IBM

Washington
Systems
Center

Technical
Bulletin

An Introduction To The IBM 3084 Processor Complex Structure And Capabilities

By: R. M. Schardt

ZZ05–0402–00
December 1982

Washington Systems Center
Gaithersburg, Maryland
Technical Bulletin

An Introduction To The IBM 3084 Processor Complex Structure And Capabilities

R. M. Schardt

This Technical Bulletin is for IBM Internal Use Only. The intent is to provide insight into the capability and performance of IBM products for IBM marketing representatives and systems engineers.

This document has not been subjected to any formal review and has not been checked for technical accuracy. Contents must be evaluated for applicability to a particular installation. You may discuss pertinent information from this Technical Bulletin with a customer, and you can abstract pertinent information for presentation to your customers. Any code included, however, is for internal information purposes only and may not be given to customers. If included code is identified as "incidental programming" its use must conform to the guidelines in GI Section 20.2 and the Systems Engineering Services Section of the DP Sales Manual. IBM personnel can make reproductions of this information for their own use. Responsibility for maintaining appropriate security accompanies such reproductions.

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

References in this publication to IBM products, programs or services do not imply that IBM intends to make these available outside the United States.

DAPS Code 0894
ZZ05-0402-00
December 1982

IBM Internal Use Only

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

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

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

PREFACE

This Technical Bulletin contains a foil presentation and background information that highlights the significant features of the 3084 Processor. It is intended for use by IBM Marketing Representatives and System Engineers.

The emphasis in this document is on the 3084 operating in single image mode under MVS/XA. When the 3084 is operating in partitioned mode, it is running as two independent 3081K systems. Therefore, information about the 3084 operating in partitioned mode can be found in the 3081 Design Presentation, Washington Systems Center Technical Bulletin ZZ05-0381. If a combined 3084 Design/Capabilities presentation is required, the foils from the 3081 Design Presentation can be merged with the foils in this document. It is assumed that the reader of this document is thoroughly familiar with 3081 Design Presentation.

The text contained in this document is classified 'IBM Internal Use Only' and as such cannot be distributed to non-IBM employees. The foils are unclassified and copyrighted. If you wish to provide handouts to your audience, the handout masters in the appendix may be reproduced and distributed. However, no permission is granted to non-IBM employees to reproduce and/or distribute this material.

Your comments and suggestions regarding the content and usefulness of this guide are solicited.

THIS PAGE INTENTIONALLY LEFT BLANK

CONTENTS

3084 Presentation Guide **1**
Introduction 1
Presentation Topics 2
Overview 4
3084 Operating Modes 11
3084 Channel Subsystem 21
3084 Consoles 42
3082 Model Q Processor Controller 48
3082 Q Recovery Logic 55
Summary 58

Appendix A. 3084 Power Boundaries **59**

Appendix B. Reference Documentation **61**
Announcement Material 61
SRLs 61
Washington Systems Center Technical Bulletins 62
Customer Documents 62
Internal Use Only Documents 62
Other 62

Appendix C. Presentation Foils **63**

Appendix D. Foil Handout Masters **110**

THIS PAGE INTENTIONALLY LEFT BLANK

3084 PRESENTATION GUIDE

INTRODUCTION

What follows in not intended to be a conventional presentation script. Instead it is designed to provide the presenter with some guidance relative to the significant points that should be made during the presentation. In addition, background information is included where appropriate.

This presentation assumes that the audience has seen the 3081 Design Presentation, ZZ05-0381.

The use of two foil projectors is recommended for this presentation. Two adjacent foils shown on the same page indicates that both should be displayed simultaneously.

PRESENTATION TOPICS

Foil 0.0 - Presentation Title

Foil 0.1 - Presentation Topics

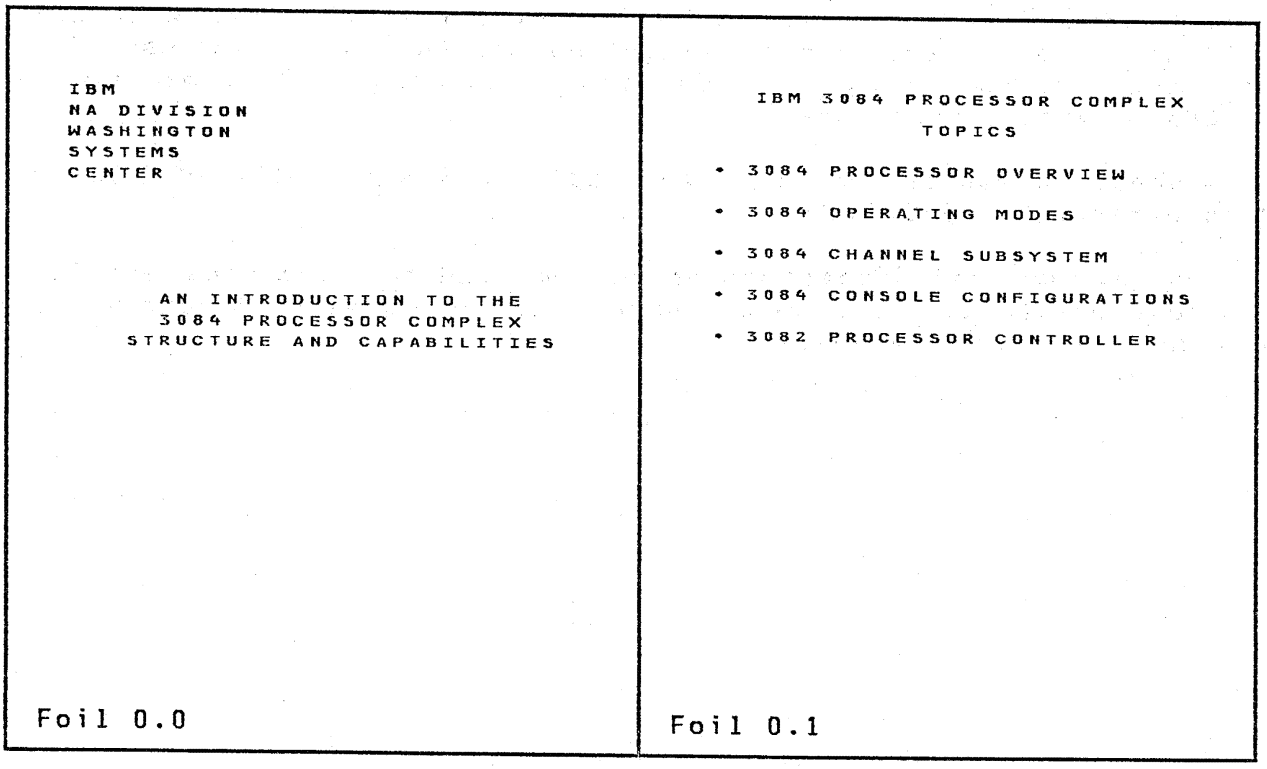

This presentation is designed for use by IBM Marketing Representatives and System Engineers. Its intent is to provide additional in-sight into the structure and capabilities of the 3084 Processor. The topics covered in this presentation include:

3084 Processor Overview - A review of the announcement highlights.

3084 Operating Modes - A discussion of the 3084 operating in Single Image Mode and Partitioned Mode.

3084 Channel Subsystem - A description of the logical and physical structure of the channel subsystem when operating under MVS/XA. Included is an introduction to channel configuration considerations.

3084 Console Configurations - A description of the structure and functions of the 3084 Consoles. Included are some considerations for configuring the 3084 consoles.

3082 Model Q Processor Controller - A discussion of how the 3082 Q operates in Single Image Mode. Included is a description of the 3082 Q structure and error recovery.

OVERVIEW

This series of foils present a system overview. It is a review of the announcement material.

Foil 1.1 - 3084 Overview

Foil 1.2 - 3084 Structure

IBM 3084 PROCESSOR COMPLEX

3084 OVERVIEW

- FOUR INTEGRATED CENTRAL PROCESSORS
  - 26NS CYCLE
- 32, 48, OR 64 MEGABYTES OF REAL STORAGE
- 48 CHANNELS
  - DATA STREAMING STANDARD ON ALL BLOCK CHANNELS
  - MAXIMUM 8 BYTE CHANNELS
- OPERATING MODES
  - SINGLE IMAGE
    - MVS/XA ONLY
  - PARTITIONED
    - MVS/XA, MVS/SP OR VM/SP

Foil 1.1

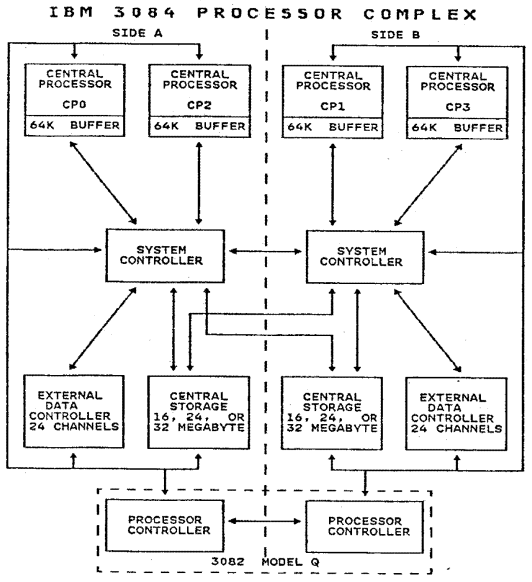

Foil 1.2

The 3084 is a tightly-coupled processor. It has four integrated central processors. The 3084 is organized into two symmetric sides; Side A and Side B. Each side is essentially a 3081 K Processor. Therefore, the 3084 can be viewed as being the result of merging two 3081 K systems.

Storage is available in 32, 48 or 64 megabytes. 3084 storage must be symmetric. That is, a 32 megabyte system must have 16 megabytes installed on each side. A 48 megabyte system must have 24 megabytes on each side. And a 64 megabyte system must have 32 megabytes on each side.

There is no channel option on the 3084. The system comes with 48 channels, 24 on each side.

The 3084 can operate in either of two modes, single image mode and partitioned mode[1]. In single image mode, the entire 3084 (both sides) operates under a single copy of MVS/XA. MVS/XA is the only System Control Program (SCP) that supports the 3084 in single image mode.

In partitioned mode, the 3084 is split into two, independent systems. Each side operates as a stand alone 3081 K processor. In partitioned mode, each side is supported by the same software that supports the 3081 K. This includes MVS/XA, MVS/SP, and VM/SP.

---

[1] Single image mode and partitioned mode are discussed in more detail in the 3084 Operating Modes section of this document.

Foil 1.3 - 3084 Overview - Upgradeability

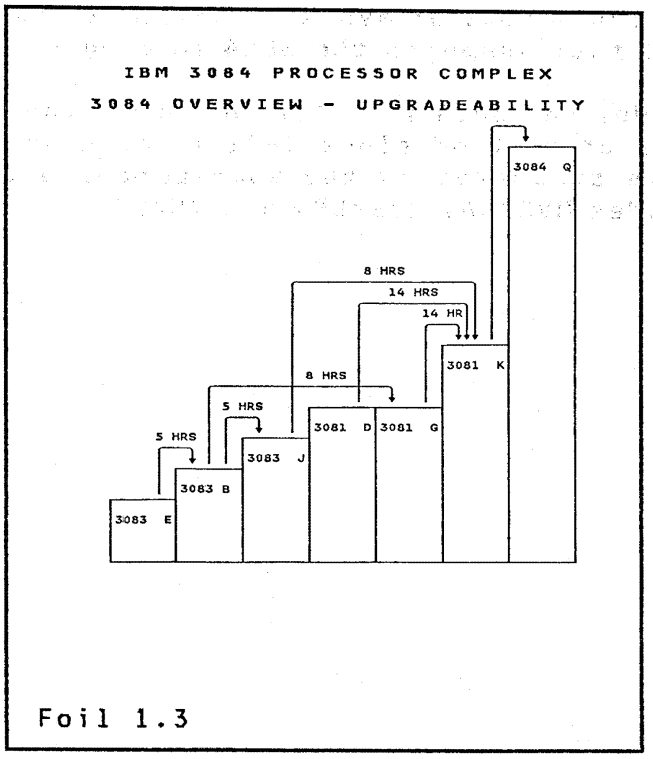

This foil illustrates the upgradeability paths for the 308X Processors. As is shown, the 3084 can be field upgraded from an installed 3081 K. In fact, the only way an installation can get a 3084 is to order an upgrade to a 3081 K.

The times shown for the various processor upgrades do not include the time required to upgrade storage. If the storage is being increased at the same time the processor is being upgraded, additional time will be required. Because of the various 3081 K configurations that can be upgraded to the 3084, upgrade times to the 3084 will be determined on a customer by customer basis. For example, it will take longer to upgrade a 16 megabyte, 16 channel 3081 K then it will take to upgrade a 32 megabyte, 24 channel 3081 K.

Foil 1.4 - Overview

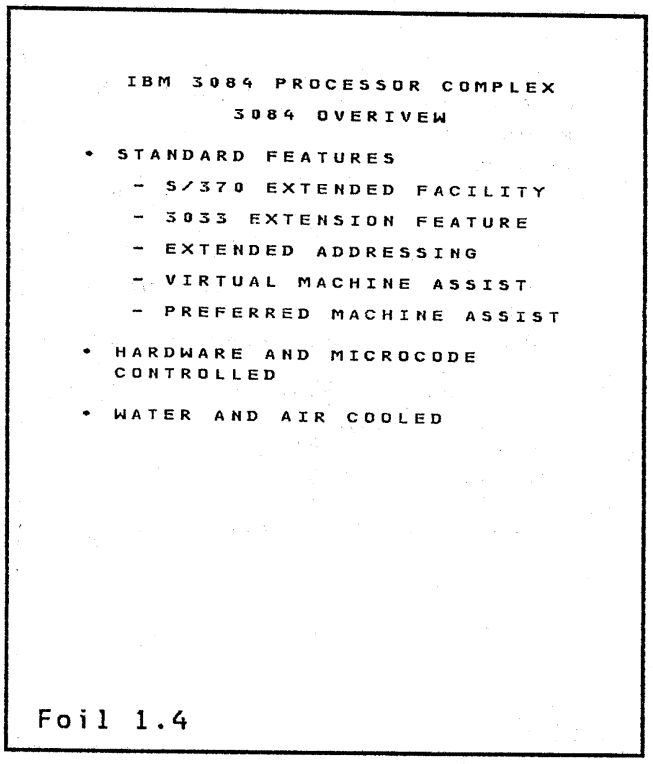

Key points:

1. The System/370 Extended Facility, the 3033 Extension Feature, Extended Addressing, Virtual Machine Assist, and Preferred Machine Assist are standard.

2. The processor is both hardware and microcode controlled. Microcode is used extensively on the 3084. The real advantage of microcode lies in the ease with which it can be modified. This has major implications for the initial cost and for the cost and time required to add features and engineering changes to the system. So in the 3084, where ease of modification is desirable, those functions are likely to be microcoded. Other functions that are not likely to be modified are implemented in hardware. Therefore, for reasons of flexibility and cost, microcode is used to a greater degree in the 3084 Processor.

3. The 3084 Processor employs both water and air cooling. The Thermal Conduction Modules (TCMs)[2] and certain power regulators are water cooled. The remaining components are air cooled.

---

[2] TCMs are described in subsequent foils.

Foil 1.5 - 3084 Overview - Technology

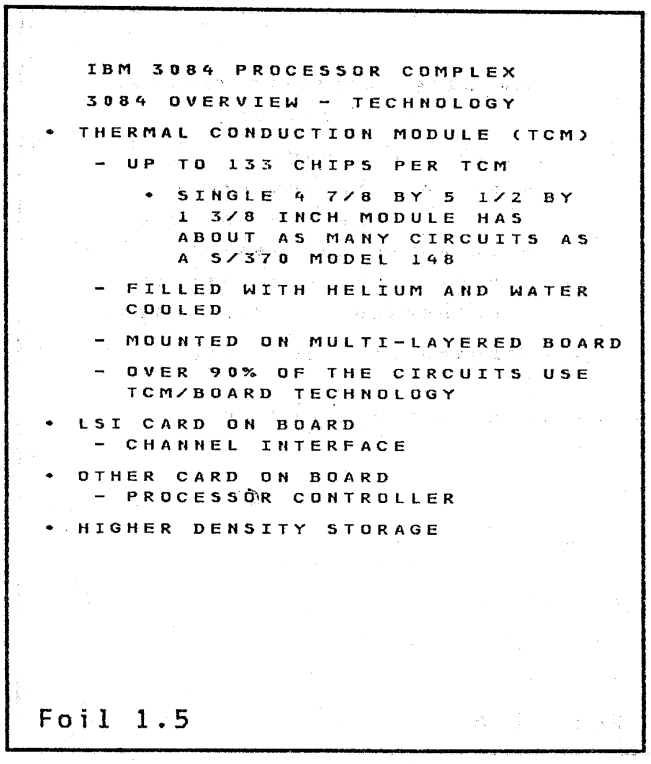

The 3084 Processing Complex uses a mixture of technologies. These foils provide an overview of these technologies.

Key points:

1. TCM Technology. The circuits in the 3084 Processor employ a LSI chip, similar to that used in the 4300. The chips are packaged in modules called Thermal Conduction Modules (TCM). Each TCM measures 4 7/8 by 5 1/2 by 1 3/8 inches. These modules can contain up to 133 chips. Each chip can contain up to 704 individual circuits. If fully utilized, the maximum number of circuits per module is about 83K[3]. The maximum number of circuits used within the 3084 TCM is about 45K. This is equivalent to about the number of circuits in a IBM System 370 Model 148.

   Each TCM is sealed, filled with Helium and cooled by water. Helium was selected because it is about six times more efficient in its ability to conduct heat than is air. This helps improve the efficiency of the TCM cooling process. In addition, Helium reduces the oxidation process thus helping to prolong the life of the internal parts.

---

[3] The average number of circuits actually used within a TCM is about 30K.

There are fifty-four TCMs in 3084, twenty-seven on each side. They are mounted on multilayer (18 layers) boards, four boards per side. These boards provide cableless interconnections between the TCM circuits. A maximum of 9 TCMs are mounted on a single board.

This combination of LSI technology and TCM/board packaging provides substantially improved reliability along with reduced space, power and cooling requirements. Over 90% of the 3084 Processor circuits use the TCM/board technology.

2. Card On Board Technology. Various 'card on board' (COB) technologies are also used in the 3084 Processor. These include:

   a. Higher density storage chip. The 3084 storage chips are 16K bit Dynamic RAM technology. This compares to 4K bit chips used in the 3033.

   b. LSI COB. These are LSI chips packaged 'card on board'. This is the technology used in the 4300 Processor. They are used in the channel interface circuits. These cards are located in the Processor Controller.

   c. Other COB. There is a broad range of other COB technologies used throughout the Processor Controller.

   One should keep in mind that system availability is a function of more than just the technology employed. Additional factors include other hardware components (i.e., memory, cables, power supplies), microcode, software, I/O subsystems, and operator error recovery procedures. Therefore one should not assume that an increase in technology reliability will provide a proportional increase in system availability.

Foil 1.6 - 3084 Overview - Floor Plan

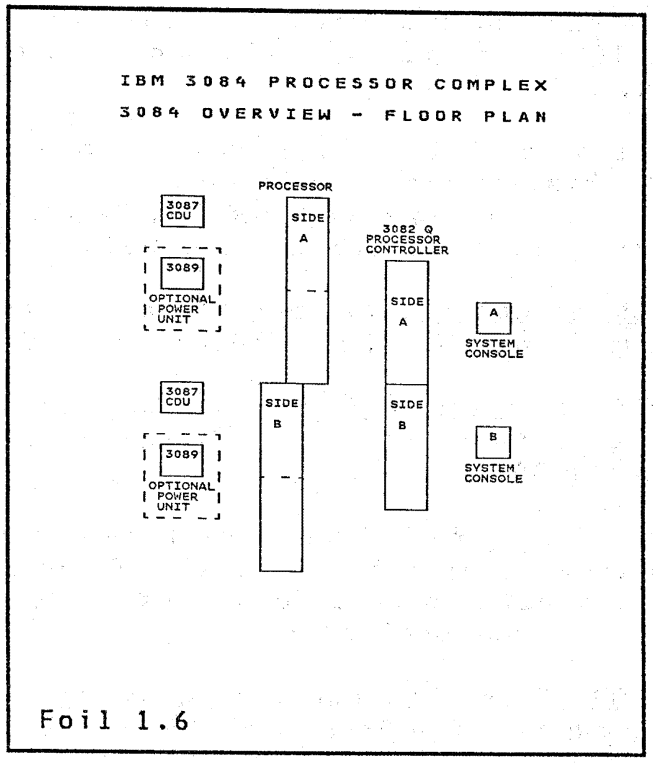

This foil shows the recommended 3084 floor plan. In single image mode, all elements, except one side of the 3082 Q, are active. This includes both system consoles. The inactive side of the 3082 Q, generally the B Side, is in standby mode. Should any components of the active side fail[4], the inactive component can be called into operation. In partitioned mode, each side runs as an independent system. Each side has its own power supplies so that one side can be powered down without effecting the other side[5].

---

4 Details of the 3082 Q operation, including its enhanced recovery, are contained in the Processor Controller section of this document.

5 The Appendix contains a layout of the power boundaries for each side of a 3084.

3084 OPERATING MODES

Foil 2.1 - 3084 Operating Modes

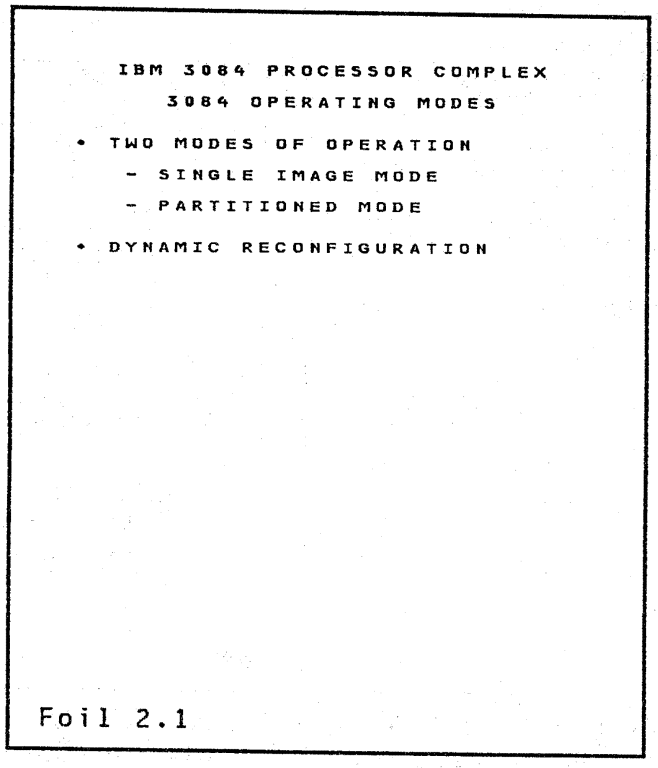

The 3084 can operate in either Single Image Mode or Partitioned Mode. This section of the presentation addresses the characteristics and capabilities of each mode. This includes a discussion of the structure of the system in each mode, along with some availability and recovery considerations. Also included is a discussion of how modes can be changed dynamically under MVS/XA without requiring an IPL.

Foil 2.2 - 3084 Single Image Mode (Description)

Foil 2.3 - 3084 Single Image Mode (Structure)

IBM 3084 PROCESSOR COMPLEX

3084 SINGLE IMAGE MODE

- ENTIRE 3084 PROCESSOR COMPLEX OPERATING UNDER A SINGLE COPY OF MVS/XA
  - DYNAMIC CHANNEL SUBSYSTEM
    - ALL PROCESSORS HAVE ACCESS TO ALL DEVICES
  - EACH SIDE HAS DIRECT ACCESS TO ALL OF CENTRAL STORAGE
  - SYSTEM AVAILABILITY
  - BOTH SIDES MUST BE AT SAME EC LEVEL

Foil 2.2

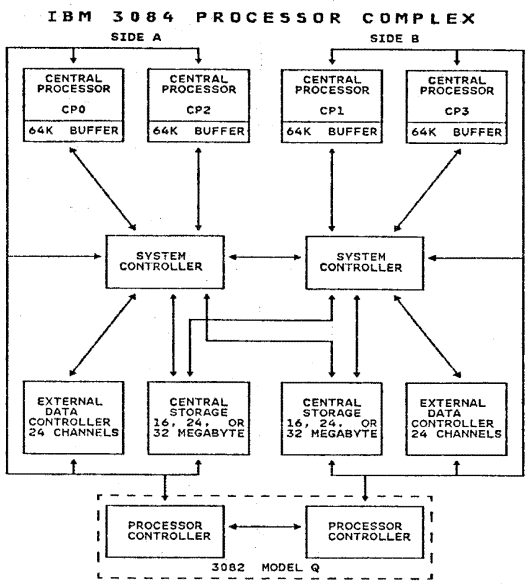

Foil 2.3

The 3084 is organized into two sides; Side A and Side B. Each side is essentially a 3081 K processor. Each side is identical with respect to central processors, storage, and channels. For storage, a 64 megabyte system must have 32 megabytes on each side. A 48 megabyte system must have 24 megabytes on each side, and a 32 megabyte system must have 16 megabytes on each side. The 48 3084 channels are symmetric by definition, 24 on each side. The central processors are also symmetric by definition; CP0 and CP2 on Side A, CP1 and CP3 on Side B. The only exception to the symmetry rule are the Channel To Channel Adapters. A different number of CTCs (up to two) can be attached to each side.

The 3084 can be operated in two modes; single image mode and partitioned mode. In single image mode, both sides (all four central processors) operate under control of a single copy of MVS/XA[6]. In partitioned mode, each side operates as an independent system with the processing power that is approximately equivalent to a 3081 K.

---

6 MVS/XA is the only System Control Program (SCP) support for the 3084 operating in single image mode.

Single Image Mode

In single image mode, a single copy of MVS/XA manages the entire 3084 Processor Complex. There is a single Dynamic Channel Subsystem[7]. That is, all four central processors have access to every I/O device, regardless of the EXDC to which the device is physically attached.

On the surface, it would appear that the System Controller has been redesigned for the 3084 so that it could interact with another System Controller. Actually, the original design of the System Controller had the capability to communicate with another System Controller. The 3084 is the first system that uses that capability.

Each System Controller has direct access to all of central storage. This means that if a failure occurs that disables one side, the other side still has access to all of central storage and, therefore, can continue operation.

Single Image Mode Recovery Overview

The reliability of the 308X circuit technology has proven to be extremely high. Therefore, the potential for hard circuit failures is relatively low. However, there are always a series of 'what if' questions that arise about component failures in new systems. The following paragraphs address these potential questions.

When operating in single image mode, individual elements within the 3084 can fail[8], and the system can continue operation. For example:

Central Processor Failure

The central processor will be placed in the checkstopped state, MVS ACR will be invoked, and the system will continue with the remaining processors.

EXDC Failure

If the Channel Processing Element (CPE) fails, the entire EXDC is inoperable and will be removed from the system. Assuming the system was configured properly, operation can continue through the remaining EXDC[7]. If either a Data Server Element or Interface Adapter Element fails, only the failing element will be removed from

---

7 Refer to the Channel Subsystem section of this document for additional information.

8 For the purpose of this discussion, the term 'failure' refers to a solid failure or an intermittent failure that is occurring frequently enough to exceed its failure rate threshold. Generally, this threshold is exceeded when the 3082 determines more time is being spent in recovery than productive work.

the system. The rest of the EXDC can continue operation.

System Controller Failure

A failure of either System Controller is catastrophic. An IML/IPL will be required. To recover, the side that does not contain the failing System Controller can be IMLed and IPLed in partitioned mode and operation can continue.

Storage Failure

Recovery is the same as in other 308X systems. Individual 4K frames in which a failure occurs, are made unavailable to the system. A storage failure in either the Hardware System Area or in the MVS nucleus, is a catastrophic failure and will require an IML.

3082 Model Q Failure

In single image mode, one side of the 3082 is operational and the other side is in stand-by mode[9]. Normally, Side A will be active and Side B will be in stand-by Mode. If an element on the active fails, the corresponding element on the inactive side can be made operational and system operation can continue.[9].

3087 CDU Failure

The 3087 Coolant Distribution Unit (CDU) contains two pumps. One pump is active while the other is in standby mode. If the active pump fails, the standby pump is activated.

The 3084 has two 3087 CDUs. Each CDU supports only the side to which it is attached. When running in single image mode, if both pumps in one of the 3087 CDUs fail, the entire system will come down. In this case, the side with the good 3087 could be re-IPLed in partitioned mode. The other side would be inoperative until the 3087 was repaired. When running in partitioned mode and both pumps in a single 3087 CDU fail, only the side with the

---

[9] Details of 3082 recovery are covered in the Processor Controller section of this document.

failing CDU will come down. The other side can continue running.

3089 PDU Failure

The 3089 Power Distribution Unit (PDU) is optional with 3084, as it is with all 308X processors. Each side of the 3084 is powered separately. Therefore, each side requires a separate power source. If the 3089 option is selected, then two are required; one to power each side.

When running in single image mode, if the power source to one side is lost, the entire system will come down. For maximum availability, consideration should be given to a power source configuration such that multiple sources are available to each side of the 3084. With this type of arrangement, a single source could fail but the system could continue operating.

When operating in partitioned mode, if the power source to one side fails, only that side will come down. The other side can continue operating.

To operate in single image mode, both sides must be at the same EC level. During the IML process, a check is made to ensure the EC level is the same on both sides. If the EC levels differ, the IML is terminated and the system cannot be IMLed in single image mode until the EC levels are brought to the same level. If the EC levels differ, the 3084 can operate in partitioned mode.

Foil 2.4 - 3084 Partitioned Mode (Description)

Foil 2.5 - 3084 Partitioned Mode (Structure)

IBM 3084 PROCESSOR COMPLEX

3084 PARTITIONED MODE

- 3084 SPLIT INTO TWO INDEPENDENT SYSTEMS
  - EACH SIDE CAN RUN EITHER MVS/XA, MVS/SP, OR VM/SP
  - MUST BE SYMMETRIC
    - CENTRAL PROCESSORS
    - CHANNELS
    - STORAGE

Foil 2.4

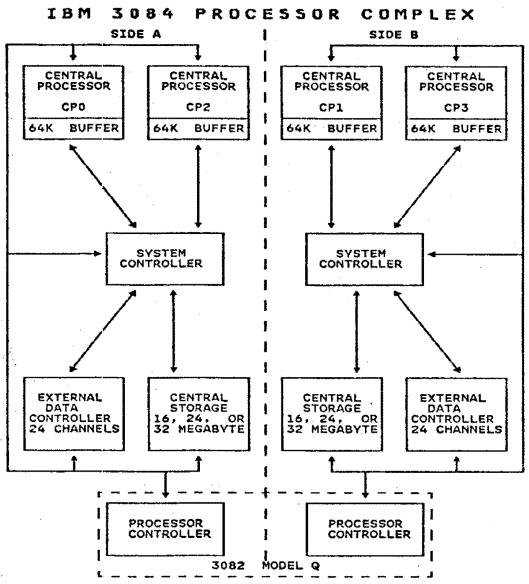

Foil 2.5

Partitioned Mode

In partitioned mode, the 3084 is split into two symmetric, independent systems. This includes separate power boundaries. For example, Side A, including its Processor Controller, could be powered off, while Side B is running[10].

In partitioned mode each side operates as a 3081 K[11]. Therefore, in partitioned mode, each side is supported by the same software that supports the 3081 K. This includes MVS/SP, VM/SP, and MVS/XA.

Each side of the 3084, running in partitioned mode, has the same availability characteristics as a 3081 K. There is no additional exposure over that of a stand alone 3081 K.

There is a small difference in processor performance when comparing the Instruction Execution Rate of each side of a 3084 to a stand alone 3081 K. This difference is due to the fact that the System Controller in the 3081 K has two paths to each Basic Storage Controller[11]. In the

---

10 See the Appendix for a layout of the 3084 power boundaries.

11 Refer to the Washington Systems Center Technical Bulletin, 3081 Design Presentation, ZZ05-0381.

3084, there is only one path to each Basic Storage Controller. The second path in the 3084 is used by the System Controller to access the other side's central storage. When operating in partitioned mode this second path is disabled. The result is that the relatively small amount of overlap in storage access, that is possible in the 3081 K, is not possible in the 3084 operating in partitioned mode.

Foil 2.6 - 3084 Dynamic Reconfiguration

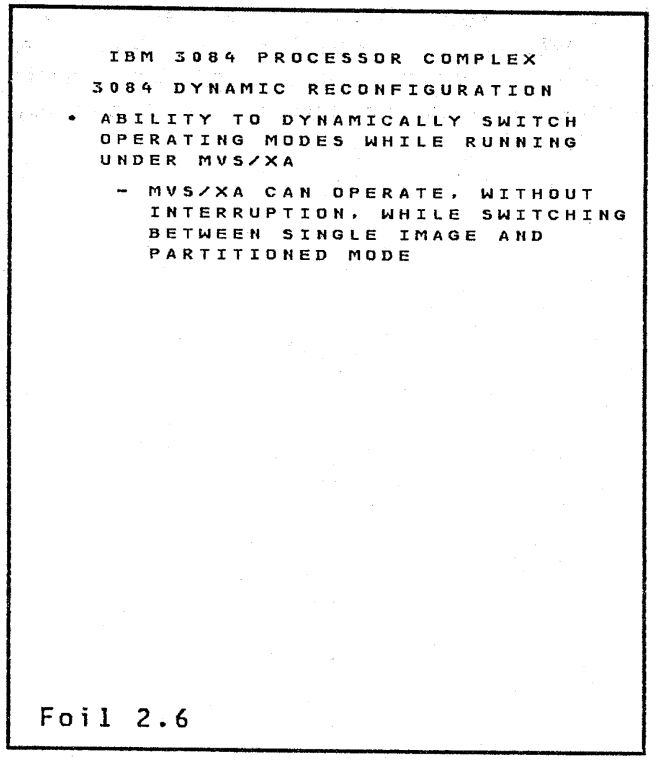

The 3084 is capable of dynamically changing operating modes. This is accomplished by using a series of CONFIGuration commands under MVS/XA[12] and by issuing a VARY PHYSICAL command to the 3082 Processor Controller from the system console. These commands provide the capability of configuring an entire side into or out of a single image mode environment[13].

