IBM

Large
Systems
Technical
Support

Technical
Bulletin

IBM 3090 Processor Complex: Planning and Installation Reference

D. D. Boos
B. E. Kasch

National Technical Support
Washington Systems Center

GG66-3090-01
November 1987

Washington Systems Center
Gaithersburg, MD
Technical Bulletin

IBM 3090 Processor Complex: Planning and Installation Reference

D. D. Boos
B. E. Kasch

GG66-3090-01
November 1987

The information contained in this document has not been submitted to any formal IBM test and is distributed on an "as is" basis **without any warranty either expressed or implied.** The use of this information or the implementation of any of these techniques is a customer responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. While each item may have been reviewed by IBM for accuracy in a specific situation, there is no guarantee that the same or similar results will be obtained elsewhere. Customers attempting to adapt these techniques to their own environments do so at their own risk.

Reference to PTF numbers that have not been released through the PUT process does not imply general availability. The purpose of including these reference numbers is to alert IBM CPU customers to specific information relative to the implementation of the PTF when it becomes available to each customer according to the normal IBM PTF distribution rules.

Performance data contained in this document was determined in a controlled environment; and therefore, the results which may be obtained in other operating environments may vary significantly. Users of this document should verify the applicable data for their specific environment.

In this document, any references made to an IBM licensed program are not intended to state or imply that only IBM's licensed program may be used; any functionally equivalent program may be used instead.

It is possible that this material may contain reference to, or information about, IBM products (machines and programs), programming, or services that are not announced in your country. Such references or information must not be construed to mean that IBM intends to announce such IBM products, programming, or services in your country.

Publications are not stocked at the address given below; requests for IBM publications should be made to your IBM representative or to the IBM branch office serving your locality.

A form for reader's comments is provided at the back of this publication. If the form has been removed, comments may be addressed to: Washington Systems Center, 18100 Frederick Pike, Gaithersburg, MD 20879.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation whatever. You may, of course, continue to use the information you supply.

Abstract

This Technical Bulletin contains information pertinent to the IBM 3090 family of processors. It is intended for IBM and Customer personnel who have planning responsibilities for the installation of any of these processors.

The overall design of this document is to provide information about each of the processors and the operating systems and major products that support these processors.

This document is structured into seven major chapters: Family Comparisons, Logical Components, Software Considerations, Installation Topics, Partitioning Topics, Virtual Storage Considerations and Bibliography.

A brief description of each chapter follows:

1. Family Comparisons - This chapter contains a single page on each of the processors in the IBM 3090 family listing the standard and optional features of each. Also included are comparison summaries of the various models.

2. Logical Components - This chapter focuses on the logical components of the IBM 3090 processor family. A description of each of these components and how they are implemented in each member of the IBM 3090 processor family is provided. The unique implementations include: uniprocessor (120E, 150, 150E, 180, 180E), two-way (dyadic) processor (200, 200E), three-way (triadic) processor (300E), four-way processor (400, 400E), and six-way processor (600E).

3. Software Considerations - This chapter contains an overview of the supporting software for the 3090 family. It describes the changes incorporated into MVS/XA 2.1.3, MVS/XA 2.1.7, MVS/XA 2.2.0, MVS/370 1.3.5, VM/XA Systems Facility, VM/High Performance Option, VM/XA System Product, and the major supporting products (JES, EREP, IOCP, and RMF).

4. Installation Topics - This chapter focuses on the major installation considerations that are pertinent to all models of the IBM 3090 family.

5. Partitioning Topics - This chapter contains considerations that are primarily applicable to the "partitionable" processors in the 3090 family. These are the Models 400, 400E, and 600E.

6. Virtual Storage Assessment - This chapter provides considerations for the use of virtual storage, why it is important not to run out, some recommendations for the current use of virtual storage, a list of IBM products that provide Virtual Storage Constraint Relief (VSCR), and the highlights of the amount of relief expected from some of these products.

7. Bibliography - This chapter contains a bibliography for the 3090 family of processors. There is a description of each document that should assist the reader in choosing which documents should be used for specific points of reference.

Preface

The primary intent of this document is to serve as a composite source of planning and installation reference material for IBM 3090 Processors. The document does not obsolete previously available Systems Center technical bulletins on 3090 processors. Information contained in this document consists of original material and material obtained from other sources, including technical bulletins, flashes, and standard IBM publications.

The approach taken by the authors in the preparation of this document was to provide a handy reference for 3090 processor planning and installation information that would suffice for most situations. If additional detail on a particular topic is required, then existing Systems Center technical bulletins and standard IBM publications should be used. A list of such documents is provided in Chapter 7, Bibliography.

Summary of Amendments

This edition contains additional information, since the April 1987 version, about the IBM 3090 family of processors. These additions fall into the following categories:

- New 3090 Models and features:
  - 3090 Model 120E
  - Additional Expanded Storage on Models 200E, 300E, 400E, and 600E
  - 4.5Mb/sec Block Channel Support for all 3090 models
  - The Multiple High Performance Guests Support Feature on all E-models
- New Software Support:
  - VM/SP System Product Releases 1 and 2
  - MVS/SP 2.2.0 VIO support for Expanded Storage
  - IOCP support for additional channels and 4.5Mb/sec channel support
- SEC 223630 Contents:
  - Function and Features
  - Operational Enhancements and Changes
- Miscellanous changes:
  - HSA Expansion
  - Increased Extended SQA Requirements
  - Setting Machine Check Thresholds for Storage Errors
  - Procedures for attaching 3480 to Processor Controller (PCE)
  - Steps required to enable 4.5Mb/sec channel support
  - Partitioning the 3090 with an MSS attached
- Additions to the Bibliography:
  - Technical Bulletins
  - Operations Guides
  - Recovery Guides
  - Educational Offerings

Contents

Figures

Chapter 1. IBM 3090 Family Comparisons

Model 150

Uniprocessor

Standard Features

Channels	16
Central Storage (MB)	32
Expanded Storage (MB)	0
Vector Features	0

Optional Features

Channels	8
Central Storage (MB)	32
Expanded Storage (MB)	0
Vector Features	1

Miscellaneous Hardware Information

CP ID's	1
Cycle time	18.5ns
Version code	05
Installation Time	14 system hours (approx.)
Upgradeable to:	
Model 180 (as available)	6 system hours (approx.)
Model 180E	9 system hours (approx.)

System Control Program Support

MVS/XA	2.2, 2.1.7, 2.1.3
MVS/370	1.3.5
VM/XA Systems Facility	R2, R1
VM/XA System Product	R1, R2
VM/SP HPO	5.0, 4.2, 3.6
EREP	3.3, 3.2, 3.1 (w/feature 3)

System Requirements

3090 Processor Unit	1 (150)
3092 Processor Controller	1 (Model 1)
3097 PCDU	1 (Model 1 or 2)
3089 Model 3	1
3370-A02 Direct Access	2 (w/String Switch feature #8150)
3180 Model 145 Display	2
3864 Model 2 Modem	1 (w/Auto-call Unit feature #5801)

Model 180

Uniprocessor

Standard Features

Channels	16
Central Storage (MB)	32
Expanded Storage (MB)	0
Vector Features	0

Optional Features

Channels	8,16
Central Storage (MB)	32
Expanded Storage (MB)	64,128,192,256
Vector Features	1

Miscellaneous Hardware Information

CP ID's	1
Cycle time	18.5ns
Version code	10
Installation Time	14 system hours (approx.)
Upgradeable to:	
Model 200 (as available)	8 system hours (approx.)
Model 200E	13 system hours (approx.)

System Control Program Support

MVS/XA	2.2, 2.1.7, 2.1.3
MVS/370	1.3.5
VM/XA Systems Facility	R2, R1
VM/XA System Product	R1, R2
VM/SP HPO	5.0, 4.2, 3.6
EREP	3.3, 3.2, 3.1 (w/feature 3)

System Requirements

3090 Processor Unit	1 (180)
3092 Processor Controller	1 (Model 1)
3097 PCDU	1 (Model 1 or 2)
3089 Model 3	1 (see note below)
3370-A02 Direct Access	2 (w/String Switch feature #8150)
3180 Model 145 Display	2
3864 Model 2 Modem	1 (w/Auto-call Unit feature #5801)

Note: Requires 2 3089s if Vector Facility and greater than 128MB of expanded storage are installed.

Model 200

Two-way (Dyadic) Processor

Standard Features

Channels	32
Central Storage (MB)	64
Expanded Storage (MB)	0
Vector Features	0

Optional Features

Channels	8,16
Central Storage (MB)	0
Expanded Storage (MB)	64,128,192,256
Vector Features	1,2

Miscellaneous Hardware Information

CP ID's	1,2
Cycle time	18.5ns
Version code	20
Installation Time	14 system hours (approx.)
Upgradeable to:	
Model 400 (as available)	20 system hours (approx.)
Model 300E	14 system hours (approx.)
Model 400E	22 system hours (approx.)

System Control Program Support

MVS/XA	2.2, 2.1.7, 2.1.3
MVS/370	1.3.5
VM/XA Systems Facility	R2, R1
VM/XA System Product	R1, R2
VM/SP HPO	5.0, 4.2, 3.6
EREP	3.3, 3.2, 3.1 (w/feature 3)

System Requirements

3090 Processor Unit	1 (200)
3092 Processor Controller	1 (Model 1)
3097 PCDU	1 (Model 1 or 2)
3089 Model 3	2
3370-A02 Direct Access	2 (w/String Switch feature #8150)
3180 Model 145 Display	2
3864 Model 2 Modem	1 (w/Auto-call Unit feature #5801)

Model 400

Four-way Processor	Partitionable
Standard Features	
Channels	64
Central Storage (MB)	128
Expanded Storage (MB)	0
Vector Features	0
Optional Features	
Channels	16,32
Central Storage (MB)	0
Expanded Storage (MB)	128,256,384,512
Vector Features	1,2,3,4
Miscellaneous Hardware Information	
CP ID's	1,2,3,4
Cycle time	18.5ns
Version code	40
Installation Time	17 system hours (approx.)
Upgradeable to:	
Model 600E	24 system hours (approx.)
System Control Program Support	
MVS/XA	2.2(SI/PP modes), 2.1.7(SI/PP modes), 2.1.3(PP mode)
MVS/370	1.3.5(PP mode)
VM/XA Systems Facility	R2(SI/PP modes), R1(PP mode)
VM/XA System Product	R1(SI/PP mode), R2(SI/PP modes)
VM/SP HPO	5.0(PP mode), 4.2(PP mode), 3.6(PP mode)
EREP	3.3, 3.2, 3.1 (w/feature 3)
System Requirements	
3090 Processor Unit	1 (400)
3092 Processor Controller	1 (Model 2)
3097 PCDU	2 (Model 1 or 2)
3089 Model 3	4
3370-A02 Direct Access	2 (w/String Switch feature #8150)
3180 Model 145 Display	3
3864 Model 2 Modem	2 (w/Auto-call Unit feature #5801)

Model 120E

Uniprocessor

Standard Features

Channels	16
Central Storage (MB)	32
Expanded Storage (MB)	0
Vector Features	0

Optional Features

Channels	8
Central Storage (MB)	0
Expanded Storage (MB)	64,128
Vector Features	1

Miscellaneous Hardware Information

CP ID's	1
Cycle time	18.5ns
Version code	03
Installation Time	11 system hours (approx.)
Upgradeable to:	
Model 150E	16 system hours (approx.)

System Control Program Support

MVS/XA	2.2, 2.1.7, 2.1.3
MVS/370	1.3.5
VM/XA Systems Facility	R2, R1
VM/XA System Product	R1, R2
VM/SP HPO	5.0, 4.2,
EREP	3.3, 3.2

System Requirements

3090 Processor Unit	1 (120E)
3092 Processor Controller	1 (Model 3)
3097 PCDU	1 (Model 1 or 2)
3089 Model 3	1
3370-A02 Direct Access	1 (w/o String Switch feature #8150)
3180 Model 145 Display	2
3864 Model 2 Modem	1 (w/Auto-call Unit feature #5801)

Model 150E

Uniprocessor

Standard Features

Channels	16
Central Storage (MB)	32
Expanded Storage (MB)	0
Vector Features	0

Optional Features

Channels	8
Central Storage (MB)	32
Expanded Storage (MB)	64,128
Vector Features	1

Miscellaneous Hardware Information

CP ID	1
Cycle time	17.75ns (see note)
Version code	06 (see note)
Installation Time	11 system hours (approx.)
Upgradeable to:	
Model 180E	9 system hours (approx.)

System Control Program Support

MVS/XA	2.2, 2.1.7, 2.1.3
MVS/370	1.3.5
VM/XA Systems Facility	R2, R1
VM/XA System Product	R1, R2
VM/SP HPO	5.0, 4.2
EREP	3.3, 3.2

System Requirements

3090 Processor Unit	1 (150E)
3092 Processor Controller	1 (Model 1)
3097 PCDU	1 (Model 1 or 2)
3089 Model 3	1
3370-A02 Direct Access	2 (w/String Switch feature #8150)
3180 Model 145 Display	2
3864 Model 2 Modem	1 (w/Auto-call Unit feature #5801)

Note: Effective upon completion of step 2 modification for two-step install systems.

Model 180E

Uniprocessor

Standard Features

Channels	16
Central Storage (MB)	32
Expanded Storage (MB)	0
Vector Features	0

Optional Features

Channels	8,16
Central Storage (MB)	32
Expanded Storage (MB)	64,128,192,256
Vector Features	1

Miscellaneous Hardware Information

CP ID	1
Cycle time	17.2ns (see note 1)
Version code	11 (see note 1)
Installation Time	11 system hours (approx.)
Upgradeable to:	
Model 200E	9 system hours (approx.)

System Control Program Support

MVS/XA	2.2, 2.1.7, 2.1.3
MVS/370	1.3.5
VM/XA Systems Facility	R2, R1
VM/XA System Product	R1, R2
VM/SP HPO	5.0, 4.2
EREP	3.3, 3.2

System Requirements

3090 Processor Unit	1 (180E)
3092 Processor Controller	1 (Model 1)
3097 PCDU	1 (Model 1 or 2)
3089 Model 3	1 (see note 2)
3370-A02 Direct Access	2 (w/String Switch feature #8150)
3180 Model 145 Display	2
3864 Model 2 Modem	1 (w/Auto-call Unit feature #5801)

Note 1: Effective upon completion of step 2 modification for two-step install systems.

Note 2: Requires 2 3089s if Vector Facility and greater than 128MB of expanded storage are installed.

Model 200E

Two-way (Dyadic) Processor

Standard Features

Channels	32
Central Storage (MB)	64
Expanded Storage (MB)	0
Vector Features	0

Optional Features

Channels	8,16,32
Central Storage (MB)	64
Expanded Storage (MB)	64,128,192,256,512,1024
Vector Features	1,2

Miscellaneous Hardware Information

CP ID's	1,2
Cycle time	17.2ns (see note)
Version code	21 (see note)
Installation Time	12 system hours (approx.)
Upgradeable to:	
Model 300E	11 system hours (approx.)
Model 400E	18 system hours (approx.)

System Control Program Support

MVS/XA	2.2, 2.1.7, 2.1.3
MVS/370	1.3.5
VM/XA Systems Facility	R2, R1
VM/XA System Product	R1, R2
VM/SP HPO	5.0, 4.2
EREP	3.3, 3.2

System Requirements

3090 Processor Unit	1 (200E)
3092 Processor Controller	1 (Model 1)
3097 PCDU	1 (Model 1 or 2)
3089 Model 3	2
3370-A02 Direct Access	2 (w/String Switch feature #8150)
3180 Model 145 Display	2
3864 Model 2 Modem	1 (w/Auto-call Unit feature #5801)

Note: Effective upon completion of step 2 modification for two-step install systems.

Model 300E

Three-way Processor

Standard Features

Channels	32
Central Storage (MB)	64
Expanded Storage (MB)	0
Vector Features	0

Optional Features

Channels	8,16,32
Central Storage (MB)	64
Expanded Storage (MB)	64,128,192,256,512,1024
Vector Features	1,2,3

Miscellaneous Hardware Information

CP ID's	0,1,2
Cycle time	17.2ns
Version code	31
Installation Time	12 system hours (approx.)
Upgradeable to:	
Model 600E	18 system hours (approx.)

System Control Program Support

MVS/XA	2.2, 2.1.7
VM/XA Systems Facility	R2
VM/XA System Product	R1, R2
EREP	3.3, 3.2

System Requirements

3090 Processor Unit	1 (300E)
3092 Processor Controller	1 (Model 1)
3097 PCDU	1 (Model 1 or 2)
3089 Model 3	2
3370-A02 Direct Access	2 (w/String Switch feature #8150)
3180 Model 145 Display	2
3864 Model 2 Modem	1 (w/Auto-call Unit feature #5801)

Model 400E

Four-way Processor	Partitionable
Standard Features	
Channels	64
Central Storage (MB)	128
Expanded Storage (MB)	0
Vector Features	0
Optional Features	
Channels	16,32,64
Central Storage (MB)	128
Expanded Storage (MB)	128,256,384,512,1024,2048
Vector Features	1,2,3,4
Miscellaneous Hardware Information	
CP ID's	1,2,3,4
Cycle time	17.2ns (see note)
Version code	41 (see note)
Installation Time	14 system hours (approx.)
Upgradeable to:	
Model 600E	16 system hours (approx.)
System Control Program Support	
MVS/XA	2.2(SI/PP modes), 2.1.7(SI/PP modes), 2.1.3(PP mode)
MVS/370	1.3.5(PP mode)
VM/XA Systems Facility	R2(SI/PP modes), R1(PP mode)
VM/XA System Product	R1(SI/PP modes), R2(SI/PP modes)
VM/SP HPO	5.0(PP mode), 4.2(PP mode)
EREP	3.3, 3.2
System Requirements	
3090 Processor Unit	1 (400E)
3092 Processor Controller	1 (Model 2)
3097 PCDU	2 (Model 1 or 2)
3089 Model 3	4
3370-A02 Direct Access	2 (w/String Switch feature #8150)
3180 Model 145 Display	3
3864 Model 2 Modem	2 (w/Auto-call Unit feature #5801)

Note: Effective upon completion of step 2 modification for two-step install systems.

Model 600E

Six-way Processor	Partitionable
Standard Features	
Channels	64
Central Storage (MB)	128
Expanded Storage (MB)	0
Vector Features	0
Optional Features	
Channels	16,32,64
Central Storage (MB)	128
Expanded Storage (MB)	128,256,384,512,1024,2048
Vector Features	1,2,3,4,5,6
Miscellaneous Hardware Information	
CP ID's	0,1,2,3,4,5
Cycle time	17.2ns
Version code	61
Installation Time	17 system hours (approx.)
System Control Program Support	
MVS/XA	2.2(SI/PP modes), 2.1.7(SI/PP modes)
VM/XA Systems Facility	R2(SI/PP modes)
VM/XA System Product	R1(SI/PP modes), R2(SI/PP modes)
EREP	3.3, 3.2
System Requirements	
3090 Processor Unit	1 (600E)
3092 Processor Controller	1 (Model 2)
3097 PCDU	2 (Model 1 or 2)
3089 Model 3	4
3370-A02 Direct Access	2 (w/String Switch feature #8150)
3180 Model 145 Display	3
3864 Model 2 Modem	2 (w/Auto-call Unit feature #5801)

Comparison Summaries

This section provides comparison summaries from the information provided in the preceding charts on the various models of the IBM 3090 processor family.

Model	Central Processors	Central Storage	Expanded Storage	Channels	Vector Facilities
150	1	32-64	0	16-24	0 or 1
180	1	32-64	0-256	16-32	0 or 1
200	2	64	0-256	32-48	0 to 2
400	4	128	0-512	64-96	0 to 4
120E	1	32	0-128	16-24	0 or 1
150E	1	32-64	0-128	16-24	0 or 1
180E	1	32-64	0-256	16-32	0 or 1
200E	2	64-128	0-1024	32-64	0 to 2
300E	3	64-128	0-1024	32-64	0 to 3
400E	4	128-256	0-2048	64-128	0 to 4
600E	6	128-256	0-2048	64-128	0 to 6

Figure 1. Standard and Optional Features

Model	CP IDs	Version Code	Storage Increment	Cycle Time	Install Time
150	1	05	2 Mb	18.5ns	14 hrs
180	1	10	2 Mb	18.5ns	14 hrs
200	1,2	20	2 Mb	18.5ns	14 hrs
400	1,2,3,4	40	2 Mb	18.5ns	17 hrs
120E	1	03	2 Mb	18.5ns	11 hrs
150E	1	06	2 Mb	17.75ns	11 hrs
180E	1	11	4 Mb	17.2ns	11 hrs
200E	1,2	21	4 Mb	17.2ns	12 hrs
300E	0,1,2	31	4 Mb	17.2ns	12 hrs
400E	1,2,3,4	41	4 Mb	17.2ns	14 hrs
600E	0,1,2,3,4,5	61	4 Mb	17.2ns	17 hrs

Figure 2. Miscellaneous Hardware Information

Modes of Operation: System control programming support is dependent on the mode in which the IBM 3090 Processor Complex is operating. For example, IBM 3090 Processor Complex Models 400, 400E, and 600E in Single Image mode, only operate in 370-XA mode and require a 370-XA mode system control program. When a IBM 3090 Model 400 or 400E is partitioned, it may be initialized in S/370 mode or in 370-XA mode and requires the appropriate system control program support for that mode. Figure 3 on page 13 defines the possible modes of operation in which each of the IBM 3090 models can be run.

Model	SI/PP	S/370	370/XA	Notes
150	N/A	Y	Y	
180	N/A	Y	Y	
200	N/A	Y	Y	
400	SI	N	Y	
400	PP	Y	Y	Each side in either mode
120E	N/A	Y	Y	
150E	N/A	Y	Y	
180E	N/A	Y	Y	
200E	N/A	Y	Y	
300E	N/A	N	Y	3-way requires 370-XA
400E	SI	N	Y	
400E	PP	Y	Y	Each side in either mode
600E	SI	N	Y	6-way requires 370-XA
600E	PP	N	Y	3-way requires 370-XA

Figure 3. Modes of Operation

Single Image Mode

- MVS/SP 2.2.0 and 2.1.7 provide support for all Models of the 3090 in both single image and partitioned modes of operation, Vector Facility options, and expanded storage. Full single image reconfiguration support is provided.

- VM/XA SF Release 2 provides support for all Models of the 3090 in both single image and partitioned modes of operation and for Vector Facility options. It does not use expanded storage for its own use, but allows it to be dedicated to a V=V or V=R guest. With an SPE (VM28091), VM/XA SF Release 2 allows expanded storage to be dedicated to a single guest or VM/XA SF R2, or partitioned between guests and VM/XA SF R2. Full single image dynamic reconfiguration support is **not** provided.

- VM/XA SP Releases 1 and 2 provide support for all Models of the 3090 in both single image and partitioned modes of operation and for Vector Facility options. They allow expanded storage to be dedicated to a single guest or VM/XA SP, or partitioned between guests and VM/XA SP. Full single image dynamic reconfiguration support is **not** provided. Planned availability of VM/XA SP Release 1 is March, 1988 (Release 2 is 1Q89).

Partitioned Mode

- MVS/SP 2.1.3 provides only partitioned mode support for Models 400 and 400E. With the Vector Facility Enhancement (VFE) installed, the Vector Facility options are supported. Expanded storage is also supported.

- MVS/SP 1.3.5 provides only partitioned mode support for Models 400 and 400E. Provides support for a maximum of 16 channels per channel set, and does **not** support expanded storage **or** Vector Facility options.

- VM/XA SF Release 1 provides only partitioned mode support for Models 400 and 400E and does **not** support the Vector Facility. It does not use expanded storage, but allows it to be dedicated to a V=V or V=R guest.

- VM/SP HPO Releases 4.2 and 5.0 provide only partitioned mode support for Models 400 and 400E, and support expanded storage as a high speed paging device. The maximum number of channels per channel set is 32. Vector Facility options are supported by HPO Releases 4.2 and 5.0.

Model	MVS/XA			MVS/370
	2.1.3	2.1.7	2.2	1.3.5
150	Y	Y	Y	Y
180	Y	Y	Y	Y
200	Y	Y	Y	Y
400	Y(P)	Y(S/P)	Y(S/P)	Y(P)
120E	Y	Y	Y	Y
150E	Y	Y	Y	Y
180E	Y	Y	Y	Y
200E	Y	Y	Y	Y
300E	N	Y	Y	N
400E	Y(P)	Y(S/P)	Y(S/P)	Y(P)
600E	N	Y(S/P)	Y(S/P)	N

Model	VM/XA SF		VM/XA SP		VM/SP HPO	
	R1	R2	R1	R2	4.2	5.0
150	Y	Y	Y	Y	Y	Y
180	Y	Y	Y	Y	Y	Y
200	Y	Y	Y	Y	Y	Y
400	Y(P)	Y(S/P)	Y(S/P)	Y(S/P)	Y(P)	Y(P)
120E	Y	Y	Y	Y	Y	Y
150E	Y	Y	Y	Y	Y	Y
180E	Y	Y	Y	Y	Y	Y
200E	Y	Y	Y	Y	Y	Y
300E	N	Y	Y	Y	N	N
400E	Y(P)	Y(S/P)	Y(S/P)	Y(S/P)	Y(P)	Y(P)
600E	N	Y(S/P)	Y(S/P)	Y(S/P)	N	N

Y = Supported S = Single Image Mode
N = Not Supported P = Partitioned Mode

Figure 4. System Control Program Support

MVS/SP 2.1.7 and MVS/SP 2.2.0 recommended service for E-Models. Be sure to contact the IBM Support Center for the latest service levels:

- OY01085 - RSM Page movement above/below 16Mb
- OY02659 - GRS Performance
- OY03215 - Reconfiguration
- OY03324 - SRM Constants
- OY03325 - SRM Constants
- OY06003 - SRM Constants (added 120E)
- OY03621 - RSM Real Storage larger than 128Mb
- OY03623 - SRM Real Storage larger than 128Mb
- OY03626 - SPVR CTL Real Storage larger than 128Mb
- OZ97887 - MCH Correctable storage key errors
- OZ97930 - DFP Subchannel Logout Handling
- OZ97996 - IOS Subchannel Logout Handling
- OZ97999 - IOS Subchannel Logout Handling

MVS/SP 2.1.3 recommended service for E-Models. Be sure to contact the IBM Support Center for the latest service levels:

- OY01085 - RSM Page movement above/below 16Mb
- OY03325 - SRM Constants
- OY06003 - SRM Constants (added 120E)
- OZ97887 - MCH Correctable storage key errors
- OZ97930 - DFP Subchannel Logout Handling
- OZ97996 - IOS Subchannel Logout Handling

- OZ97999 - IOS Subchannel Logout Handling
- OZ97978 - Reconfiguration
- OZ97979 - IPL

MVS/SP 1.3.5 recommended service for E-Models is OY03325 / OY06003 (120E) SRM constants.

The current version of IOCP (MVS, VM, Standalone) supports the E-Models. APAR OY03274 (available 4Q87) provides IOCP support for:

- 4.5Mb/Sec operation in data streaming mode on all block multiplexer channels on all 3090 models
- 64 Channel capability on models 200E and 300E (CHPIDs 30-3F)
- 128 Channel capability on models 400E and 600E (CHPIDs 30-3F and 70-7F)

Model	3092 Model	DASD 3370-A02	3097 Qty	3089-3 Qty	3180-145 Qty	Modem 3864-2
150	1	2	1	1	2-5	1
180	1	2	1	1 or 2[1]	2-5	1
200	1	2	1	2	2-5	1
400	2	2	2	4	3-6	2
120E	3	1	1	1	2-5	1
150E	1	2	1	1	2-5	1
180E	1	2	1	1 or 2[1]	2-5	1
200E	1	2	1	2	2-5	1
300E	1	2	1	2	2-5	1
400E	2	2	2	4	3-6	2
600E	2	2	2	4	3-6	2

Figure 5. System Requirements

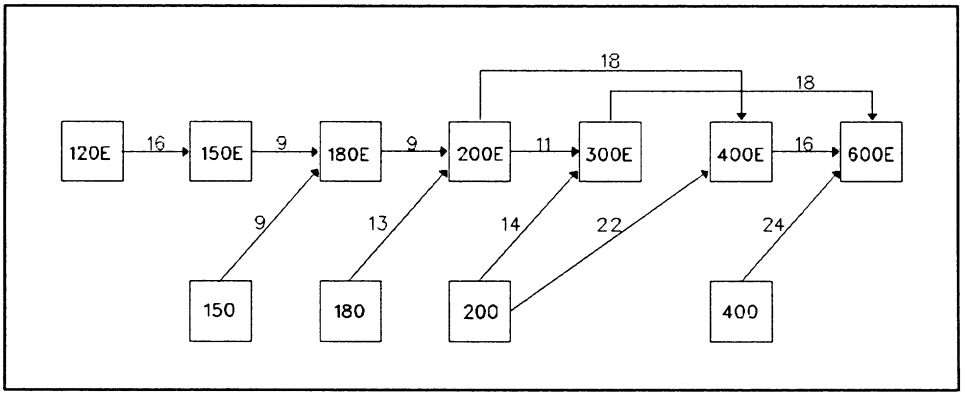

Figure 6. Model Upgrade Paths (Approximate time in system hours)

Model upgrade times are expressed in system hours and assume that any prerequisite Engineering Changes (ECs) are previously installed.

---

[1] Requires 2 3089s if Vector Facility and greater than 128MB of expanded storage are installed.

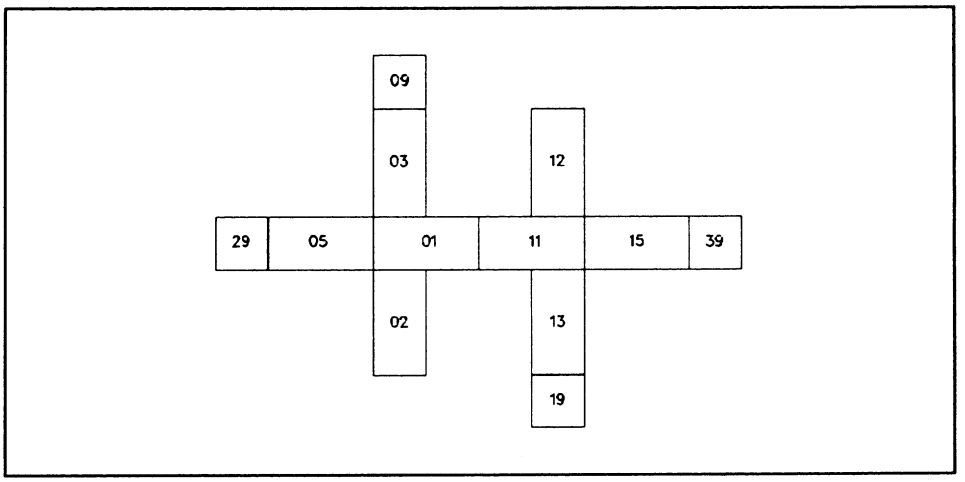

Figure 7. Frame layout of Full Featured Processor Unit Model 400E/600E.

Models	Minimum Frames Required
150/180/200	01-02-03-09
120E/150E/180E/200E	01-02-03-09
300E	01-02-03-09-05
400/400E	01-02-03-09 and 11-12-13-19
600E	01-02-03-09-05 and 11-12-13-19-15

Figure 8. Minimum Required Frames

Models	Options	Additional Frame(s)
120E	Vector Facility	05
150E	Vector Facility	05
180/180E	Vector Facility	05
200/200E	Vector Facility	05
400/400E	Vector Facility	05 and/or 15
200E	3rd Channel Group (64 Max)	05 and 29
300E	3rd Channel Group (64 Max)	29
400E	3rd Channel Group (128 Max)	05, 29, 15 and 39
600E	3rd Channel Group (128 Max)	29 and 39
200/200E	200(E) to 300E MES	05 (FC 7330 on 200(E))
400/400E	400(E) to 600E MES	05 (FC 7330 on 400(E)) and 15 (FC 7331 on 400(E))

Figure 9. Options Requiring Additional Frames

Notice that the expansion frames, 05 and/or 15, are required for:

1. Model upgrades resulting in 300E or 600E models. The third processor (CP) per side is located in the expansion frame(s).

2. Vector Facility options are located in the expansion frame(s).

3. Any 3090 E-model having the third Channel Group option. The additional frames, 29 and/or 39, attach to the expansion frame(s).

The expansion frames are standard components for the Models 300E and 600E and are included with orders for **full model** installs. Model upgrades (MESs) resulting in Model 300E or 600E must have the appropriate expansion frame features ordered for installation on the model being upgraded (200, 200E, 400, 400E) if these frames are not already installed to support Vector Facilities and/or third Channel Groups.

Chapter 2. IBM 3090 Logical Components

This chapter focuses on the logical components of the IBM 3090 processor family. A description of each of these components and their relationship with other components is included. These are the building blocks used for each of the unique implementations of the IBM 3090 processor family. The unique implementations include: uniprocessor (120E, 150, 150E, 180, 180E), two-way (dyadic) processor (200, 200E), three-way (triadic) processor (300E), four-way processor (400, 400E), and six-way processor (600E).

Logical Components

An example of the logical components used in the six-way implementation of the Model 600E is shown below.

