

GC28-1160-4  
File No. S370-34

**Program Product**

**MVS/Extended Architecture  
Planning: Recovery and  
Reconfiguration**

**MVS/System Products:**

<b>JES2 Version 2</b>	<b>5740-XC6</b>
<b>JES3 Version 2</b>	<b>5665-291</b>

**IBM**

### **Fifth Edition (June, 1987)**

This is a major revision of, and obsoletes, GC28-1160-3. See the Summary of Amendments following the Contents for a summary of the changes made to this manual. Technical changes or additions to the text and illustrations are indicated by a vertical line to the left of the change.

This edition applies to Version 2 Release 2.0 of MVS/System Product (5665-291 and 5740-XC6), and to Data Facility Product (5665-284 and 5665-XA2), and to all subsequent releases until otherwise indicated in new editions or Technical Newsletters. Changes are made periodically to the information herein; before using this publication in connection with the operation of IBM systems, consult the latest *IBM System/370 Bibliography*, GC20-0001, for the editions that are applicable and current.

References in this publication to IBM products or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM product in this publication is not intended to state or imply that only IBM's product may be used. Any functionally equivalent product may be used instead.

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 readers' comments is provided at the back of this publication. If the form has been removed, comments may be addressed to IBM Corporation, Information Development, Department D58, Building 921-2, PO Box 390, Poughkeepsie, N.Y. 12602. IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

## Preface

### What to Expect From This Publication

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 installation's requirements. This publication does not contain ready-to-use procedures; rather 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. This publication addresses recovery and reconfiguration considerations for both UPs and MPs but not to the same degree for the two types of processor complex. It talks about recovery for both UPs and MPs. However, it talks about reconfiguration for MPs only, since there's very little on a UP system that can be taken offline without bringing down the whole system. For example, only on an MP can a CPU or a channel controller (e.g., an EXDC on a 308x) be taken offline.

This publication does **not** address software recovery from software errors. For example, recovery procedures for the following subsystems/components are outside the scope of this publication:

- Global Resource Serialization
- JES2
- JES3
- CICS
- IMS
- ACF/VTAM

### How This Publication is Organized

The contents of each chapter are described in the following paragraphs.

**Chapter 1: Introduction to Recovery and Reconfiguration** provides overview information concerning the processes of recovery and reconfiguration.

**Chapter 2: Pre-Installation Planning for Reconfiguration** provides guidelines to an installation on how to set up its I/O configuration.

**Chapter 3: Recovery** describes what the hardware and software facilities do to recover from a hardware failure. This information helps the installation understand the system's attempt to recover from a hardware failure and the effect the recovery attempt has on the system.

**Chapter 4: Reconfiguration** describes the process of adding hardware units to, or removing hardware units from, a configuration.

## Bibliography

The following publications are either referred to in this manual or contain further information on a topic described in this manual.

### MVS/Extended Architecture

*Message Library: System Messages, Volume 1, GC28-1376*

*Message Library: System Messages, Volume 2, GC28-1377*

*Message Library: System Codes, GC28-1157*

*Operations: System Commands, GC28-1206*

*System Programming Library: Initialization and Tuning, GC28-1149*

*System Programming Library: System Modifications, GC28-1152*

*Installation: System Generation, GC26-4009*

*MVS Configuration Program Guide and Reference, GC28-1335*

*IBM System/370 Extended Architecture Principles of Operation, SA22-7085*

*Input/Output Configuration Program User's Guide and Reference, GC28-1027*

*IBM 3090 Processor Input/Output Configuration Program User's Guide and Reference, SC38-0038*

*IBM Disk Storage Management Guide: Error Handling, GA26-1672*

*EREP User's Guide and Reference, GC28-1378*

*Data Facility Data Set Services User's Guide and Reference, SC26-3949*

*Device Support Facilities, GC35-0033*

*IBM 3081, 3083, 3084 Messages for the System Console, GC38-0035*

*IBM 3081 Operator's Guide for the System Console, GC38-0034*

*IBM 3083 Operator's Guide for the System Console, GC38-0036*

*IBM 3084 Operator's Guide for the System Console, GC38-0037*

*IBM 3090 Channel Characteristics, SA22-7120*

(Includes all base and E models)

*IBM 3090 Functional Characteristics, SA22-7121*

(Includes all base and E models)

*IBM 3090 Processor Complex: Operator Controls for the System Console for Models 150/150E/180/180E/200/200E/400/400E, SC38-0040*

*IBM 3090 Processor Complex: Operator Tasks for the System Console for Models 200 and 200E, SC38-0041*

*IBM 3090 Processor Complex Operator Tasks for the System Console for Models 150/150E/180/180E, SC38-0049*

*IBM 3090 Processor Complex: Operator Tasks for the System Console for Models 400 and 400E, SC38-0050*

*IBM 3090 Processor Complex Recovery Guide, SC38-0051*

*IBM 3090 Processor Complex Operator Tasks for the System Console for Model 300E, SC38-0054*

*IBM 3090 Processor Complex Operator Tasks for the System Console for Model 600E, SC38-0056*

*Note:* The following package applies to the 3090 Model 200 (without the Vector Facility), although much of the controls and task information should also be useful for a 3090 Model 400.

*IBM 3090 Processor Complex Operator Training Package (3 volumes), GG24-1740 through GG24-1742*

*IBM 3081 Functional Characteristics, GA22-7076*

*IBM 3083 Functional Characteristics, GA22-7083*

*IBM 3084 Functional Characteristics, GA22-7088*

*IBM 4381 Processor Model Group 3 Functional Characteristics , GA24-4021*

*IBM 4381 Processor Operations Manual, GA24-3949*



# Contents

<b>Chapter 1. Introduction to Recovery and Reconfiguration</b>	1-1
Terminology and Definitions	1-1
Introduction to Recovery	1-3
Hardware Recovery	1-3
System Recovery	1-3
Operator Intervention	1-3
Introduction to Reconfiguration	1-3
<b>Chapter 2. Pre-Installation Planning for Reconfiguration</b>	2-1
I/O Configuration Considerations	2-1
Channel Subsystem Considerations for 308x Processors	2-2
Channel Subsystem Considerations for 3090 Processors	2-2
Channel Designations on 3090 Models 200, 200E, and 300E	2-3
Channel Designations on 3090 Models 400, 400E, 600E	2-3
Configuration Guidelines	2-4
Channel Subsystem Considerations for a 4381 Model 3 Processor	2-4
Configuring Devices for a Nonpartitionable Processor Complex	2-5
I/O Configuration Guidelines for a Partitionable Processor Complex in Single-Image Mode	2-10
String Switching	2-10
Master Console Configuration Guidelines	2-12
IOCP Considerations	2-12
MVS/XA Configuration Program or SYSGEN Considerations	2-13
RSU Parameter	2-14
RSU Implementation	2-14
RSU Parameter Specifications	2-15
<b>Chapter 3. Recovery</b>	3-1
Hardware/Operating System Recovery Actions	3-2
Information Provided with Machine Checks	3-4
Central Processing Unit Errors	3-4
Soft CPU Errors	3-5
Hard CPU Errors	3-5
Terminating CPU Errors	3-6
Vector Facility Recovery	3-7
Vector Facility Source Error Machine Check	3-7
Vector Facility Failure Machine Check	3-8
Alternate CPU Recovery (ACR)	3-8
Terminating Errors on Multiple CPUs	3-9
Service Processor Damage	3-9
Storage Errors	3-9
Soft Storage Errors	3-10
Hard Storage Errors	3-10
Effects of Storage Errors	3-11

Storage Element Failure	3-12
Channel Subsystem Errors	3-12
Channel Report Words (CRWs)	3-12
Channel Path Recovery	3-12
Channel Path Alert Conditions	3-14
Subchannel Recovery	3-14
Monitoring Facility Recovery	3-14
I/O Errors	3-14
Master Console Failure	3-15
Missing Interrupts	3-15
Unconditional Reserve/Alternate Path Recovery (APR)	3-16
Hot I/O	3-16
Sample Parameters for Hot I/O Recovery in Parmlib Member IECIOSxx	3-17
3880/3380 Considerations	3-17
3380 Enable/Disable Switch	3-17
Recovery from an Out-of-Sync Condition	3-18
DASD Maintenance and Recovery	3-18
Operator Recovery Actions	3-19
Recovery by CPU Restart	3-19
Continuing a Vector Job If a Vector Facility is Offline	3-20
Hardware Instruction Tracing (Loop Trace)	3-20
Recovery from Wait States	3-21
Disabled Wait States	3-21
Enabled Wait States	3-21
Uncoded Wait States	3-21
Spin Loop Recovery	3-21
Spin Loops	3-22
Operator Notification	3-22
Processing Messages at the System Console	3-23
ACR Considerations	3-23
Recovery Actions	3-24
Example of Recovery Procedure for Spin Loop Message	3-24
Recovery for X'09x' Wait State	3-25
Additional Recovery Actions	3-26
Restart Procedures	3-26
Procedure to Restart from Message IEE331A	3-26
Procedure to Restart from Wait State 091, 092, 095, 097, or 09E	3-26
Determining the Cause of a Spin Loop	3-27
Analysis of Excessive Spin LOGREC Records	3-27

<b>Chapter 4. Reconfiguration</b>	4-1
Logical and Physical Reconfiguration	4-2
General Considerations for Reconfiguration	4-2
Reconfiguration Support According to Processor Types	4-2
Recommended Sequence for Partitioning and Merging	4-3
DISPLAY Command Considerations	4-4
D U Command	4-4
D M Command	4-4
Program Properties Table Considerations	4-5
Real Storage Reconfiguration	4-6
Extended Storage Reconfiguration on Partitionable 3090 Models	4-9
Processor Reconfiguration	4-10
Reconfiguring a Processor with a Vector Facility	4-10



Removing the Last Vector Facility	4-11
Channel Measurements	4-11
Vector Facility Reconfiguration Examples	4-12
Channel Path Reconfiguration	4-14
I/O Device Reconfiguration	4-15
Examples of Partitioning and Merging a 3084	4-15
Partitioning from Single-Image Mode to Physically Partitioned Mode (Side B to Be Configured Offline)	4-17
Merging from Physically Partitioned Mode to Single-Image Mode (Side B To Be Configured Online)	4-25
Examples of Partitioning and Merging a Partitionable 3090	4-32
Partitioning from Single-Image Mode to Physically Partitioned Mode (Side 1 to Be Configured Offline)	4-33
Merging from Physically Partitioned Mode to Single-Image Mode (Side 1 To Be Configured Online)	4-42

**Index** X-1



## Figures

- 2-1. 308x Channel Subsystem 2-2
- 2-2. 3090 Channel Subsystem Configuration 2-3
- 2-3. Channel Paths and Channel Elements on Partitionable 3090 Systems 2-3
- 2-4. 4381 Dual-Processor Channel Subsystem Configuration 2-4
- 2-5. DASD Configuration for Maximum Availability with a 308x Complex 2-6
- 2-6. Tape Configuration for Maximum Availability (308x Complex) 2-7
- 2-7. Configuration of a 3725 for Maximum Availability (308x Complex) 2-8
- 2-8. Unit Record or Local TP Device Configuration for Maximum Availability (308x Complex) 2-9
- 2-9. DASD Configuration for Maximum Availability (3084 in Single-Image Mode) 2-11
- 2-10. Calculation of RSU Value for the 3084 and Reconfigurable 3090 Models 2-15
- 3-1. Operating System Handling of Machine Checks 3-3
- 3-2. Recovery Actions for Each Message Insert or Wait State 3-24
- 4-1. A Logical View of Real Storage (3084 Example) 4-6
- 4-2. A Physical View of Real Storage (3084) 4-7
- 4-3. Real Storage Differences Between 308x and 3090 Systems 4-8
- 4-4. Single-Image Mode of a 3084 4-17
- 4-5. Storage Layout in Single-Image Mode (3084) 4-18
- 4-6. Storage Layout - SE1 Configured Offline (3084) 4-21
- 4-7. Storage Layout - SE1 and SE3 Configured Offline (3084) 4-22
- 4-8. Physically Partitioned Mode of the 3084 4-23
- 4-9. Examples of D M Displays - 3084 System in Physically Partitioned Mode 4-24
- 4-10. Storage Layout - SE1 and SE3 Configured Offline (3084) 4-25
- 4-11. Storage Layout - SE3 Configured Online (3084) 4-27
- 4-12. Storage Layout - SE1 and SE3 Configured Online 4-28
- 4-13. Single-Image Mode of a 3084 4-30
- 4-14. Examples of D M Displays - 3084 System in Single-Image Mode 4-31
- 4-15. Differences Between the 3084 and the 3090 Models 400, 400E, and 600E 4-32
- 4-16. Single-Image Mode of a 3090 Model 400 4-33
- 4-17. Sample Real Storage Layout of 3090 Model 400 Before Partitioning 4-34
- 4-18. Real Storage Layout - SE2 Configured Offline (3090 Model 400) 4-37
- 4-19. Real Storage Layout - SE3 Configured Offline (3090 Model 400) 4-39
- 4-20. Physically Partitioned Mode of the 3090 Model 400 4-40
- 4-21. Examples of D M Displays - 3090 Model 400 System in Physically Partitioned Mode 4-41
- 4-22. Single-Image Mode of a 3090 Model 400 4-45

4-23. Examples of D M Displays - 3090 Model 400 System in Single-Image  
Mode 4-46

## Summary of Amendments

### **Summary of Amendments for GC28-1160-4 for MVS/System Product Version 2 Release 2.0**

This major revision includes the following new and changed information:

- Information in the various chapters to reflect support of the 3090 models 200E, 300E, and 600E.
- In Chapter 2 a description and summary table of channel path IDs and channel element IDs for the 3090 Models 400, 400E, and 600E.
- In Chapter 2 a description of the real storage increment size for the 3084 and the 3090 models 400, 400E, and 600E.
- In Chapter 3 additional details on subchannel recovery.
- In Chapter 3 additional comparison of 3090 hardware instruction tracing with that on a 308x.
- In Chapter 4 a table that describes real storage for the 3081, 3084, and various 3090 models: the storage element IDs, storage element size, storage increment size, storage subincrement size, and maximum storage size.
- Minor technical and editorial corrections throughout the manual.

Changes are indicated by change bars (|) in the left margin.

### **Summary of Amendments for GC28-1160-3 for MVS/System Product Version 2 Release 1.7**

This major revision contains changes to support reconfiguration of the 3090 Model 400 in MVS/System Product Version 2 Release 1.7. The changes include:

- The addition of a new CONFIG parameter to configure extended storage elements offline and online.
- Changes to the reconfiguration examples and DISPLAY M examples in Chapter 4 to illustrate differences between the 3084 and the 3090 Model 400 during partitioning and merging.
- Change of message ID for the D M display message. It used to be IEE490I; it is now IEE174I.

- In Chapter 4, a description of the recommended sequence of issuing CONFIG commands, aimed at allowing channel measurements during partitioning and merging and at reducing resource contention.
- In Chapter 4, addition of a procedure to get a storage element to come online when it does not respond in the normal manner to CONFIG STOR(E=X),ONLINE.
- In Chapter 4, addition of a separate section of partitioning and merging examples to illustrate the reconfiguration of the 3090 Model 400.
- In Chapter 4, changes to the 3084 sample D M displays to reflect programming changes that make the 3084 displays more consistent with the 3090 Model 400 displays.

Changes are indicated by change bars (|) in the left margin.

**Summary of Amendments  
for GC28-1160-2  
for MVS/System Product Version 2 Release 1.3**

This major revision, which supports MVS/System Product Version 2 Release 1.3, includes the following new and changed information:

- In Chapter 2, additional guidance on how to increase availability for devices configured across 3090 channel elements.
- In Chapter 4, a table of real storage differences between a 3090 and a 308x.
- For MVS/System Product Version 2 Release 1.3 Vector Facility Enhancement:
  - In Chapter 3, modifications to the machine check handler processing.
  - In Chapter 4, a new VF parameter for the CONFIG command, and sample CONFIG commands and D M displays that illustrate the use of this parameter.
- Minor technical and editorial corrections.

Changes are indicated by change bars (|) in the left margin.

## **Chapter 1. Introduction to Recovery and Reconfiguration**

To each installation, the approaches to recovery and reconfiguration may be different depending on its recovery philosophy. For example, one installation decides that continuous system operation (as long as work can be done) is one of its priorities. In this case, after a malfunction the installation keeps the system operational as long as possible, even at the risk of lost diagnostic data, and defers maintenance. Another installation, however, decides that immediate repair of a failing unit is one of its priorities. In this case after a malfunction, the installation takes as much of the system as is necessary out of operation to perform the maintenance.

An installation should base its recovery and reconfiguration procedures on its operational priorities. Each installation may need several procedures because the operational priorities may change with workload or time-of-day. For example, priorities and procedures may change during a shift to accommodate a heavier or lighter workload. Also, priorities and procedures that apply to first shift may not apply to third shift.

### **Terminology and Definitions**

To resolve any conflicts concerning terms used in this publication, the following list defines the meaning and usage of those terms.

**Configuration** - a set of hardware units that can support a single operating system.

**Dual Processor** - A non-partitionable multiprocessor that has two CPUs, each having its own integrated channel paths. That is, each CPU's channel paths work only with that CPU and cannot be accessed by the other CPU.

**Hardware Unit** - a CPU, storage element, channel path, device, etc.

**Master Console** - the console used for communications between the operator and the software system.

**MP or Multiprocessor** - a processor complex that has more than one CPU.

**Partition** - one of the configurations formed by partitioning.

Partitioning - the process of forming multiple configurations from one.

**Note:** A processor complex that supports partitioning is termed partitionable; a processor complex that does not support partitioning is termed nonpartitionable. These types of processor complex are partitionable: the 3084 and the 3090 Models 400, 400E, and 600E.

Physical Partition - a hardware implementation of a partition.

Physically Partitioned Mode - the state of a processor complex when its hardware units are divided into multiple configurations.

Processor - a central processing unit (CPU)

Processor Complex - the maximum set of hardware units that support a single operating system.

Service Processor (equivalent to "processor controller) - that part of a processor complex that provides for the maintenance of the complex and may perform:

- Some or all of the functions associated with operator facilities
- Recovery actions associated with machine-check handling
- Reconfiguration operations

Side - equivalent to the term "Physical Partition".

Single-Image Mode - the state of a processor complex when all of its hardware resources are in a single configuration.

System - the interactive combination of a configuration and the operating system (software).

System Console - the console used by the operator to enter hardware commands and to receive hardware messages.

UP or Uniprocessor - a processor complex that has one CPU.

VF or Vector Facility - an optional processing facility to do vector mathematics, available for all 3090 models. There can be one or more Vector Facilities for each processor complex, but only one Vector Facility is associated with each central processor unit.



## Introduction to Recovery

Recovery is the attempt by the hardware, the operating system, the operator, or any combination of the three, to correct system malfunctions and return the system to a state in which it can do productive work.

### Hardware Recovery

Many temporary hardware errors are recovered by the hardware and do not require operating system or operator intervention.

### System Recovery

System recovery involves both the hardware and the operating system, because many hardware malfunctions are communicated to the operating system for retry and recovery. When a malfunction occurs, it may cause the machine check handler (MCH), alternate CPU recovery (ACR), I/O Supervisor (IOS), or another operating system function to be invoked. That operating system function may retry in an attempt to recover or may determine that recovery from the particular malfunction is not feasible and configure the failing unit offline.

### Operator Intervention

For some hardware failures, operator intervention is required to attempt recovery to keep the system in operation. For example, if a channel path fails, the operator can configure it offline. If the system is operating in single-image mode, the operator can reconfigure to physically partitioned mode to allow maintenance to be performed on a partition. In addition, other system errors may require operator intervention to attempt recovery. Those errors include wait states, loops, excessive spin loop timeouts, Hot I/O, etc. (Refer to Chapter 3 for detailed information.)

## Introduction to Reconfiguration

Reconfiguration is the process of adding hardware units to, or removing hardware units from, a configuration. Operational units (for example, CPUs, storage elements, and channel paths) can be added to the configuration (configured online) to make them available to the system. Failing units can be removed from the configuration (configured offline) to make them unavailable to the system and (possibly) allow the system to continue operation.

One facet of system reconfiguration is **partitioning** — changing from single-image mode to physically partitioned mode. This capability is available only on a 3084 and 3090 Models 400, 400E, and 600E systems. An installation can use partitioning as an operational convenience or as an aid to recovery. In the first case, partitioning allows one operating system to run in one partition and another one in the other partition. For example, an installation could run MVS/XA in one partition; MVS/370, VM/SP, or a test system in the other partition. In the second case, an installation could give the failing partition to service personnel for diagnosis and repair and still have the system continue operation.

Another facet of system reconfiguration on a 3084 and the 3090 Models 400, 400E, and 600E is **merging** -- changing from physically partitioned mode to single-image mode. An installation uses merging to create a single more powerful MP system from the two systems that exist in physically partitioned (or PP) mode. One of the two PP-mode systems must be stopped and the partition (side) on which it was running be merged with the MVS/XA system that is running on the other partition (side).

## Chapter 2. Pre-Installation Planning for Reconfiguration

Prior to placing a system into operation, an installation should plan the configuration for maximum availability, or to state it in terms of reconfiguration — to provide the maximum capability for reconfiguration. Also, since reconfiguration is a consideration for recovery, the planning serves a dual purpose.

There are two aspects of the planning — one focuses on the I/O configuration and the other on the operating system. For the I/O configuration, an installation should focus on such things as device attachment (which devices are attached to which paths) and master console configuration. For the operating system, an installation should focus on such things as system generation considerations and SYS1.PARMLIB updates.

### I/O Configuration Considerations

This section deals with the following areas:

- Channel subsystem considerations for the various processor types
- Sample I/O configurations illustrated on a 308x system
- I/O configuration guidelines for a 3084 complex in single-image mode
- Master console configuration guidelines
- IOCP considerations

I/O configuration planning will enhance system availability and recovery. This involves the consideration of the number of paths to each device and the hardware elements in each path. **MULTIPLE PATHS TO A DEVICE SHOULD INCLUDE AS FEW COMMON HARDWARE ELEMENTS AS POSSIBLE TO MINIMIZE THE EFFECT OF A MALFUNCTION; THAT IS, TO PREVENT A SINGLE MALFUNCTION FROM DISABLING ALL THE PATHS TO A DEVICE.** Multiple paths to some devices may require the installation of two-channel switches, or switching devices such as a 3814 or 2914.

The illustrations that follow in this chapter show examples of I/O configurations that will give an installation maximum availability in case of a hardware element failure. Each illustration shows the hardware elements in the path to a device and shows how the device should be connected.

## Channel Subsystem Considerations for 308x Processors

The hardware elements in a path from a 308x to a device are:

- External data controller (EXDC)
- Data server element (DSE)
- Control unit
- Channel path (CHP)

Figure 2-1 shows the relation among the hardware elements that make up the 308x channel subsystem.

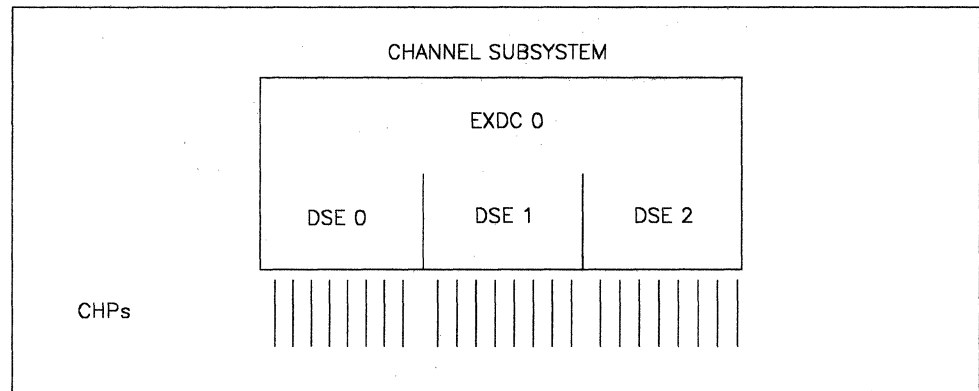


Figure 2-1. 308x Channel Subsystem

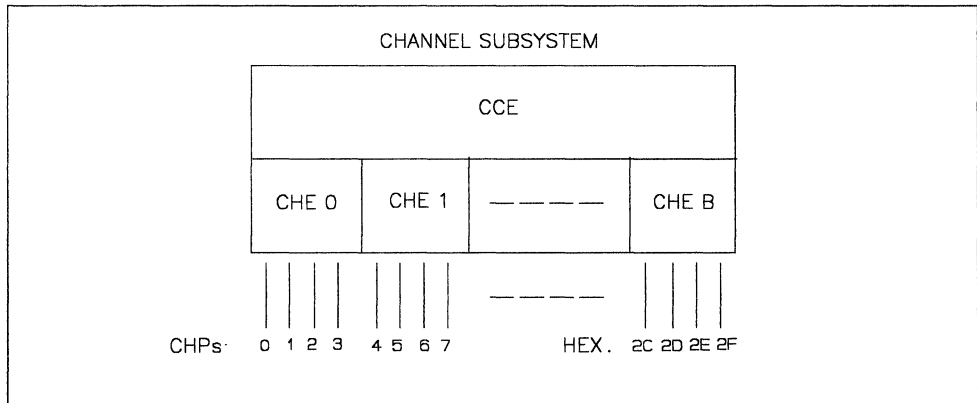
To maintain maximum availability in case of a CHP or DSE failure, an installation should configure multi-path devices to CHPIDs on different DSEs. On systems with multiple EXDCs (i.e., 3084), the installation should configure multi-path devices across the EXDCs.

## Channel Subsystem Considerations for 3090 Processors

The hardware elements in a path from a 3090 to a device are:

- Channel control element (CCE)
- Channel element (CHE)
- Channel path (CHP)
- Control unit

Figure 2-2 shows the relation among the hardware elements that make up the 3090 channel subsystem.



**Figure 2-2. 3090 Channel Subsystem Configuration**

**Channel Designations on 3090 Models 200, 200E, and 300E**

The nonpartitionable models of the 3090 (for example, models 200, 200E, and 300E) have 32 standard CHPs numbered 00 -- 1F. The model 200 can have 8 or 16 additional CHPs, numbered 20 -- 2F, and the models 200E and 300E can have up to 32 additional CHPs, numbered 20 -- 3F. Note that unlike the 308x CHPIDs, the CHPs on the nonpartitionable models of the 3090 are consecutively numbered with no gaps in the numbering.

A channel element (CHE) has four consecutively numbered channel paths. The CHEs numbered 0 -- 7 relate to the 32 standard CHPs on all three nonpartitionable models. For the optional CHPs, the CHEs are numbered 8 -- B on the model 200, and 8 -- F on the models 200E and 300E. Figure 2-3 summarizes CHP and CHE numbering for several 3090 models.

**Channel Designations on 3090 Models 400, 400E, 600E**

Figure 2-3 shows the channel path (CHP) designations and the channel element (CHE) designations for the 3090 models 400, 400E, and 600E.

