C with one IntCF remains a model 92C. Addition of IntCF does not affect software licensing charges based on aggregate system capacity.
- **Integrated Coupling Migration Facility (ICMF):** Provides coupling facility functionality in an LPAR with simulated rather than physical ESCHs. All MVS/ESA or OS/390 systems in the Parallel Sysplex controlled by ICMF must be on the same physical system. Intended for Parallel Sysplex test environment.
- **Interpreted SIE:** See Extensions for Virtual Machines.
- **Interpretive Execution Facility (IEF):** Provides hardware support for certain virtual machine operations. Provides START INTERPRETIVE EXECUTION (SIE) instruction which is used to dispatch virtual machines under VM and LPARs under MLPF.
- **Inter-System Coupling Channel (ISCH):** Special highspeed channels used to connect MVS/ESA and OS/390 systems to a S/390 Coupling Facility in a Parallel Sysplex. ISCHs for the F-7C and F-7E processor systems employ 9/125-micron single-mode fiber optic cables.
- **Interval timer:** Provides program-controlled external interruptions of the IP for purposes of signaling the end of an interval of time. The external interruption occurs when the interval timer decrements from zero to a negative value.
- **I/O Device Self-Description:** Provides support for I/O devices characterized as “self-describing components (SDC).” Allows an SDC to describe its characteristics, its topological position in an ESCON I/O configuration, and the identity of all its nodes.
- **I/O Error Alert:** Allows a channel to be alerted when a control unit malfunction that could affect the control unit’s ability to continue operation occurs.
- **I/O Interface Reset:** Minimizes shared DASD hang conditions by providing a facility for an operating system to reset its I/O configuration prior to entering a disabled wait state after certain check conditions. This function is also called SCP-Initiated Reset.
- **I/O Power Sequence Control:** Provides optional features which allow a system to power on and off peripheral control units.

- **Key-controlled Storage Protection:** Provides store and fetch protection against unauthorized access of information in main storage through the use of storage keys. A storage protection violation results in data not being stored in the protected area. A fetch protection violation results in data not being retrieved from a protected area. Both types of protection violations generate an interrupt.
- **Limited Channel Logout:** Assists model-independent recovery from channel errors by providing four bytes of channel status information.
- **Logging Volume Reduction:** Provides COMPARE UNTIL SUBSTRING EQUAL (CUSE) interruptible instruction. Exploited by IMS/ESA version 3 and above to reduce the volume of logging information for DL/I data bases.
- **Logical String Assist:** Provides three instructions COMPARE LOGICAL STRING (CLST), MOVE STRING (MVST), and SEARCH STRING (SRST) which are useful in C language applications using null ending characters as string delimiters.
- **Memory Array Alternation:** Provides alternate processor storage memory chips which dynamically and non-disruptively replace memory chips which might fail.
- **Monitoring:** Provides MONITOR CALL (MC) instruction. Provides the ability to record certain events during program execution based on selected criteria.
- **Move Page Facility 1:** Provides MOVE PAGE (MVPG) instruction to move a 4 KB page of data between expanded and main storage. Both the source and destination must be valid. Introduced in ESA/370. See also Move Page Facility 2.
- **Move Page Facility 2:** Extends MOVE PAGE instruction of Move Page Facility 1 to provide improved performance under certain conditions. The extended instruction also allows use of an access key for either of the operands. Provides options for VM problem program use. Introduced in ESA/390. See also Move Page Facility 1.
- **Multiple High Performance Guests:** Support for multiple preferred guests of VM. One V=R preferred guest is supported. Six V=F preferred guests are supported (five preferred guests are supported if a V=R guest is running.) The preferred guests, which exhibit higher performance than if they were run non-preferred (V=V), can run concurrently with non-preferred guests.
- **Multiple Logical Processor Feature (MLPF):** Supports logical partitioning in a processor complex, and multiple preferred guests for VM systems. The former allows multiple operating systems to run concurrently on a single hardware system; each operating system runs in a logical partition (LPAR) with dedicated but reconfigurable storage, dedicated or shared IPs, and dedicated or shared channels. Other features of MLPF are shown below.
  - 15 LPARs: Up to 15 LPARs can be active concurrently.
  - 16 GB Main Storage Support: Of the total processor storage in a system, up to 16 GB can be used as main storage by LPARs. No more than 2 GB can be used as main storage for a single LPAR.
  - Dynamic Storage Reconfiguration (DSR): Provides facilities to configure main and expanded storage to and from active LPARs non-disruptively. Storage configured to an active LPAR must be contiguous to and at a higher address than the LPAR's current storage.
  - Enhanced DSR: Provides facilities to configure main and expanded storage to and from active LPARs non-disruptively. The storage to be configured to an active LPAR may be at any address, i.e., need not be contiguous to or at a higher address than the LPAR's current storage.
  - LPAR Management Time Reporting: Includes LPAR time management information in the data available to various Resource Measurement Facility (RMF) reports. Inclusion of this information facilitates capacity planning and resource usage decisions.
  - LPAR Preferred Path: Extends the preferred path function to the LPAR environment. Allows specification of a preferred channel path for a device in LPAR mode.
  - LPAR Processor Weight Management: Provides control mechanisms to manage the amount of IP service delivered to logical partitions by using the following general management criteria:
    - An LPAR's unused IP cycles are made available to other logical IPs using the proportions defined for the other logical IPs.
    - When the service ratio of an LPAR goes under 1 % of the entire system, the LPAR is provided with the approximate service ratio of 1 %.
    - When the resource capping is specified to an HCCF-LPAR, the actual service ratio may sometimes exceed the user-defined value because it is necessary to dispatch service to the HCCF-LPAR regularly.