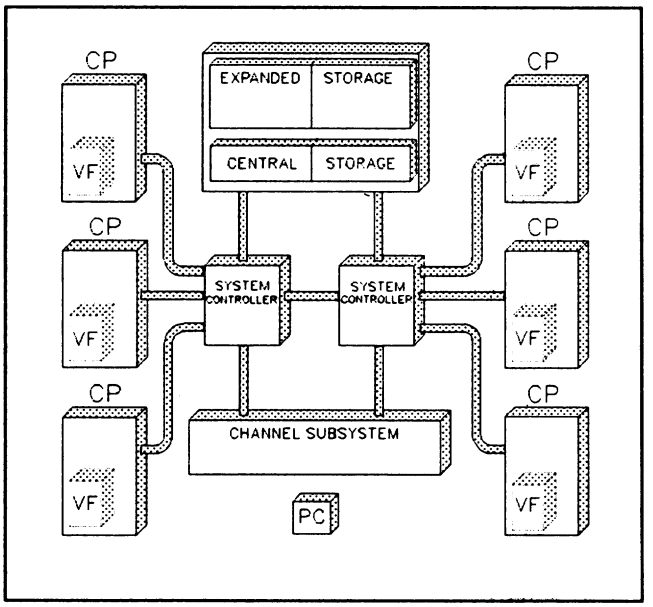

Figure 10. Logical Components of IBM 3090 Model 600E

Central Processors

Each central processor in the 3090 processor complex is microcode controlled and contains an instruction element (IE), an execution element (EE), a buffer control element (BCE), also known as a high-speed buffer (or cache), a control store element (CSE), and optionally a vector element (VE).

The packaging technology is the TCM (Thermal Conduction Module) mounted on boards. The circuit technology is ECL (Emitter Coupled Logic).

Instruction Element

The instruction element (IE) controls sequencing of all instructions. The IE fetches and decodes instructions, computes operand addresses, fetches the initial operands for the execution element, and queues the instructions to the execution element. The instruction element can process multiple instruction streams at the same time.

Execution Element

The execution element (EE) executes instructions set up by the instruction element and stores the results. The execution element performs logical decisions, arithmetic functions and many control functions. Execution element operation is overlapped with instruction element operation for increased performance.

Buffer Control Element

The buffer control element (BCE) handles all central processor references to and from central storage, performs dynamic address translation, and controls the high-speed buffer (HSB). The BCE includes:

- 64Kb high-speed buffer
- Buffer Directory
- Translation Lookaside Buffer (TLB)
- Dynamic Address Translation (DAT) hardware

Vector Facilities

The vector facility is available as an optional feature for each of the central processors of the IBM 3090. The vector facility is an extension of the central processor's instruction and execution elements that increases the throughput in certain engineering and scientific applications that use vectors. A Central Processor with the vector facility installed performs vector arithmetic and logical operations on as many as 128 sets of operands with a single instruction.

Multiple High Performance Guests Support

The multiple high performance guests support feature (optionally available on all E-Models) is a prerequisite for support of the VM/XA SP Enhancement for Multiple Preferred Guests. It allows the support of multiple high performance guests running concurrently with other virtual machines. In addition to the V=R preferred guest virtual machine, as many as three V=F preferred guests are supported. SIE assist supports devices dedicated to a V=F guest. VMA, under SIE, supports V=F VM/SP and VM/SP HPO guests.

System Control Element

The system control element (SCE) accepts and processes storage requests from the central processors and the channel subsystem. The SCE prioritizes and coordinates requests and ensures that the requestor receives the most recent copy of the data. Error checking and reporting on data movement is also carried out by the SCE.

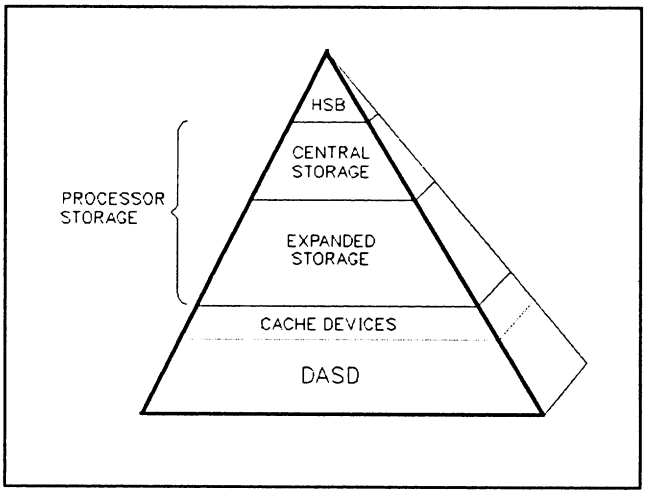

Figure 11. Storage Hierarchy

Processor Storage

The IBM 3090 processors have three levels of storage: high-speed buffer storage associated with each central processor, central storage, and optional expanded storage. The high-speed buffer for each central processor is 64Kb in size. All accesses to central storage by the central processors go through the high speed buffer. Accesses to central storage as a result of I/O operations do not go through the high speed buffer, but instead go directly to central storage.

When the IBM 3090 was first introduced, the memory chip technology was 64K-bit chips for central storage and 288K-bit chips for expanded storage. With IBM 3090 E Models, central and expanded storage are using IBM's advanced design one-million-bit (megabit) chip technology. This memory chip operates almost twice as fast and occupies about a third less space than the original IBM one-megabit chip.

Central Storage

Automatic error detection and correction is integral to the design of central storage. Single-bit errors detected during central storage data transfer operations are corrected. Certain multiple-bit errors are flagged for follow-on action. Some double bit errors can be corrected by the hardware and subsequent notification given to the operating system. This provides the operating system with an opportunity to deallocate the failing page frame, thereby avoiding the necessity of correcting subsequent double bit errors when referencing the page frame, and also eliminating the possibility that the correctable double bit error could turn into an uncorrectable multiple bit error.

Central storage also contains the Hardware System Area (HSA). The HSA is allocated from the top of central storage during the Power-On-Reset (POR) process. The HSA is assigned error free storage and contains system microcode and control blocks (representing the channel paths, control units and I/O devices) used by the channel subsystem to control I/O operations. It also contains the hardware I/O trace information when the I/O trace function is activated.

The size of the HSA is dependent, in part, on the number of control units and I/O devices in the configuration. The maximum number of I/O devices (UCW's) supported is equal to 4096 minus the number of channel paths defined using the I/O Configuration Program (IOCP).

During system initialization for all models, at least 384K bytes of central storage are assigned to the HSA. The expansion of HSA for all purposes can be up to 3Mb. The amount of central storage available for use by the operating system is reduced by an amount equal to the size of the HSA. For additional information on the HSA, see "HSA Expansion" on page 92.

Expanded Storage

In addition to central storage, expanded storage is also available. Expanded storage may be thought of as part of a hierarchy of storage used to contain instructions and data. It resides logically between central storage and auxiliary (DASD) storage in terms of performance and capacity. Expanded storage provides improved system performance by reducing the paging load to channel-attached paging devices in storage constrained/heavy paging environments.

Expanded storage connects through the expanded storage controller to the central storage memory controller. Expanded storage is not byte addressable and no instructions or data can be accessed directly from it. All access to expanded storage is through central storage. Data movement, in 4Kb page sizes, between central storage and expanded storage occurs in microseconds and is synchronous with processor operations. No data path exists between expanded storage and the I/O subsystem. All I/O is done to/from central storage.

Data movement between central storage and expanded storage is managed by the SCP[2] and is transparent to the user. MVS/SP 2.1.3 and later releases manage expanded storage similar to central storage. VM/SP HPO 3.6, 4.2 and 5.0 manage expanded storage as a high speed area for preferred paging, swapping or both. VM/XA Systems Facility Release 1 does not use expanded storage for its own use, but allows it to be dedicated to a V=V or V=R guest. With SPE VM28091, on PUT 8705, VM/XA SF Release 2 allows expanded storage to be dedicated to a single guest or to VM/XA SF R2 for paging or partitioned between guests and VM/XA SF R2. VM/XA System Product Releases 1 and 2 allow expanded storage to be dedicated to a single guest or VM/XA SP, or partitioned between guests and VM/XA SP. MVS/SP V1 does not support expanded storage.

Error checking and correction codes are used within expanded storage to detect errors. Single and double-bit errors are detected and corrected. Triple bit errors are detected. Uncorrectable errors are flagged and presented to the SCP for handling.

---

2 For additional details on the support of expanded storage by the SCP, see Chapter 3, "Software Considerations."

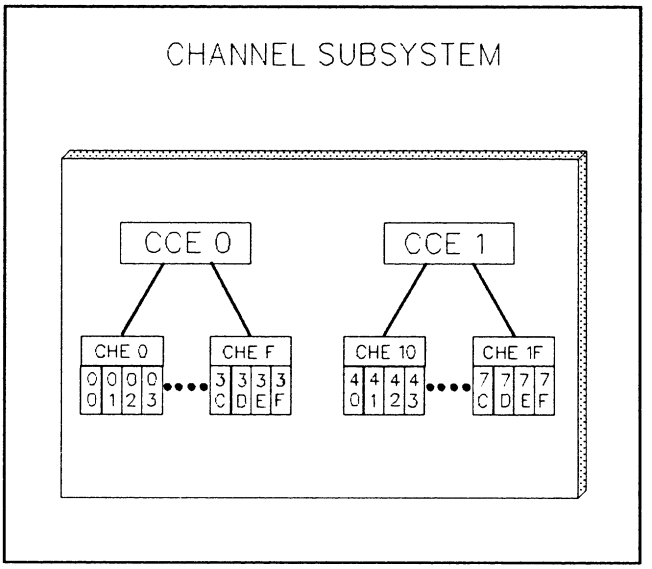

Figure 12. Channel Subsystem. This example is a Model 600E in Single Image mode.

Channels

Models 120E, 150, 150E, 180, 180E, 200, 200E, and 300E contain one channel subsystem. Models 400, 400E, and 600E contain two channel subsystems (one on each side). In Single Image mode they function together as a single dynamic channel subsystem to the control program.

A channel subsystem consists of one channel control element (CCE) and as many as 64 channel paths. Channel paths are physically packaged in groups of four, with each group of four paths controlled by a channel element (CHE). Channel path grouping for standard and optional channel groups for all models in the 3090 family is shown in Figure 27 on page 85.

All channel paths can operate in block multiplex mode and up to 4 channels (4 on each side for Models 400, 400E, and 600E) can be defined (via IOCP) to operate in byte multiplex mode. Both data streaming and non-data streaming devices may attach to any block multiplex channel path and may be intermixed.

Each block multiplex channel path is capable of a 3.0Mb per second data rate (up to 4.5Mb per second with SEC 223630 installed) when operating in data streaming mode, and a 1.5Mb per second data rate in DC-Interlock (high-speed transfer) mode.

Each byte multiplexer channel path is capable of operating with an effective data rate generally in the range of 90 to 300Kb per second for data transfer burst sizes of four bytes or more. Configurations consisting of control units with larger transfer burst sizes can achieve higher performance.

The channel subsystem on all 3090 Models, except 300E and 600E, may operate in either S/370 mode or 370-XA mode. In 370-XA mode, all channel paths are accessible by any processor. In S/370 mode, channel paths are assigned to one or two channel sets. A channel set is associated with a central processor. The 3090

implementation supports up to 32 channels per channel set. Software support of 32 channels per channel set is SCP dependent[3].

The channel subsystem handles all I/O operations for the central processors. It controls communications between a configured channel and the control unit and device. The channel, control unit, and device configurations are defined to the channel subsystem by the I/O configuration data set (IOCDS) that is selected during hardware system initialization. The IOCDS is created by the I/O Configuration Program (IOCP) and is stored on the 3370 Direct Access Storage devices that are attached to the Processor Controller.

Status-Verification Facility

The channel subsystem of the IBM 3090's also supports the "Status-Verification Facility." This facility provides the system with a means of indicating that it has been presented with device status of good parity, but containing a combination of bits not valid at the time of status presentation. This facility applies to 370-XA mode operation only, and is controlled (set active/inactive) via the "STADET=" keyword on the IOCP IODEVICE macro.

If the facility is active, and a device presents illegal status, the channel terminates the I/O request and generates an interface control check interruption. The LOGREC record written as a result will indicate "Device - status check," as the reason for the interface control check.

See the description of the IODEVICE Macro in section "IOCP Support" on page 79 for further details.

---

3 MVS/SP 1.3.5 supports only 16 channels per channel set. VM/SP HPO 3.6, 4.2 and 5.0 support 32 channels per channel set.

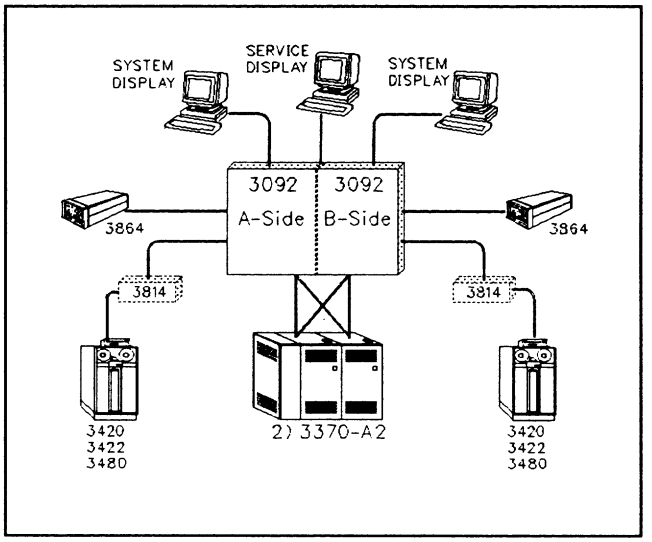

Figure 13. Processor controller. This example is a 3092 Model 2.

Processor Controller

The IBM 3092 Processor Controller is a standard component of all 3090 Processor Complex models. The 3092 Model 1 is required for the Models 150, 150E, 180, 180E, 200, 200E, and 300E. The 3092 Model 2 is required for the Model 400, 400E, and 600E and the 3092 Model 3 is required for the Model 120E.

The IBM 3092 Processor Controller monitors and controls the operations of the 3090 Processor Complex. The 3092 Processor Controller is involved in system initialization, power sequencing, continuous monitoring during system operation, and error processing. The 3092 Processor Controller contains the switches used for power on, power off, emergency power off, and indicators for power status and service mode.

During initialization of the 3090 Processor Complex, the 3092 Processor Controller validates areas of central storage as error-free data locations, records failing storage locations, and assigns the Hardware System Area to contiguous error-free locations in central storage. The 3092 Processor Controller controls power sequencing to all processor complex components and to all interconnected I/O control units that are under power sequence control (IOPS).

The 3092 Processor Controller also provides the operator interface to the hardware processor complex through the display frames, menus and commands available at the system console. Using the system console, the operator can control and monitor the state of the hardware system. This includes the capability to select elements to be in the configuration, determine the mode of system operation (S/370 mode or 370-XA mode), IPL the SCP, alter/display locations of central storage, and monitor system activity through the use of System Activity Displays (SAD).

During processing, the 3092 Processor Controller continuously monitors the state of the system, including power and thermal conditions. Should an abnormal

condition arise, an over/under voltage condition for example, or a rise or drop in the coolant temperature above/below a certain limit, a message is sent to the system console informing the operator of the condition, and specifying what further action should be taken.

The 3092 Processor Controller also provides extensive error recording, recovery, and diagnostic support for the processor unit. Error handling by the processor controller provides both automatic recovery from many hardware malfunctions, and reporting of machine or channel-check interruptions. The 3092 Processor Controller logs errors as they occur and then analyzes them. Failure symptoms, saved at the time of occurrence, are analyzed concurrent with other processor controller functions. During the analysis process, the 3092 Processor Controller attempts to isolate the failure to a specific failing FRU (Field Replaceable Unit) or group of FRUs.

The 3092 Processor Controller contains two processors (A-side and B-side), except for the 3092 Model 3 used with the 3090 Model 120E that has one processor. One processor is active and the other processor is backup. For Models 400, 400E and 600E in Single Image mode, one processor is active and the other processor is backup. For Models 400, 400E, and 600E in Physical Partitioned mode, each side of the processor controller controls one half of the complex. There is, therefore, no processor controller backup in PP mode.

PCE Attached Devices

The 3092 Processor Controller requires the following corequisite devices:

- Two IBM 3370 Model A2 DASD devices, each with a string-switch feature, except for the 3092 Model 3 that requires only one 3370 Model A2 unit (string-switch not required). For a 3092 Model 1, and for a 3092 Model 2 in a single-image configuration:
  - One 3370 is active and contains system microcode, IOCDS datasets, analysis routines used to isolate system failures, and other code for controlling system operation.
  - One 3370 acts as a backup, and is a duplicate of the active 3370.

    Note: For a 3092 Model 2 in a physically partitioned configuration, one 3370 is active for each side. For a 3092 Model 3, only one 3370 is required.

- The processor controller also requires access to one of the following:
  - An IBM 3803 Model 2 Tape Control Unit (or equivalent) and its associated IBM 3420 Model 4, 6, 8 tape drive (or equivalent 6250 bpi drive) with a maximum data rate of 1.25 megabytes per second.
  - An IBM 3422 Magnetic Tape Subsystem (or equivalent) using 6250 bpi. This support was made available via SEC 223630. The 3422 must be at EC level A46632 or higher.

- An IBM 3480 Magnetic Tape Unit Model B11 or B22 tape drive (or equivalent 38K bpi cartridge density drive) operating with DC Interlock protocol at a maximum data rate of 1.25 megabytes per second.

  Note: Models 400, 400E, and 600E require access to a tape drive from each side.

The tape drive is used to install engineering changes, MES upgrades, and for reloading the 3370's, if necessary. ***Valid addresses for the tape control unit and drive are X'00-9F' and X'B0-FF'. Addresses X'A0-AF' cannot be used if the tape is to be accessed by the 3092 Processor Controller.***

Note: It is advisable to have this tape switchable to a channel on some processor which has it defined by IOCP. This allows service personnel to run diagnostic tests on the tape drive when required for the tape control unit or drive.

The 3090 Models 120E, 150, 150E, 180, 180E, 200, 200E, and 300E require two IBM 3180 Display Stations Model 145[4]; one as a system console, and one as a service console. The operating system requires at least one additional channel-attached display for an operator console. The 3090 Models 400, 400E, and 600E require three IBM 3180 Display Stations Model 145; two as system consoles, and one as a service console. The operating system requires two additional channel-attached displays for operator consoles in a physically partitioned configuration. Refer to "Console Considerations" on page 88 for additional recommendations on operating system consoles.

Optionally, an IBM 3287 Printer Model 1 or 2, or IBM 4224 Printer Model 201 or 202, or equivalent, may be attached. The printer is optional and may be used for printing 3180 display screens.

The IBM 3092 Processor Controller also contains a Remote Support Facility (RSF). The Remote Support Facility requires an IBM 3864 Model 2 modem with an auto-call unit (feature code #5801) or equivalent, plus a switched telephone line. For the Models 400, 400E, and 600E, a modem and telephone line are required on each side of the Processor Controller. The modem and telephone line are used to access the IBM remote support system.

---

4 The 3180 Models 140 and 145 are interchangeable.

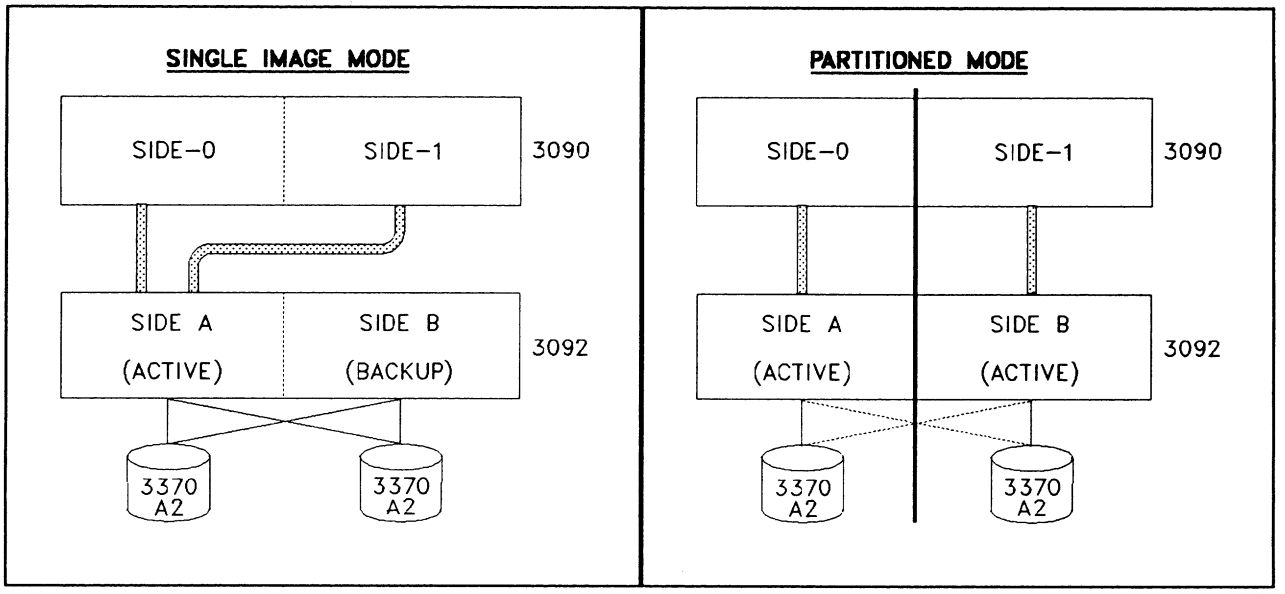

Figure 14. PCE Operating states of a 3092 Model 2.

PCE Operating States - Single Image

This section applies to the 3092 Model 2 Processor Controller for models 400, 400E, 600E in Single Image mode and the 3092 Model 1 Processor Controller. The 3092 Model 3 for the 3090 Model 120E contains a single processor that is always active.

There are various operating states associated with an individual side of the 3092 Processor Controller. The "active" state means that the side is currently controlling the 3090 Processor Complex. The "backup" state means that the side is not currently controlling the 3090 Processor Complex, but is available to assume the "active" state should the active side fail. The "offline" state means that the side is not available to assume the "active" or "backup" state, because of either a repair action or patch application. During system initialization, if both sides are functional, the A-side becomes the active side and the B-side becomes the backup side.

During operation, with both sides of the 3092 Processor Controller functional and available, there is constant communication between the two sides about their individual states. This constant communication between the two sides allows either the "active" or "backup" side to detect that the other side has failed. Thus, an "active" side failure can be detected either by the active side itself, or by the backup side. A "backup" side failure can be detected either by the backup side itself, or by the active side.

Consider the following failure scenarios:

- If the active side fails and the failure is not hardware related, the active side will initiate a recovery process. The recovery process will be retried several times. If recovery is successful, operation continues with the currently "active" side. If recovery is unsuccessful, or if the error was hardware related, the failing "active" side will attempt to notify the backup side to take over the

active assignment. If this notification is not possible, then the backup side should detect the active side failure because of a break in communication with the active side.

If the backup side detects the failure of the active side, it will initiate an automatic switchover and assume the active function.

- If the backup side fails, the "exerciser" which is continuously running in the backup side will attempt to notify the active side of the failure. If notification is successful, the active side will take the backup side offline. If the failure precludes notification (e.g., a hardware failure), then the active side should detect the backup side failure because of a break in communication between the two sides.

  If the active side detects the backup side failure, then it takes the backup side offline.

In either failure situation above, an automatic call will be placed, with operator approval, to the Remote Support System using the Remote Support Facility. A modem and auto-call unit are always associated with the active side of the 3092 Processor Controller, either because of switchover (Model 1) or both sides have one (Model 2).

PCE Operating States - Partitioned

This section applies to the 3092 Model 2 Processor Controller for models 400, 400E, 600E in Physical Partitioned mode.

In Physical Partitioned mode each side of the complex is controlled by one side of the PCE. There is no "backup" and therefore, no recovery from solid PCE hardware failures is possible.

Remote Support Facility

The Remote Support Facility of the 3090 processors allows automating problem diagnosis and isolation, and dispatching of the IBM Service Representative. The Remote Support Facility can be invoked automatically by the system as well as manually by the user. All machine-generated or remote initiated calls must be authorized by the operator before being dialed (unless pre-authorized from the RSFCNF frame).

If the 3092 Processor Controller detects a hard system error or exceeds thresholds for soft errors, it will invoke the Remote Support Facility. The Remote Support Facility will automatically dial the Remote Support System, and when connected, will transmit error data to a remote support application running on the Remote Support System. The remote support application analyzes the error type and notifies the IBM dispatch system to dispatch a service representative, along with any parts needed to repair the system.

Problem Analysis

The user can also initiate a call to the Remote Support System using the Remote Support Facility. In this situation, the 3092 Processor Controller has not detected a problem, however the user perceives there is a problem with the system, and wishes to report it to IBM.

To attempt recovery before initiating the request for remote service, the user should invoke the problem analysis (PA) facility from the system console index frame. If there was a power malfunction, a set of problem analysis frames displays information on the status of power boundaries and lists suggested recovery actions. For other malfunctions, status information is displayed for customer-selected categories including:

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

Each procedure in the problem analysis gives current status information for that type of malfunction, and then lists possible recovery actions. The IOPD frames may also be used when troubleshooting IFCC and I/O device problems. If the attempt at recovery fails and if remote service is required, a selection from the frame invokes the RSF authorization frame to begin the process of placing a call.

Service Update

The Remote Support Facility is also used during the service update process. A service update is scheduled once a week and its purpose is three-fold: 1) to test the Remote Support Facility communications link; 2) to receive any high impact patches; and 3) to transmit machine unique data to an IBM support system. The machine unique data consists of the patch history, Remote Support Facility configuration data, the Remote Support Facility call transaction log, and the System Status Recording (SSR) console file.

System Status Recording

The System Status Recording file is intended to provide a statistical data base of IBM 3090 processor operational history. This information is primarily for the use of IBM 3090 Product Engineering, to provide trend reports and statistical summaries of overall 3090 product behavior. In those cases when an Engineering evaluation of a specific IBM 3090 processor is required, the SSR file can provide supplemental information to Product Engineering for developing support plans.

The System Status Recording file currently collects the following types of data:

- Power On/Off occurred
- IPL invoked
- Restart invoked
- Disabled Wait State Entered
- 3092 Warmstart occurred
- 3092 Switchover occurred
- 3092 Reconfiguration action occurred
- RSF Call Request issued
- Selected 3092 internal process failure messages

Single Points of Failure

A single point of failure is a hardware failure that causes the system control program (SCP) to fail. The SCP, especially an SCP like MVS/XA, can be a powerful tool in maintaining system availability. A system includes an active SCP, a processor complex, and configured I/O devices. When the SCP fails (SCP not active), the system is also down.

Although single points of failure are generally discussed as hardware failures that cause the SCP to fail, software failures can also cause the SCP to fail.

Single points of failure include the following:

- Most power/thermal failures
- Uncorrectable errors (UEs) in processor storage when the UE is in a critical storage area (like HSA)
- I/O device failures, when the device is a critical resource
- Central storage controller failure
- System control element (SCE) failure
- Central processor failure in the Models 120E, 150, 150E, 180, and 180E.
- Software failures, such as application programs

If a single point of failure occurs in Model 400/400E/600E in a Single Image configuration, the operator can always partition the processor complex and IPL the system on the operational side in a Physically Partitioned configuration. The operator should perform the least disruptive recovery action, but partitioning is always possible.

Logical Components Summary

The following figures show the conceptual layout of the unique design implementations of each member of the IBM 3090 processor family.

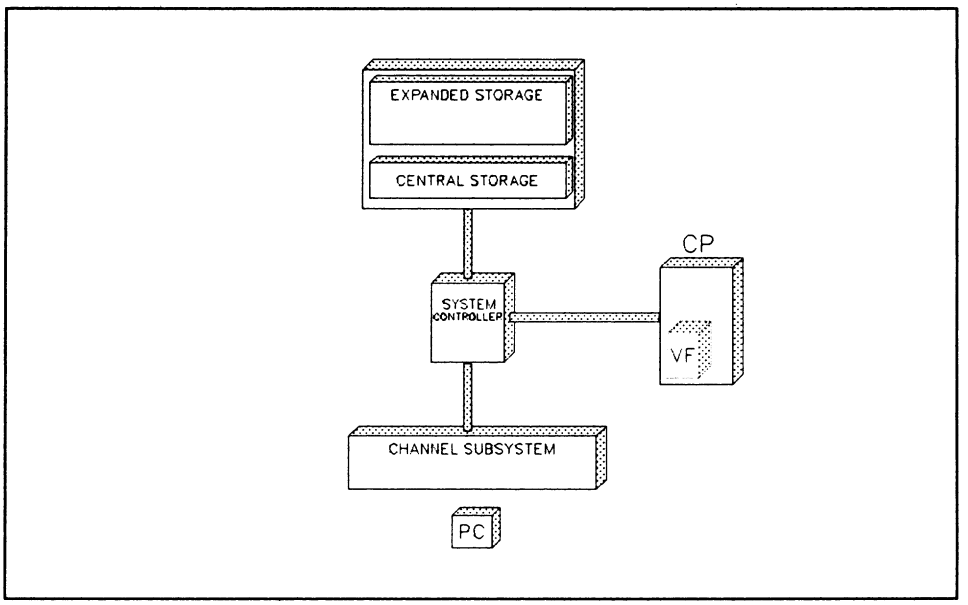

Figure 15. Uniprocessor Design - Models 120E, 150/150E, 180/180E

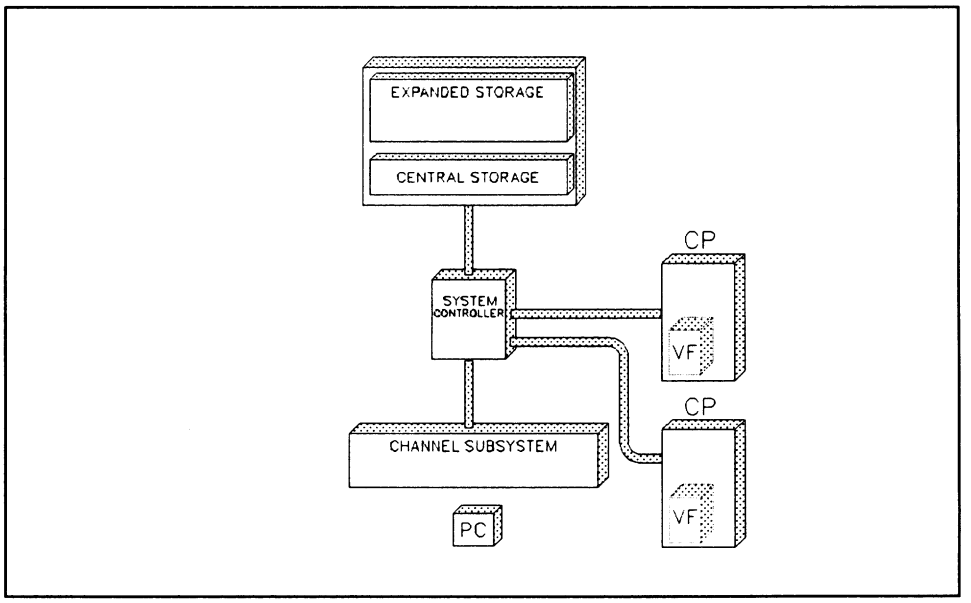

Figure 16. Two-Way (Dyadic) Processor Design - Models 200/200E

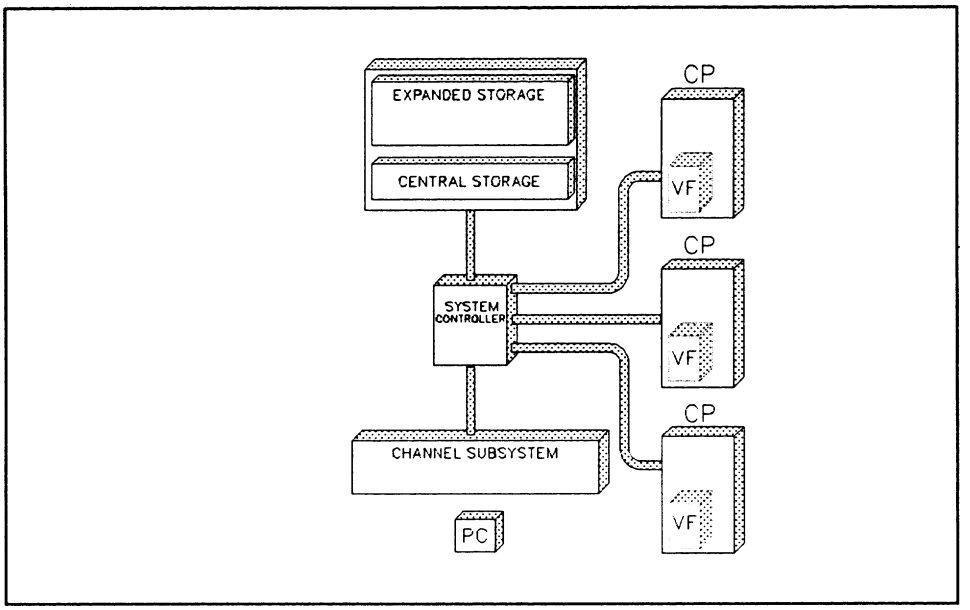

Figure 17. Three-Way (Triadic) Processor Design - Model 300E

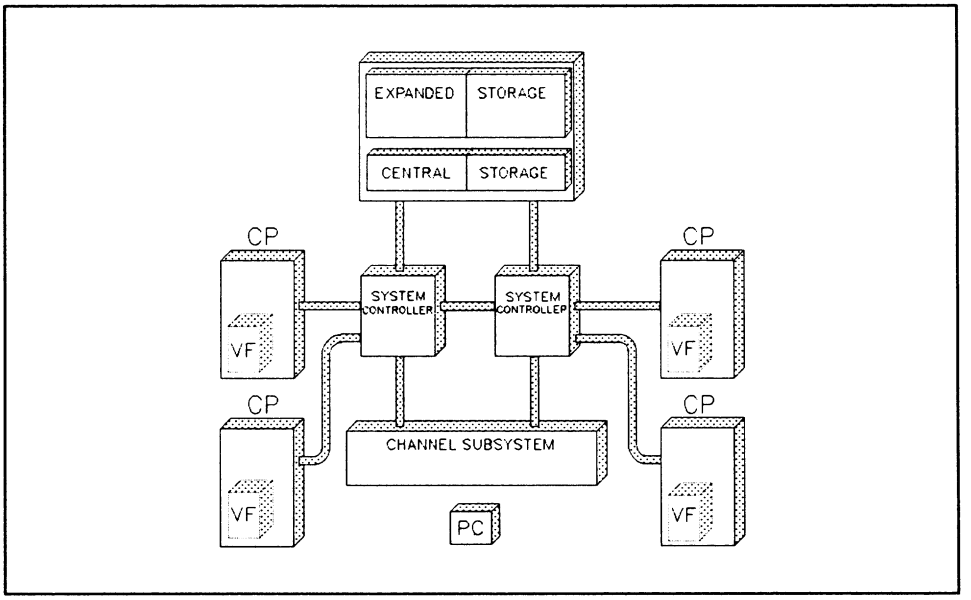

Figure 18. Four-Way Processor Design - Models 400/400E

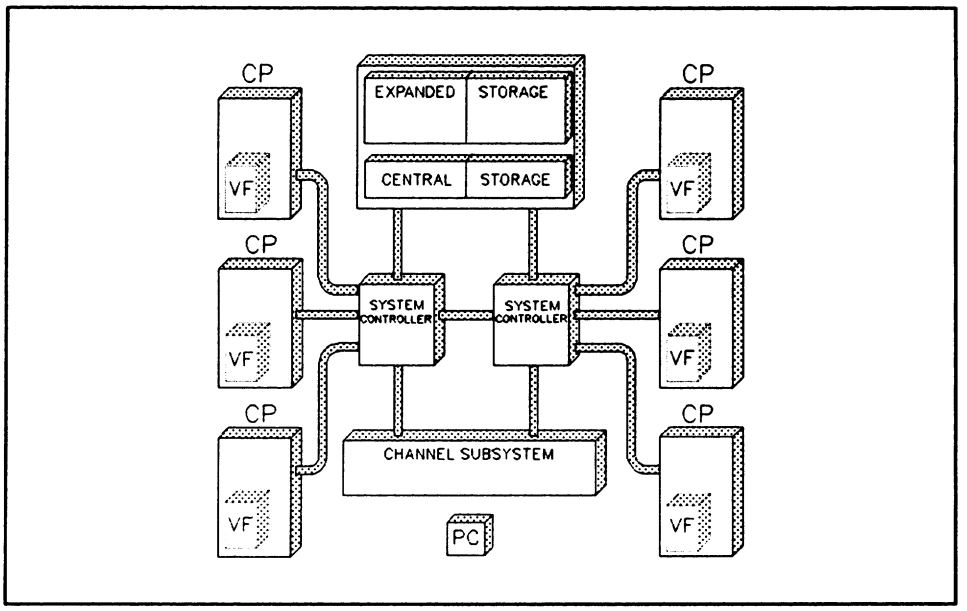

Figure 19. Six-Way Processor Design - Model 600E

Chapter 3. Software Considerations

This chapter contains an overview of supporting software for the 3090 family. It shows the changes incorporated into MVS/XA, MVS/370, VM/HPO, VM/XA SF and the major supporting products (JES, EREP, IOCP and RMF levels).