Model Number	Model 400	Model 400E	Model 600E
Max Number of CHPs	96	128	128
CHPIDs - Side A	0 - 2F	0 - 3F	0 - 3F
CHPIDs - Side B	40 - 6F	40 - 7F	40 - 7F
Max Number of CHEs	24	32	32
CHE IDs - Side A	0 - B	0 - F	0 - F
CHE IDs - Side B	10 - 1B	10 - 1F	10 - 1F

**Figure 2-3. Channel Paths and Channel Elements on Partitionable 3090 Systems**

## Configuration Guidelines

When you configure a device with multiple paths to the same system, the following guidelines allow for maximum availability of the device:


1. Attach each path from the device to a separate CHE.
2. Distribute paths across both odd and even numbered CHEs. Thus, one path from the device could include any of these channel paths: 0--3, 8--B, 10--13, 18--1B. The other path (assuming two per device) could be selected from channel paths 4--7, C--F, 14--17, or 1C--1F.
3. If optional channel paths (20--2F) are installed, distribute the paths to a device across both optional and standard channel paths. When doing this, also follow the configuration pattern of odd and even CHEs (point #2 in this list).
4. For a Model 400, 400E, or 600E, distribute the paths across channel control elements (CCEs). (A CCE is analogous to an EXDC on a 308x.)

## Channel Subsystem Considerations for a 4381 Model 3 Processor

The hardware elements in the path from a 4381 dual processor to a device are:

- Channel path (CHP)
- Control unit

Figure 2-4 shows the relation among the hardware elements that make up the 4381 dual-processor channel subsystem.



**Figure 2-4. 4381 Dual-Processor Channel Subsystem Configuration**

*Note:* The Model 3 and 14 have 6 standard channels on each processor, with 6 optional channels: 3 extra for each processor, as shown. The model 24 (not shown) has up to 24 channels, 12 on each processor.

Where possible, I/O devices should be connected to channel paths on both CPUs of the processor complex. It is particularly important for recovery that system DASD devices be accessible from either CPU and that at least one operator console be attached to each CPU. This is especially necessary if a CPU failed, since the associated channel paths would become unusable.

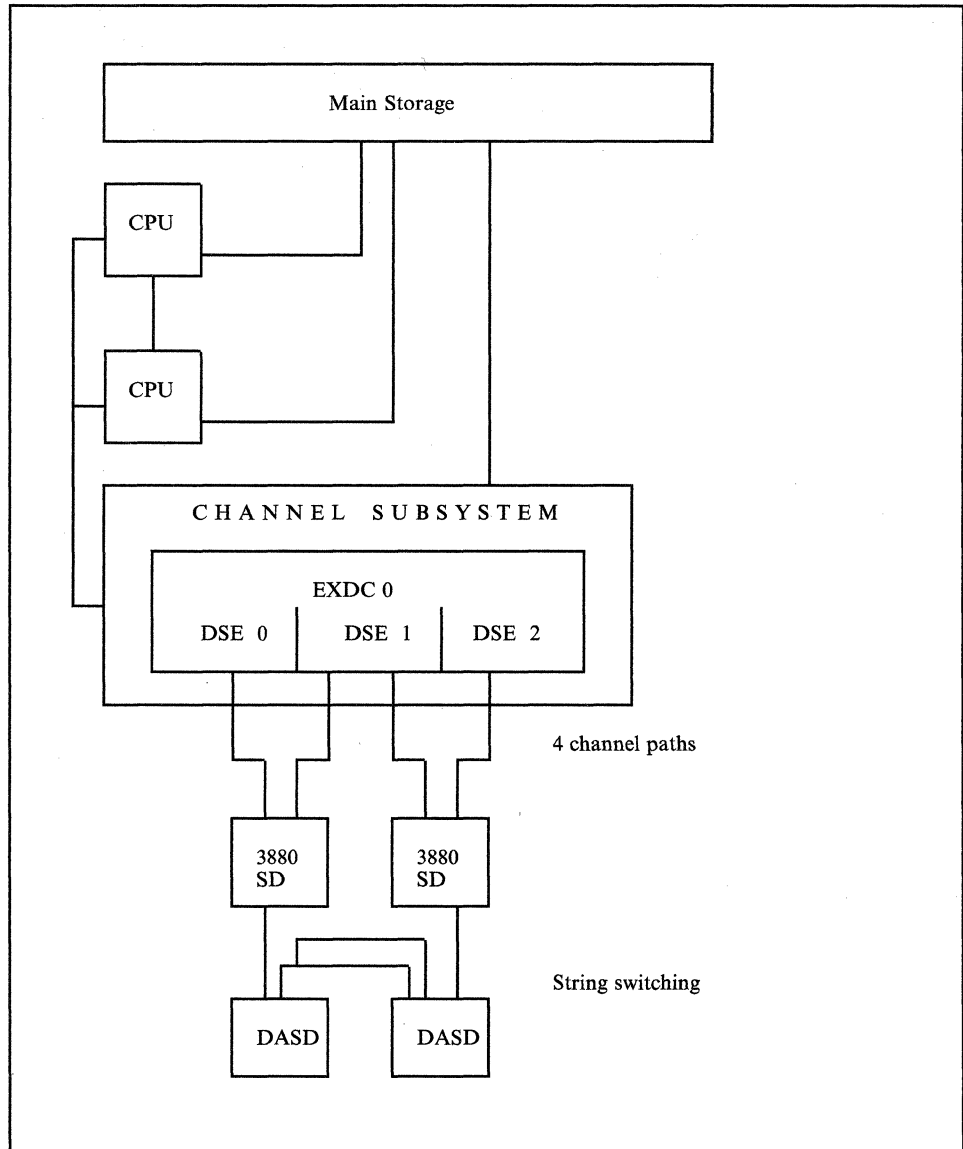
It is desirable to attach an MVS operator console to both CPUs. The MVS operator console for 4381 Model 3 is connected through a local channel adapter only to channel path zero on CPU zero. It is desirable to have at least one operator console device attached through a local 3274 control unit on a channel path belonging to CPU 1. This configuration would still allow the operator to communicate with MVS if either CPU 0 or channel path zero were to fail.

### **Configuring Devices for a Nonpartitionable Processor Complex**

The following device-type configurations are shown:

- DASD
- Tape
- 3725
- Unit record, local displays, etc.

Figure 2-5 on page 2-6 through Figure 2-8 on page 2-9 show, respectively, the DASD, tape, 3725, and unit record configurations for maximum availability on a 308x complex. Although these figures illustrate the 308x, the general concepts apply to all MVS/XA-supported processor types.



**Figure 2-5. DASD Configuration for Maximum Availability with a 308x Complex**

String switching (or a 3380 model AA4) is chosen so all devices can be accessed, even if a control unit fails.



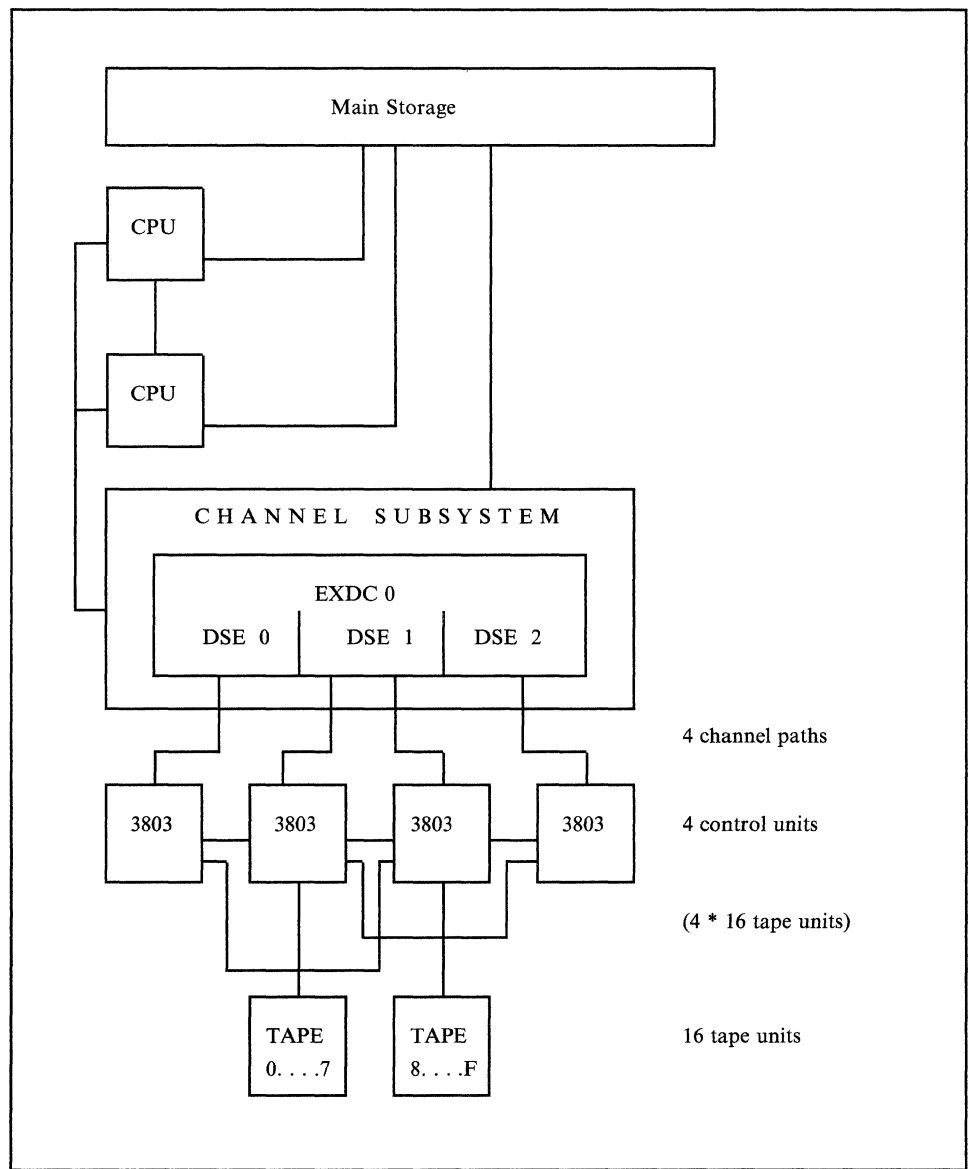


Figure 2-6. Tape Configuration for Maximum Availability (308x Complex)

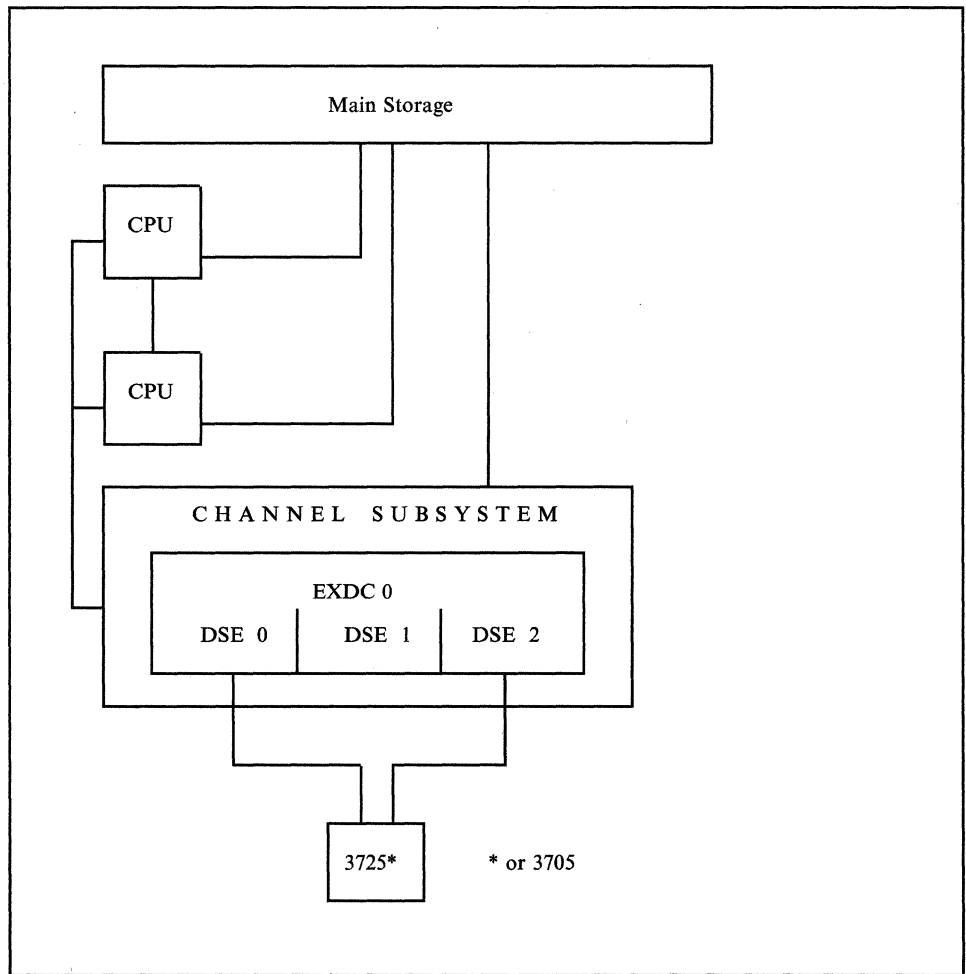
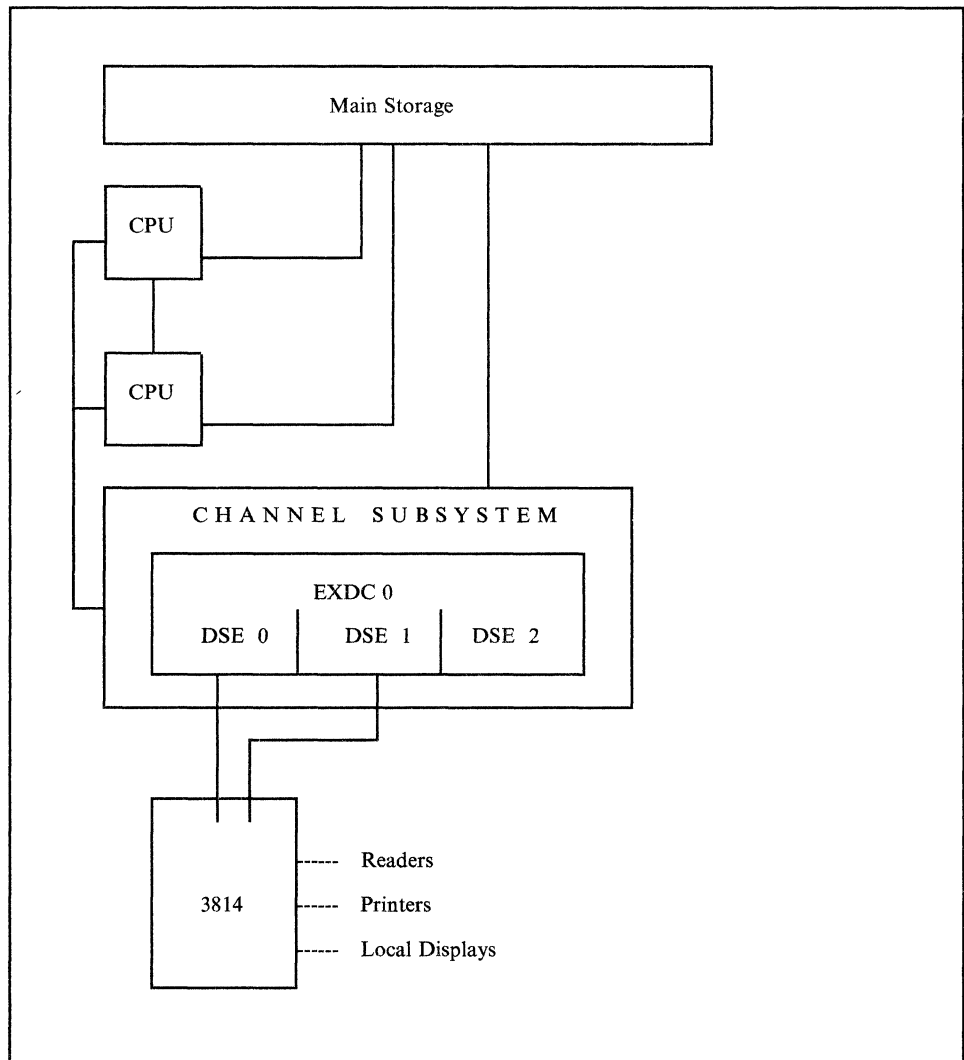


Figure 2-7. Configuration of a 3725 for Maximum Availability (308x Complex)



**Figure 2-8. Unit Record or Local TP Device Configuration for Maximum Availability (308x Complex)**

## I/O Configuration Guidelines for a Partitionable Processor Complex in Single-Image Mode

The partitionable processor complexes are the 3084 and the 3090 models 400, 400E, and 600E. When an installation plans its I/O configuration for one of these complexes in single-image mode, the plan should address not only that mode, but physically partitioned mode as well. The resulting I/O configuration should provide maximum availability in either mode of operation.

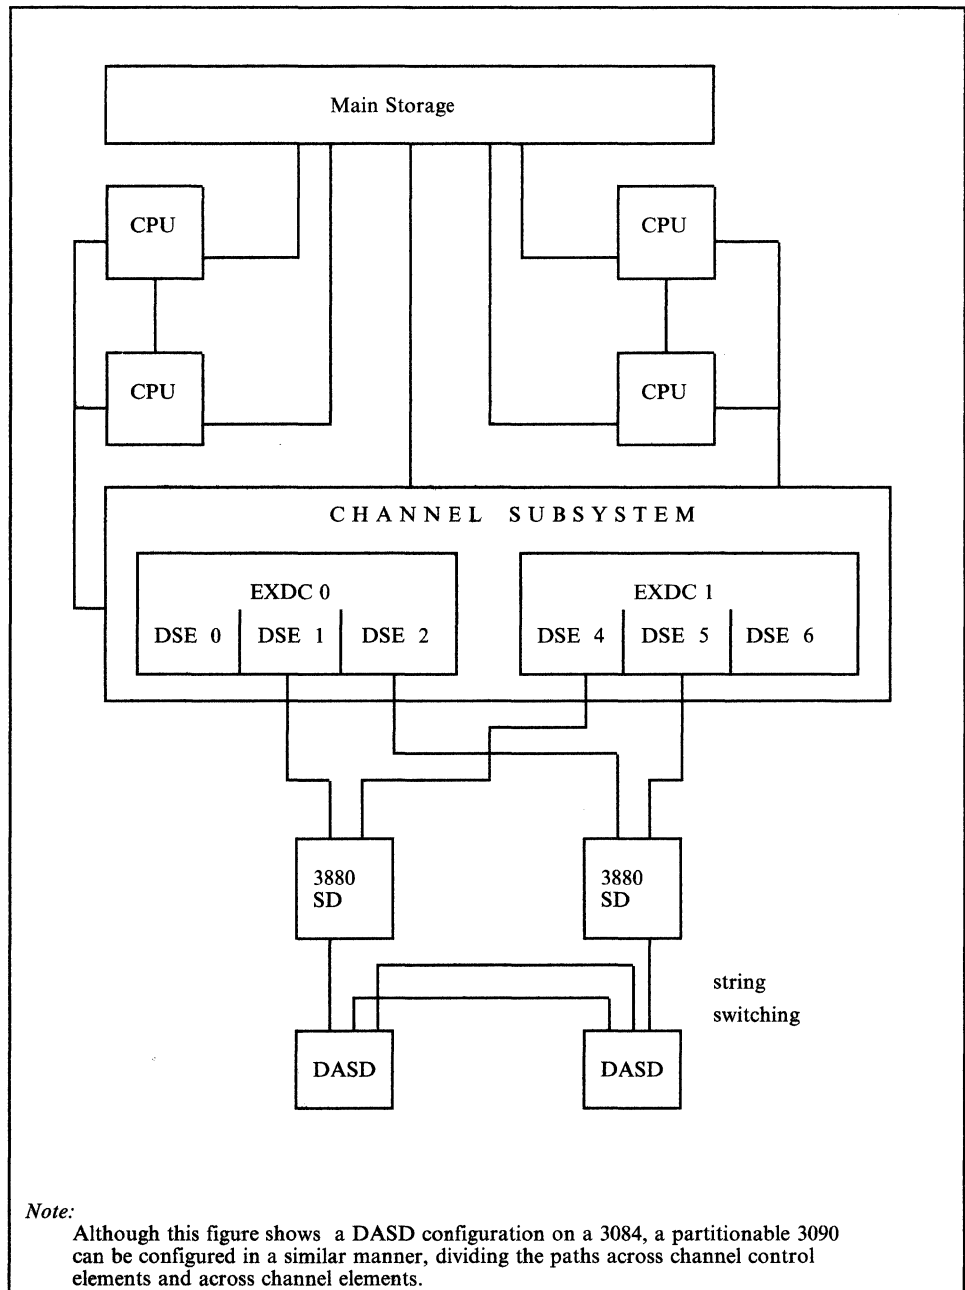
Some general recommendations for an I/O configuration are:

- Attach control units symmetrically, whenever possible, so they can be accessed from both sides.
- Attach critical unit record or local TP device control units through a 3814 (2914) switch, so they can be switched to either side.
- Attach a tape or DASD device to one channel path from each side and operate with two channel paths online. (To provide the same availability in physically partitioned mode, attach the device to four channel paths, two from each side.)

Figure 2-9 illustrates a configuration for single-image mode for maximum availability. The example shows a DASD configuration on a 3084. However, tape, TP-controller, and unit record configurations on a 3084 (and all device configurations on a 3090) should be done similarly, using features like two-channel switches and switching units, to attach the devices to different channel paths on different sides.

### String Switching

String switching features are used so that all devices can be accessed if a control unit fails.



**Figure 2-9. DASD Configuration for Maximum Availability (3084 in Single-Image Mode)**

## Master Console Configuration Guidelines

When attaching the master console (and its alternate) to a complex, an installation should implement the following guidelines. Following these guidelines provides a high degree of access to the consoles and helps increase the availability of the system:

- Dedicate a control unit to the master console and a different one to its alternate.
- Ensure that in a string of control units on the channel path, the control unit for the master/alternate console is the first terminal control unit on the channel path. Also, ensure that the control unit is set for *high* priority (performed by the service rep).
- Attach the master console and its alternate such that they share the least number of common hardware elements. For systems that can be physically partitioned, attach the master and alternate consoles to channel paths on different sides.

If the master and alternate consoles are NOT on dedicated control units, the ability of the operator to communicate with the operating system in certain recovery situations is impacted.

During system recovery processing for situations such as Hot I/O and Spin Loop Timeout, the Disabled Console Communications Facility (DCCF) is used to communicate with the operator. If DCCF is unable to issue a message to the MVS master console or its alternate, it attempts to issue the message to the system console. If the message cannot be issued to the system console, either the entire system or one CPU (depending on the problem) will be put into a restartable wait state. To recover from the wait state, the operator must use recovery procedures which may require modification of real storage. By providing the master console with its own dedicated control unit, the chances of encountering a restartable wait state, as a result of DCCF processing, are reduced significantly, and the operator may not need to display or modify real storage.

## IOCP Considerations

When operating in single-image mode, the 'Dual Write' function allows an installation to define duplicate copies of a new I/O configuration on each partition with a single execution of IOCP. The installation can then physically partition the complex and have the new I/O configuration available to either partition. In this way, for example, the installation can use either partition to test the new I/O configuration, or act as a common back-up.

## MVS/XA Configuration Program or SYSGEN Considerations

The MVS configuration program is to be used by installations that have installed MVS/System Product Version 2 Release 2.0 (MVS/SP2.2.0) or a later release. These installations use the MVS configuration program to:

- Define new I/O configurations or eligible device tables
- Replace existing I/O configurations or eligible device tables
- Define the consoles that the nucleus initialization program (NIP) can use
- Migrate I/O configurations or eligible device tables that were previously defined through the SYSGEN process so they can be used on the MVS/SP2.2.0 or a subsequent release.

FEATURE=SHAREDUP is an optional parameter on the IODEVICE statement in the MVS/XA configuration program and in the pre-SP2.2.0 sysgen program. The specification of this feature:

- Eliminates the overhead of the hardware device reserve/release logic when a device is attached only to partitions operating in single-image mode.
- Indicates that the device reserve/release logic is to be used only when operating in physically partitioned mode and allows the sharing of the device between partitions.

*Note:* Specify FEATURE=SHARED (not FEATURE=SHAREDUP) if a device is attached to more than one processor complex.

FEATURE=ALTCTRL is an optional parameter on the IODEVICE statement in the MVS/XA configuration program. The specification of this feature allows a device to be accessed through an alternate control unit.

For additional information regarding SHAREDUP and ALTCTRL, refer to *MVS/XA Configuration Program Guide and Reference* or the *MVS/XA System Generation* manual.

## RSU Parameter

The RSU parameter specifies the number of storage increments that the operating system tries to keep available for storage reconfiguration. (The size of a storage increment depends on the processor model and sometimes on the system engineering change (SEC) level.) At system initialization (IPL-time), when the RSU parameter is processed, the operating system assigns the number of storage increments specified in the RSU parameter to 'non-preferred' status (non-preferred for long-term page fixes for a non-swappable job). 'Non-preferred' storage is also called 'reconfigurable' storage. The operating system uses storage frames from both non-preferred and preferred storage to satisfy normal page allocation requests and requests for short-term page fixes.

Normally, the operating system assigns long-term fixed pages for a non-swappable job only to storage frames in the preferred area. However, if a long-term fixed page for a non-swappable job requires storage space but the preferred area is full, the operating system may convert some non-preferred storage to preferred storage. If some non-preferred storage is converted to preferred, the amount of storage available for reconfiguration is less than that specified in the RSU parameter. The operating system informs the operator of the condition by issuing message IAR005I.

When the operating system is requested to configure storage offline, it attempts to free the amount of real storage required to support the request. The physical real storage and the address ranges assigned to that storage cannot be configured offline either logically or physically until the required amount of storage is available.

The RSU parameter is specified in the IEASYSxx member of SYS1.PARMLIB or as an IPL parameter. The default value assigned to the RSU parameter is 0 — indicating that all storage is designated as preferred (that is, non-reconfigurable). (See *SPL: Initialization and Tuning* for detailed information.)

## RSU Implementation

At IPL-time, the specification of  $RSU = x$  is satisfied from the total amount of installed real storage (both online and offline). Therefore, a specification of  $RSU = 0$  indicates that the operating system will designate ALL installed real storage as preferred. As a result, an installation must specify the proper RSU value at IPL-time to be able to physically partition a complex during the life of the IPL. This is true regardless of whether the system is IPLed in single image or physically partitioned mode.

The requirement to always specify the proper RSU value may appear to penalize a complex operating in physically partitioned mode for long periods of time. However, such is not the case. Reconfigurable real storage is allocated logically as though the total amount of installed real storage was online at IPL-time. For example, if a 3084 Q64 is initialized in physically partitioned mode with  $RSU = 4$



(assuming a real storage increment size of 8MB), the operating system allocates the required 32MB of reconfigurable real storage from real storage addresses not owned by the initialized partition. If an installation subsequently configures the offline partition into the system to operate in single-image mode, the real storage owned by that partition is automatically designated as reconfigurable and the RSU requirement is met.