---

12 The CONFIG Command is new for MVS/XA. It performs functions similar to the MVS VARY Command.

13 Actually, MVS can be reconfigured by a single MVS command. This could be accomplished by defining procedures that contained all the required MVS commands and then simply executing the appropriate procedure.

Foil 2.3 - 3084 Single Image Mode Structure

Foil 2.5 - 3084 Partitioned Mode Structure

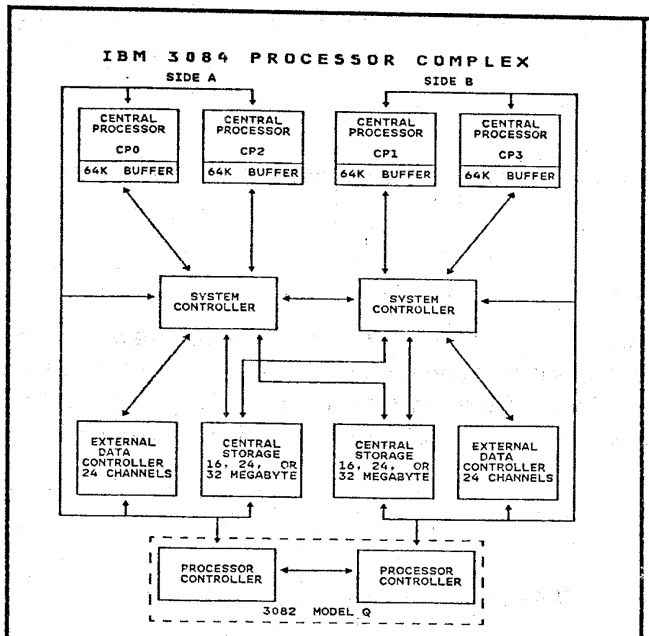

Foil 2.3

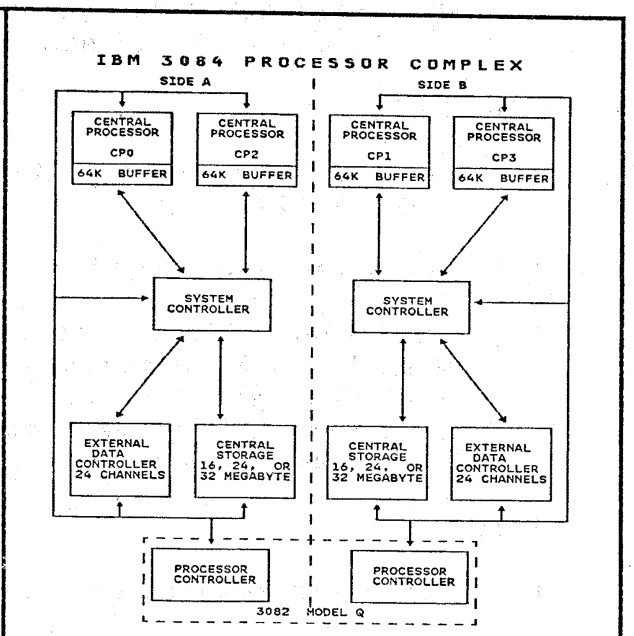

Foil 2.5

Here are a few examples of reconfiguration sequences that can be accomplished.

- Example #1

  1. The 3084 is operating in single image mode under MVS/XA.
  2. Side B is CONFIGed offline. When the CONFIG commands complete, the VARY PHYSICAL SIDE B OFF command is issued from the system console.
  3. MVS/XA continues to run on Side A uninterrupted.
  4. When the VARY PHYSICAL command completes, Side B can be taken for maintenance, powered down for engineering changes, or can be IPLed under MVS/SP, MVS/XA or VM/SP. MVS/XA will continue to operate on Side A uninterrupted.

- Example #2

  1. MVS/XA is running on Side B. VM/SP is running on Side A.
  2. VM/SP is quiesced.

3. Side A is CONFIGed on from MVS/XA. When the CONFIG commands complete, the VARY PHYSICAL SIDE A ON is issued from the system console.

4. When the VARY PHYSICAL command completes, MVS/XA is operating in single image mode.

- Example #3

1. The 3084 is in Partitioned Mode. Each side is operating under MVS/XA.

2. One side is quiesced.

3. The other side issues a CONFIG 'side on' command sequence. When the CONFIG command sequence completes, a VARY PHYSICAL SIDE ON command is issued from the system console. When the VARY PHYSICAL command completes, the 3084 will be operating in single image mode.

- Example #4

1. MVS/SP is running on Side B. VM/SP is running on Side A.

2. Side B (MVS) is quiesced.

3. MVS/XA is IPLed on Side B.

4. Side A (VM) is quiesced.

5. MVS/XA (Side B) initiates a CONFIG side on command sequence to Side A. When the CONFIG commands complete, a VARY PHYSICAL SIDE A ON is issued from the system console.

6. When VARY PHYSICAL command completes, MVS/XA is operating in single image mode.

These examples illustrate the flexibility of the 3084 operating under MVS/XA. The system can move from single image to partitioned mode and back to single image mode without requiring an IPL. During peak load, the system can run in single image mode. When the load is light, the system can be partitioned and another system can be IPLed for testing, etc. When the load increases, the system can again be configured in single image mode. All this can be accomplished while at the same time providing uninterrupted service from MVS/XA.

3084 CHANNEL SUBSYSTEM

From the users point of view, the MVS/XA channel subsystem looks the same regardless of whether it is running on a 3083 E or a 3084 Q. There is a single channel subsystem. The difference is that, unlike the other 308X processors, the 3084 has two EXDCs. The 3084 channel subsystem, therefore, must recognize and manage both EXDCs.

The 3084 channel subsystem can operate in either 370 or 370-XA Mode. 370 Mode is supported only when the 3084 is operating in partitioned mode. Because 370 Mode operation is covered in the 3081 Design Presentation, it is not addressed here.

This section addresses the operation of the 3084 channel subsystem operating in an MVS/XA environment in single image mode. Included are some considerations for channel configuration.

Foil 3.1 - 3084 Functional Elements

Foil 3.2 - 3084 Channel Subsystem Functions

IBM 3084 PROCESSOR COMPLEX

3084 CHANNEL SUBSYSTEM FUNCTIONS

- MANAGES THE TRANSFER OF DATA BETWEEN I/O CONTROL UNITS AND CENTRAL STORAGE
  - 48 CHANNELS (24 PER SIDE)
    - UP TO EIGHT BYTE CHANNELS CAN BE DEFINED (4 PER SIDE)
  - 3.0 MEG BLOCK CHANNELS
    - CAN ATTACH STANDARD AND DATA STREAMING DEVICES
  - 144 MEG AGGREGATE DATA RATE

Foil 3.1

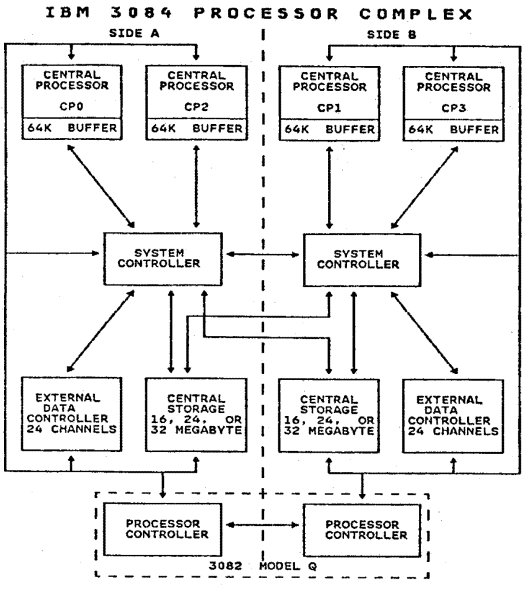

Foil 3.2

The 3084 channel subsystem manages data transfer between the I/O control units and Central Storage. The functions provided include the following:

- Creating data paths between the I/O control units and Central Storage.
- Executing asynchronous portion of I/O instructions.
- Monitoring and controlling data transfer.
- Informing the Central Processor when I/O operations have completed.

Operating under MVS/XA in single image mode, the 3084 channel subsystem consists of two External Data Controllers (EXDC), one on each side. Each EXDC controls 24 data paths for a total of 48. Up to four byte channels can be defined to each EXDC, for a total of eight for the system.

Block channels have a maximum data rate of 3.0 megabytes per second in data streaming mode. Both standard and data streaming devices can be attached to the block channels. Maximum EXDC aggregate data rate is 72 megabytes with 70% utilization on each channel. That is equivalent to all 24 channels servicing data streaming devices at the 3.0 megabyte data rate. The total aggregate data rate for the 3084 operating in single image mode is 144 megabytes, 72 megabytes per EXDC.

Foil 3.3 - 3084 Channel Subsystem - 370-XA Mode

Foil 3.4 - 3084 Channel Subsystem Structure - 370-XA Mode

IBM 3084 PROCESSOR COMPLEX

EXTERNAL DATA CONTROLLER (EXDC)

3084 CHANNEL SUBSYSTEM - 370-XA MODE

- DYNAMIC CHANNEL SUBSYSTEM
  - ALL CENTRAL PROCESSORS CAN ACCESS ALL DEVICES DIRECTLY
    - NO CHANNEL SETS
    - NO ASYMMETRIC DEVICES
  - I/O INTERRUPT CAN BE TAKEN BY ANY CENTRAL PROCESSOR
  - UP TO 8 INTERRUPT SUBCLASSES
  - CHANNEL PATH SELECTION
    - UP TO 4 PATHS
  - 3380 DYNAMIC PATH SELECTION
  - UP TO 4080 UCW'S
  - ALL NEW I/O INSTRUCTIONS
  - CHANNEL MONITORING FACILITY
- TERMINOLOGY
  - SUBCHANNEL/DEVICE
  - UCW
    - NO SHARED UCW'S

Foil 3.3

IBM 3084 PROCESSOR COMPLEX
3084 CHANNEL SUBSYSTEM STRUCTURE
370-XA MODE

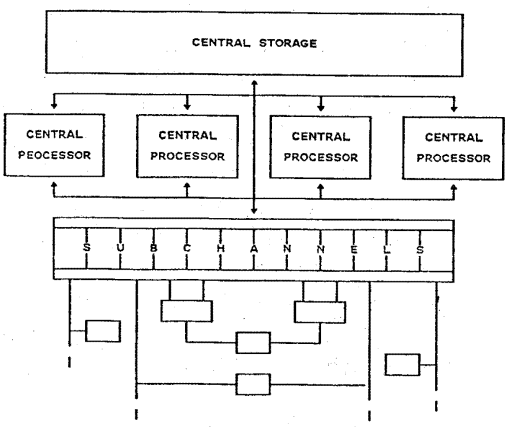

Foil 3.4

The following is a review of the capabilities of the 3084 channel subsystem operating under MVS/XA.

1. The main feature of the MVS/XA channel subsystem is the 'floating channel'. By 'floating channel' we mean that any central processor has direct access to any I/O device. This means that there is no longer any need for channel sets or asymmetric devices in 370-XA Mode.

2. Unlike 370 Mode where the central processor that issues an I/O command is the only processor that can accept the resulting I/O interrupt, in 370-XA Mode, any central processor can accept an I/O interrupt from any I/O operation regardless of whether it started the operation.

   In addition, under program control, a device may be assigned one of up to eight interrupt subclasses. These subclasses determine the priority at which the I/O interrupt from the device will be serviced by the central processor. MVS will support the use of two interrupt subclasses. One subclass will be provided for paging I/O. A second, lower priority subclass will be assigned to all other I/O requests.

3. Selecting which channel will be used to start an I/O operation is done by the channel subsystem. Up to four physical channels

can be specified to a given device. By moving the channel selection function into the channel subsystem, overlap between the channel and central processor is increased and software path lengths are reduced. This new channel selection process is described in subsequent foils.

4. Under MVS/XA, certain models of the 3380 are capable of dynamic path selection and dynamic path reselection. Dynamic path selection is the channel selection function of the channel subsystem working with the 3880 control unit to dynamically determine which path[14] will be used to start the operation. In 370 Mode once a path has been selected, all activity relative to that request must take place on that path. With 3380 dynamic path reselection, I/O requests are no longer confined to the originally selected path. In fact, any of the possible paths to a device may be used. For example, an I/O operation may be started down one path and the data transfer may use an entirely different path. The decision of which path to use is made at the time a path is required. Any of the possible paths that are available can be used.

   A major advantage of dynamic path selection/reselection is that it can simplify the I/O configuration process. Because there are up to four possible paths and these paths are selected at the time they are required, concerns about load balancing and path utilizations are somewhat minimized. For example, in 370 Mode, if the path utilization were 50%, there is a 50% probability that the path will be unavailable when needed. On the other hand, in 370-XA Mode, even if each of the four possible paths had a utilization of 50%, the probability that a path will be unavailable is only 6.25% (50% X 50% X 50% X 50%).

5. Up to 4080 UCWs can be defined. This is the same number provided in 370 mode. However, because in 370 mode a unique UCW is required for each path to a device, multiple UCWs are required for those devices having optional and/or alternate paths. In 370-XA mode only one UCW is required per device, regardless of the number of paths to the device. The MVS support for 370-XA mode provides for the definition of up to 4096 UCBs.

   The 370-XA architecture actually provides the capability for defining up to 65,536 UCWs. However, current 3084 implementation limits this number to 4080.

   Support for up to 256 channel paths is provided by 370-XA architecture. The 3084 has implemented 48 channels.

6. 370-XA Mode includes a complete new set of I/O instructions. In 370-XA Mode, none of the 370 I/O instructions are valid. For

---

[14] The term 'path' refers to channel, control unit, and head of string.

example, the SIO and SIOF instructions are replaced by a Start Subchannel Command (SSCH).

7. The channel has a monitoring facility which can time significant events during an I/O operation. For example, time spent in the channel queue and device service time are accumulated. This information will be collected and reported by RMF. More specific information about the Channel Monitoring Facility is presented later in this section of the presentation.

8. In 370-XA Mode, each device has its own subchannel. Because of this, the terms 'subchannel' and 'device' are used interchangeably. Also each subchannel has its own UCW. Therefore, in 370-XA Mode there are no 'shared' UCWs as there are in 370 Mode.

Foil 3.5 - 370-XA Mode I/O Queuing

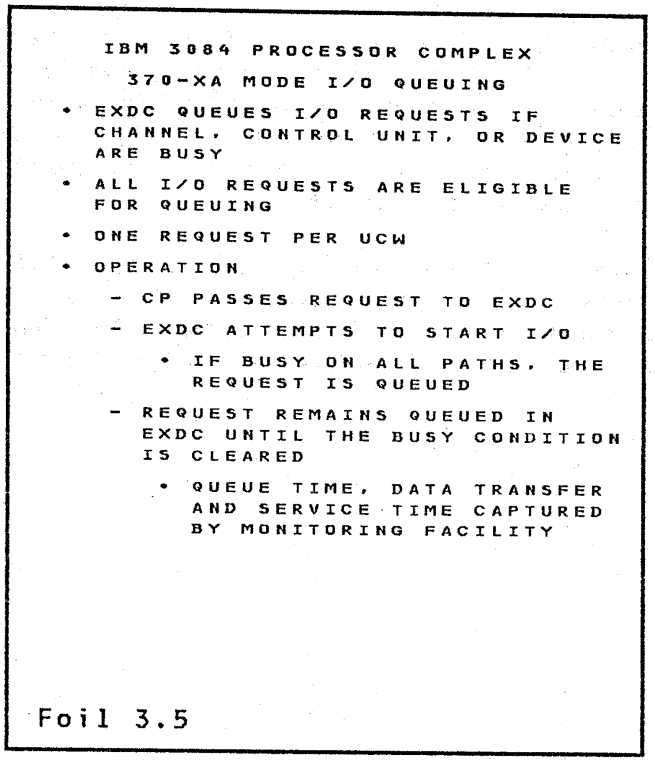

370-XA Mode I/O Queuing

I/O queuing performs the same function in 370-XA Mode as it does in 370 Mode. However, the operation differs in several ways. These differences include:

1. In 370 Mode I/O queuing is operational only for the SIOF instruction. The SIOF instruction is invalid in 370-XA Mode. The Start Subchannel (SSCH) instruction replaces both SIO and SIOF. SSCH is the only instruction used to initiate I/O operations in 370-XA Mode and, therefore, all I/O requests are eligible for I/O queuing in 370-XA Mode.

   As in 370 Mode, one request per UCW can be queued in the channel subsystem. That is all that is necessary because, as is the case in 370 Mode, the operating system will not attempt to start more than one request to a given device at a time.

2. The central processor executes the SSCH instruction and passes the request to the EXDC. As in 370 Mode, the EXDC attempts to start the request. If none of the four possible paths are available because of either channel, control unit, or device/head of string busy, the request is queued by the EXDC.

   The major difference in 370-XA Mode is that the request remains queued in the EXDC until the busy condition clears. That is,

until the I/O operation can be started. There is no 10ms timeout like that found in 370 Mode.

In 370 Mode, the 10ms timeout was implemented so that data could be collected about delays caused by control unit and device contention. In 370-XA Mode, this information is provided in a different manner.

370-XA Mode has a channel monitoring facility that has the capability of time stamping events during an I/O operation. For each I/O operation, actual queue time (if any), data transfer time, and total service time can be determined. This information is accumulated on a device basis and is used by RMF to produce device activity reports.

Foil 3.6 - 370-XA Mode Channel Monitoring Facility

Foil 3.7 - 370-XA Mode Channel Monitoring Facility

IBM 3084 PROCESSOR COMPLEX

370-XA CHANNEL MONITORING FACILITY

- CAN DETERMINE DURATION OF SIGNIFICANT EVENTS DURING AN I/O OPERATION
- COLLECTED FOR EACH DEVICE
- DATA ACCUMULATED IN CENTRAL STORAGE AND IS USED BY RMF
  - DEVICE DATA COLLECTED
    - TOTAL START SUBCHANNELS
    - SAMPLE START SUBCHANNLS
    - PLUS:

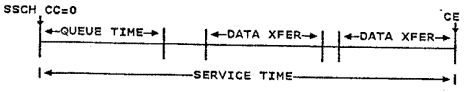

Foil 3.6

IBM 3084 PROCESSOR COMPLEX

370-XA CHANNEL MONITORING FACILITY

- LOGICAL CONTROL UNIT DATA COLLECTED
  - NUMBER OF ATTEMPTS TO START AN I/O OPERATION
  - NUMBER OF SUCCESSFUL ATTEMPTS TO START AN I/O OPERATION
  - NUMBER OF CU BUSY
  - NUMBER OF DEVICE BUSY
  - NUMBER REQUESTS QUEUED TO THE LOGICAL CONTROL UNIT

Foil 3.7

The Channel Monitoring Facility is a timer driven function that provides the ability to determine the duration of significant events that occur during an I/O operation. Data is accumulated on a device basis and is stored in Central Storage. The data will be used by RMF to produce I/O and device related reports.

The monitoring facility can be activated and deactivated on a device basis by the software. Normally, the monitoring will be active.

The data that can be collected includes:

Total SSCH — This is the total number of SSCHs issued to a given device. This includes SSCHs executed when the monitoring facility was ON and OFF.

Sample SSCH — This is a count of SSCHs issued when the monitoring facility was ON.

Queue Time — This is the elapsed time from the time the I/O request is placed on the channel queue until the I/O operation actually is started (start of Initial Selection). It includes time waiting because of channel, control unit, and/or device/head of string busy.

Data Transfer Time This is the time when the device is actually connected to the control unit and channel and is transferring data.

Service Time This is the elapsed time from the beginning of Queue Time until Channel End is presented to the channel.

In addition to the above, data is also provided on a logical control unit level. It includes:

1. Number of attempts to start an I/O operation.
2. Number of successful attempts to start an I/O operation.
3. Number of unsuccessful attempts because of control unit busy.
4. Number of unsuccessful attempts because of device/head of string busy.
5. Number of requests queued to this logical control unit.

This data is maintained in the channel and can be retrieved by software. This data will also be used by RMF.

Foil 3.8 - 370-XA Mode Channel Path Selection (Functions)

Foil 3.9 - 370-XA Mode Channel Path Selection (Example)

IBM 3084 PROCESSOR COMPLEX

370-XA MODE CHANNEL PATH SELECTION

- START SUBCHANNEL ISSUED BY SOFTWARE
  - INDICATES SUBCHANNEL (DEVICE) TO BE STARTED
- CHANNEL PATH SELECTION DONE BY MICROCODE AT LOGICAL CONTROL UNIT LEVEL
  - ROTATE
  - PREFERRED PATH
- MORE SIGNIFICANT FOR NON-3380 DASD DEVICES

Foil 3.8

IBM 3084 PROCESSOR COMPLEX

370-XA MODE CHANNEL PATH SELECTION

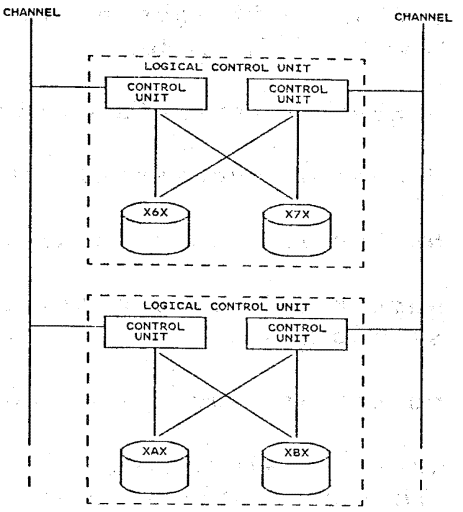

Foil 3.9

In previous foils, the capabilities of dynamic path selection were discussed. We learned that I/O operations to 3380s with dynamic path selection are no longer tied to a specific path. An operation could be started down one path and completed on another. Part of the dynamic path selection process is channel selection. That is, determining which channel will be selected initially to try to start the I/O operation. In these foils we will learn, in a little more detail, how the the channel subsystem decides which of the four possible channels to attempt to start a given I/O request. First we will discuss how 370-XA channel selection works on all 308X systems. And then, we will discuss the specifics of how it works on a 3084.

In 370-XA Mode, all I/O operations are started with the new I/O instruction, START SUBCHANNEL. Essentially, this instruction tells the channel subsystem to start an operation to a specific device (subchannel). It is the responsibility of the channel subsystem (i.e., channel microcode) to select the channel that will be used for the I/O operation.

Channel selection in 370-XA Mode is done at the logical control unit level. A logical control unit is similar in concept to a logical channel[15]. It includes all the devices that share the same control unit(s).

The default channel selection is rotate. That is, I/O operations are started by rotating the requests among the up to four possible channels. For example, in the foil shown, the first request would be tried on left channel and if it were successful the next request would be tried on the right channel.

In addition to the rotate function, a preferred channel can be specified for each device. The preferred channel is the channel that will always be tried first. Again referring to the foil, the X6X devices could be given the left channel as their preferred channel, while the X7X devices could be given the right channel as their preferred channel. In this situation, all requests to the X6X devices would be tried down the left channel first; if unsuccessful, the right channel would be tried. Likewise, requests to the X7X devices would be tried down the right channel first; if unsuccessful, the left channel would be tried. Note that each device can be given a preferred path.

The preferred channel function is particularly useful in those situations where the user would like to isolate activity to selected devices while at the same time maintaining maximum availability.

Comparing the 370-XA microcode channel selection to the MVS 370 Mode software channel selection, one finds the 370-XA channel selection provides more control at a finer level of granularity. 370 Mode channel selection is done at a logical channel level. 370-XA Mode is done at a logical control unit level.

Channel selection is the same regardless of device type (i.e., 3380 with DPS or other DASD). However, the implications of which channel is selected are much more significant for those devices that do not support dynamic path selection. This is because whichever channel is selected must be used for the entire I/O operation. Dynamic path selection is somewhat 'self tuning' because it selects an available path at the time the path is needed. However, path selection for devices not capable of dynamic pathing need additional user control to ensure an acceptable level of utilization to meet performance requirements.

---

[15] It should be noted that logical channels no longer exist in 370-XA mode.

Foil 3.10 - External Data Controller (EXDC) Structure

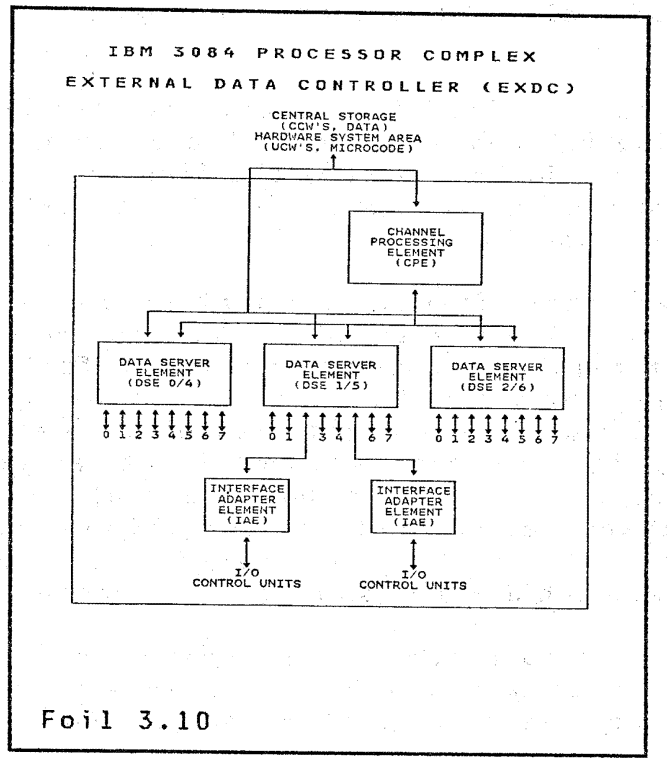

When the 3084 is running in single image mode under MVS/XA, the channel subsystem is logically the same as it is on any 308X system. That is, any central processor has direct access to any I/O device. And any processor can field an I/O interrupt whether or not that processor initiated the I/O operation. The difference on the 3084 is that it has two EXDCs. Therefore, the microcode will work slightly different.

However, before we get into those differences, lets take a few moments to review the physical structure of the 308X EXDC.

The EXDC is composed of three major elements; Interface Adapter Element (IAE), Data Server Element (DSE), and Channel Processing Element (CPE).

The IAE controls the data movement to and from the I/O control units. IAEs are physically located in the 3082. They provide the physical connection to the cables that connect to the I/O control units. There are 48 IAEs in the 3084, eight in every DSE. The IAE can operate at a data rate of 3.0 megabytes per second. The IAE is hardware and can function as either a byte or block channel. The DSE microcode determines if the IAE is a byte or block channel.

Each IAE has a two digit, hardware assigned channel path identifier (CHPID). The first digit is the DSE number to which it is attached (0, 1, and 2 on Side A, and 4, 5, and 6 on Side B). The second digit is the IAE number (0 through 7).

Up to four byte channels can be defined per EXDC. Byte channels can be assigned to DSE 0, 1, 4 and 5 only. Byte channels cannot be defined to DSE 2 or 6. The byte channels must be assigned to one of the first four path identifiers within the DSE (i.e., 00 through 03, 10 through 13, 40 through 43, and 50 through 53).

The primary function of the DSE is to control the data transfer between the IAE and Central Storage. When the data transfer is complete, the DSE analyzes the ending status and passes it along with the channel status to the CPE.

There are no preassigned channel priorities. IAEs are continuously polled and each channel is assured service via a time slicing technique. This means that high speed devices can be attached to any channel. It is no longer necessary to restrict high speed devices to low address (high priority) channels as was necessary on the 303X processors.

The DSE is capable of handling all eight channels transferring data at a rate of 3.0 megabytes (24 megabyte aggregate data rate). The DSEs are microcode controlled and each DSE is packaged in a single TCM.

In general, the CPE provides the management functions for the EXDC. It receives I/O instructions from the Central Processor, passes the command to the appropriate DSE, handles command modifications (e.g., chaining), and finally receives the ending status from the DSE and presents it to the Central Processor. When contention for use of the storage bus exists between DSEs, the CPE decides which DSE get the use of the bus.

Foil 3.11 - 3084 XA Mode Channel Path Selection (Single EXDC)

Foil 3.12 - 3084 XA Mode Channel Path Selection (Multiple EXDCs)

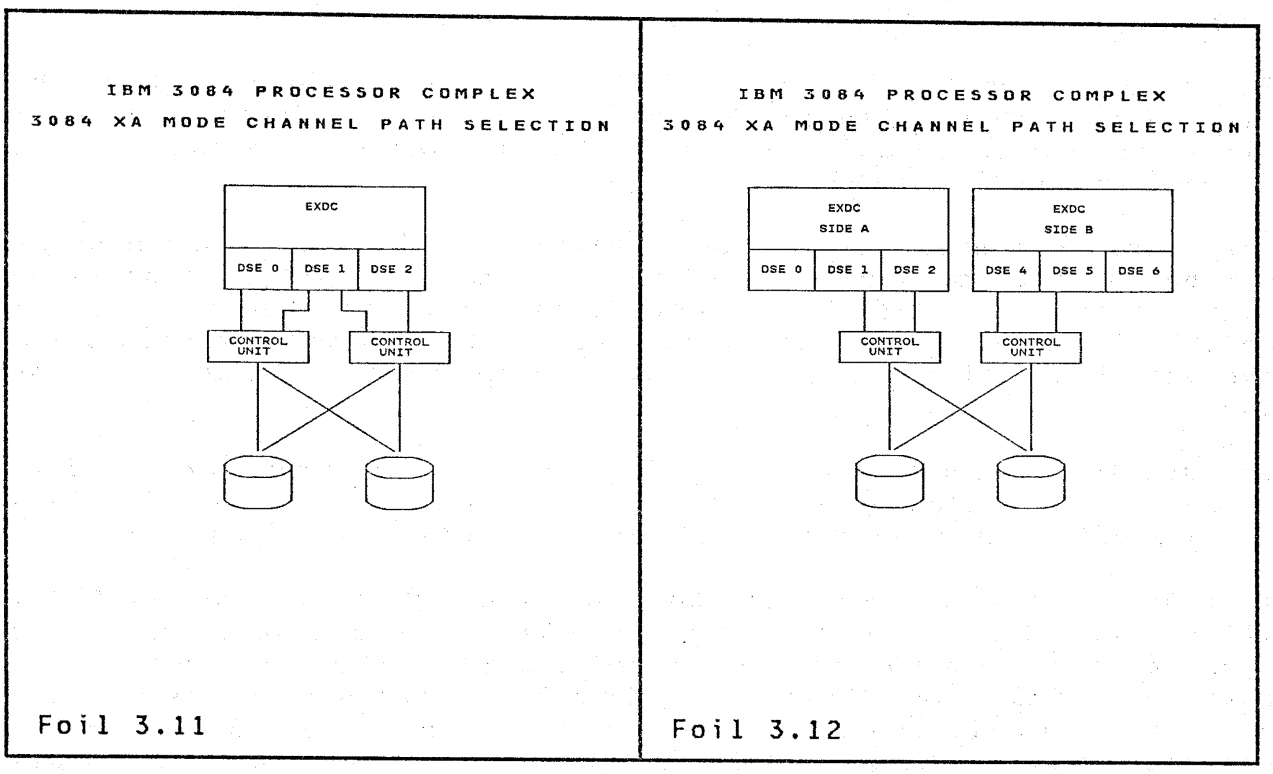

The intent of these two foils is to illustrate how 370-XA path selection works with a single EXDC and how it works with two EXDCs. These foils are not intended to suggest how channels should be configured. Channel configurations will be covered later in this section.

In MVS/XA, up to four channel paths can be defined to a device. Normally, the channel subsystem will use a rotate algorithm to determine which path will be used for a given I/O request. The paths to a given device are sequenced based upon how they were defined to the IOCP. The channel remembers which path was used in the last request to the device. The channel will attempt to start the next request to the device on the next path in the sequence. If all the paths to a given device are defined from a single EXDC, the central processor that initiates the request, sends the request to that EXDC. The EXDC will then attempt to start the request using the rotate algorithm described above.

However, if the paths to a device are defined to two EXDCs, as could be the case in the 3084, then the path selection works slightly different. For example, suppose a device has two paths defined to it from each EXDC. When the Start Subchannel command is executed, the request is broadcast to both EXDCs. The EXDC that responds first will attempt to start the request via the paths defined to it. Requests will be

rotated within the paths defined to this EXDC. If all the paths to the device, defined to this EXDC, are busy, the other EXDC is signaled with a request to start the operation. The second EXDC also rotates the requests among the paths defined to it.

In addition to rotate, one of the up to four possible paths to a device can be defined as a 'preferred path'. If a preferred path is defined, the channel subsystem will always try that path first. If the preferred path is busy, the remaining paths are tried using the rotate algorithm.

In the 3084, even though paths may be defined to both EXDCs, only one preferred path can be defined. When an I/O request is initiated to a device that has paths defined to both EXDCs, and has a preferred path, only the EXDC with the preferred path is notified. If all the paths defined to this EXDC are busy, the other EXDC is signaled and requested to start the operation.

Let's summarize 370-XA path selection on the 3084. Because paths can be defined to a specific device from two EXDCs, path selection is a two stage process; selecting the EXDC, and then selecting the path. If a device has a preferred path or if all defined paths to a device are from a single EXDC, then only that EXDC is signaled when an I/O operation is initiated. If no preferred paths have been defined, and there are paths from both EXDCs, the request is broadcast to both EXDCs. Whichever EXDC responds first, that EXDC will attempt to start the operation. If unsuccessful, the other EXDC is signaled.