MVS/XA Support

The purpose of the section is to describe the support incorporated into various MVS/XA releases and features in support of the 3090 processors. The reader should not conclude from the following descriptions that all MVS/XA releases support all 3090 processors in all modes of operation. To determine which 3090 models are supported by the various MVS/XA releases, the reader should refer to Chapter 1, "IBM 3090 Family Comparisons." Before finalizing any plans the reader should verify the appropriate support level with the local IBM representative.

Initial support for the 3090 processors was incorporated into MVS/SP 2.1.3. This support addressed primarily the 3090 Model 200 without the optional Vector Facility feature. The 3090 Vector Facility optional feature support was incorporated into MVS/SP 2.1.3 VFE (Vector Facility Enhancement). Support for the 3090 uniprocessor models (Model 150 and 180) was incorporated into MVS/SP 2.1.3, MVS/SP 2.1.3 VFE and MVS/SP 2.1.3 AE (Availability Enhancement)[5] features via PTFs made available prior to general availability of the uniprocessor models.

Support for the 3090 Model 400 was incorporated into MVS/SP 2.1.7 (both single image and partitioned mode operation). MVS/SP 2.1.7 also integrated the Vector Facility Enhancement feature (MVS/SP 2.1.3 VFE) and the Availability Enhancement feature (MVS/SP 2.1.3 AE). Prior to general availability of the 3090 Model 400, PTFs were also made available to MVS/SP 2.1.3 (including the VFE and AE enhancement features) to support the 3090 Model 400 in partitioned mode only.

Support for the E-Models of the 3090 Processor family was made available as PTFs to MVS/SP 2.1.7 and MVS/SP 2.1.3 (including the VFE and AE enhancement features). The MVS/SP 2.1.7 E-Model changes provide the support for increased central storage (greater than 128 MB), increased expanded storage (greater than 512 MB), additional channels (greater than 96), new version codes, SRM constants, and Models 300E and 600E (3-way and 6-way support). The MVS/SP 2.1.3 E-Model changes provide the support for the new version codes for Models 120E, 150E, 180E, 200E, and 400E (partitioned mode only).

MVS/SP 2.2.0 incorporates all 3090 support made available in the prior releases (and PTFs) of MVS/XA. MVS/SP 2.2.0, via PTFs (May, 1988), allows VIO paging to expanded storage.

---

[5] The MVS/SP 2.1.3 Availability Enhancement provides operating system support for the Extended Recovery Facility (XRF). It will not be addressed in this document.

MVS/SP 2.1.3

Overview

MVS/SP 2.1.3 support of the 3090 processors occurred in several functional areas. The SRM was updated over time to add the 3090 models to the CPU adjustment table. This table is used by SRM in its resource management algorithms which are affected by processor speed.

Real Storage Management, Systems Resources Management, and Stand-alone Dump were all updated to support expanded storage. A RAS[6] item (frame deallocation) required changes to RSM and the Machine Check handler. Expanded storage support and frame deallocation support will be described in more detail later.

The following additional product versions and/or features were the minimum level required for initial 3090 operation under MVS/XA 2.1.3.

- RMF 3.3
- EREP 3.1 with Feature Number 3 (FMID = FER3130)
- DFP 1.2 PTF (UZ72308) for SYSGEN toleration of a new IOCP keyword on the IODEVICE macro
- IOCP Program - IOPIOCP V1L0 (FMID = HIO1104)

RMF V3.3 provided support for 3090 processor-dependent statistics in addition to expanded storage. RMF V3.3 was a complete replacement of the previous level of RMF. RMF V3.4 incorporates all the functions of the previous releases of RMF Version 3, and included data set support and new reports for Monitor III, as well as support for the 3090.

DFP V1.2 required PTF UZ72308 to support a new IOCP keyword (STADET=) on the SYSGEN IODEVICE macro. There were no changes required to executable DFP V1.2 code for the 3090 processor.

EREP V3.1 with feature 3 contained the initial support for the 3090. Feature 3 contains the formatter exits for the 3090 model dependent information.

The IOCP program for the 3090 processor is a new IOCP program. While its base was the existing 308X IOCP program, this new program is only applicable to the 3090 processor. Since it was derived from the existing 308X program, the macros and most of the keywords/values are the same. While some changes will have to made to an existing 308X IOCP input stream to support the 3090 processor, these changes will be minor.

---

[6] Reliability, Availability, Serviceability

Expanded Storage Support

Expanded Storage (ES) is optional processor storage which may be used to improve the responsiveness and throughput characteristics of the system. It is managed by MVS/XA in a fashion similar to central storage and provides a high performance alternative to DASD paging. Expanded storage may be thought of as part of a hierarchy of storage used to contain instructions and data. It resides logically between central storage and auxiliary (DASD) storage in terms of performance and capacity.

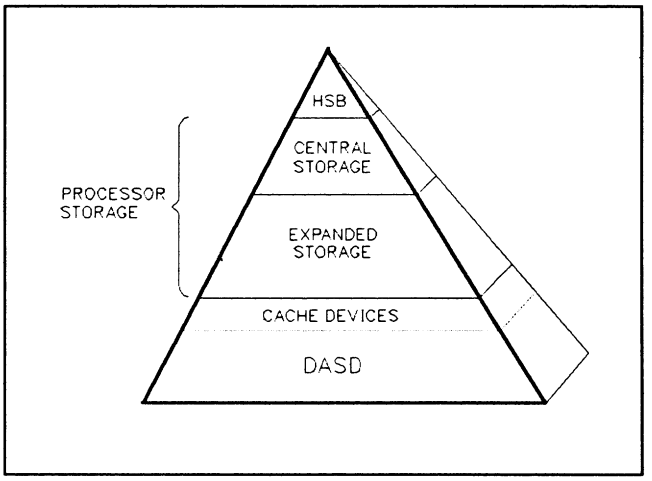

Figure 20. Storage Hierarchy

Expanded storage is managed by the Real Storage Manager (RSM) and System Resource Manager (SRM) components of MVS/XA. There are two phases to the RSM implementation for expanded storage. The first can be called the *selection* phase. Here RSM determines which pages should be sent to expanded storage and which pages should go to auxiliary storage. The second phase can be called the *migration* phase. It is the process of removing unreferenced pages from expanded storage and placing them on auxiliary storage.

Whenever a page is removed from central storage RSM must decide whether the page should be directed to expanded storage or auxiliary storage. There are a number of reasons the page may be leaving central storage, including SRM control decisions (swapping), and replenishment of the available frame queue (page stealing). Rather than unconditionally schedule the page for transfer to auxiliary storage, the page is considered as a candidate for expanded storage. The objective is to use expanded storage for work which requires a high level of responsiveness. The pages selected for expanded storage should also have a good probability for near term reuse. This is similar to the objectives for the use of central storage. The decision is based on the characteristics of the work and the contention for both central and expanded storage.

Expanded Storage Initialization: At system initialization time a Resource Initialization Module (RIM) allocates and initializes a set of control blocks and queues related to expanded storage. Each 4K of expanded storage is represented by an E-frame in a manner similar to a page frame of central storage. Queues are maintained of available E-frames and in-use E-frames.

The Selection Process: A table, called the Criteria Age Table, is also built during initialization reflecting a relative responsiveness requirement for broad categories of work. Each entry in the table establishes a relative priority for placing pages associated with that category of work in expanded storage. Categories are established based on the nature of the work (TSO, batch, etc.) and the type of pages (common, working set, etc.). Defaults are provided which are expected to suffice for the majority of cases. The user may override the default values by using the IEAOPTxx member of the 'SYS1.PARMLIB' dataset. The 'relative priority' contained in this table is used to decide whether a candidate page should be placed in expanded storage. The value is referred to as 'criteria age'.

Storage Isolation: Storage isolation allows an installation to protect the working set of a critical application. Storage isolation now applies to working set pages residing in both central and expanded storage. The migration process will try and honor the storage isolation values and not migrate an application's pages, if that would cause it to drop below its target working set size.

VIO Datasets: VIO datasets themselves are not maintained in expanded storage. However, stolen VIO pages that are part of the working-set of an address space can be sent to expanded storage.

Access to Expanded Storage: During system operation, when it becomes necessary to bring a page into central storage due to a page fault or swap-in request, RSM checks to see if the page resides in expanded storage or on auxiliary storage. If the page resides on auxiliary storage, the ASM is requested to page it in. If the page resides in expanded storage however, RSM transfers it directly from expanded storage to central storage. If no errors were encountered during the transfer process, the E-frame is freed and the address space is dispatched.

Error Handling: If a permanent error occurs during the transfer of a page from expanded storage to central storage, RSM determines if a good copy of the page resides on auxiliary storage. If auxiliary storage contains a good copy, it is used to satisfy the request, the E-frame which contained the error is marked unavailable, and the operation continues. If a good copy did not exist on auxiliary storage, the requestor is ABENDED (just as if there had been an error on a paging device). If a page was being transferred to expanded storage when the error was detected, RSM simply assigns another E-frame to contain the page, marks the E-frame containing the error unavailable, and the operation continues.

The Migration Process: Expanded storage is large, but finite, so a method must be provided to control over-commitment.

The selection process described earlier placed pages in expanded storage based on estimates of reuse. In time, pages will collect in expanded storage which are not being re-referenced. When the number of available E-frames falls below a threshold monitored by RSM, a process called 'migration' is begun. Once started, migration continues until the queue of available E-frames rises above a high threshold. The process involves transferring pages which have gone unreferenced for the longest period of time to auxiliary storage. You will recall that there is no capability to do I/O directly to/from expanded storage. Therefore, the migration process first transfers pages from expanded storage to central storage, then interfaces with the ASM to have the pages written to auxiliary storage. The selection of pages for migration is based on an approximate Least Recently Used (LRU)

algorithm. Note that the transfer of the page to auxiliary storage is not necessary if a valid copy already exists (unchanged page) on auxiliary storage. In this case the overhead of the physical I/O is avoided.

If possible, storage isolation targets will be honored by the migration process. If migration is unable to bring the number of available expanded storage frames above the high threshold, it will make another pass ignoring the storage isolation targets.

RMF Support for Expanded Storage: RMF V3.3 provides statistics on the use of expanded storage. RMF is discussed in detail starting on page 45. Both Monitor I and Monitor II were updated to report on expanded storage.

Operator Command Changes

D M=ESTOR Command: The 'D M=ESTOR' command reports on the amount of expanded storage available for use. The command also reports on the amount of expanded storage unavailable due to uncorrectable storage errors.

Display Matrix Command: Minor changes were made to the syntax of the 'Display Matrix' command to make it match the syntax of the CONFIG command. The command format to display the status of channel paths to a device was changed. The old format was 'D M=nnn'. The new format is 'D M=DEV(nnn)'. A range of addresses, or list of addresses, can also be specified. Also, the CHP and CPU keywords now allow operands so the command can be tailored to show only pertinent data. Operational procedures should be updated to reflect these changes.

Frame Deallocation

Frame deallocation supports a hardware RAS enhancement. When the hardware detects certain double bit errors associated with central storage it corrects the error and presents a machine check. The Machine Check Handler calls the Real Storage Manager to attempt to free the storage. If the storage frame in error backed a valid pageable page, RSM will obtain a new frame into which the data will be moved. If the storage frame does not back a valid pageable page, frame deallocation will not be scheduled. The frame will, instead, be marked as a bad frame. The frame will remain allocated to the page and will be marked offline when the page is freed.

Frame deallocation causes frames with correctable errors to be marked offline so they will not be reallocated. Previously, frames with data errors could only be taken offline when they were backing unchanged pageable pages, or later when being freed (e.g.. when the pages were FREEMAINed at task termination).

This support enhances system operation by taking the frame offline as soon as possible and avoiding continued hardware correction of the frame in error. There is also a possibility that the double bit error could turn into a triple bit error. A triple bit error in a storage frame would probably result in termination of the work using the frame.

Dispatcher Enhancement

The MVS dispatcher, in an attempt to stay as responsive as possible, consumes some additional CPU time when the system is running at less than 100% utilization. If no ready work is found after scanning the dispatching queue, the dispatcher proceeds to make one additional pass to insure that no units of work anywhere in the system are able to use the processor. This additional pass of the dispatching queue is known as the dispatcher 'recursive scan'.

This recursive scan consumes some amount of CPU time. The amount varies with the number of address spaces on the queue, and the frequency of entering the wait state. The effect is greatest at low utilizations, but can still be significant at fairly high utilizations. It is important to note that this increase in CPU utilization does not affect the systems capability to do useful work. The recursive scan is done enabled, and the dispatcher can be preempted by newly readied work. It can be thought of as additional time spent before entering the wait state.

Prior to MVS/SP 2.1.3, on systems that used logical swapping, there was an additional *Low Utilization Effect*. A logically swapped address space was placed on the ASCB ready queue at a priority lower than the active address spaces. The dispatcher scanned all of the logically swapped address spaces looking for units of work on both the first scan and on the recursive scan. Measurements made with sophisticated instrumentation (hardware monitor) indicated that this could be a major part of the low utilization effects in TSO systems. This is because TSO systems often have a high number of logically swapped address spaces, with a relatively small number of active address spaces. However, this search of logically swapped address spaces had a minimal impact on system performance because it was only done when there were no known units of work waiting for the processor.

The additional CPU time could, however, cause a distortion in performance measurements and capacity planning efforts. From a performance measurement point of view, this effect caused an understatement of Internal Throughput. This was especially true with tightly coupled processors. Remaining (unused) processor capacity was also understated, which may distort capacity planning calculations.

MVS/SP 2.1.3 incorporates a change to the dispatcher to bypass these logically swapped users during the first and recursive scans. This change can produce a noticeable decrease in processor utilization with heavy TSO workloads, especially when running at relatively low utilizations.

Available Frame Queue Management

The SRM controls the thresholds used to replenish the available queue of real storage frames used for allocation of GETMAINs, page faults and swap-in requests. Two constants are initialized at IPL time to control the replenishment of frames. The available frame queue low threshold signals that page replacement needs to run when the number of frames on the available frame queue equals the threshold. When the number of frames equals the available frame queue high threshold, RSM signals SRM that the replacement process can stop.

The labels for the constants and the values that are used at initialization have been included in Part 5 (System Resources Manager), Section 6, of the *MVS/XA Initialization and Tuning Guide, GC28-1149*, since some installations have felt that the constants used were not appropriate for them.

The SRM does change the thresholds during system execution if insufficient frames are available to allow a swap-in of an address space. The thresholds are increased by the swap-in working set size of the address space until the frames are accumulated and allocated. After the frames are allocated, the thresholds are decreased by the swap-in working set size.

In MVS SP 2.1.3, if no expanded storage is available for use, the real storage available frame queue thresholds are increased and are managed based on the amount of time that the available frame queue is unable to satisfy demand. The available frame queue low threshold is initialized to approximately one and one half times the value used in earlier systems. This allows for swapping in a ready address space with less delay to accumulate the frames for the working set. It decreases the delay to load programs which may need several pages at once. The thresholds are no longer simply constants but are adjusted up and down based on the amount of time that there are insufficient frames to satisfy new requests for storage. The goal is that the available frame queue should be a source of available frames nearly 100% of the time. The SRM will increase the values as appropriate based on both swap-in demands and on real storage availability when no expanded storage is in use.

If expanded storage is available for use, the expanded storage available frame queue is managed as described above. The amount of time that the expanded storage available frame queue is unable to provide a source for real storage page replacement controls the setting of the expanded storage thresholds. The expanded storage available frame queue is the resource to be managed since a real frame can always be obtained provided there are available expanded storage frames. A suitable real storage frame is chosen and the contents transferred to expanded storage, freeing up the real frame. The real storage frame thresholds are managed as in previous systems when expanded storage is part of the configuration.

Spin-Loop Recording

Support has been included in MVS/SP 2.1.3 to allow the Excessive Spin Notification routine (IEEVEXSN) to initiate LOGREC recording[7] on a CP that is causing an excessive spin condition. The recorded data may allow subsequent identification and analysis of the routine causing the spin. After it initiates LOGREC recording, IEEVEXSN proceeds to issue the message for which it was called.

This support also changes the 09X wait-state PSW so that it contains (in its sixth byte) the logical ID (4x) of the CP that is causing the spin. **This is not the CP that is in the 09X wait state.** For example, if CP 0 is in a '092' wait state because of CP 3, the wait-state PSW for CP 0 would be X'000A0000 00430092'.

---

[7] This support was previously made available for MVS/SP 2.1.1 and 2.1.2 via PTF maintenance. See APAR II01984 description for additional information.

If Alternate CPU Recovery (ACR) is to be the response to this wait-state condition, the 4x in the PSW allows the operator to identify the CP to be stopped without having to refer to storage location X'40C' for the pointer to the CP ID. This pointer, however, will still be put at X'40C'.

This support also changes the excessive-spin loop factor from 20 to 40 seconds. This is done to limit recording to only those events that are the cause of the most serious spin-loops.

Reason for Support: This support is intended to collect data on the causes of excessive spin conditions and to provide limited diagnostic data. If the spin conditions occur frequently enough to affect system operation, the data can be used to determine what function was in control and, if necessary, to develop a trap to further isolate and correct the cause of the spin loop. Since the data is collected in a software record, it does not affect the processing of an excessive spin condition.

Conditions Under which IEEVEXSN Initiates Recording: IEEVEXSN will initiate recording on a CP causing an excessive spin only when the following conditions are met:

1. The call to IEEVEXSN was for a reason other than message IEA490A. No recording is done for IEA490A because this message indicates that hardware is the cause of the spin-loop rather than the software routine running at that particular time.

2. The call to IEEVEXSN was for a reason other than a caller's inability to obtain the restart resource, and IEEVEXSN can obtain the restart resource.

3. At least five minutes have passed since IEEVEXSN last issued a restart. This prevents multiple records for the residual effects of a single problem.

Contents of New LOGREC Record: An excessive-spin LOGREC record is identified by the values X'071' in the system abend code field, X'10000000' in the abend reason code field, and 'EX SPIN' in the jobname field. Although the abend code (X'071') usually indicates an operator initiated restart, no operator action is involved in the generation of the spin loop LOGREC record.

The VRADATA portion of the SDWA will contain:

- Identification text 'EXCESS SPIN RESTART TO RECORD'.
- The 16 FRR addresses on the stack saved by the RESTART FLIH in PSARSAV.
- An index value designated 'INDEX=x' where x is a number from 0 to 16. This number indicates which of the 16 addresses represents the current FRR. If x=0, there are no FRRs on the stack.
- The 16 control registers on entry to IEAVTEXS.
- If RTM was in control when IEEVEXSN issued the restart, the original completion code, reason code, and cross memory registers from the RT1W control block.
- If RTM was not in control when the restart was issued, the excessive spin length factor is saved.

Analysis of Excessive Spin LOGREC Records: The excessive-spin LOGREC records will help to identify the MVS routine that is running on the CP that is causing the spin condition. This will enable installations to take further diagnostic action and identify the correct MVS component if IBM Support is desired. The key data for identifying the MVS component is the FRR data in the Variable Recording Area (VRADATA) of the SDWA[8].

Analysis of these LOGREC records to identify the cause of the excessive spin should consist of the following:

1. Locate the 16 addresses from the stack that was current when recording was initiated on the target CP. These addresses follow the identification text 'EXCESS SPIN RESTART TO RECORD' at the start of the VRADATA area.

2. Identify the current FRR on that stack from the INDEX=x value that follows the 16 addresses. The 'x' in this field is a number from 0 to 16. If it is 0, there are no FRRs on the stack. Otherwise 'x' is an index indicating which of the 16 addresses points to the current FRR. For example, if x=2, the second FRR in the list is the current one.

3. The component occupying this address can then be identified from a storage map. Traps can then be set to determine which function within that component is causing the excessive spin.

SMF Changes

The following SMF record types were changed in MVS/SP 2.1.3.

- The SMF Type 22 (Configuration) record was changed to record information on the amount of expanded storage on the system. For each contiguous block of online expanded storage, SMF record type 22 will contain one expanded storage section. Each expanded storage section is ten bytes in length. For additional information pertaining to the content of the ten byte section, see the SMF manual. Programs using the mapping macro for this record do not need to be recompiled unless they need to access the new data.

- The SMF type 90 (System Status record) subtype 3 (SETDMN section) was changed to support the new maximum multiprogramming level supported by SRM. The maximum multiprogramming level was increased from 256 to 999. The fields in this section have been increased from one to two bytes to handle the larger value. **Any program using SMF type 90 subtype 3 records must be recompiled.**

---

[8] System Diagnostic Work Area

The following SMF records built by RMF were also changed. **Programs using these records will have to be recompiled.**

- The SMF Type 71 (Paging Activity) record was updated to include a Swap Placement Data Section.

- The SMF Type 72 (Workload Activity) record was updated to include a new field (SMF72ADJ). This field contains the CPU time per service unit (in microseconds) multiplied by 16. This field is used to derive the number of service units per CPU second.

- The SMF Type 74 (Device Activity) record was updated to include new fields reporting on device busy delay time and control unit busy delay time.

- The SMF Type 78 (Monitor I Activity) record was updated to include additional I/O statistics reported by the channel subsystem on 3090 processors.

- The SMF Type 79 (Monitor II Activity) had the following changes:

  - The Sub-type 4 (SPAG) section was updated to include information pertaining to expanded storage.

  - The Sub-type 10 (DDMN) data section had the minimum and maximum multiprogramming level fields expanded from one byte to two bytes and placed at the end of the section.

  - A new Sub-type (Sub-type 14 - I/O Queueing Configuration Control Section) was added. This new section contains logical control unit queueing information.

RMF 3.3

RMF 3.3 provides support for 3090 processor dependent statistics in addition to expanded storage. Additional changes were made to the I/O activity reports to take advantage of processor model dependent measurements. RMF 3.3 is a complete replacement for the previous level of RMF.

RMF 3.4 also supports the 3090 processor and incorporates all the functions of the previous releases of RMF 3. Also included were enhancements to Monitor III (new reports and dataset support).

The data for the channel activity report is now collected by the processor controller. The processor controller samples much more frequently than the operating system did using the STCPS (Store Channel Path Status) instruction. If for some reason the processor controller is not collecting the data, the old method of STCPS will be used.

I/O QUEUEING ACTIVITY

LCU	CONTENTION RATE	DELAY Q LNGTH	% ALL CH PATH BUSY	CONTROL UNITS	CHAN PATHS	CHPID TAKEN	% CU BUSY
02E	2.462	0.01	5.87	091	09	6.445	34.10
				A91	15	11.086	8.24
03C	.026	0.00	0.01	0E3	0E	1.091	2.12
				AE3	13	1.056	3.30

Figure 21. Sample I/O Queueing Activity Report for 3090 processor

The I/O Queueing Activity Report provides information, grouped by logical control unit (LCU), on the I/O configuration and activity rate, queue lengths, and percentages when one or more I/O components were busy. An LCU consists of all the devices accessed through a common set of control units. The format and contents of the report are model dependent. Figure 21 reflects the report output if running on a 3090 processor.

D I R E C T A C C E S S D E V I C E A C T I V I T Y

DEV NUM	VOLUME SERIAL	LCU	DEVICE ACTIVITY RATE	AVG RESP TIME	AVG IOSQ TIME	**AVG CUB DELAY**	**AVG DB DELAY**	....
440	335003	024	3.564	112	0	**0.3**	**23.4**	
441	335004	024	7.906	70	0	**0.1**	**0.0**	
442	335005	024	0.011	61	0	**0.2**	**18.8**	
.	..	.	. .	.	.			
.	..	.	. .	.	.			

Figure 22. Sample I/O Device Activity Report

The Direct Access Device Activity portion of the 'I/O Device Activity Report' now reports the average device busy delay and the average control unit busy

delay for each device. This data will allow you to find bottlenecks on a logical control unit or a single volume.

AVG CUB DELAY is the average number of milliseconds of delay an I/O request encountered because the control unit was busy. The data field heading appears in the reports for all processors, but data appears in the field only when RMF is running on a 3090 processor. The field is left blank for 308X and 4381 processors.

AVG DB DELAY is the average number of milliseconds of delay that I/O requests to this device encountered because the device was busy. The data is collected by the hardware and obtained from the model dependent portion of the SCHIB (Sub-CHannel Information Block).

RMF Paging Activity Report

```
                 P A G I N G   A C T I V I T Y

                                          | REAL STORAGE MOVEMENT
          PAGE OUT                        | TOTAL RATE      51.74
 ------------------------------           |
 . . .                                    | ---------------------
                                  %       | RATE OF PAGE MOVEMENT
 . . .          NON    TOTAL    TOTL      | TO ES          465.74  (*1)
      SWAP     SWAP    RATE     SUM       |
 . . .                                    | ---------------------
                                          | MIGRATION RATE  15.75  (*2)
 . . .                                    |
                                          | ---------------------
 . . .                                    |    AVAILABLE ESF       (*3)
               0.00    0.00       0       |   MIN       57
 . . .         -----   -----   ----       |   MAX     1619
               0.00    0.00       0       |   AVG      343
 . . .                                    |
               8.36    8.36      39       | ---------------------
 . . .                                    |   HIGH UIC  MIGR AGE   (*4)
      12.17    0.80   12.96      61       |   MIN   18       499
 . . .-----   -----   -----    ----       |   MAX   74       684
      12.17    9.16   21.33     100       |   AVG   43.8     585.9
 . . .                                    |
                                          | ---------------------
 . . .         8.36    8.36      39       |   ESF CONFIGURATION    (*5)
      12.17    0.80   12.96      61       |   INSTALLED    ONLINE
 . . .-----   -----   -----    ----       |    32768       32768
      12.17    9.16   21.33     100       |
```

Figure 23. Sample Paging Activity Report (Main Storage Paging Rates)

The paging activity report provides information about the demands made on the system paging facilities. Starting with RMF 3.3, RMF will provide information in support of expanded storage on the 3090 processors. The RMF Paging Activity Record (SMF type 71) has been changed to include this new information. New fields plus an overlay structure to support the swap activity have been implemented.

The Paging rate section remains basically the same with additional fields to provide information about page transfer activity from real storage to expanded storage. The following is an explanation of these new fields.

- **RATE OF PAGE MOVEMENT (*1)** - This field indicates the rate per second of page movement from real storage to expanded storage. It includes all page movement, local and swap. If you subtract the MIGRATION RATE from the RATE OF PAGE MOVEMENT you have a measurement that indicates the reduction in I/O activity to auxiliary storage.

- **MIGRATION RATE (*2)** - This field indicates the rate per second at which pages are migrating from expanded storage to auxiliary storage. It includes both local and swap pages. A low migration rate is desired, but remember a low migration rate may exist because of under utilization of expanded storage. If a high migration rate exists you have over committed expanded storage. This may have an adverse effect on performance.

- **AVAILABLE ESF (*3)** - This field contains three counters that track expanded storage frames (the MIN, MAX, and AVERAGE number of frames available). If the migration rate to auxiliary storage is low and the MIN counter is reporting a high number of available frames one could conclude that expanded storage is under utilized. A low migration rate to auxiliary storage and the MIN counter recording a small number of available frames would indicate effective use of the available expanded storage.

- **HIGH UIC (*4)** - This field contains three counters that keep track of the UIC (Unreferenced Interval Count). The MIN, MAX, and the AVERAGE age of real storage frames are reported.

- **MIGR AGE (*4)** - This field contains three counters that keep track of the MIN, MAX, and AVERAGE age of expanded storage frames.

- **ESF CONFIGURATION (*5)** - This field contains two counters, one that shows the number of INSTALLED expanded storage frames, and one that shows the number of ONLINE expanded storage frames. If these fields are not equal then one or more expanded storage frames have been taken offline due to an error.

Swap Placement Activity Report: The swap placement section contains information about the demands made on the system's swapping facilities. It reports data for logical swapping, physical swapping to auxiliary storage and physical swapping to expanded storage. Column headings (from left to right in the report) describe how the SRM handles the swap requests and how effective the SRM is in the way it satisfies these requests. If expanded storage is not available or not used, the rightmost five column headings still appear, but data fields for count, rate and percentage are blank. The swap placement activity report gives the reasons for swaps such as terminal I/O wait, long wait, detected wait, unilateral swap-out and enqueue exchange, exchange on recommendation value, requested, auxiliary storage shortage, real pageable storage shortage and transition to non-swappable. If no swaps occurred during the measurement interval for a particular swap reason, the report omits that swap reason and its associated counts, rates and percentages.

The following is an explanation of the column headings:

- **TOTAL** - field contains the total number of swap requests. It represents the number of logical and physical swap candidates. The rate is calculated by dividing this number by the interval time in seconds.

- **AUX STOR TOTAL** - field contains the total number of swap requests directly or indirectly placed on auxiliary storage. It represents the number of AUXILIARY STORAGE DIRECT swaps plus the number of AUXILIARY STORAGE VIA TRANSITION swaps. The percent is obtained by dividing this number by the TOTAL number of swaps. This number represents the physical swaps to auxiliary storage and should be compared to the total swap count. A small number compared to the TOTAL swap count means that there is a high effective swap rate between LOGICAL and/or EXTENDED storage.

- **AUX STOR DIRECT** - field contains the number of swap requests directly placed on auxiliary storage. It is part of the AUX STOR TOTAL count. It represents the number of swap requests not logically swapped or directed to expanded storage. The percent is obtained by dividing this number by the AUX STOR TOTAL field.