- **Multiple Logical Processor Feature (MLPF) (Cont.):**
  - LPAR Support for Logical IP Vary Off: Allows logical partitions to reallocate utilization by the IPs in such a manner that an activated LPAR is able to release an IP, using an SCP command, and make the IP available for allocation to another LPAR. The physical IP resource distributed to each online logical IP in an LPAR is determined by dividing the physical IP share for the LPAR by the number of logical IPs online. Each time a logical IP is varied offline or online the distribution is automatically recalculated.
  - MLPF Policy: Provides the LPAR environment pre-scheduled on demand by a program. Also provides the combination of operations to define, control, and reconfigure the MLPF or LPAR automatically.
  - Resource Capping: Provides control mechanisms to limit or cap the amount of IP service which can be delivered to a logical partition.
  - Time Machine: Provides an integrated hardware-based mechanism for finding and analyzing code which may not properly process the year portions of dates due to use of 2-digit year fields.
  - Time Warp: Provides an integrated mechanism for allowing code which is not Year 2000 compliant to properly process dates beyond the year 1999. Intended as a temporary measure until the code is remedied.
- **Multiprocessing**: Permits the use of a multiprocessor configuration highlighted in the following subfunctions.
  - CPU Address Identification: Provides STORE CPU ADDRESS (STAP) instruction and additional external interrupt conditions. One IP can signal certain conditions to another specified IP using conventional SIGNAL PROCESSOR (SIGP) instruction. The other IP can determine the signaling IP using STAP.
  - CPU Signaling and Response: Provides SIGNAL PROCESSOR (SIGP) instruction and the means for one IP to communicate with another by sending, receiving, acting on certain architected orders, and responding to the requesting IP.
  - Prefixing: Provides a means of assigning different 4 KB blocks of main storage for each IP's real addresses 0 through 4095 (the Prefixed Storage Area or PSA). A prefixed address is known as an absolute address.
  - Shared Main Storage: Allows all IP's access to a common main storage.
  - TOD Clock Synchronization: With the assistance of a clock-synchronization program, this subfunction provides the capability for all IPs in a multiprocessing configuration to store the same TOD clock value, even though there may be multiple TOD clocks in the configuration. The facility works by synchronizing the stepping rates of all TOD clocks and by providing an error condition signaled by an external interruption, should one of the clocks become unsynchronized.
- **Open Systems Connection Feature (OCF):** Provides integrated LAN adapters for Ethernet, Fiber Distributed Data Interface (FDDI), and Asynchronous Transfer Mode (ATM).
- **Open Systems Adapter (OSA):** See Open Systems Connection Feature (OCF).
- **Page Protection:** Provides a bit in each page table entry to prevent storing data into a 4 KB page of main storage. Enhances key-controlled storage protection.
- **PER Extensions:** Extensions to PER 2 that provide added information about PER events.
- **Perform Locked Operation Facility:** Consists of PERFORM LOCKED OPERATION (PLO) instruction which is useful in avoiding the overhead of programmed locks and in locking the front and back of queues.
- **PMC Connectivity:** Allows one or more F-7, F-9, and/or HCF7 systems to be operated from a single-point-of-control Processor Management Console (PMC).
- **Private Space:** Prevents the use of TLB entries for common segments and the application of low-address protection and fetch protection override to the specified address space by setting a bit in the segment-table.
- **Processor Availability Facility (PAF):** Provides a hardware-based mechanism for work in progress on a failed processor to be transferred to a running processor non-disruptively.
- **PSW Key Handling:** Provides SET PSW KEY FROM ADDRESS (SPKA) and INSERT PSW KEY (IPK) instructions to allow enhanced manipulation of the storage keys in PSWs.
- **Recovery Extensions:** Extensions that enhance recovery functions, specifically:
  - Clear channel function, which performs an I/O system reset in a channel when Clear Channel (CLRCH) instruction is executed.
  - Machine check extensions including the machine-check external damage-code validity bit and a detailed indication of the external damage's cause.
  - Limited channel logout extensions, which implement an additional logout bit to indicate whether the I/O interface is operative and another additional logout bit to indicate whether the logout is valid.
- **Redundant Components:** Certain hardware components have N+1 or dual redundancy. If one component fails, the other(s) take over the load dynamically and non-disruptively, and the system continues to run. The following types of components have N+1 or dual redundancy:
  - Fans (N+1 redundancy)
  - Power Supplies (N+1 redundancy)
  - Refrigeration Units (dual redundancy)