## RSU Parameter Specifications

An installation specifies a value for the RSU parameter according to the type of complex and the mode of operation of that complex. The recommended values to ensure the least system overhead and maximum capability for reconfiguration are as follows:

- Uniprocessor (for example a 3090 Model 180E) - specify RSU=0
- Nonpartitionable multiprocessor (for example a 3090 Model 200E) - specify RSU=0
- Single-image mode with the capability to physically partition the complex, specify RSU according to the following formula:

$$\text{RSU} = \frac{\text{Total amount of installed real storage}}{2 * \text{real storage increment size}}$$

On a 3084 the increment size is either 4M if the system has 64M or less of real storage, or is 8M if the system has more than 64M. On a 3090 the increment size is 2M on the Model 400, and 4M on Models 400E and 600E.

*Note:* Figure 2-10 lists RSU values calculated from real storage sizes and storage increment sizes for the 3084 and the reconfigurable 3090 models.

Processor Type	Installed Real Storage (MB)	Storage Increment (MB)	RSU Value
3084	32	4	4
3084	48	4	6
3084	64	4	8
3084	96	8	6
3084	128	8	8
3090 Mod 400	128	2	32
3090 Mod 400E	128	4	16
3090 Mod 400E	256	4	32
3090 Mod 600E	128	4	16
3090 Mod 600E	256	4	32

**Figure 2-10. Calculation of RSU Value for the 3084 and Reconfigurable 3090 Models**



## Chapter 3. Recovery

Recovery is the attempt by the hardware, operating system, the operator, or any combination of the three, to correct system malfunctions and return the system to a state in which it can do productive work. Some recovery actions are 'automatic'; that is, the hardware recovers from a malfunction without any intervention by the operator or any action by the operating system. Other recovery situations require overt actions by the system and/or the operator. For example, to keep the system in operation, the operator or the system can configure offline a failing component, such as a storage element, a processor, or a channel path. The system continues processing, possibly with some degradation.

The process of recovery includes the following:

- Hardware/operating system communication and corrective actions
- Operator/operating system communication and recovery actions

This chapter describes the following categories of hardware malfunctions:

- Central processing unit errors
- Service processor damage
- Storage errors
- Channel subsystem errors

For each of the preceding categories, the discussion includes the effect on system operation and the recovery actions taken (if any).

This chapter also includes some recommended operator actions for responding to such events as wait states, loops, spin loops, missing interrupts, etc.

Additionally, this chapter presents some recommendations for DASD maintenance and recovery.

# Hardware/Operating System Recovery Actions

The following categories of hardware errors are discussed:

- Central processing unit errors
- Service processor damage
- Storage errors
- Channel subsystem errors

When any of the preceding errors, except for some I/O errors, occurs, the hardware notifies the operating system with a machine check interruption. Machine check interruptions fall into one of three classes depending on the severity of the error. The classes are:

- soft (or repressible) errors - least severe type. Generally do not affect the operation of the task currently in control. Soft errors can be disabled (repressed) so that they do not cause a machine check interruption.
- hard errors - malfunctions that affect the execution of the current instruction or invalidate the contents of hardware areas (such as registers).
- terminating errors - malfunctions that affect the operation of a CPU.

Hard and terminating errors are also referred to as "exigent" errors.

Refer to Figure 3-1 for an illustration of how the operating system handles machine checks.

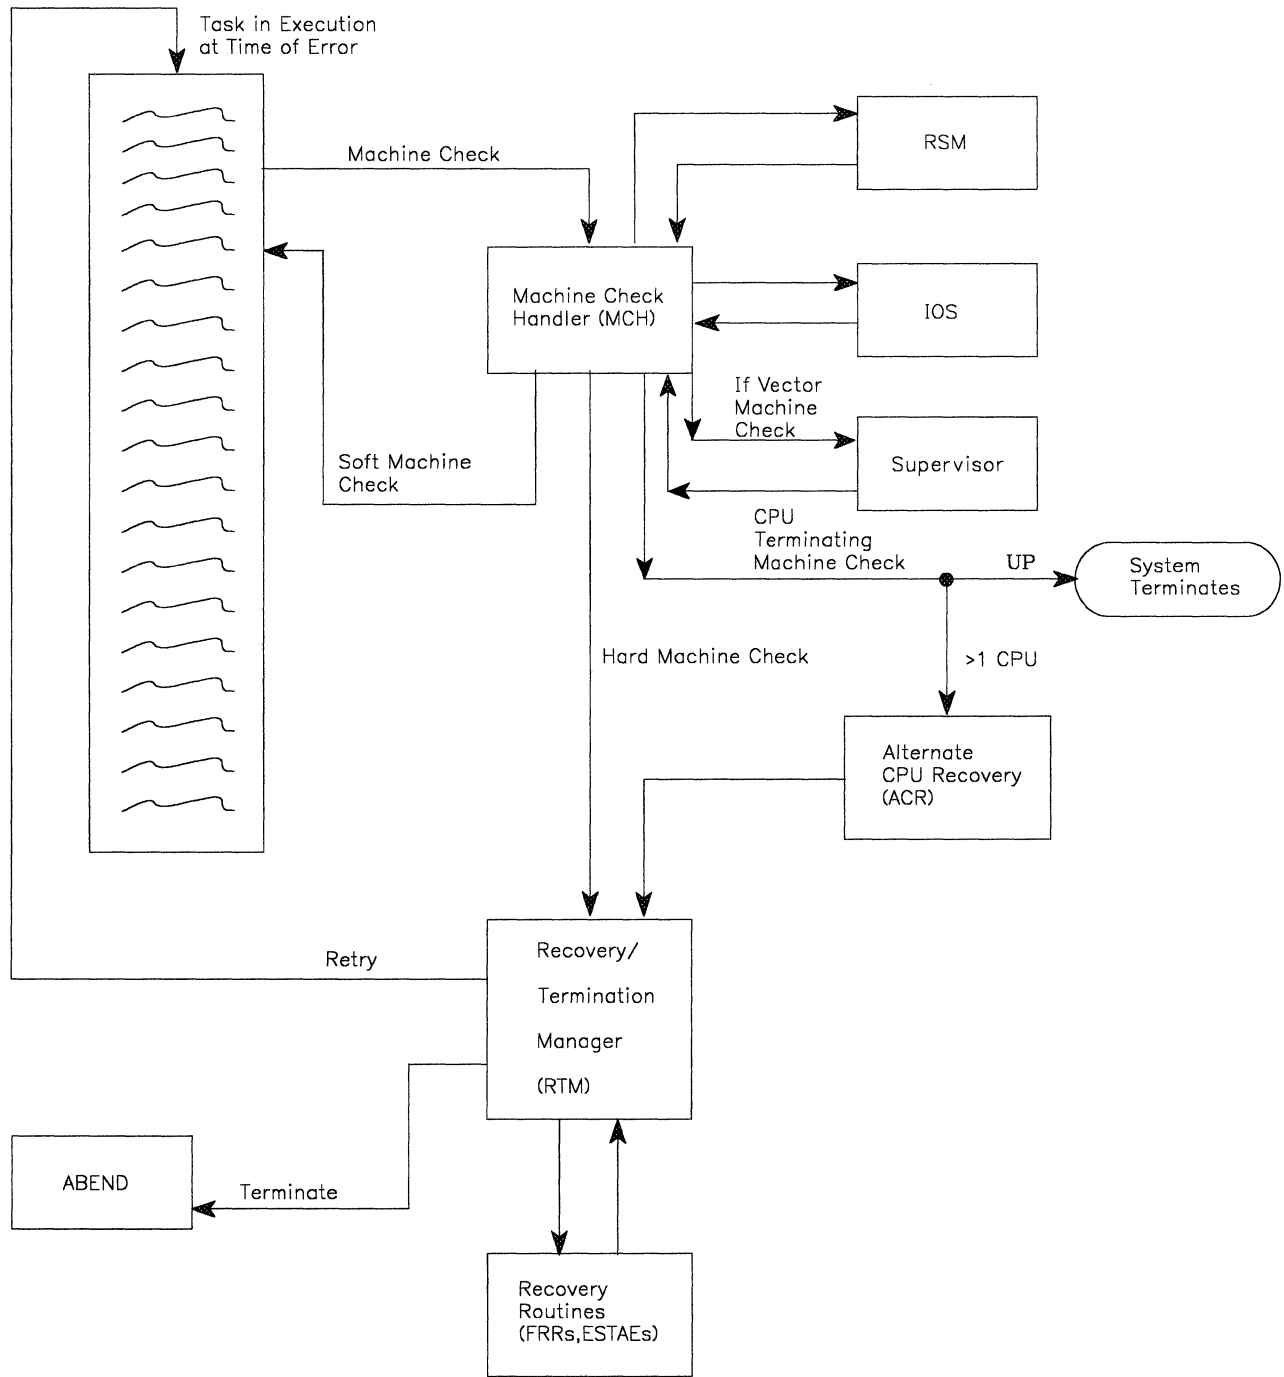


Figure 3-1. Operating System Handling of Machine Checks

## Information Provided with Machine Checks

When the hardware detects a failure, it stores the following types of information about the failure:

- The machine check interrupt code (MCIC), which contains information about the severity of the error, the time of the error (in relation to the current instruction stream), and an indication of whether the processor has successfully stored additional information about the error. The interrupt code is the major interface between the hardware and the operating system that uses the MCIC to determine what action to take.
- The model-independent fixed logout, which contains the values of the general purpose, floating point, and control registers at the time of error. It also contains the CPU timer and clock comparator values.
- The model-dependent extended logout, which contains diagnostic information needed by service personnel. The operating system does not use this information; it writes the information to SYS1.LOGREC along with the other information pertaining to the error.
- The machine check old PSW, which contains the PSW at the time of error.

The format and content of these storage areas are described in detail in *IBM System/370 Extended Architecture Principles of Operation*.

After storing the preceding types of information, the hardware gives control to the machine check handler (MCH) by loading the machine check new PSW. MCH gathers the information about the error into a buffer for later recording to SYS1.LOGREC. Next, MCH assesses the severity of the error by checking the machine check interrupt code and determines the appropriate course of action.

## Central Processing Unit Errors

Central processing unit (CPU) errors result from a malfunction of a hardware element such as a timing facility, instruction-processing hardware, or microcode. When a CPU error occurs, the recovery processing has, in general, two stages depending on the severity and type of error:

1. When possible, the hardware retries the failing operation a certain number of times. If the retry works, the hardware may issue a recovery machine check interruption (which is repressible) so that the operating system can record the error to SYS1.LOGREC. After recording, the operating system returns control to the interrupted task.
2. If the error is too severe for hardware retry or the retries fail, the hardware issues either a hard or terminating machine check interruption. The machine check handling routines determine the severity of the error and take the appropriate action that may range from terminating the interrupted task to terminating the entire system.

## Soft CPU Errors

The CPU errors that can result in a soft machine check are:

- System Recovery (SR) — a malfunction has occurred, but the hardware has successfully corrected or circumvented it.
- Degradation (DG) — continuous degradation of system performance has been detected.

The operating system does not inform the operator about the occurrence of soft machine checks until the “threshold” for a given type is reached. The threshold set by MCH for each soft machine check is four. The operator can change the threshold for either an SR or a DG machine check with the MODE command. When a threshold for a type of machine check is reached, MCH issues message IGF931E.

When a threshold is reached, the operator can allow the system to remain in quiet mode or enter the MODE command to reenable system recovery or degradation machine checks by setting a new threshold value. If the operator specifies RECORD=ALL for a particular type of machine check, the system does not enter quiet mode; it records all instances of system recovery and degradation machine checks in SYS1.LOGREC. The operating system issues message IGF931E when the number of machine checks is a multiple of the threshold. For example, if REPORT=3 is specified, message IGF931E appears after the third, sixth, ninth, twelfth machine checks, and so on.

Numerous IGF931E messages appearing on the operator’s console might indicate a performance degradation. In this case, the installation might want to configure offline the processor that is experiencing the errors. Maintenance on the offline processor can be done by service personnel as indicated by installation procedures.

## Hard CPU Errors

A hard machine check indicates that the current instruction could not complete. When MCH receives a hard machine check, it records the error on SYS1.LOGREC, issues message IGF972E, and passes control to the Recovery/Termination Manager (RTM) that either terminates the interrupted task or retries the interrupted task at a pre-defined retry point. Even though the task may be terminated, the system usually continues to run.

The CPU errors that cause hard machine checks are:

- System Damage (SD) — a malfunction has caused the processor to lose control over the operation it was performing to the extent that the cause of the error cannot be determined.
- Instruction Processing Damage (PD) — a malfunction has occurred in the processing of an instruction.
- Invalid PSW or Registers (IV) — the hardware was unable to store the PSW or registers at the time of error, as indicated by validity bits in the MCIC. Any error (even a soft machine check) associated with these validity bits is

treated as a hard machine check because the operating system does not have a valid address to use to resume operation. The error goes through recovery processing.

- Timing Facility Damage — damage to the TOD clock (TC), processor timer (PT), or clock comparator (CC) has been detected.

To overcome the effects of numerous hard machine checks, the MODE command allows the operator to define machine check thresholds for each type which, when reached, cause the failing processor to be configured offline by ACR. (The default threshold is five machine checks in five minutes.) Thus, the operator can control whether, and to what extent, the system monitors the frequency of hard machine checks, and can define a separate threshold and time interval for each.

If installation thresholds have been established but numerous IGF972E messages are generated, (RECOVERY INITIATED FOR PROCESSOR FAILURE ON CPUx), the installation should consider configuring CPUx offline prior to the expiration of the threshold.

## Terminating CPU Errors

A terminating machine check occurs when the operating system or the hardware considers a failure severe enough that a processor cannot continue operation.

In a UP environment, the operating system terminates with a disabled wait state (such as A01, A26), and issues the following message:

```
IGF910W UNRECOVERABLE MACHINE FAILURE, RE-IPL SYSTEM
```

In a multiprocessor environment, the action taken is as follows:

- If the hardware determines that a processor cannot continue operation, it places the processor in a check-stop state and attempts to signal the other processor(s) by issuing a malfunction alert (MFA) external interrupt. The hardware issues an MFA when:
  - it cannot store the machine check logout data about the error
  - it cannot load the machine check new PSW
  - it is disabled for machine checks
- If the operating system determines that a processor cannot continue operation, it attempts to signal the other processor(s) by issuing a SIGP emergency-signal (EMS) instruction to cause an external interrupt. The operating system issues an EMS when:
  - MCH is processing one machine check when another machine check occurs that cannot be handled
  - A hard-machine-check threshold (installation option), established by issuing the MODE command, has been reached
  - Channel subsystem damage is detected
  - The content of the MCIC is invalid



When a processor receives either an MFA or EMS external interruption (relative to the preceding stated conditions), the External Interruption handler gives control to MCH. MCH, in turn, invokes Alternate CPU Recovery (ACR) processing which takes the malfunctioning processor offline and initiates recovery processing for that processor.

In a multiprocessor environment, an MFA or EMS is received by all the other online processors. On the first processor to receive the signal, MCH tests and sets a flag before starting to process the error. When the other processors receive the interruption, MCH on those processors sees that the error is already being processed and returns to the interrupted task.

## Vector Facility Recovery

If the Vector Facility on a 3090 has a malfunction, it will present one of two types of machine check:

- Vector facility source error
- Vector facility failure

These machine checks are represented by two newly defined bits in the machine check interrupt code (MCIC).

### Vector Facility Source Error Machine Check

A vector facility source error is a hard machine check presented with the PD bit and the VF source bit set in the MCIC. Vector facility source errors are not counted as processor damage machine checks for threshold purposes against the CPU but are counted toward a separate vector source (VS) threshold count. The VS threshold can be set by the operator via the MODE command VS parameter.

For a vector facility source error, the MCH performs these steps:

- Tries to save the vector environment for possible retries during RTM processing (that is, during handling of an '0F3' ABEND.)
- Increases by one the threshold count of vector source machine checks.
- Routes the current work to RTM with an '0F3' ABEND code, for possible recovery or ABEND processing.
- If the threshold of vector source machine checks has not been reached (default is five in five minutes), the MCH takes no further action.
- If, however, the vector source threshold has been reached, the MCH indicates in the common system data area (CSD) that a Vector Facility is logically offline, requests the service processor to physically disconnect the Vector Facility, and issues this message:

```
IGF970I VFn NOW OFFLINE. UNRECOVERABLE ERROR  
DETECTED.
```

## Vector Facility Failure Machine Check

A Vector Facility failure is a soft machine check presented with the Vector Facility failure bit set in the MCIC. In this case, if the interrupted task is a vector task, its vector status (such as vector registers and clock) are invalidated, and the Vector Facility (but not the CPU) is taken offline. The interrupted unit of work is terminated only if it attempts to issue another vector operation. In this case, the work is terminated because, even if there are other operational Vector Facilities, the user's vector status at the time of failure cannot be guaranteed.

## Alternate CPU Recovery (ACR)

ACR is a function that is initiated on an operative CPU when that CPU receives a signal that another CPU has had a terminating error. ACR has two major functions:

1. To configure offline the malfunctioning CPU
2. To initiate the release of system resources held on the malfunctioning CPU

If the failing CPU has a Vector Facility, the Vector Facility is also taken offline.

ACR initiates the release of any resources held on the failing CPU by invoking RTM which initiates the functional recovery routines (FRRs) of the work on the failing CPU. ACR allows the operating system to continue its normal operation on the remaining CPU(s) although the task that was interrupted by the error on the failing CPU may be terminated.

When ACR is complete, it sets up message IEA858E stating that ACR is complete and identifying the CPU that was configured offline. At this point, the operator can try to configure the failing CPU back online using a CONFIG CPU(x),ONLINE command. The configuration 'online' may, or may not, be successful depending on the error that caused the CPU to be configured offline.

Some hardware malfunctions may cause a subsequent CONFIG CPU,ONLINE command to that CPU to fail, or may cause the problem to reoccur when the CPU is brought back online. In these cases, hardware service is necessary before the CPU can be successfully brought back into the system.

However, if a CPU was configured offline because a threshold was reached or because of an operating system problem, a subsequent request to configure the CPU back online may work. Since online processing does hardware resets and rebuilds the CPU-related control blocks, the cause of the problem may be eliminated.

On the 4381 Model 3 and Model Group 14 processors, the channel paths physically attached to the failing CPU will be configured offline during ACR. In addition to a malfunction alert (MFA) external interrupt, the hardware also presents a channel report pending machine check interrupt which indicates permanent errors for all channel paths attached to the failing CPU. In response to the channel errors, MVS physically takes offline all channel paths attached to the failing CPU. If the failing CPU is returned online in response to a CONFIG CPU,ONLINE command, the operator should also use the CONFIG CHP,ONLINE command to try to configure online the channel paths.

## Terminating Errors on Multiple CPUs

In a multiprocessor environment, failure of some hardware elements may cause a terminating error on more than one CPU. It is also possible that a terminating error may occur on a CPU while ACR is still processing a terminating error on another CPU. In either case, MCH issues message IGF973W indicating that an ACR is already in progress and puts the entire system into a '050' nonrestartable wait state.

## Service Processor Damage

In a 308x complex (excluding 3081D), when the system detects that the service processor is failing, the system:

- Generates a service processor damage machine check
- Informs the subsystems (such as IMS) so that they can perform an orderly shutdown

In a 3081D complex *only*, when the system detects the unique hardware malfunction called 'service processor stall', it:

- Generates a service processor damage machine check
- Logically but not physically configures one processor offline

In all 308x complexes the system issues message IEA470W so that the operator can perform an orderly shutdown of the system. Processing can continue until functions of the service processor are required; at that time, the system becomes inoperative. To recover, the operator performs an IML.

*Note:* This is also true for a 3084 complex in single image mode if the backup service processor is not available, or if a successful switchover to the backup does not occur when the active service processor goes down.

## Storage Errors

The hardware detects and corrects storage errors where possible. The machine check handler (MCH) is informed of the error by a machine check interrupt, and MCH invokes recovery routines through RSM. If the storage error is detected during an I/O operation, however, the operation is terminated with either a channel data check or a channel control check, depending on whether the error was encountered during data transfer or CCW/IDAW fetching. No machine check interrupt is generated in this case. Error recovery procedures (ERPs) recover from this type of error.

## Soft Storage Errors

The soft storage errors are system recovery (SR) errors with the 'storage error corrected' flag set in the MCIC to indicate that the storage controller was able to repair the error.

When a 'storage error corrected' condition occurs, MVS attempts to stop using the affected frame. This action eliminates performance degradation that would result from the hardware's correction of later occurrences of the same error. It also minimizes the chance that the same problem will later occur as a 'storage error uncorrected'. If the frame contains pageable data, MVS moves that data to another frame, and the original frame is marked offline. If the data in the frame cannot be moved, the frame is marked 'pending offline', and is subsequently taken offline if the frame is released or if its contents are made pageable. (Before MVS takes a frame offline, it tests the frame and if it has no errors, the frame is returned to available status.)

The threshold for SR machine checks affects the ability of MVS to deal with 'storage error corrected' conditions. When the threshold for SR machine checks is reached, MVS disables SR machine checks. This action prevents subsequent 'storage error corrected' from being presented. MVS then does not take any action to remove the affected frame.

Because the default threshold for SR machine checks is 4, you should consider using the MODE command to raise the SR threshold to 50 for all the CPUs. The increased SR threshold allows MVS recovery functions to handle more 'storage error corrected' for any given IPL. If the revised threshold is eventually reached, MVS issues message IGF931E to inform the operator, and disables this class of machine check.

You can raise the SR threshold to 50 by means of this operator command:

```
MODE SR,RECORD = 50
```

Note that although this recovery technique applies to all systems supported by MVS/XA, it is especially pertinent to 3090 systems. Because the 3090 performs double-bit error correction, a larger percentage of its storage errors is presented as 'storage error corrected'.

## Hard Storage Errors

This section deals with these types of hard storage errors:

- Storage error uncorrected — indicates that the hardware could not repair a storage error.
- Key in storage error uncorrected — indicates that the hardware could not repair a storage key that was in error.

When a hard storage error occurs, MCH invokes the real storage manager (RSM) to attempt recovery. If RSM cannot repair the error, it either takes the storage frame (4K) offline or marks it pending offline (which means that RSM will take the frame offline when the frame becomes free). MCH processing issues message IGF971E that indicates which processor is handling the error; and if possible, the address of the storage. If the operator receives message IGF971E for numerous storage addresses within an identifiable range, configuring that range offline using a CONFIG STOR command may be warranted.

Because a 'storage error uncorrected' condition represents the potential loss of critical data, MVS in most cases will terminate the affected unit of work. If the recovery routines in this termination complete successfully, and cause the freeing of the affected storage frame, the frame is marked offline and system processing continues. The recovery processing, however, could try to refer to the storage that originally caused the machine check, thus causing further errors. Such action could result in the PD threshold for machine checks being reached, thus taking a CPU offline.

You can reduce the chance of having a storage error take a CPU offline by using the following MODE command to raise the threshold for PD machine checks on all CPUs to 25 machine checks in 5 minutes:

```
MODE PD,RECORD=25,INTERVAL=300
```

## Effects of Storage Errors

Errors in critical areas of storage may cause the hardware system or the operating system to become inoperative. Those areas of storage and the effect of an error are as follows:

**Hardware Storage Area (HSA):** An uncorrectable storage error in the HSA causes the system to enter a check-stop state. The system can be recovered by these two steps:

1. Power-on-reset or SYSIML CLEAR
2. IPL

**HIGH SPEED BUFFER:** A processor high speed buffer error can result in the loss of the processor and possibly the system. The real storage frame corresponding to any changed data in the high speed buffer is marked with an uncorrectable storage error. Since the high speed buffer may contain critical system data, recovery may require an IPL.

**NUCLEUS:** A storage error in nucleus pages requires an IPL for recovery. If the IPL fails, recovery requires either a power-on-reset or SYSIML CLEAR, followed by IPL.

**LPA/SQA/LSQA:** A storage error in SQA could have the same effects as a nucleus storage error.

For a storage error in LPA, the operating system handles recovery. Normally, only the associated job is terminated with the remainder of the system unaffected.

## Storage Element Failure

If a storage element fails and sufficient usable storage is available, the operator can recover by:

1. Releasing the configuration (via the CONFIG frame), then deselecting the failing storage element (if this function is supported).
2. Issuing a storage validation function - either SYSIML CLEAR or a power-on-reset.
3. Re-IPLing.

## Channel Subsystem Errors

If the channel subsystem fails, the hardware generates a 'channel subsystem damage' machine check interrupt. MCH invokes IOS to handle the interrupt.

IOS puts the entire system into an A19 nonrestartable wait state and issues message IOS019W.

## Channel Report Words (CRWs)

When the channel subsystem detects an error, it

- builds a CRW that describes the error
- queues the CRW for retrieval by IOS
- generates a machine check interrupt with 'CRW pending' indicator set in the machine check interrupt code (MCIC).

MCH invokes IOS to handle the interrupt.

IOS retrieves the CRW by issuing the Store CRW (STCRW) instruction and records the CRW in SYS1.LOGREC. The CRW contains a code that indicates the source of the error: the channel path, the subchannel, channel configuration alert, or the monitoring facility. (For additional information on CRWs, see *System/370 Extended Architecture Principles of Operation*.)

## Channel Path Recovery

If the CRW indicates that a channel path caused the machine check, IOS attempts to recover the channel path or route I/O down an alternate channel path. (If multiple CRWs indicate errors on different channel paths, a failure in the hardware elements common to those channel paths may be indicated.)

The channel path conditions identified in the CRW are:

- A permanent error on the channel path; a system reset to the channel path has not been done (reserved devices are still reserved and the path groups for devices that have dynamic pathing active are still intact)
- A permanent error on the channel path; a system reset to the channel path has been done (devices reserved on this path are no longer reserved and the path groups for devices that have dynamic pathing active are not intact)
- The channel path is in a terminal error condition
- The channel path is recovered (initialized); a system reset to the channel path has been done (devices reserved on this path are no longer reserved and the path groups for devices that have dynamic pathing active are not intact)

The channel path conditions fall into two categories: expected and unexpected. An expected channel path condition occurs as a result of a previous recovery action taken for an unexpected channel path error, and indicates the result of the action. An unexpected channel path error occurs with no warning.

The permanent errors can be either expected or unexpected. The terminal error condition is only unexpected; it is never the result of a previous recovery action. The initialized condition can only be expected; it means that a previous recovery action has successfully recovered the channel path and the channel path is available for use.

A permanent error condition means that the channel path cannot be recovered.

A terminal error condition means that the channel path is not permanently lost but cannot be used in its current condition. IOS attempts to recover the channel path by issuing the Reset Channel Path (RCHP) instruction to initiate hardware recovery processing. This action by IOS results in another CRW with “expected” error status.

A recovered or initialized condition means that a previous recovery action has been successful in recovering the channel path.

During channel path recovery processing, IOS communicates with the operator by issuing IOSxxx messages. These messages may be issued to:

- request a specific operator action during or after recovery processing
- inform the operator of the recovery status
- inform the operator of the actions taken by IOS

## Channel Path Alert Conditions

IOS communicates with the operator when two other indicators are set in a CRW — ‘configuration alert temporary’ or ‘channel path temporary’. In either case, IOS performs no recovery processing.