- **AUX STOR VIA TRANSITION** - field contains the number of swap requests indirectly placed on auxiliary storage. It is part of the AUX STOR TOTAL count. It represents the AUX STOR TOTAL swaps minus the AUX STOR DIRECT swaps. It is the number of swaps placed on auxiliary storage that were either logically swapped from real storage to auxiliary storage or directed from real storage to expanded storage and then migrated to auxiliary storage. The percent is obtained by dividing this number by the AUX STOR TOTAL column.

- **LOG SWAP** - field contains the number of logical swap candidates. It is part of the TOTAL swap count. It represents the TOTAL swaps minus the AUX STOR DIRECT and EXT STOR DIRECT swaps. The percent is obtained by dividing this number by the TOTAL swaps.

- **LOG SWAP EFFECTIVE** - field contains the number of logical swap candidates that were logically swapped and never physically swapped. It represents the LOG SWAP / EXT STOR EFFECTIVE count minus the EXT STOR EFFECTIVE count. The percent is obtained by dividing this number by the number of LOG SWAP candidates.

- **EXT STOR DIRECT** - field contains the number of logical swap candidates that were originally directed to expanded storage. It is part of the EXT STOR TOTAL swap count. The percent is obtained by dividing this number by the number of TOTAL swaps.

- **EXT STOR TOTAL** - field contains the number of swap candidates that were directed to expanded storage. They include those swaps originally directed to expanded storage and those logically swapped that were physically swapped and directed to expanded storage. It represents the MIGRATED FROM EXT STOR swaps plus the number of logical swaps physically

swapped to expanded storage. The percent is obtained by dividing this number by the number of TOTAL swaps.

- **MIGRATED FROM EXT STOR** - field contains the number of swaps that were migrated from expanded storage to auxiliary storage. It is part of the EXT STOR TOTAL swap count and indicates the number of swaps, due to trimming, migrated from expanded storage to auxiliary storage. The percent is obtained by dividing this number by the number of EXT STOR TOTAL swaps.

- **EXT STOR EFFECTIVE** - field contains the number of swaps that were directed to expanded storage and not migrated to auxiliary storage. It represents the EXT STOR TOTAL swap count minus the MIGRATED FROM EXT STOR swap count. This number represents the effective use of expanded storage. The percent is obtained by dividing this number by the number of EXT STOR TOTAL swaps. A high EXT STOR EFFECTIVE percent is desired.

- **LOG SWAP /EXT STOR EFFECTIVE** - field contains the number of swaps that either remained logically swapped in real storage, or after being directed to expanded storage, remained there and were not migrated to auxiliary storage. This value indicates that I/O to auxiliary storage was avoided when the swap request was satisfied. It represents the TOTAL swap count minus the AUX STOR TOTAL swap count. The percent is obtained by dividing this number by the number of TOTAL swaps. A high LOG SWAP /EXT STOR EFFECTIVE percent is desired.

RMF Monitor II & Expanded Storage: RMF Monitor II also reports on expanded storage use.

```
                              CPU= 26  UIC=254  PDT=  87  DPR= 13  SPAG

PGS-SWPD PRIVATE    VIO  TAR        HI  ES  MIG  ESF  MIG
. . .   IN   OUT    I+O  CWS  AFC UIC RTE AGE  AVL  RTE
        ---  ---    ---  ---  --- --- --- ---  ---  ---
. . .   7.0   16     23    0  266 170  00 *** 1246  0.0
       10.4   47    150    0  240 190 147 ***  975  0.0
```

Figure 24. Sample Monitor II Paging (SPAG) Report

The Monitor II paging report (SPAG) has added support for the paging rate to expanded storage, the migration age, the amount of expanded storage currently available and the migration rate. In order to have room for reporting this new information, RMF no longer reports the reclaim counts for LPA, CSA and private pages, nor does it report the average high UIC.

```
                         CPU= 69  UIC=142  PDT=  16  DPR=  4  ASD

                  . . . DP  DP  RS         ESF  WS   TX . .
JOBNAME    DMN . . . PR  PO   F   ESF  +RS  IN   SX . .
*MASTER*     0 . . . FF   1  70   12    82   0    0 . .
PCAUTH       6 . . . 76  33  24    2    26  12    0 . .
TRACE        6 . . . 76  34  56   36    92  12    0 . .
```

Figure 25. Sample Monitor II Address Space Data (ASD) Report

The Monitor II Address Space State Data report no longer reports the transaction time (which is available elsewhere), in order to report on the use of expanded storage. This is for use in analyzing storage isolation effectiveness.

ESF is a one to three digit number indicating the number of expanded storage frames occupied by virtual storage pages associated with this address space. If no expanded storage is installed, this field is blank.

ESF + RS is a one to three digit number indicating the total number of real storage frames and expanded storage frames occupied by virtual pages associated with this address space. If no expanded storage is installed, this field is blank.

RMF Trace Support (Expanded Storage Fields): RMF also added trace support for four new expanded storage fields. These fields are valid and will contain data only when running on a 3090 processor.

- RCEESINU - the number of in-use expanded storage frames.
- RCEESWRT - number of pages written to expanded storage frames.
- RCENWSF - total number of secondary and non-working set pages migrated to auxiliary storage.
- RCEWSDNE - number of primary working set pages migrated to auxiliary storage.

MVS/SP 2.1.3 VFE (Vector Facility Enhancement)

Support for the 3090 processors Vector Facility was introduced in MVS/SP 2.1.3 VFE (Vector Facility Enhancement)[9]. Each CP of a 3090 processor may have an optional Vector Facility feature attached to enhance its performance in processing vectors of fixed and floating point numbers. The Vector Facility is an extension of the central processor's instruction and execution elements which allows the CP to execute vector arithmetic and logical operations on vectors of any length, from one element on up to 128 sets of operands, with a single instruction. The support provided by MVS/SP 2.1.3 VFE allows an installation to manage jobs using vector instructions by dynamically recognizing a vector user and setting up the proper environment.

The following is a brief description[10] of the changes introduced in support of the Vector Facility.

Operator Commands: The Vector Facility is a semi-independent element. This means the central processor (CP) can execute without the Vector Facility being online (only scalar instructions will work), but the Vector Facility cannot execute without the CP being online. Therefore, the 'CONFIG' and 'DISPLAY' operator commands were updated to provide operational control over and to report on Vector Facility status.

Task Management: As mentioned earlier, the system dynamically recognizes a vector user and sets up the proper environment. No external indication need be provided to identify a task as potentially using the Vector Facility.

When a task initially starts, it is considered to NOT be a Vector Facility user. The task, therefore, is eligible to execute on any of the CPs (unless specifically marked otherwise in the Program Properties Table). Later, if the task issues a vector instruction, a program check (X'19') occurs since the task was initially ineligible to issue vector instructions. The program check interrupt handler determines that the task attempted to use the Vector Facility for the first time and causes the proper environment for vector operation to be established. The task is marked as authorized to execute vector instructions and also assigned vector affinity. Vector affinity means that the task is now eligible to be dispatched only on CPs having a Vector Facility attached.

Conversely, the system also monitors the task's use of the Vector Facility while the task is dispatched. If the task has executed a long time[11] without using the Vector Facility, then the system marks the task as ineligible to execute vector instructions and restores the tasks original affinity. Unless the original affinity

---

9 The MVS/SP 2.1.3 Vector Facility Enhancement (VFE) and the MVS/SP 2.1.3 Availability Enhancement (AE) are mutually exclusive. The MVS/SP 2.1.3 Availability Enhancement provides operating system support for the Extended Recovery Facility (XRF). Both the Vector Facility Enhancement and the Availability Enhancement have been integrated into MVS/SP 2.1.7.

10 Additional information on Vector Facility support can be found in WSC technical bulletin *MVS/SP 2.1.3 Vector Facility Enhancement Overview*, LG66-0227.

11 A long time is currently defined as 50 milliseconds.

specified otherwise, the task is now eligible to execute on all CPs of the processor complex.

Even though the task is now marked ineligible to execute vector instructions, it is still a Vector Facility user. If the task resumes executing vector instructions, a program check (X'19') will again occur, the tasks vector environment will be restored, and task will again be marked as executable only on CPs having a Vector Facility attached. This process of marking the task eligible/ineligible and setting/resetting the affinity can occur many times before the task finally completes.

The capability of dynamically detecting a Vector Facility user and also recognizing that, while the task may still be a Vector Facility user it is not currently using the Vector Facility, reduces system overhead in managing the Vector Facility. For example, the system need only save and restore the vector environment for tasks that use the Vector Facility (not all tasks in the system), and then only if the task continues to use the Vector Facility during successive dispatch periods.

I/O Enablement/Disablement: The design of MVS/XA allows for any CP to handle I/O interrupts. One CP is initially enabled to handle I/O interrupts. If it cannot keep up with the I/O interrupt rate, another CP will be enabled. Conversely, if more CPs than needed are enabled to handle I/O interrupts, the system will disable one of the CPs. Support was added to this process to take into account CPs having a Vector Facility attached. The process works as before but with the following enhancements:

1. When **enabling** a CP for I/O interrupts, the system first tries to find one **without** a Vector Facility.

2. When **disabling** a CP for I/O interrupts, the system first tries to find one **with** a Vector Facility.

SMF Changes: The following SMF record changes were introduced by MVS/SP 2.1.3 VFE:

- An incompatible update was made to SMF record type 6. For a detailed description of the type 6 record format, see *MVS/Extended Architecture System Programming Library: System Management Facilities*, (GC28-1153-4).

- SMF enhancements to provide an installation with support to record new valuable Vector Facility accounting and capacity planning information resulted in additions to the following records:

  - Record type 22 now includes the online status of the Vector Facility.

  - Record type 30 includes Vector Facility usage time and affinity time for the initiator and job/step processing.

  - Record type 70 includes the status of the Vector Facility for Resource Measurement Facilities (RMF).

RMF: RMF reports overall vector affinity time on the CPU report. Vector facility support is provided by RMF V3.3 and V3.4 and subsequent releases. RMF Monitor III was not changed for the Vector Facility The current CPU analysis can be used to indicate when contention for the Vector Facility is causing a significant delay.

Miscellaneous Changes: The following additional changes were also made in support of the Vector Facility:

- Checkpoint/Restart

  No external changes were made to checkpoint/restart. However, before checkpoint/restart can be used for a Vector Facility task, PTF UZ40733 must be applied to DFP 1.2 and PTF UZ40795 must be applied to DFP 2.1.

- IPCS

  IPCS was changed to format the new Vector Facility control blocks. This formatter is also used by Print Dump (AMDPRDMP).

- Print Dump

  Print Dump (AMDPRDMP) now formats the new Vector Facility control blocks. It also uses the IPCS formatter mentioned above.

- Assembler H Version 2

  The assembler now recognizes the new Vector Facility instructions. See the description of APAR PP49757 for the maintenance providing this support.

- AMASPZAP (Superzap)

  AMASPZAP now recognizes the Vector Facility instructions and prints their mnemonics when processing a DUMPT request.

- EREP

  EREP 3.2 Vector Facility support is provided in MVS/XA by PTF UR90065 and in VM by PTF UV90146.

- Stand-alone Dump

  A new level of stand-alone dump was provided. Note that this level (supplied with MVS/2.1.3 VFE) will not support previous levels of the operating system, including MVS/SP 2.1.3 and MVS/SP 2.1.3 AE.

Incompatibilities: The following system incompatibilities with previous releases of MVS/SP V2 (including MVS/SP 2.1.3) were introduced by MVS/SP VFE:

- The Installation Channel Path Table (ICHPT) was moved above 16 megabytes virtual.
- The SYSEVENT routine was moved above 16 megabytes virtual. If the user has been using the SYSEVENT macro interface to call SRM SYSEVENT processing, there will be no perceived incompatibility.

MVS/SP 2.1.7

The major items in MVS/SP 2.1.7 are:

- Support for the IBM 3090 Model 400 processor in Single Image mode
- Reconfiguration support for the IBM 3090 Model 400 processor
- The inclusion of the MVS/SP 2.1.3 Availability Enhancement (AE) and the MVS/SP 2.1.3 Vector Facility Enhancement (VFE)

MVS/SP 2.1.7 adds support for the IBM 3090 Model 400 processor in 370-XA mode. The support includes:

- Changes to the display produced for CONFIG ONLINE/OFFLINE commands to include side information. A side consists of channel paths, processors, expanded storage elements, real storage elements, and vector facilities that can support a single operating system.
- Changes to the display produced for a DISPLAY M command to include side information.
- Changes to CONFIG channel path processing to allow all channel paths on a side to be reconfigured with a single command. This support is also available through the CONFIGxx member of SYS1.PARMLIB.
- Changes in the order of processing of CONFIG commands specified in the CONFIGxx member of SYS1.PARMLIB.
- Vector Facilities go online/offline with the associated Central Processor (CP).
- MVS/SP 2.1.7 supports the reconfiguration of expanded storage elements. The support includes:
  - A new operand - ESTOR(E=id) - on the CONFIG and the DISPLAY M commands; and in the CONFIGxx member of SYS1.PARMLIB.
  - Changes to the display produced for CONFIG ONLINE/OFFLINE commands to include status of expanded storage elements.

MVS/SP 2.1.3 Availability Enhancement (AE) and MVS/SP 2.1.3 Vector Facility Enhancement (VFE) were shipped as mutually exclusive on MVS/SP 2.1.3. In MVS/SP 2.1.7, they are both included which permits an installation to build and maintain one common system. MVS/SP 2.1.7 installs on an MVS/SP 2.1.3 base with SMP or SMP/E without the necessity of a full SYSGEN. If the installation has either MVS/SP 2.1.3 VFE or AE installed, as part of the product install they must run an SMP delete job to back out VFE or AE. This job is provided and documented in the Program Directory that is shipped with the product tape.

The MVS/SP 2.1.7 product tape was initially packaged with all prerequisites, MVS/SP 2.1.0 thru MVS/SP 2.1.3, at an 8509 PUT Service level. However, because of APAR/PTF fixes that had been integrated into the MVS/SP 2.1.7 product, it was recommended that the user be at least at the 8602 PUT level.

Compatibility

The priority sequencing of the dispatching queue has been eliminated in MVS/SP 2.1.7. As a result, the sequence field in the ASCB is now reserved and SMF record type 79 no longer contains the dispatching position value.

The virtual storage requirements below 16 megabytes for MVS/SP 2.1.7, without the XRF capability being exercised, are unchanged from that of MVS/SP 2.1.3.

In MVS/SP 2.1.7, SMF record types 71, 78, 79, and 90 have incompatible changes as follows:

- A new swap placement data section was added to record type 71. The former swap count field remains, but is unused.
- A new subtype for I/O queueing activity on an IBM 3090 processor was added to record types 78 and 79.
- Larger domain numbers are now contained in two new range fields of record type 79.
- Record type 90, subtype 3, was changed to increase the multiprogramming level from 255 to 999.

Note: **The following incompatibilities appeared originally in MVS/SP 2.1.3 VFE and are included here for completeness:**

- The Installation Channel Path Table (ICHPT) was moved above 16 megabytes virtual.
- The SYSEVENT routine was moved above 16 megabytes virtual. If the user has been using the SYSEVENT macro interface to call SRM SYSEVENT processing, there will be no perceived incompatibility.
- An incompatible update to SMF record type 6. For a detailed description of the type 6 record format, see *MVS/Extended Architecture System Programming Library: System Management Facilities*, (GC28-1153-4).
- System Management Facilities (SMF) enhancements to provide an installation with support to record new valuable Vector Facility accounting and capacity planning information resulted in additions to the following records:
  - Record type 22 now includes the online status of the Vector Facility.
  - Record type 30 includes Vector Facility usage time and affinity time for the initiator and for the job/step processing.
  - Record type 70 includes the status of the Vector Facility for Resource Measurement Facilities (RMF).

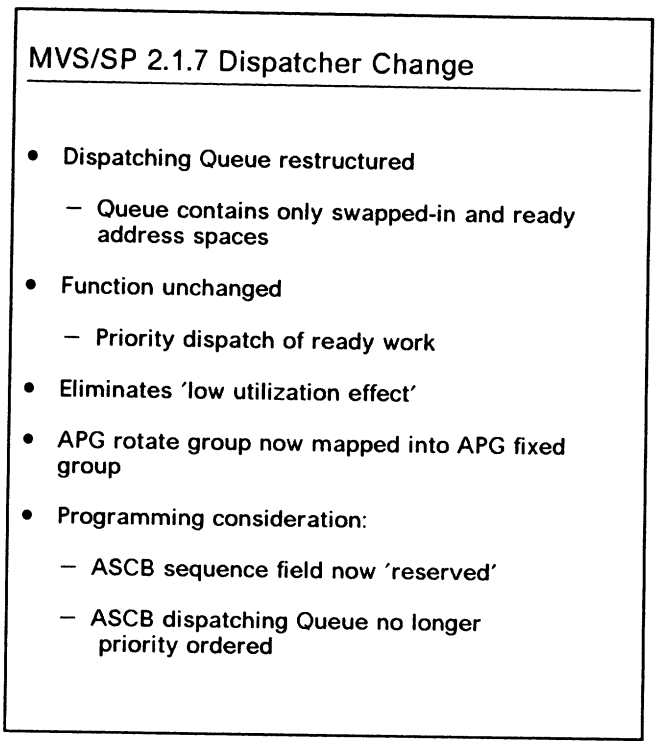

Figure 26. MVS/SP 2.1.7 Dispatcher Change

Dispatcher Change

The dispatching queue has been restructured with MVS/SP 2.1.7. The ASCBs on the dispatching queue (which is now called the swapped-in queue) are no longer in dispatching priority order. However, the function of the dispatcher has remained unchanged. The dispatch order of ready work is still:

- Global SRBs
- ASCBs with work - local SRBs and tasks

The MVS dispatcher has been changed in the MVS/SP 2.1.7 release to accommodate the growth in processing power of large IBM processors. In releases before MVS/SP 2.1.7 the overhead of the dispatcher is proportional to the multiprogramming level and hence grows linearly with the power of the processor. In MVS/SP 2.1.7 the dispatcher has been changed to reduce the overhead. Furthermore, this enhanced dispatcher does not compromise current processor performance. The MVS/SP 2.1.7 level of MVS has been shown to be equivalent in performance to the previous MVS/SP 2.1.3 release in all tested environments, including BATCH, IMS, CICS, and TSO, on IBM 308X and IBM 3090 processors.

Hence, MVS/SP 2.1.7 has completely eliminated the "Low Utilization Effect" in the MVS dispatcher. Briefly, this effect is the overhead expended in scanning for dispatchable work units before entering the wait state. This extra searching is no longer necessary with the enhanced dispatcher. This is the second, and final, step in the removal of the "Low Utilization Effect." The first step was taken in

MVS/SP 2.1.3 when logically swapped address spaces were removed from the search path.

Since the "Low Utilization Effect" has now been eliminated from MVS, the need for the BR15 methodology has been eliminated as well. There is now no significant difference in internal throughput between the MVS/SP 2.1.7 system running the BR15 jobs and the MVS/SP 2.1.7 system not running the jobs. This should make the capacity planning decisions of the customer measurably easier since the measured CPU utilization via RMF or SAD is now a more accurate indicator of the working utilization.

Since the system dispatches address spaces of equal priority in the order which they become ready, the rotate priority group (IPS - Rx) no longer has significance. When the system encounters an Rx specification in IEAIPSxx, it accepts it as Fx, the first fixed priority. MVS/XA 2.1.7 removes the rotate algorithm by using the rotate priority position (A) as the first fixed priority. Therefore, there are now six fixed priorities (A - F) specified as Fx, Fx0, Fx1, Fx2, Fx3, Fx4, respectively.

The ASCB sequence field (ASCBSEQN) has been changed to a reserved field, but set to zero. As a result, RMF 3.4.1 places a zero in the dispatching position field (R791SEQN) of SMF type 79 subtype 1 record.

Supporting Products

MVS/XA DFP

MVS/XA Data Facility Product 1.2, including its prerequisites, is a co-requisite for MVS/SP 2.1.7. MVS/XA DFP 2.1 can also be used as a co-requisite for MVS/SP 2.1.7. If the Availability Enhancement (AE) is to be used, MVS/XA DFP 2.1 or later must be used as the co-requisite.

JES2

The JES2 component of MVS/SP-JES2 2.1.7 is functionally equivalent to the JES2 component distributed with MVS/SP-JES2 2.1.5 (the 31-bit JES2). Only the BCP component has been changed for MVS/SP-JES2 2.1.7. The JES2 components that are supported with the MVS/SP 2.1.7 BCP are:

- JES2 1.3.3 - Support terminated 6/30/87
- JES2 1.3.4 - Support terminates 12/31/87
- JES2 1.3.6
- JES2 2.1.5
- JES2 2.2.0

JES3

The JES3 component of MVS/SP-JES3 2.1.7 is functionally equivalent to the JES3 component distributed with MVS/SP-JES3 2.1.5 (the 31-bit JES3). Only the BCP component has been changed for MVS/SP-JES3 2.1.7. The JES3 components that are supported with the MVS/SP 2.1.7 BCP are:

- JES3 1.3.1 - Support terminated 6/30/87
- JES3 1.3.4
- JES3 2.1.5

IOCP

PTFs for the MVS and VM versions of IOCP are required to support the IBM 3090 Model 400. PTF UZ90453 (APAR OZ91644) applies to IOCP for MVS/SP 2.1.7, MVS/SP 2.1.3, and MVS/SP 1.3.5. For VM/SP HPO the APAR is VM26002, and for VM/XA Systems Facility IOCP the APAR is VM26003.

EREP

IBM Environmental Record Editing and Printing Program (EREP) support for the IBM 3090 Model 400 is provided in both Release 3.1 with Feature 3 and Release 3.2. IBM 3090 Vector Facility support in MVS EREP 3.2 is provided via PTF UR90065, and in VM EREP via PTF UV90146.

RMF 3.4.1

Resource Measurement Facility (RMF) Version 3 Release 4.1 is a supplemental package to the RMF Version 3 Release 4 base product; therefore, RMF Version 3 Release 4 must be installed before installing RMF 3.4.1.

RMF Version 3 Release 4.1 supports MVS/SP-JES2 2.1.7 (5740-XC6) and MVS/SP-JES3 2.1.7 (5665-291), and introduces no new function to Version 3 Release 4. RMF releases prior to Version 3 Release 4.1 will not run correctly at the MVS/SP 2.1.7 level.

Changed Monitor II Reports: The address space state data (ASD) report and the address space data by jobname (ASDJ) report no longer show the position of address spaces on the dispatching queue. The DP PO field is removed from these reports.

Compatibility: RMF Version 3 Release 4.1 is upwardly compatible with RMF Version 3 Release 3 and Version 3 Release 4 and introduces no new function or interfaces. RMF Version 3 Release 4.1 is the level of RMF required to function with MVS/SP-JES2 2.1.7 or MVS/SP-JES3 2.1.7. SMF record type 79 no longer contains a dispatching position value. However, field R791SEQN of the type 79 subtype 1 record now contains zeros since its source field, ASCBSEQN, contains zeros.

Before Checkpoint/Restart can be used for a task using the Vector Facility, PTF UZ40733 must be applied to MVS/XA DFP 1.2, and PTF UZ40795 to MVS/XA DFP 2.1.

MVS/SP 2.1.7 users who have the Resource Access Control Facility (RACF) licensed program (5740-XXH) installed must be at RACF Release 1.5 or higher.

MVS/SP 2.1.7 users of both RACF and Hierarchical Storage Manager (DFHSM) (5740-XRB) must install PTF UZ67673 for DFHSM, which provides the "always call" interface with DFHSM.

MVS/SP 2.2.0

MVS/SP 2.2.0 contains new function in the base control program (BCP) and in the JES2 component. The JES3 component, which has not been updated in this release, remains functionally equivalent to the JES3 component of MVS/SP 2.1.5.

In addition to all IBM 3090 support previously provided in prior MVS/XA releases, the updates to the base control program (BCP) in MVS/SP 2.2.0 include:

- **Defining I/O Configurations** - allows installations to specify customization data, such as I/O configurations, console definitions, PFK definitions, PPT definitions, and user SVCs without having to (1) modify system code or (2) go through the SYSGEN process.

- **JES3/MCS Coexistence** - improves usability by providing a common interface for MVS, JES2, and JES3 console commands and messages. (JES3 support for this function is provided in MVS/SP-JES3 2.2.1.)

- **Constraint Relief** as a result of:

  - Optionally moving several SWA control blocks to virtual storage above 16 megabytes. (JES3 support for this function is provided in MVS/SP-JES3 2.2.1.)

  - Increasing the maximum allowable number of DD statements for each job step from 1635 to 3273.

  - Allowing a data set expiration date beyond the year 1999.

  - Paging enhancements in the auxiliary storage manager (ASM).

- **Usability Enhancements** in:

  - System resource manager (SRM) that provide better control and tuning capabilities.

  - System management facility (SMF) that allow the installation to selectively include/exclude subtypes within record types.

- **IPCS Enhancements** - include several RAS and usability features, such as new dump analysis subcommands and exit facilities. (JES3 support for this function is provided in MVS/SP-JES3 2.2.1.

- **IOS Formatter** - includes new operands for dump diagnosis.

- **Data in Virtual** - provides a new interface for application processing of data.

- **Symptom Recording Service** - a service provided to record symptomatic and environmental data about software errors or exceptional conditions.

- **RAS Enhancements** in:
  - Reusable ASIDs - reduces the number of scheduled IPLs in installations executing applications that use cross memory services.
  - Other RAS enhancements, such as enhanced SVC dumps.
- **Processing Improvements** in the following:
  - Scheduler JCL Facility (SJF) - improves the search used for information in JCL descriptor tables (JDTs).
  - Global Resource Serialization - improves communication between systems.
- **Previously-provided function** that has already been made available as a PTF, such as the enhanced RSM trace.

VIO to Expanded Storage

A new capability has been added to MVS/XA 2.2.0[12] which will allow some customers to further exploit the availability of large amounts of expanded storage. This support will be available as PTFs in May, 1988.

The benefit of VIO to expanded storage is primarily a reduction in the elapsed time required to run a job or an interactive transaction. The improvement in application run time will be a function of the I/Os that are requested to or from DASD (either VIO paging or nonVIO) and due to application design, are not overlapped with processor execution potential. For many applications, the unoverlapped I/O time is very significant component of their total processing time.

Compatibility with Current VIO Support: If an application is using VIO today no change is necessary to take advantage of the new support. Applications not using VIO would need to specify an appropriate VIO unitname in the JCL specification.

Expanded storage **will not be used** for batch jobs using Checkpoint/Restart (VIO journaling). Such jobs will continue to run without change and will have their VIO activity directed to the paging subsystem as before.

Use of Expanded Storage: The design insures that expanded storage will only be used for VIO pages when no significant increase in demand or swap paging would result. This is controlled by the specification of a high Criteria Age value for VIO pages. The system supplied default is 900 seconds which is much higher then the default values for any other category of pages and will therefore prevent VIO pages being directed to expanded storage long before all other categories of pages. The user may change this default by modifying the IEAOPT member of PARMLIB.

---

12 Not available for prior levels of MVS/XA.

In some cases it may be desirable to insure that a specific application be allowed to retain it's VIO pages in expanded storage. The existing Storage Isolation capability can be used to define a working set large enough to include the VIO pages.

If it becomes necessary to migrate VIO pages in expanded storage to DASD, those paging datasets specified as 'eligible for VIO' will be used. Auxiliary storage 'slots' will be allocated contiguously to minimize seek activity.

Considerations: When there is not sufficient expanded storage, the VIO activity will go to the paging subsystem. This implies that it should be configured to handle the load that would result. Note that this support is not selective. If VIO is allowed on a system, any application may choose to use it by a simple JCL specification. If the installation finds it necessary to restrict the VIO to certain applications or control space allocation, it will be necessary to implement user exits to exercise that control.

RMF Reporting: RMF will provide support for VIO to expanded storage in Monitor I and Monitor II. The Monitor I paging activity report will include the rate of VIO pages to expanded storage in the total rate of transfers of pages to expanded storage. Any VIO pages migrated from expanded storage will be included in the migration rate, also reported in the paging activity report. Monitor I will record the number of VIO pages moved to expanded storage, the number moved from expanded storage to central storage to satisfy a read request for a VIO dataset, and the number of VIO pages migrated from expanded storage. These numbers will be recorded in separate fields in the type 71 SMF record provide by RMF as the source of the paging activity report.

Monitor II reports the rate of transfer of pages to expanded storage and the migration rate from expanded storage in the session paging activity report. The report will include VIO pages moved to expanded storage and the rate of VIO pages migrated in the existing migration rates.

Installation Planning: Installation information will be provided with the PUT tape. At that time, APARs OY08834, OY08835, OY08836, and OY08837 for MVS/SP 2.2.0 and OY09186 for RMF 3.5 will be available through INFO/ACCESS and normal IBM Support Center procedures.

MVS/370 Support

Overview

The following product versions and/or features are the minimum level necessary to run MVS/370 on the 3090 Processor Complex. The reader should refer to Figure 4 in Chapter 1, "IBM 3090 Family Comparisons" to determine which models of the 3090 Processor Complexes support MVS/370 operation.

- MVS/SP 1.3.5
- EREP 3.1 Feature Number 3 (FMID = FER3130)
- DFP[13] 1.x PTF (UZ72306) for SYSGEN toleration of a new IOCP keyword on the IODEVICE macro
- IOCP Program - IOPIOCP V1L0 (FMID = HIO1104)

MVS/SP 1.3.5 has as a prerequisite MVS/SP 1.3.3. The installation of MVS/SP 1.3.5 is similar to that of MVS/SP 1.3.3.

DFP V1.x is the same level as that required for MVS/SP 1.3.x. There are no changes to executable DFP code to support the 3090 processor. The PTF is needed to support a new IOCP keyword on the SYSGEN IODEVICE macro.

EREP V3.1 is a prerequisite for an EREP feature (Feature 3) containing formatter exits for 3090 model dependent information.

The IOCP program for the 3090 processor is a new IOCP program. While its base was the existing 308X IOCP program, this new program is only applicable to the 3090 processor. Since it was derived from the existing 308X program, the macros and most of the keywords/values are the same. While some changes will have to be made to an existing IOCP input stream to support the 3090 processor, these changes will be minor.

The processor support in MVS/SP 1.3.5 occurred in several functional areas. Base SCP support is that needed to allow the MVS/SP operating system to run on the 3090 processor. SRM added the model number and version code for the 3090 to its CPU adjustment table. SRM uses this table in its resource management algorithms that are affected by processor speed. The IOS Channel Check Handler was updated to support the logging of Interface Control Checks and Channel Control Checks in the fixed logout area of the 3090 processor.

---

[13] DFP/370, program number 5665-295, is not a prerequisite for MVS/SP 1.3.5 operation on a 3090 processor. User's who have Data Facility/Extended Function, program number 5740-XYQ, installed can continue to use DF/EF, in conjunction with MVS/SP 1.3.5, on the 3090 processor. User's should be aware that support provided by PTF UZ72306 for DFP/370 is not available for DF/EF.

Frame Deallocation

Frame deallocation supports a hardware RAS enhancement. When the hardware detects certain double bit storage errors, it corrects the error and presents a machine check. The Machine Check Handler will call the Real Storage Manager to attempt to free the storage. If the storage frame backs a valid pageable page, RSM will obtain a new frame into which the data will be moved. If the storage frame does not back a valid pageable page, frame deallocation will not be scheduled. The frame will instead, be marked as a bad frame. The frame will remain allocated to the page, and when the page is freed, it will be marked offline.

Frame deallocation will cause frames with correctable errors to be marked offline so they will not be reallocated. Previously, frames with data errors could only be taken offline when they were backing unchanged pageable pages, or later when being freed (e.g.. when the pages were FREEMAINed at task termination).

Expanded Storage

MVS/SP VERSION 1 DOES NOT SUPPORT EXPANDED STORAGE.

Vector Facility

MVS/SP VERSION 1 DOES NOT SUPPORT THE VECTOR FACILITY.

VM Support

VM/XA Systems Facility

VM support for the IBM 3090 processors in S/370-XA mode is provided by the Virtual Machine/Extended Architecture Systems Facility (VM/XA SF). Two releases of VM/XA SF are currently offered, Release 1 and Release 2. VM/XA SF R1 supports the uniprocessor models, the dyadic models, and the model 400/400E operating in partitioned mode. VM/XA SF R1 does not support the Model 300E, the Model 400/400E in single image mode, or the Model 600E in either partitioned or single image mode.

VM/XA SF R2 supports all models of the 3090 processors and also incorporates single image mode support for the models 300E, 400/400E and 600E, as well as continuing support of partitioned systems (400/400E and 600E).

VM/XA SF support of the 3090 processors occurs in four general areas:

- Basic processor support
- Expanded storage support
- Virtual extended channel support
- Vector Facility support

Basic Processor Support

Basic Processor Support consists of the following items:

- Processor Recognition. VM/XA SF will recognize the CPU identification code of the 3090 processors and tailor the scheduling parameters to the unique performance characteristics of the processor.
- Real Processor Controller Support. VM/XA SF will use the Service Call instruction to obtain certain services from the processor controller.
- Virtual Processor Controller Support. Virtual processor controller support has been extended to allow for a guest Service Call instruction requesting the READ or WRITE DATE function for reading or writing the I/O Configuration Data Sets (IOCDS).
- Machine Check Handler. Support was added to handle the new machine check interruption codes of the 3090 processors.