- **Remote Support:** Allows a system to be diagnosed and serviced from a remote support center. Telecommunication links in the console subsystem connect the system to the support center.
- **Scalar Square Root:** Provides a hardware implementation of the square root function.
- **SCP-initiated Reset:** See I/O Interface Reset.
- **Segment Protection:** Provides a means to prevent an instruction from storing data in a segment of main storage.
- **Service Signal:** With this facility, the processor console can cause an external interruption to signal information to a program.
- **Set-Address Space Control Fast (SACF):** Enhances Set Address Space Control (SAC) instruction. Changes address-space mode settings without serialization or check-point synchronization to improve throughput performance.
- **SIGP Assists:** See Extensions for Virtual Machines.
- **Sorting instructions:** Provides COMPARE AND FORM CODEWORD (CFC) and UPDATE TREE (UPT) instructions useful for sorting data. Exploited by DFSORT Release 7 and above.
- **Storage-key 4 KB block:** Consists of the following subfunctions.
  - Single-key 4 KB block: Allows a single storage key to be associated with each 4 KB block of main storage.
  - Storage-key exception control: Provides a mechanism to prevent a special operation exception by INSERT STORAGE KEY (ISK), RESET REFERENCE BIT (RRB), and SET STORAGE KEY (SSK) instructions.
- **Storage-key instruction extensions:** SET STORAGE KEY EXTENDED (SSKE), INSERT STORAGE KEY EXTENDED (ISKE), and RESET REFERENCE BIT EXTENDED (RRBE) instructions. These are similar to SSK and ISK instructions, except that the extended instructions support 31-bit addresses and operate on the storage key for a 4 KB main storage block.
- **Storage-key instructions:** SET STORAGE KEY (SSK), INSERT STORAGE KEY (ISK), and RESET REFERENCE BIT (RRB) instructions which provide for protection of each 2 KB block of main storage through the use of storage keys.
- **Subspace Group Facility:** Provides a more robust protection mechanism than SSSP for use by CICS. CICS runs both CICS subsystem programs and one or more application programs in a single address space. With Subspace Group Facility and the CICS storage isolation facility, each application program is assigned storage private to itself; this private storage cannot be accidentally or maliciously modified by any other application program. This protection mechanism is exploited by CICS/ESA Version 4 Release 1 and above.
- **Subsystem Storage Protection (SSSP):** Allows subsystem storage to be isolated from application programs, preventing any application programs from overwriting (or overlaying) subsystem storage. SSSP is exploited by CICS/ESA Version 3 and above.
- **Suppression on Protection:** Allows an instruction which would otherwise have a page protection exception condition to be suppressed rather than terminated with the exception. Information about the protected page is placed in low storage and is available to the system during execution. Exploited by AIX/ESA's Copy on Write function.
- **Sysplex Timer:** See External Timer Attachment Feature (ETAF)
- **System/370 architecture:** The architecture for computer systems and programs defined in *IBM System/370 Principles of Operation*.
- **System/370 Extended Facility:** Extended facility enhancements in S/370 can be MVS-dependent or non-MVS-dependent as shown below.
  - MVS-dependent portion of System/370 Extended Facility includes:
    - Supervisor Call (SVC) Assist instruction, which reduces the overhead time needed to enter MVS supervisory services routines.
    - Fix Page instruction, Add Functional Recovery Routine instruction, six tracing instructions, and four lock-handling instructions, all of which improve processor performance.
  - Non-MVS-dependent portion of System/370 Extended Facility includes:
    - Low address protection (LAP), which prevents instructions from storing data into the area of main storage (addresses 0 through 511) used by an IP during interrupt processing.
    - INVALIDATE PAGE TABLE ENTRY (IPTE) instruction and common segment bit, which increase DAT efficiency.
    - TEST PROTECTION (TPROT) instruction, which tests potential protection violations without causing a protection exception.
- **System/370 I/O instructions:** Architecturally-defined instructions and facilities to perform I/O operations in System/370 environment.
- **Test Block:** Provides TEST BLOCK (TB) instruction which tests whether a 4 KB block of main storage is usable (free of errors).
- **Time-of-day (TOD) clock:** Provides a consistent and high-resolution measurement of elapsed time that can be used by a program for determining the date and time of day. The TOD clock's cycle is approximately 143 years. The clock is incremented such that, logically, a one is added in bit position 51 every microsecond.
- **Tracing (ASN, branch, and explicit):** Provides address-space number (ASN), branch, and explicit tracing for use in software problem determination.