- For ‘channel path temporary’, IOS issues message IOS162A to inform the operator that the channel subsystem could not identify the device requesting service.
- For ‘configuration alert temporary’, IOS issues message IOS163A to inform the operator that the channel subsystem could not associate a valid subchannel with the device requesting service.

## Subchannel Recovery

If the CRW indicates that a subchannel caused the machine check, IOS examines the error recovery code in the CRW. If the CRW indicates that the subchannel is available, the channel subsystem has recovered from a previous malfunction. I/O functions in progress and presentation of status by the device have not been affected. No program action is required.

If the CRW indicates that the subchannel is “installed parameter initialized”, IOS determines if the device associated with the subchannel is still valid. If it is, IOS reenables the subchannel. If, however, the device related to the subchannel is not valid, IOS marks the device as unusable and issues message IOS151I.

## Monitoring Facility Recovery

IOS does not do any recovery for the monitoring facility. IOS schedules the system resource manager’s (SRM’s) recovery routine.

## I/O Errors

The channel subsystem generates an I/O interrupt for the following I/O error conditions:

- if the device is not operational on any path
- for any device status errors (for example, unit check)
- for any subchannel status errors (for example, interface control check, channel control check)

IOS processing of the interrupt may be:

- invoking a driver exit
- interfacing with attention routines and volume verification processing
- invoking a device-dependent ERP for error recovery processing



- unconditional reserve processing
- redriving the I/O request on a channel path other than the one that generated the interrupt
- entering a restartable wait state (115) if a paging device is not operational and waiting for the operator to RESTART a processor. (The RESTART reason is ignored.)
- issuing message IOS050I to inform the operator that a subchannel status error occurred.

## Master Console Failure

If the master console becomes unavailable (cannot be accessed), the operating system normally switches automatically to an alternate without a re-IPL. If the alternate master console cannot be used, MVS tries to write to the hardware system console. (Refer to "Processing Messages at the System Console" on page 3-23.)

## Missing Interrupts

A missing interrupt condition exists when IOS expects an interrupt that does not occur within a specified time interval. For example, the IBM-default time interval between checks for missing interrupts varies from 15 seconds for paging DASD devices (other than 3330V) to 12 minutes for the 3330V and 3851 devices. An installation can define, in parmlib member IECIOSxx, time intervals for all devices in its I/O configuration and override the IBM-supplied defaults. (Refer to *SPL: Initialization and Tuning* for additional information.)

The missing interruption handler (MIH) determines whether an expected interrupt has failed to occur within a specified time interval. Some possible missing interrupt conditions are:

- an idle UCB with I/O requests queued that should be started
- an outstanding I/O operation that should have completed
- an outstanding mount for a tape or disk

If an expected condition does not occur, MIH informs the operator and tries to correct the situation before system performance is affected. In addition, missing interrupt incidents are recorded in SYS1.LOGREC.

For missing interrupts, MIH issues message IOS071I or IOS076E to inform the operator of the particular condition that exists. Message IOS076E also describes the operator actions required to reset some of the conditions.

*Note:* If there are missing interrupts on the devices that contain SYSRES or page volumes, the operator may not receive any message, because the MIH message writer and the CONSOLE address space (Comm Task) are pageable. The operator can learn about the missing interrupts by initiating the RESTART function with REASON 1.

For recurring missing interrupts, MIH issues message IOS075E together with message IOS076E or IOS077E to inform the operator of the recurring condition on a particular device.

## Unconditional Reserve/Alternate Path Recovery (APR)

Alternate path recovery (APR) permits recovery from control unit or channel path failures that cause a DASD device, or a string of DASD devices, to no longer be accessible to the system. APR is performed only after IOS guarantees ownership of the device; that is, the device is reserved to this system.

APR issues an UNCONDITIONAL RESERVE command along each online path, one at a time, to the device. If an alternate path is available, APR issues message IEA428I. If no alternate paths are available, APR issues message IEA429I. IOS then boxes the device and terminates all subsequent requests to that device with a permanent error.

If IOS cannot guarantee ownership of the device, it issues message IEA427A, which gives the operator three recovery options. However, before replying with one of the options, the operator should ensure that the device is owned (reserved) to this system.

## Hot I/O

A Hot I/O condition occurs when a device, control unit, or channel path causes continuous unsolicited I/O interrupts. If the Hot I/O condition goes undetected, it can cause the system to enter a loop or it can exhaust the system queue area (SQA). IOS attempts to recover from a Hot I/O condition so that a re-IPL is not required. For diagnostic purposes, IOS records all Hot I/O incidents on SYS1.LOGREC.

IOS first tries recovery at the device level by issuing the Clear Subchannel instruction in an attempt to clear the Hot I/O condition. If the condition is cleared, processing continues normally. If the condition persists, the next recovery action is determined by one of the following:

- the parameters the installation defined in parmlib member IECIOSxx for Hot I/O recovery
- operator response to the appropriate Hot I/O message or restartable wait state:
  - IOS110A or wait state 110 (non-DASD),
  - IOS111A or wait state 111 (unreserved DASD), or
  - IOS112A or wait state 112 (reserved DASD).

Because IPLs related to Hot I/O are generally caused by incorrect operator actions, an installation should use the IECIOSxx parmlib member to make Hot I/O recovery more automatic and reduce the need for immediate operator intervention. The following sample parameters, when defined in the IECIOSxx parmlib member, tell IOS how to handle automatic recovery from Hot I/O for three classes of devices: non-DASD, non-reserved DASD, and reserved DASD. (Additional information on Hot I/O parameter specification is discussed in *SPL: Initialization and Tuning*, and *SPL: System Modifications*.)

### Sample Parameters for Hot I/O Recovery in Parmlib Member IECIOSxx

The following entries are an example of how to specify the Hot I/O recovery parameters in the IECIOSxx parmlib member. As of SP2.2.0, the values shown are also the IBM default values.

HOTIO	DVTHRS=100	
		Specifies 100 repeated interrupts as the threshold for IOS recognizing the condition.
HOTIO	DFLT110=(BOX,)	
		Box the non-DASD device on the first occurrence of this condition and prompt the operator for the recursive condition.
HOTIO	DFLT111=(CHPK,BOX)	
		Attempt channel path recovery for non-reserved DASD on first occurrence. On recursion, box the device.
HOTIO	DFLT112=(CHPK,OPER)	
		Attempt channel path recovery for reserved DASD on first occurrence, but prompt the operator for the recursive condition.

### 3880/3380 Considerations

The 3880/3380 AA4 is designed to allow concurrent maintenance at the storage director (SD) level. Prior to attempting concurrent maintenance, all paths from all processor complexes through the failing SD to the devices must be varied offline. Failure to vary all paths offline may result in various error symptoms, including Interface Control Checks, Path Inoperative conditions, and out-of-sync conditions between the 3380 array and the operating system.

Prior to returning a repaired SD to the system, an IML of the SD or a power down-up sequence must be performed to establish a correct copy of the 3380 array for the repaired SD. The operator should issue VARY PATH ONLINE commands for all paths to all devices through the repaired SD.

### 3380 Enable/Disable Switch

The Enable/Disable switch on the 3380 A box should NEVER be set to 'disable' when any paths to the device are online. Setting the switch to 'disable' could cause an 'out-of-sync' condition between the array and the operating system. This out-of-sync condition can occur whenever the dynamic path group information

maintained in the 3380 'A' box is reset without notification to the operating system. Any of these operator actions could cause an out-of-synch condition:

- IMLing the 3880 control unit
- Disabling the 3880 interface switch
- Disabling the 3380 interface switch

In addition, certain 3880/3380 hardware failures can affect the arrays.

### **Recovery from an Out-of-Synch Condition**

Array out-of-synch conditions may be indicated by missing interrupts or path-inoperative I/O errors. MVS/XA provides automatic detection and recovery through the dynamic pathing validation support. This code detects potential out-of-synch conditions (e.g., an MIH condition), and then validates the physical path group information. If the dynamic pathing validation code finds a mismatch between the hardware and software path group information, it invokes recovery to rebuild the dynamic path selection arrays.

If MVS/XA cannot rebuild the arrays, the operator will usually see repeated occurrences of IOS077E messages, with a 'START PENDING' insert, on all processor complexes sharing the 3380. To attempt to recover from an out-of-synch condition, the operator must issue a VARY device ONLINE command on the system where the out-of-synch condition exists. VARY device ONLINE commands issued on systems that do not have the out-of-synch condition will not cause additional problems, but will not re-synchronize the array. If the VARY device ONLINE commands do not re-synchronize the array, a re-IPL of all sharing processor complexes will re-synchronize the array.

### **DASD Maintenance and Recovery**

DASD can experience failures such as defective disk surfaces, drives, and actuators. When these failures occur, data becomes inaccessible to the operating system and could be lost. To prevent the loss of the data, an installation should consider the use of the EREP System Exception Report in conjunction with Device Support Facilities (DSF) to monitor possible error conditions and correct any before they cause outages. For additional information on the use of the System Exception Report and DSF, refer to the following publications:

*IBM Disk Storage Management Guide: Error Handling  
Device Support Facilities  
EREP User's Guide and Reference*

When a DASD error does occur (for example, a defective track, volume or actuator), an installation can use Data Facility Data Set Services (DFDSS) to retrieve the data from the defective areas and copy it to a back-up DASD. Refer to *Data Facility Data Set Services User's Guide and Reference* for detailed information.

## Operator Recovery Actions

This section deals with recovery actions available to the operator. It includes these topics:

- Recovery by CPU restart
- Continuing a vector job if a Vector Facility is offline
- Hardware instruction tracing
- Recovery from wait states
- Excessive spin loop recovery

### Recovery by CPU Restart

The operator can initiate recovery from some system incidents, such as loops and uncodded wait states, by issuing a restart to the processor that has the problem. The RESTART REASON that is entered as part of the restart process directs MVS to perform one of two recovery actions:

1. RESTART REASON 0 - Message IEA500A is displayed on the master console to identify the current unit of work. The operator can reply either RESUME to allow the current unit of work in progress to continue or ABEND to terminate the current unit of work with a X'071' abend.

If the operating system cannot communicate with the master console (or its alternate) to issue message IEA500A, it terminates the current unit of work with a X'071' abend.

2. RESTART REASON 1 - the operating system:
  - Interrupts the current unit of work
  - Detects and attempts to repair errors in critical system areas
  - Writes a record to SYS1.LOGREC (with completion code X'071' and reason code 4) when repair actions were taken.
  - Reports the results of some of the actions taken in message IEA501I
  - Returns control to the interrupted unit of work

Refer to *System Commands* for additional information concerning the restart function.

*Note:* Restart of the CPU in a restartable wait state ignores the restart reason.

## Continuing a Vector Job If a Vector Facility is Offline

If a Vector Facility goes offline while a job is running on the Vector Facility, the job is redispached to another Vector Facility, if there is one available. If, however, there is no other Vector Facility available, the job is swapped out and message IRA700I is issued:

```
IRA700I jobname xxxxxxxx WAITING FOR AVAILABILITY OF VF
```

In this case the operator has these choices:

- Issue CONFIG VF(x) to bring the Vector Facility back online. The job is then swapped in.
- Cancel the job.
- Do nothing, in which case the job may time out (depending on the SMF Job Wait Time specified) and be cancelled by MVS.

With no Vector Facility online, other jobs that try to do a vector operation will be swapped out. If no Vector Facility is brought online and these jobs remain swapped out for the interval specified as the SMF Job Wait Time (JWT), the jobs will be terminated with ABEND code 522. You can prevent these “time-outs” of swapped-out vector jobs by specifying TIME=1440 on the appropriate JOB or EXEC statement, or by means of a user-provided exit routine.

## Hardware Instruction Tracing (Loop Trace)

When a loop occurs on a 308x, the operator can activate instruction tracing only on the selected target processor from the system console. But *all* processors in the complex are left in the manual state at the completion of the trace. On the 308x the tracing is called a “loop trace”. The trace records a pre-set number of instructions. The recorded data, included in a dump taken after the completion of the trace, can be used for problem determination.

On a 3090 the tracing is called an “instruction trace”. Tracing occurs in a round-robin sequence on *all* 3090 CPUs that are configured online, starting with the target CPU. At the completion of the instruction trace, the CPUs are *not* left in the manual state. They are in the state they were in when the trace was started.

To resume normal operation after the completion of the loop trace, the operator must START all processors.

## Recovery from Wait States

A system wait state is entered when bit 14 of the current PSW is set to 1 with the right half of the PSW containing the wait state code. Wait states indicate that a processor is not currently executing instructions. Three types of wait states exist:

- disabled wait states
- enabled wait states
- uncoded wait states

### Disabled Wait States

Disabled wait states are used to:

- terminate the system when an unrecoverable error is detected (non-restartable wait states).
- communicate to the operator a condition that requires operator action when normal communication by a message is not possible. In this case the wait state is restartable.

The operator should refer to *System Codes* to determine if a disabled wait state is restartable. If it is, the operator should perform the indicated actions.

### Enabled Wait States

Enabled wait states usually indicate that the system is waiting for:

- work
- an operator action or response
- a system resource to be freed

To recover from an enabled wait state, the operator should follow the procedures documented in *System Codes* in the chapter "Uncoded Wait States".

### Uncoded Wait States

When a condition arises such that the right half of the PSW does not match any of the wait state codes documented in *System Codes*, an uncoded wait state has occurred. To recover from an uncoded wait state, the operator should follow the procedures documented in *System Codes* in the chapter "Uncoded Wait States".

## Spin Loop Recovery

A loop is a sequence of instructions that is being repeatedly executed by a processor. A processor in a loop may control resources needed by other processor(s). This usually results in degradation of system performance because the other processor(s) may enter a wait state or a loop while the resources are unavailable. Some possible characteristics of a loop are:

- CPU utilization remains at 100% on the System Activity Display (SAD)
- Wait indicator is off

To stop or document the loop, the operator should follow the procedures documented in *System Codes* in the chapter "Loops."

## Spin Loops

A spin loop is a situation in which one processor in a multiprocessor environment is unable to communicate with another processor or requires a resource currently held by another processor. The processor that has attempted communication is the 'detecting' or 'spinning' processor. The processor that has failed to respond is the 'disabled' or 'failing' processor.

The 'detecting' processor attempts communication with the 'disabled' processor for a period of time that is determined by an operating system threshold called the "excessive spin length factor". Because the execution rate of a processor depends on the actual instructions executed, the time required to exceed the threshold depends on various situations.

When the 'detecting' processor exceeds the threshold, a Spin Loop Timeout situation exists. The detecting processor invokes the Excessive Spin Notification Facility to notify the operator of the spin loop. The notification is usually a message (IEE331A or IEA490A). If the message cannot be issued, the Excessive Spin Notification Facility loads a '09x' restartable wait state on the 'detecting' processor.

During system error recovery, valid reasons exist for long periods of disabled processing. The threshold value was chosen to be greater than the maximum disabled time attributable to most valid system loops. However, a temporary spin loop condition, which does not recur after the retry option, indicates a single processor was disabled for more than the threshold time. Frequent temporary spin loop conditions may indicate a hardware or software problem; therefore, an installation should determine the cause and correct the problem.

## Operator Notification

The operating system attempts to notify the operator of a spin loop condition by sending the appropriate DCCF message to the master console. If the message cannot be written to the master console, the operating system then tries to send it to the alternate. If MVS can't access the alternate, it tries to communicate via the system console. If the operating system cannot access any of these consoles, it loads a '09x' restartable wait state.

If the master console (and its alternate) are configured as recommended in Chapter 2, the probability of the operator receiving the message and being able to respond prior to operator response timeout (approximately two minutes) is very high. However, if another device is attached to the same control unit as the master console, the following could occur:

- Operating system sends the message
- Before the operator can respond to the message, another device attached to the control unit is accessed



- Operator tries to respond to the message but control unit is busy
- Control unit remains busy longer than operator response timeout
- The operating system tries to communicate via the system console.
- If unsuccessful, the operating system loads X'09x' restartable wait state

Also, if the operator does not respond to the message prior to operator response timeout, the operating system tries to write the message to the system console. If not successful, the operating system loads a '09x' restartable wait state.

### Processing Messages at the System Console

When MVS is successful in writing an excessive spin-loop (or any DCCF) message to the system console, the alarm sounds and the existing screen image on the system console is replaced with an indication that an SCP message is pending. The operator displays the message by entering F SYMSG (on a 308X) OR F SCPMSG (on a 3090).

The *response* line appears below the SCP message but does not contain the characters *R0*, as would the response line for a message on an MVS console. Do **not** enter the *R0*; simply enter the response indicated for the message.

There is no "timeout" interval for entering the reply to a DCCF message on the system console. The message remains pending on the System Message Facility screen image until the operator replies to it.

### ACR Considerations

ACR is one of the recovery options available for spin loop conditions. An operator reply of ACR causes the current unit of work on the failing processor to be terminated with a '0F3' abend, and the failing processor to be configured offline.

ACR processing can resolve the cause of most spin loops, because:

- Configuring offline the disabled processor releases the resource waited for by the 'spinning' processor.
- The attempt to signal the disabled processor ceases when it is configured offline.

Except for those spin loops caused by the SIGP circuitry, an operator can usually configure online the offline processor by issuing a CONFIG CPU(x), ONLINE command after ACR is complete. For additional information, refer to the previous topic, "Alternate CPU Recovery (ACR)".

The option to use ACR to take a CPU offline does not necessarily mean that the problem causing the excessive spin loop is related to a CPU. Most excessive spin loops result from problems in either software or non-CPU hardware. In these cases, an ACR response provides recovery through executing the software recovery routines set up for the the CPU being configured offline.

These routines try to recover by:

- Terminating the work that was active on the “failing” CPU
- Freeing resources held on that CPU
- Deleting or refreshing queues and control blocks

In many cases these actions will resolve the cause of the excessive spin. Therefore, unless there’s another indication of a CPU problem, the operator should configure the CPU back online when ACR is complete.

## Recovery Actions

Figure 3-2 indicates the recommended recovery actions for each text insert in messages IEE331A and IEA490A and the associated ‘09x’ wait state.

Associated ‘09x’ Wait State	IEE331A Message Insert	Primary (See Note) Action	Secondary Action
1	RISGNL RESPONSE	Continue	ACR
2	LOCK RELEASE	Continue	ACR
N/A	RESTART RESOURCE	Continue	ACR
5	ADDRESS SPACE TO QUIESCE	Continue	ACR
7	INTERSECT RELEASE	Continue	ACR
E	SUCCESSFUL BIND BREAK RELEASE	Continue	ACR
<hr/>			
	<b>IEA490A</b> Message Insert		
3	(NOT OPERATIONAL)	ACR	
8	(EQUIPMENT CHECK)	ACR	
9	(OPERATOR INTERVENING)	Start Stopped Processor, Continue	
A	(CHECK STOP)	ACR	
B	(NOT READY)	ACR	
C	(BUSY CONDITION)	ACR	
D	(RECEIVER CHECK)	ACR	

*Note:* The “continue” option consists of either the reply ‘U’ to message IEE331A or restarting the processor that is in the ‘09x’ wait state.

**Figure 3-2. Recovery Actions for Each Message Insert or Wait State**

### *Example of Recovery Procedure for Spin Loop Message*

IEE331A PROCESSOR(y) IS IN AN EXCESSIVE DISABLED SPIN LOOP WAITING FOR LOCK RELEASE. REPLY ‘U’ TO CONTINUE SPIN, OR STOP PROCESSOR(n) AND REPLY ‘ACR’. (AFTER STOPPING THE PROCESSOR, DO NOT START IT.)

Reply ‘U’ on the console displaying the message to continue in the spin loop. If the message recurs, reply ACR.

## Recovery for X'09x' Wait State

The symptoms may be:

- Audible alarm sounds on system console
- Message display ceases on both MCS and JES3 consoles
- System Activity Display (SAD) shows one processor with 0% CPU utilization and all other processors at 100% utilization

The possible recovery options are:

1. To continue in the spin loop, restart the CPU in the '09x' wait state (The restart reason is ignored on a restart of a CPU in a '09x' wait state.)
2. To initiate ACR at the system console:
  - Stop all processors
  - Select CPUy for the purpose of displaying and altering storage (y = id of CPU in '09x' wait state)
  - Display location '30E' in PSA of CPUy
  - Store 'AA' in location '30E'
  - Identify the failing processor. This can be done in either of two ways:
    - a. The failing processor's logical ID (e.g., 4n) is in the sixth byte of the '09x' wait state PSW. For example, if CPU 0 was in an '092' wait state because of a lock held on CPU3 (the failing processor), the wait state PSW for CPU 0 would be X'000A0000 00430092.
    - b. Display location '40C' in the PSA of CPUy. Display contents of the address obtained from location '40C' to identify the failing processor: for example, 00000000 = CP0, 00000001 = CP1, 00000002 = CP2, 00000003 = CP3, etc.
  - Start all processors except the failing processor and the processor in the '09x' wait state
  - Restart the processor in the '09x' wait state (The restart reason is ignored.)
  - After ACR processing is complete, enter CONFIG CPU(n),ONLINE at the master console.

## Additional Recovery Actions

There is another recovery procedure that is available as an operator response to excessive spin loops with message IEE331A or the equivalent '09x' wait state (except when the Restart Resource is the reason--see Figure 3-2 on page 3-24). Because this procedure is more complex than the normal responses to message IEE331A and the equivalent '09x' wait state code, it should be used only by operators who have received extensive training in the restart functions provided by the system.

The procedure consists of terminating the work and initiating recovery on the CPU that caused the spin. The operator would restart the CPU, rather than use ACR to remove it. In some cases this procedure will remove the cause of the excessive spin without the operator having to remove the CPU. In other cases the spin-loop message or wait state will recur, making it necessary to use the ACR option.

To use the Restart function to initiate recovery on the CPU causing the spin, the operator follows two different procedures, depending on whether the IEE331A message or '09x' wait state was issued.

## Restart Procedures

### *Procedure to Restart from Message IEE331A*

1. Reply 'U' to the message.
2. From the system console, initiate a Restart with reason code 0 for the CPU that the message identified as the cause of the excessive spin.

*Note:* For additional information on restart, see the previous topic in this chapter titled "Recovery by CPU Restart".

### *Procedure to Restart from Wait State 091, 092, 095, 097, or 09E:*

1. Obtain the logical CPU id (4n) from the sixth byte of the '09x' wait-state PSW. This is the CPU identified as the cause of the excessive spin loop.
2. From the system console, restart the CPU in the '09x' wait state (the restart reason code is ignored here).
3. Similarly, restart with reason code 0 the CPU (n) that was identified in the sixth byte of the '09x' wait state PSW.

*Note:* For additional information on restart, see the previous topic in this chapter titled "Recovery by CPU Restart".

## Determining the Cause of a Spin Loop

For some types of spin loops, the excessive spin notification facility initiates the building and recording of a LOGREC record on the processor that is causing the spin loop. This record is a standard system diagnostic work area (SDWA). In this SDWA the primary information on the cause of the spin loop is in the variable recording area (VRA) at offset X'194'. The spin loop record is identified by a completion code and a reason code. The system completion code is X'94071000' in fields SDWAABCC and SDWASABC. The reason code is X'10' in fields SDWACRC and SDWAOCRC. (Refer to the *Debugging Handbook* for the location of these fields in the SDWA.)

The VRA contains the following information on the spin loop:

- Identification text "EXCESS SPIN RESTART TO RECORD".
- The sixteen FRR addresses on the stack from the disabled (failing) CPU.
- An index value "INDEX = x", where x is a number between 0 and 16. This number indicates which of the 16 addresses represents the current FRR. If x=0 there are no current FRRs on the stack, unless IEAVESPR is the first FRR stack entry.

If the first stack entry is module IEAVESPR, then the current FRR is the index value + 1.

- The sixteen control registers on the failing CPU.
- The original completion code, reason code, and cross memory registers from the RTIW control block, if RTM was in control when Excessive Spin initiated the recording.
- The excessive spin length factor if RTM was not in control when recording was initiated.

## Analysis of Excessive Spin LOGREC Records

You can use the excessive-spin LOGREC records to identify the MVS routine that is running on the processor that is causing the spin condition. You should follow these steps:

1. Locate the sixteen FRR addresses from the stack that was current when the target processor was restarted. (These addresses appear after the identification text EXCESS SPIN RESTART TO RECORD at the beginning of the VRA.)
2. Identify the current FRR on the stack from the INDEX = x value that follows the sixteen addresses. (The x in this field is a number that can range from 0 to 16.) If x=0 there are no current FRRs on the stack. Otherwise, its value is an index that indicates which of the 16 addresses points to the current FRR. For example, if x=2, the second FRR address points to the current FRR.

3. Determine from a storage map the component that owns the FRR at this address. Traps can be set to determine which routine within the component is causing the excessive spin.

The LOGREC record provides data for analysis of the excessive spin loop without forcing a re-IPL of the system, as would occur with a standalone dump. However, a recovery usually limits the amount of data that can be collected to identify the cause of the spin loop. If an installation chooses to debug the spin loop problem and the related LOGREC record contains insufficient data, the operator should, on receipt of the next spin-loop message or wait state:

- Perform loop trace or instruction trace
- Take a stand-alone dump
- Re-IPL MVS

## Chapter 4. Reconfiguration

Reconfiguration is the process of adding hardware units to, or removing hardware units from, a configuration. When units are configured online (added to the configuration), they become available for system use; when configured offline (removed from the configuration), they become unavailable for system use. An installation can use reconfiguration to:

- adapt a system to changing work load environments by configuring operative units online or offline as required.
- perform concurrent maintenance (maintenance on a part of a complex while the other part continues normal operation).
- (possibly) allow a system to continue operation by configuring failing units offline.

A hardware unit (or units) may be removed from online status before the complex is initialized (that is, before the system is IPLed). An operator can deselect by means of the hardware system console such units as processor(s), storage element(s), or channel path(s). However, an operator should never deselect a unit during system operation by use of the hardware system console, because the operating system is NOT notified of the removal.

During system operation, some instances of reconfiguration are automatic; that is, the operating system configures failing units offline without any operator intervention. Other instances require operator intervention; that is, an operator can issue a CONFIG command so that a CONFIGxx member of SYS1.PARMLIB causes reconfiguration or can issue explicit CONFIG and/or VARY commands to configure units online or offline.

## Logical and Physical Reconfiguration

Logical reconfiguration is the process that allows or prevents the use of a resource by the operating system. Physical reconfiguration is the process that allows or prevents the use of a resource by the hardware. By issuing a CONFIG command from the master console, an operator can cause the logical and physical reconfiguration of any of the following system elements (if applicable to the particular type of processor):

- CPUs
- Storage increments
- Storage elements
- Channel paths
- Vector Facilities
- Extended storage elements

*Note:* Physical reconfiguration may not be supported for all hardware units by all processor models. Refer to the applicable *Functional Characteristics* manual for detailed information. For example on a 4381 Model 3, CONFIG CPU does only a logical reconfiguration, but CONFIG CHANNEL PATH does both a logical and physical reconfiguration.