Expanded Storage Support

VM/XA Systems Facility Release 1 does not use expanded storage for its own use, but allows it to be dedicated to a V=V or V=R guest.

VM/XA Systems Facility Release 2 allows the dedication of expanded storage to a single guest, the dedicated use of expanded storage by VM/XA SF R2 for paging or the partitioning of expanded storage between guests and VM/XA SF R2. This enhanced use of expanded storage, along with a block paging facility is available now as corrective service (VM28091) and as normal maintenance on PUT tape 8705.

Expanded storage can be dedicated to a guest via a new directory control statement 'XSTORE', or the new 'ATTACH XSTORE' command. Once the expanded storage has been dedicated to a guest, VM/XA SF will allow that guest to transfer pages to and from expanded storage. The expanded storage can be taken away from a guest via the 'DETACH XSTORE' command.

Virtual Extended Channel Support

Unlike VM/SP HPO, the VM/XA Systems Facility will allow a S/370 mode guest to use virtual channels from X'00' to X'1F'. Since VM/SP HPO is the only S/370 mode operating system that supports 32 channel addresses, this facility allows for a VM/SP HPO system, generated for 32 channels to be run in a virtual machine.

Vector Facility Support

VM/XA SF Release 2 supports the use of 1 to 6 Vector Facilities. Each Vector Facility may be dedicated to the V=R user (when its associated processor is dedicated), or shared among V=V guests. Guests which may use the Vector Facility include MVS/XA, VM/SP HPO running CMS, and CMS (with and without EDAC[14]) running first level on VM/XA SF. VM/XA SF R2 will also support the use of the Vector Facility by guest VM/XA SF systems.

Installations planning on installing VM/XA SF Release 2 are reminded that System Engineering Change (SEC) 223620 is mandatory when the Vector Facility is installed. For additional information on the contents of SEC 223620, see "SEC 223620" on page 97.

---

[14] IBM Program Offering *CMS/Extended Data Array Capability, 5798-DWA.*

VM/XA System Product

VM/XA System Product (VM/XA SP) supersedes all releases of VM/XA Systems Facility (VM/XA SF) and provides a migration path for VM/XA Systems Facility users, as well as VM/SP and VM/SP HPO users. All VM/XA Systems Facility functions are supported including single-image N-way processing, vector processors, partitionable expanded storage, and block paging.

Current offerings include VM/XA System Product Release 1, VM/XA System Product Release 1 Enhancement for Multiple Preferred Guests and VM/XA System Product Release 2.

VM/XA System Product Release 1

VM/XA System Product Release 1, in addition to including all the functions of VM/XA Systems Facility, also provides a bimodal CMS. This CMS has been enhanced to take advantage of the capabilities of 370-XA architecture including 31-bit addressability for both programs and data. It will operate in either System/370 mode with 24-bit addressing or 370-XA mode with 24 or 31-bit addressing capability.

Programs running on CMS will now be able to take advantage of 370-XA virtual machine sizes that are greater than 16 megabytes (up to the current 999 megabyte CP restriction) through the following enhancements:

- A new program linkage mechanism, to replace the existing SVCs 201 and 202, has been provided via macros (CMSCALL and CMSRET). These macros allow programs to call other programs or CMS services anywhere within the addressing range of the CMS virtual machine. Programs using SVCs 201, 202 and 203 will continue to function below the 16 megabyte line.

- Program handling functions have been extended to support 31-bit addressing.

- Program load and module generation facilities have been enhanced to support module relocation at load time.

- The CMS storage management facility has been restructured to support and take advantage of 31-bit addressing. Management of free storage is improved by the use of a subpooling scheme which allows storage related to a common subpool to be manipulated as a single entity. New macros (CMSSTOR and SUBPOOL) have been provided to assist assembler applications in gaining access to these facilities.

- The CMS programming interfaces have been enhanced to support programs which require more than 16 megabytes of storage for execution and data.

The enhancements to the CMS programming interfaces enable an application to be coded so that it can be independent of whether it is running in a S/370 virtual machine or a 370-XA virtual machine. In addition, programs running in a 370-XA virtual machine can choose to execute with either 24-bit addressing, 31-bit addressing, or dynamically switch between the two.

Additional Extensions Include:

- Enhanced support for Discontiguous Saved Segments (DCSS):

  - CMS has been updated to support the management of shared segments. This support enables shared segments to be safely loaded within a user's virtual machine. This support has removed the restriction that a DCSS must reside at an address greater than the size of the user's virtual machine.

- Enhancements to the MVS development, test and execution environment provided by CMS:

  - The OS macros shipped with CMS have been replaced with the equivalent versions provided in MVS/SP release 2.2.0 and DFP 2.3.0.

  - The OS simulation capabilities of CMS have been upgraded to support a selected set of MVS/XA macros and to support 370-XA architecture when executing in 370-XA mode. The AMODE and RMODE operands on the simulated macros are recognized in both S/370 and 370-XA virtual machines.

    - The OS GETMAIN and FREEMAIN macro functions have been enhanced to utilize the new CMS storage management facilities.

    - The OS Linkage Editor function on VM has been replaced with the MVS/XA Linkage Editor contained in DFP 2.3.0.

  - Some previous functional restrictions for CMS on VM/XA SF have been removed. These include full support for the CMS productivity aids, RDRLIST, PEEK, NOTE and SENDFILE commands. In addition CMS/DOS and CMS support for VSAM are now supported in S/370 mode virtual machines.

VM/XA SP R1 provides a number of additional enhancements, including the following:

- Performance Monitoring Facility

  A native VM monitoring facility, including significant enhancements over prior VM/370 monitors, is provided in the VM/XA System Product environment. The monitor facility provides the necessary data which will enable the user to analyze the utilization of, and contention, for the major system resources such as processors, storage, I/O devices and the paging subsystem.

- VM/XA Real Time Monitor Support

  Support as provided by the VM/XA Real Time Monitor/Systems Facility (RTM/SF) on VM/XA Systems Facility Release 2 has been upgraded to support VM/XA SP R1.

  RTM/SF is a system monitoring aid that can be used to help detect and diagnose problems, analyze system performance, and provide operators with

awareness of machine operations. RTM/SF is command, timer, and table driven, and presents information in realtime on any VM/XA SP supported terminal or display.

- Logged-on User Limit Enhancements

  Function is provided to allow installations to set the maximum number of users allowed to log onto the system. The limit can be adjusted to fit the needs of the installation, or the limit can be removed entirely. The CP command SET MAXUsers can be used to adjust or remove the user logon limit. The CP command Query MAXUsers can be used to display the current limit. The SYSMAXU macro can be used at SYSGEN time to adjust or remove the limit.

- Resource Access Control Facility (RACF) Support

  The data security functions as provided in VM/SP (with or without VM/SP HPO) are included in VM/XA SP R1 to provide compatibility for those users who have implemented RACF data security procedures on their current VM/SP systems.

- Additional Device Support

  VM/XA SP R1 includes support for additional devices. A complete list of supported devices will be included in the *VM/XA System Product General Information Manual (GC23-0362)*, which became available 3Q87.

- User Class Enhancements

  VM/XA SP R1 provides for up to 32 unique user privilege classes. This capability is functionally equivalent with VM/SP (and HPO) Releases 4 and 5 and provides improved flexibility and control of user access to command and DIAGNOSE facilities. Installations may dynamically alter the privilege class structure or change the privilege class of any command.

- SPOOL File Limit Relief

  Installations may now support a maximum of 9,999 SPOOL files per user, subject to limitations of warmstart and storage space. Prior to this support installations on VM/XA SF were limited to 9,999 SPOOL files per system. This constraint relief is similar to that provided in VM/SP HPO Release 5.

In general, VM/XA SP R1 is upward compatible from VM/XA SF R2. Programs written in higher level languages are source and object level compatible on CMS (i.e., CMS supports the same languages for both systems). Programs using 24-bit virtual storage addresses continue to execute in virtual storage below 16 megabytes without change Application programs that currently execute on VM/XA SF should execute without change unless dependent on those functions listed below.

- EXCEPTIONS - Programs or EXECs that depend on one or more of the following will need to be examined for compatibility.
  - INTERNAL
    - CP or CMS control block structure
  - EXTERNAL
    - Applications that use undocumented interfaces
    - Programs that depend on the syntax or output format of CP commands. The syntax of most command invocations remains upward compatible. The format of many command responses has changed, particularly for Query and SPOOLing commands.
    - Programs that operate on SPOOL files owned by another user must now refer to SPOOL files by USERID and SPOOLID. CP SPOOLing commands are enhanced for this support.
    - DIAGNOSE CODES - All diagnose codes available in VM/XA SF are supported in VM/XA SP R1.
    - MESSAGES - There are a small percentage of messages where the text has been improved for clarity and consistency, which should have minor effects on the users. Essentially, message text will remain the same for approximately 95 of the 370-XA messages. Approximately half of the 370-XA messages have changed message numbers. Return codes have also experienced extensive changes. Additionally, VM/XA SP has many 4-digit message numbers, therefore some command invocations will now result in 4-digit return codes. Applications that depend on 10-character message header lengths and/or 3-digit CP command return codes should be examined. Systems using programmed operator facilities must examine their routing tables for any changes necessary.

An additional compatibility objective of VM/SP R1 was to provide a growth path for VM/SP HPO users with maximum compatibility to ease migrations. However, due to factors such as architectural differences, full compatibility is not possible. VM/XA SP Release 1 has made every effort to fulfill the migration objective by keeping, where possible, command syntax, responses and message text the same. Common DIAGNOSEs are functionally equivalent.

General availability of VM/XA System Product R1 is planned for March 1988.

VM/XA System Product R1 Enhancement

VM/XA System Product Release 1 Enhancement for Multiple Preferred Guests provides the capability to run multiple high-performance guests (up to 4) running concurrently with other virtual machines. In addition to the Virtual = Real (V = R) preferred guest virtual machine, three additional Virtual = Fixed (V = F) preferred guests are now supported along with multiple concurrent CMS users and test or maintenance virtual machines. All guest operating systems which are

supported by VM/XA as V=R preferred virtual machines are supported as V=F guests. Preferred guest recovery is limited to the V=R guest. This enhancement for multiple preferred guests is only supported on IBM 3090 Processor Complex Enhanced models that have the Multiple High Performance Guests Support feature installed.

For example, MVS/370 and MVS/XA production systems could run on the same processor complex, along with VM/SP or VM/SP HPO guests running System/370 CMS-intensive workloads. Multiple MVS/XA systems could run on the same processor complex to provide application isolation or to alleviate MVS subsystem virtual storage constraints, and at the same time use the new 370-XA capable CMS. Multiple VM/SP or VM/SP HPO systems could run on one processor complex to migrate CMS applications to use the 370-XA architecture.

The VM/XA System Product Release 1 Enhancement for Multiple Preferred Guests can be beneficial in the following scenarios.

- **Migrations**: Users migrating operating systems or applications can run the old and new systems or applications on the same processor complex simultaneously, as opposed to requiring separate processor complexes, while achieving performance slightly lower than that experienced natively.

- **Consolidation**: Users with the requirement to combine multiple production system images on one processor complex can consolidate without the complexity involved in merging them into one operating system image.

- **Production and Test**: Users can run multiple production systems simultaneously with test or development systems on the same processor complex, with high performance for the production systems.

- **Multiple System Images**: Users with the requirement to run multiple production system images on one processor complex for application isolation or virtual storage constraint relief, can do so with acceptable system performance without incurring duplicate hardware or software costs. For example, this approach could be used for IMS users or for large VM/SP HPO establishments with virtual storage constraints. Also, multiple XRF systems could be run on one processor complex for test purposes.

VM/XA SP R1 Enhancement also provides enhanced control and use of processor resources. Management of central storage and expanded storage is enhanced to extend the performance characteristics of the V=R guest to the V=F guests. The use of the Start Interpretive Execution (SIE) emulation facility is expanded to use the SIE Assist for those devices dedicated to a V=F guest. Allocation of processor cycles or the IBM 3090 Vector Facility continues to be flexible and dynamic, with the added capability to dedicate processors with or without the IBM 3090 Vector Facility to any preferred or non-preferred guest.

- Central and expanded storage

  When the VM/XA SP system is generated, the total amount of central storage to be assigned to the V=R and V=F guests is defined. Specification of central storage for each individual guest is made at LOGON of the guest. This specification may be changed, and the guest re-IPLed, if sufficient contiguous

central storage is available. The virtual storage of a V = F guest is mapped into contiguous central storage, and remains allocated and reserved until the guest logs off.

The allocation of expanded storage is made more flexible by this support. Portions of expanded storage may be dedicated to each V = R, V = F or V = V guest, while the remainder can be used by the VM/XA control program to support paging for other users such as CMS.

- SIE support

  The use of the SIE assist is expanded to support those devices dedicated to a V = F guest. Also, the support of the Virtual Machine Assist (VMA) under SIE is available to the V = F guest to improve the performance of VM/SP and VM/SP HPO guests.

- Processor allocation

  VM/XA SP's event driven, priority scheduler, can be used to allocate resources such as processor cycles based on a specified "share." The specification can be changed dynamically, to allow tuning of the total system as requirements change, with no disruption to the guests. Also, specified processors may be defined as "dedicated" to any preferred or non-preferred virtual machines for their sole use, or changed to "undedicated" to enable sharing of the processors by multiple preferred and/or V = V guests. Similarly, the IBM 3090 Vector Facility is dedicated to any guest if the processor on which it is installed is dedicated to the guest, or the IBM 3090 Vector Facility can be shared among multiple V = R, V = F or V = V guests if it is installed on a processor which is shared among those guests.

The VM/XA System Product Release 1 Enhancement for Multiple Preferred Guests is supported on the IBM 3090 Processor Complex Enhanced Models with the Multiple High Performance Guests Support hardware feature installed. Limited availability on selected 3090 Enhanced Model processors and guest SCPs is planned for March '88. General availability is planned for 3Q88.

VM/XA System Product Release 2

VM/XA System Product Release 2 is a replacement for VM/XA Systems Facility and VM/XA System Product Release 1. It contains all the functions and features of VM/XA System Product Release 1. In addition, it will contain the VM/SP Group Control System (GCS) which supports the installation and operation of VTAM and NetView Release 1. The inclusion of native SNA support provides growth options for current VM/SP HPO users with SNA networks.

The inclusion of SNA support allows VM/XA System Product R2 to participate in and control SNA networks without requiring a guest such as VM/SP HPO or VCNA to handle the SNA functions. This support is similar to that provided in VM/SP and VM/SP HPO R4 in that the VTAM and NetView programs operate in System/370 mode only.

General availability of VM/XA System Product Release 2 is planned for 1Q89.

VM/SP High Performance Option

VM support for the IBM 3090 processors in S/370 mode is provided by Virtual Machine/System Product High Performance Option (VM/SP HPO). Three releases of VM/SP HPO containing 3090 processor support are currently offered; HPO 3.6, HPO 4.2, and HPO 5.0. The reader should refer to Figure 4 on page 14 in the "Comparison Summaries" section to determine which HPO release supports which 3090 models and modes of operation. Prior to finalizing any plans, this support should be verified with the local IBM representative.

There are four major components to the HPO support for the 3090 processors.

- Basic Processor Support
- Real Extended Channel Support
- Real Expanded Storage Support
- Vector Facility Support (HPO 4.2 and HPO 5.0)

Basic Processor Support

Basic Processor Support provides the necessary changes to the Control Program to enable VM/SP HPO to operate on a 3090 processor.

Recovery processing supports the new Machine Check interruption codes of the 3090 processors.

The Control Program has been changed to use the Service Call instruction, where appropriate, instead of the Diagnose instruction, to request services from the service processor. This is done for the VARY ON/OFF PROCESSOR command, and for reading and writing the IOCDS. The READ SCP INFORMATION service call will be simulated by the control program when it is issued from a virtual machine.

The 3090 processors have both Preferred Machine Assist and Control Switch Assist, and both features are supported by VM/SP HPO. MVS guests running with Preferred Machine Assist must be at a level that supports the 3090 processors.

Extended Channel Support

Extended Channel Support provides the capability of using more than 32 channels on a 3090 Model 200/200E or 3090 Model 400/400E in partitioned mode. It does this by supporting up to 32 channels per channel set. This requires that real device addresses in VM/SP HPO be extended from 3 to 4 digits, where the channel portion of the address can now range from X'00' to X'1F'. This support is for real channels only. Virtual machines are still restricted to having 16 channels, X'0' to X'F'.

The implementation of Extended Channel support affects every command, command response, system message, sysgen macro, and utility control statement that contains a real device address. Every command, macro or control statement that accepted a 1 to 3 digit real device address will now accept a 1 to 4 digit real device address. Every message or command response that contained a 3 digit real device address will now contain a 4 digit real device address. Users who have

EXECs or programs that issue CP commands and examine the responses should review those EXECs and programs to determine the extent to which they must be modified to handle the 4 digit device addresses. The necessary fields in the Monitor records have also been expanded to accommodate the 4 digit addresses. Users with their own programs to read Monitor files should also examine those programs to determine the extent to which they must be modified.

There are additional considerations concerning Extended Channel Support:

- All real channels above channel X'F' must be included in DMKRIO in order to be used. They may not be preferred channels to an MVS guest running with Preferred Machine Assist (PMA).

- Since we cannot have virtual channel addresses greater than X'F', and the CP ATTACH command that will attach a range of devices to a virtual machine sets the virtual address equal to the real address, this form of the ATTACH command cannot be used for devices on channels X'10' to X'1F'. Devices on these channels must be attached to virtual machines, one at a time.

- In order to allow an HPO 3.6 or 4.2 system to run in a virtual machine, or on a processor that does not have extended channels, it is recommended that you include in DMKRIO enough channels with addresses from X'0' to X'F' to support the necessary SYSOWN devices to allow the system to function without extended channels.

- When configuring the channels on a 3090 Model 200 processor, you should keep in mind that several stand-alone programs do not support extended channels, and therefore can only use devices on channels X'0'-X'F' when running in stand-alone mode. These programs are:

  - HPO Starter System
  - MVS Stand-alone Dump
  - Stand-alone DSF

Expanded Storage Support

Expanded storage on 3090 processors is storage that can be used by VM/SP HPO as a high speed area for preferred paging, swapping or both. Expanded storage cannot be dedicated to a guest. Internally, VM/SP HPO maps this storage with the same structure of control blocks that were used in previous releases to map paging and swapping space on DASD devices (ALOCBLOKs and RECBLOKs). Installations control how much expanded storage is to be used for paging and swapping via the SYSPAG or SYSXSTOR macros in module DMKSYS. This is done when the system is generated. On a 3090 Model 400/400E, expanded storage must be configured with exactly one-half of the total expanded storage on each side.

System Generation of Expanded Storage

Expanded storage can be configured into a system via either of the following DMKSYS macros:

- SYSPAG
- SYSXSTOR

The primary differences in these two macros are:

1. The SYSPAG macro may be used for all system areas (SW, PP, PG, PM, PS and DU) and all devices supported by those areas. The SYSXSTOR macro may be used to specify only expanded storage for the system areas supporting expanded storage (SW and PP).

2. With the SYSPAG macro, the amount of expanded storage to be configured is specified as a number of 4-megabyte increments. Internally in CP control blocks, the increments are numbered starting at 0. All messages relating to expanded storage increments display increment numbers starting at 1.

   With the SYSXSTOR macro, the amount of expanded storage is specified in megabytes. By eliminating specification of particular physical increments of expanded storage, the SYSXSTOR macro makes it possible to run a common configured system on either side of a 3090 Model 400/400E.

The SYSPAG macro was introduced in HPO 3.4, but the SYSXSTOR macro was introduced later, as APAR VM26294 to HPO Releases 3.6 and 4.2. Interim documentation was provided in *VM/SP HPO Enhancements to Paging Storage*, GC23-0382-00. Both SYSPAG and SYSXSTOR are included in the base of HPO 5.0. Syntax and coding considerations can be found in the *VM/SP HPO Planning Guide and Reference.*

Other Changes for Expanded Storage

Since the proper use of expanded storage can have a significant effect on system performance, many system performance indicators have been modified to provide information on the use of expanded storage. Monitor records have been expanded to include counts of expanded storage paging and swapping. Both VMMAP and VM/RTM have added fields and reports. In addition, The QUERY PSTOR command has been added to provide information on the amount of expanded storage in use.

For system debugging purposes, it may be necessary to examine the contents of expanded storage. For this reason, DDR has been enhanced to display or print the contents of expanded storage. To use DDR to display the contents of expanded storage you simply specify the input device as PSTOR.

Vector Facility Support

HPO 4.2 and HPO 5.0 provide support for the Vector Facility Feature on 3090 processors. Support is provided in the following areas:

- Access to the vector facility by multiple CMS users. This support provides for the saving and restoring of the real vector registers so that the real vector facility can be shared among many CMS users. In order to minimize the saving and restoring of the vector registers, this is only done when the user of the vector facility changes. If a vector user is re-dispatched and his registers are still loaded in the vector facility, no saving or restoring will be done. In addition, vector register save areas are obtained from above the 16MB line. For HPO 4.2, the vector support in CP is provided in the base level. The CMS portion of the vector support is provided on PUT 8604. For HPO 5.0, both the CP and CMS support is provided in the base level.

- Operator Control. New commands are available to vary the vector facility online and offline, and query commands have been added so that the operator may identify vector facility users.

- Vector Facility Usage Data. Usage of the vector facility is recorded in the monitor data, accounting data, and in the response to the INDICATE command.

Chapter 4. Installation Topics

IOCP Support

The IOCP process for 3090 processors is the same as that used on the 308X processors. The primary purpose of the process is to convert I/O configuration data into a form that is usable by the channel subsystem. The process involves reading configuration data (channel path, control unit, and device) supplied by the user, and constructing tables and control blocks that are written to the IOCDS datasets located on the 3370 DASD attached to the 3092 processor controller. The I/O configuration definition is loaded into the processor Hardware System Area (HSA) during the power-on-reset process. Additionally, the IOCP process also produces hardcopy reports of the generated configuration.

The IOCP process for 3090 processors supports three basic operating environments: MVS (both MVS/SP V1 and V2), VM (VM/SP HPO and VM/XA Systems Facility), and stand-alone. The software versions are shipped with the corresponding level of software to support the 3090 processors, and the stand-alone version is shipped with the hardware.

The IOCP program for the 3090 processor is a new IOCP program. The program name is IOPIOCP[15] and the Function Modification IDentifier (FMID) is HIO1104. The program can be run on a non-3090 processor to facilitate checking out the IOCP input stream. A subsequent update to IOCP incorporated Model 400 support (support for Side-1 CHPIDs). Support for 64 channel paths on the Models 200E/300E and 128 channel paths on the Models 400E/600E, as well as the 4.5Mb per second block multiplexer channel transfer support became available during 4Q87. The following list of IOCP APARs provides the above support.

- Model 400/400E Support (Side-1 Channel Paths)

  - UZ90453 (OZ91644) - MVS/SP 2.1.7, 2.1.3, 1.3.5
  - VM26002 - VM/SP HPO
  - VM26003 - VM/XA

- Additional Channel Path Support (64 on 200E/300E & 128 on 400E/600E) and 4.5Mb per second block multiplexer channel transfer support

  - OY03274

---

[15] The program name for the 308X IOCP is ICPIOCP.

I/O Configuration Datasets

IOCDS datasets are located on the 3370 DASD attached to the 3092 Processor Controller. IOCDS datasets can be read/written and each can also be individually write protected. Non-partitionable 3090 processors have four IOCDS datasets (A0 - A3). Partitionable 3090 processors have eight IOCDS datasets (A0 - A3 and B0 - B3). In partitioned mode, each side can access only those IOCDS datasets belonging to the side (e.g... Side-0 has access to only A0 - A3, and Side-1 has access to only B0 - B3). All eight IOCDS datasets are accessible in single image mode. In addition, single image mode operation supports a dual-write capability to corresponding IOCDS datasets (e.g... A0/B0, A1/B1, etc....).

The dual-write function allows the installation to define duplicate copies of an IOCDS on each side of a partitionable processor with a single execution of IOCP. The processor controller marks the dual-written pairs of IOCDS datasets as being equal (containing the same information and written at the same time). The benefits of this dual-write function are two-fold:

1. Minimal effort is required to ensure synchronization of a single image IOCDS. Should the processor complex ever be brought up in partitioned mode (either planned or unplanned as part of a recovery scenario), an IOCDS containing the same single image configuration would be available to either side. A later dynamic merge to single image mode could then be accomplished without requiring a system interruption (POR) to select the IOCDS containing the single image configuration.

2. Dual-writing the IOCDS datasets will prevent loss of RMF data that could occur as a result of dynamic partitioning of the processor complex. RMF, during its initialization, reads the active IOCDS and extracts information for the following reports:

   - I/O DEVICE ACTIVITY REPORT - The LCU number to which the device belongs. RMF issues message ERB265I if this data is unavailable.

   - I/O QUEUEING ACTIVITY REPORT - RMF issues message ERB260I if this report cannot be produced.

   - CHANNEL PATH ACTIVITY REPORT - Block or byte indication in the TYPE field will be blank. RMF issues message ERB265I if this data is unavailable.

     An example should suffice. Suppose that the A0/B0 IOCDS datasets were dual-written and that the processor is currently in single image mode with B0 as the active IOCDS. The processor complex is subsequently dynamically partitioned with Side-1 as the off-going side. Since the active IOCDS (B0) is associated with Side-1 (off-going side), it cannot be accessed from Side-0[16]. Should RMF now be stopped (or fail) and be restarted on the on-going side (Side-0 in this case), RMF would be allowed to read the A0 IOCDS dataset and continue reporting the above information. If the active

---

16 Message *The Active IOCDS is only valid on the off-going side. 64449* should have been displayed on the system console.

IOCDS (A0/B0 in this case) had not been dual-written, RMF would not be allowed to read the A0 IOCDS, and the above information and/or reports would not be produced. Additional information on the dual-write function can be found in *3090 Processor Complex Recovery Guide, SC38-0051-2* or later version.

Active IOCDS Dataset Considerations

The following considerations apply to 'active' IOCDS dataset access and the 'ACTIVE IOCDS' field on the CONFIG frame.

- Active IOCDS Dataset Access
  - The active IOCDS becomes inaccessible until the next POR
    1. When the active IOCDS is updated (partitionable and non-partitionable models).
    2. During single-Image to physical-partitioned reconfiguration when the active IOCDS is associated with the off-going side and it was not dual-written.
    3. During physical-partitioned to single-image reconfiguration when the active IOCDS is associated with the on-coming side and it was dual-written but was updated while in the physical-partitioned configuration.
- Active IOCDS Field (when partitioning)
  - The active IOCDS field is always blanked
    - For the off-going side.
    - For the on-going side when a 'POR Required' state exists:
      1. When the vary-side-off is initiated -or-
      2. As a result of the initiation of the vary-side-off.
  - The active IOCDS field on the on-going side continues to show the active IOCDS if a 'POR Complete' state remains after the vary-side-off is completed.
    - If the active IOCDS is associated with the off-going side the field on the on-going side is intensified.
- Active IOCDS Field (when merging)
  - The active IOCDS field remains unchanged if the active IOCDS on the on-going side is associated with the on-going side.
  - The active IOCDS field is de-intensified if the active IOCDS on the on-going side is associated with the on-coming side, AND the active IOCDS was 'dual-written' AND not modified while the system was in physical partitioned mode.
  - In all other cases, the active IOCDS field is blanked.

IOCP Macros

Some changes were introduced to the IOCP macros in support of 3090 processors. A brief look at each of the macros and a discussion of the changes follows.

ID Macro: There are no changes to the ID macro. Users are reminded that the first eight characters of the parameter specified by the MSG1= keyword become the Installation name associated with the IOCDS dataset. Users should take advantage of this capability and implement a naming scheme for their various I/O configurations.

CHPID Macro: No new keywords were introduced on the CHPID macro. However, there are some differences between 3090 support and 308X support.

Channel path ID's on the 3090 are contiguous. There is no gap in the CHPID number assignment as there was on the 308X. Non-partitionable 3090 models and Side-0 of partitionable models support CHPIDs starting with CHPID X'00' and continuing through CHPID X'3F' (total of 64 channel paths). Side-1 of partitionable 3090 models support CHPIDs starting with CHPID X'40' and continuing through CHPID X'7F'.

Byte channel path operation is restricted to a specific subset of the supported channel paths. Non-partitionable 3090 processors support a maximum of four byte channel paths (paths 00, 01, 04, 05). Partitionable 3090 processors support a maximum of eight byte channel paths (paths 00, 01, 04, 05, and 40, 41, 44, 45).

The 3090 processors also support 32 channels per channel set[17], in 370-mode as opposed to the 16 channels per channel set supported on the 308X processors. The two channel sets (0 and 1) can have a combined total of 64 channel paths, but no individual channel set can have more than 32 channels assigned to it. *Users should be aware that if the system is to run in 370 mode under MVS/SP V1, and more than 16 channels per channel set are defined to IOCP, IOCP will not flag this as an error.*

The channel number assigned to the channel path can now be specified as two hexadecimal digits in the range of 00 through 1F. This is to accommodate 32 channels per channel set. The leading zero does not have to be specified for channel numbers 0 through F. Channel numbers within a channel set must be unique. However, there is no requirement for contiguous channel numbers within a channel set (as there was on 308X processors).

---

17 Support for 32 channels per channel set is provided ONLY by VM/SP HPO R3.6, 4.2, and 5.0. MVS/SP V1 does not support greater than 16 channels per channel set.

CNTLUNIT Macro: No new keywords or parameters were originally introduced on the CNTLUNIT macro. However, IOCP support of the 3090 allows for a maximum definition of 48 control units per channel path, whereas the 308X process allowed only 16. **Note that this is a limit change for definition only.** ***It does not change the number of control units that can be physically attached to a channel path.*** The number of control units attachable to a channel path remains at eight, unless a switching unit such as the 3814/2914 is being used, in which case up to sixteen control units may be attached, subject to the cable length restrictions of the control units.

With this increase in the number of control units that can be defined per channel path, it should not be necessary to use the cluster technique[18] as was sometimes necessary in the 308X IOCP process. The cluster technique was used to circumvent the definition limit of 16 control units per channel path for those Installations utilizing a common IOCP across multiple configurations.

A new parameter for the PROTOCL= keyword on the CNTLUNIT macro was introduced with support for the 4.5Mb per second block multiplexer channel transfer option. Data transfer in data streaming mode at 4.5Mb per second for supported control units is enabled by specifying PROTOCL=S4 on the CNTLUNIT macro. PROTOCL=S indicates data streaming operation at a 3Mb per second transfer rate.

IODEVICE Macro: A new keyword, STADET=, has been added to the IOCP IODEVICE macro, and is used to control the use of the "Status Verification Facility." This parameter applies to XA mode only. The default setting is "Y" (facility active).

As of this writing, STADET support is disabled in the 3090 channel microcode. Users are encouraged to code STADET=N on every IODEVICE macro in their IOCP input stream rather than take the default (Y) setting, even though the facility has been deactivated. Coding STADET=N will ensure that there will be no surprises should the facility be re-activated in the future. If re-activated, user's can, at their discretion, regenerate the IOCP and specify STADET=Y for those devices for which the facility is to be made active.

3090/308X IOCP Compatibility

The IOCP input for the 3090 and 308X is similar but not compatible. The reasons for the incompatibilities are CHPID numbers and macro keywords and parameters. The format of the IOCDS datasets is also different between 308X and 3090 processor complexes.

If the 3090 processor is installed in an environment that also contains a 308X processor, then two IOCP input streams will have to be maintained, one for the 3090 and one for the 308X. Should the environment also contain a 4381 operating in XA mode, then three IOCP input streams will have to be maintained.

---

[18] The cluster technique is described in WSC Technical Bulletin *3081 Processor Complex - IOCP and Channel Configuration*, GG22-9209.

Combined SYSGEN/IOCP Input Stream

The MVS/XA SYSGEN 'IODEVICE' macro was updated to support the new IOCP 'STADET=' keyword. This support was originally distributed with PTF UZ72308. As long as four[19] digit device numbers are not specified in the ADDRESS= parameter of the IOCP 'IODEVICE' macro, a combined MVS/XA SYSGEN/IOCP input stream will work.

The MVS/370 SYSGEN 'IODEVICE' macro (DFP/370 Program Product (5665-295) version only) was also updated to support the new IOCP 'STADET=' keyword. This support was originally distributed with PTF UZ72306.

The DFP/370 Program Product 'IODEVICE' macro also includes support for the IOCP 'PATH=' keyword. However, no support was added for the IOCP 'UNITADD=' keyword. Therefore, a combined MVS/370 SYSGEN/IOCP input stream is feasible only if MVS/370 channel numbers (0 to F)[20] are specified on the CHPID macros, and no unsupported IOCP parameters are specified on the IODEVICE macro.

The following matrix shows which IOCP keywords are supported on the SYSGEN IODEVICE macro for MVS/XA (DFP/XA) and MVS/370 (DFP/370 and non-DFP/370). A combined IOCP/MVS SYSGEN input stream is feasible if the IODEVICE macro supports the IOCP keyword(s), or if the IOCP input stream does not contain keywords that are not supported by the IODEVICE macro.