Path selection is done by rotating the requests among the paths defined to a particular EXDC. If all the paths to a device are defined from a single EXDC, then the requests will be rotated among those paths. If a device has paths defined from two EXDCs, whichever EXDC gets the request will rotate the request among the paths defined to it. If unsuccessful, the other EXDC is signaled, and it rotates the request among its paths.

From a performance standpoint, the effects of configuring across one or two EXDCs should be about the same. This is because in each case the load will tend to balance itself across the defined paths. The question of defining paths from one or two EXDCs should be centered around availability considerations and the modes under which the 3084 will operate.

Foil 3.13 - 3084 Channel Configuration (Single Image Mode)

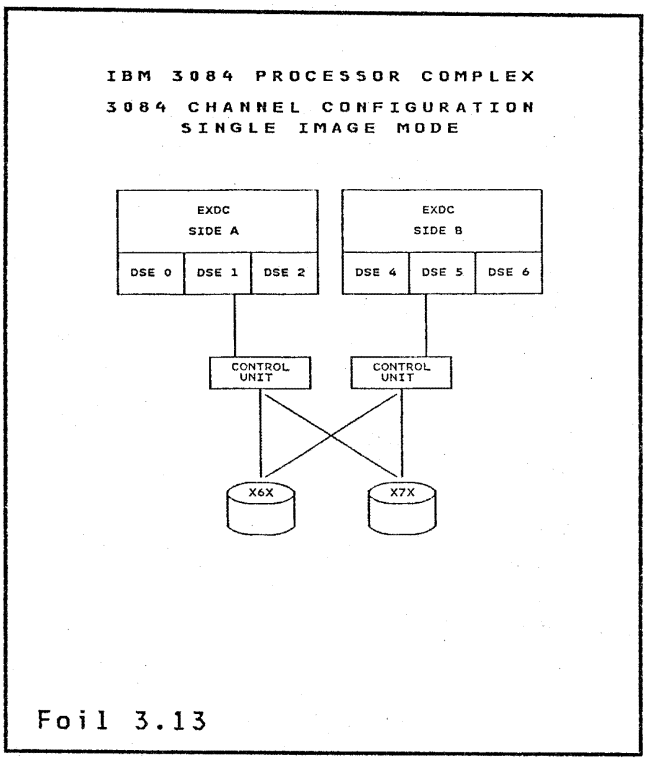

There are a great many variables that effect how an installation will configure its channels. It is not the intent of this document to cover all the considerations for configuring an I/O subsystem. Rather, the intent is to introduce the reader to some considerations that may be unique to configuring a 3084.

How the 3084 channels are configured depends not only on availability and performance considerations, but also on what mode(s) the system will be running. Considerations are somewhat different depending upon whether the system will run in single image mode only, partitioned mode only, or a combination of both. As was stated earlier, running in partitioned mode is the same as running as two stand alone 3081 Ks. Because much has already been documented on configuring 3081 channels, it will not be repeated here[16]. Instead we will concentrate on systems running in single image mode only and a combination of single image and partitioned modes.

The fact is that there isn't anything new about configuring a single image mode 3084. The same concepts and considerations that applied to the 3033MP, apply to the 3084. As was the case with the 3033MP, the

---

16 Refer to 3081 Channel Characteristics and Configuration Guide, GA22-7077, and 3081 IOCP and Channel Configuration, GG22-9209, and 3081 Channel Performance Evaluation Methodology, ZZ05-0380.

main consideration is whether the system will be split and run as two independent processors.

In general, the following guidelines should be followed.

1. Symmetry should always be considered, that is, an equal number of paths from each side (EXDC) should be configured.

2. For availability reasons, select Channel Path IDs from different DSEs if more that one path is defined from a side.

For example, if the 3084 is going to run in single image mode only[17], then the DASD configuration could consist of a channel path from each EXDC. This configuration provides for availability in that if an EXDC, channel, or control unit fails, there still is a path available to all devices.

---

[17] Actually, it is anticipated that all 3084s will run partitioned at least for some period of time. For example, when doing maintenance the installation will probably choose to take one side at a time.

Foil 3.14 - 3084 Mixed Mode Physical Configuration

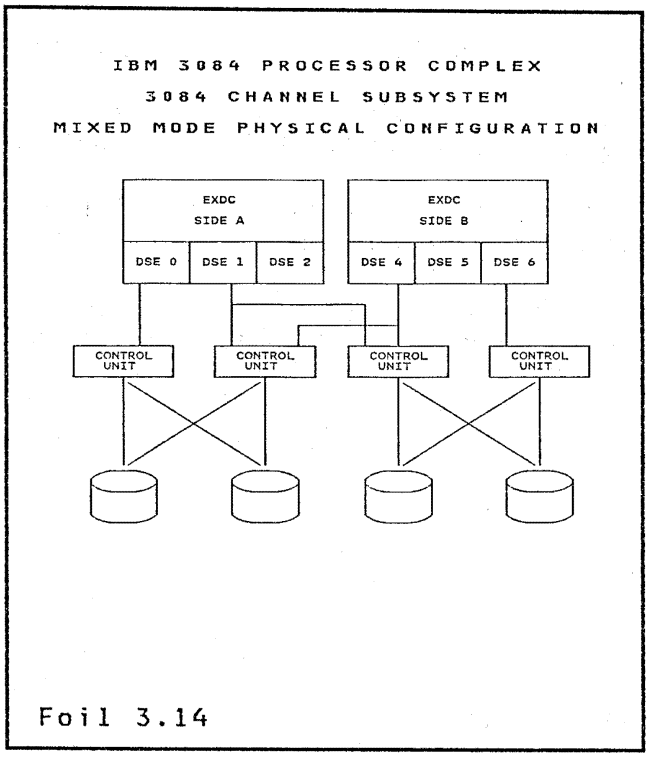

If the 3084 will be run in partitioned mode for any length of time, then four paths, two from each side (EXDC) is probably more appropriate. This provides two paths to all the devices when the system is running in partitioned mode. For example, using the diagram in the above foil, one could physically cable the control units as shown.

Foil 3.15 - 3084 Configuration - Mixed Mode Running Single Image

Foil 3.16 - 3084 Configuration - Mixed Mode Running Partitioned

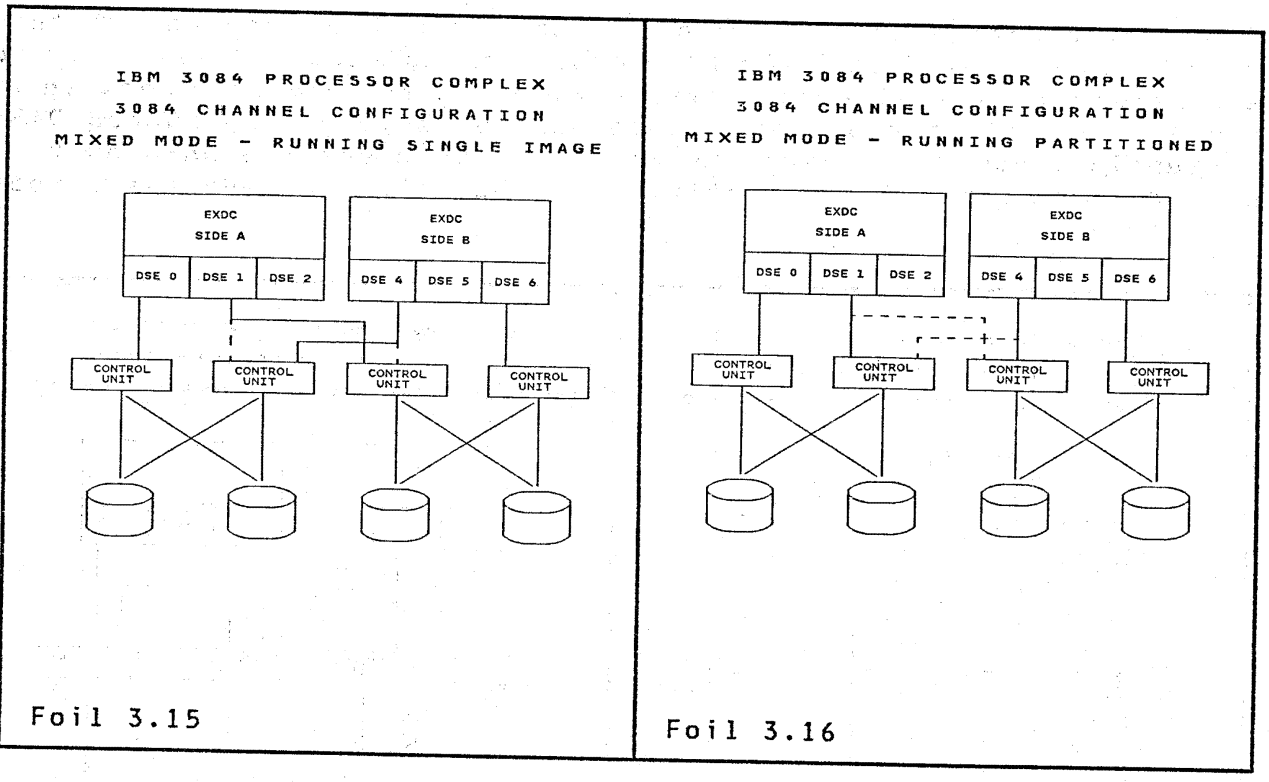

However, when running in single image mode, the path from DSE 1 could be configured offline to the left hand pair of control units, while the path from DSE 4 could be configured offline to the right hand pair of control units. This would give each pair of control units two paths (that's all they would ever need), one path from each EXDC.

When running in partitioned mode, the left hand pair of control units could use the paths from DSE 0 and 1, while the right hand pair of control units could use the paths from DSE 4 and 6. Again, each pair of control units has two paths.

Foil 3.17 - 3084 Configuration Shared Mixed Mode - Single Image

Foil 3.18 - 3084 Configuration Shared Mixed Mode - Partitioned

The configuration considerations, that have been covered to this point, have assumed that the DASD could be isolated and dedicated to each side when partitioned. This is the preferred approach because it simplifies DASD performance management by eliminating shared DASD contention. However, there are environments where the DASD must be shared when the 3084 is partitioned. The following addresses that environment.

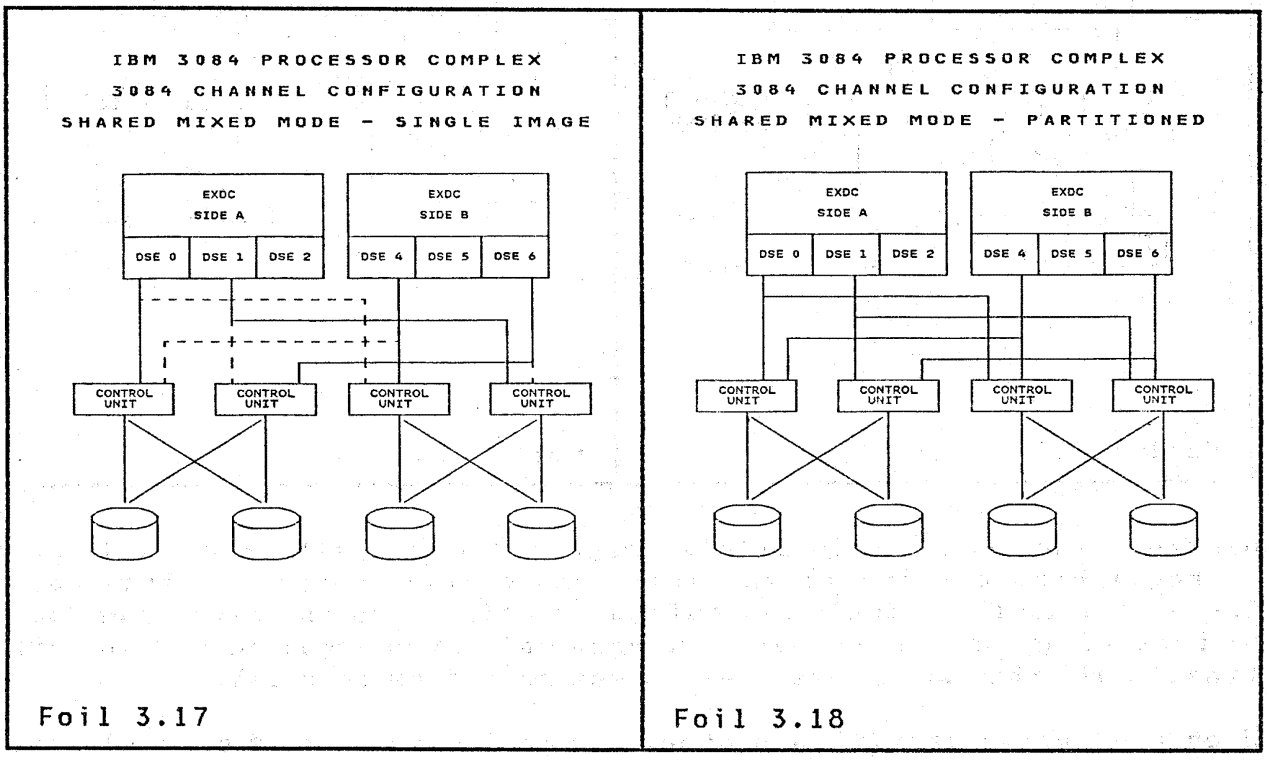

As you can see, this physical configuration is slightly different from the previous configuration in that there are two physical paths to each control unit. When running in single image mode, one path to each control unit is configured offline in such a way as to provide a path from each EXDC to each pair of control units.

When running in partitioned mode, all physical paths are active. This provides two paths to each pair of control units from each side.

Foil 3.19 - 3084 Configuration - Additional Considerations

Foil 3.20 - 3084 Configuration - Additional Considerations

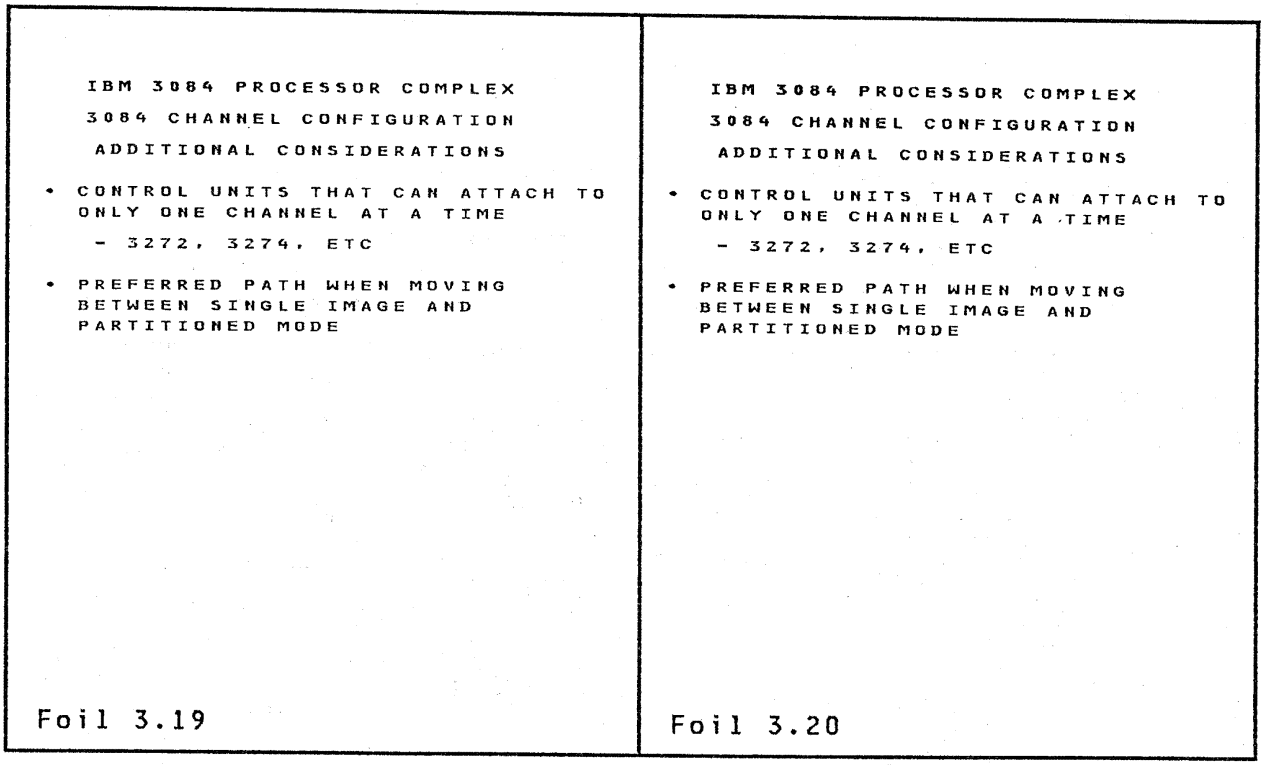

Although asymmetric devices do not exist in 370-XA mode, an installation must still concern itself with devices that have control units that can attach to only one channel at a time, e.g., 3274. For these kinds of devices a switch, such as the 3814, should be used so that the device can be switched to another path. Again , if running predominately in single image mode, this second path should be from the other EXDC.

When configuring DASD devices, a preferred path may be beneficial. Because only one preferred can be specified to a device, consideration should be given to ensure that the preferred path is not assigned to the side that will be configured offline when going from single image mode to partitioned mode.

As you can see, these are general guidelines for maintaining maximum system availability. They are applicable in most environments. As additional operational experience with the 3084 is gained, additional guidelines will be provided.

3084 CONSOLES

This section describes the 3084 Consoles. It includes a review of the consoles available to the 308X processors plus some considerations for configuring 3084 consoles.

Foil 4.1 - 3084 Consoles Per Side (Hardware Layout)

Foil 4.2 - 3084 Consoles (Functional Description)

IBM 3084 PROCESSOR CUMPLEX
3084 CONSOLES

- OPERATOR'S CONSOLE
  - SOFTWARE INTERFACE
  - ANY CONSOLE MVS OR VM SUPPORT
  - IF 3278 OR 3279, CAN BE UP TO 1500 METERS FROM PROCESSOR
- SYSTEM CONSOLE
  - SYSTEMS FUNCTIONS
    - IPL, ALTER/DISPLAY, CONFIGURATION FRAMES, ETC.
  - SYSTEM ACTIVITY DISPLAY
  - NO INTERFACE TO SOFTWARE
  - 3278 MODEL 2A
    - UP TO 1500 METERS FROM PROCESSOR
- SERVICE SUPPORT CONSOLE
  - CE MAINTENANCE CONSOLE
  - CAN PERFORM FUNCTIONS OF BOTH SYSTEM AND OPERATOR'S CONSOLE
  - SIMILAR TO 3278 MODEL 2A
- PROGRAMMING SUPPORT CONSOLE
  - RETAIN ACCESS
  - OPTIONAL
- SYSTEM/SERVICE SUPPORT CONSOLE HARDCOPY PRINTER
  - OPTIONAL

Foil 4.1

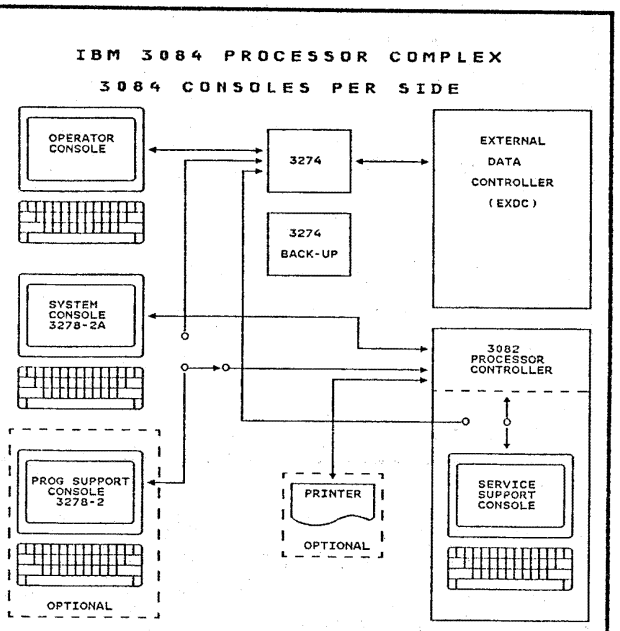

Foil 4.2

Our discussion of the 3084 consoles will be done in two stages. First, we will review each of the consoles and their function. Then we will address some configuration considerations in both single image and partitioned mode.

To review the consoles and their functions, we will look at the consoles that can be attached to each side of the 3084.

Key Points:

1. Operator's Console - Interface to the software. Any device currently supported by the software (i.e., MVS/XA for single image mode, and MVS/XA, MVS/SP, or VM/SP in partitioned mode) can be used as an operator's console. However, 3278 and/or 3279 devices are recommended because they can be placed up to 1500 meters

from the processor. This is particularly important for those installations who choose to operate their systems from a control center.

Operator's consoles do not come with the 3084. They must be ordered separately. Therefore, it is the responsibility of each installation to ensure that operator's consoles are available when installing a 3084. This includes any control units that may be required.

The operator's console configuration is attached to a channel of the 3084 Processor and provides the user with the normal interface to the SCP software. If the the operator console is a 3278 or 3279, a 3274 Control Unit is required. A second 3274 is shown on the foil to remind the audience that a backup control unit is recommended for maximum availability.

When configuring a 3084 for single image mode, a minimum of two MVS consoles are recommended. For maximum availability, each should be attached to a different EXDC.

2. System Console - Performs hardware functions only. This includes such things as IML, IPL, Alter/Display, System Activity Display, etc[18]. The System Console is a 3278 Model 2A. It can be placed up to 1500 meters from the processor.

   The 3084 has two System Consoles; one attached to each side. Normally, when operating the 3084 in single image mode, the operator will use both System Consoles.

   The System Console is attached directly to the Processor Controller. It is not channel attached. Therefore, the System Console cannot interact with the software (i.e. cannot be used as an operator's console). The System Console can assume the functions of the Service Support Console in a back-up situation.

3. Service Support Console - This console is used by the CE to run diagnostics, execute the stand alone version of IOCP, and other service related activities that require communication with the Processor Controller. It, like the System Console, is connected directly to the Processor Controller. In addition, the Service Support Console can perform all the functions of the System Console and, when attached to the 3274 Control Unit, can operate as an operator's console.

   The 3084 has two Service Support Consoles. When running in single image mode, one Service Support Console is active, the other is in standby mode.

---

[18] The System Activity Display provides displays similar to those provided by the 3033 Processor (e.g., processor and channel utilizations).

Minimum service configuration requires that the CE have have access to an operator's console for concurrent execution of software diagnostic and maintenance programs. The recommended way to satisfy this requirement is to channel attach the Service Support Console through a 3274 Control Unit. However, any console supported by the software could satisfy this requirement.

Changing from Service Support to operator functions is accomplished via a switch located within the Processor Controller. The Service Support Consoles are similar to the 3278-2A. One is housed in each side of the 3082 Q.

4. Programming Support Console (Optional). An optional Programming Support Console can be attached to the system. It attaches to the TP Modem through the Processor Controller for access to the RETAIN system. If attached to a 3274 Control Unit, the Programming Support Console can be used as an operator's console. The Programming Support Console must be a 3278 Model 2.

   If an installation chooses not to have a Program Support Console, no functions are lost because RETAIN can be accessed from either the System or Service Support Consoles. The Program Support Console provides the convenience of being able to place the RETAIN access function in a more convenient location.

5. Printer (Optional) - An optional printer can be attached to the Processor Controller. This printer can be either a 3287, 3268, or 3230 (refer to the sales manual for specific models required). It provides the capability of printing any frame displayed on either the System or Service Support Console. Although this printer is optional, actual experience indicates that it is an invaluable tool for both FE and customer operations. For this reason it is recommended that the printer be attached and dedicated to the Processor Controller on all 308X Systems.

As can be seen from the above, the console arrangement on the 3084 Processor is different from the 303X system. The 3036 provides the combined functions of operator, system, and service console functions. Operator, system and service functions are physically separate in the 3084 Processor. The system and service functions have been redesigned and enhanced and are now performed by the Processor Controller. This arrangement satisfies the requirement of the CE needing a service console near the processor complex, while at the same time allowing the system and operator consoles to be some distance away.

Foil 4.3 - 3084 Console Configuration - Single Image Mode

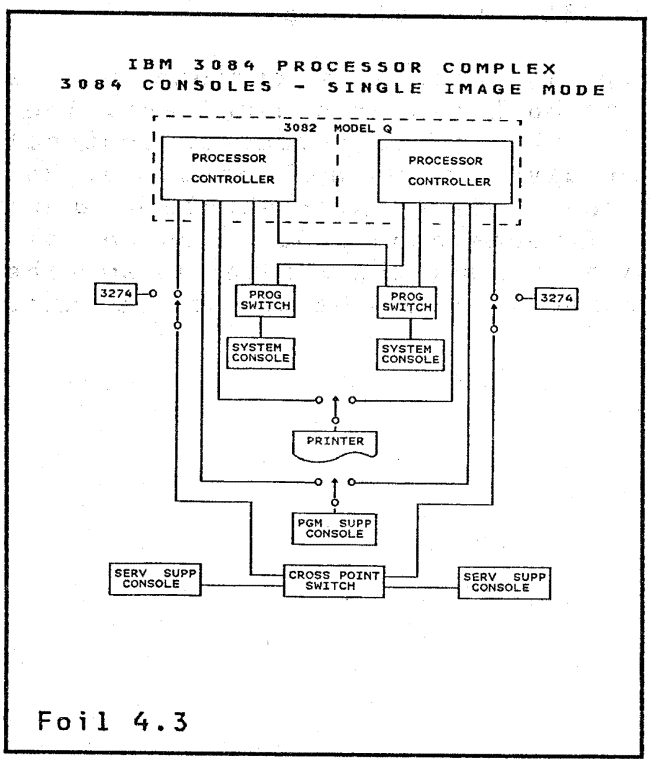

Before discussing console configurations, let's look at the physical attachment of the 3084 consoles. Note that both system consoles can be attached to either processor controller via a programmable switch. This switch is under control of the 3082 microcode. When operating in single image mode, both system consoles are attached to the active processor controller[19]. Both system consoles are used when operating in single image mode. For example, one system console may be used to show the system activity display for one side, the other system console may display the activity of the other side.

When operating in partitioned mode, each system console is attached to its respective processor controller.

Both Service Support Consoles are attached through a cross point switch. With this switch either Service Support Console can be connected to either processor controller. This provides a back-up capability should one of the Service Support Consoles fail.

It is possible to attach a single Program Support Console, through a user supplied switch, to either processor controller. Normally, the Program Support Console would be switched to the active processor

---

[19] When operating in single image mode, only one processor controller is active. The other is in stand by mode. Processor controller operation is discussed in the Processor Controller section of this document.

controller. As is the case on all other 308X Processors, the Program Support Console can be channel attached through a 3274 and used as an operator's console, TSO terminal, etc.

Like the Program Support Console, an optional printer can be attached through a switch, to both system controllers. The printer would be normally attached to the active processor controller. Technically, this printer could also be channel attached. However, experience strongly suggests that this printer should be dedicated to the processor controller. If something appears on one of the consoles and you want a hardcopy of it, and the printer is switched to the channel, the display will probably be gone before you can switch the printer.

Foil 4.4 - 3084 Console Configuration - Partitioned Mode

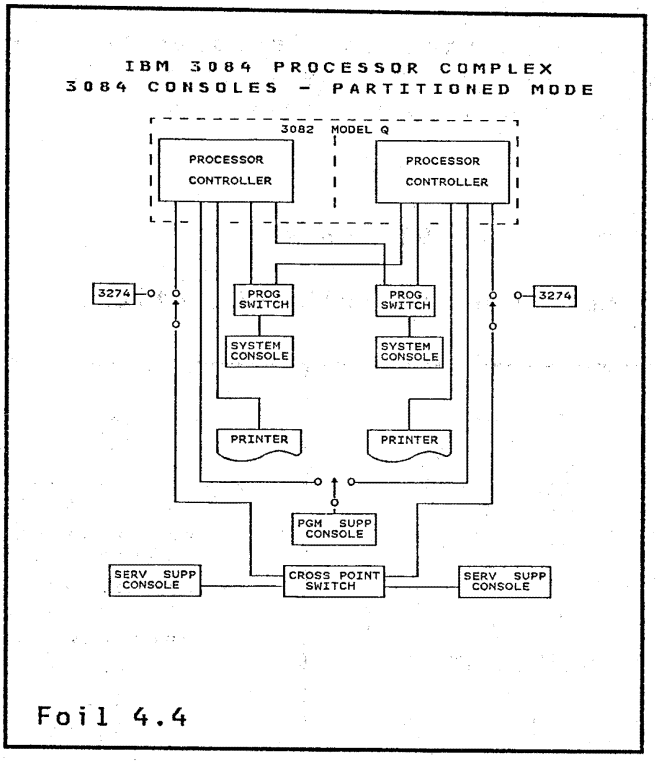

Foil 4.4

In partitioned mode, the programmable switches are set by the 3082 microcode so that each system console is attached to its side only. This establishes a configuration exactly like a stand alone 3081 K.

One could attach a Program Support Console to each side of the 3084, but it really isn't necessary. It is highly unlikely that you will need to access RETAIN simultaneously from both sides. A single Program Support Console switchable to either side, via a user supplied switch, should be adequate for either single image or partitioned operation.

The printer is another story. If the 3084 will be run primarily in partitioned mode, a printer should be dedicated to each side. Again, if you have a display that you want printed and the printer is not switched to the system, you will probably lose the display before you can switch the printer.

The service support consoles are switchable the same way in partitioned mode as they are in single image mode. That is, either console can be switched to either processor controller through the cross point switch.

To sum up the 3084 console configurations, the major consideration is which mode(s) the system will be running. From an equipment standpoint, the real difference is whether you attach one or two printers to the 3082 Q.

3082 MODEL Q PROCESSOR CONTROLLER

This section of the presentation discusses the structure and operation of the 3082 Model Q Processor Controller. In addition, the enhanced 3082 Q error recovery is presented.

Foil 5.1 - 3084 Functional Elements

Foil 5.2 - 3082 Q Processor Controller Functions

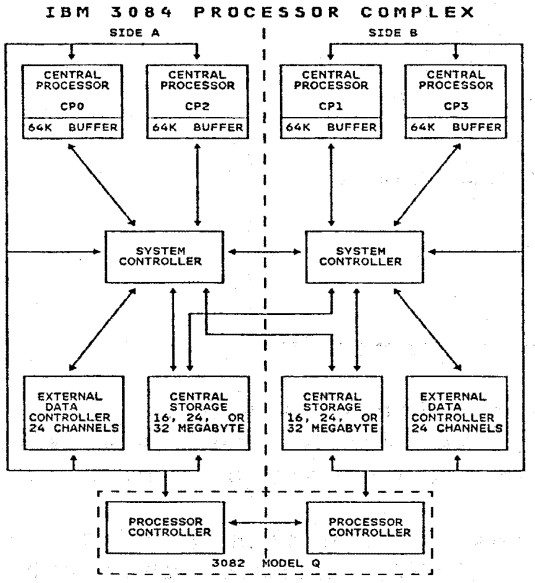

Foil 5.1

IBM 3084 PROCESSOR COMPLEX

3082 MODEL Q

- MONITORING AND SYSTEM SUPPORT FACILITY (MSSF)
  - POWER SEQUENCING, MONITORING AND CONTROL
  - MICROCODE LOADING AND INITIALIZATION
  - CONFIGURATION CONTROL
    - LOGICAL AND PHYSICAL
  - ERROR HANDLING
    - RECOVERY
    - RECORDING
    - DATA ANALYSIS
  - INTERACTIVE CE INTERFACE
- ENHANCED 3082 RECOVERY IN SINGLE IMAGE MODE

Foil 5.2

The 3082 Model Q Processor Controller (PC) provides monitoring and system support functions for the 3084 Processor[20]. The 3082 Q is essentially a combination of two 3082 Model 24s. When operating in single image mode, one half of the 3082 Q, usually Side A, is active, while the other half is in standby mode. In partitioned mode, each half of the 3082 Q supports the side to which it is attached.

---

[20] The functions performed by the Processor Controller are sometimes referred to as "MSSF". This is an acronym for Monitoring and System Support Facility. MSSF is sometimes used to refer to the Processor Controller.

Key functions performed by the 3082 include the following.

1. Power on/off sequencing for the processor complex is controlled by the 3082. During the power on sequence, all voltages are checked and adjusted automatically if required. Once operational, voltage, current and temperature are monitored. The 3082 can power a single regulator on or off when necessary. The 3082 also collects error data for FRU (Field Replaceable Unit) isolation.

2. All microcode for the 3084 processor is stored within the 3082 on the IPC files. Each IPC file should have the identical microcode. When operating in single image mode, all microcode is loaded from the active side. When operating in partitioned mode, microcode for each side is loaded from its own IPC. This microcode supports Central Processor, EXDC, and Processor Controller functions. Most of this microcode is associated with the Processor Controller. During IML, the CP and EXDC microcode is loaded into the Hardware System Area of Central Storage. From the Hardware System Area it will be loaded into the CPs and EXDC(s).

3. Configuration Control - Prior to 308X systems, logical configuration was accomplished by issuing an MVS command (VARY CPU, Storage, etc.) and physical reconfiguration was accomplished by setting switches. In the 3084 system both logical and physical reconfiguration can be accomplished via MVS commands

4. Error Handling - Error handling functions are concentrated in the 3082. In addition to the expected functions of error recovery and error recording, the 3082 has the ability to analyze error data and identify failing units. This analysis is done in parallel with normal processor operation. Refer to the '3081 Design Presentation' for details of 3082 error analysis.

5. Each side of the 3082 Q contains a Service Support Console. This console provides the CE with a vehicle for interacting with the system. Error data can be viewed, diagnostics run, etc. The Service Support Console is also used to run the stand alone version of IOCP.