In addition, an operator can issue a VARY command from the master console to cause the logical reconfiguration of I/O devices or I/O paths. (An I/O path is the logical route between a processor and a device.) For detailed information concerning the syntax and use of the CONFIG and VARY commands, refer to *System Commands*.

## General Considerations for Reconfiguration

This section contains the following topics:

- Degree of reconfiguration support according to processor types
- Recommended sequence for partitioning and merging
- DISPLAY command considerations
- Program properties table considerations

### Reconfiguration Support According to Processor Types

In any installation, the operating system supports reconfiguration as noted in the following paragraphs. However, a particular processor complex might not support all the specified reconfiguration options. Refer to the applicable *Functional Characteristics* publication for the processor-dependent information.

- Uniprocessor or UP (for example, a 3090 Model 180E) - Depending on the processor type, an installation can configure offline a storage element, storage increments, channel paths, and devices. In a UP system, the main purpose of reconfiguration is to configure offline failing units to allow the system to continue operation.



- Nonpartitionable multiprocessor (for example, a 3081, 4381 MG14, or a 3090 Model 200E or 300E) - Depending on the processor type, an installation can configure offline a CPU, a Vector Facility, a storage element, storage increments, channel paths, and devices. Again, the main purpose of reconfiguration is to configure failing units offline to allow the system to continue operation.
- Partitionable multiprocessor (for example, a 3084 or a 3090 Model 400E or 600E) - an installation has the maximum capability for reconfiguration: single-image mode to physically partitioned mode; or physically partitioned mode to single-image mode. (Examples of both processes are presented later in this chapter.) In addition, an installation can configure offline multiple hardware units (channel paths, Vector Facilities, CPUs, extended storage elements, real storage, and devices) to allow the system to continue operation.

### Recommended Sequence for Partitioning and Merging

The order in which CONFIG commands are issued can affect the function and performance of the system. The recommended order in which elements should be taken offline during partitioning are: channel paths, CPUs, extended storage (if applicable), and real storage. This is the sequence of processing of the CONFIGxx member of SYS1.PARMLIB.

**Channel paths** should go offline first to reduce the load on the CPUs and to allow the capturing of channel activity data.

**CPUs** should go offline next to reduce the workload before real storage goes offline.

*Note:* Because Vector Facilities go offline with their associated CPUs, it is not necessary to issue CONFIG VF commands as part of the partitioning process.

**Extended storage elements** (if applicable) should go offline next while real storage is still available to be used for the migration of data from extended storage to auxiliary storage.

**Real storage elements** should be the last to go offline.

During merging the reverse order should be used to bring elements back online.

## DISPLAY Command Considerations

An operator can use two different forms of the DISPLAY command to display information concerning the status of the hardware configuration. The two forms are D U and D M. The information in the display could prove useful when attempting reconfiguration. (For detailed information concerning the syntax and use of the DISPLAY command, refer to *System Commands*.)

### D U Command

By issuing a D U command, an operator can display (in message IEE450I) the online/offline status of a device or set of devices. By specifying ALLOC on a D U command, an operator can display (in message IEE106I) the jobname(s) to which a device is allocated. If an operator is attempting to vary that device offline, the vary cannot complete until the device is unallocated; that is, the jobs must complete or they must be cancelled.

### D M Command

By issuing a D M command, an operator can display (in message IEE174I) the status of specified hardware units. The display of the status of real storage can assist an operator during reconfiguration. The display includes: storage offline, storage 'pending' offline, and reconfigurable storage sections. For storage 'pending' offline, the display includes the ASID and jobname of the current user of the storage. Storage in use cannot be configured offline until it is free; that is, the using job must complete or it must be cancelled.

An operator can also use the D M=CONFIG(xx) command to display (in message IEE097I) the deviation between the current hardware configuration and the one in a specified CONFIGxx member of SYS1.PARMLIB. The deviation display can be used to determine which units to configure online/offline to satisfy shift requirements or a changing work load.

The D M=CPU command, as part of its CPU display, also gives Vector Facility status.

The D M=ESTOR command shows the amount of extended storage that is in each of the following states: offline, reconfigurable, pending, offline, and belonging to another configuration. The D M=ESTOR(E), in comparison, gives the status of extended storage for each element.

The D M=SIDE command displays the total resources on each side of a partitionable processor complex.

## Program Properties Table Considerations

The operating system normally attempts to assign requests for long-term fixed pages to preferred storage frames when the requesting job was initiated non-swappable. However, an authorized job can be initiated as swappable and during execution issue a SYSEVENT to make itself non-swappable for a short period of time. The job may request long-term fixed pages that are assigned to non-preferred storage. Usually this does not present a problem because the job shortly makes itself swappable again. The storage that backs the long-term fixed pages can be freed by swapping out the job when the storage is required for storage reconfiguration.

However, an installation may encounter a long-running job that makes itself non-swappable for long periods of time and also makes requests for short-term fixed pages that cannot be freed until the job ends normally. Some of those requests may be satisfied from non-preferred storage. Since the frames cannot be freed by paging them out or by swapping out the job, storage reconfiguration may not be possible.

An installation can resolve the foregoing problem by including such jobs in the PPT and setting the appropriate flag bits. (Refer to *SPL: Initialization and Tuning* for detailed information.)

## Real Storage Reconfiguration

The real storage that is shared by all processors in a configuration is logically divided into storage increments (See example in Figure 4-1). Each real storage increment is composed of two subincrements - one subincrement contains the even-numbered frames of the increment (for example, 0K, 8K, 16K, etc.); the other contains the odd-numbered frames (for example, 4K, 12K, 20K, etc.).

64M		64M
60M	HSA and Preferred	60M
56M	Reconfigurable	56M
52M	Reconfigurable	52M
48M	Reconfigurable	48M
44M	Reconfigurable	44M
40M	Reconfigurable	40M
36M	Reconfigurable	36M
32M	Reconfigurable	32M
28M	Reconfigurable	28M
24M	Preferred	24M
20M	Preferred	20M
16M	Preferred	16M
12M	SQA and Preferred	12M
8M	Preferred	8M
4M	Preferred	4M
0M	V = R and Preferred	0M

Figure 4-1. A Logical View of Real Storage (3084 Example)

Real storage is physically divided into storage elements (SEs). The subincrements of an increment are likely to reside in different real storage elements (See Figure 4-2).

SE0	SE2	SE1	SE3
28-32M Even Frames Reconfigurable	28-32M Odd Frames Reconfigurable	60-64M Even Frames HSA and Preferred	60-64M Odd Frames HSA and Preferred
24-28M Even Frames Preferred	24-28M Odd Frames Preferred	56-60M Even Frames Reconfigurable	56-60M Odd Frames Reconfigurable
20-24M Even Frames Preferred	20-24M Odd Frames Preferred	52-56M Even Frames Reconfigurable	52-56M Odd Frames Reconfigurable
16-20M Even Frames Preferred	16-20M Odd Frames Preferred	48-52M Even Frames Reconfigurable	48-52M Odd Frames Reconfigurable
12-16M Even Frames SQA and Preferred	12-16M Odd Frames Preferred	44-48M Even Frames Reconfigurable	44-48M Odd Frames Reconfigurable
8-12M Even Frames Preferred	8-12M Odd Frames Preferred	40-44M Even Frames Reconfigurable	40-44M Odd Frames Reconfigurable
4-8M Even Frames Preferred	4-8M Odd Frames Preferred	36-40M Even Frames Reconfigurable	36-40M Odd Frames Reconfigurable
0-4M Even Frames V = R and Preferred	0-4M Odd Frames V = R and Preferred	32-36M Even Frames Reconfigurable	32-36M Odd Frames Reconfigurable

Figure 4-2. A Physical View of Real Storage (3084)

Refer to Figure 4-3 to see the differences between 308x and 3090 real storage sizes and IDs. (For other differences between the 3084 and the partitionable 3090 models, see Figure 4-15 on page 4-32.)

Processor Type	Storage Element ID	Storage Element Size	Increment Size	Subincrement Size	Max. Storage Size
3081 (see note)	0, 2	8M, 16M, 32M	4M, 8M	2M, 4M	64M
3084 (see note)	0, 2, 1, 3	16M, 32M	4M, 8M	2M, 4M	128M
3090					
Mod 200	0, 1	32M	2M	1M	64M
Mod 200E	0, 1	32M, 64M	4M	2M	128M
Mod 300E	0, 1	32M, 64M	4M	2M	128M
Mod 400	0, 1, 2, 3	32M, 64M	2M	1M	256M
Mod 400E	0, 1, 2, 3	32M, 64M	4M	2M	256M
Mod 600E	0, 1, 2, 3	32M, 64M	4M	2M	256M

**Figure 4-3. Real Storage Differences Between 308x and 3090 Systems**

*Note:* If 308x installed storage is 96M or larger, the increment size is 8M. For 308x storage of 64M or less, the increment size is 4M.

To reconfigure real storage ranges or amounts (if the processor type supports this function), an operator at the master console would issue a CONFIG STOR ONLINE/OFFLINE command.

The following real storage increments cannot be configured offline:

- the increment containing absolute address 0
- the highest addressable increment available at IPL-time
- any increment containing preferred real storage frames

A storage element can be configured offline only if:

- It contains only non-preferred storage frames
- The preferred storage subincrements in this storage element can be moved to another storage element containing reconfigurable storage subincrements. (The operating system requests the service processor move the data and addresses.)

When reconfiguring from single-image mode to physically partitioned mode, an operator must be able to configure offline the real storage elements owned by the side going offline.

When configuring a real storage element offline, the operator may see message IEE575A indicating that real storage configuration is waiting to complete. The message may be cancelled in a short period of time (typically less than a minute) and may be displayed several times as the operating system configures the real storage element offline. If the message remains outstanding for a long period of time, it indicates that the operating system cannot find sufficient reconfigurable real storage to satisfy the configuration request. The operator should issue the `D M=STOR` command to identify the job using the real storage that cannot be freed. The operator can then take one of two actions:

1. Cancel the jobs that are using the storage to allow the storage configuration to complete
2. Reply 'C' to message IEE575A to terminate the storage configuration process

If the operator takes action 2, any real storage already configured offline remains offline.

The operator should document the names of the jobs using the real storage and give them to the system programming staff for possible inclusion in the program properties table.

## Extended Storage Reconfiguration on Partitionable 3090 Models

The 3090 partitionable models allow the operator to reconfigure extended storage elements by means of the `CONFIG ESTOR(E=id)` command, either from the master console or as part of a `CONFIGxx` member of `SYS1.PARMLIB`. During partitioning this command should be issued **before** real storage is removed, because the migration of data from extended storage to auxiliary storage uses real storage.

A separate `CONFIG` command must be issued for each extended storage element to go offline or to go online. There can be as many as four elements, two per side, numbered 0--3. These numbers are specified, one per command, in the `E=id` keyword of the `CONFIG` command.

To determine the status of installed extended storage before issuing the `CONFIG` command, the operator can use the `D M=ESTOR(E)` command.

# Processor Reconfiguration

When an operator configures a processor offline:

- The operating system stops dispatching work to that processor.
- The processor enters the stopped state.
- The processor is then taken offline first logically, then physically.

*Note:* The operating system rejects a CONFIG CPU(x),OFFLINE command when:

- The target processor is the only online processor.
- The target processor is the only processor with an operative timer.
- An ACR condition occurs during offline processing.
- Any active jobs have CPU affinity with the target processor. Message IEE718I is issued listing the currently scheduled jobs with CPU affinity. The operator can prevent the operating system from scheduling any additional jobs, by replying YES to message IEE718D. The operator can either wait for the active jobs to complete or cancel them, and then reissue the CONFIG CPU(x),OFFLINE command.

## Reconfiguring a Processor with a Vector Facility

If the processor is a 3090 and it has an associated Vector Facility, this processing occurs:

- When CONFIG CPU(x),OFF is issued, the Vector Facility associated with CPU x is taken logically and physically offline.
- When the processor is brought back online through use of a CONFIG CPU(x) or CONFIG CPU(x),ONLINE command (without VFON or VFOFF being specified), an associated Vector Facility will be in the physical and logical state it had before the processor went offline. That is, if the Vector Facility was online before its processor went offline, it will still be online.

For examples of the use of the CONFIG VF(x) command, and the CONFIG CPU(x),ONLINE command with VFON and VFOFF specified, see "Vector Facility Reconfiguration Examples" on page 4-12.

*Note:* To take a 3090 Vector Facility offline for repair or physical maintenance, it is necessary to take offline the side (partitionable model) or to shut down the entire system (nonpartitionable model) to set up a maintenance configuration. The "x" designation in the CONFIG VF(x) command should be the same as that of the associated CPU.

After the repair, issue the CONFIG CPU(x),ONLINE,VFON command, which would bring back online first the CPU, then its associated Vector Facility. Another way to bring the CPU and associated Vector Facility online would be to issue two commands: CONFIG CPU(x),ONLINE followed by CONFIG VF(x),ONLINE.



## Removing the Last Vector Facility

If a CONFIG command specifies the removal of the last Vector Facility in the system, and vector jobs are scheduled, the following message will appear on the master console:

```
IEE176I CONFIG {CPU(x) | VF(x)},OFFLINE COMMAND WOULD REMOVE  
LAST VF, dd VF JOBS SCHEDULED. JOBNAMEs ARE: jobname, [jobname...]
```

```
IEE177D REPLY 'U' TO SUSPEND VF JOBS. REPLY 'C' TO CANCEL CONFIG  
COMMAND
```

When you reply 'U' to the IEE177D message, the vector jobs will be put into a vector wait. If they were submitted with TIME = 1440 specified in their JCL, they will not time out and will not be cancelled. Message IEE700I will appear for each vector job in a vector wait:

```
IRA700I jobname WAITING FOR AVAILABILITY OF VF
```

Later, when a Vector Facility is brought online, the vector jobs will continue on that Vector Facility.

## Channel Measurements

When reconfiguring from single-image mode to physically partitioned mode or from physically partitioned mode to single-image mode, an installation should note the following points concerning channel measurements.

1. When reconfiguring from single-image mode to physically partitioned mode, an installation should configure channel paths offline before the processors to prevent SRM from suspending channel measurements.
2. When reconfiguring from physically partitioned mode to single-image mode, the TOD clocks in both partitions must be synchronized before SRM starts channel measurements. Therefore, an installation should configure online the processors before the channel paths.

If the processors are configured online before the channel paths, SRM suspends channel measurements for approximately 16 seconds. If, however, the channel paths are configured online before the processors, SRM suspends channel measurements from the time the first channel path is configured online until after the processors are configured online and the TOD clocks are synchronized.

*Note:* The order used by the CONFIGxx parmlib member preserves channel measurement.

## Vector Facility Reconfiguration Examples

This section illustrates various CONFIG commands that can take a Vector Facility offline or bring it back online.

### Example 1:

This example shows that CPU x is to be brought online physically and logically, and if CPU x has a Vector Facility, the logical and physical status of its Vector Facility (online or offline) is to be the same as it was when the CPU was last online.

```
Issue CONFIG CPU(x)
      or
CONFIG CPU(x),ONLINE
```

When this command is completed, these messages are displayed on the master console if the Vector Facility comes online:

```
IEE504I CPU(x) ONLINE
IEE504I VF(x) ONLINE
```

If, however, the CPU has no associated Vector Facility, or if the Vector Facility does not come online, the message is:

```
IEE504I CPU(x) ONLINE
```

### Example 2:

This example shows that CPU x is to be brought online logically and physically, along with its Vector Facility. This action might be part of a merging procedure to bring a side back online after partitioning the system. (If CPU x does not have a Vector Facility, the CPU is brought online anyway.)

```
Issue CONFIG CPU(x),ONLINE,VFON
```

These messages are then displayed on the master console:

```
IEE504I CPU(x) ONLINE
IEE504I VF(x) ONLINE
```

If the CPU has no associated Vector Facility, the message is:

```
IEE504I CPU(x) ONLINE
IEE506I VF(x) NOT RECONFIGURED - CPU HAS NO VF
```

### Example 3:

This example illustrates bringing CPU x logically and physically online, but keeping its Vector Facility logically and physically offline.

```
Issue CONFIG CPU(x),ONLINE,VFOFF
```

Upon successful completion of processing, these messages appear:

```
IEE504I CPU(x) ONLINE  
IEE505I VF(x) OFFLINE
```

If, however, the CPU x does not have a Vector Facility, these messages appear:

```
IEE504I CPU(x) ONLINE  
IEE506I VF(x) NOT RECONFIGURED - CPU HAS NO VF
```

**Example 4:**

This example shows how to take CPU x offline physically and logically, and take its Vector Facility logically offline, so that no software can access the Vector Facility.

```
Issue CONFIG CPU(x),OFFLINE
```

Upon successful completion of processing, this message is issued:

```
IEE505I CPU(x) OFFLINE
```

**Example 5:**

This example shows how to bring the Vector Facility for CPU x logically and physically online, if CPU x is already logically and physically online. One possible use might be to try to bring a Vector Facility back online, in an attempt to recover, after machine checks had caused the MCH to take the Vector Facility offline automatically.

```
Issue CONFIG VF(x) or CONFIG VF(x),ONLINE
```

When the Vector Facility is brought online, this message is issued:

```
IEE504I VF(x) ONLINE
```

If, however, CPU x is offline, this message is issued instead:

```
IEE506I VF(x) NOT RECONFIGURED - CPU NOT ONLINE
```

If CPU x is online but does not have a Vector Facility, this message appears:

```
IEE506I VF(x) NOT RECONFIGURED - CPU HAS NO VF
```

**Example 6:**

This example shows how to take the Vector Facility for CPU x logically and physically offline, if CPU x is logically and physically online. One possible use might be take a Vector Facility offline after a re-IPL, if the Vector Facility had been repeatedly causing errors before the re-IPL.

```
Issue CONFIG VF(x),OFFLINE
```

When the Vector Facility goes offline, this message is issued:

IEE505I VF(x) OFFLINE

If, however, CPU x is offline when the operator issued the CONFIG, this message appears instead:

IEE506I VF(x) NOT RECONFIGURED - CPU NOT ONLINE

If CPU x is online but does not have a Vector Facility, the following is displayed:

IEE506I VF(x) NOT RECONFIGURED - CPU HAS NO VF

## Channel Path Reconfiguration

To reconfigure channel paths, an operator issues a CONFIG CHP ONLINE/OFFLINE command at the master console. An operator can reconfigure channel paths on an individual basis. However, when configuring from single-image mode to physically partitioned mode, or physically partitioned mode to single-image mode, an operator can reconfigure all the channel paths owned by a side with a single command: CONFIG CHP(ALL,x). (x is the identifier of the side, either 0 or 1 for the 3090, A or B for the 3084.

Offline processing determines which devices are connected to a channel path and if that path is the last path to a device. To configure offline the last path to a device, an operator can use:

- the UNCOND operand to configure offline the last path to an unallocated, online device
- the FORCE operand to configure offline the last path to a device regardless of the state of the device. (Refer to *MVS/XA System Commands* for cautions on the use of the FORCE operand of CONFIG.)

To ensure that the specification of FORCE is intentional, the operating system issues message IEE800D requesting that the operator reply YES or NO to continue or negate the execution of the CONFIG command with the FORCE operand.

Online processing determines which devices are connected to a particular channel path and updates their applicable control blocks so they can use the newly configured-online channel path.

## I/O Device Reconfiguration

To reconfigure I/O devices, an operator issues a VARY ONLINE/OFFLINE command at the master console. If an operator issues a VARY OFFLINE command for a device that is currently in use, the operating system marks the device 'pending offline'. The operating system makes no further allocations to the device unless the volume mounted on the device is specifically requested.

Since vary offline processing cannot complete until a device is unallocated, an operator can either wait until the jobs using the device complete or cancel them.

*Note:* If a partitionable complex is being reconfigured from single image to partitioned mode, and a tape mount is pending, the tape drive(s) might not start after they are mounted and the system has been partitioned. The problem can be circumvented by the operator issuing a VARY device online command for the tape drive(s).

*Note:* When partitioning, before issuing the CONFIG CHP(ALL,n),OFFLINE,UNCOND command, complete or cancel any mounts that may be affected by this command.

## Examples of Partitioning and Merging a 3084

Two examples of reconfiguration are presented in this section: configuring from single-image mode to physically partitioned mode and configuring from physically partitioned mode to single-image mode.

These examples show:

- The required commands used to partition and to merge
- The messages that are issued during processing
- How the contents of real storage are handled during the processing

In each of the examples, you should assume the following conditions:

- 48 channel paths
- installed storage of 64M consisting of four 16M storage elements with 4M storage increments
- RSU = 8 specified at IPL
- V = R area contained in storage increment 0-4M
- SQA contained in 12-16M even frames
- storage ranges 0-28M, and 60-64M not reconfigurable
- storage range 28-60M reconfigurable
- HSA contained in storage increment 60-64M

Notes:

### 1. Side Terminology

*The hardware and the operating system use different terms for the two sides of a 3084. Some hardware messages and displays refer to 'Side A' and 'Side B'; while the operating system refers to 'Side 0' and 'Side 1'. Side 0 and Side A are synonymous, as are Side 1 and Side B.*

*Also, the hardware and the operating system use different names for the reconfiguration commands. For example, the hardware messages that reflect the issuance of an MVS CONFIG command indicate that the service processor received a VARY command for a physical unit.*

### 2. Configuration Switch

*When configuring a 3084 from either single-image mode to physically partitioned mode or physically partitioned mode to single-image mode, an operator may be told to change the CONFIGURATION switch by messages issued on the system console.*

*An operator should never change the CONFIGURATION switch during normal system operation unless instructed to do so by a message on the system console. Otherwise, changing the switch may cause a system outage.*

## Partitioning from Single-Image Mode to Physically Partitioned Mode (Side B to Be Configured Offline)

Prior to configuring from single-image mode to physically partitioned mode, the 3084 processor complex appears as shown in Figure 4-4.

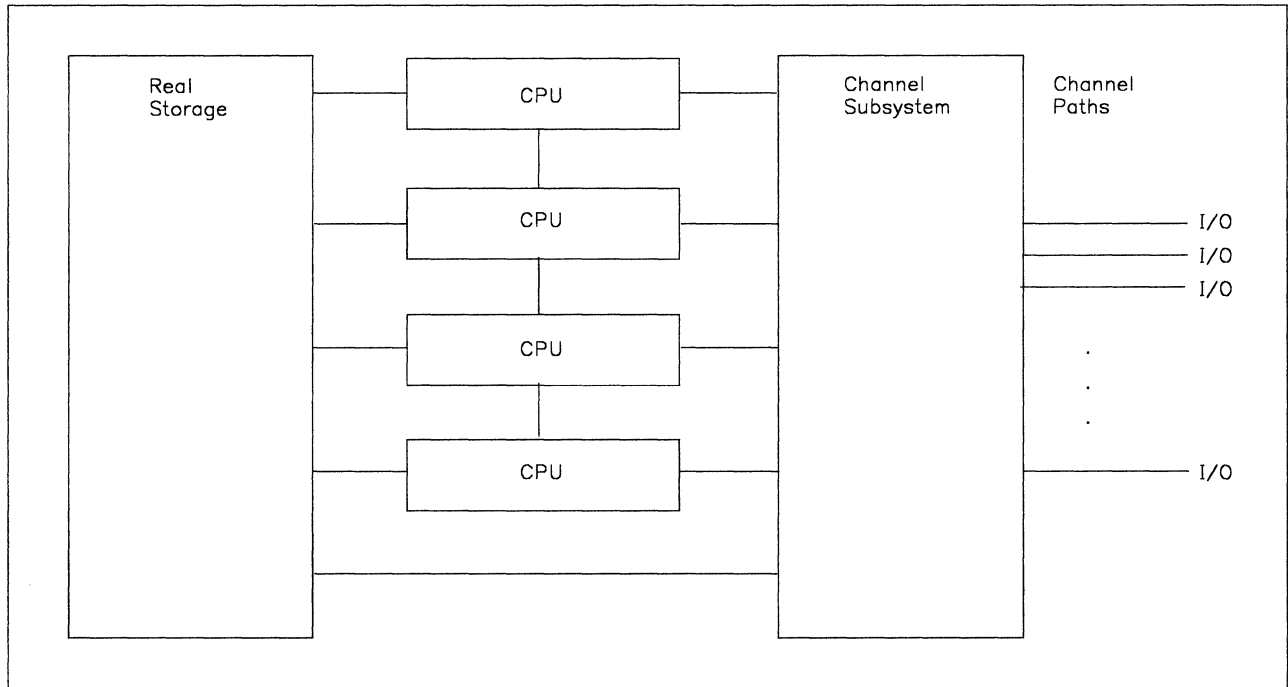


Figure 4-4. Single-Image Mode of a 3084

Assume the 3084 storage layout of four storage elements (SE0 through SE3), as shown in Figure 4-5.

SE0	SE2	SE1	SE3
28-32M Even Frames Reconfigurable	28-32M Odd Frames Reconfigurable	60-64M Even Frames HSA and Preferred	60-64M Odd Frames HSA and Preferred
24-28M Even Frames Preferred	24-28M Odd Frames Preferred	56-60M Even Frames Reconfigurable	56-60M Odd Frames Reconfigurable
20-24M Even Frames Preferred	20-24M Odd Frames Preferred	52-56M Even Frames Reconfigurable	52-56M Odd Frames Reconfigurable
16-20M Even Frames Preferred	16-20M Odd Frames Preferred	48-52M Even Frames Reconfigurable	48-52M Odd Frames Reconfigurable
12-16M Even Frames SQA and Preferred	12-16M Odd Frames Preferred	44-48M Even Frames Reconfigurable	44-48M Odd Frames Reconfigurable
8-12M Even Frames Preferred	8-12M Odd Frames Preferred	40-44M Even Frames Reconfigurable	40-44M Odd Frames Reconfigurable
4-8M Even Frames Preferred	4-8M Odd Frames Preferred	36-40M Even Frames Reconfigurable	36-40M Odd Frames Reconfigurable
0-4M Even Frames V = R and Preferred	0-4M Odd Frames V = R and Preferred	32-36M Even Frames Reconfigurable	32-36M Odd Frames Reconfigurable

**Figure 4-5. Storage Layout in Single-Image Mode (3084)**

An operator could use the following sequence of commands at the MVS console to physically partition a system. The reconfiguration is presented in the following order:

- Configure channel paths offline
- Configure CPUs offline
- Configure real storage offline



1. **Enter: CONFIG CHP(ALL,1),OFFLINE,UNCOND**

*Note:* Before issuing the CONFIG CHP(ALL,n),OFFLINE,UNCOND command, complete or cancel any mounts that may be affected by this command.