	CUNUMBR	TIMEOUT	PATH	UNITADD	STADET
DFP/XA	Y	Y	Y	Y	Y
DFP/370	Y	Y	Y	N	Y
non-DFP/370	Y	Y	N	Y	N

19 IOCP supports specification of four digit device numbers. MVS/XA supports only three digit device numbers. Note that IOCP will not issue a warning or error message if a four digit value is specified for the ADDRESS= parameter on the IODEVICE macro.

20 Even though the 3090 processor supports a maximum of 32 channels per channel set, MVS/SP Version 1 supports only 16 channels per channel set.

I/O Configuration

Prior to discussing general I/O configuration guidelines for the IBM 3090 Processor Complexes, a brief review of channel path grouping is in order.

Models 120E, 150/150E, 180/180E 200/200E, 300E and Side-0 of Models 400/400E/600E												/// /// ///	Side-1 of Models 600E/400E/400			
150	180	200	400	120E	150E	180E	200E	300E	400E	600E	CHPIDs	GROUP	CHPIDs	600E	400E	400
S	S	S	S	S	S	S	S	S	S	S	00,01,02,03	0	40,41,42,43	S	S	S
S	S	S	S	S	S	S	S	S	S	S	04,05,06,07	1	44,45,46,47	S	S	S
S	S	S	S	S	S	S	S	S	S	S	08,09,0A,0B	2	48,49,4A,4B	S	S	S
S	S	S	S	S	S	S	S	S	S	S	0C,0D,0E,0F	3	4C,4D,4E,4F	S	S	S
O	O	S	S	O	O	O	S	S	S	S	10,11,12,13	4	50,51,52,53	S	S	S
O	O	S	S	O	O	O	S	S	S	S	14,15,16,17	5	54,55,56,57	S	S	S
–	O	S	S	–	–	O	S	S	S	S	18,19,1A,1B	6	58,59,5A,5B	S	S	S
–	O	S	S	–	–	O	S	S	S	S	1C,1D,1E,1F	7	5C,5D,5E,5F	S	S	S
–	–	O	O	–	–	–	O	O	O	O	20,21,22,23	8	60,61,62,63	O	O	O
–	–	O	O	–	–	–	O	O	O	O	24,25,26,27	9	64,65,66,67	O	O	O
–	–	O	O	–	–	–	O	O	O	O	28,29,2A,2B	A	68,69,6A,6B	O	O	O
–	–	O	O	–	–	–	O	O	O	O	2C,2D,2E,2F	B	6C,6D,6E,6F	O	O	O
–	–	–	–	–	–	–	O	O	O	O	30,31,32,33	C	70,71,72,73	O	O	–
–	–	–	–	–	–	–	O	O	O	O	34,35,36,37	D	74,75,76,77	O	O	–
–	–	–	–	–	–	–	O	O	O	O	38,39,3A,3B	E	78,79,7A,7B	O	O	–
–	–	–	–	–	–	–	O	O	O	O	3C,3D,3E,3F	F	7C,7D,7E,7F	O	O	–

S = standard O = optional – = not applicable

Figure 27. 3090 Channel Path Grouping

Channel paths on the 3090 Processor Complexes are internally packaged in groups of four[21]. A number of such channel groups (depending on model) make up a standard configuration. Optional configurations (again depending on model) are made up of additional channel groups. Figure 27 depicts the standard and optional channel groups (along with the respective CHPID numbers assigned to the group) available on the various 3090 Processor Complexes.

[21] While internal packaging is in groups of four, the granularity for adding optional channel paths to a system is in units of 8 or 16 channel paths (depending on model and optional channel group being added).

General configuration guidelines call for configuring devices across control units and channel paths to minimize the potential that a failure in the data transmission path (control unit, channel path, channel subsystem element) will prevent access to the device.

Devices supporting multiple active interfaces (e.g.. DASD, TAPE, 3725 with 2-processor switch, etc...) should be configured with a primary and backup path(s) and operate with a minimum of two active paths.

Devices supporting a single active interface (e.g.. unit record, graphics, 327X controllers) should also be configured with a primary and a backup path. The use of the primary/backup path should be controlled through a switching unit, such as the IBM 3814 Switching Management System or the IBM 2914 Switching Unit. The use of a switching unit will facilitate recovery in case of channel path failure.

In both cases above, the primary and alternate or backup paths should be configured from different channel groups. In the case of MVS/370, if the configuration supports two channel sets, then the primary and alternate or backup paths should be assigned across the channel sets.

Considerations for Single-Image Operation

For SI operation, additional consideration should be given to configuring devices so that access from each side of the partitionable system is possible. Access is considered either symmetric or asymmetric. The terms are defined as follows:

1. **Symmetric access:** Symmetric access implies that the device can be attached to multiple (two or more) channel paths and actively support I/O operations on these multiple paths. Symmetry is obtained by configuring at least one path to the device from each side of the partitionable processor. DASD and Tape devices are common examples of devices that can be configured for symmetric access.

2. **Asymmetric access:** Asymmetric access implies that the device, while having attachment capability to more than one channel path (via a two channel interface switch, for example) can operate only on a single path. The term 'asymmetric access' is still applicable even though the user is operating under the control of MVS/XA and the S/370-XA Channel Subsystem. Although there is no CP-path affinity in S/370-XA mode operation, access is asymmetric with respect to the channel subsystem elements[22] of the two sides of a partitionable processor. This fact must be considered for recovery purposes and when operational plans include partitioning (either dynamic or static) of the partitionable processor. Consoles, graphics devices, communications controllers and printers are examples of asymmetric devices.

Symmetrically accessed devices should be configured with a minimum of one path from each side and should operate with a path from each side active. Whether two paths (one from each side) or four paths (two from each side)

---

22 **Partitionable models of the 3090 processors contain two channel subsystems (one on each side). In SI mode the two hardware channel subsystems appear as a single dynamic channel subsystem to the control program.**

should be configured depends to some extent on the use of the system. When the system is configured for SI operation, two paths may be sufficient (one from each side). However, if the role of the system is to also act as a backup (via partitioning) for another system, then more than one path from the same side may be required.

Asymmetrically accessed critical devices should also be configured with a path from each side (through a switching device or two channel interface connection). The additional path provides backup in case of path or side failure, and also in those instances where static partitioning is used to recover from a single-image failure.

If operational plans call for dynamic partitioning of the system (for example, after primary production hours in order to provide a test system), consideration should be given to the active channel path for devices with asymmetric access. The active path for these devices should be assigned to the side that will not be partitioned off. For example, if operational procedures call for partitioning off Side-1 of the system to be used for test purposes, then devices with asymmetric access required to support applications that must continue to run on Side-0, should normally have the active path assigned to Side-0. Applications using these devices can thus continue to run uninterrupted during the partitioning and subsequent merge process.

One general SI configuration point to keep in mind is that **devices available from both sides of a partitionable processor, having the same device number, must be the same device.** This point has been overlooked by some users migrating to SI operation who have not had prior SI experience. One situation that has been encountered multiple times involves operator consoles. In this situation, addressing standards dictated that the master console address be the same across all the systems (for example, address 101). Later, when one of the systems was upgraded to a partitionable processor due to run in PP mode interim to SI mode, it's master console was also assigned address 101. Figure 28 is an example of IOCP statements used to define such configuration.

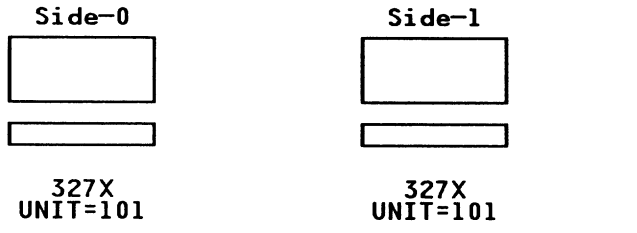

Figure 28. Example of Invalid Single-Image Console Definition

As long as the system is run in PP mode the system will operate correctly. Path 08 is only available from Side-0 and path 48 from Side-1. The SCP in each side communicates with a master console at address 101 (down path 08 for Side-0, and down path 48 for Side-1).

When the system is operated in SI mode a problem arises. When the MVS/XA system initializes, it finds two paths (08 and 48) available to device 101. The system has no way of determining that the device at the end of these two paths is not the same device. Which physical device eventually receives the activity directed to device 101 is dependent on which sides (0 or 1) channel subsystem handles the request. If Side-0's channel subsystem handles the request, the activity will be sent to the device on path 08. Conversely, if Side-1's channel subsystem handles the request the activity will be sent to the device on path 48. Thus, console messages would be spread in some random fashion across both devices, making operation difficult. Such a configuration would work in SI mode if one of the two devices were always offline (interface disabled). For correct SI mode operation, ensure that all devices have unique device numbers.

Console Considerations

The MVS master and alternate console should consist of a dedicated control unit/display combination. That is, a control unit with a single display attached for the master console, and a control unit with a single display attached for the alternate console. This arrangement ensures maximum availability and accessibility during recovery operations.

When planning configurations for non-partitionable systems, two such units are sufficient. For partitionable systems operating in PP mode, four such units are required (a master and alternate per side). For partitionable systems operating in SI mode, three such units are sufficient. One unit configured directly to each side and the third unit configured so that it can be switched to either side. In the event of a SI mode failure and subsequent recovery running PP mode (until SI operation can be resumed), the third unit can be used to provide the alternate console for production operation. This would leave the side experiencing the failure with a console to be used in testing the side (once repair has been made) prior to a merge back to SI mode operation.

Console Configuration Guidelines: Consoles should be configured according to the following guidelines:

1. Configure the master and its alternate console on different channel paths, with the different paths coming from separate channel groups.

2. For SI operation configure the master and alternate console across sides of the partitionable processor.

3. If multiple 3x74 type control units are configured on the channel paths containing the master/alternate consoles, then ensure that the control units attaching the master/alternate consoles are configured ahead of (higher in channel priority) other 3x74 type control units on the channel path.

Additionally, for MVS/370 operation, if the configuration supports two channel sets, configure the master console and the alternate console across the two channel sets.

Miscellaneous Topics

DCCF System Console Support

Disabled Console Communications Facility (DCCF) support in MVS/XA was enhanced to allow displaying messages on the hardware system console of 308X and 309X processors when the MVS master console and its first alternate are not available. Previously, there was no capability in MVS/XA to communicate with the operator via the system console. Figure 29 shows the order of search for displaying messages output by DCCF both before and after the enhancement.

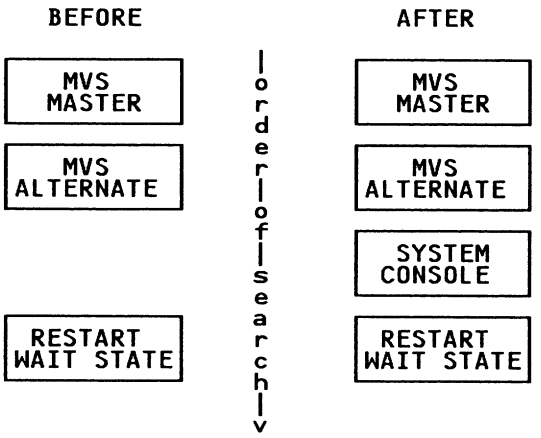

Figure 29. DCCF Message Display - Order of Search

This support increases the system's ability to communicate with the operator via messages, thus reducing the number of restartable wait states. DCCF will route messages to the system console under the following circumstances:

1. The DCCF message requires a reply and the MVS master/alternate consoles (first and second choices) are not available.

2. An operator reply to a DCCF message sent to the MVS master/alternate console was not read within 125 seconds.

Each message issued to the system console is preceded by ringing the Processor Controller alarm to attract the operator's attention. The current activity on the system console is not preempted. A message, indicating that a message is pending and the system message facility frame should be displayed, is broadcast to the system console and the service support console. To view the messages issued through DCCF on the system console, the operator must invoke the system message facility frame of the processor (SCPMSG for 3090 processors and SYSMSG for 308X processors). The operator's reply must not be preceded by "R 0, " (in contrast with the reply on the master/alternate console, where characters "R 0, " are displayed by DCCF to indicate the reply line). If the message

cannot be output to the consoles (MVS master, MVS alternate, or system console) DCCF will put the system into a restartable wait state.

While there is a 125 second timeout associated with DCCF messages sent to the MVS master/alternate console, after which the system is put in a restartable wait state, this timeout does not apply to DCCF messages sent to the system console. In this case, the system waits (not necessarily a wait state) for the operator reply.

Note: Messages issued via DCCF are not displayed on the system console when MVS/XA is running as a guest on a VM/XA SF system.

On machines which do not support communication with the system console, or if a DCCF message can not be displayed on the system console for any other reason, a restartable wait state will be loaded as it was prior to this enhancement.

While this enhanced support has the potential for reducing the number of restartable wait states for recovery situations using DCCF when the master and alternate consoles are not accessible, it does NOT eliminate the need to provide a dedicated control unit/display combination for the master and alternate console. The reasons for this are twofold:

1. The system console is not always located near the MVS/XA master operators console (even though this is recommended).

2. Not all callers of DCCF request use of the system console as an option for displaying messages. An example of such callers are the MVS/XA restart routines.

For these reasons users should continue to provide a separate control unit/display combination for the MVS/XA master console and a similar arrangement for the MVS/XA alternate console. Only by dedicating these resources to the master and alternate consoles can a consistent operator interface be provided during recovery operations.

The above enhanced support is described by APARs OZ91846, OZ92108, and OZ92109 and was originally provided by the following PTFs.

MVS/SP 2.1.3	- UZ90481, UZ90484
MVS/SP 2.1.3 VFE	- UZ90482, UZ90485
MVS/SP 2.1.3 AE	- UZ90486
MVS/SP 2.1.7	- UZ90487

Processor Identification

There is a single unique 5-digit processor serial number associated with 3090 processors. Individual CP's are numbered 0 through 5 (depending on model), and are identified by the first four bits of the CPU identification number (bits 8-11 of the doubleword stored by the STIDP[23] instruction). For CP IDs and

---

[23] Additional information on the STIDP instruction may be found in *IBM System/370 Principles of Operation* GA22-7077, and *IBM System/370 Extended Architecture Principles of Operations* SA22-7085.

processor complex version codes the reader should refer to Chapter 1, "IBM 3090 Family Comparisons" in section "Comparison Summaries" on page 12.

User-written or non-IBM code that may be dependent on the CPU number, model (CPU type) or version code, must be identified and changed to recognize the new values.

TOD Clock Setting

Each central processor (CP) of a 3090 processor complex has a TOD (Time-of-Day) clock associated with it. If the TOD clock requires setting at IPL time, it will be necessary to enable TOD clock setting from the system console for each online CP in the processor complex. For example, on a Model 600E in single image mode the TOD enable function would have to be done six times. TOD enablement can be accomplished in one of three ways:

- Using the OPRCTL frame
  - Enable the S= field and press ENTER
- Using the TOD enable key function
  - Press the ALT and TOD keys at the same time
- Using the TOD enable service language command
  - Key in "TOD" at the system console, and press ENTER

When varying a processor online that was previously offline, or during a dynamic merge from partitioned mode to single image mode, it will also be necessary to enable TOD clock setting for each CP that is being brought online.

3180 Model 140/145 Setup Consideration

As noted in the *3090 Processor Complex Operator Controls for the System Console, SC38-0040* the console status line (line 25) provides a SIDE:MODE indicator. The SIDE indicator displays the side of the processor controller (A or B) to which the display is attached. The mode indicator is used to report on side status, transitional states for the processor controller, and also availability of a backup processor controller side. The latter indication involves the use of both upper and lower case characters.

For example, operating a 3090 Model 400/400E or 600E in PP mode, the normal display is "A:a MODE PP" on displays attached to the A side of the processor controller, and "B:a MODE PP" on displays attached to the "B" side of the processor controller. The lowercase "a" indicates that there is NO backup processor controller side available. In SI mode, the status line would normally show "A:A MODE SI" indicating the availability of a backup processor controller side.

In order for the mode indicator to be displayed in lower case, the individual 3180 display must be setup for dual case operation. Dual case operation is achieved by depressing the 3180 'SETUP' key and then depressing the "FWD A,a" key until the correct case is displayed, then depressing the 'SETUP' key again to exit from setup mode.

HSA Expansion

The Hardware System Area is an area of real storage set aside for use by the hardware. The HSA contains system microcode, control blocks describing the hardware I/O configuration, as well as entries for the I/O trace function when it is activated. The size of HSA therefore, is a function of the amount of system microcode, the IOCDS selected at POR, and the number of I/O trace units specified (default of 2 on the 3090).

The HSA is allocated from the top of real storage down in increments of 32Kb within a 512Kb section. When a 512Kb section fills up, or there is not enough room for the next piece to be allocated in HSA, a new section is allocated for HSA starting at the next lower storage increment. An HSA section does not necessarily have to be 512Kb in size before the next HSA section is allocated. HSA expansion for all purposes can be up to 3Mb (six 512Kb sections). The amount of central storage available for use by the SCP is reduced by an amount equal to the combined sizes of the HSA sections.

The following example shows three sections of allocated HSA in a 128Mb system with both a 2Mb and a 4Mb storage increment size. Enhanced 3090 Models 180E and above have a storage increment size of 4MB. All other 3090 processors (E-models and non-E-models) have a storage increment size of 2Mb,

Although only 3 HSAs are shown, remember that there can be up to 6 of them.

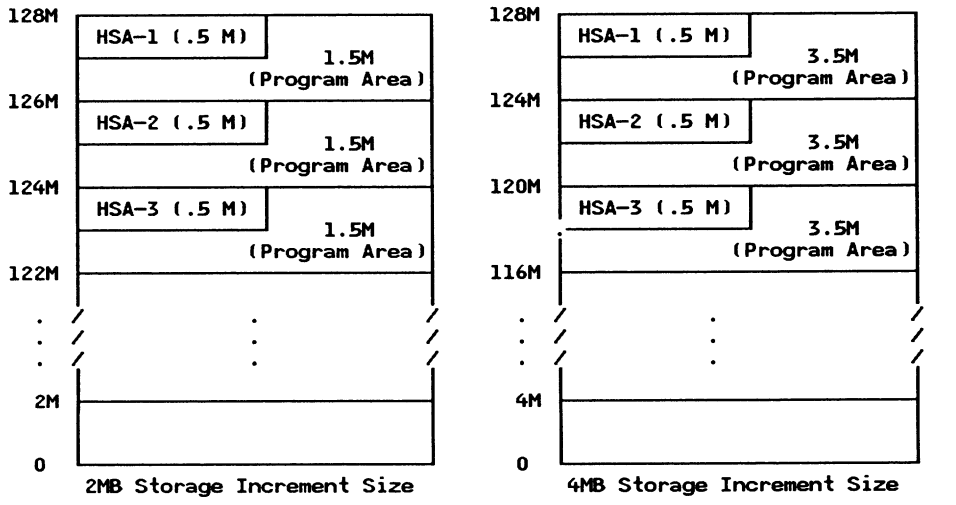

Figure 30. 3090 Processor HSA Allocation

Increased Extended SQA (ESQA) Requirement

When adding expanded storage to a 3090 system, users should ensure that enough storage has been set aside in Extended SQA to hold the control blocks describing expanded storage frames. There is a 32 byte ESTE control block for each frame of expanded storage. Failure to specify enough ESQA storage could result in an "064" wait state during IPL.

The amount of ESQA storage (in number of 64K blocks) is specified by the Extended SQA subparameter on the SQA=(a,b) statement in the IEASYSxx member of SYS1.PARMLIB. The subparameter 'a' specifies the size of SQA, located below 16 megabytes. The subparameter 'b' specifies the size of Extended SQA, located above 16 megabytes. These values are added to the system's minimum virtual SQA of four 64K blocks (256K) and minimum virtual Extended SQA of approximately 8 megabytes. Both SQA and Extended SQA are fixed in real storage as they are used.

The following table lists the amount of storage required (in megabytes) and the number of 64K blocks required for the various increments of expanded storage available on 3090 processors.

ES Size	Storage in MB	# of 64K Blocks
64Mb	.5Mb	8
128Mb	1.0Mb	16
192Mb	1.5Mb	24
256Mb	2.0Mb	32
384Mb	3.0Mb	48
512Mb	4.0Mb	64
1024Mb	8.0Mb	128
2048Mb	16.0MB	256

Any plans for adding expanded storage to a 3090 system should include a step to re-assess the number of 64K blocks specified for Extended SQA, and if necessary, increasing the number.

Potential Real Storage Usage Increase

When additional processor storage is added to the 3090 complex (i.e., model upgrades or MES for additional central or expanded storage options), there is additional real storage required for page frame table entries (PFTE's) that reflect central storage and for extended storage table entries (ESTE's) that reflect expanded storage. This is most often **discovered** when looking at an RMF report and finding that the size of the NUCLEUS or the number of fixed frames (or both) has increased.

For MVS/SP 1.3.5, there is a sixteen (16) byte PFTE for each frame of real storage. These are located in the NUCLEUS below the 16Mb line. For a system that increases central storage by 64Mb, there would be an increase in the NUCLEUS of about 256 Kb (64 frames).

For MVS/XA, there is a thirty-two (32) byte PFTE for each frame of real storage. These are located in the Extended NUCLEUS above the 16Mb line. There is also a thirty-two (32) byte ESTE for each frame of expanded storage. These are located in Extended SQA above the 16Mb line and are page-fixed. For a system that increases central storage by 64Mb, there would be an increase in the Extended NUCLEUS of about 512Kb (128 frames). If expanded storage was increased by 64Mb, the increase in Extended SQA (fixed) would be about 512Kb (128 frames). Storage increases of 128Mb would, therefore, require an increase of about 1Mb (256 frames) in either ENUC or ESQA.

Error Recovery Enhancements

In 1986 IBM enhanced the CPU and channel error recovery capability of the IBM 3090 Processor Complex to:

- Increase system availability
- Increase the potential for deferred maintenance actions
- Enhance measurement of system recovery.

Bit assignments have been added to the LOGREC records (MCH and SLH) to more precisely define the results of hardware and software error recovery actions and to allow for the elimination of redundant records.

MCH Record Processing

Prior to this support, storage key errors that were presented by the hardware as uncorrected but were subsequently recovered by MVS/XA were still reported as hard errors in the Machine Check Handler (MCH) LOGREC record. These errors also caused the IPD (Instruction Processing Damage) threshold count to be incremented. MVS/XA added a new LRB flag (LRBMRECV) to indicate that a storage key error was corrected and no longer increments the IPD threshold if:

- the storage key was reset,
- the interrupted instruction can be retried, and
- there are no other related errors.

Whenever the LRBMRECV bit (byte 3, bit 5 in the MCH record) is ON, a "soft" error is indicated. In this case, because the IPD threshold will not be incremented, these corrected errors will not contribute to taking a CPU offline because of threshold processing. It should be noted that there is no change to the other hardware information in the MCH record for this error, just the added bit if MVS/XA was able to recover from the error.

SLH Record Processing

To eliminate redundant LOGREC recording of SLH (Subchannel Logout Handler) records for the same recovery attempt, changes were made to the IBM 3090 channel microcode and to MVS/XA to reflect the status of hardware and software recovery actions. This, in effect, means that retries will not be recorded unless the conditions change. For example, an IFCC changes to a CCC or an alternate path is used. There will not be an indication of the number of retries

from the LOGREC record, however, all system messages will be as before. All incidents of errors will be recorded.

The results of these recovery actions are indicated in the SLH LOGREC records as:

- SOFT - Failure was successfully recovered by the system. A time dependent application may be impacted.
- DEGRADE - Failure was successfully recovered by the system, but hardware resources may have been lost; performance may be degraded; or a time dependent application may be impacted.
- HARD - Failure is not successfully recovered by the system. One or more jobs or the operating system may have been lost and/or impacted. Hardware resources may have been lost.

These indicators are stored in the SLH at byte 5 (Record Dependent Switches byte 2), bits 6 and 7 with the following meanings:

X'01' - HARD
X'02' - DEGRADE
X'03' - SOFT

Required IBM 3090 Hardware Levels

These enhancements were initially provided in the following Models of the IBM 3090:

- Included in all E-Models
- Included in all Model 400s
- Model 180s and 200s with SEC 218590 or later
- Model 150s with SEC 219109 or later

Required MVS/XA Software Support

The required MVS/XA maintenance support for these enhancements is provided by the following APARs:

- OZ97887 - MCH
- OZ97930 - DFP
- OZ97996 - IOS
- OZ97999 - EXCP

These changes to the MCH and SLH record formats are documented in TNL GN28-0863 to the *MVS/XA SPL: SYS1.LOGREC Error Recording, GC28-1162* manual.

Setting Machine Check Threshholds for Storage Errors

When the hardware detects a storage error, it notifies the SCP by presenting a machine-check interrupt. If the hardware was able to correct the error, the machine-check interrupt indicates a "soft" System Recovery (SR) machine check. If the hardware was not able to correct the error, the machine-check interrupt indicates a "hard" Instruction Processing Damage (PD) machine check.

The Real Storage Manager in MVS/XA 2.1.3 and later releases contains enhanced recovery processing for storage errors that may make it desirable to raise the thresholds for both SR and PD machine checks.

Corrected Storage Errors

When an SR machine check for a corrected storage error is presented, MVS/XA attempts to stop using the affected frame. This is done to minimize the chance that the problem will later re-occur as an uncorrected storage error and to avoid any potential degradation due to hardware corrective action on subsequent errors for the same frame. If the frame contains pageable data, MVS/XA moves the data to another frame and the original frame is marked offline. If the data in the frame cannot be moved, that frame is marked "pending offline" and is subsequently taken offline if the frame is released, or its contents are made pageable.

Prior to taking the frame(s) offline, MVS/XA tests the frame(s). The testing of the frame(s) is done asynchronously (under an SRB) rather than as part of the actual machine check processing. Depending on the results of the test, MVS/XA (RSM) will either put the frame back on the available queue or take it offline.

A low threshold for SR machine checks affects MVS/XA's ability to react to these corrected storage errors. When the threshold for SR machine checks is reached, MVS/XA disables SR machine checks. This results in any subsequent corrected storage errors not being presented to MVS/XA and therefore no action can be taken to remove the affected frame.

The current default threshold for SR machine checks is four (4). When the threshold is reached, MVS/XA notifies the operator via message IGF931I that the limit has been reached and disables this class of machine checks. This results in any subsequent corrected storage errors not being presented to MVS/XA and therefore no action can be taken to remove the affected frame.

Users should consider raising the SR threshold to 50 for each CP of the processor complex. This will allow MVS/XA recovery processing to handle more corrected storage error conditions for a given IPL. To raise the limit issue the following MVS/XA operator command or include it in the command member (COMMNDxx) of PARMLIB.

```
MODE SR,RECORD=50,CPU=ALL
```

Although this recommendation applies to all systems supported by MVS/SP 2.1.3 or later, it is especially recommended on 3090 systems. Since the 3090 provides double-bit error correction, a larger percentage of storage errors should be presented as corrected.

Uncorrected Storage Errors

Uncorrected storage errors represent the potential loss of critical data. In most cases, MVS/XA will terminate the affected unit of work. If recovery processing related to this termination completes successfully and results in the freeing of the affected frame, the frame is marked offline and system processing continues. However, the recovery processing itself may attempt to reference the storage that caused the machine check and consequently cause additional errors. This, in turn, could result in the threshold for PD machine checks being reached and a CP being taken offline.

To reduce the chances of having an uncorrected storage error cause the loss of a CP, users should also consider raising the threshold for PD machine checks on all CPs in the processor complex, from 5 in 5 minutes (default) to 16[24] in 5 minutes. To raise the limit issue the following MVS/XA operator command or include the command in the command member (COMMNDxx) of PARMLIB.

```
MODE PD,RECORD=16,INTERVAL=300,CPU=ALL
```

Note: The above recommendations for both SR and PD machine check threshold settings will be incorporated into the fix for APAR OY08020 as the defaults.

SEC 223620

SEC 223620 was the general availability EC level[25] for the 3090 Enhanced Model Processors. The SEC contains a number of operational enhancements as well as 3480 tape support for the 3092 Processor Controller. The SEC was incorporated into MES model upgrades during April 1987 and with normal SEC shipments for previously available 3090 Processors (non-E-models) during May 1987.

Operational Enhancements/Changes

The following operational enhancements/changes have been incorporated into SEC 223620. User's should ensure that operational procedures are updated to reflect these changes.

- Rotatable SAD (System Activity Display) Frame Enhancement

  This enhancement provides the capability to rotate between multiple SAD frames (up to sixteen). The roll rate is distinct from the refresh rate and can be set by the operator. The default roll rate is 10 seconds.

---

[24] This is a change from the previous recommendation to raise the limit from 5 in 5 minutes to 25 in 5 minutes.

[25] One-step install 3090 Enhanced Models were originally shipped at SEC level 223620. 3090 Enhanced Models installed via the two-step install process will be at SEC 223620 level upon completion of the step 2 modification.

- SAD Display "Except Key" CP Utilization Enhancement

  This enhancement allows definition of CP state activities (supervisor, problem or both) with the exception of activity associated with the above states for a selected key. This enhancement is beneficial to those system control programs employing an active wait instead of an actual enabled wait state PSW during idle periods.

- IFRST Command Enhancement

  A value of "ALL" must now be specified if the purpose of issuing the IFRST command is to reset all channel interfaces. This change was implemented to ensure positive action by the system operator if all channel interfaces are to be reset. The previous implementation of this command defaulted to "ALL" if no CHPID number was specified with the command. As a result, failure to specify a CHPID number caused all channel interfaces to be reset, resulting in a system outage. Entering the IFRST command without specifying a CHPID number or "ALL" will now result in an error message.

  Note: The IFRST command has several aliases: INTFRST, IFRST, RSTINTF, and RSTINF. The above information applies to the command aliases as well.

- CONFIG Frame Change

  The "A= ACTION" field on the CONFIG frame has been updated to include the selection of a maximum configuration. This option is now selected via A3. The previous A3 option (Select IOCDS Mgmt.) is now option A4. Operator procedures should be updated to reflect this change.

- I/O Power Sequence Delay Option

  The I/O Power Sequence Delay option is an optional delay specified in minutes and seconds between the completion of a power-on-reset and the start of auto-IPL. The purpose is to allow I/O power sequencing to complete before IPL is initiated. The option is enabled and/or disabled via the operator SYSDEF frame on the system console.

Support for 3480 Attachment to the Processor Controller

This support became available in 2Q87. It was included in System Engineering Change (SEC) 223620 and incorporated into the 3090 E-models. This SEC also began shipments with MES model upgrades during April 1987 and with normal SEC shipments for previously available models during May 1987.

This SEC **adds** the 3480 support and does not remove or enhance the previous support for 3420 tape attachment to the Processor Controller.

This new support is for 3480 Models B11 or B22 tape drive (via 3480 Model A11/A22 Control Unit). The control unit can be up to 400 feet from the Processor Controller. It uses DCI (DC Interlock) protocol with support for a maximum data rate of 1.25 megabytes per second.

Note: Before attempting to access the tape drive from the 3092 PCE, ensure that:

1. The tape drive that will be accessed can be assigned any valid S/370 control unit/device address in the range X'00-9F' and X'B0-FF'. This range EXCLUDES control unit addresses X'A0-AF'.

2. The tape drive to be accessed by the 3092 PCE is varied offline to ALL other processor complexes in the installation, i.e.:

   a. MVS: - 'VARY PATH(ddd,xx),OFFLINE'
   b. MVS: - 'VARY ddd,OFFLINE'
   c. VM: - 'CP RESET' if "unsupported" device in DMKRIO
   d. VM: - 'CP DETACH' device from VM

3. If tape access is provided using a switching device (e.g., 2914, 3814), standard considerations should also be given when switching the 3803 or 3480 control unit interface:

   - Repeat the VARY OFFLINE, RESET, DETACH steps for all devices on that control unit interface.
   - Switch the interface.

4. **Procedure 1 - To set DCI mode in the 3480 CU.**

   If a 3480 Model A11/B11 tape subsystem or 3480 Model A22/B22 tape subsystem is used and the interface is not already in DC Interlock mode, then DC Interlock data transfer mode must be set using the following steps:

   a. Record the current setting of the "Type" thumbwheel switch for the interface that is to be used with the 3092. It will be used to restore the original channel data transfer type setting as described in Procedure #2 (see below).

   b. Set the channel interface Enable/Disable switch to the Disable position on the 3480 tape interface that is connected to the 3092. The green Disable light must come on before continuing to the next step.

   c. Set the "Type" thumbwheel switch to DC Interlock: either switch position 2 or switch position 3 depending on the CU address. (Consult the "CU Address Instructions" on the 3480 Control Unit Operator Setup Panel.)

   d. Set the channel interface Enable/Disable switch to the Enable position.

5. The 3480 Tape Control Unit Interface path is now available for use from the 3092.

6. **Procedure 2 - Restoring the original 'CHANNEL ADDRESS TYPE' setting.**

   If a 3480 Model A11/B11 tape subsystem or 3480 Model A22/B22 tape subsystem was used and the channel "DATA TRANSFER TYPE" was changed to DC Interlock in Procedure #1 (see above), then perform the following steps:

   a. Set the channel interface Enable/Disable switch to the Disable position on the 3480 tape CU interface that is connected to the 3092. The green Disable light must come on before continuing to the next step.

   b. Set the "Type" thumbwheel switch to the setting which was recorded earlier in Procedure #1.

   c. Set the channel interface Enable/Disable switch to the Enable position.