- **Translation:** Converts a virtual main storage address into an absolute main storage address. The virtual address is first translated into a real address through the use of main-storage-resident segment tables and page tables. Prefixing is then applied to generate an absolute address. In access-register mode, a virtual address in a base register may designate a virtual address in another address space or data space (i.e., using another set of segment tables than those in use for the address space executing the instruction). This facility provides the subfunctions shown below.
  - Dynamic Address Translation (DAT): Causes the hardware to control the translation of virtual addresses into absolute addresses.
  - Extended Control (EC) mode: Provides a PSW format compatible with System/370 architecture and facilities. Required by Dynamic Address Translation (DAT) and Dual Address Space (DAS). Contrast with Basic Control mode (BC-mode).
  - Program Event Recording 2 (PER 2): Provides monitoring and trapping means in software problem determination. Actions subject to monitoring include the following:
    - Execution of a successful branch with, optionally, the branch target being within a designated area of storage
    - Instruction fetches from a specified storage area
    - Alteration of a specified storage area with, optionally, the storage being within designated address spaces
    - Execution of STORE USING REAL ADDRESS instruction

    PER 2 also provides the following functions:
    - Set-system-mask suppression: Provides a facility, enabled by the SSM-suppression-control bit (bit 1 of control register 0), to cause a special-operation exception upon attemped execution of SET SYSTEM MASK (SSM) instruction. SSM changes bits 0-7 of the PSW resulting in possible mode changes of several critical hardware facilities.
    - Store status: Places the contents of the following fields in permanently assigned areas in the first 512 bytes of absolute storage: CPU timer, clock comparator, current PSW, prefix, access registers, floating-point registers, general registers, and control registers. Store status can be performed under operator control or program control.
- **TSCF target support:** Provides a console subsystem external interface to allow an F-7C or F-7E system to be controlled from a NetView focal point system.
- **Virtual Machine Assists under SIE:** Provides the functions of Virtual Machine Assist (VMA) for System/370 mode guests in an interpretive execution environment. VMA reduces overhead for System/370 mode guests by reducing the amount of time the Control Program (CP) spends in real supervisor state when processing certain privileged instructions issued by the guest.
- **3033 Extension:** Includes the following functions.
  - Dual Address Space (DAS): Assists communications between virtual address spaces by providing twelve instructions, two address spaces usable concurrently by a program, a mechanism to change to other address spaces, subroutine linkages implemented through tables, and multiple access keys.
  - Start I/O Fast (SIOF) Queuing: By means of START I/O FAST RELEASE (SIOF) instruction and supporting channel facilities, this function allows completion of the instruction even though the I/O operation may still be pending in the subchannel until conditions allow I/O initiation to the device.
  - Suspend and Resume: By means of RESUME I/O (RIO) instruction and other supporting facilities, this function allows channel programs to be suspended and later resumed without having to re-initiate the I/O operation.
- **31-bit IDAWs:** 31-bit addressing in Indirect-Data-Address Words (IDAWs), allowing I/O operations to transfer data to and from all storage addresses.
- **31-bit Real Addressing:** Ensures that the following architectural fields provide 31-bit addresses (real or absolute addresses, as appropriate) regardless of the current addressing mode.
  - Prefix register
  - Primary segment-table origin in control register 1
  - Linkage-table origin in control register 5
  - Secondary segment-table origin in control register 7
  - ASN-first-table origin in control register 14
  - Page-table origin in the segment-table entry
  - Page-frame real address in the page-table entry
  - ASN-second-table origin in the AFT entry
  - Segment-table origin, linkage-table origin, and authority-table origin in the AST entry
  - Entry-table origin in the linkage-table entry
  - Address in format-1 CCWs