The following messages are displayed on the master console:

```
IEE503I    CHP(ALL,1),OFFLINE
IEE712I    CONFIG PROCESSING COMPLETE
```

As each channel path is configured offline, this message is displayed on the system console:

```
VARY CHAN PATH nn OFF RECEIVED BY MSSF. RESULT=0020.
```

Once the operating system determines that all channel paths associated with the EXDC are offline, it configures the EXDC offline and displays the following message on the system console:

```
VARY I/O SIDE B OFF RECEIVED BY MSSF. RESULT=0020.
```

2. **Enter: CONFIG CPU(1),OFFLINE**

The following messages are displayed on the master console:

```
IEE505I    CPU(1),OFFLINE
IEE712I    CONFIG PROCESSING COMPLETE
```

After CPU1 is configured offline, the following message is displayed on the system console:

```
VARY CPU 01 OFF RECEIVED BY MSSF. RESULT=0020.
```

3. **Enter: CONFIG CPU(3),OFFLINE**

The following messages are displayed on the master console:

```
IEE505I    CPU(3),OFFLINE
IEE712I    CONFIG PROCESSING COMPLETE
```

After CPU3 is configured offline, the following message is displayed on the system console:

```
VARY CPU 03 OFF RECEIVED BY MSSF. RESULT=0020.
```

4. Enter: CONFIG STOR(E=1),OFFLINE

The following messages are displayed on the master console:

```
IEE510I  1 OF EVERY 2 FRAMES OF REAL STORAGE LOCATIONS 32M TO 36M
OFFLINE

IEE510I  1 OF EVERY 2 FRAMES OF REAL STORAGE LOCATIONS 36M TO 40M
OFFLINE

IEE510I  1 OF EVERY 2 FRAMES OF REAL STORAGE LOCATIONS 40M TO 44M
OFFLINE

IEE510I  1 OF EVERY 2 FRAMES OF REAL STORAGE LOCATIONS 44M TO 48M
OFFLINE

IEE510I  1 OF EVERY 2 FRAMES OF REAL STORAGE LOCATIONS 48M TO 52M
OFFLINE

IEE510I  1 OF EVERY 2 FRAMES OF REAL STORAGE LOCATIONS 52M TO 56M
OFFLINE

IEE510I  REAL STORAGE LOCATIONS 56M TO 60M OFFLINE (See the following note.)

IEE712I  CONFIG PROCESSING COMPLETE
```

*Note:* Because SE1 contained some HSA and preferred storage (the even frames in the 60M--64M range in Figure 4-5), the reconfiguration process consists of swapping this group of frames with the odd frames in the 56M--60M range in SE3. This is why all the storage from 56M--60M goes offline rather than just half this range.

After configuring SE1 offline, storage appears as shown in Figure 4-6.

SE0	SE2	SE1	SE3
28-32M Even Frames Reconfigurable	28-32M Odd Frames Reconfigurable	56-60M Odd Frames	60-64M Odd Frames HSA and Preferred
24-28M Even Frames Preferred	24-28M Odd Frames Preferred	56-60M Even Frames	60-64M Even Frames HSA and Preferred
20-24M Even Frames Preferred	20-24M Odd Frames Preferred	52-56M Even Frames	52-56M Odd Frames Reconfigurable
16-20M Even Frames Preferred	16-20M Odd Frames Preferred	48-52M Even Frames	48-52M Odd Frames Reconfigurable
12-16M Even Frames SQA and Preferred	12-16M Odd Frames Preferred	44-48M Even Frames	44-48M Odd Frames Reconfigurable
8-12M Even Frames Preferred	8-12M Odd Frames Preferred	40-44M Even Frames	40-44M Odd Frames Reconfigurable
4-8M Even Frames Preferred	4-8M Odd Frames Preferred	36-40M Even Frames	36-40M Odd Frames Reconfigurable
0-4M Even Frames V = R and Preferred	0-4M Odd Frames V = R and Preferred	32-36M Even Frames	32-36M Odd Frames Reconfigurable

OFFLINE

Figure 4-6. Storage Layout - SE1 Configured Offline (3084)

After SE1 is configured offline, the following message is displayed on the system console:

VARY STOR ELEM 01 OFF RECEIVED BY MSSF. RESULT = 0020.

5. Enter: **CONFIG STOR(E = 3),OFFLINE**

The following messages are displayed on the master console:

```

IEE510I  REAL STORAGE LOCATIONS 28M TO 32M OFFLINE
IEE510I  1 OF EVERY 2 FRAMES OF REAL STORAGE LOCATIONS 32M TO 36M
          OFFLINE
IEE510I  1 OF EVERY 2 FRAMES OF REAL STORAGE LOCATIONS 36M TO 40M
          OFFLINE
IEE510I  1 OF EVERY 2 FRAMES OF REAL STORAGE LOCATIONS 40M TO 44M
          OFFLINE
IEE510I  1 OF EVERY 2 FRAMES OF REAL STORAGE LOCATIONS 44M TO 48M
          OFFLINE
IEE510I  1 OF EVERY 2 FRAMES OF REAL STORAGE LOCATIONS 48M TO 52M
          OFFLINE
IEE510I  1 OF EVERY 2 FRAMES OF REAL STORAGE LOCATIONS 52M TO 56M
          OFFLINE
IEE712I  CONFIG PROCESSING COMPLETE

```

After SE3 is configured offline, the following message is displayed on the system console:

VARY STOR ELEM 03 OFF RECEIVED BY MSSF. RESULT = 0020.

After configuring SE3 offline, storage appears as shown in Figure 4-7.

SE0	SE2	SE1	SE3
60-64M Even Frames HSA and Preferred	60-64M Odd Frames HSA and Preferred	56-60M Odd Frames	28-32M Odd Frames
24-28M Even Frames Preferred	24-28M Odd Frames Preferred	56-60M Even Frames	28-32M Even Frames
20-24M Even Frames Preferred	20-24M Odd Frames Preferred	52-56M Even Frames	52-56M Odd Frames
16-20M Even Frames Preferred	16-20M Odd Frames Preferred	48-52M Even Frames	48-52M Odd Frames
12-16M Even Frames SQA and Preferred	12-16M Odd Frames Preferred	44-48M Even Frames	44-48M Odd Frames
8-12M Even Frames Preferred	8-12M Odd Frames Preferred	40-44M Even Frames	40-44M Odd Frames
4-8M Even Frames Preferred	4-8M Odd Frames Preferred	36-40M Even Frames	36-40M Odd Frames
0-4M Even Frames V = R and Preferred	0-4M Odd Frames V = R and Preferred	32-36M Even Frames	32-36M Odd Frames
		OFFLINE	OFFLINE

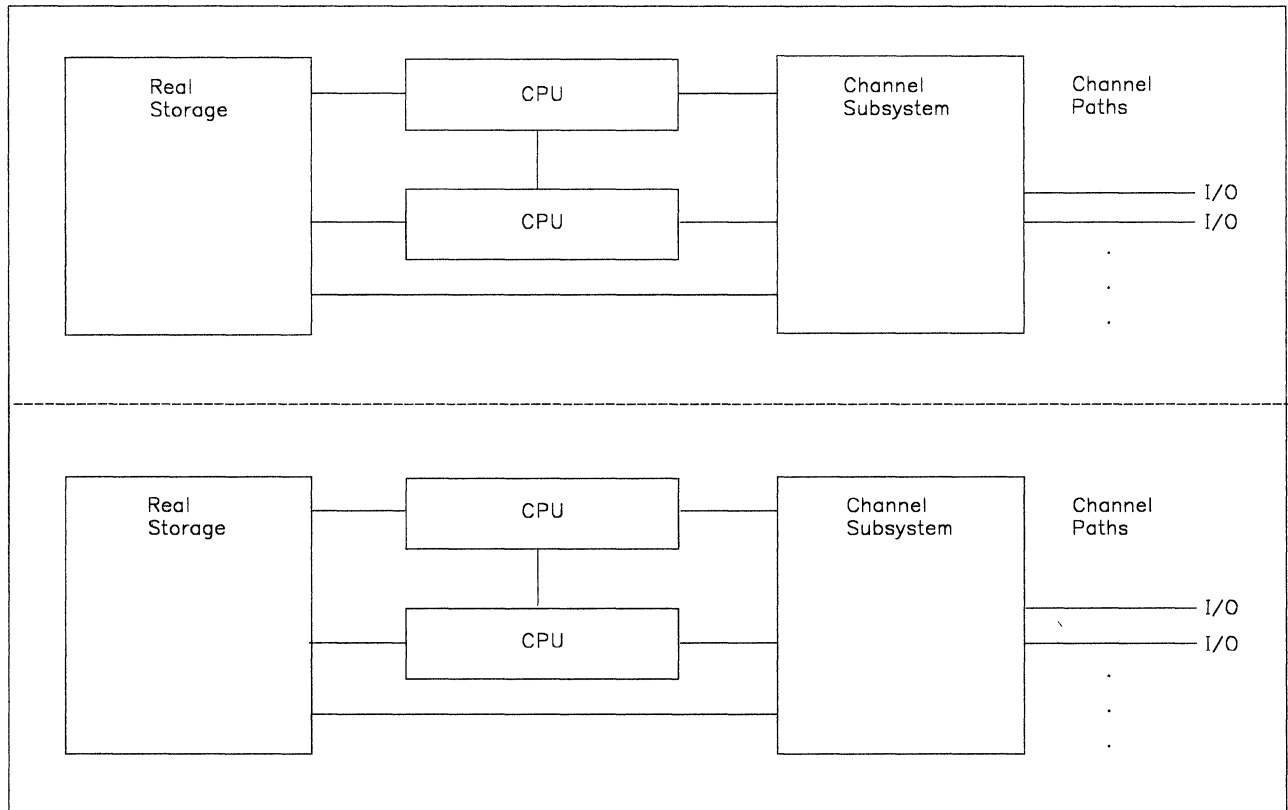
Figure 4-7. Storage Layout - SE1 and SE3 Configured Offline (3084)

Once the operating system determines that no elements (CPUs, real storage, CHPs) remain configured to Side B, it configures Side B offline and displays the following message on the system console:

VARY SIDE B OFF RECEIVED BY MSSF. RESULT = 0020.

6. **Enter: VARYPHY SIDEB,OFF** at the system console. If message SET CONFIGURATION SWITCH TO PP is displayed on the system console, change the CONFIGURATION switch to PP. When VARYPHY processing completes, message REQUEST COMPLETED is displayed on the system console.

After configuring from single-image mode to physically partitioned mode, the processor complex appears as shown in Figure 4-8.



**Figure 4-8. Physically Partitioned Mode of the 3084**

At the master console on Side A, the operator can verify the physical partitioning of the system using the series of D M commands shown in Figure 4-9.

```

D M=SIDE

IEE174I hh.mm.ss MATRIX DISPLAY
SIDE STATUS
SIDE:                0                      1
STATUS:              ONLINE                UNAVAILABLE
I/O ENGINE:         0
CPU:                 0 2
CHP:                 0-7 10-17 20-27
STOR(E=x):          0 2

TOTAL STOR: 64M
*=OFFLINE
-----

D M=CPU

IEE174I hh.mm.ss MATRIX DISPLAY
PROCESSOR STATUS
CPU  STATUS      SERIAL
 0  ONLINE      0123453084
 1  OFFLINE
 2  ONLINE      2123453084
 3  OFFLINE
-----

D M=CHP

IEE174I hh.mm.ss MATRIX DISPLAY
CHANNEL PATH STATUS
 0 1 2 3 4 5 6 7 8 9 A B C D E F
0 + + + + + + + . . . . .
1 + + + + + + + . . . . .
2 + + + + + + + . . . . .
***** SYMBOL EXPLANATIONS *****
. DOES NOT EXIST          * LOGICALLY OFF, PHYSICALLY ONLINE
- LOGICALLY & PHYSICALLY OFFLINE + LOGICALLY & PHYSICALLY ONLINE
-----

D M=STOR(E)

IEE174I hh.mm.ss MATRIX DISPLAY
STORAGE ELEMENT STATUS
0: OWNED STORAGE=16M  STATUS=ONLINE
STOR(E=1) IS PART OF ANOTHER CONFIGURATION - NO STATUS OBTAINED
2: OWNED STORAGE=16M  STATUS=ONLINE
STOR(E=3) IS PART OF ANOTHER CONFIGURATION - NO STATUS OBTAINED

```

Figure 4-9. Examples of D M Displays - 3084 System in Physically Partitioned Mode

At this point, the installation has two separate systems, each monitored by its own service processor and each with its own system and service consoles. One system, which consists of Side A, continues productive work. Side B is now just a collection of hardware units. The operator must perform the following steps on the side-B system console before Side B can do any productive work:

- IML
- Define the configuration by use of the Configuration frame
- Power-on-reset (selecting an IOCDS for the I/O configuration and either 370 mode or 370-XA mode)
- IPL

### Merging from Physically Partitioned Mode to Single-Image Mode (Side B To Be Configured Online)

The process of configuring from physically partitioned mode to single-image mode is essentially the reverse of configuring from single-image mode to physically partitioned mode.

Prior to configuring SE1 and SE3 online, assume the storage layout is as shown in Figure 4-10.

SE0	SE2	SE1	SE3
60-64M Even Frames HSA and Preferred	60-64M Odd Frames HSA and Preferred	56-60M Odd Frames	28-32M Odd Frames
24-28M Even Frames Preferred	24-28M Odd Frames Preferred	56-60M Even Frames	28-32M Even Frames
20-24M Even Frames Preferred	20-24M Odd Frames Preferred	52-56M Even Frames	52-56M Odd Frames
16-20M Even Frames Preferred	16-20M Odd Frames Preferred	48-52M Even Frames	48-52M Odd Frames
12-16M Even Frames SQA and Preferred	12-16M Odd Frames Preferred	44-48M Even Frames	44-48M Odd Frames
8-12M Even Frames Preferred	8-12M Odd Frames Preferred	40-44M Even Frames	40-44M Odd Frames
4-8M Even Frames Preferred	4-8M Odd Frames Preferred	36-40M Even Frames	36-40M Odd Frames
0-4M Even Frames V = R and Preferred	0-4M Odd Frames V = R and Preferred	32-36M Even Frames	32-36M Odd Frames
		OFFLINE	OFFLINE

Figure 4-10. Storage Layout - SE1 and SE3 Configured Offline (3084)

The reconfiguration sequence is presented in the following order. All steps except the first one are done on the master console:

- Vary Side B online (on Side A system console)
  - Configure storage online
  - Configure CPUs online
  - Configure channel paths online
1. **Enter: VARYPHY SIDEB,ON** from the system console on Side A, in order to tell the processor controller that Side B is to be varied online as part of the MP configuration. If message SET CONFIGURATION SWITCH TO MP is displayed on the system console, change the CONFIGURATION switch to MP. When VARYPHY processing completes, message REQUEST COMPLETED is displayed on the system console.
  2. **Enter: CONFIG STOR(E=3),ONLINE** (This step and the remaining steps are entered on the master console.)

The following messages are displayed on the master console:

```
IEE524I    REAL STORAGE LOCATIONS 28M TO 32M ONLINE
IEE524I    1 OF EVERY 2 FRAMES OF REAL STORAGE LOCATIONS 32M TO 36M
           ONLINE
IEE524I    1 OF EVERY 2 FRAMES OF REAL STORAGE LOCATIONS 36M TO 40M
           ONLINE
IEE524I    1 OF EVERY 2 FRAMES OF REAL STORAGE LOCATIONS 40M TO 44M
           ONLINE
IEE524I    1 OF EVERY 2 FRAMES OF REAL STORAGE LOCATIONS 44M TO 48M
           ONLINE
IEE524I    1 OF EVERY 2 FRAMES OF REAL STORAGE LOCATIONS 48M TO 52M
           ONLINE
IEE524I    1 OF EVERY 2 FRAMES OF REAL STORAGE LOCATIONS 52M TO 56M
           ONLINE
IEE712I    CONFIG PROCESSING COMPLETE
```

Since the operating system has requested the MSSF to vary Side B online, the following message is displayed on the system console:

```
VARY SIDE B ON RECEIVED BY MSSF. RESULT=0020.
```

At this point, the communications link between the two system controllers (SC0 and SC1) is established.

After SE3 is configured online, the following message is displayed on the system console:

```
VARY STOR ELEM 03 ON RECEIVED BY MSSF. RESULT=0020.
```

Storage now appears as shown in Figure 4-11.



SE0	SE2	SE1	SE3
60-64M Even Frames HSA and Preferred	60-64M Odd Frames HSA and Preferred	56-60M Odd Frames	28-32M Odd Frames Reconfigurable
24-28M Even Frames Preferred	24-28M Odd Frames Preferred	56-60M Even Frames	28-32M Even Frames Reconfigurable
20-24M Even Frames Preferred	20-24M Odd Frames Preferred	52-56M Even Frames	52-56M Odd Frames Reconfigurable
16-20M Even Frames Preferred	16-20M Odd Frames Preferred	48-52M Even Frames	48-52M Odd Frames Reconfigurable
12-16M Even Frames SQA and Preferred	12-16M Odd Frames Preferred	44-48M Even Frames	44-48M Odd Frames Reconfigurable
8-12M Even Frames Preferred	8-12M Odd Frames Preferred	40-44M Even Frames	40-44M Odd Frames Reconfigurable
4-8M Even Frames Preferred	4-8M Odd Frames Preferred	36-40M Even Frames	36-40M Odd Frames Reconfigurable
0-4M Even Frames V = R and Preferred	0-4M Odd Frames V = R and Preferred	32-36M Even Frames	32-36M Odd Frames Reconfigurable

OFFLINE

Figure 4-11. Storage Layout - SE3 Configured Online (3084)

3. Enter: CONFIG STOR(E = 1),ONLINE

The following messages are displayed on the master console:

```

IEE524I  1 OF EVERY 2 FRAMES OF REAL STORAGE LOCATIONS 32M TO 36M
          ONLINE

IEE524I  1 OF EVERY 2 FRAMES OF REAL STORAGE LOCATIONS 36M TO 40M
          ONLINE

IEE524I  1 OF EVERY 2 FRAMES OF REAL STORAGE LOCATIONS 40M TO 44M
          ONLINE

IEE524I  1 OF EVERY 2 FRAMES OF REAL STORAGE LOCATIONS 44M TO 48M
          ONLINE

IEE524I  1 OF EVERY 2 FRAMES OF REAL STORAGE LOCATIONS 48M TO 52M
          ONLINE

IEE524I  1 OF EVERY 2 FRAMES OF REAL STORAGE LOCATIONS 52M TO 56M
          ONLINE

IEE524I  REAL STORAGE LOCATIONS 56M TO 60M ONLINE

IEE712I  CONFIG PROCESSING COMPLETE
  
```

After SE1 is configured online, the following message is displayed on the system console:

VARY STOR ELEM 01 ON RECEIVED BY MSSF. RESULT = 0020.

Storage now appears as shown in Figure 4-12.

SE0	SE2	SE1	SE3
60-64M Even Frames HSA and Preferred	60-64M Odd Frames HSA and Preferred	56-60M Odd Frames Reconfigurable	28-32M Odd Frames Reconfigurable
24-28M Even Frames Preferred	24-28M Odd Frames Preferred	56-60M Even Frames Reconfigurable	28-32M Even Frames Reconfigurable
20-24M Even Frames Preferred	20-24M Odd Frames Preferred	52-56M Even Frames Reconfigurable	52-56M Odd Frames Reconfigurable
16-20M Even Frames Preferred	16-20M Odd Frames Preferred	48-52M Even Frames Reconfigurable	48-52M Odd Frames Reconfigurable
12-16M Even Frames SQA and Preferred	12-16M Odd Frames Preferred	44-48M Even Frames Reconfigurable	44-48M Odd Frames Reconfigurable
8-12M Even Frames Preferred	8-12M Odd Frames Preferred	40-44M Even Frames Reconfigurable	40-44M Odd Frames Reconfigurable
4-8M Even Frames Preferred	4-8M Odd Frames Preferred	36-40M Even Frames Reconfigurable	36-40M Odd Frames Reconfigurable
0-4M Even Frames V = R and Preferred	0-4M Odd Frames V = R and Preferred	32-36M Even Frames Reconfigurable	32-36M Odd Frames Reconfigurable

Figure 4-12. Storage Layout - SE1 and SE3 Configured Online

If any storage in the storage elements has no assigned addresses (for example, when the system was partitioned at IPL), the operator can issue a CONFIG STOR(0M-64M) command to ensure that all storage has assigned addresses.

4. Enter: CONFIG CPU(3),ONLINE

After CPU3 is configured physically online, the following message is displayed on the system console:

VARY CPU 03 ON RECEIVED BY MSSF. RESULT = 0020.

If the TOD clocks are not synchronized, message IEA889A is displayed on the master console requesting that the TOD clock security switch be depressed to allow the synchronization of the TOD clocks. When the clocks are synchronized, the following messages are displayed on the master console:

```
IEE504I    CPU(3),ONLINE
IEE712I    CONFIG PROCESSING COMPLETE
```

5. **Enter: CONFIG CPU(1),ONLINE**

The following messages are displayed on the master console:

IEE504I CPU(1),ONLINE

IEE712I CONFIG PROCESSING COMPLETE

After CPU1 is configured online, the following message is displayed on the system console:

VARY CPU 01 ON RECEIVED BY MSSF. RESULT=0020.

6. **Enter: CONFIG CHP(ALL,1),ON**

The following messages are displayed on the master console:

IEE502I CHP(ALL,1),ONLINE

IEE712I CONFIG PROCESSING COMPLETE

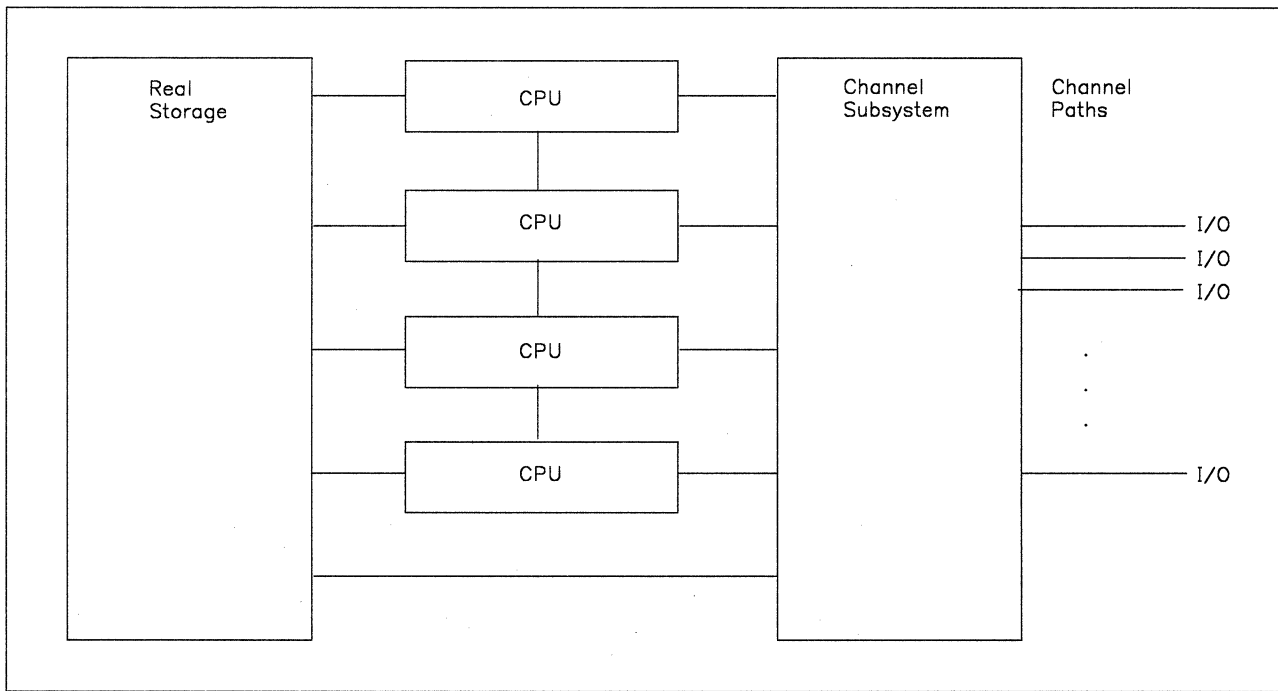
The operating system configures the EXDC online and the following message is displayed on the system console:

VARY I/O SIDE B ON RECEIVED BY MSSF. RESULT=0020.

Then, the operating system configures online the individual channel paths owned by the EXDC on SIDE B. As the individual channel paths are configured online, the following message is displayed on the system console:

VARY CHAN PATH nn ON RECEIVED BY MSSF. RESULT=0020.

After all the channel paths are configured online, the system should be operating in single-image mode. After configuring from physically partitioned mode to single-image mode, the processor complex appears as shown in Figure 4-13.



**Figure 4-13. Single-Image Mode of a 3084**

Figure 4-14 shows a series of D M commands, directed at the various elements, that can be used to verify single-image mode.

D M=SIDE

```
IEE174I hh.mm.ss MATRIX DISPLAY
SIDE STATUS
SIDE:          0                1
STATUS:        ONLINE          ONLINE
I/O ENGINE:    0                1
CPU:           0 2              1 3
CHP:           0-7 10-17 20-27 40-47 50-57 60-67
STOR(E=x):     0 2              1 3

TOTAL STOR:    64M
*=OFFLINE
```

---

D M=CPU

```
IEE174I hh.mm.ss MATRIX DISPLAY
PROCESSOR STATUS
CPU  STATUS      SERIAL
0   ONLINE      0123453084
1   ONLINE      1123453084
2   ONLINE      2123453084
3   ONLINE      3123453084
```

---

D M=CHP

```
IEE174I hh.mm.ss MATRIX DISPLAY
CHANNEL PATH STATUS
  0 1 2 3 4 5 6 7 8 9 A B C D E F
0 + + + + + + + + . . . . .
1 + + + + + + + + . . . . .
2 + + + + + + + + . . . . .
4 + + + + + + + + . . . . .
5 + + + + + + + + . . . . .
6 + + + + + + + + . . . . .
***** SYMBOL EXPLANATIONS *****
. DOES NOT EXIST          * LOGICALLY OFF, PHYSICALLY ONLINE
- LOGICALLY & PHYSICALLY OFFLINE + LOGICALLY & PHYSICALLY ONLINE
```

---

D M=STOR(E)

```
IEE174I hh.mm.ss MATRIX DISPLAY
STORAGE ELEMENT STATUS
0 : OWNED STORAGE = 16M STATUS = ONLINE
1 : OWNED STORAGE = 16M STATUS = ONLINE
2 : OWNED STORAGE = 16M STATUS = ONLINE
3 : OWNED STORAGE = 16M STATUS = ONLINE
```

Figure 4-14. Examples of D M Displays - 3084 System in Single-Image Mode

## Examples of Partitioning and Merging a Partitionable 3090

Two examples of reconfiguration are presented in this section: configuring from single-image mode to physically partitioned mode and configuring from physically partitioned mode to single-image mode.

The following examples show:

- The required commands use to partition and merge
- The messages that are issued during processing
- How the contents of real storage are handled during processing

Before describing the examples themselves, this section lists significant differences between the 3090 Models 400, 400E, 600E, and the 3084. By noting the differences between these machine types, you should be able to understand the assumptions that underlie the examples. The differences are listed in Figure 4-15.