SEC 223630

SEC 223630, which applies to all 3090 Processor Models and became available in October 1987, incorporates support for previously announced additions to functions and features (some of which are only available on 3090 Enhanced Model Processors), as well as a number of operational enhancements.

Function/Feature Support

- 4.5Mb/Second Block Multiplexer Channel Support (All 3090 Models)

  Support is provided for block multiplexer channel operation at a 4.5Mb/second transfer rate. The 4.5Mb/second transfer rate in data streaming mode for supported control units is enabled by specifying a parameter on the IOCP CNTLUNIT macro during IOCDS generation.

- Additional Expanded Storage Support (200E/300E/400E/600E)

  Support is provided for up to 512Mb of expanded storage on the 3090 Models 200E/300E and up to 1024Mb of expanded storage on the 3090 Models 400E/600E.

- Third Channel Group Support (200E/300E/400E/600E)

  Support is provided for up to 64 channel paths on the 3090 Models 200E/300E and up to 128 channel paths on the 3090 Models 400E/600E.

- Stand-alone IOCP (All 3090 Models)

  – Third Channel Group Support

    Stand-alone IOCP has been updated to support the definition of 64 CHPIDs per side (for those 3090 Enhanced Models supporting attachment of the third channel group). Definition for CHPIDs in the range 00-7F (00-3F on Side-0 and 40-7F on Side-1) are now supported. Previously IOCP supported the definitions of CHPIDs 00-2F and 40-6F.

- 4.5Mb/Second Data Transfer Support

  Stand-alone IOCP has been updated to support definition of control units capable of operating at a 4.5Mb per second data transfer rate. Data transfer at 4.5Mb per second in data streaming mode is defined by specifying PROTOCL=S4 on the IOCP CNTLUNIT macro during IOCDS generation. Activation of the updated IOCDS requires a processor complex power-on-reset. Additional information on the use of this parameter can be found in TNL SN22-5138 to the *3090 Processor Complex - IOCP User's Guide and Reference, SC38-0038.*

- 3422 PCE Tape Attachment (All 3090 Models)

  Support is provided for attachment of the 3422 tape to the processor controller. The processor controller now supports attachment of 3480, 3420 or 3422 tape units. The 3422 must be at EC level A46632 or higher.

- SIE Support of VMA Functions (3090 Enhanced Models Only)

  Start Interpretive Execution support of Virtual Machine Assist functions is provided. Enhanced SIE support of VMA functions is applicable to 3090 Enhanced Model processors only. It provides for a performance improvement when running VM/SP or VM/SP HPO systems as guests of VM/XA Systems Facility or VM/XA System Product.

Operational Enhancements/Changes

- Configuration (CONFIG) Frame (All 3090 Models)

  - The terminology 'PROCESSOR STORAGE' on the CONFIG frame has been changed to 'CENTRAL STORAGE'. This change was made to make the terminology more precise.

  - The selection groups on the CONFIG frame for Central and Expanded Storage are renumbered so that

    - PMA0-3 are selected using F0-F3 (previously F1-F4)
    - ESA0-3 are selected using G0-G3 (previously G1-G4)

- Operator Control (OPRCTL) Frame (All 3090 Models)

  The OPRCTL now includes a new option ('D') indicating selection of PCE stall detection. When enabled, a Service Processor Damage (SPD) machine check is generated when both the active and backup (if available) processor controller element sides become unusable.

  Note: A backup PCE is not available under the following conditions:

  - Never on 3092 Model 3 (contains only a single processor)
  - On 3092 Model 2 when a partitionable processor (Models 400/400E/600E) is run in partitioned mode.
  - On 3092 Model 1 if one side of the PCE has failed or is offline for maintenance.

- Channel Configuration (CHNCFA) Frame (All 3090 Models)

  - The CHNCFA frame can now display up to 64 channel paths per side (for those 3090 Models supporting attachment of the third channel group). The previous maximum was 48 channel paths per side. A second panel is added to enable the display of 128 channel paths for partitionable systems operating in single-image mode.

  - The CHNCFA frame has also been enhanced to display an "I" on the 'CONFIGURED' line under each CHPID that has exceeded the interface control check (IFCC) threshold. When the threshold is reset, the "I" is also reset.

- Active IOCDS (All 3090 Models)

  The active IOCDS will now be automatically *write protected* at power-on-reset time. This change was made as a precaution to prevent inadvertant alteration of the active IOCDS after the system has been power-on-reset. The "write protect" status of the active IOCDS can still be released after the power-on-reset is performed by use of the A3 option on the IOCDSM frame.

- Problem Analysis (PA) Facility Improvement (All 3090 Models)

  The performance of the PA facility has been improved. The time delay between the selection of a PA frame and the appearance of the frame on the display has been reduced to approximately 10 seconds.

4.5Mb/sec Block Channel Support

Hardware Availability

- Available for all 3090 models

  - New 3090s - beginning October 30,1987
  - SEC 223630 for installed 3090s beginning October 30,1987

- Available for all 3088 models

  - New 3088s - beginning October 30,1987
  - EC A50442 for installed 3088s beginning October 30,1987

- Available for all 3814 models

  - Via Specify Feature #9001, beginning September 1,1987
  - Feature #9001 is for information only, no parts or instructions are shipped.

- Available for 3990 Model 3, and 3880 Models 21 and 23

  - 3990-3 availability is 3Q88
  - Feature #6140 for 3880 Models 21 and 23 beginning October 1987 for new and MES orders

Steps Required to Enable 4.5Mb/sec Support

1. Install SEC 223630 on 3090.

   SEC 223630 provides support for 4.5Mb per second transfer in data streaming mode.

2. Place 3990, 3880-21, 3880-23 properly on the channel following the 4.5Mb channel configuration guidelines listed below.

   The installation of Feature #6140 on 3880 Models 21 and 23 may occur during this step or during steps 3 or 4.

3. Enable 4.5Mb/sec in IOCP for 4.5Mb capable control units.

   To specify 4.5Mb transfer for a control unit, use PROTOCL=S4 on the IOCP CNTLUNIT macro. Ensure the specified control units have been properly positioned during Step 2 above. Then do a Power-On-Reset to load the new IOCDS.

4. Enable the 4.5Mb/sec control unit channel interface setting on the 3990-3, 3880-21, or 3880-23 that you specified in the previous IOCP step.

   Steps 1, 2, and 3 must be completed prior to this step.

4.5Mb Channel Configuration Guidelines

The 4.5Mb/sec control units must physically precede slower control units on the block channel but do not need to be first in selection priority.

The 3880 models 21 and 23 with feature #6140 must follow all 3990 control units on the channel and must precede all other storage control units.

The 3990 and all control units supporting 4.5Mb/sec data transfer rate require all blue cables from the channel to the control unit.

Reduced Diameter Channel (Blue) Cables

The 3990 and all control units supporting 4.5Mb/sec data transfer rate require all blue cables. The reduced diameter (blue) cables have less line loss and less skew than the older gray cables. Additionally,

- The coaxial portion of the blue and the currently produced gray cables are the same.
- The PVC coverings or jacketing for both cables are the same except for the color and thickness.
- Gray channel cable previously produced were not produced to current tolerances and in some situations may not be functionally compatible with 4.5 Mb requirements. Therefore, in order to avoid any potential performance problems on 4.5 Mb control units, the use of reduced diameter cables (blue) is mandatory.

Interfaces and Switches

Feature #6140 for the 3880 provides a channel interface switch for each channel interface on the control unit. This switch can be independently set to 3.0 or 4.5Mb for each channel interface. Thus, one or more interfaces could transfer to the 3090 at the 4.5Mb rate and the remaining interfaces to processors supporting only a 3.0Mb rate.

To transfer at the 4.5Mb rate to a 3090, not only must the 3880 interface switch be set to 4.5, but the 3090 IOCP must also be updated to specify the 4.5 rate for the control unit involved. This is done by specifying PROTOCL=S4 on the IOCP CNTLUNIT macro for the affected control unit (or units).

If a 3088 (running at 4.5Mb) is attached to a 3044 (or another processor which does not have 4.5Mb capability), the 3088 will operate at the speed of the slower of the two connected channels.

The 3044 channel extender has a maximum data rate of 1.25Mb per second when operating with devices that use data streaming channel-control unit protocol, the 3088, in this case. The data rate of the 3044 will "gate" the flow of data between the two channels.

The 3814 will operate at the rate of the control unit and channel.

Software Requirements

The 3090 Processor Complex 4.5Mb/sec data transfer in data streaming mode for supported control units is enabled by specifying a parameter on the IOCP CNTLUNIT macro during IOCDS generation.

Data transfer at 4.5Mb per second in data streaming mode is defined by specifying PROTOCL=S4 on the IOCP CNTLUNIT macro during IOCDS generation. Activation of the updated IOCDS requires a processor complex power-on-reset. Additional information on the use of this parameter can be found in TNL SN22-5138 to the *3090 Processor Complex - IOCP User's Guide and Reference*, SC38-0038.

Support for the new 'S4' parameter on the IOCP CNTLUNIT macro under MVS is provided by the fix for APAR OY03274.

IOCP APAR OY03274 - This APAR provides support for the definition of 64 channel paths/side and 4.5Mb/sec protocol. The following text is intended to correct/clarify the APAR "header" text:

- Support is provided for the definition of 4.5Mb/sec control units on BOTH the 3090 "previously available" as well as "enhanced" E-models (the APAR text only indicates 3090-E models). This APAR should be installed on all systems which have FMID HIO1104 installed. Although FMID HIO1104 includes support for both the 308X (ICPIOCP) and 3090 (IOPIOCP), the function of 308X IOCP is NOT affected by this APAR.

- The function provided by this APAR is "synchronized" with the function contained in 3090 MEC 227303. This MEC is installed as part of SEC 223630B. (The MEC number designates the level of standalone IOCP. The installed MEC level can be viewed from the ECSTAT frame on the system console.)

- While it is recommended to keep the "standalone" and "SCP" versions of IOCP "in sync," system operation should not be adversely affected if the two levels are not in sync except for the following considerations:

  1. If APAR OY03274 is installed before SEC 223630B, standalone IOCP cannot modify or generate reports for IOCDS's which have been generated by the SCP's IOCP subsequent to the application of this APAR.

  2. If SEC 223630B is installed before the APAR, the SCP version of IOCP cannot generate reports for IOCDS's which have been created by standalone IOCP subsequent to the installation of SEC 223630B.

- Regardless of the order in which the APAR and SEC 223630B are installed, the Power-on-Reset (POR) and SYSIML ('O8' from OPRCTL frame) functions can still be performed using IOCDS's created:

  - BEFORE the installation of the APAR and/or SEC

  - AFTER the installation of the APAR (with SEC not yet installed)

    Note: This does not allow for the using of the new functions provided by SEC 223630 (4.5Mb or 3rd Channel Group support).

  - AFTER the installation of the SEC (with APAR not yet installed)

PCE Stall Detect

When a Service Processor Damage machine check occurs, it is important to understand the state of the processor complex.

- There is no active PCE.

- There is no backup PCE.

- Normal system monitoring and recovery are not available.

- MVS/XA will stay up and running for an indeterminate period of time.

If PCE Stall Detect is enabled (OPRCTL frame), notification of this failure will be through MVS/XA message:

IEA470W The Processor Controller has failed. Some critical system functions have failed. An orderly shutdown of the system should be immediately attempted in order to minimize the impact of this failure.

The system no longer has:

- Any monitoring for power/thermal conditions
- Ability to record any information about errors
- Any hardware error recovery
- Any MVS recovery for hardware errors

MVS/XA will continue to operate until functions of the Processor Controller are required. Error recovery requires the PCE. Should an error occur when the 3092 is not operational, information about the error will not be recorded for future analysis. Errors that are normally recoverable through the services of the 3092 will bring the system down.

Functions Not Available

Operator actions that require the services of the Processor Controller can no longer be performed:

- IPL
- Restart
- POR
- START and STOP
- Dynamic or Static Partitioning

There are some commands that should not be used after a PCE Stall Detect. If issued these commands will not complete and will stay Enqueued on certain MVS resources that are used for serialization for reconfiguration and display commands (i.e., SYSZMSF, SYSZVARY PATH, SYSZVARY CPU).

Do not use the following commands:

- CONFIG commands
- D M=
  - STOR
  - ESTOR
  - SIDE
  - IO
  - CPU
  - CHP
  - CONFIG(nn)
- QUIESCE

The remainder of the D M command options will function normally provided the above commands are not issued.

QUIESCE will cause an immediate outage. The only way to get out of a "quiesced" state is to Restart the system and the Restart function is not available.

Considerations after a PCE Stall

If PCE Stall Detect is enabled or disabled, the same two options are available; shut down or keeping running. If disabled, the major difference is that MVS/XA does not notify you of the failure. The failure would not be recognized until the operator noticed that the SAD was not being updated, or that the system console did not respond or that the PCE status line contained an error indicator. In some cases it could be hours before the failure is detected.

Considerations:

- Shut down
  - Data validity - If all applications are brought down in an orderly manner, all the records currently in storage will be written to the data base and its status will be known.
  - After shut down, no further application/data recovery required.
  - Good time to repair PCE - It's 2:00AM, let's get the processor controller fixed before 8:00 when my heavy workload starts.
  - Backup system available for critical workload.
- Keep running
  - It's 2:00 in the afternoon and I can shut down at 6:00PM.
  - The processor complex has not experienced an error in 17 months.
  - Parts not available for 2 hours.
  - There is a deadline to meet.
  - Backup system currently running higher priority work.
  - Work in progress or long running jobs have a chance of running to completion.

If you decide to keep running you have a better chance of succeeding if you avoid any commands that try to communicate with the PCE.

- Don't issue CONFIG or D M commands.
- Don't Quiesce for any reason.
- SLIP Traps should be deactivated. Should a SLIP trap spring, the system will go into a wait state and no Restart can be performed.
- Do answer DCCF messages quickly. They must be answered prior to the 125 second timeout which puts the system into a restartable wait state, which requires the PCE in order to restart.

Chapter 5. Partitioning Topics

Partitioning Considerations

Partitioning is the reconfiguration of a 3090 Model 400/400E or 600E processor complex from a Single Image configuration to a Physically Partitioned configuration. It involves deconfiguring all of the elements associated with a particular side, in order that the side may be used for some other purpose. The granularity of partitioning is on a side basis and can be accomplished either statically or dynamically.

Static partitioning is the process of reconfiguring a processor complex from single image configuration to physically partitioned configuration without System Control Program involvement. Static partitioning is disruptive to both sides of the processor complex leaving each in a POR required state. Dynamic partitioning is the process of reconfiguring a processor complex from a single image configuration to a physically partitioned configuration with System Control Program involvement[26]. Dynamic partitioning is non-disruptive to one-half of the processor complex (on-going side) and disruptive to the other half of the processor complex (off-going side, leaving it in a POR required state).

The following are situations where partitioning might be used:

1. **Partitioning for a test system**. In this situation, one half of the system is to be used as a test system after prime production hours. This would normally be a planned partitioning event and dynamic in nature.

2. **Partitioning to provide a backup system**. In this situation, one half of the system is to be used to back up an existing processor in in the complex. This would normally be an unplanned partitioning event, dynamic in nature, occurring only when one of the other systems in a complex fails.

3. **Partitioning as part of single image recovery**. In this situation, a failure occurred during single image operation and the system cannot be re-initialized in single image mode (due to a hardware failure associated with one of the sides). This would be an unplanned partitioning event, static in nature.

4. **Partitioning for maintenance**. Depending on the situation, this could either be a planned (deferred maintenance) or unplanned (needs to be fixed now) event. In either case, the event would normally be dynamic in nature.

---

[26] MVS/XA 2.1.7 and 2.2 are the only IBM System Control Programs which support dynamic partitioning

Planning is essential in all cases involving partitioning. From a total system workload perspective, there are essentially two choices:

1. Continue to run the total workload on one-half of the system at a reduced level of service for all applications.

2. Run only the critical applications at an acceptable level of service.

If your choice is the latter, then clearly it is important that someone identify critical versus non-critical applications.

The I/O configuration must be carefully planned to ensure paths to critical I/O devices exist from both sides of the system. This may entail installing additional two/four channel switch features on those devices having this option or installing additional switching units (such as the IBM 3814 or 2914).

Devices supporting multiple active interfaces (DASD, tape, 3725 w/2-processor switch, etc...) should have one or more paths configured from each side, and operate with at least one active path from each side. Devices supporting a single active interface (graphics, unit record, 3725 w/o 2-processor switch, etc...) should have a path configured from each side, with the active path controlled through a switching unit (such as the IBM 3814 or 2914). The master and alternate console should be configured across sides and should also consist of a dedicated control unit/display combination (a single display on the control unit). The reason for this dedicated combination is to guarantee access to these consoles during recovery processing using DCCF (Disabled Console Communication Facility).

The proper operating system controls must be enacted to ensure success of the dynamic partitioning process. This involves specifying the correct value for the RSU= parameter in the IEASYSxx member of SYS1.PARMLIB, and may also involve updating of the program properties table for certain applications. Both of these affect the ability of the system to configure central storage offline during the dynamic partitioning process.

Last, but not least, operational procedures must be written, tested, and practiced to ensure that operational personnel are familiar with the process.

The length of time it takes to dynamically partition the system is workload and configuration dependent. The greater the system activity and the more complex the configuration, the longer it will take the system to isolate the elements (channel paths, processors, and processor storage) of the off-going side so they can be taken offline. The elements belonging to the off-going side must be offline before the system can be dynamically partitioned.

The MVS/XA CONFIG commands to isolate the elements of the off-going side can be entered individually by the operator or can be taken from a SYS1.PARMLIB member (CONFIGxx). The recommended order in which elements should be taken offline during dynamic partitioning is channel paths, processors, expanded storage, central storage. For a dynamic merge the reverse sequence should be used to bring the elements back online.

Channel Paths

In the case of channel paths, the operator will have to deal with locally attached devices supporting only a single active interface (3x74 control units, unit record devices, graphics terminals, etc...) currently online to the off-going side. Under normal operation the operating system does not allow the last path to online or 'online and allocated' devices to be taken offline. While there are command options to override this, their use must be carefully thought out because of the potentially disruptive effect they could have on the application.

It is improbable that single-path devices configured to the off-going side and in use by an application can be configured to the on-going side without being disruptive to the application. For a planned dynamic partitioning event (one in which the off-going side is preordained) this is not a major problem, as the active paths for these devices can be configured to the on-going side. If the opposite situation should arise however, and the only active path is configured to the off-going side, the application(s) using these devices will have to be taken down in order to take the channel paths offline.

Devices supporting multiple active interfaces (DASD, tape, 3725 w/2-processor switch, etc...) do not present a problem, provided there is at least one active path from each side of the system.

Processors

Processors should go offline fairly quickly. Exceptions include the following:

- Active jobs in the system running with affinity to an off-going processor.
- Vector jobs are running, and taking the target processor(s) offline would remove the last Vector Facility in the system.

In both of the above cases, messages (IEE718I for jobs with affinity and IEE176I/IEE177D for Vector) will be issued informing the operator of the situation and requesting action.

Processor Storage

In order to dynamically partition the system, all processor storage belonging to the off-going side (both expanded and central) must be configured offline. This requires that these areas be cleared of any pages before the storage can go offline.

Pages residing in the expanded storage arrays of the off-going side are migrated to auxiliary storage. This reconfiguration migration process involves the use of central storage, hence the recommendation that expanded storage should be CONFIGured offline before central storage. Migration involves first transferring the pages to central storage and then writing the pages to auxiliary storage using the Auxiliary Storage Manager. No attempt is made to locate free slots in the expanded storage of the on-going side and transferring the pages to these slots.

For central storage, the operation is more complex and in reality begins during the MVS/XA initialization process at IPL. Dynamic partitioning, as imple-

mented on 3090 partitionable processors (as well as 3084 processors) requires that all central storage belonging to the off-going side be taken offline before the dynamic partitioning process can complete. Furthermore, in order to take the storage offline, it must be kept free of long term fixed pages. Long term fixed pages include SQA pages, common area fixed pages, and LSQA or private area fixed pages associated with non-swappable jobs. Thus two things are needed; an indication of the amount of storage to be kept available for reconfiguration, and internal processing capability to keep this area free of long term fixed pages.

The Reconfigurable Storage Unit (RSU) parameter identifies the amount of storage to be kept available for reconfiguration. The value specified is in terms of storage increments. The size[27] of a storage increment is processor model dependent.

During the processing of the RSU parameter at system initialization, MVS/XA assigns the number of storage increments specified in the RSU parameter to non-preferred[28] status (non-preferred for long term page fixes for non-swappable jobs). The remainder of the storage increments are assigned preferred status (preferred for long term page fixes for non-swappable jobs). Both preferred and non-preferred storage is used to satisfy normal page allocation requests and requests for short-term page fixes.

The important distinction between 'preferred' and 'non-preferred' storage is that requests by non-swappable jobs for long-term page fixes are always assigned to 'preferred' real storage frames. The intent here is to ensure 'non-preferred' storage will be available to take offline when storage reconfiguration is required. 'Non-preferred' storage frames contain pages that are not fixed and can be paged out to free the backing real storage, pages that are considered to be 'short-term' fixed and will be freed in a short amount of time or, if long-term fixed pages reside in these frames, the job(s) holding the long-term fixed frame(s) can be swapped out thus releasing the frames.

When MVS/XA has isolated the amount of central storage required to support the reconfiguration, the physical storage and the address ranges assigned to the storage can be varied offline both logically and physically.

Specifying the RSU Parameter

The RSU parameter is specified in the IEASYSxx member of SYS1.PARMLIB. The default value assigned is '0'. Use the following formula to determine the value to be specified.

$$RSU = \frac{\text{Total Amount of Installed Storage}}{2 * \text{Storage Increment Size}}$$

---

27 The storage increment size on the 3090 Model 400 is 2MB and on the Models 400E/600E is 4MB.

28 Non-preferred storage is also called reconfigurable storage.

Prior to MVS/SP 2.1.1, requests for reconfigurable storage (RSU=x) were satisfied from the storage address ranges available (configured) at IPL-time. If RSU=0 was specified or defaulted, then no reconfigurable storage was allocated. However, if the user later varied storage online, the additional storage was automatically marked as non-preferred (reconfigurable). This implementation allowed the user to initialize the processor complex in partitioned mode with RSU=0 specified. When the processor complex was subsequently dynamically merged for single image operation, the storage belonging to the 'on-coming' side was marked 'non-preferred'. From then on, the processor complex could be dynamically partitioned just as though the proper RSU= value had been specified. The requirement that one-half of the total installed central storage be defined as non-preferred to support dynamic partitioning was satisfied transparent to the user.

MVS/SP 2.1.1 changed the way reconfigurable storage was assigned. The assignment of reconfigurable storage is now treated 'globally' across ALL INSTALLED STORAGE regardless of the mode of operation. As a result of this change the proper RSU value MUST be specified if dynamic partitioning is to be performed. A value of '0', either specified or defaulted, will cause MVS/XA to designate all installed storage as preferred. As a result, the installation will not be able to dynamically partition the processor complex. Failure to specify the proper amount of reconfigurable storage (one-half of total installed central storage) will also result in not being able to dynamically partition the processor complex.

The proper value depends on the 3090 model (which determines the physical Storage Increment size) and the amount of Central Storage installed. The following table defines the correct values for each model/CS configuration to support dynamic partitioning:

3090 Model	Total CS Installed	SI Size	RSU Value
400	128 MB	2 MB	32
400E	128 MB	4 MB	16
400E	256 MB	4 MB	32
600E	128 MB	4 MB	16
600E	256 MB	4 MB	32

For 3090 models 120E, 150/150E, 180/180E, 200/200E, 300E, and Models 400, 400E, and 600E operating ONLY in partitioned (PP) mode, the RSU parameter should be specified as "0". These processor complexes are not partitionable and therefore do not require that reconfigurable storage be reserved for dynamic partitioning

For additional information on the RSU parameter, the reader should refer to *MVS/Extended Architecture SPL: Initialization and Tuning, GC28-1149* and *MVS Extended Architecture Planning: Recovery and Reconfiguration, GC28-1160.*

The requirement to always specify an RSU value for dynamic partitioning purposes may appear to penalize a side operating in partitioned mode for long periods of time. In reality it does not. The non-preferred storage is allocated LOGICALLY as though the total amount of installed storage were available. On a subsequent dynamic merge, the storage belonging to the on-coming side will be automatically marked non-preferred and the RSU requirement will be satisfied.

Message IAR005I: Message *IAR005I PREFERRED AREA HAS EXPANDED. RECONFIGURABILITY MAY BE IMPAIRED.* is issued when the MVS/XA system converts a storage unit initially set up as non-preferred to preferred status. This could occur when the condition arises that a request for a long-term page fix requires storage space but the preferred area is full. This conversion from non-preferred to preferred (also known as dynamic expansion of the preferred area) lowers the number of reconfigurable storage units specified by the RSU parameter. Should this occur, there is a real possibility that the system cannot be dynamically partitioned. Once a storage unit has been converted from non-preferred to preferred status, it remains in that state until the system is re-initialized (IPLed). The operator should report any occurrence of this message to the system staff so that appropriate follow-up action may be taken.

Message IEE575A: During the dynamic partitioning process, the operator may see MVS/XA action message *IEE575A CONFIG (STORAGE|ESTOR) WAITING TO COMPLETE _ REPLY C TO CANCEL* appear and then be internally cancelled in a short period of time (typically less than a minute). The message may be displayed several times as the system takes central storage offline. If the message remains outstanding for a long period of time, it usually means that the system has encountered a fixed page, which cannot be freed, in an area of non-preferred storage that is being taken offline to satisfy the partitioning request. This could occur because a job started out as swappable, short term fixed some pages (which could end up in the non-preferred area of storage), and then made itself non-swappable. Another situation where this could occur is with outstanding tape mounts during job initiation. Pages could have been fixed (short term) in the non-preferred area of storage during job initiation. While the tape mount is pending, the job is set non-swappable. Since the job is non-swappable there is nothing the system can do to free the page. Only mounting the tape(s) or cancelling the job will free up the page or pages.

The job or ASID owning the particular frame can be identified with the 'D M=STOR' command. Once the job owning the page has been identified then action can be taken. It may be necessary to cancel the job or take other action (such as mounting the tape) in order for storage reconfiguration to complete. If cancelling the job is not desirable, then MSG IEE575A should be replied to with a 'C' to terminate the storage reconfiguration process. The appropriate 'CONFIG STOR(E=x), ONLINE' command should then be executed to return any storage that was taken offline. The operator should document the name of the job holding the storage and give it to the system programming staff. It will probably be necessary to include this job in the Program Properties Table or take other appropriate action to ensure that requests for fixed pages go into the preferred area of storage, so as not to affect the dynamic partitioning process.

Note: The 'D M=STOR' command must be issued before replying to MSG IEE575A in order for the jobname and ASID of the job holding the page to be identified.

Message 64461 (on system console): If message *'B1 STOPS SCP AND CAUSES POR REQD. ENTER X1 TO CONFIRM, B1 TO DESELECT. 64461'* appears on the system console during dynamic partitioning, it means that at least one element from the off-going side has not been deconfigured. ***A response of 'X1' to this message will cause a system wide outage. It will result in termination of the SCP on the on-going side.*** The operator should enter a D M=SIDE command to determine which elements have not been deconfigured and to reissue the appropriate MVS/XA CONFIGuration commands for those elements.

Note: Depending on the EC level of the processor complex the text for the above message may also appear as *'ENTER X1 TO CONTINUE (STOPS SCP AND CAUSES POR REQD) OR B1 TO CANCEL. 64461'*.

Message 64449 (on system console): During dynamic partitioning, message *'THE ACTIVE IOCDS IS ONLY VALID ON THE OFF-GOING SIDE. 64449'* may be displayed on the system console. This message indicates that the active IOCDS (the one used during the last power-on-reset) belongs to the off-going side. No access to this IOCDS is possible from the on-going side. Normal system operation is not affected as the I/O configuration contained in the IOCDS was loaded into the HSA during the last power-on-reset. In this situation, a new IOCDS must be selected (one accessible from the on-going side) if the on-going side should require a power-on-reset before the system has been reconfigured for single image operation.

Program Properties Table

While the system normally attempts to assign long-term fixed storage requests to 'preferred' storage frames when the job making the request was initiated as 'non-swappable', some jobs, executing with 'authorized' status, are initiated as 'swappable' and later issue a SYSEVENT to mark themselves as 'non-swappable' for a short period of time. Normally this does not present a problem even if 'long-term' page fixes are assigned to 'non-preferred' storage since the job will shortly mark itself 'swappable' again, and the storage frames used to back the 'long-term' fixed pages can be freed by swapping out the job.

Jobs may be encountered, however, that make requests for short-term fixed storage and then mark themselves non-swappable for long periods of time. These requests may be allocated to non-preferred storage (the job was swappable when the request was made). Since the frames cannot be freed by paging them out or swapping out the job (the job has made itself non-swappable), storage reconfiguration may not be able to complete, thus impacting the systems ability to complete the dynamic partitioning process.

This situation can be resolved by setting the appropriate flag bits in the program properties table or changing the job or application to issue the TRANSWAP SYSEVENT. Information on program properties table settings and TRANSWAP SYSEVENT can be found in *MVS/Extended Architecture System Programming Library: System Modifications, GC28-1152.*

DFHSM (Data Facility Hierarchical Storage Manager)

Under normal circumstances, DFHSM is initiated as a swappable address space and LSQA pages may be allocated to non-preferred storage frames. If a SETSYS NOSWAP command is issued as part of the DFHSM start-up procedure, the DFHSM task is set non-swappable. A later attempt to configure storage offline as part of dynamic partitioning may fail if non-preferred storage has been polluted by DFHSM fixed frames.

This problem is addressed by DFHSM APAR OY01177, which changes the SYSEVENT in ARCCPSET to default to TRANSWAP. TRANSWAP will initiate a swap of the address space, cause it to be marked as non-swappable, and on the subsequent swap-in, fixed pages will be acquired from the preferred area of storage.

Partitioning the 3090 with an MSS Attached

If you plan to reconfigure a 3090 to partitioned or single-image mode, and if your 3090 has two channel paths to the MSS defined as an MP in MSS Table Create, you will need the maintenance for APARs OZ75779 and OZ76112 installed on the MSSE/XA component of your MVS/XA system. Without this maintenance, unpredictable results can occur when reconfiguring from one mode to the other.

The "Z S,CONFIG" command must be issued for each dynamic partition/merge sequence. additional information on this command are shown below. The command results in an outstanding WTOR message (ICB540A) that the operator must respond to during the partitioning or merging sequence. An outline of the sequence of events follows:

Partitioning	Merging
- Start the partitioning process (MVS CONFIG commands)	- Issue 'Z S,CONFIG' command
- After CHPIDs are offline issue 'Z S,CONFIG' command	- Outstanding WTOR (ICB540A)
- Outstanding WTOR (ICB540A)	- Reply 'A' to message ICB540A
- Reply 'D' to message ICB540A	- Wait for message ICB542I (MSS ASSOCIATE COMPLETE, BOTH INTERFACES ASSOCIATED)
- Wait for message ICB543I (MSS DISASSOCIATE AND PURGE COMPLETE FOR OTHER PROCESSOR)	- Start the 3090 merge process (MVS CONFIG commands)
- Continue with partitioning process	

Figure 31. 'Z S,CONFIG' Command Example

MSS Command - Z S,CONFIG

The maintenance for APAR OZ75779 allows the operator to control the sequence of the MSS ASSOCIATE and DISASSOCIATE orders required for the MSS to change the status of the two 3090 host connections. This operator control was made available by adding a new parameter to the 'HALT S' (Z S) command. The syntax of the command is Z S,CONFIG or Z S,CF. The following message will appear in response to the command:

```
ICB540A MSS RECONFIGURATION. REPLY A, D, OR U TO ASSOCIATE,
        DISASSOCIATE, OR IGNORE
```

There are three new messages associated with the new command:

```
ICB541I MSS DISASSOCIATE NOT ALLOWED, PATHS NOT RECONFIGURED

ICB542I MSS ASSOCIATE COMPLETE, BOTH INTERFACES ASSOCIATED

ICB543I MSS DISASSOCIATE AND PURGE COMPLETE FOR OTHER PROCESSOR
```

With this new command, when operators want to reconfigure from single-image to partitioned mode, they perform the normal sequence of MVS CONFIG commands to vary processors, storage and channel paths offline. After all channel paths are configured offline, they use the Z S,CONFIG command to issue the DISASSOCIATE (and an imbedded PURGE) order to the MSC.

When the operators want to reconfigure from partitioned to single-image mode, they use the Z S,CONFIG command to issue an ASSOCIATE order before any channel paths are configured online. The operators can then start the 3090 merge process by issuing the normal sequence of MVS CONFIG commands. See the sequence of commands in Figure 31 on page 117.

Note: Assuming the 3090 has two channel paths to the MSC, and these paths have been defined as an MP in MSS Table Create, there are two occasions when the new command will cause unpredictable results:

1. When running in single-image mode, and one of the channel paths to the MSS is configured offline, a DISASSOCIATE order must *not* be issued.

2. When running in partitioned mode, and only one channel path is online to the MSS, an ASSOCIATE order must *not* be issued.