6. When running in single image mode, the 3082 Q provides an additional level of recovery. This is because most of the elements within 3082 Q are duplexed in single image mode. Should a failure occur within an active element, a switch can be made to the standby element. More specific information on this recovery process is included later in this section.

Foil 5.3 - 3082 Q Components Per Side

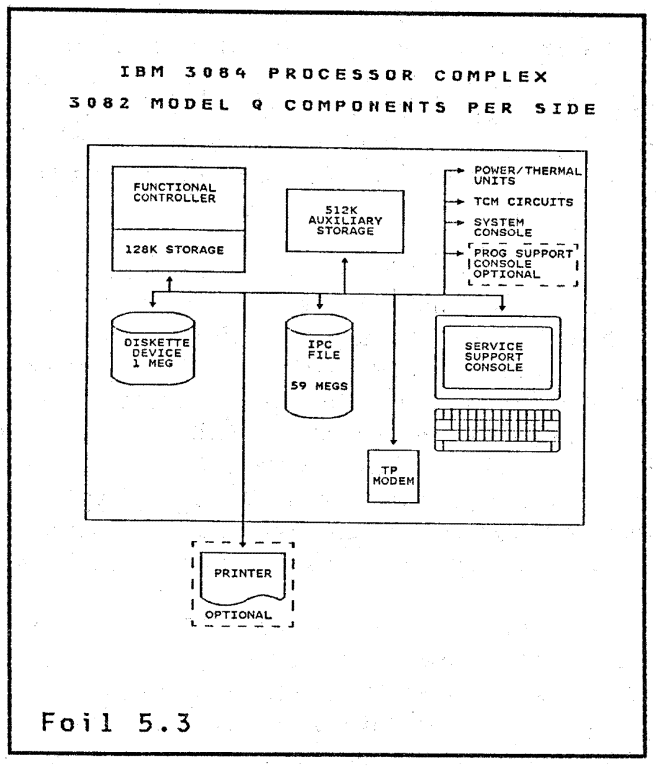

Before getting into the details of how the 3082 Q operates, it may be worth a few minutes to review the 3082 Q components and their functions.

Each side of the 3082 Q consists of the following components.

1. Functional Controller - A microprocessor with 128K of native storage attached to it.

2. Programmed I/O Bus - Controls the various devices that attach to the PC. Functionally it provides services similar to those of a channel.

3. Auxiliary Storage - 512k bytes of high speed storage that appear to the processor as a zero latency file. This storage is used for various tables and data that the PC updates and must have relatively fast access to.

4. Integrated Processor Controller (IPC) File - Fixed (non-removable) file with a capacity of 59 megabytes. This file is used to store system microcode, error analysis and diagnostic microcode, EC history, feature information, error logout data, IOCP images, etc.

5. Diskette Device - Removable diskette with a 1 megabyte capacity. This file is used for initial load or refresh of the IPC file and

for introducing microcode ECs into the system and other service related functions.

6. TP Modem - Provides attachment, via customer supplied telephone line, to RETAIN and remote support location. Through this link, all functions can be performed that would normally be performed at the System and Service Support Consoles.

7. Service Support Console - A console similar to the 3278-2A that is housed within the Processor Controller. Its functions are described in the '3084 Consoles' section.

8. TCM and Power/Thermal Sensors - The Processor Controller is attached to sensors throughout the 3084 system. These sensors are used for monitoring and during error recording and recovery.

9. Printer (optional) - The function and use of this printer was discussed in the section on 3084 Consoles.

Foil 5.4 - 3082 Q - Single Image Mode

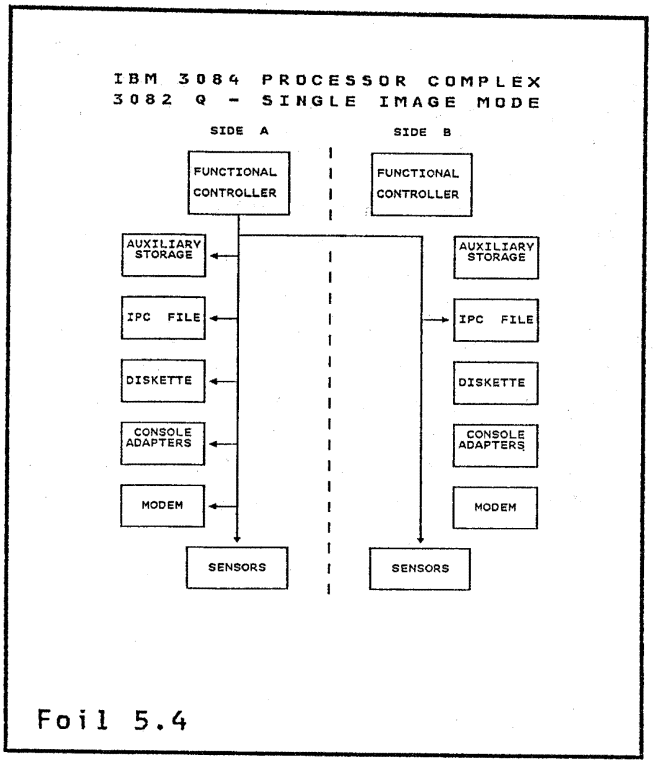

When operating in single image mode, the 3082 Q is logically configured as shown. Generally, the functional controller on the A side will be active and the B side will be in standby mode. All the components on both the A and B side are actually connected to the active functional controller. However, with the exception of the IPC File and the Sensors, only one of the two like components will be active. For example, only one of the two Auxiliary Storage Units will be active. The other is in backup mode.

Because each set of Sensors is hard wired to its side, both sets are active and connected to the active functional controller.

Both IPC Files are normally active. If an update is required to the IPC File, that update is written to both files. This is done to ensure that both IPC Files contain the same information so that if one should fail, the system can continue with the remaining file.

AS we have seen, keeping critical data on both files identical, allows one file to back up the other. System microcode is stored on the IPC file. Because either IPC File may required to support the entire system in single image mode, this microcode must be the same on both files. This is the reason the system checks the Engineering Change (EC) level during IML to ensure the EC level is the same on both sides.

If an active component within the 3082 fails, the standby component will be switched in. For example, should the active Auxiliary Storage fail, the standby Auxiliary Storage would become active.

Only one set of console adapters is active. These adapters are used to connect the various consoles to the 3082. Consoles connected to the 3082 include the system console, service support console, and optionally, the program support console. As was shown on previous foils, these consoles are switched to the active side of the 3082 when running in single image mode. Actually, the consoles are switched to the console adapters on the active side.

Foil 5.5 - 3082 Q - Partitioned Mode

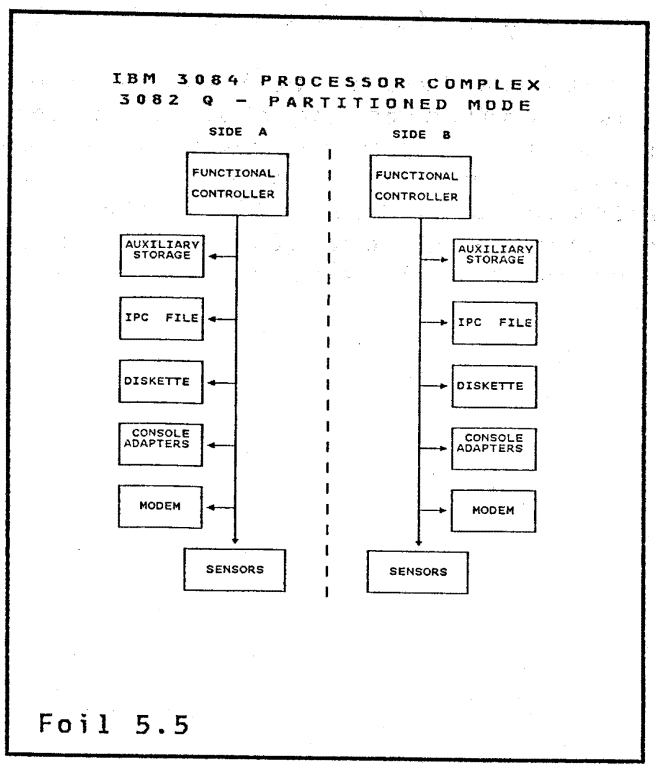

In partitioned mode, the 3082 Q is split into two independent sides. Each side is equivalent to a 3082 Model 24. Each side is powered separately. So that when running partitioned, one side could be powered down without effecting the other side.

3082 Q Recovery Logic

There are really several aspects of 3084 recovery logic that should identified so that the differences between the 3084 and other 308X processors can be understood. They are:

1. 3084 recovery logic in partitioned mode.

2. 3084 recovery logic in single image mode.

3. 3082 recovery logic in partitioned mode.

4. 3082 recovery logic in single image mode.

3084 and 3082 recovery logic, when operating in partitioned mode,is identical to that of a 3081 K. This recovery process is discussed in detail in the '3081 Design Presentation'. If you are not familiar with that information, you should review it before continuing here.

3084 single image mode recovery logic is also the same as in the 3081. There is an operational difference in that one half of the 3082 Q is doing the monitoring and recovery for the entire 3084. In order to do this, both sets of sensors are connected to the active half of the 3082. Other than this, 3084 single image mode recovery logic is the same as described in the 3081 Design Presentation.

The real difference in the 3084 Processor Complex over other 308X systems is the recovery capabilities the 3082 Q has when operating in single image mode. We have seen that when operating in single image mode, half of the 3082 Q is active, the other half is in standby mode. This backup of 3082 components exists only in the 3084 when operating in single image mode and therefore, provides additional recovery capabilities.

Foil 5.6 - 3082 Q Single Image Error Recovery Logic

Foil 5.7 - 3082 Q Single Image Configuration

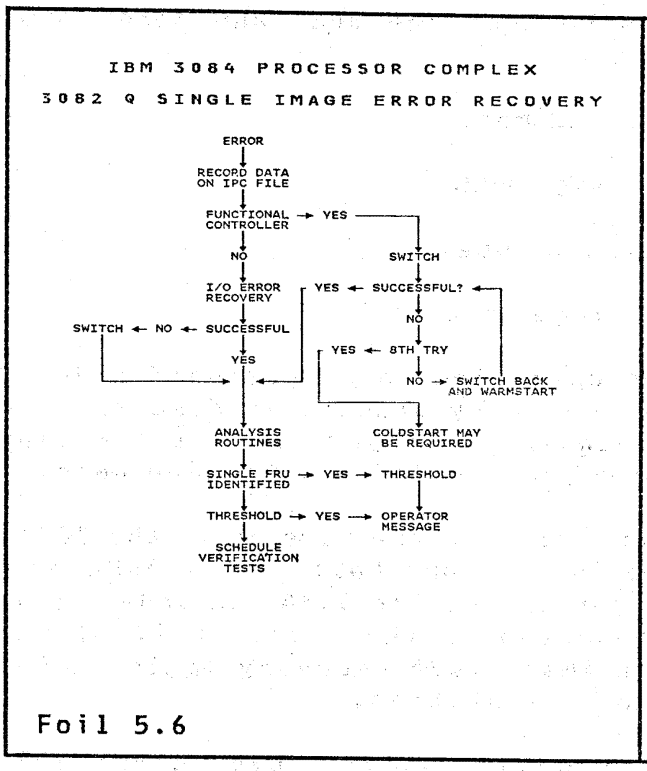

Foil 5.6

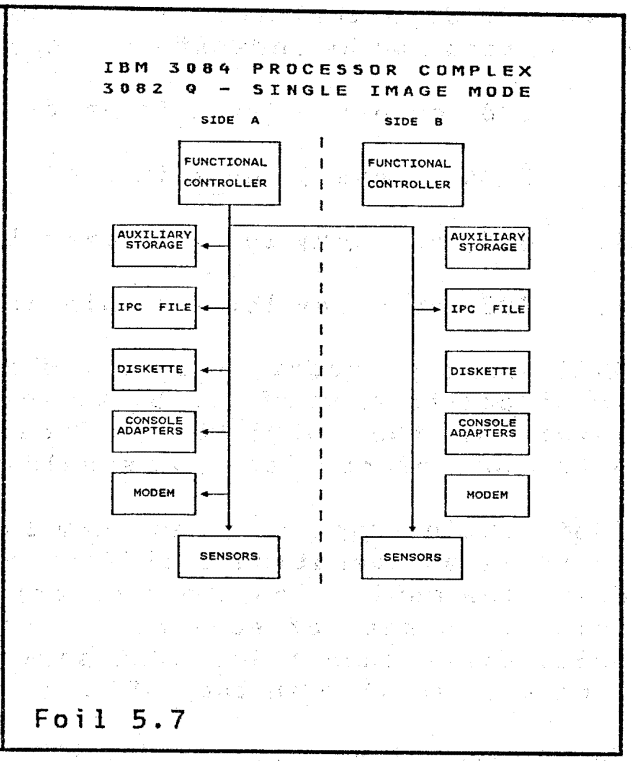

Foil 5.7

As you can see, all components of the 3082 Q, except the inactive functional controller, are connected to the active functional controller. All components of the 3082 have backup except the sensors. Because the sensors are wired to each side, both sets of sensors must be operational. A failure of either sensor is a catastrophic failure and will require an IML.

Normally, when the 3084 is powered on, the A Side and its components become active. Should a failure occur, a backup component may be switched in to replace the failing unit.

If the functional controller fails, an immediate switch to the backup functional controller takes place. Assuming the switch is successful, the now active functional controller takes control and analysis routines are scheduled for the failing functional controller. If the switch was not successful, a switch back to the original functional controller takes place and a warmstart is scheduled to it[21]. Up to a total of eight retries are completed before the failure is considered

---

[21] A warmstart consists of reloading the functional controller microcode. 3082 warmstarts are designed to be independent of 3084 operation and therefore, if successful, should not effect system operation.

catastrophic. During these retries, multiple warmstarts will be attempted to each functional controller. If the retries are not successful, the error is considered catastrophic and a coldstart may be required.

If one of the devices attached to the programmed I/O bus (channel) fails, the microcoded error recovery routines for that device are invoked. If recovery is successful, the device remains active. If recovery is unsuccessful, the backup device becomes active and analysis routines are scheduled for the failing device.

The additional recovery capabilities for the 3082 Q, significantly reduces the potential for a single failure to effect the entire system. These 3082 capabilities along with the recovery capabilities of the 3084 that were described earlier in the document, enhance the availability characteristics of the 3084 Processor Complex.

SUMMARY

This document has attempted to introduce the capabilities of the 3084. As additional knowledge and experience are gained, this document will be updated.

APPENDIX A. 3084 POWER BOUNDARIES

The following diagram shows the power boundaries of each side of the 3084 Processor Complex. There are four major power domains within each side. Elements that share the same power and therefore are powered up and down together are:

1. Central Processor 0/1 and the System Controller.

2. Central Processor 2/3 and the EXDC

3. Central Storage

   a. Central Storage is powered at the BSM level. Two BSMs are powered together. Refer to the '3081 Design Presentation' for more information.

4. Processor Controller.

   a. The IAEs are physically located in the Processor Controller and are powered with it. The IAEs are powered in groups of eight.

3084 POWER BOUNDARIES

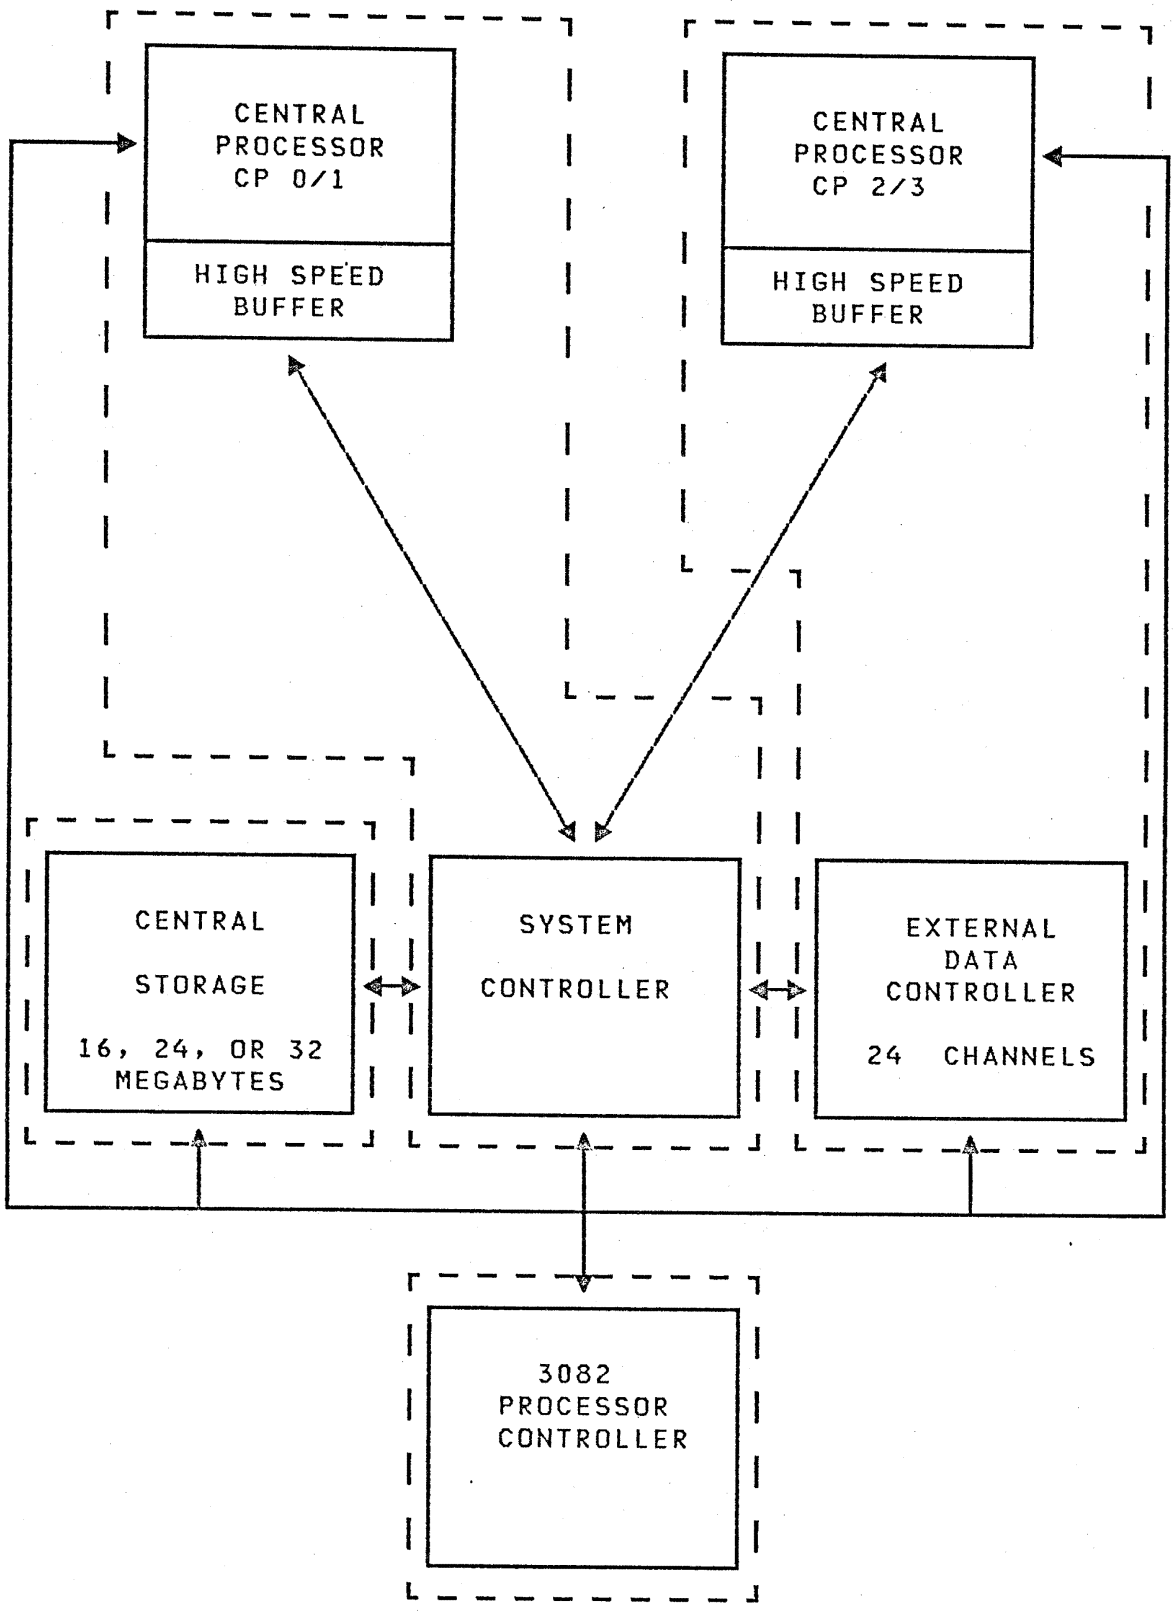

APPENDIX B. REFERENCE DOCUMENTATION

The following documents provide additional information about the IBM 3081 Processor Complex.

ANNOUNCEMENT MATERIAL

- 30XX Processor Evolution Presentation - G320-6525
- Tightly Coupled Processors Presentation - G320-6412
- Engineering/Scientific Computing: A Perspective - G320-6524-0
- Water Cooled, Large-Scale Processors - G320-6517
- IBM Large Systems Design - G320-6522

SRLS

- 3081 Functional Characteristics SRL - GA22-7076
- 3081 Channel Characteristics and Configuration Guide - GA22-7077
- MVS IOCP Users Guide - GC28-1027
- VM IOCP Users Guide - GC20-1843
- 3081 Operator's Guide - GC38-0034
- Guide to Writing MVS Recovery Procedures - GC28-0952
- MVS Recovery and Reconfiguration - GC28-1060

WASHINGTON SYSTEMS CENTER TECHNICAL BULLETINS

Customer Documents

- MVS/370 Tightly-Coupled Concepts - GG22-9276
- 3081 Operator Training Guide - GG22-9265
- 3081 Quick Reference Guide - GG22-9213
- 3081 IOCP and Channel Configuration - GG22-9209
- MVS Tuning Perspective - GG22-9023-04
- 3033 MP/AP Installation Notebook - GG22-9210-00

Internal Use Only Documents

- 3081 Design Presentation - ZZ05-0381
- 3033 Processor Presentation - ZZ05-0045
- 3081 Processor Complex Performance - ZZ05-0379
- 3081 K Performance Study - ZZ05-0394
- Large Systems Performance Reference - ZZ05-0400
- 3081 Channel Performance Evaluation Methodology - ZZ05-0380

OTHER

- IBM Journal of Research and Development Vol. 25, No. 5, September 1981
- IBM Journal of Research and Development Vol. 26, No. 1, January 1982

APPENDIX C. PRESENTATION FOILS

The following pages contain the presentation foil masters.

IBM
NA DIVISION
WASHINGTON
SYSTEMS
CENTER

AN INTRODUCTION TO THE 3084 PROCESSOR COMPLEX STRUCTURE AND CAPABILITIES

IBM 3084 PROCESSOR COMPLEX

TOPICS

- 3084 PROCESSOR OVERVIEW
- 3084 OPERATING MODES
- 3084 CHANNEL SUBSYSTEM
- 3084 CONSOLE CONFIGURATIONS
- 3082 PROCESSOR CONTROLLER

IBM 3084 PROCESSOR COMPLEX

3084 OVERVIEW

- FOUR INTEGRATED CENTRAL PROCESSORS
  - 26NS CYCLE
- 32, 48, OR 64 MEGABYTES OF REAL STORAGE
- 48 CHANNELS
  - DATA STREAMING STANDARD ON ALL BLOCK CHANNELS
  - MAXIMUM 8 BYTE CHANNELS
- OPERATING MODES
  - SINGLE IMAGE
    - MVS/XA ONLY
  - PARTITIONED
    - MVS/XA, MVS/SP OR VM/SP

IBM 3084 PROCESSOR COMPLEX

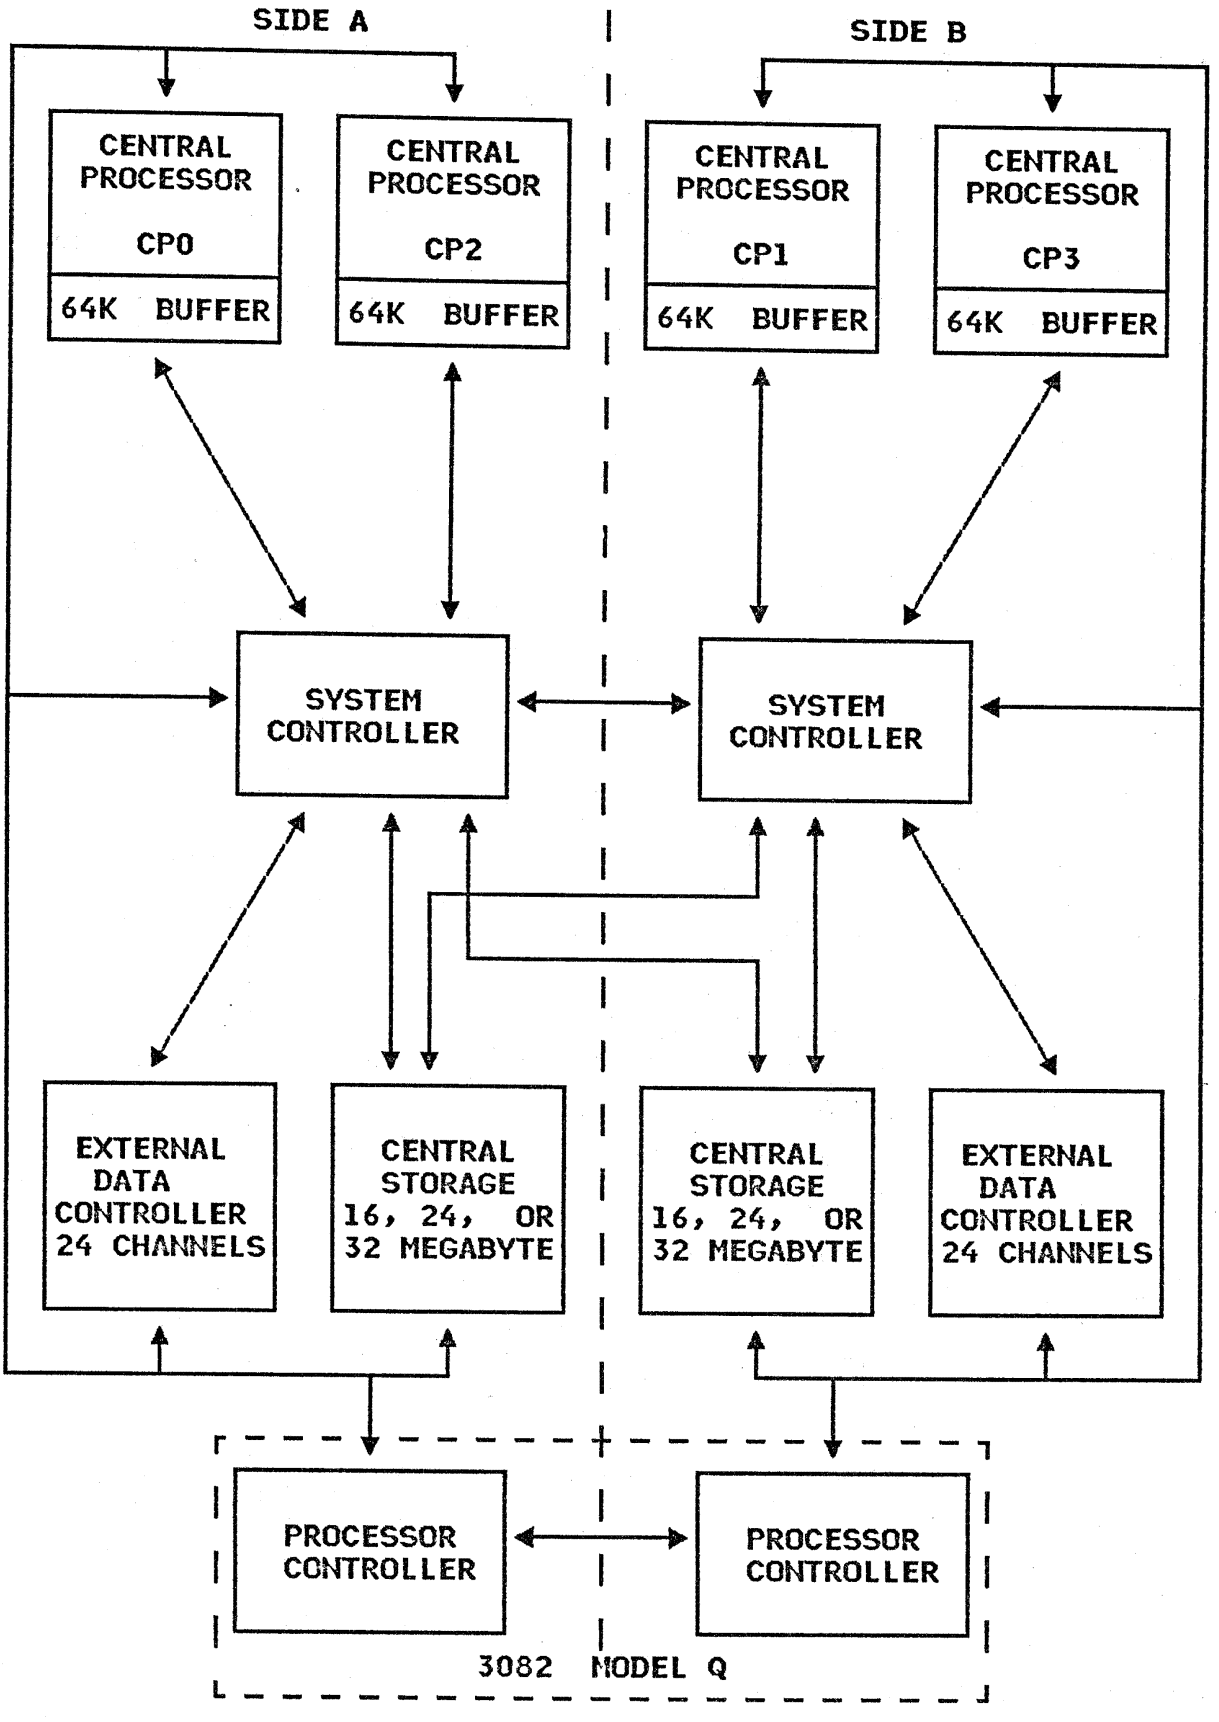

IBM 3084 PROCESSOR COMPLEX

3084 OVERVIEW - UPGRADEABILITY

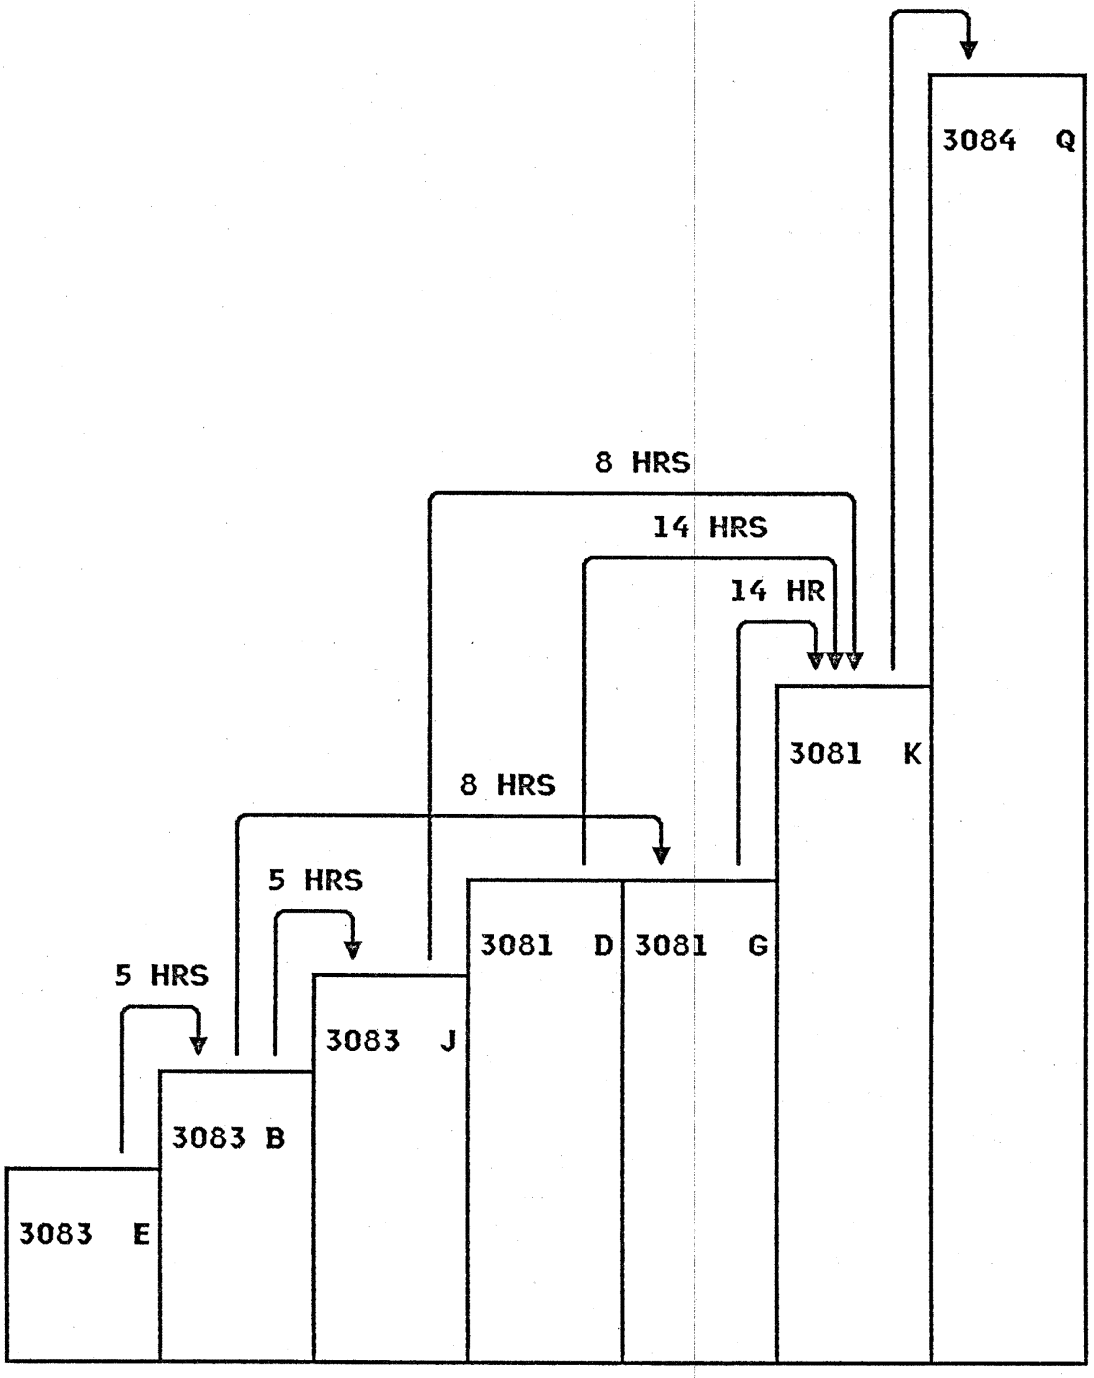

IBM 3084 PROCESSOR COMPLEX

3084 OVERIVEW

- STANDARD FEATURES
  - S/370 EXTENDED FACILITY
  - 3033 EXTENSION FEATURE
  - EXTENDED ADDRESSING
  - VIRTUAL MACHINE ASSIST
  - PREFERRED MACHINE ASSIST
- HARDWARE AND MICROCODE CONTROLLED
- WATER AND AIR COOLED