CHAPTER 7 COOLANT DISTRIBUTION UNIT

7.1 Overview

The F-7C Turbo Models and the F-7E Turbo Models adopt the cooling method of water cooling on the HDM, while the other components are air-cooled in the same manner as the F-7C Base Models and the F-7C CF Models. The Coolant Distribution Unit (CDU) is the main component for the said water cooling. Specifically, the CDU makes water of low temperature by exchanging heat between an air-chilled closed loop circulating CFC alternative refrigerant and an open loop circulating water to the HDM. This is similar to the cooling method optionally deployed by the G-8 Processor Group.

7.2 Function

The purpose of the CDU is to supply a controlled constant amount of water flow, at a constant temperature, to the HDM. In doing so, it removes the heat from the returned water, recirculating cool water back to the HDM. Control and monitoring of CDU operation is done by the service processor (SVP). Abnormal status information, such as loss of water pressure or temperature out-of-bounds conditions, are reported to the SVP, and appropriate actions are taken to prevent system damage from over-temperature conditions.

The closed loop circulates CFC alternative refrigerant R134a (HFC134A), while the open loop circulates water with 200 ppm BTA (1,2,3-Benzotriazole: $C_6H_5N_3$) solution.

Figure 7-1 shows the cooling schematics in theory.

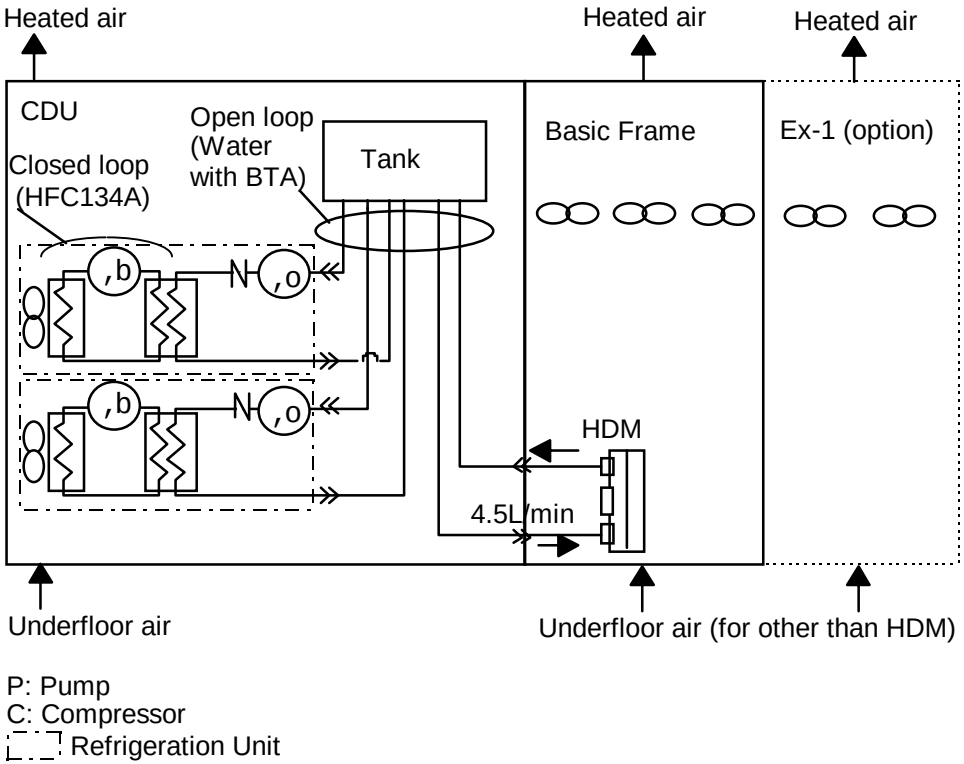

Figure 7-1 Function of CDU in F-7C Turbo Models and F-7E Turbo Models