MSS EC Requirements (MVS/XA Operation)

When running a 3850 Mass Storage (MSS) in an MVS/XA environment, the 3850 should be at E/C 737576 with microcode patch S81809 installed.

Chapter 6. Virtual Storage Assessment

Virtual Storage Review

Virtual Storage is a resource that is required to install and execute programs. Virtual storage assessment is an essential part of planning an IBM 3090 installation. In each system there is a virtual storage requirement for each program, including data. The virtual storage requirement for the system must account for all the programs and data that must be directly addressed. For the nucleus and common area of virtual storage (SQA, LPA, CSA), program and data requirements are generally additive. For the Private area of virtual storage, the largest Private area determines the virtual storage requirement. See diagram below.

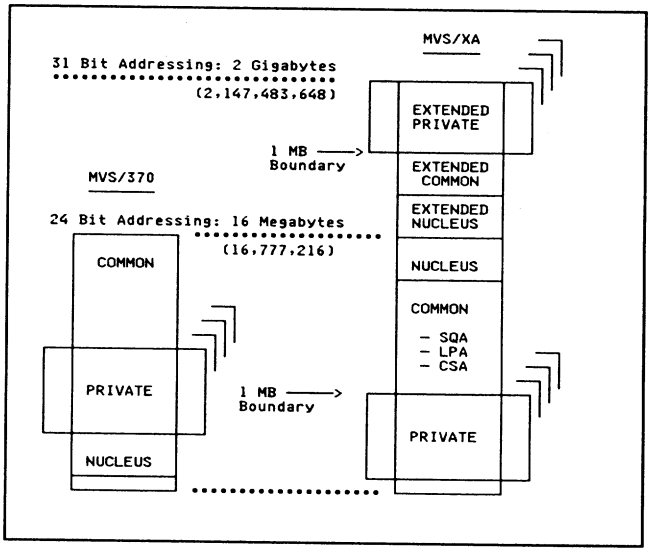

As programs, data, and subsystems grow, and as more subsystems are combined in one MVS system, the addressing limits may be reached. The 370-XA 31-bit addressing provides the framework for extending the addressing limit to 2 gigabytes -- from 16 megabytes in S/370 24-bit addressing. Next it is up to the programs to implement the new 31-bit addressing. New (31-bit) and old (24-bit) programs may coexist but must communicate to each other in 24-bit addressing. MVS/XA and MVS/XA DFP provide some virtual storage constraint relief by using less of the common area that is "below the 16Mb line" -- i.e., in the 24-bit range of addressing. The relatively new program, DB2, is almost entirely above the 16Mb line. VTAM Version 3 provides considerable VSCR (Virtual Storage Constraint Relief). MVS/XA DFP 2.1 provides considerable relief by allowing a separate Private address space for catalogs and reducing the amount of PLPA and CSA (SPK0) that is required. Additional programs providing VSCR are listed at the end of this section. New compilers offer the potential for application programs to exploit the 31-bit addressing.

The actual amount of virtual storage constraint relief that an installation realizes upon installing one of these products may be environment dependent. If previous virtual storage tuning has been done, for example, less than the documented relief

may be attained. The installation should evaluate the relief characteristics of each product in light of the specific environment.

There are two basic approaches to the implementation of virtual storage constraint relief. One is a vertical approach -- the implementation of 31-bit support, starting with MVS/XA and MVS/XA DFP. Relief is attained by moving code and data above the 16Mb line. The other approach is a horizontal approach. Relief is attained horizontally by splitting code and data from one Private address space into multiple Private address spaces. The horizontal approach applies to MVS/370, as well. An important example of the horizontal approach has been implemented by installations using CICS. There are a number of ways the split has been performed. A popular way is splitting off a terminal owning region and using the MRO facility to communicate to the CICS code and data in the other CICS address space(s). In some cases this has made the code easier to maintain and more reliable. CICS also provides vertical relief. IMS/VS using MSC, can provide relief by the horizontal approach by separating the Data Communication code from the Data Base code activity. This is commonly known as a FRONTEND / BACKEND split, and reduces the size of the private area requirements.

Virtual Storage Considerations

As program function increases and workload activity increases, virtual storage requirements increase. Virtual storage requirements should be tracked and future requirements should be estimated. This information can help plan when to install products that provide VSCR. Running out of virtual storage is more severe than running out of most other resources. In most cases running out of a resource means reducing the rate of doing work at the installation. Running out of virtual storage may mean that the system stops or will not IPL. Virtual storage constraint can also prevent the utilization of the full capacity of a large IBM 3090 processor complex.

Planning for an IBM 3090 system installation should include virtual storage tracking and virtual storage projections **prior** to the installation of the processor complex.

A virtual storage analysis is especially important when workloads from two or more systems are to be merged. Particular attention should be given to estimating the **total** workload that is planned for execution on the processor complex. For a better understanding of virtual storage alternatives, product-specific information, and tuning techniques, see *Virtual Storage Tuning Cookbook*, G320-0597.

For help in making an assessment of the virtual storage requirements for large processor installations, refer to *Virtual Storage Considerations When Combining Workloads on Large Processors*, GG22-9412.

Virtual storage requirements may also be a factor in backup scenarios and must be included in developing the backup plans.

From a system tracking point of view, consider in particular:

- SQA Overflow
  - To minimize CSA fragmentation: allocate enough SQA.
  - To minimize the size of common: do not allocate more SQA than necessary.
  - Recommendation: allocate SQA so that it overflows into CSA up to 64Kb (1 segment).
- CSA Unallocated - Allow enough unallocated CSA to accommodate fragmentation and error recovery.
- Know what is in PLPA and why it is there.
- CSA usage by subsystem (generally, Storage Protect Key)
  - SPK = 0 Supervisor, ... RACF, etc..
  - SPK = 1 JES
  - SPK = 2 VSPC
  - SPK = 3 not used by IBM
  - SPK = 4 not used by IBM
  - SPK = 5 Data Management
  - SPK = 6 TCAM, VTAM NCCF, NPDA, NLDM
  - SPK = 7 IMS
  - SPK = 8 V = V USER PROGRAMS
  - SPK = 9-F V = R USER PROGRAMS
- The installation may want to limit the size of the Private region that users can request. This can be done by supplying a routine in the SMF IEFUSI exit to define region size and GETMAIN limits for the Private region both below and above the 16Mb line. Reference: *MVS/XA SPL: System Modifications*, GC28-1152.
- Identify each Private requirement with less than 2Mb available. This 2Mb amount would allow moving the line 1Mb and still having 1Mb available.
- Learn what to expect in the rise and fall of virtual storage requirements during the daily cycle. Also compare virtual storage usage shortly after an IPL to usage a "long time" after an IPL.
- Take a close look at virtual storage requirements just before and just after installing a new program, product or release level.
- Track subsystem (IMS, CICS, ...) private address space usage. Work with the subsystem specialists regarding Private and CSA requirements, as appropriate.

Some Products Providing VSCR

PRODUCT	V.R.L	NUMBER	AVAILABILITY	
MVS/SP (JES2)	2.1	5740-XC6	The Base for 31-bit Products.	
MVS/SP (JES3)	2.1	5665-291	The Base for 31-bit Products.	
MVS/XA DFP	1.1	5665-284	(DFP co-requisite.)	
	2.1	5665-XA2	2/04/86	
	2.2	5665-XA2	3/04/86	
	2.3	5665-XA2	3Q87	
ACF/VTAM	3.1	5665-289	5/10/85	
APL 2	1.2	5668-899	12/06/85	
ASSEMBLER H	2.1	5668-962	3/31/83	
CICS	1.6	5740-XX1	12/31/82	
	1.7	5740-XX1	12/20/85	
COBOL II	1.	5668-958	12/28/84	Compiler and Library
COBOL II	1.	5668-940	12/28/84	Compiler only
DB2	2.	5740-XYR	3/07/86	
DFSORT	8.	5740-SM1	3/28/86	
DFDSS	2.2	5665-327	3Q86	
FORTRAN	4.	5748-F03	10/26/84	Compiler and Library
FORTRAN	4.	5748-LM3	10/26/84	Library only
GDDM	4.	5748-XXH	10/16/84	
IMS	1.3	5740-XX2	6/29/84	
IMS	2.1	5665-332	3/28/86	
ISPF	2.2	5665-319	10/01/85	
ISPF/PDF	2.2	5665-317	10/01/85	
JES2	2.1.5	5740-XC6	12/27/85	
JES3	2.1.5	5665-291	12/27/85	
NCCF	2.2	5665-316	6/07/85	
NLDM	3.	5668-971	6/07/85	
NPDA	3.2	5665-321	6/07/85	
PL/I	5.1	5734-PL1	10/01/85	
QMF	2.	5668-972	6/28/85	
RACF	1.7	5740-XXH	12/27/85	
RMF	3.4.1	5665-274	3Q86	(supports MVS/SP2.1.7)
TSO/E	2.1	5665-285	12/18/84	

Some Product Highlights (MVS/XA Systems)

MVS/XA DFP 2.1 provides 540Kb of VSCR in PLPA, and approximately 20Kb of VSCR per ICF catalog (CSA, SPK0). The common VTOC access facility is moved to the ELPA (Extended Pageable Link Pack Area), and ICF catalog management modules and data are moved to a separate address space.

MVS/XA DFP 2.3 provides approximately 350Kb of VSCR compared to DFP/XA 2.2. The source of this relief in the movement of modules from LPA to ELPA.

MVS/XA 2.2.0 and MVS/XA DFP 2.3 have been changed to support the Scheduler Work Area (SWA) usage of virtual storage above the 16Mb line. The expected VSCR is 200 to 800Kb, depending on the installation configuration.

VTAM V3 moves all buffers (except IOBUF), most modules, and most dynamic storage areas above the 16Mb line. For CICS or IMS, the CSA SPK6 for VTAM is significantly reduced.

JES2 2.1.2 obtains storage for some RJE device control blocks dynamically, and moves the IO copy of the checkpoint record to Extended Private. See *MVS/SP JES2 Virtual Storage Tuning Guidelines*, GG66-0203.

JES2 2.1.5 provides about 30Kb of VSCR in CSA, and 400Kb to 2000Kb of VSCR in Private by moving control blocks and data areas (PCEs and DCTs) above the 16Mb line. See *MVS/SP JES2 Migration Considerations*, GG66-0236.

JES3 2.1.5 moves most JES3 modules, control blocks, and major buffer pools above the 16Mb line from Private and CSA (SPK1). VSCR results are highly dependent on the environment. See *MVS/XA JES3 Conversion Notebook*, SC28-1501.

Benchmarks relative to JES3 1.3.4 have shown the following ranges for VSCR:

```
JES3 Global Private    700Kb to 5000Kb
CSA                     60Kb to 1000Kb
PLPA                    60Kb
```

CICS 1.7 provides VSCR through (1) automatic terminal installation facility, (2) RACF 1.7 security tables above the 16Mb line, (3) command level programs written in PL/I 5.1, (4) COBOL II subprograms that do not contain CICS commands or CICS structure dependencies, and (5) Assembler H V2 support. The CICS message table module, trace table, and dynamic transaction backout buffers are above the 16Mb line. There may be some increase in the virtual storage size of CICS modules and tables.

IMS 1.3 moves the Message Format Services pool above the 16Mb line. DL/I (LSO = S) and DBRC are moved to additional address spaces. (No VSCR for Fast Path in this release.)

IMS 2.1 moves much of Fast Path above the 16Mb line. Fast Path VSCR is expected to range from 1.5Mb to 6.3Mb. There is a small increase in virtual

storage for non-Fast Path IMS 2.1. See the Storage Estimates section of the *IMS/VS Version 2 System Administration Guide.*

With DFDSS 2.2 the buffers for COPY, COPYDUMP, DEFRAG, DUMP, PRINT, and RESTORE may be placed above the 16Mb line -- reducing the calling program's Private area.

For DFHSM 2.2.1 the minimum recommended Private area is 4Mb. There is no increase in CSA required for this release.

RACF 1.7 increases PLPA slightly (+ 75Kb). CSA (SPK0) requirements are reduced by moving data and index buffers to ECSA.

RMF has most data and modules above the 16Mb line.

ISPF 2.1.2 provides VSCR of 342Kb to 381Kb in PLPA over the previous release.

ISPF/PDF 2.1.2 provides VSCR of up to 363Kb in PLPA over 2.1.0. Note that VS requirements prior to ISPF 2.1.2 and ISPF/PDF 2.1.2 had generally been increasing.

TSO/E 2.1.2 provides VSCR of 155Kb to 205Kb over R1, 200Kb to 255Kb over R2.1, and 200Kb to 350Kb over R2.1 with session manager.

GDDM R3 or later has all 250Kb of code in ELPA.

Chapter 7. Bibliography

Related Reference Materials

The following reference materials are either referred to in this Technical Bulletin or contain further information on a topic described in this document.

IBM 3090 Processor Complex Documentation:

- *IBM System/370 3090 Processor Complex Installation Manual - Physical Planning*, GC22-7074

  This publication contains information necessary for the planning and physical installation of the 3090 Processor Complex. It is intended to be used by installation planning and marketing personnel and by customers in planning their installations.

- *IBM 3090 Processor Complex Site Readiness Checklists*, GC22-7075

  This checklist aids installation planning representatives and customers to verify site readiness for the installation of a 3090 Processor Complex.

- *IBM 3090 Processor Complex Physical Planning Scale Template (Metric)*, GX22-7107
- *IBM 3090 Processor Complex Physical Planning Scale Template (English)*, GX22-7108

  Physical Planning Acetate Templates (3), 8 1/2 x 11 inches for planning machine room layouts. Metric scale: 10 mm = 0.5 m. English scale: 1/4″ = 1′.

- *IBM 3090 Full-Size Physical Planning Template*, SX22-7112

  This transparent template contains the one-to-one diagrams of the floor space requirements of each of the units in the 3090 Processor Complex. These templates can be used to locate within a customer's raised floor room the various machine units as well as floor cutouts and caster locations.

- *IBM 3090 Full-Size Physical Planning Template for Vector Feature*, SX22-7113

  This transparent template contains the one-to-one diagram of the floor space requirement of frame 05 of the 3090 Processor Complex. This template can be used to locate within a customer's raised floor room the optional frame 05. This template should be used with the 3090 Full-Size Physical Planning Template, SX22-7112.

- *IBM System/370 Extended Architecture Principles of Operation*, SA22-7085

  This publication provides a detailed definition of the machine functions performed by systems operating in System/370 extended-architecture (370-XA) mode. It is provided principally for Assembler Language programmers, although anyone concerned with the functional details of systems operating in 370-XA mode can find it useful.

- *IBM 3090 Processor Complex Functional Characteristics*, SA22-7121

  This publication contains a broad overview of Configuration, Function, Architecture, and Features for the 3090 Processor Complex.

- *IBM 3090 Processor Complex Channel Characteristics and Configuration Guide*, SA22-7120

  This publication contains channel operation information and channel performance and configuration considerations for the 3090 Processor Complex.

- *IBM 3090 Processor Complex IOCP User's Guide and Reference*, SC38-0038

  This publication is intended for systems programmers and service representatives who are responsible for defining, installing, and configuring the channels or the channel paths, control units, and I/O devices on the 3090 Processor Complex. It describes the MVS version, the VM version, and the stand-alone version of IOCP.

Technical Bulletins:

- *IBM 3090 Processor Complex Planning and Installation Reference*, GG66-3090

  This publication serves primarily as a composite source of planning and installation reference material for the 3090 Family of Processors. It is intended for internal and external personnel who have planning responsibilities for the installation of any of these processors. The overall design is to provide information about each of the processors and the operating systems and major products that support these processors. This publication is structured into seven major chapters: Family Comparisons, Logical Components, Software Considerations, Installation Topics, Partitioning Topics, Virtual Storage Considerations and Bibliography. This publication is a component of GBOF-2208.

- *IBM 3090 Model 200 Processor Complex Installation Notebook*, GG66-0213

  This publication contains information pertinent to the installation of a 3090 Model 200 Processor Complex. The overall design is for use both as a presentation guide and a technical reference manual. Use of this dual design permits information to be imported to the general audience, while at the same time assisting those whose responsibility it is to present this information. This publication is structured into four major sections: hardware overview, software support, installation considerations, and operational considerations. Full size copies of the visuals used in the hardware overview and operational considerations sections are provided in appendices A and B. For those who are presenters, two adjacent visuals shown on the same page indicates that both should be displayed simultaneously. This publication is a component of GBOF-2208.

- *IBM 3090 Model 400 Installation Considerations*, GG66-0252

  This publication contains information pertinent to the installation of the 3090 Model 400 Processor Complex. It is intended for internal and customer personnel who have planning responsibilities for these installations. The overall design is for use both as a presentation guide and a technical reference manual. Use of this dual design permits information to be imparted to the general audience, while at the same time assisting those who have the responsibility to present this information. This publication is structured into six major chapters: Hardware Considerations, Software Considerations, Operational Differences, Reconfiguration, Recovery Considerations and Virtual Storage Assessment. Full size copies of the visuals used in these chapters are provided in the appendices. This publication is a component of GBOF-2208.

- *MVS/XA Storage Estimates and Processor Performance*, GG22-9397

  This publication is intended to provide account teams and their customers with an awareness of the requirement for and the value of processor storage for MVS/XA on the 308X and 3090 processors. This report discusses, at a high level, how processor storage is used, how to estimate the amount of storage

required, and the performance benefits of additional processor storage for various operating environments. This publication is a component of GBOF-2208.

- *MVS/SP 2.1.3VFE Overview*, LG66-0227

  This publication describes the changes in MVS/SP Version 2 (MVS/XA) to support the 3090 Vector Facility Enhancement. These changes are contained in MVS/SP 2.1.3 Vector Facility Enhancement (MVS/SP 2.1.3 VFE). This technical bulletin contains a foil presentation as well as the text for these foils. The audience for this technical bulletin is systems programmers of installations who are installing a 3090 Vector Facility.

- *IBM 3090 Engineering/Scientific Performance*, GG66-0245

  This publication contains the results of numerous measurements of the 3090 executing jobs with high floating point content (i.e., Engineering/Scientific jobs). These measurements were made using MVS/SP Version 2 (MVS/XA) and VM/SP HPO 4.2. It discusses the factors affecting Engineering/Scientific job performance. Emphasis is placed on the 3090 Vector Facility. It also discusses the metrics of performance measurement. In particular, it explains the use of application performance criteria, rather than the MFLOPS terminology, as the basis for product measurement. This publication is a component of GBOF-2208.

- *MVS Virtual Storage Tuning Cookbook*, G320-0597

  This publication was compiled to help the systems staff in virtual storage tuning efforts. It attempts to describe how each subsystem uses virtual storage, the tools available for measuring utilization, rules of thumb for efficient storage use, and techniques for tuning the subsystem for efficient storage use. It assumes an environment of MVS and either JES2 or JES3.

- *Virtual Storage Considerations When Combining Workloads on Large Processors*, GG22-9412

  This document is intended to help the capacity planner make an assessment of the virtual storage requirements for large processor installations when combining, growing, and forecasting future workloads.

- *JES2 Virtual Storage Tuning Guidelines*, GG66-0203

  This publication is designed to assist an installation having problems with JES2's use of virtual storage. It consists of virtual storage tuning guidelines and more detailed explanation of JES2 virtual storage usage. It covers the JES2 component of MVS/SP JES2 1.3.3 (or MVS/SP JES2 2.1.1), FMID HJE2329, and the JES2 component of MVS/SP JES2 1.3.4 (or MVS/SP JES2 2.1.2), FMID HJE2330. The primary audience for this publication is assumed to be MVS systems programmers responsible for the installation, maintenance, and migration of the JES2 component of MVS/SP. This publication is a component of GBOF-2208.

- *3090 Vector Facility Support Under VM/HPO*, GG24-3095

  This publication assists experienced users in planning and using a 3090 Vector Facility with VM/SP HPO 4.2. It is a component of GBOF-2201 and is an ORDER NUMBER SUBSCRIPTION ONLY document.

Operations:

- *IBM 3090 Processor Complex Messages for the System Console*, SC38-0039

  This document lists and defines the messages displayed during the operation of the System Console for the 3090. The primary user is the System Console operator.

- *IBM 3090 Processor Complex Operator Controls for the System Console Models 120E, 150, 150E, 180, 180E, 200, 200E, 400, and 400E*, SC38-0040
- *IBM 3090 Processor Complex Operator Controls for the System Console Models 300E and 600E*, SC38-0056

  These publications include descriptions of the controls used by the operator. They are written for managers, programmers, and operations personnel.

  Note: Publication SC38-0056 is an ORDER NUMBER SUBSCRIPTION ONLY document.

- *IBM 3090 Processor Complex Operator Tasks for the System Console Models 200 and 200E*, SC38-0041
- *IBM 3090 Processor Complex Operator Tasks for the System Console Models 120E, 150, 180, 150E and 180E*, SC38-0049
- *IBM 3090 Processor Complex Operator Tasks for the System Console Models 400 and 400E*, SC38-0050
- *IBM 3090 Processor Complex Operator Tasks for the System Console Model 300E*, SC38-0054
- *IBM 3090 Processor Complex Operator Tasks for the System Console Model 600E*, SC38-0055

  These publications include descriptions of each task normally performed by a system console operator. A brief description of the 3090 Processor Complex and the 3180 Display station is also included. The primary user is the system console operator. The system programmer is the secondary user.

  Note: All of these publications, except SC38-0041, are ORDER NUMBER SUBSCRIPTION ONLY documents.

- *IBM 3090 Model 150E/180E/200E/300E: Operations Training Package - Volume I*, GG24-3163
- *IBM 3090 Model 150E/180E/200E/300E: Operations Training Package - Volume II*, GG24-3164
- *IBM 3090 Model 150E/180E/200E/300E: Operations Training Package - Volume III*, GG24-3165

  These publications provide operations training material for the 3090 Models 150E, 180E, 200E, and 300E. The material includes presentation script, class hand-outs and machine exercises. They are intended for a presenter or for individual self-study. Volume I contains the presenter's materials which may also be used for self-study. Volume II contains the student materials. This is the student package. Volume III includes the presentation foil masters and is a co-requisite of Volume I for class presentation. These publications are a component of GBOF-0421 and are ORDER NUMBER SUBSCRIPTION ONLY documents.

- *IBM 3090 Model 400/400E/600E: Operations Training Package - Volume I*, GG24-3160
- *IBM 3090 Model 400/400E/600E: Operations Training Package - Volume II*, GG24-3161
- *IBM 3090 Model 400/400E/600E: Operations Training Package - Volume III*, GG24-3162

  These publications provide operations training material for the 3090 Models 400, 400E, and 600E. The material includes presentation script, class hand-outs and machine exercises. They are intended for a presenter or for individual self-study. Volume I contains the presenter's materials which may also be used for self-study. Volume II contains the student materials. This is the student package. Volume III includes the presentation foil masters and is a co-requisite of Volume I for class presentation. These

publications are a component of GBOF-0421 and are ORDER NUMBER SUBSCRIPTION ONLY documents.

Recovery:

- *IBM 3090 Processor Complex Recovery Concepts*, GG24-3077

  This publication provides an introduction to the recovery concepts of a 3090 Processor Complex. It is intended for systems programmers and operating staff and should be used as a first introduction on how the 3090 recovery concepts is implemented. This publication is a component of GBOF-0421 and is an ORDER NUMBER SUBSCRIPTION ONLY document.

- *IBM 3090 Processor Complex Recovery Guide*, SC38-0051

  This publication is a guide to hardware recovery on systems that include the IBM 3090 Processor Complex. Recovery is the ability to do useful work on the system after a failure occurs. This publication is written for operators of the IBM 3090 Processor Complex and customer technical support personnel. A general knowledge of the IBM System/370 and operator training on the 3090 is required and assumed.

- *MVS/XA Planning: Recovery and Reconfiguration*, GC28-1160

  The emphasis of this publication is on maintaining system availability after an abnormal event. This publication is intended for the programmers and planners who develop recovery and reconfiguration procedures tailored to their installations requirements. It contains hardware and software information and guidelines needed to develop procedures that the installation can use to control the system after an error situation has resulted in a loss of system availability or any hardware unit.

- *IBM 3090 Processor Complex: RSF PA and IOPD*, GG24-3076

  This publication provides presentation material for the 3090 in the following areas: Remote Support Facility, Problem Analysis and I/O Problem Determination. It includes presentation script and foil masters. This publication is a component of GBOF-0421 and is an ORDER NUMBER SUBSCRIPTION ONLY document.

- *IBM 3090 Model 400E/600E System Recovery Procedures*, GG24-3096

  This publication provides guidelines and procedures for recovery and reconfiguration of 3090 Model 400E/600E Processor Complexes. Recovery procedures for 3880/3380 DASD are also included. Specific topics include recovering from HOT IO, MIH, and Spin-Loop conditions as well as many other abnormal situations. It is intended for systems programmers and operating staff and can be tailored for customer needs. It is a component of GBOF-2208 and GBOF-0421.

Related Education Offerings:

- *MVS/XA Operator Training*, Course Code 32150

  This SRA self study course of 10 to 15 hours is also available through the Guided Learning Centers using course code Z2150 for enrollment.

  It teaches the information necessary to monitor and control the operating system in an MVS/XA environment. Along with text instructions, the course contains Computer Assisted Instruction (CAI) compo-

nents which allow the student to practice many of the functions of this operating system in a protected environment.

This course, along with the courses listed below, form a complete operator curriculum for MVS/XA:

- 308X Hardware Operator Training (32153) or 3090 Hardware Operator Training (32151)
- Controlling JES2 (32181)
- MVS/XA: System Problem Determination (32152)

MINIMUM CONFIGURATION: Most models of the IBM Personal Computer with one double-sided disk drive and 128K minimum memory, using DOS 2.0 or higher.

AUDIENCE: Any personnel who will operate MVS/XA.

COURSE MATERIALS AND PRICES: The complete course includes one text, four diskettes, and a Personal Reference Guide. An additional PRG should be ordered for each student.

		12 MONTH CHARGE	12 MONTH RENEWAL
SR21-1241	Complete Course	950.00	715.00
SR21-1243	Additional PRG	10.00 purchase	
Z2150	Guided Learning Center	390.00	

- *3090 Hardware Operator Training*, Course Code 32151

This SRA self study course of 10 to 16 hours is also available through the Guided Learning Centers using course code Z2151 for enrollment.

The course teaches the information necessary to operate the 3090 Processor Complex. The course contains a 3090 System Console simulator that allows the student to practice many of the functions of the 3090 on an IBM Personal Computer.

MINIMUM CONFIGURATION: An IBM Personal Computer, PC XT, PC AT, or 3270 PC with a minimum of 128K memory, one double-sided diskette drive, DOS 2.0 to 3.2, and a color or monochrome display.

AUDIENCE: Operators of the IBM 3090 system.

COURSE MATERIALS AND PRICES: The complete course includes four diskettes and one Personal Reference Guide.

		12 MONTH CHARGE	12 MONTH RENEWAL
SR21-1219	Complete Set	950.00	715.00
SR21-1221	Additional PRG	10.00 purchase	
Z2151	Guided Learning Center	390.00	

- *MVS/XA: System Problem Determination*, Course Code 32152

  This SRA self study course of 10 to 15 hours is also available through the Guided Learning Centers using course code Z2152 for enrollment.

  It introduces a system problem determination methodology. It is designed to support MVS/XA when used on a 308X or 3090 processor complex. The course contains Computer Assisted Instruction (CAI) components, which simulate problem determination where a student is asked to determine if the problem is caused by hardware or software.

  MINIMUM CONFIGURATION: An IBM Personal Computer, PC XT, PC AT, or 3270 PC with a minimum of 128K memory, one double-sided diskette drive, DOS 2.0 to 3.2, and a color or monochrome display.

  AUDIENCE: This course is intended for technical operations support personnel at the central site who are responsible for doing problem determination in the MVS/XA - 308X/3090 and related hardware and software environments.

  COURSE MATERIALS AND PRICES: The complete course includes a text, five diskettes and a Personal Reference Guide (PRG). An additional PRG should be ordered for each student.

		12 MONTH CHARGE	12 MONTH RENEWAL
SR21-0752	Complete Course	1250.00	940.00
SR21-0754	Additional PRG	11.50 purchase	
Z2152	Guided Learning Center	495.00	

- *3090 Installation Planning*, Course Code Y4837

  This no charge classroom course of 1 day teaches the information necessary to plan for the installation of an IBM 3090 Processor Complex, from both a hardware and software planning perspective. The course includes a product description, software support, installation considerations, and service enhancements.

  NOTE: Although this is a no-charge course, enrollments should be placed through IBM-Direct. Please do not call your IBM branch office to enroll.

  AUDIENCE: Project managers, installation managers, installation planners, and systems programmers.

  PREREQUISITES: Large systems installation experience.

  OBJECTIVES: After completing this course, a student should be able to do the following for an IBM 3090 Processor Complex:

  - Describe the hardware features
  - Describe the software support
  - Configure the complex
  - Develop an installation plan
  - Describe the remote support enhancements

- *3090 MODEL 400E/600E Planning Considerations*, Course Code H3761

This no charge classroom course of 1 day teaches the information necessary to plan for the installation of an IBM 3090 Model 400E/600E Processor Complex, from both a hardware and software planning perspective. This course discusses Model 400E/600E hardware features, software support, reconfiguration, recovery, and MVS/XA virtual storage assessment.

NOTE: Although this is a no-charge course, enrollments should be placed through IBM-Direct. Please do not call your IBM branch office to enroll.

AUDIENCE: Project managers, installation managers, installation planners, and systems programmers.

PREREQUISITES: Large systems installation experience. Prior attendance in the 3090 Installation Planning Course (Y4837) is desirable.

OBJECTIVES: After completing this course, a student should be able to do the following for an IBM 3090 Model 400E/600E Processor Complex:

- Describe the hardware features
- Describe the software support
- Develop an installation plan
- Reconfigure the complex between single-image and partitioned modes
- Describe recovery options and when to use them
- Evaluate and tune MVS/XA virtual storage

Automatic Distribution of Technical Bulletins

Large Systems Technical Support Centers
International Technical Support Centers

To provide automatic distribution of Systems and Support Center center bulletins to customers, a special procedure has been established using System Library Subscription Service (SLSS). To receive bulletins of interest automatically when they are released, the bill of forms numbers listed below should be added to your SLSS subscription. These bulletins are generally released only once and are not updated. The following GBOF numbers apply to technical bulletins published by Large Systems Technical Support (Washington and Dallas Systems Centers).

- GBOF-0064 High Availability Systems Specialist
- GBOF-0086 Printer Systems Specialist
- GBOF-2205 Data Systems Specialist
- GBOF-2206 Host Communications Specialist
- GBOF-2207 Distributed Systems Specialist
- GBOF-2208 Large Systems Specialist
- GBOF-2209 Intermediate Systems Specialist
- GBOF-2211 Systems Management Specialist
- GBOF-2212 Office Systems Specialist

Please advise your marketing representative or systems engineer of the categories of Large Systems Technical Support bulletins you wish to receive automatically.

The following GBOF numbers apply to technical bulletins published by the International Technical Support Centers (previously World Trade Systems Centers). Technical bulletins published by the ITSCs and previously distributed under the Large Systems Technical Support GBOF numbers listed above will no longer be distributed under those GBOF numbers. ITSC bulletins have been regrouped and are now distributed under the following GBOF numbers.

- GBOF-0421 ITSC Large Systems
- GBOF-0422 ITSC Data Systems
- GBOF-0423 ITSC Intermediate Systems
- GBOF-0424 ITSC Small Systems
- GBOF-0425 ITSC Communications
- GBOF-0426 ITSC Workstations
- GBOF-0427 ITSC Distributed Systems
- GBOF-0428 ITSC Office/Text Systems
- GBOF-0429 ITSC Miscellaneous & Cross Category
- GBOF-0430 ITSC Storage & I/O
- GBOF-2200 ITSC Engineering/Scientific
- GBOF-2201 ITSC VM Systems
- GBOF-2202 ITSC Distributed Communications Systems
- GBOF-2203 ITSC Consumer Systems
- GBOF-2204 ITSC Systems Design
- GBOF-2199 ITSC Application Software

Please advise your marketing representative or systems engineer of the categories of International Technical Support Center bulletins you wish to receive automatically.

Index

READER'S COMMENT FORM

Title: IBM 3090 Processor Complex:
Planning and Installation Reference
Washington Systems Center
Technical Bulletin GG66-3090-01

You may use this form to communicate your comments about this publication, its organization, or subject matter, with the understanding that IBM may use or distribute whatever information you supply in any way it believes appropriate without incurring any obligation to you.

Please state your occupation: _____

Comments:

Please mail to:

Brian Kasch
Washington Systems Center
18100 Frederick Pike
Gaithersburg, MD 20879

Cut or Fold Along Line

Reader's Comment Form

Fold and tape | Please Do Not Staple | Fold and tape

NO POSTAGE
NECESSARY
IF MAILED
IN THE
UNITED STATES

BUSINESS REPLY MAIL

FIRST CLASS PERMIT NO. 40 ARMONK, N.Y.

POSTAGE WILL BE PAID BY ADDRESSEE:

B. E. Kasch
IBM Corporation
Washington Systems Center
18100 Frederick Pike
Gaithersburg, MD 20879

Fold and tape | Please Do Not Staple | Fold and tape

IBM®

IBM 3090 Processor Complex:
Planning and Installation Reference

Printed in U.S.A.

GG66-3090-01

GG66-3090-01