IBM 3084 PROCESSOR COMPLEX

3084 OVERVIEW - TECHNOLOGY

- THERMAL CONDUCTION MODULE (TCM)
  - UP TO 133 CHIPS PER TCM
    - SINGLE 4 7/8 BY 5 1/2 BY 1 3/8 INCH MODULE HAS ABOUT AS MANY CIRCUITS AS A S/370 MODEL 148
  - FILLED WITH HELIUM AND WATER COOLED
  - MOUNTED ON MULTI-LAYERED BOARD
  - OVER 90% OF THE CIRCUITS USE TCM/BOARD TECHNOLOGY
- LSI CARD ON BOARD
  - CHANNEL INTERFACE
- OTHER CARD ON BOARD
  - PROCESSOR CONTROLLER
- HIGHER DENSITY STORAGE

IBM 3084 PROCESSOR COMPLEX

3084 OVERVIEW - FLOOR PLAN

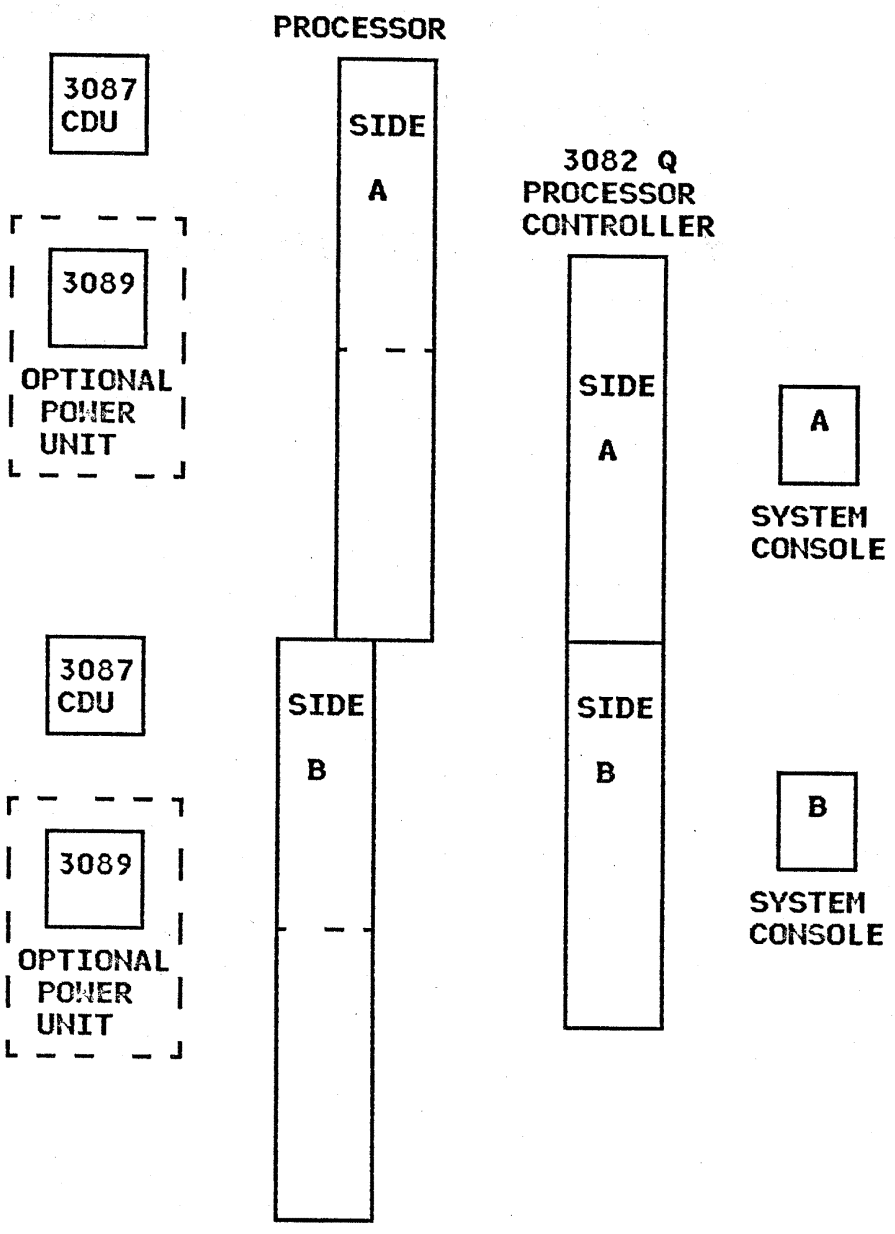

IBM 3084 PROCESSOR COMPLEX

3084 OPERATING MODES

- TWO MODES OF OPERATION
  - SINGLE IMAGE MODE
  - PARTITIONED MODE
- DYNAMIC RECONFIGURATION

IBM 3084 PROCESSOR COMPLEX

3084 SINGLE IMAGE MODE

- ENTIRE 3084 PROCESSOR COMPLEX OPERATING UNDER A SINGLE COPY OF MVS/XA
  - DYNAMIC CHANNEL SUBSYSTEM
    - ALL PROCESSORS HAVE ACCESS TO ALL DEVICES
  - EACH SIDE HAS DIRECT ACCESS TO ALL OF CENTRAL STORAGE
  - SYSTEM AVAILABILITY
  - BOTH SIDES MUST BE AT SAME EC LEVEL

IBM 3084 PROCESSOR COMPLEX

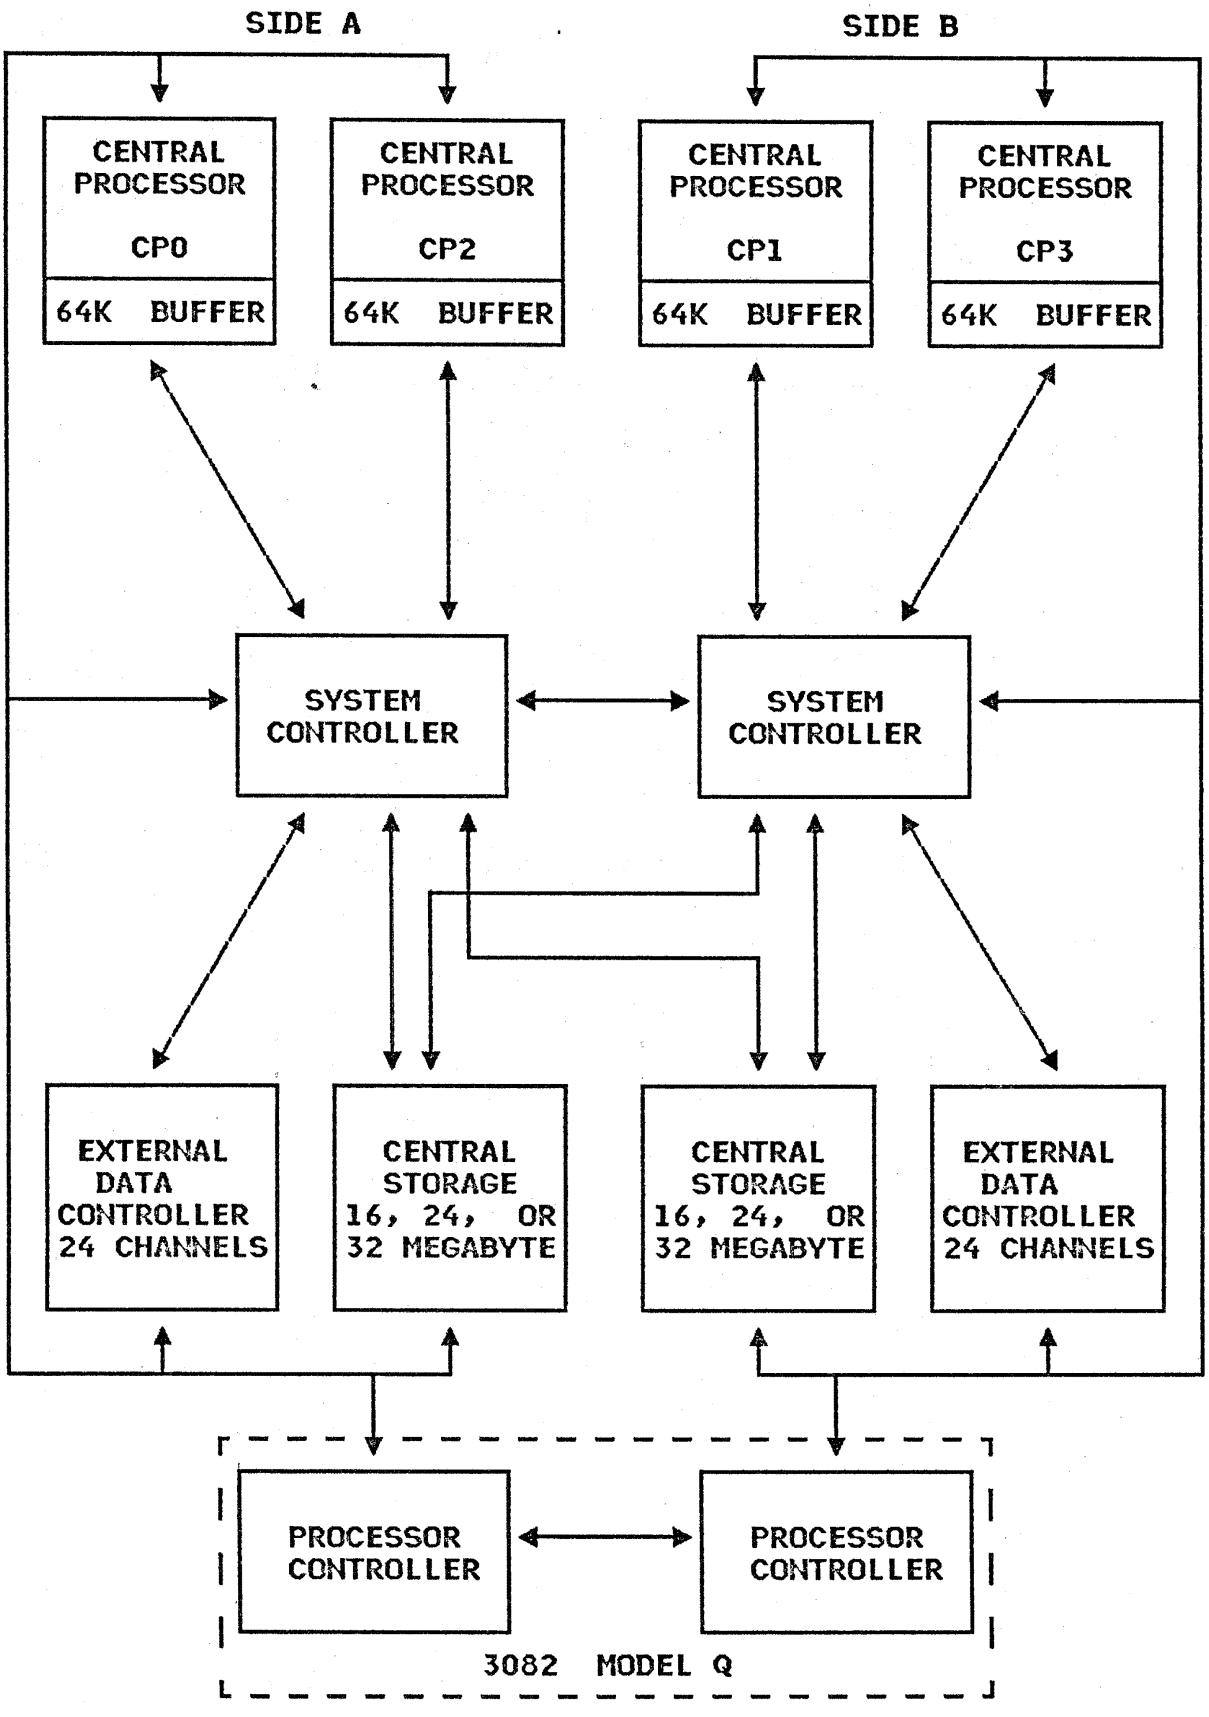

IBM 3084 PROCESSOR COMPLEX

3084 PARTITIONED MODE

- 3084 SPLIT INTO TWO INDEPENDENT SYSTEMS
  - EACH SIDE CAN RUN EITHER MVS/XA, MVS/SP, OR VM/SP
  - MUST BE SYMMETRIC
    - CENTRAL PROCESSORS
    - CHANNELS
    - STORAGE

IBM 3084 PROCESSOR COMPLEX

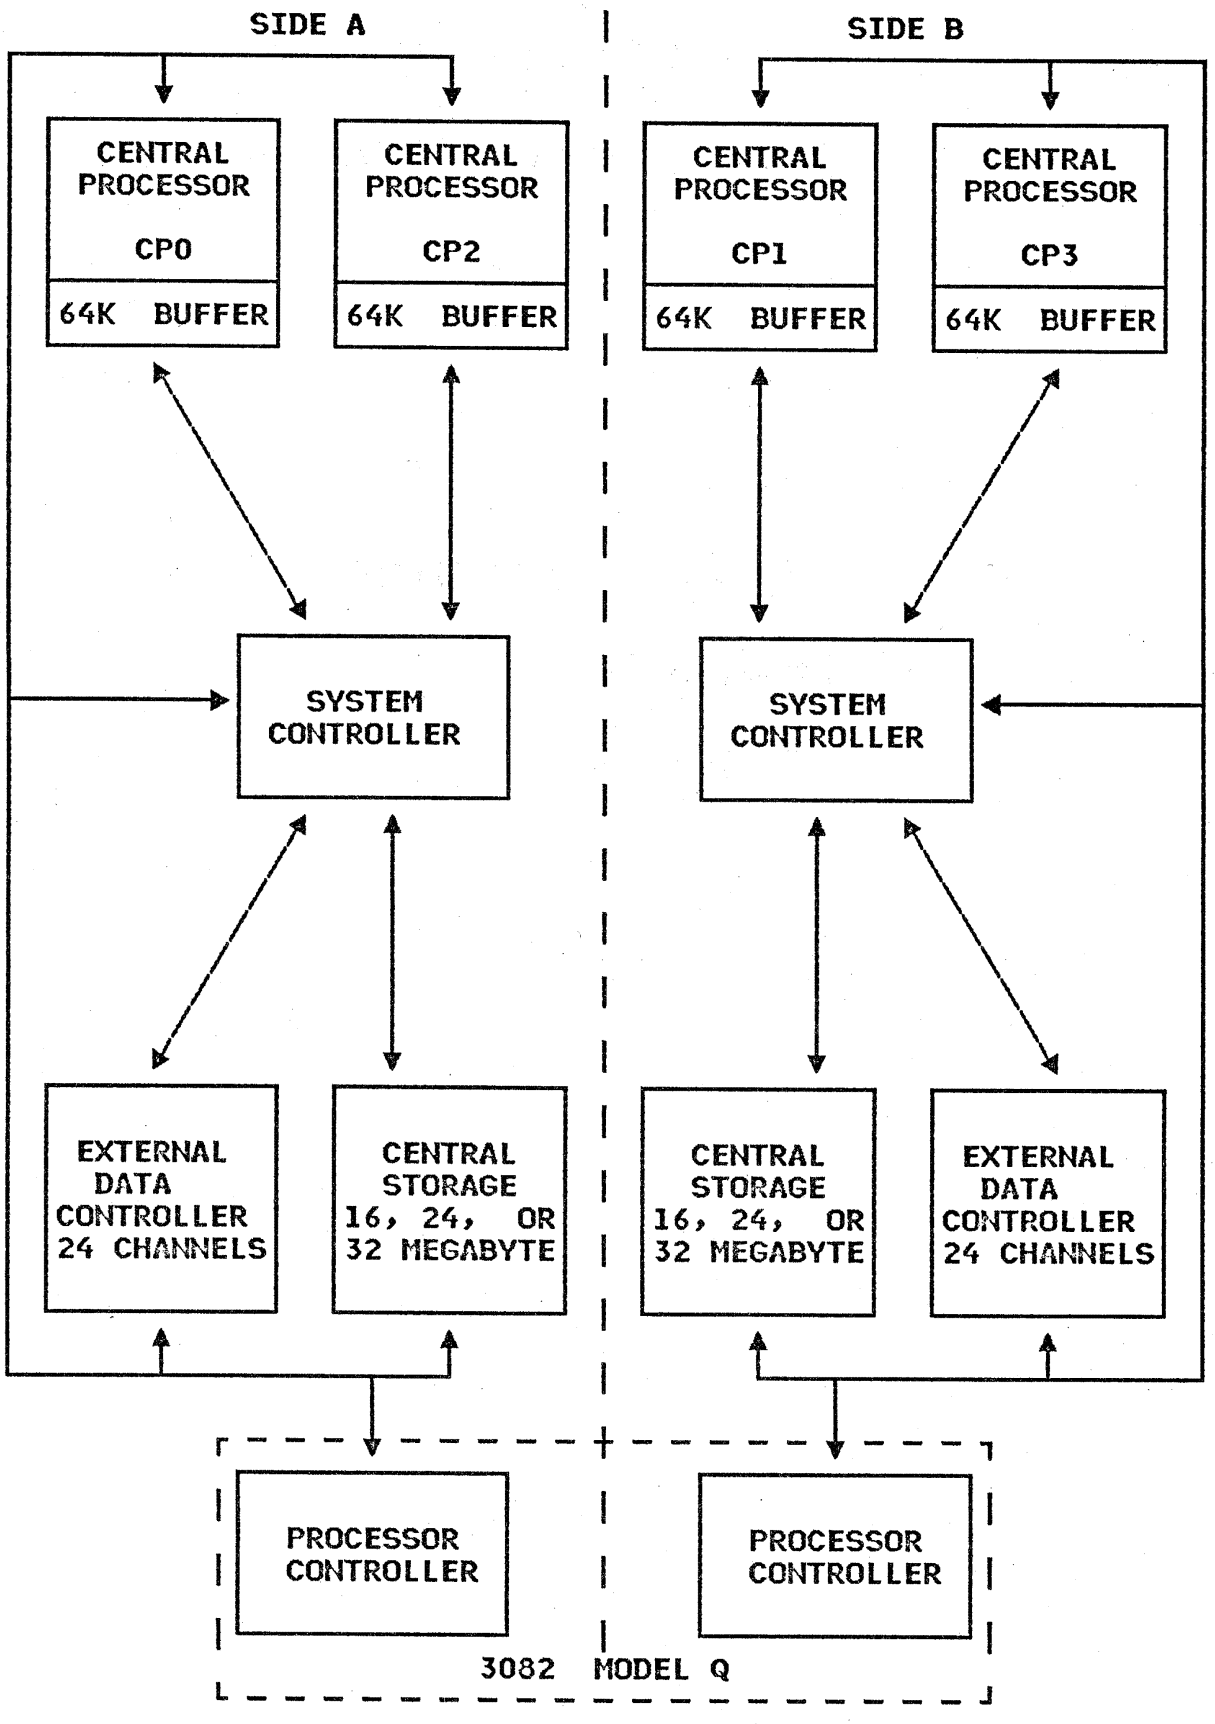

IBM 3084 PROCESSOR COMPLEX

3084 DYNAMIC RECONFIGURATION

- ABILITY TO DYNAMICALLY SWITCH OPERATING MODES WHILE RUNNING UNDER MVS/XA
  - MVS/XA CAN OPERATE, WITHOUT INTERRUPTION, WHILE SWITCHING BETWEEN SINGLE IMAGE AND PARTITIONED MODE

IBM 3084 PROCESSOR COMPLEX

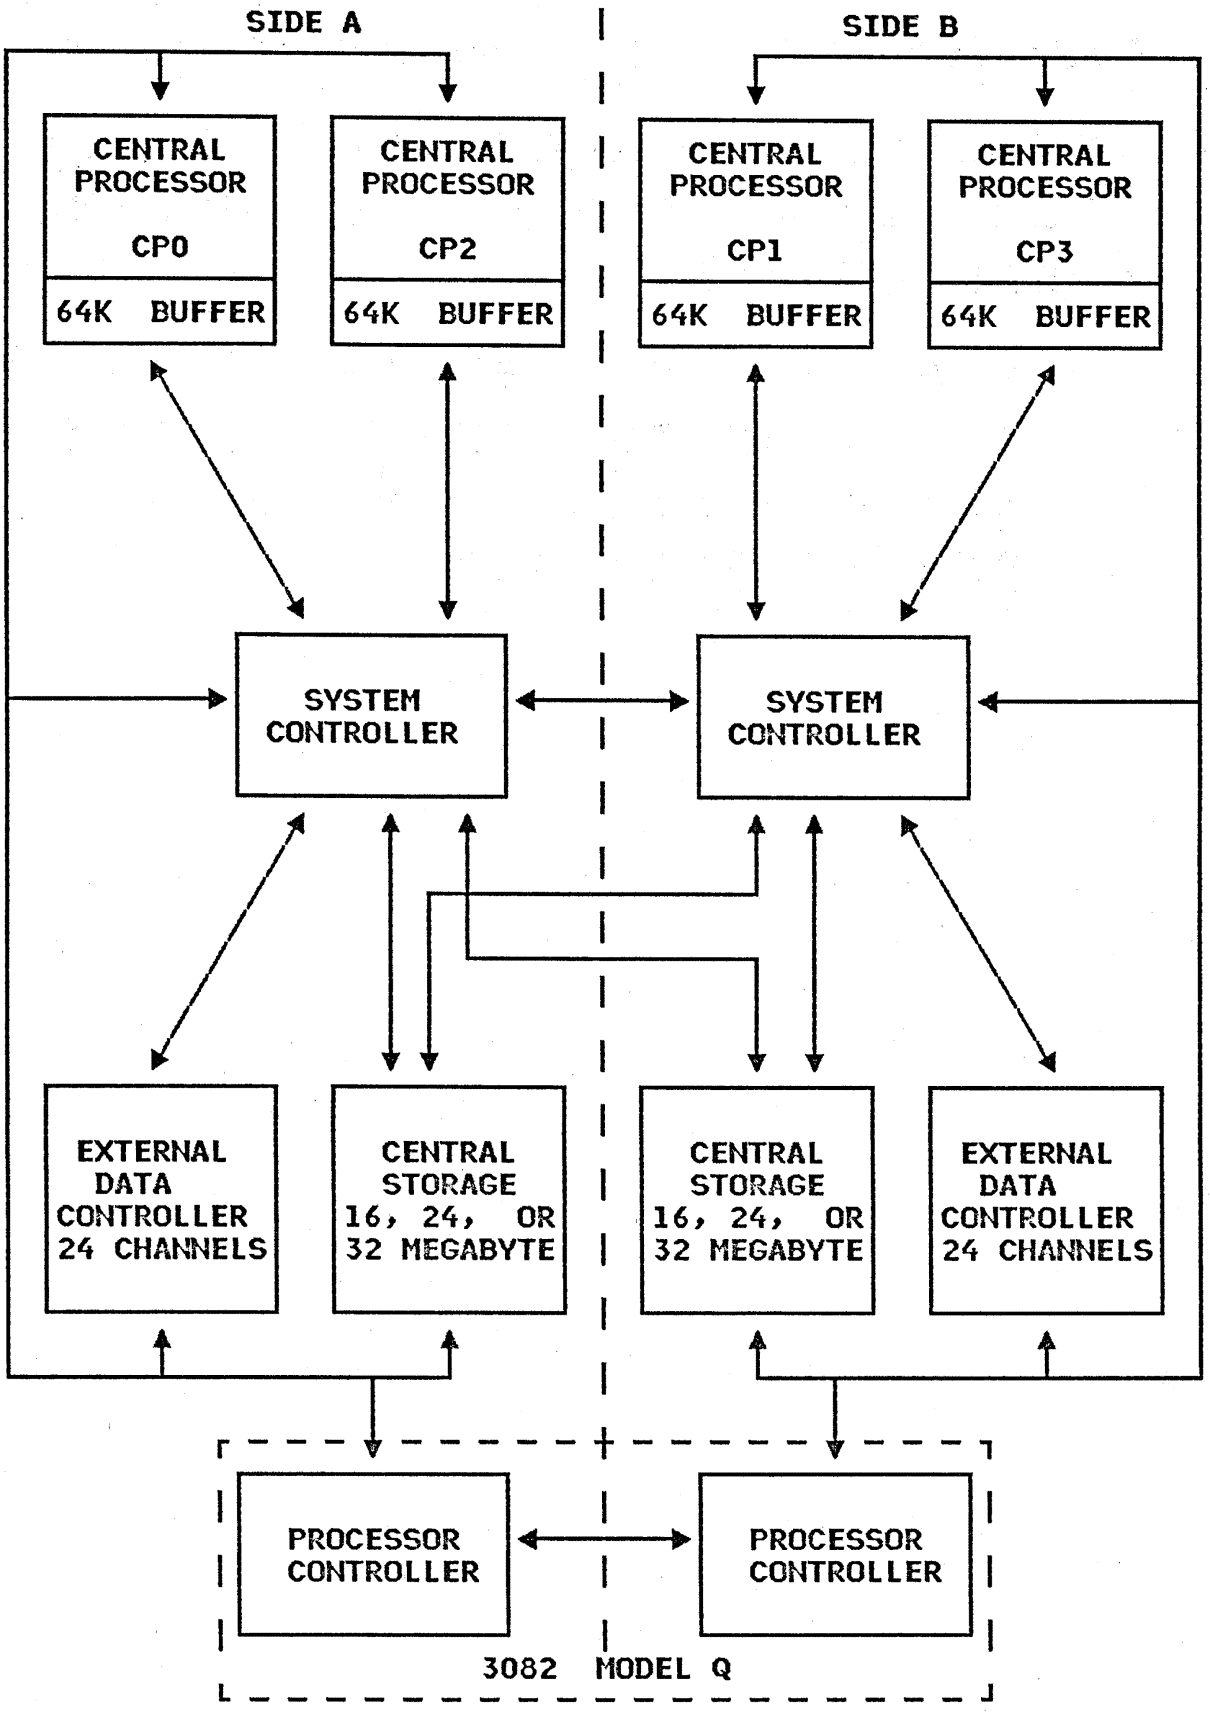

IBM 3084 PROCESSOR COMPLEX

3084 CHANNEL SUBSYSTEM FUNCTIONS

- MANAGES THE TRANSFER OF DATA BETWEEN I/O CONTROL UNITS AND CENTRAL STORAGE
  - 48 CHANNELS (24 PER SIDE)
    - UP TO EIGHT BYTE CHANNELS CAN BE DEFINED (4 PER SIDE
  - 3.0 MEG BLOCK CHANNELS
    - CAN ATTACH STANDARD AND DATA STREAMING DEVICES
  - 144 MEG AGGREGATE DATA RATE

IBM 3084 PROCESSOR COMPLEX

EXTERNAL DATA CONTROLLER (EXDC)

3084 CHANNEL SUBSYSTEM - 370-XA MODE

- DYNAMIC CHANNEL SUBSYSTEM
  - ALL CENTRAL PROCESSORS CAN ACCESS ALL DEVICES DIRECTLY
    - NO CHANNEL SETS
    - NO ASYMMETRIC DEVICES
  - I/O INTERRUPT CAN BE TAKEN BY ANY CENTRAL PROCESSOR
  - UP TO 8 INTERRUPT SUBCLASSES
  - CHANNEL PATH SELECTION
    - UP TO 4 PATHS
  - 3380 DYNAMIC PATH SELECTION
  - UP TO 4080 UCW'S
  - ALL NEW I/O INSTRUCTIONS
  - CHANNEL MONITORING FACILITY
- TERMINOLOGY
  - SUBCHANNEL/DEVICE
  - UCW
    - NO SHARED UCW'S

IBM 3084 PROCESSOR COMPLEX
3084 CHANNEL SUBSYSTEM STRUCTURE
370-XA MODE

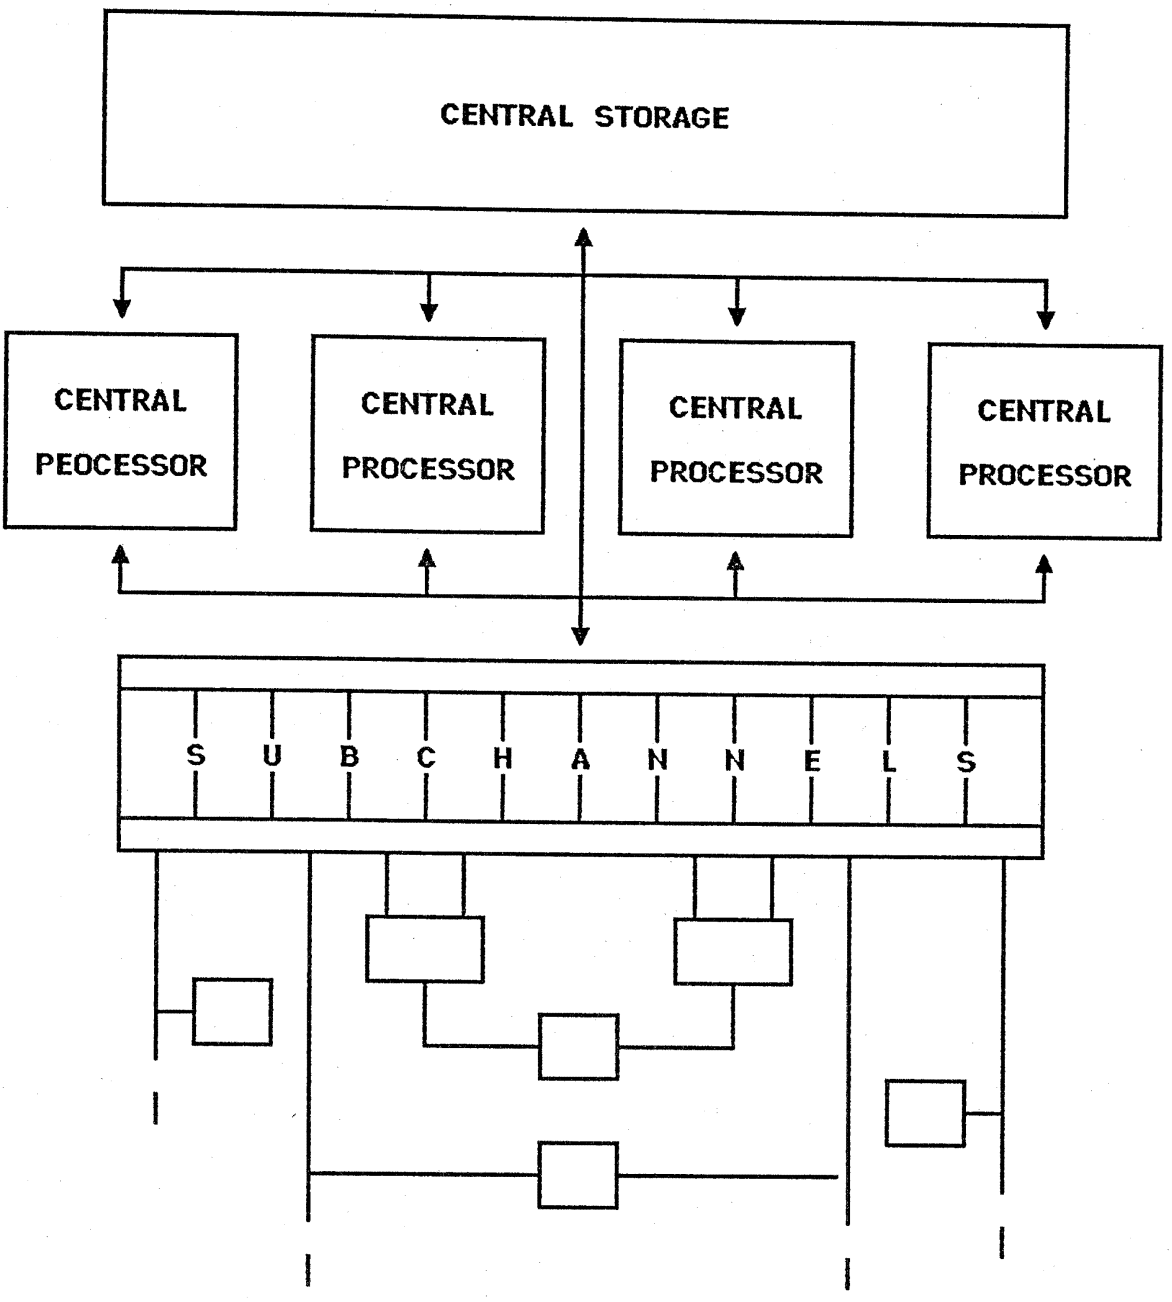

IBM 3084 PROCESSOR COMPLEX

370-XA MODE I/O QUEUING

- EXDC QUEUES I/O REQUESTS IF CHANNEL, CONTROL UNIT, OR DEVICE ARE BUSY
- ALL I/O REQUESTS ARE ELIGIBLE FOR QUEUING
- ONE REQUEST PER UCW
- OPERATION
  - CP PASSES REQUEST TO EXDC
  - EXDC ATTEMPTS TO START I/O
    - IF BUSY ON ALL PATHS, THE REQUEST IS QUEUED
  - REQUEST REMAINS QUEUED IN EXDC UNTIL THE BUSY CONDITION IS CLEARED
    - QUEUE TIME, DATA TRANSFER AND SERVICE TIME CAPTURED BY MONITORING FACILITY