CHARACTERISTIC	3084	3090 Model 400	3090 Models 400E/600E
Real Storage Subincrement Size	2MB, 4MB	1MB	2MB
Numb Subincrements per Element	8	32, 64	16, 32
Size of Real Storage Element	16MB, 32MB	32MB, 64MB	32MB, 64MB
Real Storage Ranges	64MB--128MB	128MB, 256MB	128MB, 256MB
Real Storage Element IDs	Side 0: 0, 2 Side 1: 1, 3	Side 0: 0, 1 Side 1: 2, 3	Side 0: 0, 1 Side 1: 2, 3
Extended Storage Ranges	None	0MB, 128MB, 256MB, 384MB, 512MB, 1024MB	0MB, 128MB, 256MB, 384MB, 512MB, 1024MB
Extended Storage Element IDs	None	Side 0: 0, 1 Side 1: 2, 3	Side 0: 0, 1 Side 1: 2, 3
Size of Extended Storage Elements	Not Applicable	64MB, 128MB	64MB, 128MB, 256MB
CPU IDs	Side 0: 0, 2 Side 1: 1, 3	Side 0: 1, 2 Side 1: 3, 4	Side 0: 0, 1, 2 Side 1: 3, 4, 5
Can Have Vector Facilities	No	Yes	Yes

**Figure 4-15. Differences Between the 3084 and the 3090 Models 400, 400E, and 600E**

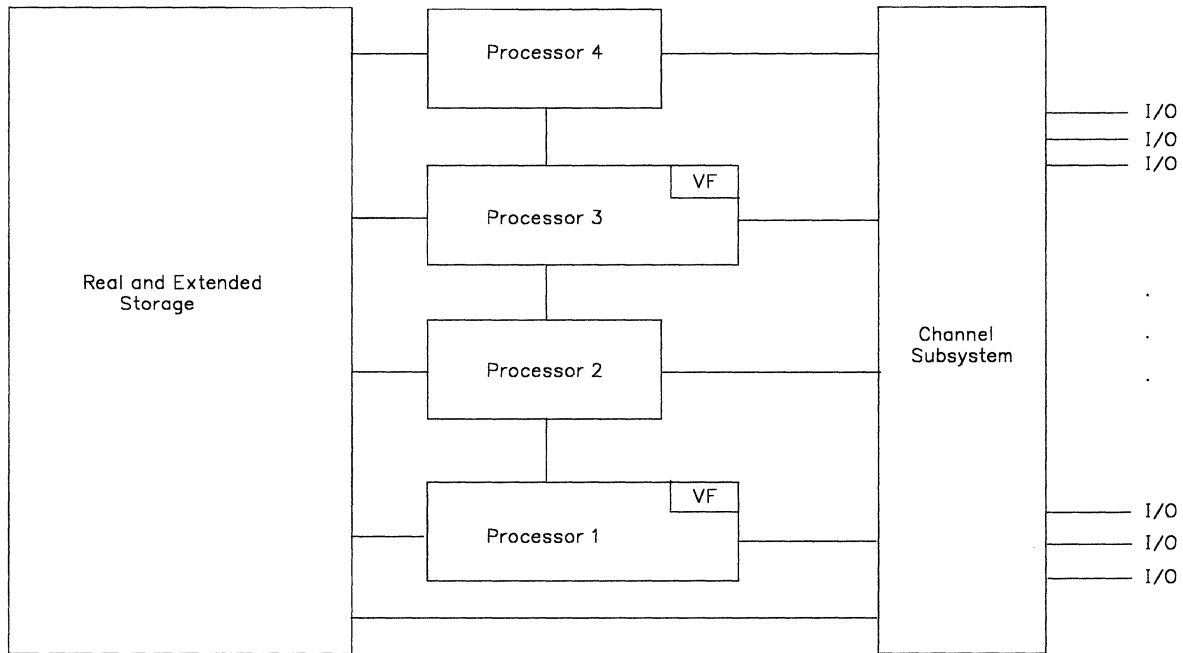
In each of these 3090 examples, assume the following:

- 96 channel paths.
- Installed real storage of 128M consisting of four 32M real storage elements with 2M storage increments.
- RSU = 32 specified at IPL.
- V = R area contained in real storage increment 0-2M.
- SQA contained in 12-16M even frames.
- Real storage ranges 0-58M and 122-128M are not reconfigurable

- Real storage range 58-122M reconfigurable.
- HSA contained in storage increment 126-128M.
- Installed extended storage of 512M, consisting of four 128M extended storage elements, two on each side.

**Partitioning from Single-Image Mode to Physically Partitioned Mode (Side 1 to Be Configured Offline)**

Prior to configuring from single-image mode to physically partitioned mode, the 3090 Model 400 processor complex appears as shown in Figure 4-16.



**Figure 4-16. Single-Image Mode of a 3090 Model 400**

Prior to configuring from single-image mode to physically partitioned mode, assume the 3090 Model 400 storage layout as shown in Figure 4-17.

Side 0

Side 1

SE0	SE1	SE2	SE3
Online	Online	Online	Online
58-60M Odd Frames Reconfigurable	62-64M Odd Frames Reconfigurable	122-124M Odd Frames Preferred	126-128M Odd Frames Preferred
58-60M Even Frames Reconfigurable	62-64M Even Frames Reconfigurable	122-124M Even Frames Preferred	126-128M Even Frames Preferred
56-58M Odd Frames Preferred	60-62M Odd Frames Reconfigurable	120-122M Odd Frames Reconfigurable	124-126M Odd Frames Preferred
56-58M Even Frames Preferred	60-62M Even Frames Reconfigurable	120-122M Even Frames Reconfigurable	124-126M Even Frames Preferred
...	...	...	...
2-4M Even Frames Preferred	6-8M Even Frames Preferred	66-68M Even Frames Reconfigurable	70-72M Even Frames Reconfigurable
0-2M Odd Frames Reconfigurable	4-6M Odd Frames Preferred	64-66M Odd Frames Reconfigurable	68-70M Odd Frames Reconfigurable
0-2M Even Frames Preferred	4-6M Even Frames Preferred	64-66M Even Frames Reconfigurable	68-70M Even Frames Reconfigurable

Notes

SE0	SE1	SE2	SE3
<b>3 reconfigurable subincrements</b>	<b>4 reconfigurable subincrements</b>	<b>30 reconfigurable subincrements</b>	<b>28 reconfigurable subincrements</b>
<b>29 preferred subincrements</b>	<b>28 preferred subincrements</b>	<b>2 preferred subincrements</b>	<b>4 preferred subincrements</b>

This figure assumes a Power-on-Reset in single-image mode and an IPL with RSU = 32.

Figure 4-17. Sample Real Storage Layout of 3090 Model 400 Before Partitioning

An operator performs the following steps to physically partition a system. All except the last step are done on the master console; the last step is done on the system console. The reconfiguration is presented in this order:

- Issue D M = SIDE to determine which resources are on each side.
- Configure channel paths offline.
- Configure CPUs and Vector Facilities offline. (Vector Facilities go offline with their CPUs.)
- Configure extended storage offline.



- Configure real storage offline.
- Define the configuration, power-on-reset, and IPL.

*Note:* Although the following example relates to a Model 400, the commands to partition either a 3090 Model 400E or 600E are identical to those shown, except that for a Model 600E the CF CPU command must also specify CPU 5.

1. **Enter: DM=SIDE** to determine the status of resources on each side of the 3090.

The system status would look like this:

```
IEE174I 05.43.18 DISPLAY M
SIDE STATUS
SIDE:           0                1
STATUS:        ONLINE          ONLINE
CPU:           1-2             3-4
VF:           1                3
CHP:           0-2F           40-6F
STOR(E=X):    0-1             2-3
ESTOR(E=X):   0-1             2-3

TOTAL STOR:    128M UNASSIGNED: 0M
TOTAL ESTOR:   512M
```

2. **Enter: CF CHP(ALL,1),OFFLINE,UNCOND**

*Note:* When partitioning, before issuing the CONFIG CHP(ALL,n),OFFLINE,UNCOND command, complete or cancel any mounts that may be affected by this command.

The following messages are displayed on the master console:

```
IEE097I 05.17.30 DEVIATION STATUS
          FROM CONFIG COMMAND
          NO DEVIATION FROM REQUESTED CONFIGURATION

IEE172I ALL CHANNEL PATHS ON SIDE 1 NOW OFFLINE

IEE503I CHP(ALL,1),OFFLINE

IEE712I CONFIG PROCESSING COMPLETE
```

3. **Enter: CF CPU(3,4),OFFLINE**

The following messages are displayed on the master console:

```
IEE097I 05.23.30 DEVIATION STATUS
          FROM CONFIG COMMAND
          NO DEVIATION FROM REQUESTED CONFIGURATION

IEE505I CPU(3),OFFLINE

IEE505I VF(3),OFFLINE

IEE505I CPU(4),OFFLINE

IEE712I CONFIG PROCESSING COMPLETE
```

The Vector Facility associated with CPU3 goes offline with its CPU. When CPU3 is later brought online, the Vector Facility will come online also.

*Note:* If the removal of a CPU would take offline the last available Vector Facility on the 3090, and vector jobs are scheduled, an operator action is needed. This is described in a previous topic in this chapter, "Removing the Last Vector Facility" on page 4-11.

4. **Enter: CF ESTOR(E = 2),OFFLINE** to configure offline an extended storage element on side 1.

The master console in response indicates that the command was accepted.

```
IEE097I    05.26.35 DEVIATION STATUS
              FROM CONFIG COMMAND
              NO DEVIATION FROM REQUESTED CONFIGURATION
IEE510I    EXTENDED STORAGE LOCATIONS 256M TO 384M OFFLINE
IEE526I    EXTENDED STORAGE ELEMENT(2) OFFLINE
IEE712I    CONFIG PROCESSING COMPLETE
```

5. **Enter: CF ESTOR(E = 3),OFFLINE** to configure the other extended storage element offline.

The master console in response indicates that the command was accepted:

```
IEE097I    05.28.45 DEVIATION STATUS
              FROM CONFIG COMMAND
              NO DEVIATION FROM REQUESTED CONFIGURATION
IEE510I    EXTENDED STORAGE LOCATIONS 384M TO 512M OFFLINE
IEE526I    EXTENDED STORAGE ELEMENT(3) OFFLINE
IEE712I    CONFIG PROCESSING COMPLETE
```

The reconfiguration of extended storage elements on side 1 is now complete.

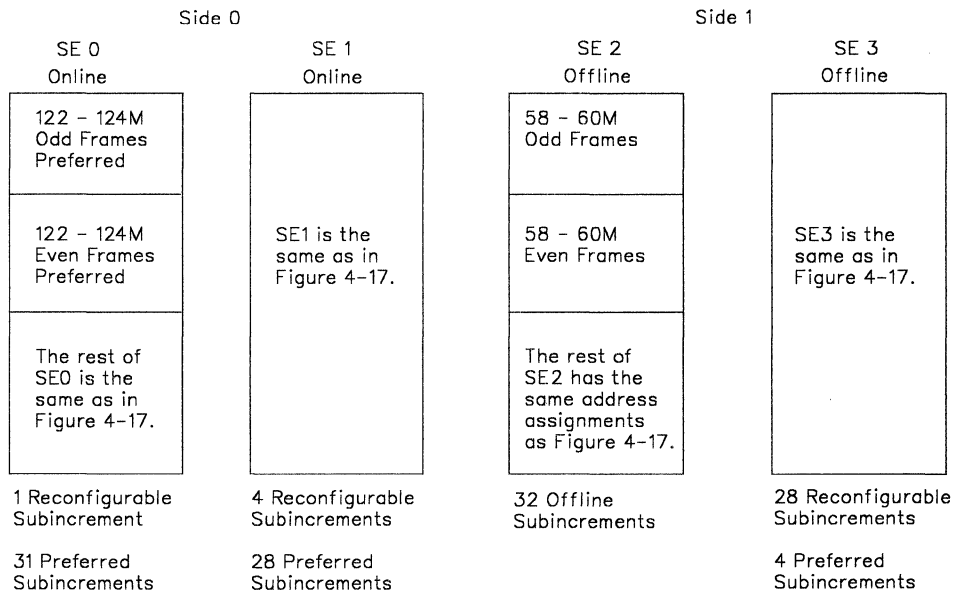
6. **Enter: CF STOR(E = 2),OFFLINE**

These messages are displayed on the master console:

```
IEE510I    REAL STORAGE LOCATIONS 58M TO 60M OFFLINE
IEE510I    REAL STORAGE LOCATIONS 64M TO 68M OFFLINE
IEE510I    REAL STORAGE LOCATIONS 72M TO 76M OFFLINE
IEE510I    REAL STORAGE LOCATIONS 80M TO 84M OFFLINE
IEE510I    REAL STORAGE LOCATIONS 88M TO 92M OFFLINE
IEE510I    REAL STORAGE LOCATIONS 96M TO 100M OFFLINE
IEE510I    REAL STORAGE LOCATIONS 104M TO 108M OFFLINE
```

- IEE510I REAL STORAGE LOCATIONS 112M TO 116M OFFLINE
- IEE510I REAL STORAGE LOCATIONS 120M TO 122M OFFLINE
- IEE526I REAL STORAGE ELEMENT(2) OFFLINE
- IEE712I CONFIG PROCESSING COMPLETE

After you configure SE2 offline, storage appears as shown in Figure 4-18. Because SE2 contains some HSA and preferred storage (the 122M--124M range in Figure 4-17), the reconfiguration process consists of swapping this group of frames with the frames in the 58M--60M range in SE0. This is why the storage from 58M--60M went offline.



**Figure 4-18. Real Storage Layout - SE2 Configured Offline (3090 Model 400)**

*Note:* The preferred even subincrements 122M to 124M and the odd subincrements 122M to 124M have been swapped into SE0 and have been replaced with the reconfigurable subincrements 58M to 60M, both odd and even.

7. Enter: **CF STOR(E=3),OFFLINE**

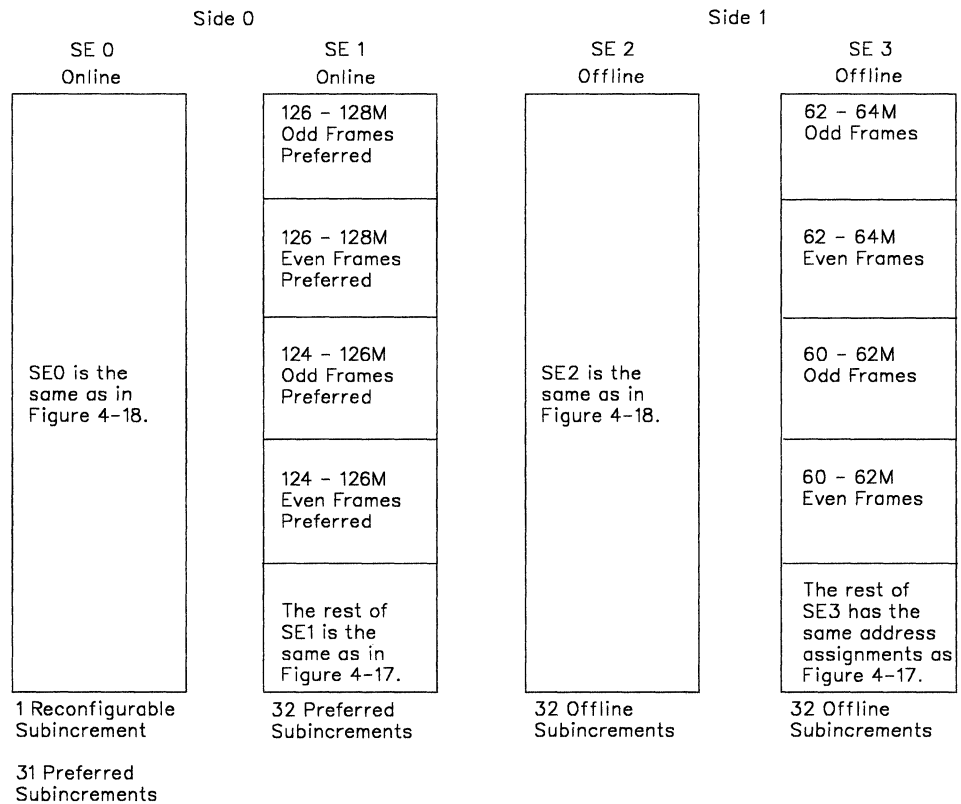
The following messages are displayed on the master console:

**IEE575A** CONFIG STORAGE WAITING TO COMPLETE - REPLY C TO CANCEL

*Note:* There can be more than one IEE575A message before the series of IEE510I messages

**IEE097I** 05.31.45 DEVIATION STATUS  
FROM CONFIG COMMAND  
NO DEVIATION FROM REQUESTED CONFIGURATION  
**IEE510I** REAL STORAGE LOCATIONS 60M TO 64M OFFLINE  
**IEE510I** REAL STORAGE LOCATIONS 68M TO 72M OFFLINE  
**IEE510I** REAL STORAGE LOCATIONS 76M TO 80M OFFLINE  
**IEE510I** REAL STORAGE LOCATIONS 84M TO 88M OFFLINE  
**IEE510I** REAL STORAGE LOCATIONS 92M TO 96M OFFLINE  
**IEE510I** REAL STORAGE LOCATIONS 100M TO 104M OFFLINE  
**IEE510I** REAL STORAGE LOCATIONS 108M TO 112M OFFLINE  
**IEE510I** REAL STORAGE LOCATIONS 116M TO 120M OFFLINE  
**IEE526I** REAL STORAGE ELEMENT(3) OFFLINE  
**IEE712I** CONFIG PROCESSING COMPLETE

After configuring SE3 offline, storage appears as shown in Figure 4-19.

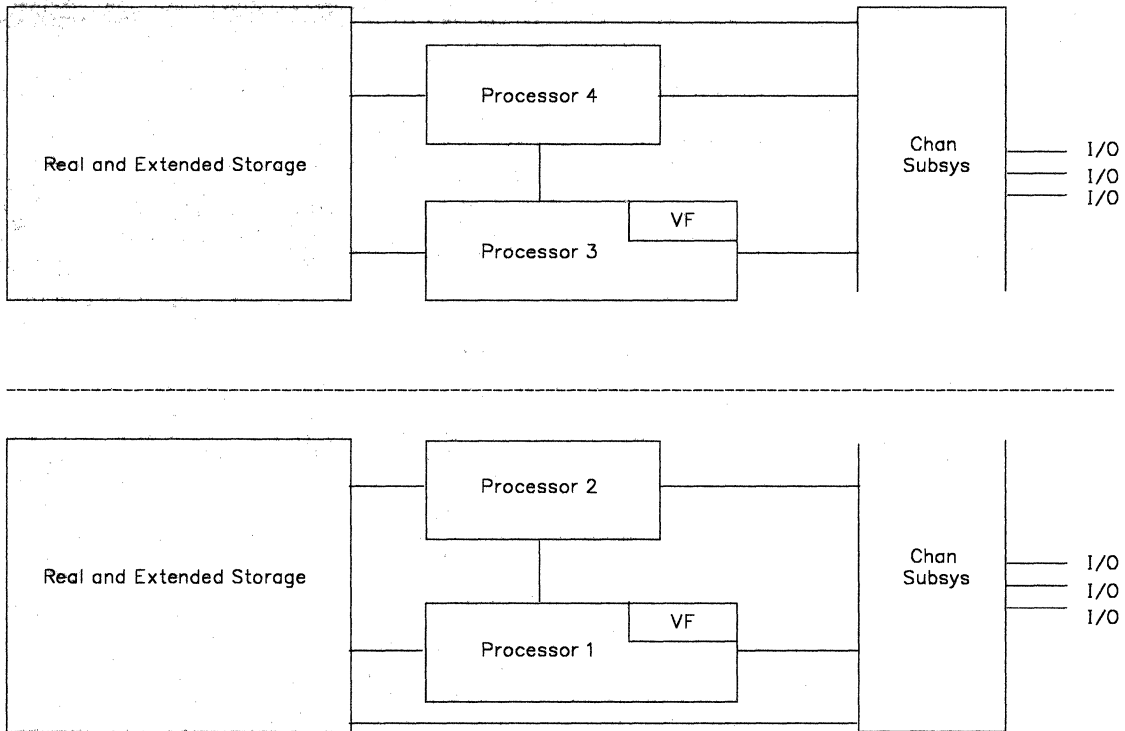


**Figure 4-19. Real Storage Layout - SE3 Configured Offline (3090 Model 400)**

*Note:* The preferred subincrements 124-126M even, 124-126M odd, 126-128M even, and 126-128M odd (see SE3 in Figure 4-17 on page 4-34) have been swapped into storage element 1, and have been replaced in SE3 with reconfigurable subincrements 60-62M even, 60-62M odd, 62-64 even, and 62-64M odd.

8. This step and the next step in partitioning involve several actions that use the system console. On the system console use the partition control frame (PARCTL) to vary side 1 offline. When processing completes, the side-1 system console becomes active.
9. To bring up the side that was taken offline, perform the following:
  - a. Define the configuration, using the Configuration frame on the side 1 system console.
  - b. Do Power-on-Reset at the side 1 system console.
  - c. IPL side 1.

At this point, the installation has two separate systems, each monitored by its own service processor and each with its own system and service consoles. (See Figure 4-20.)



**Figure 4-20. Physically Partitioned Mode of the 3090 Model 400**

At the master console on Side 0, the operator can verify the physical partitioning of the system by use of the D M=SIDE command. If more information is needed, he can then issue the D M command. Sample displays are shown in Figure 4-21.

```

D M=SIDE

IEE174I hh.mm.ss MATRIX DISPLAY
SIDE STATUS
SIDE:           0                               1
STATUS:         ONLINE                          UNAVAILABLE
CPU:            1-2
VF:             1
CHP:            0-2F
STOR(E=x):      0-1
ESTOR(E=x):     0-1

TOTAL STOR:     128M          UNASSIGNED: 0M
TOTAL ESTOR:    512M
*=OFFLINE
-----

D M
IEE174I hh.mm.ss DISPLAY M nnn
PROCESSOR STATUS
CPU STATUS      SERIAL
 1 ONLINE VFON 1700903090
 2 ONLINE      2700903090

CHANNEL PATH STATUS
  0 1 2 3 4 5 6 7 8 9 A B C D E F
 0 + + + + + + + + + + + + + + +
 1 + + + + + + + + + + + + + + +
 2 + + + + + + + + + + + + + + +
*****SYMBOL EXPLANATIONS*****
+ ONLINE          - OFFLINE          . DOES NOT EXIST

HSA STATUS
ADDRESS=7F80000 LENGTH=512K
ADDRESS=7DD0000 LENGTH=192K

STORAGE SIZE STATUS
HIGH REAL STORAGE ADDRESS IS 128M
HIGH EXTENDED STORAGE ADDRESS IS 512M

REAL STORAGE STATUS
ONLINE-NOT RECONFIGURABLE
FIRST 4K OF EVERY 8K FROM 0M TO 2M
2M-58M
122-128M
ONLINE-RECONFIGURABLE
SECOND 4K OF EVERY 8K FROM 0M TO 2M
PENDING OFFLINE
NONE
0M IN OFFLINE STORAGE ELEMENT(S)
0M UNASSIGNED
64M IN ANOTHER CONFIGURATION

```

Figure 4-21 (Part 1 of 2). Examples of D M Displays - 3090 Model 400 System in Physically Partitioned Mode

```

REAL STORAGE ELEMENT STATUS
0: OWNED STORAGE=32M UNASSIGNED STORAGE=0M STATUS=ONLINE
1: OWNED STORAGE=32M UNASSIGNED STORAGE=0M STATUS=ONLINE
STOR(E=2) IS PART OF ANOTHER CONFIGURATION-NO STATUS OBTAINED
STOR(E=3) IS PART OF ANOTHER CONFIGURATION-NO STATUS OBTAINED

EXTENDED STORAGE STATUS
ONLINE-RECONFIGURABLE
OM-256M
PENDING OFFLINE
NONE
OM IN OFFLINE STORAGE ELEMENT(S)
256M IN ANOTHER CONFIGURATION

EXTENDED STORAGE ELEMENT STATUS
0: OWNED STORAGE=128M STATUS =ONLINE
1: OWNED STORAGE=128M STATUS =ONLINE
ESTOR(E=2) IS PART OF ANOTHER CONFIGURATION - NO STATUS OBTAINED
ESTOR(E=3) IS PART OF ANOTHER CONFIGURATION - NO STATUS OBTAINED

SIDE STATUS
SIDE:           0                               1
STATUS:         ONLINE                          UNAVAILABLE
CPU:            1-2
VF:            1
CHP:           0-2F
STOR(E=x):     0-1
ESTOR(E=x):    0-1

TOTAL STOR:    128M          UNASSIGNED: 0M
TOTAL ESTOR:  512M
*=OFFLINE

```

**Figure 4-21 (Part 2 of 2). Examples of D M Displays - 3090 Model 400 System in Physically Partitioned Mode**

**Merging from Physically Partitioned Mode to Single-Image Mode (Side 1 To Be Configured Online)**

The process of merging (that is, configuring from physically partitioned mode to single-image mode) is essentially the reverse of partitioning (configuring from single-image mode to physically partitioned mode). In this example side 1 of the partitioned system is to be merged with the system running on side 0.

*Note:* Although the following example relates to a Model 400, the commands to merge either a 3090 Model 400E or 600E are identical to those shown, except that for a Model 600E the CF CPU command must also specify CPU 5.

The command sequence to implement this merge is:

- Quiesce any programming system running on side 1.
- Use the PARCTL frame to vary side 1 offline at side 1's system console.
- Use the PARCTL frame to vary side 1 online at side 0's system console.
- Configure real storage elements online



- Configure real storage online in the elements (if needed)
  - Configure extended storage elements online
  - Configure CPUs online
  - Configure channel paths online
1. **Quiesce the control program on side 1.** One way to do this is to issue the QUIESCE command at the side 1 MVS master console. When the command completes, all the processors on side 1 will be in the 'CCC' restartable wait state.
  2. **Vary side 1 offline** by use of the Partition Control frame (PARCTL) on the side 1 system console.
  3. **Vary side 1 online**, this time using the Partition Control frame (PARCTL) on the *side 0* system console.

*Note:* During the merging process the hardware will be initializing the backup processor controller DASD. This hardware action does *not* prevent MVS from configuring online the side 1 resources. When side 1 has come online, issue the following commands at the side 0 MVS master console (steps {step4} through 7 on page 4-45):

4. **Enter: CF STOR(E=2),ONLINE and CF STOR(E=3),ONLINE** for the two side 1 storage elements. The expected response to each of these commands is a series of IEE524I messages that indicate that various storage ranges have come online (see "Responses for CF STOR(E=2),ONLINE" on page 4-44 and "Responses for CF STOR(E=3),ONLINE" on page 4-44).

If, however, the storage in a specified storage element does not come online, it is because it does not have assigned storage addresses. You will receive a single message (instead of the usual series):

IEE574I NO STORAGE TO COME ONLINE IN REAL STORAGE ELEMENT(x).

In this case, issue any remaining CF STOR(E=x) command not yet issued, and enter this two-step procedure:

- D M=STOR to find out the amount (ddM UNASSIGNED) of storage that does not have assigned addresses.
- CF STOR(ddM),ONLINE to assign storage addresses to this storage.

The previously unavailable storage in storage element x should now come online.