IBM 3084 PROCESSOR COMPLEX

370-XA CHANNEL MONITORING FACILITY

- CAN DETERMINE DURATION OF SIGNIFICANT EVENTS DURING AN I/O OPERATION
- COLLECTED FOR EACH DEVICE
- DATA ACCUMULATED IN CENTRAL STORAGE AND IS USED BY RMF
  - DEVICE DATA COLLECTED
    - TOTAL START SUBCHANNELS
    - SAMPLE START SUBCHANNLS
    - PLUS:

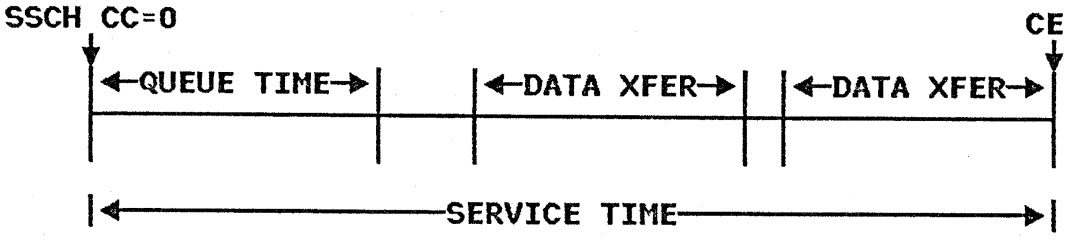

IBM 3084 PROCESSOR COMPLEX

370-XA CHANNEL MONITORING FACILITY

- LOGICAL CONTROL UNIT DATA COLLECTED
  - NUMBER OF ATTEMPTS TO START AN I/O OPERATION
  - NUMBER OF SUCCESSFUL ATTEMPTS TO START AN I/O OPERATION
  - NUMBER OF CU BUSY
  - NUMBER OF DEVICE BUSY
  - NUMBER REQUESTS QUEUED TO THE LOGICAL CONTROL UNIT

IBM 3084 PROCESSOR COMPLEX

370-XA MODE CHANNEL PATH SELECTION

- START SUBCHANNEL ISSUED BY SOFTWARE
  - INDICATES SUBCHANNEL (DEVICE) TO BE STARTED
- CHANNEL PATH SELECTION DONE BY MICROCODE AT LOGICAL CONTROL UNIT LEVEL
  - ROTATE
  - PREFERRED PATH
- MORE SIGNIFICANT FOR NON-3380 DASD DEVICES

IBM 3084 PROCESSOR COMPLEX

370-XA MODE CHANNEL PATH SELECTION

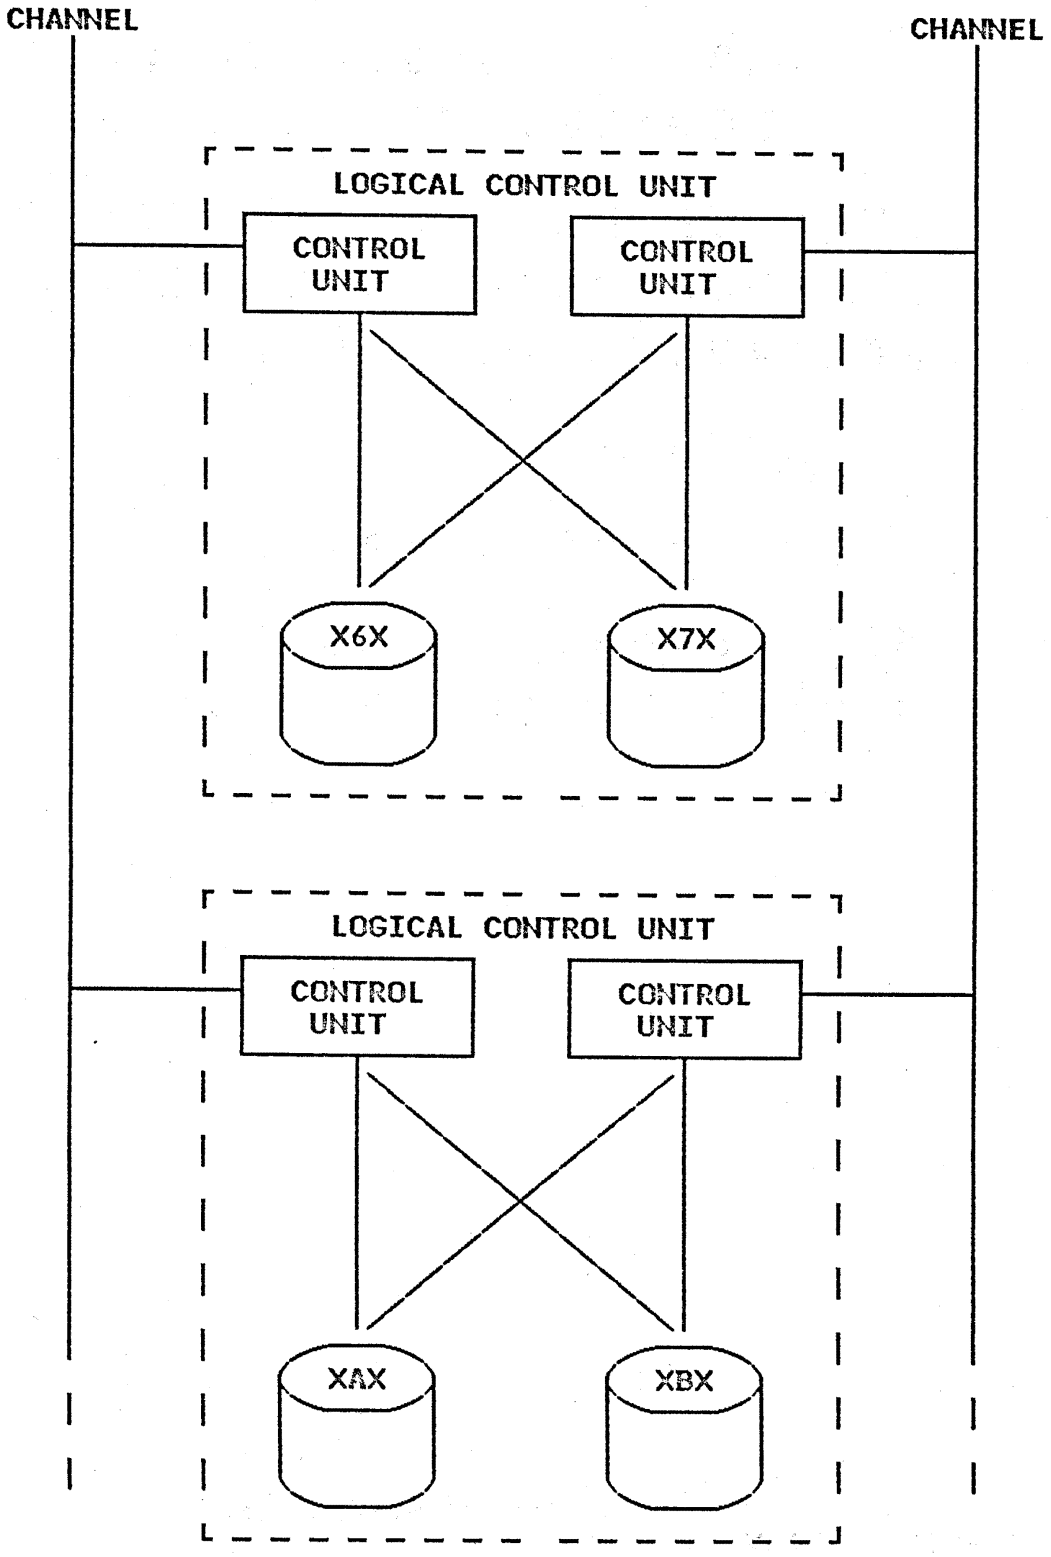

IBM 3084 PROCESSOR COMPLEX

EXTERNAL DATA CONTROLLER (EXDC)