If all real storage elements are now online, continue by configuring extended storage (see step 5 on page 4-44).

### Responses for CF STOR(E = 2),ONLINE

The response at the side 0 master console is:

```
IEE097I 05.18.30 DEVIATION STATUS
              FROM CONFIG COMMAND
              NO DEVIATION FROM REQUESTED CONFIGURATION
IEE524I REAL STORAGE LOCATIONS 58M TO 60M ONLINE
IEE524I REAL STORAGE LOCATIONS 64M TO 68M ONLINE
IEE524I REAL STORAGE LOCATIONS 72M TO 76M ONLINE
IEE524I REAL STORAGE LOCATIONS 80M TO 84M ONLINE
IEE524I REAL STORAGE LOCATIONS 88M TO 92M ONLINE
IEE524I REAL STORAGE LOCATIONS 96M TO 100M ONLINE
IEE524I REAL STORAGE LOCATIONS 104M TO 108M ONLINE
IEE524I REAL STORAGE LOCATIONS 112M TO 116M ONLINE
IEE524I REAL STORAGE LOCATIONS 120M TO 122M ONLINE
IEE526I REAL STORAGE ELEMENT(2) ONLINE
IEE712I CONFIG PROCESSING COMPLETE
```

### Responses for CF STOR(E = 3),ONLINE

The response at the side 0 master console is:

```
IEE097I 05.20.35 DEVIATION STATUS
              FROM CONFIG COMMAND
              NO DEVIATION FROM REQUESTED CONFIGURATION
IEE524I REAL STORAGE LOCATIONS 60M TO 64M ONLINE
IEE524I REAL STORAGE LOCATIONS 68M TO 72M ONLINE
IEE524I REAL STORAGE LOCATIONS 76M TO 80M ONLINE
IEE524I REAL STORAGE LOCATIONS 84M TO 88M ONLINE
IEE524I REAL STORAGE LOCATIONS 92M TO 96M ONLINE
IEE524I REAL STORAGE LOCATIONS 100M TO 104M ONLINE
IEE524I REAL STORAGE LOCATIONS 108M TO 112M ONLINE
IEE524I REAL STORAGE LOCATIONS 116M TO 120M ONLINE
IEE526I REAL STORAGE ELEMENT(3) ONLINE
IEE712I CONFIG PROCESSING COMPLETE
```

At this point all the storage in storage elements 2 and 3 is reconfigurable. All the preferred storage is in storage elements 0 and 1.

#### 5. Enter: CF ESTOR(E = 2),ONLINE

The response at the side 0 master console is:

```
IEE097I 05.17.30 DEVIATION STATUS
              FROM CONFIG COMMAND
              NO DEVIATION FROM REQUESTED CONFIGURATION
IEE524I EXTENDED STORAGE LOCATIONS 256M TO 384M ONLINE
IEE526I EXTENDED STORAGE ELEMENT(2) ONLINE
IEE712I CONFIG PROCESSING COMPLETE
```

#### 6. Enter: CF ESTOR(E = 3),ONLINE

The response at the side 0 master console is:

```
IEE097I 05.20.35 DEVIATION STATUS
              FROM CONFIG COMMAND
              NO DEVIATION FROM REQUESTED CONFIGURATION
IEE524I EXTENDED STORAGE LOCATIONS 384M TO 512M ONLINE
IEE526I EXTENDED STORAGE ELEMENT(3) ONLINE
IEE712I CONFIG PROCESSING COMPLETE
```

7. Enter: CF CPU(3,4),ONLINE

The response at the side 0 master console is:

```
*09 IEA889A DEPRESS TOD CLOCK SECURITY SWITCH
R 09,Y
IEE600I REPLY TO 09 IS;Y
*10 IEA889A DEPRESS TOD CLOCK SECURITY SWITCH
R 10,Y
IEE600I REPLY TO 10 IS;Y
IEE097I 05.21.03 DEVIATION STATUS
          FROM CONFIG COMMAND
          NO DEVIATION FROM REQUESTED CONFIGURATION
IEE504I CPU(3),ONLINE
IEE504I VF(3),ONLINE
IEE504I CPU(4),ONLINE
IEE712I CONFIG PROCESSING COMPLETE
```

8. Enter: CF CHP(ALL,1),ONLINE

The response at the side 0 master console is:

```
IEE097I 05.23.46 DEVIATION STATUS
          FROM CONFIG COMMAND
          NO DEVIATION FROM REQUESTED CONFIGURATION
IEE754I NOT ALL PATHS BROUGHT ONLINE WITH CHP(4B)
IEE754I NOT ALL PATHS BROUGHT ONLINE WITH CHP(4E)
IEE754I NOT ALL PATHS BROUGHT ONLINE WITH CHP(5D)
IEE172I ALL CHANNEL PATHS ON SIDE 1 ARE NOW ONLINE
IEE502I CHP(ALL,1),ONLINE
IEE712I CONFIG PROCESSING COMPLETE
```

After all the channel paths have been configured online, the system should be operating in single-image mode. The processor complex now appears as shown in Figure 4-22.

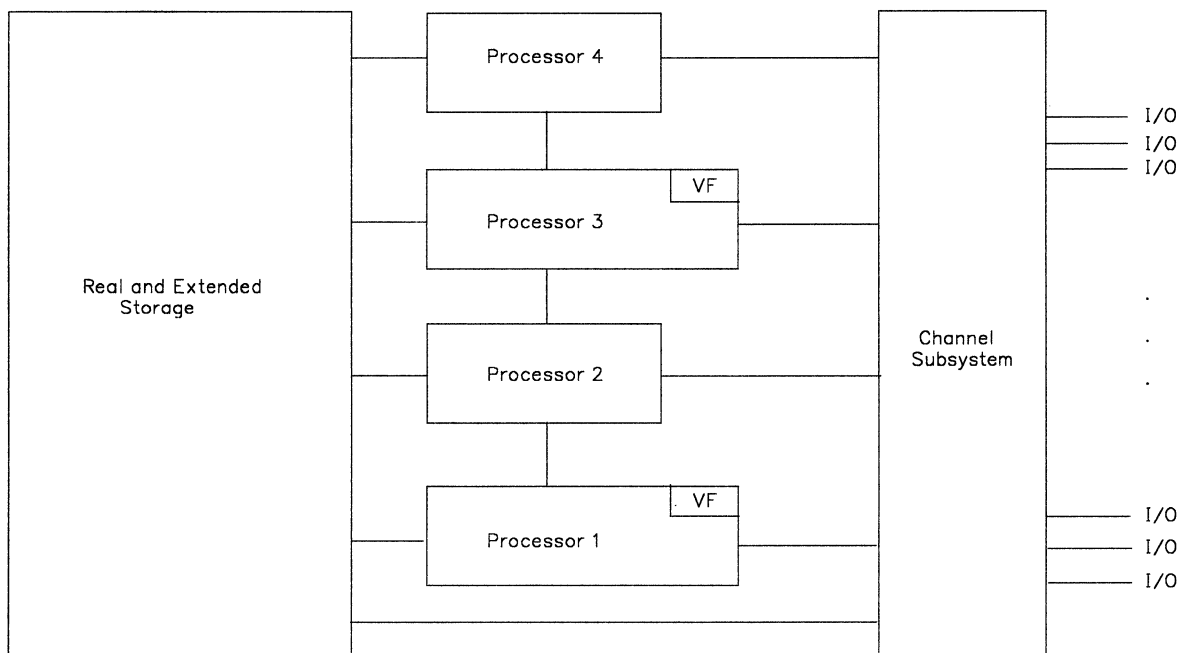


Figure 4-22. Single-Image Mode of a 3090 Model 400

9. Enter the **D M=SIDE** command to verify single image mode. If more information is needed, use the **D M** command, as shown in Figure 4-23.

```

D M=SIDE

IEE174I hh.mm.ss MATRIX DISPLAY
SIDE STATUS
SIDE:           0                1
STATUS:         ONLINE          ONLINE
CPU:            1-2             3-4
VF:             1                3
CHP:            0-2F           40-6F
STOR(E=x):     0-1             2-3
STOR(E=x):     0-1             2-3
ESTOR(E=x):    0-1             2-3

TOTAL STOR:    128M             UNASSIGNED: 0M
TOTAL ESTOR:   512M
*=OFFLINE

-----

D M

IEE174I hh.mm.ss DISPLAY M

PROCESSOR STATUS
CPU STATUS      SERIAL
 1 ONLINE VFON 1700903090
 2 ONLINE          2700903090
 3 ONLINE VFON 3700903090
 4 ONLINE          4700903090

CHANNEL PATH STATUS
 0 1 2 3 4 5 6 7 8 9 A B C D E F
0 + + + + + + + + + + + + + + +
1 + + + + + + + + + + + + + + +
2 + + + + + + + + + + + + + + +
3 + + + + + + + + + + + + + + +
4 + + + + + + + + + + + + + + +
5 + + + + + + + + + + + + + + +
6 + + + + + + + + + + + + + + +
***** SYMBOL EXPLANATIONS *****
+ ONLINE          - OFFLINE          . DOES NOT EXIST
HSA STATUS
ADDRESS=7F80000 LENGTH=512K
ADDRESS=7DD0000 LENGTH=192K

STORAGE SIZE STATUS
HIGH REAL STORAGE ADDRESS IS 128M
HIGH EXTENDED STORAGE ADDRESS IS 512K

```

Figure 4-23 (Part 1 of 2). Examples of D M Displays - 3090 Model 400 System in Single-Image Mode

```

REAL STORAGE STATUS
ONLINE-NOT RECONFIGURABLE
  FIRST 4K OF EVERY 8K FROM 0M TO 2M
  2M-58M
  122M-128M
ONLINE-RECONFIGURABLE
  SECOND 4K OF EVERY 8K FROM 0M TO 2M
  58M-122M
PENDING OFFLINE
  NONE
OM IN OFFLINE STORAGE ELEMENT(S)
OM UNASSIGNED
OM IN ANOTHER CONFIGURATION
REAL STORAGE ELEMENT STATUS
0: OWNED STORAGE=32M  UNASSIGNED STORAGE=0M  STATUS=ONLINE
1: OWNED STORAGE=32M  UNASSIGNED STORAGE=0M  STATUS=ONLINE
2: OWNED STORAGE=32M  UNASSIGNED STORAGE=0M  STATUS=ONLINE
3: OWNED STORAGE=32M  UNASSIGNED STORAGE=0M  STATUS=ONLINE

EXTENDED STORAGE STATUS
ONLINE-RECONFIGURABLE
  0M-512M
PENDING OFFLINE
  NONE
OM IN OFFLINE STORAGE ELEMENT(S)
OM IN ANOTHER CONFIGURATION

EXTENDED STORAGE ELEMENT STATUS
0: OWNED STORAGE=128M STATUS=ONLINE
1: OWNED STORAGE=128M STATUS=ONLINE
2: OWNED STORAGE=128M STATUS=ONLINE
3: OWNED STORAGE=128M STATUS=ONLINE

SIDE STATUS
SIDE:           0                               1
STATUS:        ONLINE                          ONLINE
CPU:           1-2                             3-4
VF:            1                               3
CHP:           0-2F                            40-6F
STOR(E=X):     0-1                             2-3
ESTOR(E=X):    0-1                             2-3

TOTAL STOR:    128M  UNASSIGNED: 0M
TOTAL ESTOR:   512M
*=OFFLINE

```

Figure 4-23 (Part 2 of 2). Examples of D M Displays - 3090 Model 400 System in Single-Image Mode



# Index

## A

- ABEND code 522
  - avoidance through TIME parm on JOB or EXEC statement 3-20
- ACR
  - (see also this topic under excessive spin loop recovery)
  - considerations 3-23
  - introduction to 3-8
- ALTCTRL (FEATURE = ALTCTRL) 2-13
- alternate CPU recovery (see ACR)
- alternate master console (see master console configuration guidelines)
- alternate path recovery (APR) 3-16
- APR (see alternate path recovery)

## C

- central processing unit errors 3-4
- channel measurements
  - in physical partitioning or merging 4-11
- channel path
  - alert conditions 3-14
  - reconfiguration 4-14
  - recovery 3-12
- channel report word (CRW) 3-12
- channel subsystem
  - considerations
    - 308x 2-2
    - 3090 2-2
    - 4381 2-4
  - diagram (308x) 2-2
  - errors 3-12
    - hot I/O 3-16
    - missing interrupts 3-15
  - monitoring facility recovery 3-14
- CONFIG ESTOR(E=id) 4-9
- configuration considerations
  - I/O 2-1
  - master console 2-12
- configuration switch, 3084 4-16
- configuring
  - channel paths 4-14
  - devices for a nonpartitionable processor complex 2-5
  - I/O devices 4-15
  - processors 4-10
  - storage 4-6
  - 3084 from physically partitioned mode to single image 4-25
  - 3084 from single-image mode to physically partitioned 4-17

- 3090 Model 400 from physically partitioned mode to single image 4-42
- 3090 Model 400 from single-image mode to physically partitioned 4-33
- CONFIGxx parmlib member
  - use of 4-1
- considerations
  - DASD 3-17
  - DASD (3380/3880) 3-17
  - IOCP 2-12
  - program properties table 4-5
  - reconfiguration 4-2
  - XA configuration 2-13
- control units dedicated to consoles 2-12
- CPU
  - reconfiguration 4-10
  - restarting 3-19
- CPU errors 3-4
  - hard 3-5
  - soft 3-5
  - terminating (see also ACR) 3-6
- CRW (channel report word) 3-12

## D

- D M command 4-9
  - considerations 4-4
  - examples, 3084
    - in physically partitioned mode 4-24
    - in single image mode 4-31
  - examples, 3090 Model 400
    - in single image mode 4-46
- D U command
  - considerations 4-4
- DASD
  - configuration
    - 3084 in single-image mode 2-11
  - configuration example
    - 308x 2-6
  - considerations 3-17
  - maintenance and recovery 3-18
- data facility data set services (DFDSS) 3-18
- data server element (DSE) 2-2
- DCCF 2-12
- dedicated control units for consoles 2-12
- definitions and terminology 1-1
- deselecting hardware units 4-1
- device
  - attachment 2-1
  - configuration examples
    - DASD 2-6
    - tape 2-7
    - 3725 2-8

- configuration for a nonpartitionable processor complex 2-5
- configurations
  - DASD on 3084 in single-image mode 2-11
- considerations
  - 308x 2-2
  - 3090 2-2
  - 4381 2-4
- path (308x), hardware elements in 2-2
- TP device, local 2-9
- unit record 2-9
- device support facilities 3-18
- DFDSS (data facility data set services) 3-18
- disabled console communications facility 2-12
- disabled wait states 3-21
- DISPLAY command
  - D M 4-4, 4-9
  - D M examples
    - 3084 in physically partitioned mode 4-24
    - 3084 in single-image mode 4-31
    - 3090 Model 400 in physically partitioned mode 4-41
    - 3090 Model 400 in single-image mode 4-46
- D U
  - considerations 4-4
- DSE (see data server element)
- DSF (device support facilities) 3-18
- dual write function (IOCP) 2-12

## E

- enable/disable switch, 3380 3-17
- enabled wait states 3-21
- EREP system exception report (see DASD maintenance and recovery)
- errors
  - channel subsystem 3-12
  - CPU 3-4
  - I/O (see also channel subsystem errors) 3-14
  - storage 3-9
  - storage element failure 3-12
- examples, 3084
  - D M displays
    - physically partitioned mode 4-24
    - single-image mode 4-31
  - partitioning and merging 4-15
- examples, 3090
  - partitioning and merging 4-32
- examples, 3090 Model 400
  - D M displays
    - physically partitioned mode 4-41
    - single-image mode 4-46
- excessive spin loop
  - ACR considerations 3-23
  - determining the cause 3-27
  - operator notification of spin loop 3-22
  - spin loops, introduction to 3-22
- excessive spin loop recovery

- LOGREC records, analysis of 3-27
- procedure to restart from message IEE331A 3-26
- recovery action, additional 3-26
- recovery actions for each message insert or wait state 3-24
- recovery for X'09x' wait state 3-24, 3-25
- recovery procedure example for spin loop message 3-24

- EXDC 2-2
- extended storage
  - reconfiguration (3090)
    - introduction 4-9
- external data controller 2-2

## F

- FEATURE = ALTCTRL 2-13
- FORCE operand, CONFIG CHP command 4-14

## G

- glossary of terms 1-1
- guidelines
  - master console configuration 2-12
  - 3084 I/O configuration in single-image mode 2-10
  - 3090 I/O configuration in single-image mode 2-10

## H

- hard CPU errors 3-5
- hard storage error
  - MODE command used to increase PD threshold 3-11
- hard storage errors 3-10
- hardware
  - errors, CPU 3-4
  - instruction tracing 3-20
  - recovery actions 3-2
  - recovery, introduction to 1-3
  - units, deselecting before system is IPLed 4-1
- hot I/O, condition and sample parms 3-16

## I

- I/O
  - configuration guidelines, 3084 or 3090 Model 400 in single-image mode 2-10
  - device reconfiguration 4-15
  - errors (see also channel subsystem errors) 3-14



- hot 3-16
- missing interrupts 3-15
- IEASYSxx parmlib member, RSU parameter 2-13, 2-14
- IECIOSxx parmlib member 3-15
- IEE331A (spin loop message)
  - procedure to restart from 3-26
  - recovery actions for each message insert 3-24
- instruction tracing 3-20
- interrupts
  - I/O, missing 3-15
- IOCP considerations 2-12
- IODEVICE statement 2-13

## L

- last path to a device, configuring offline 4-14
- local TP device configuration 2-9
- LOGREC records, used to determine cause of excessive spin loop 3-27
- long-term fixed pages 2-13, 4-5
- loop tracing 3-20
- loops, spin (see also excessive spin loop recovery) 3-22

## M

- machine checks
  - CPU errors, hard 3-5
  - CPU errors, soft 3-5
  - CPU errors, terminating (see also alternate CPU recovery) 3-6
  - flow through the operating system 3-3
  - information provided with machine checks 3-4
  - Vector Facility errors 3-7
- maintenance and recovery, DASD 3-18
- manual recovery (see operator recovery actions)
- master console
  - configuration guidelines 2-12
  - failure 3-15
- measurements, channel
  - in physical partitioning or merging 4-11
- merging
  - sequence, recommended 4-3
- merging examples
  - 3084 4-15, 4-25
  - 3090 Model 400 4-32, 4-42
- MIH (see missing interrupts)
- missing interrupts 3-15
- MODE command
  - used to increase PD threshold for storage errors 3-11
  - used to increase the SR threshold for storage errors 3-10
- monitoring facility recovery 3-14

- multiprocessors, terminating errors (see also ACR) 3-9

## N

- non-preferred storage, relation to RSU parm 2-13
- nonpartitionable processor complex, configuring devices 2-5
- notification, operator (in excessive spin recovery) 3-22

## O

- offline real storage, configuring (effect of RSU parm) 2-14
- operating system recovery actions 3-2
- operational priorities 1-1
- operator intervention (see also operator recovery actions) 1-3
- operator notification, in excessive spin recovery 3-22
- operator recovery actions 3-19
- out-of-sync condition, 3380 array
  - recovery from 3-18

## P

- pages, long-term fixed 4-5
- partitioning
  - as an operational convenience 1-3
  - definition of 1-3
  - example 4-15
  - sequence, preferred 4-3
  - 3084 example 4-17
  - 3090 Model 400 example 4-33
- path, last to a device, configuring 4-14
- planning, pre-installation 2-1
- PPT considerations 4-5
- pre-installation planning 2-1
- preferred storage, relation to RSU parm 2-13
- processor reconfiguration 4-10
- program properties table considerations 4-5

## R

- real storage
  - procedure to use if a storage element does not come online 4-43
- reason code 0 during CPU restart 3-19
- reason code 1 during CPU restart 3-19
- reconfigurable storage, relation to RSU parm 2-13
- reconfiguration 4-1
  - channel path 4-14

- considerations 4-2
- degrees of support 4-2
- examples of partitioning and merging a 3084 4-15
- examples of partitioning and merging a 3090 4-32
- extended storage (3090) 4-9
- I/O device 4-15
- introduction to 1-3
- logical and physical 4-2
- merging example for 3084 4-25
- merging example, 3090 Model 400 4-42
- partitioning
  - definition of 1-3
  - example 4-17
- pre-installation planning 2-1
- processor 4-10
- single-image mode to physically partitioned 4-17
- storage 4-6
- Vector Facility 4-10
- 3090 partitioning example 4-33
- recovery 3-1
  - actions for excessive spin loops 3-24
  - actions, additional, for excessive spin loops 3-26
  - channel path 3-12
  - considerations and procedures 3-1
  - CPU restart (see also alternate CPU recovery) 3-19
  - excessive spin loop
    - determining the cause 3-27
  - for '09x' wait state 3-25
  - hardware actions 3-2
  - IEE331A, procedure to restart from 3-26
  - introduction to 1-3
  - monitoring facility 3-14
  - operating system actions 3-2
  - operator actions 3-19
  - operator intervention (see also operator recovery actions) 1-3
  - restart function to recover from excessive spin loop 3-26
  - spin loop (see also excessive spin loop) 3-22
  - spin loop recovery procedure, example 3-24
  - subchannel 3-14
  - Vector Facility failure 3-7
  - X'09x' wait state 3-24
- recovery and reconfiguration, introduction to 1-3
- restart reason 0 3-19
- restart reason 1 3-19
- restarting a CPU (see also alternate CPU recovery) 3-19
- RSU parameter, in IEASYSxx parmlib member 2-13

**S**

- service processor damage/stall 3-9
- SHAREDUP (FEATURE=SHAREDUP in IODEVICE statement) 2-13
- short-term fixed pages, PPT considerations 4-5
- side terminology, 3084 4-16
- single image, 3084

- I/O configuration guidelines 2-10
- single image, 3084 or 3090 Model 400
  - DASD configuration 2-11
- soft CPU errors 3-5
- soft storage errors 3-10
  - MODE command used to increase SR threshold 3-10
- spin loop (see also excessive spin loop) 3-22
- spin loop recovery procedure, example 3-24
- storage
  - configuring offline (effect of RSU parm) 2-13
  - differences between 308x and 3090 4-8
  - errors, effects of 3-9
  - errors, hard 3-10
  - errors, soft 3-10
  - extended 4-9
  - inability to come online when CF
    - STOR(E=x), ONLINE is issued 4-43
  - increments 4-6
  - layout, 3084
    - single image mode 4-18
    - storage element 1 offline 4-21
    - storage element 3 configured online 4-27
    - storage elements 1 and 3 configured online 4-28
    - storage elements 1 and 3 offline 4-22, 4-25
  - layout, 3090 Model 400
    - single image mode 4-34
    - storage element 2 offline 4-37
    - storage element 3 offline 4-39
  - non-preferred 2-13
  - physical view of (pictorial) 4-7
  - reconfiguration 4-6
  - subincrements 4-6
- storage element
  - composition 4-7
  - failure 3-12
  - procedure to use if its storage does not come online 4-43
  - reconfiguration 4-6
- string switching 2-6
- subchannel recovery 3-14
- switches, two-channel, need for 2-1
- system recovery, introduction to 1-3
- SYS1.PARMLIB
  - IEASYSxx member (RSU parm) 2-14
  - IECIOSxx member 3-15

**T**

- tape configuration 2-7
- terminating CPU errors (see also ACR)
  - general 3-6
  - multiprocessors 3-9
- terminology and definitions 1-1
- time intervals (MIH)
  - changing the IBM-default values 3-15

TP device (local) configuration 2-9  
tracing, instruction 3-20  
two-channel switches, need for 2-1

## U

uncoded wait states 3-21  
UNCOND operand, CONFIG CHP command 4-14  
unconditional reserve (see alternate path recovery)  
unit record configuration 2-9

## V

Vector Facility  
continuing a job whose Vector Facility is  
offline 3-20  
preventing timeout of swapped out vector  
jobs 3-20  
reconfiguration 4-10  
recovery 3-7  
taking offline the last one in the system 4-11

## W

wait states, recovery from 3-21

## X

X'09x' wait state, recovery for 3-22, 3-25  
XA configuration 2-13  
XA configuration program 2-13

## Numerics

2914 or 3814 switching system, need for 2-1  
308x  
channel subsystem considerations 2-2  
3084  
configuration switch 4-16

differences from the 3090 Model 400 4-32  
differences from the 3090 Model 400E/600E 4-32  
examples

D M displays for physically partitioned  
mode 4-24

D M displays for single-image mode 4-31  
examples of partitioning and merging 4-15  
I/O configuration guidelines 2-10  
merging example 4-25  
side terminology 4-16  
single image mode 2-11  
single-image DASD configuration 2-11  
storage layout

single image mode 4-18

storage element 1 offline 4-21

storage element 3 configured online 4-27

storage elements 1 and 3 configured  
online 4-28

storage elements 1 and 3 offline 4-22, 4-25

3090

channel subsystem considerations 2-2

3090 Model 400

differences from the 3084 4-32

examples

D M displays for physically partitioned  
mode 4-41

D M displays for single-image mode 4-46  
of partitioning and merging 4-32

extended storage reconfiguration

introduction 4-9

I/O configuration guidelines 2-10

merging example 4-42

real storage layout

single image mode 4-34

single image mode DASD configuration 2-11

storage layout

storage element 2 offline 4-37

storage element 3 offline 4-39

3090 Model 400/600E

differences from the 3084 4-32

3090 Vector Facility failure 3-7

3380

considerations 3-17

enable/disable switch 3-17

out-of-synch condition, recovery 3-18

3725 configuration 2-8

3814 switching system, need for 2-9

4381

channel subsystem configuration 2-4

channel subsystem considerations 2-4



GC28-1160-4

This manual is part of a library that serves as a reference source for systems analysts, programmers, and operators of IBM systems. 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.

**Note:** *Copies of IBM publications are not stocked at the location to which this form is addressed. Please direct any requests for copies of publications, or for assistance in using your IBM system, to your IBM representative or to the IBM branch office serving your locality.*

Possible topics for comment are:

Clarity    Accuracy    Completeness    Organization    Coding    Retrieval    Legibility

If you wish a reply, give your name, company, mailing address, and date:

---

---

---

---

What is your occupation? \_\_\_\_\_

How do you use this publication? \_\_\_\_\_

Number of latest Newsletter associated with this publication: \_\_\_\_\_

Thank you for your cooperation. No postage stamp necessary if mailed in the U.S.A. (Elsewhere, an IBM office or representative will be happy to forward your comments or you may mail directly to the address in the Edition Notice on the back of the title page.)

Reader's Comment Form

Cut or Fold Along Line

Fold and tape


Please Do Not Staple

Fold and tape



NO POSTAGE  
NECESSARY  
IF MAILED  
IN THE  
UNITED STATES

**BUSINESS REPLY MAIL**  
FIRST CLASS PERMIT NO. 40 ARMONK, N.Y.



POSTAGE WILL BE PAID BY ADDRESSEE

International Business Machines Corporation  
Department D58, Building 921-2  
PO Box 390  
Poughkeepsie, New York 12602



Fold and tape


Please Do Not Staple

Fold and tape



Printed in U.S.A.

GC28-1160-04







Printed in U.S.A.