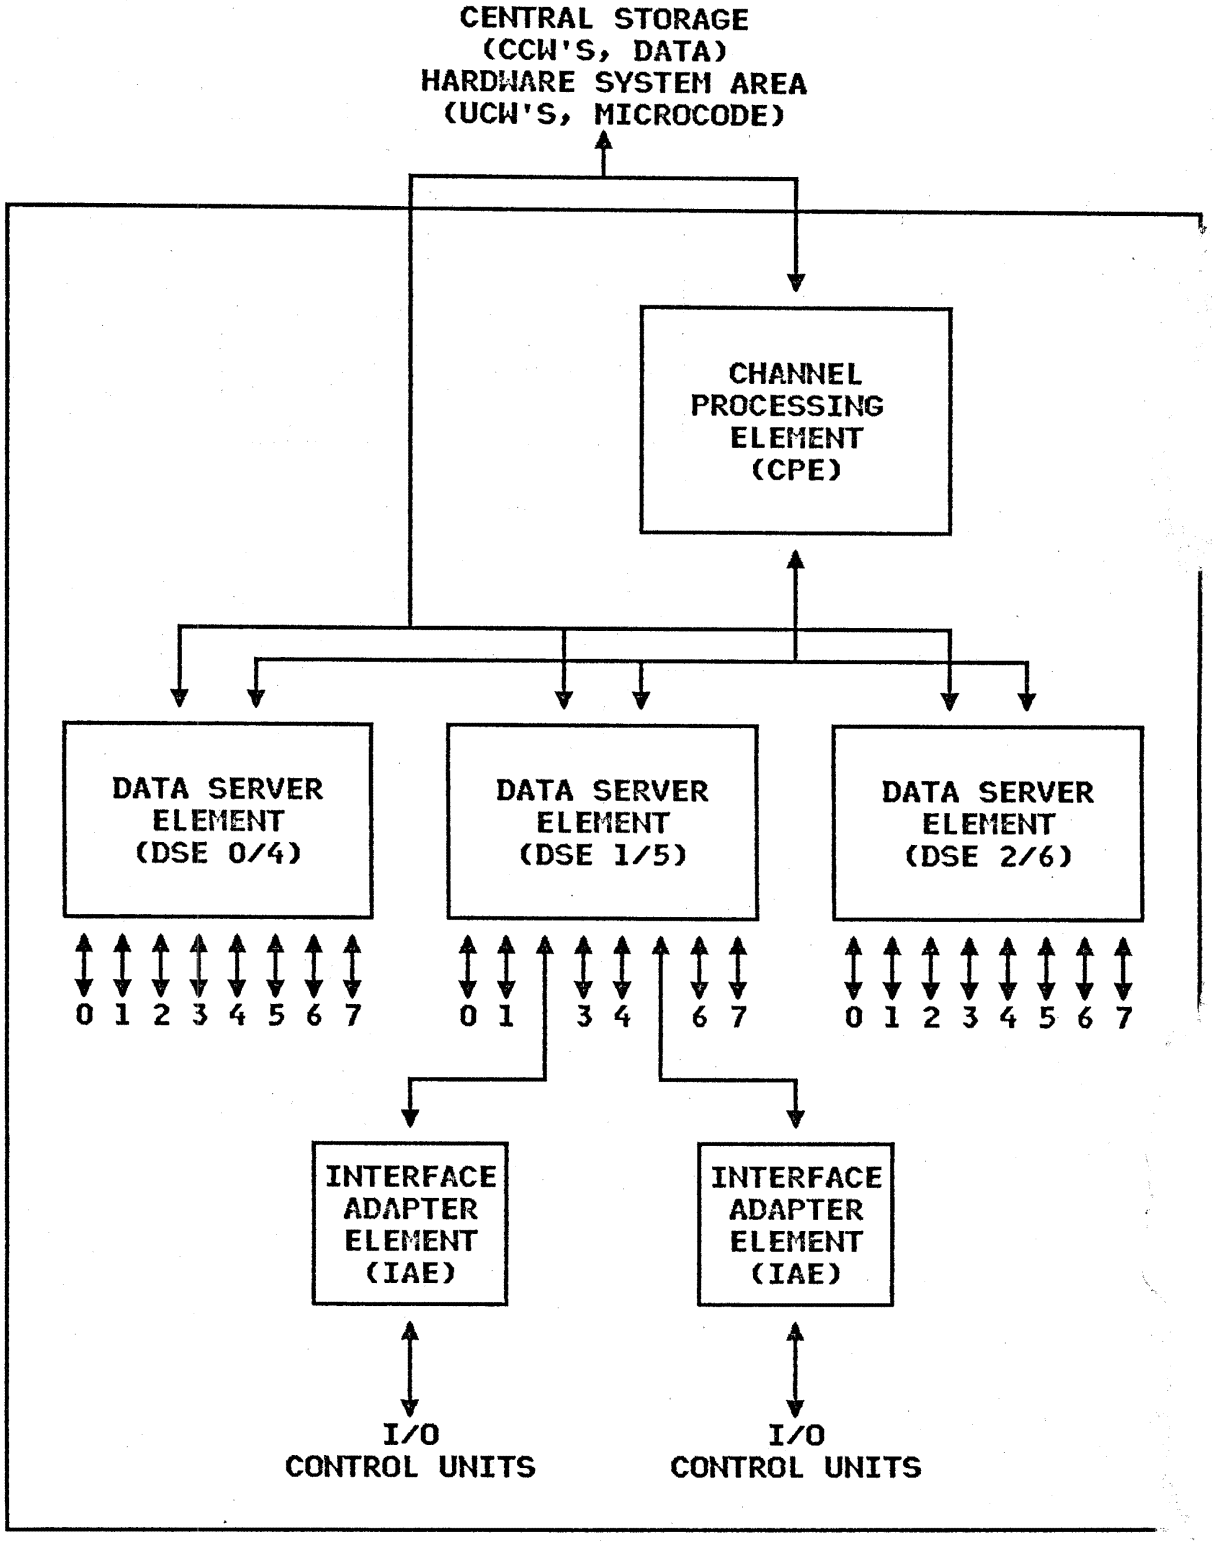

IBM 3084 PROCESSOR COMPLEX

3084 XA MODE CHANNEL PATH SELECTION

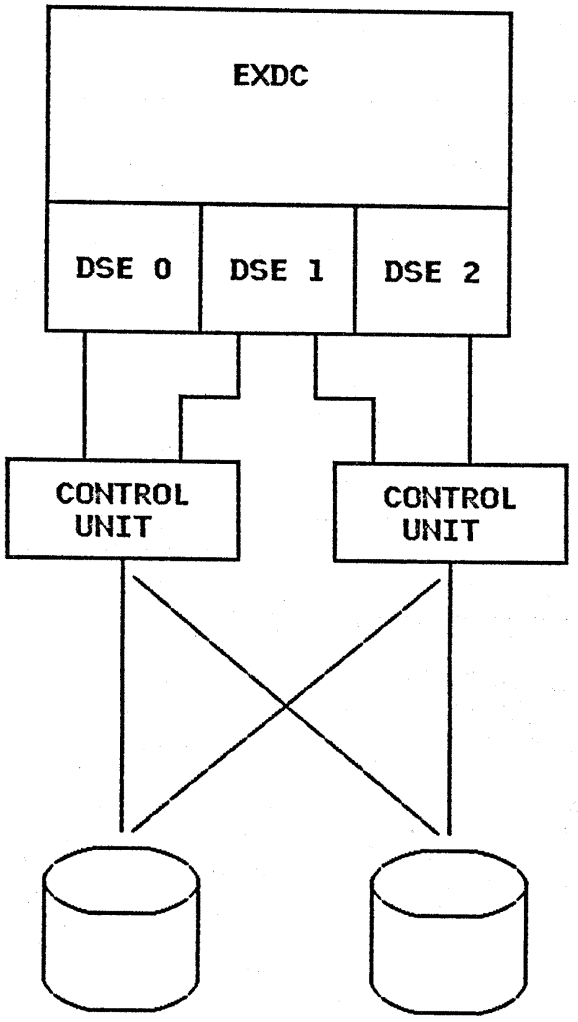

IBM 3084 PROCESSOR COMPLEX

3084 XA MODE CHANNEL PATH SELECTION

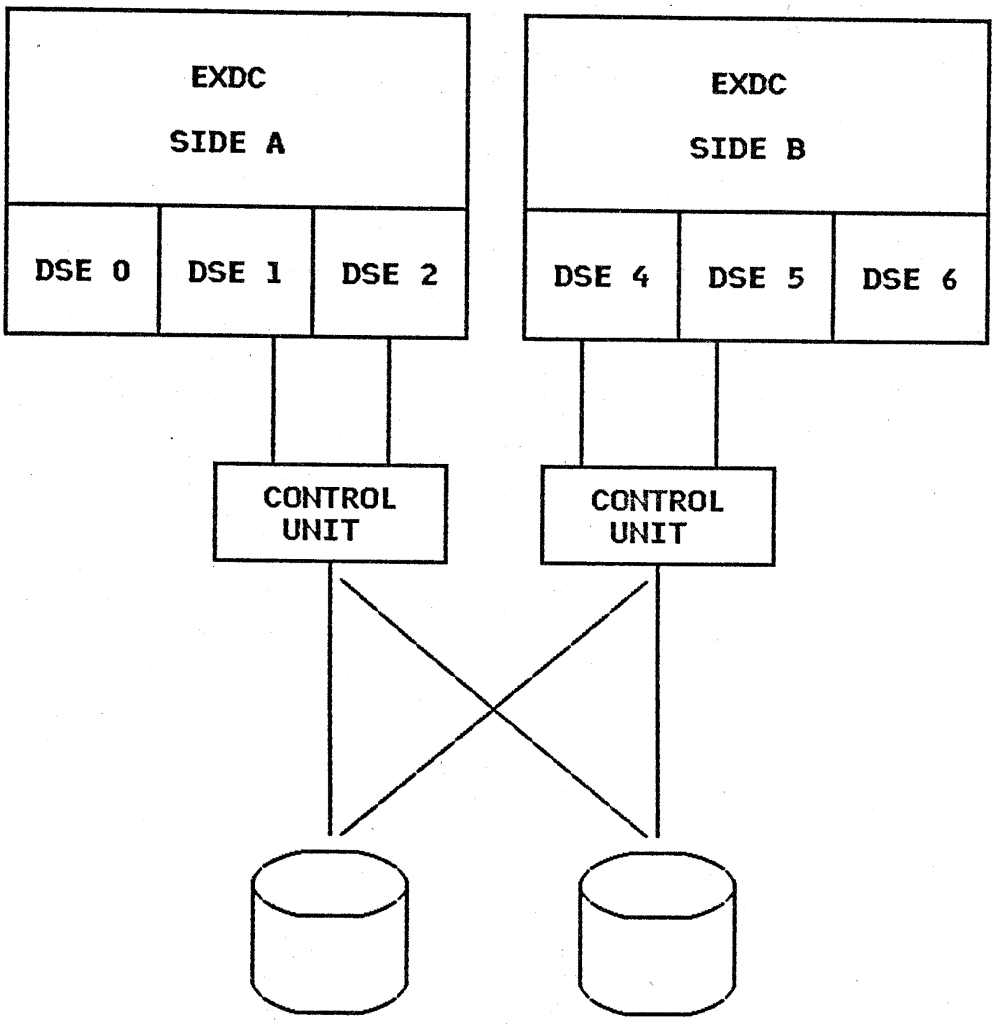

IBM 3084 PROCESSOR COMPLEX

3084 CHANNEL CONFIGURATION
SINGLE IMAGE MODE

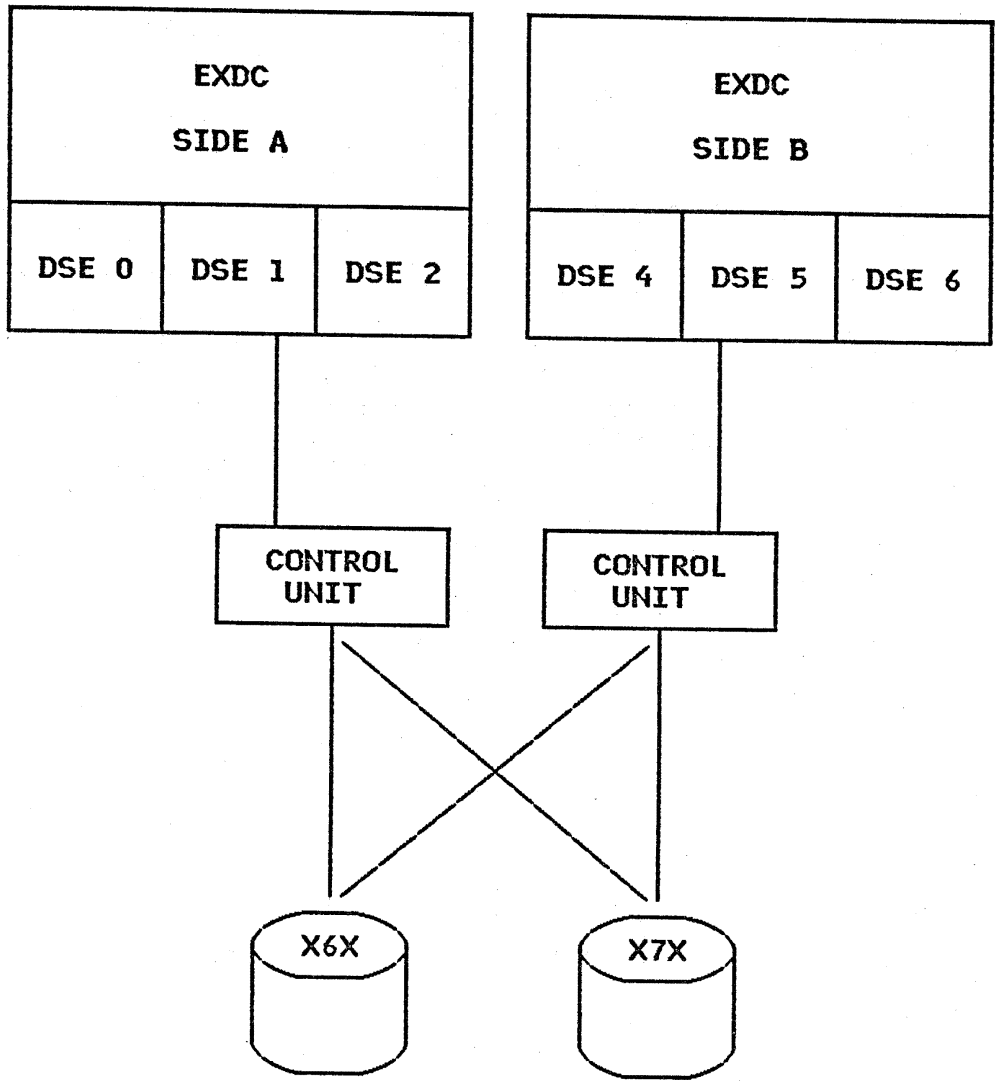

IBM 3084 PROCESSOR COMPLEX

3084 CHANNEL SUBSYSTEM

MIXED MODE PHYSICAL CONFIGURATION

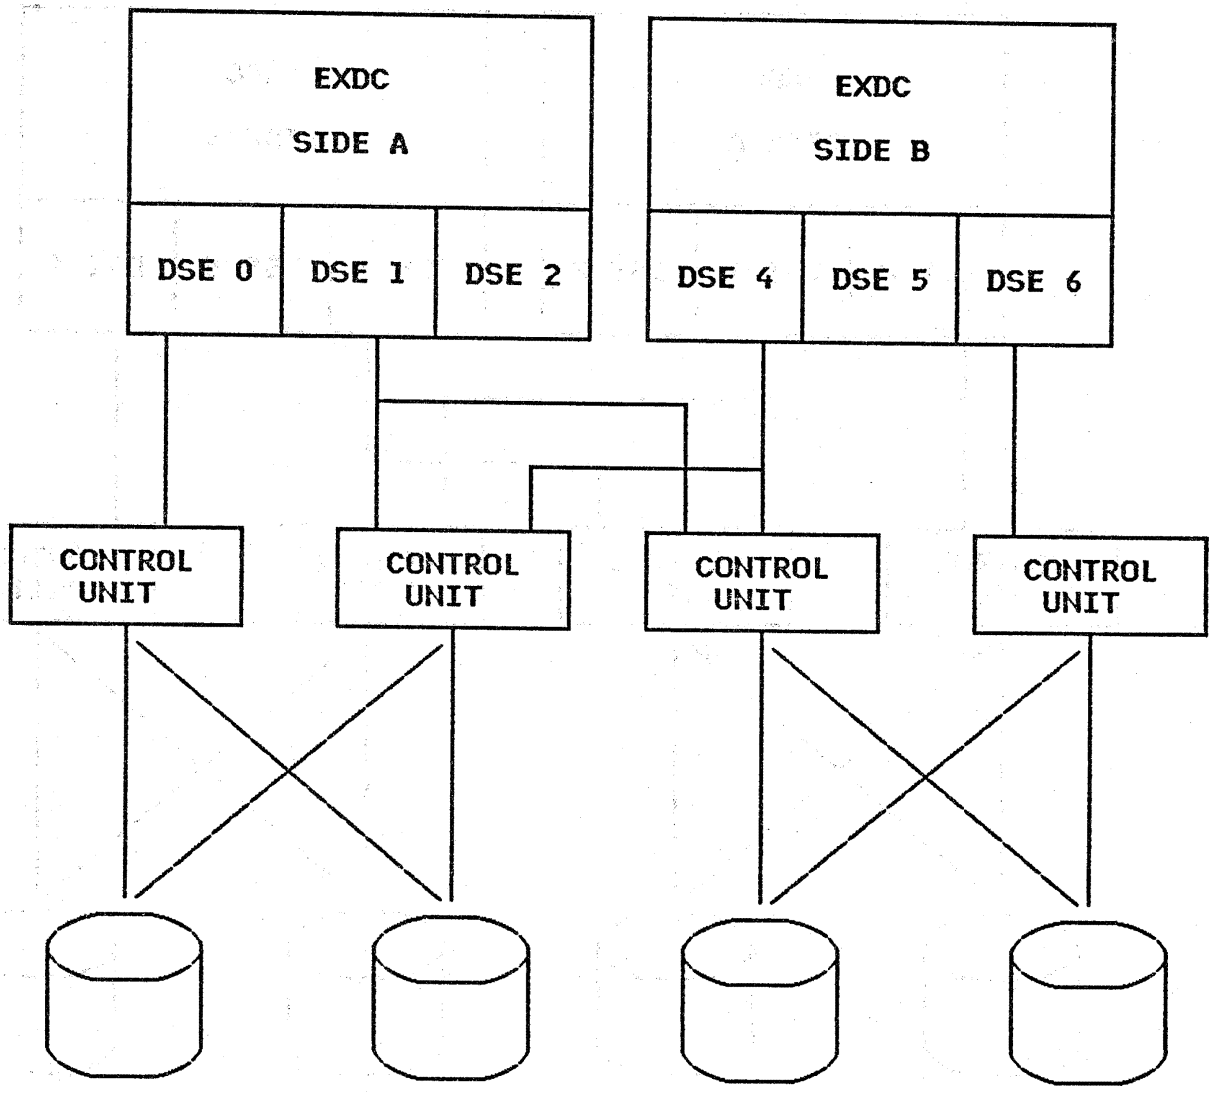

IBM 3084 PROCESSOR COMPLEX

3084 CHANNEL CONFIGURATION

MIXED MODE - RUNNING SINGLE IMAGE

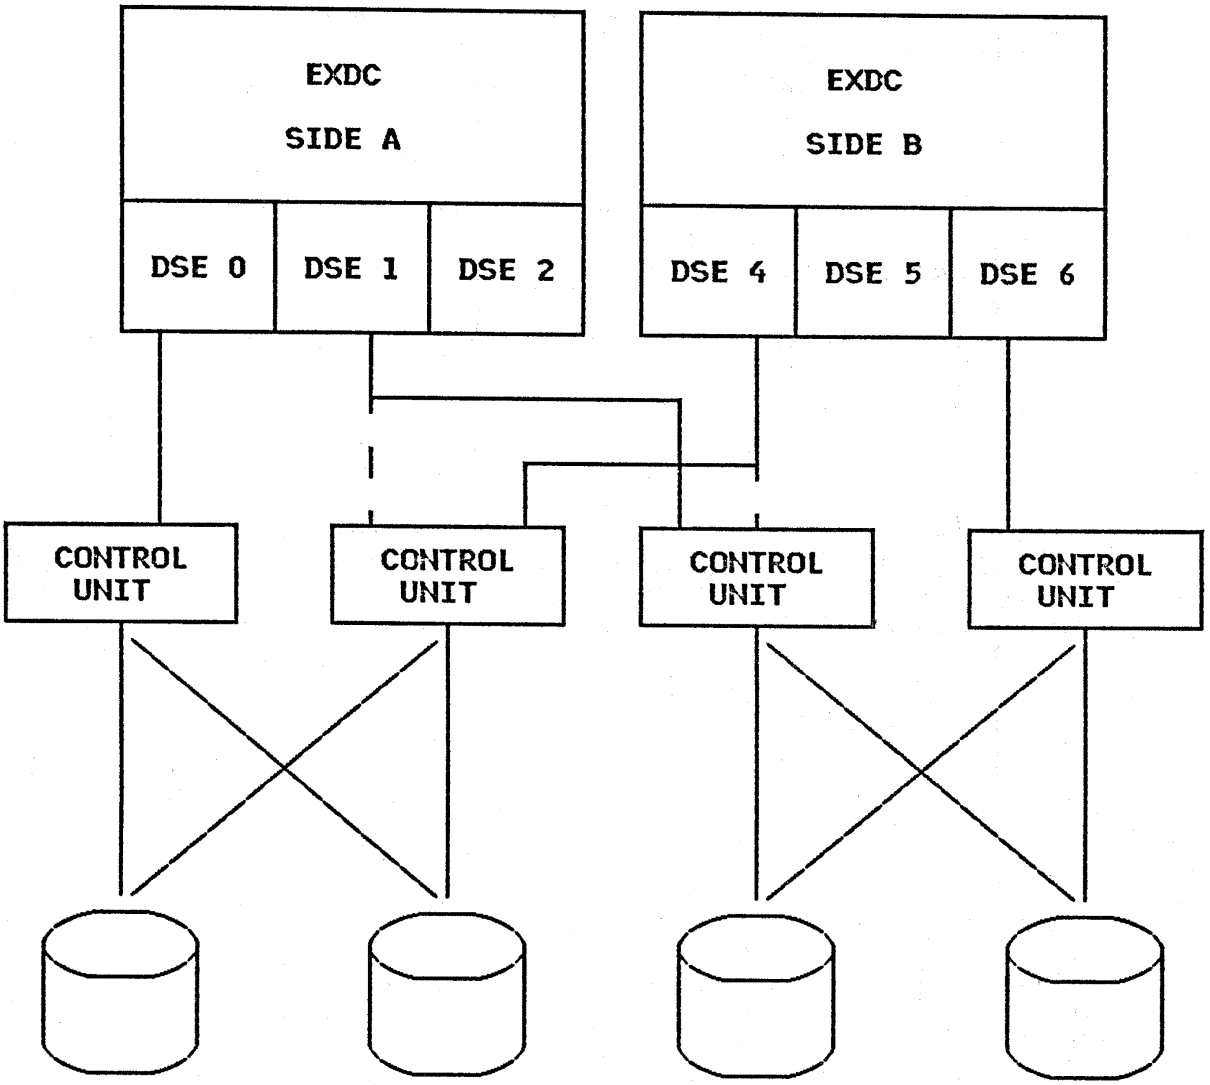

IBM 3084 PROCESSOR COMPLEX

3084 CHANNEL CONFIGURATION

MIXED MODE - RUNNING PARTITIONED

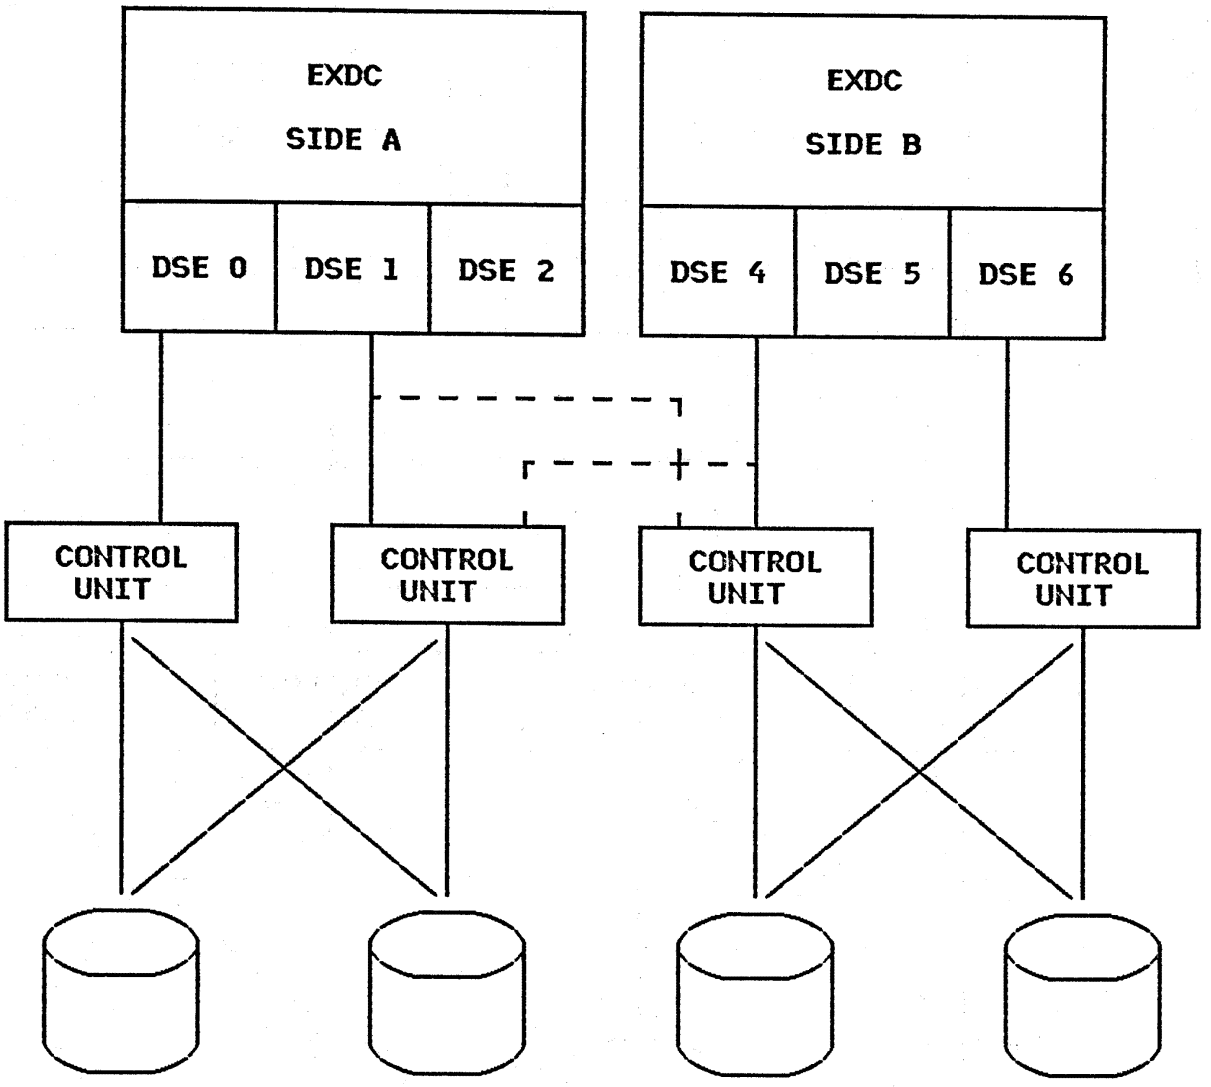

IBM 3084 PROCESSOR COMPLEX

3084 CHANNEL CONFIGURATION

SHARED MIXED MODE - SINGLE IMAGE

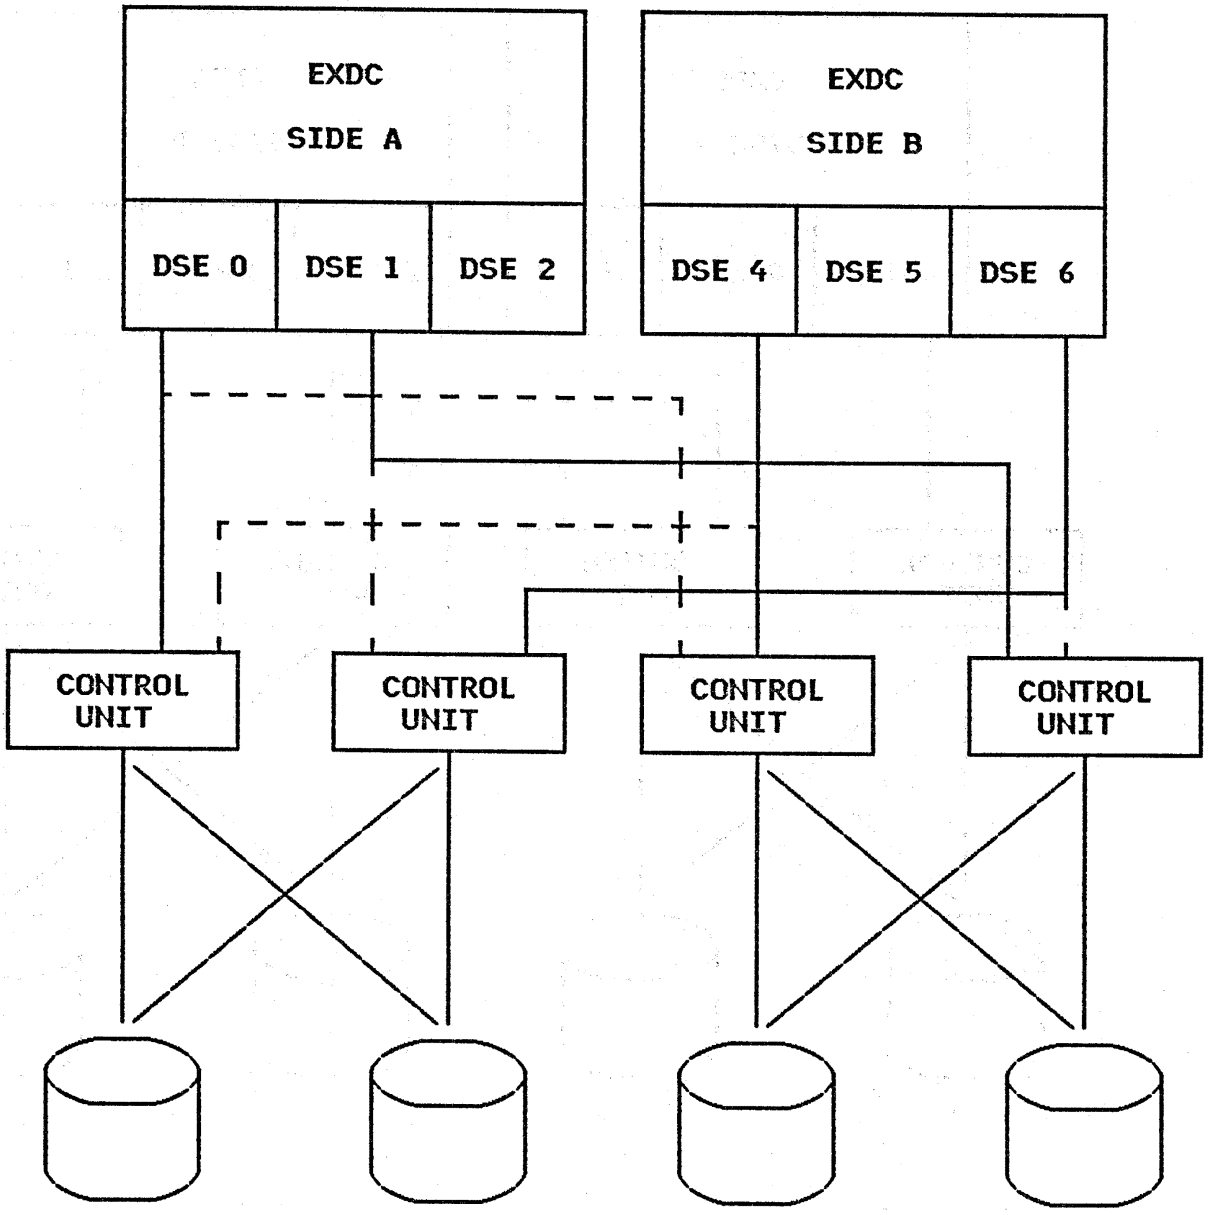

IBM 3084 PROCESSOR COMPLEX

3084 CHANNEL CONFIGURATION

SHARED MIXED MODE - PARTITIONED

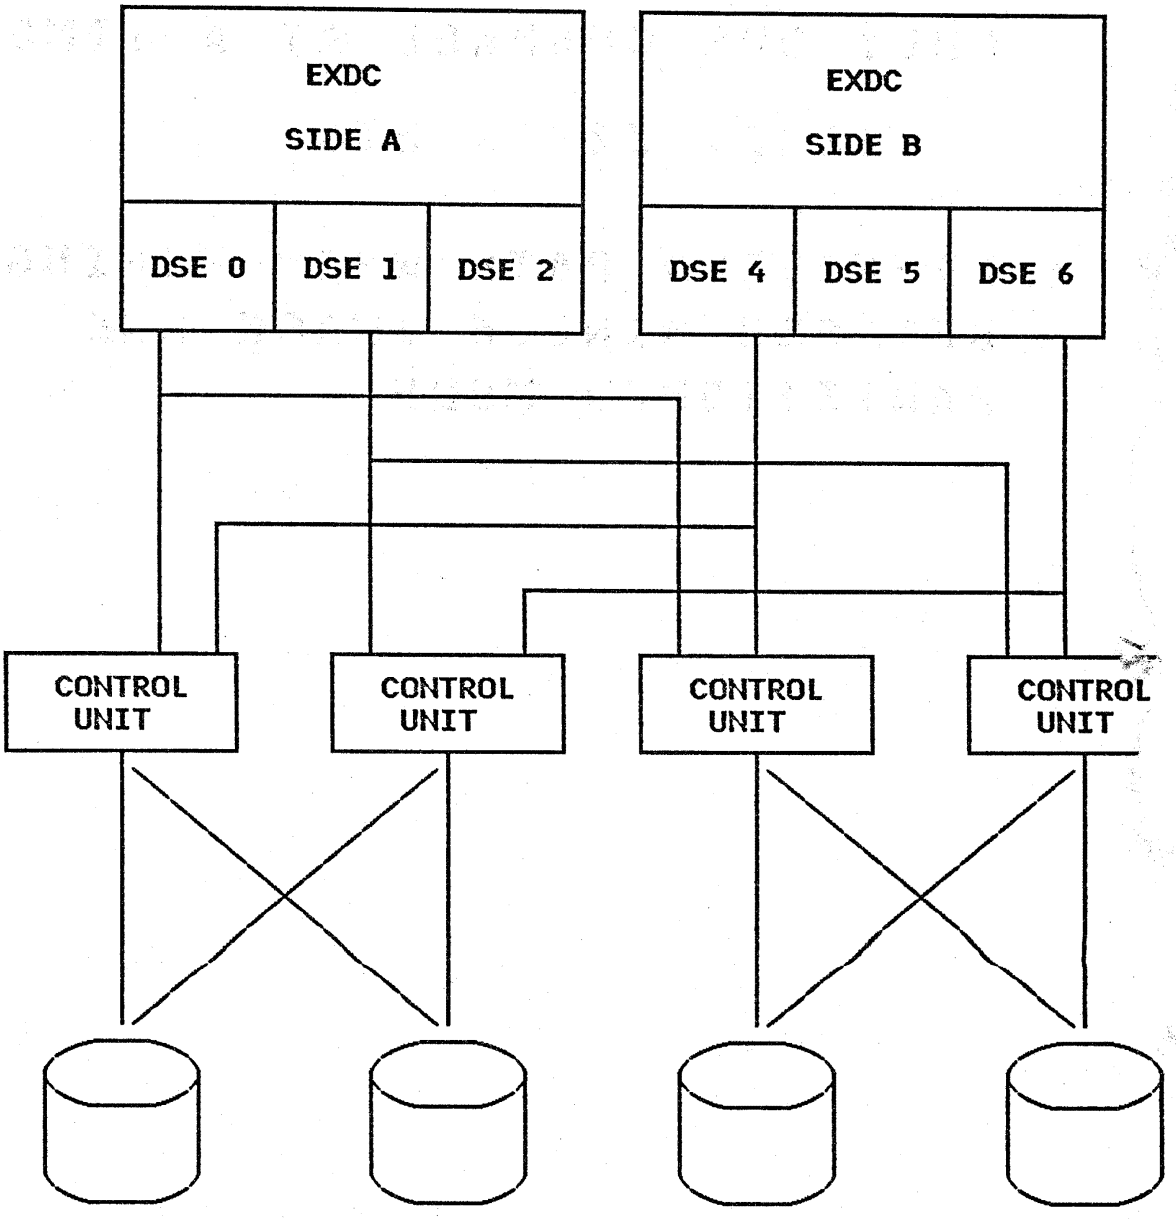

IBM 3084 PROCESSOR COMPLEX

3084 CHANNEL CONFIGURATION

ADDITIONAL CONSIDERATIONS

- CONTROL UNITS THAT CAN ATTACH TO ONLY ONE CHANNEL AT A TIME
  - 3272, 3274, ETC
- PREFERRED PATH WHEN MOVING BETWEEN SINGLE IMAGE AND PARTITIONED MODE

IBM 3084 PROCESSOR COMPLEX

3084 CHANNEL CONFIGURATION

ADDITIONAL CONSIDERATIONS

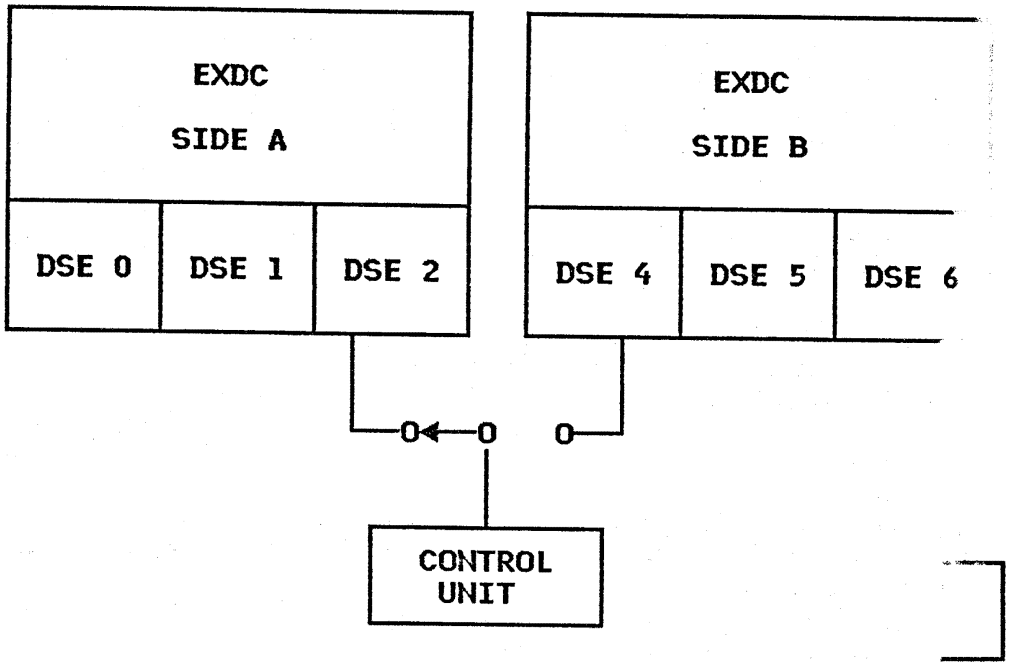

IBM 3084 PROCESSOR COMPLEX

3084 CONSOLES PER SIDE

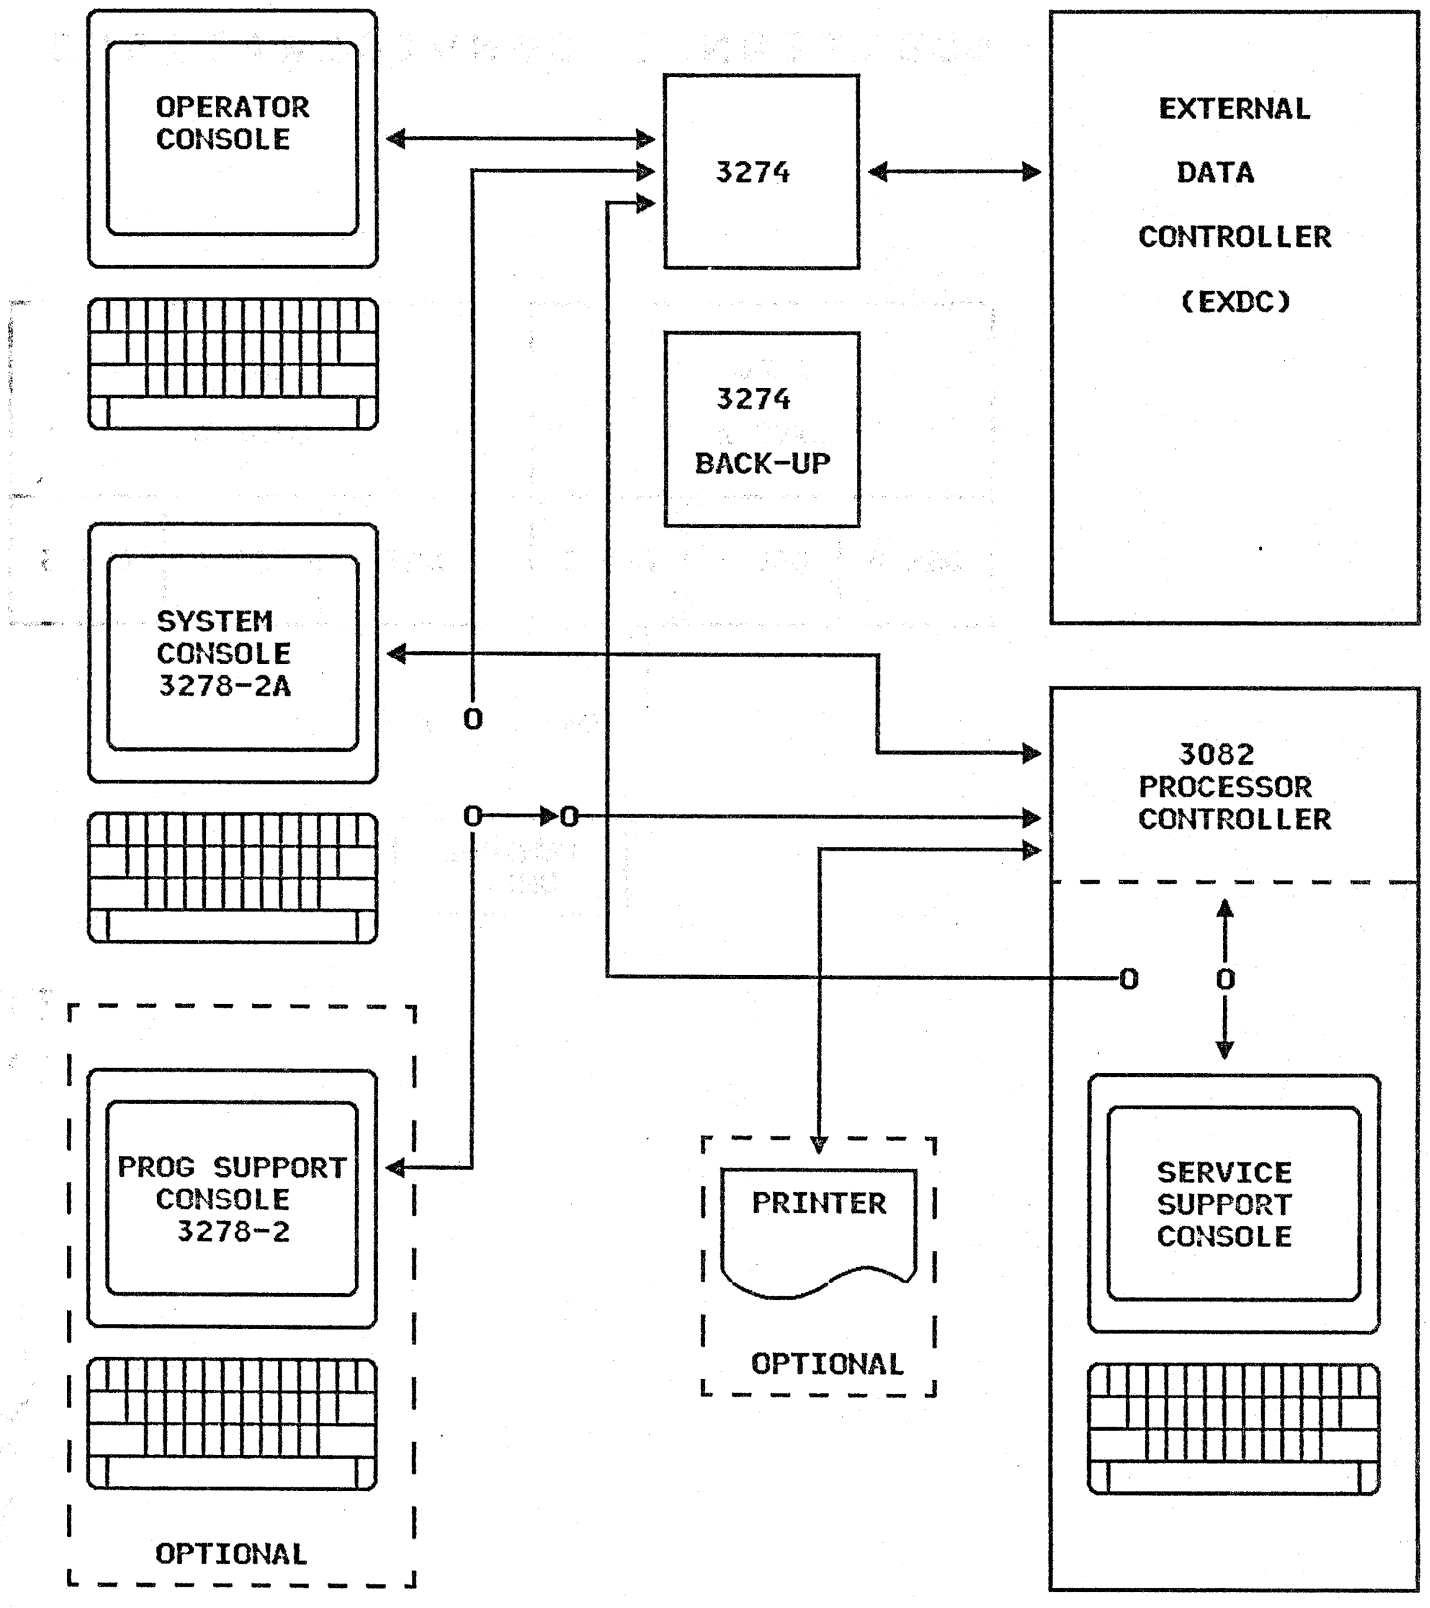

IBM 3084 PROCESSOR COMPLEX
3084 CONSOLES

- OPERATOR'S CONSOLE
  - SOFTWARE INTERFACE
  - ANY CONSOLE MVS OR VM SUPPORT
  - IF 3278 OR 3279, CAN BE UP TO 1500 METERS FROM PROCESSOR
- SYSTEM CONSOLE
  - SYSTEMS FUNCTIONS
    - IPL, ALTER/DISPLAY, CONFIGURATION FRAMES, ETC
  - SYSTEM ACTIVITY DISPLAY
  - NO INTERFACE TO SOFTWARE
  - 3278 MODEL 2A
    - UP TO 1500 METERS FROM PROCESSOR
- SERVICE SUPPORT CONSOLE
  - CE MAINTENANCE CONSOLE
  - CAN PERFORM FUNCTIONS OF BO SYSTEM AND OPERATOR'S CONSOL
  - SIMILAR TO 3278 MODEL 2A
- PROGRAMMING SUPPORT CONSOLE
  - RETAIN ACCESS
  - OPTIONAL
- SYSTEM/SERVICE SUPPORT CONSOLE HARDCOPY PRINTER
  - OPTIONAL

IBM 3084 PROCESSOR COMPLEX
3084 CONSOLES - SINGLE IMAGE MODE

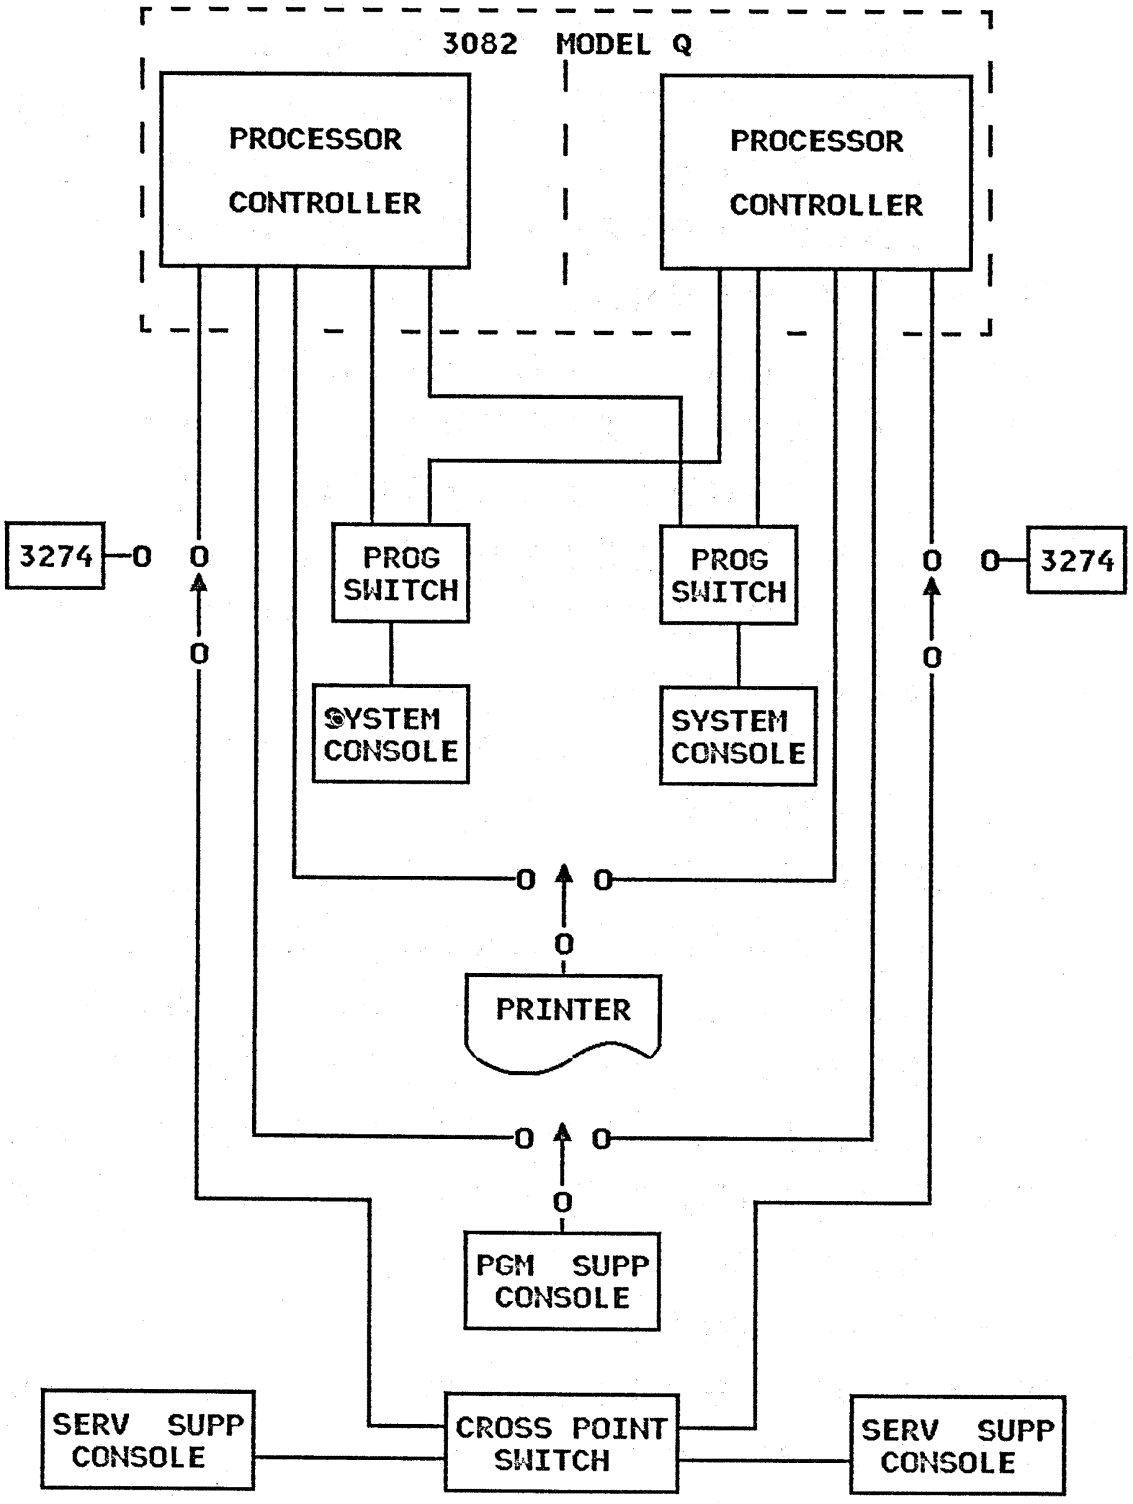

IBM 3084 PROCESSOR COMPLEX
3084 CONSOLES - PARTITIONED MODE

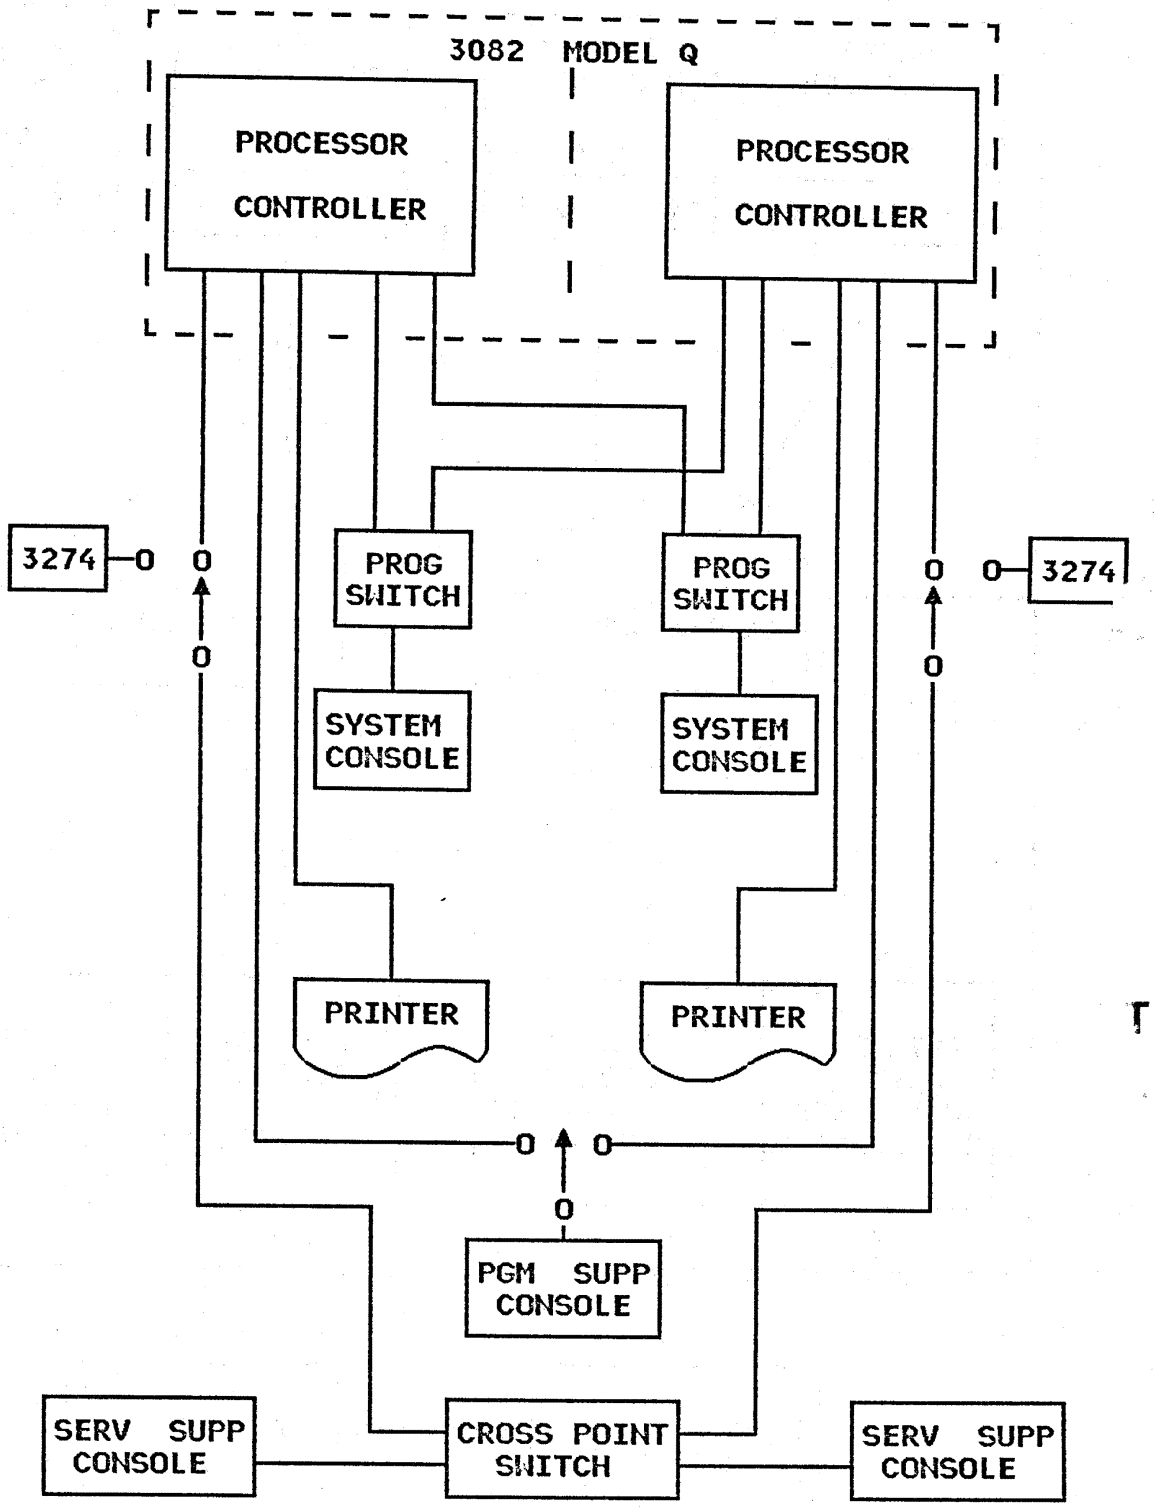

IBM 3084 PROCESSOR COMPLEX

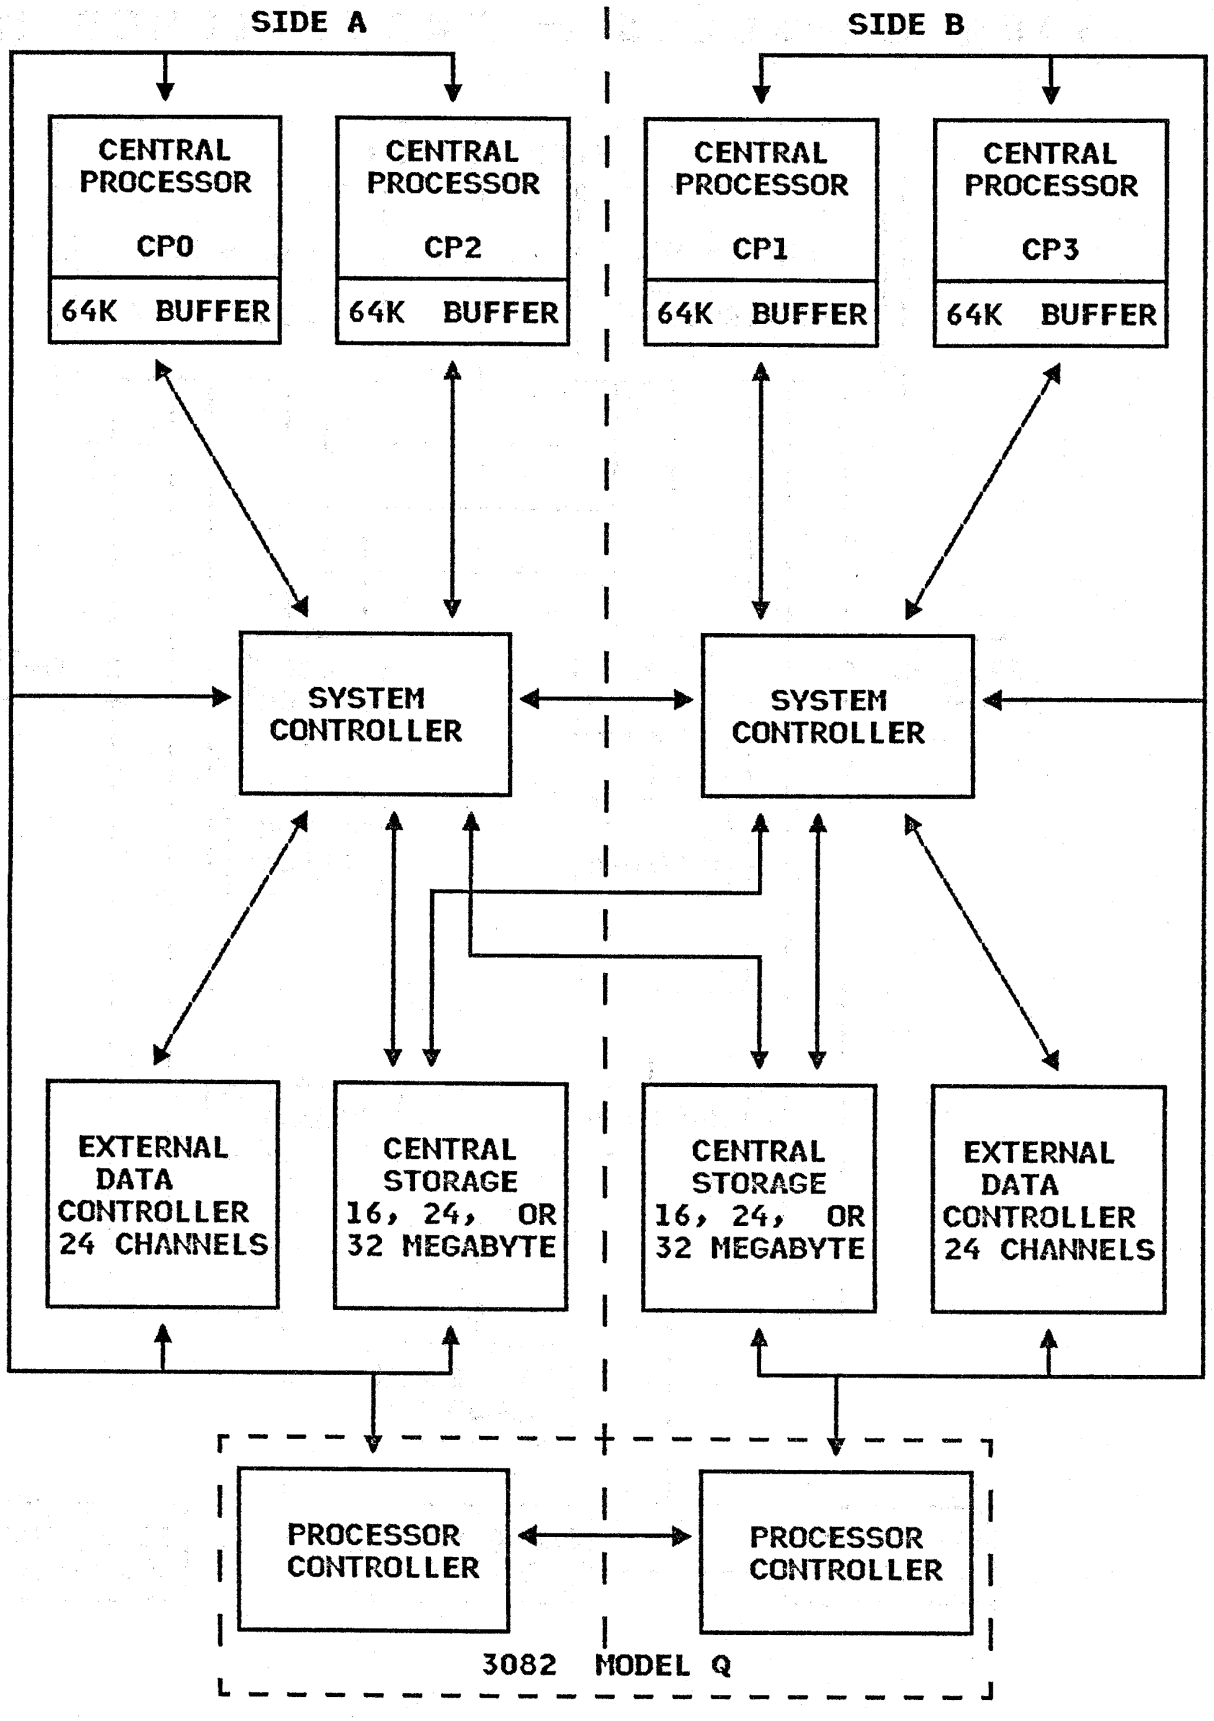

IBM 3084 PROCESSOR COMPLEX

3082 MODEL Q

- MONITORING AND SYSTEM SUPPORT FACILITY (MSSF)
  - POWER SEQUENCING, MONITORING AND CONTROL
  - MICROCODE LOADING AND INITIALIZATION
  - CONFIGURATION CONTROL
    - LOGICAL AND PHYSICAL
  - ERROR HANDLING
    - RECOVERY
    - RECORDING
    - DATA ANALYSIS
  - INTERACTIVE CE INTERFACE
- ENHANCED 3082 RECOVERY IN SINGLE IMAGE MODE

IBM 3084 PROCESSOR COMPLEX

3082 MODEL Q COMPONENTS PER SIDE

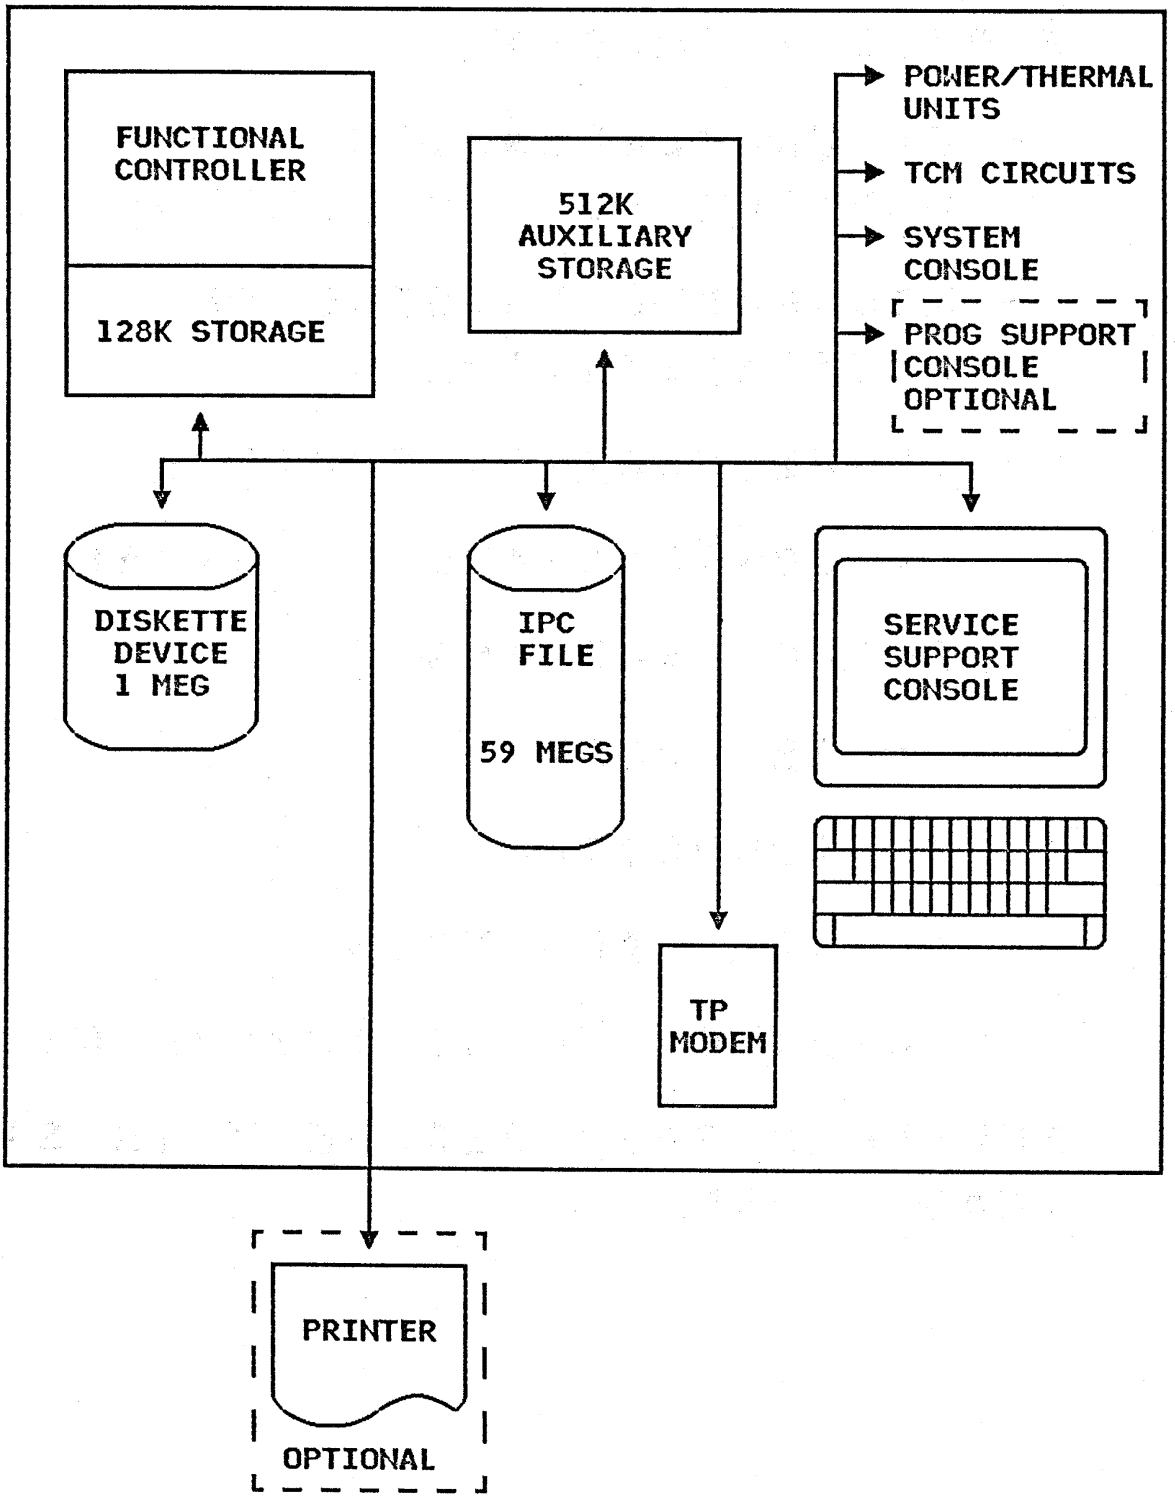

IBM 3084 PROCESSOR COMPLEX
3082 Q - SINGLE IMAGE MODE

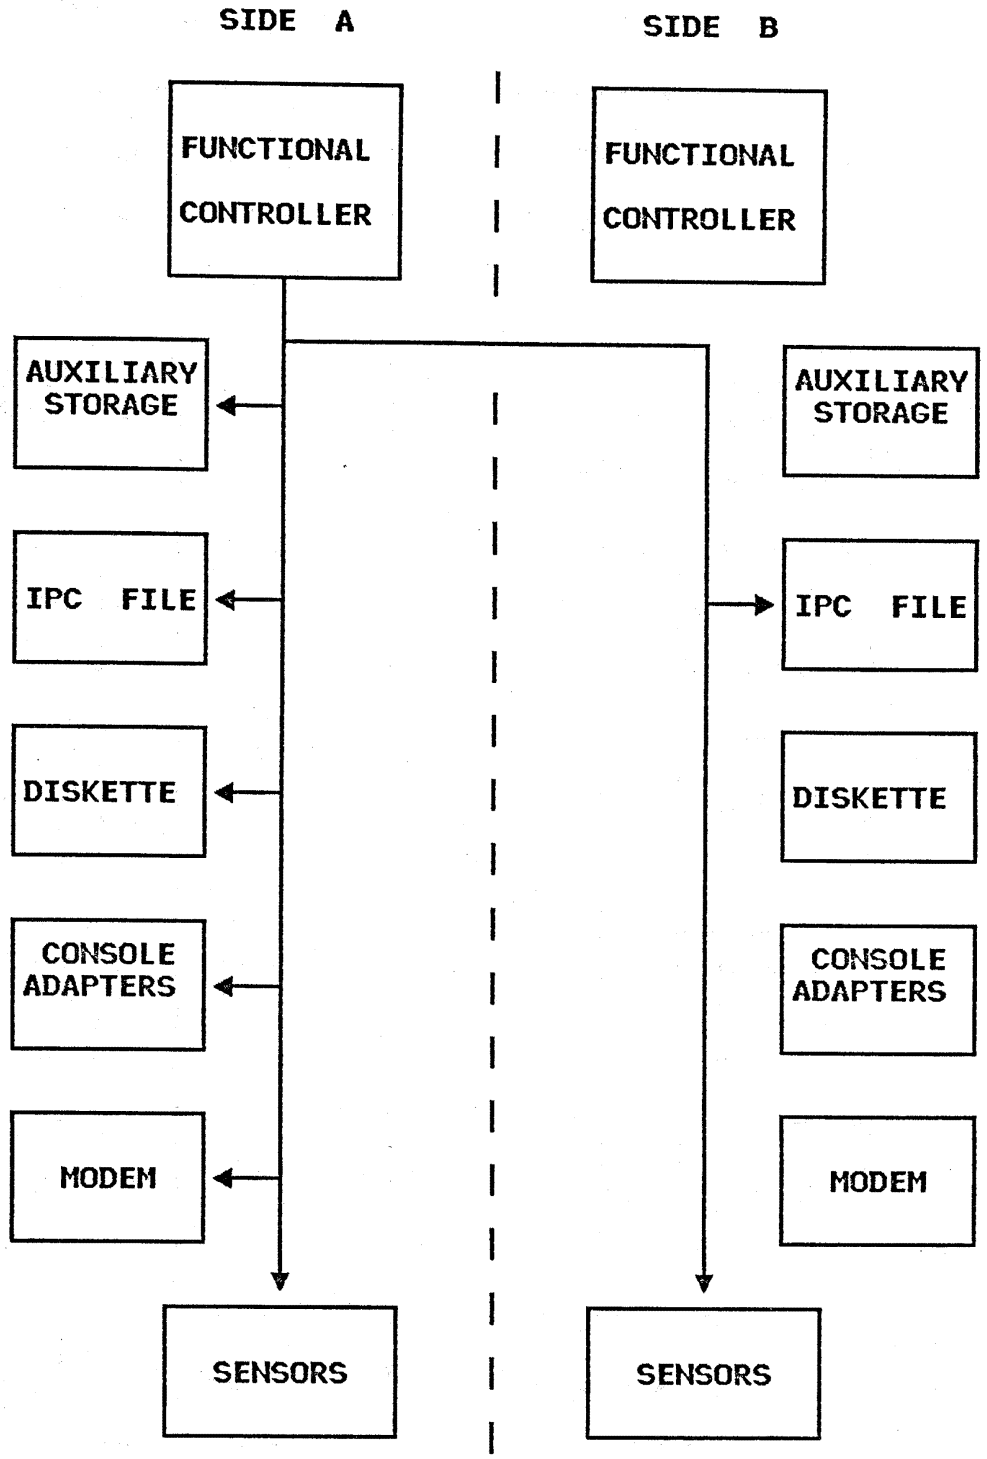

IBM 3084 PROCESSOR COMPLEX
3082 Q - PARTITIONED MODE

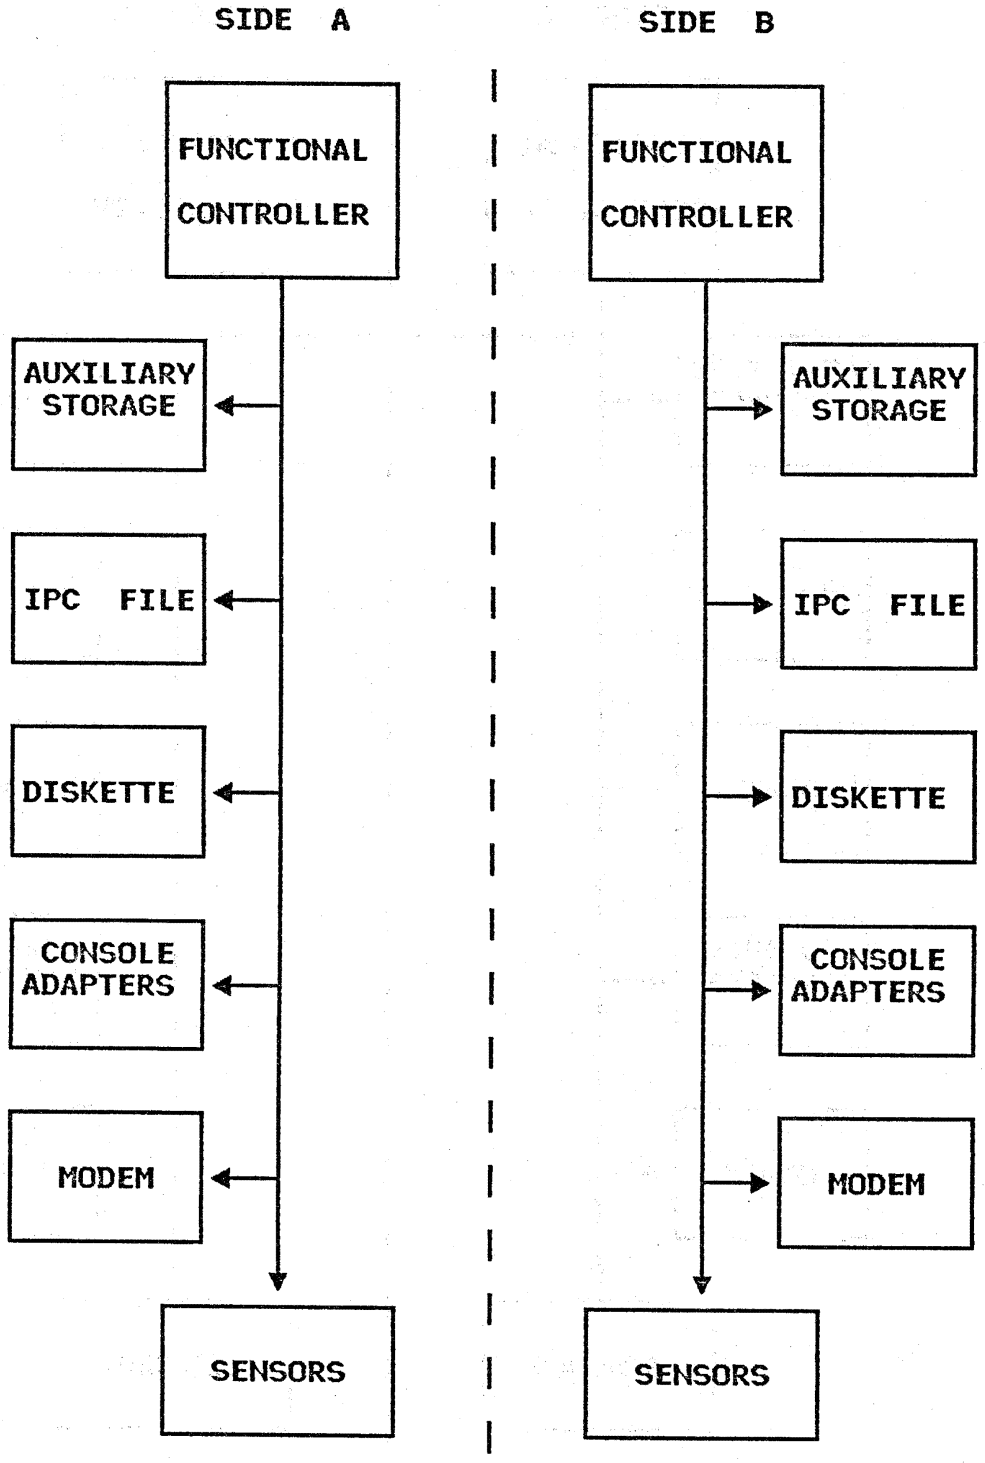

IBM 3084 PROCESSOR COMPLEX

3082 Q SINGLE IMAGE ERROR RECOVERY

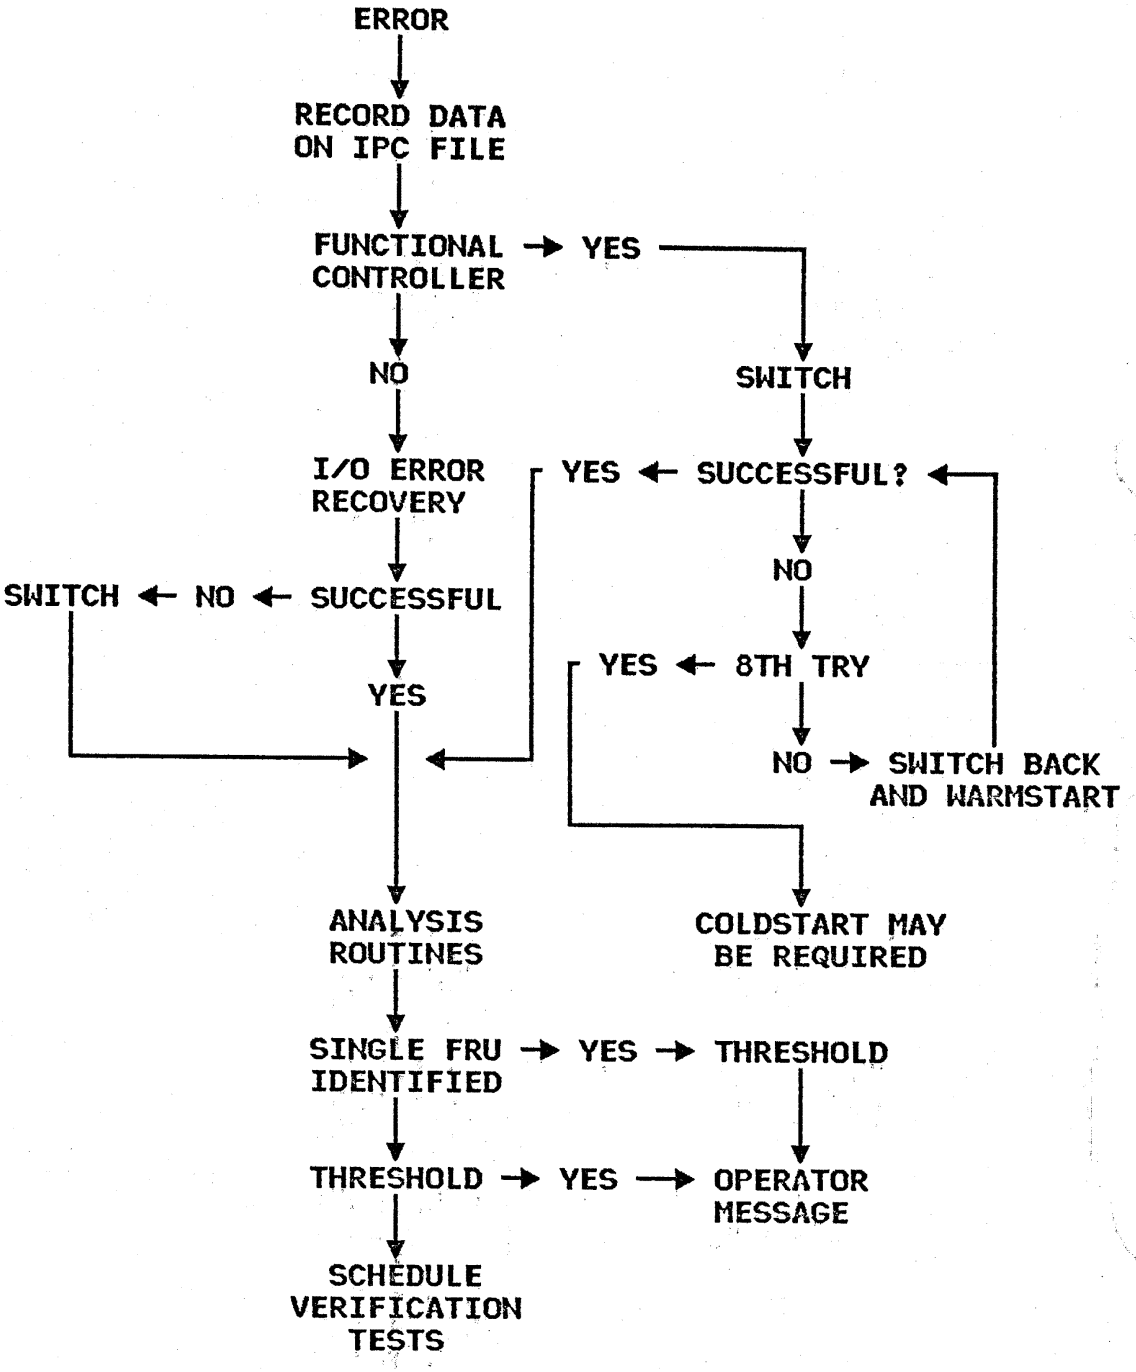

IBM 3084 PROCESSOR COMPLEX
3082 Q - SINGLE IMAGE MODE

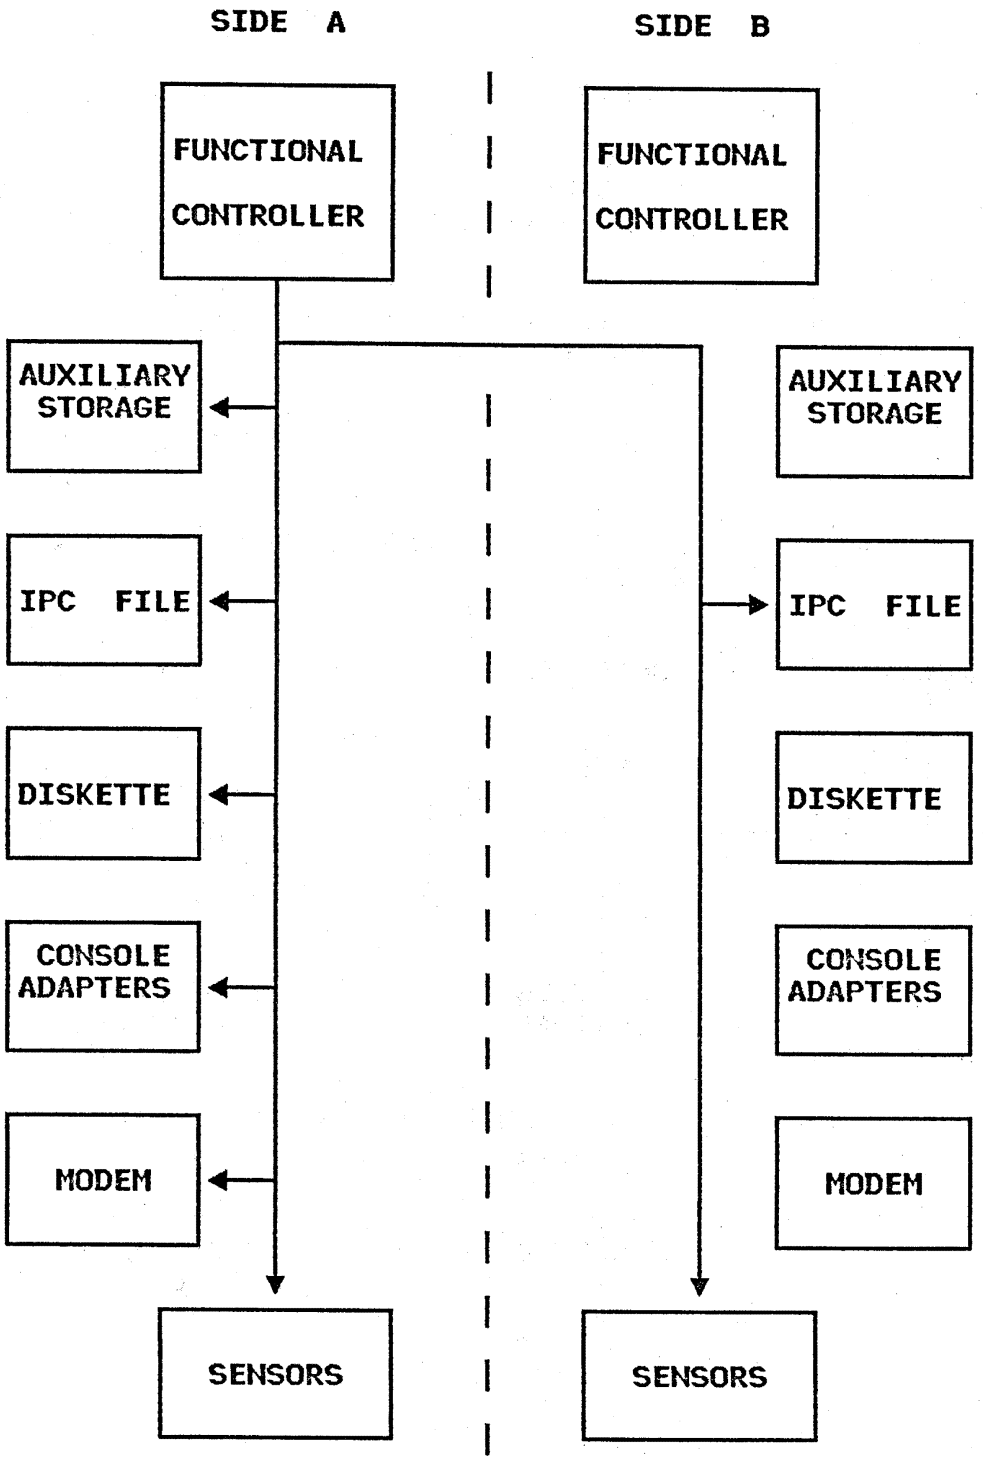

THIS PAGE INTENTIONALLY LEFT BLANK

APPENDIX D. FOIL HANDOUT MASTERS

The following pages contain the handout masters. Note that this material is unclassified and copyrighted. If you wish to provide handouts to your audience, the handout masters may be reproduced and distributed. However, no permission is granted to non-IBM employees to produce and/or distribute this material. Only the foils are unclassified. The presentation text is 'IBM Internal Use Only' and therefore cannot be distributed to non-IBM employees.

IBM
NA DIVISION
WASHINGTON
SYSTEMS
CENTER

AN INTRODUCTION TO THE 3084 PROCESSOR COMPLEX STRUCTURE AND CAPABILITIES

IBM 3084 PROCESSOR COMPLEX

TOPICS

- 3084 PROCESSOR OVERVIEW
- 3084 OPERATING MODES
- 3084 CHANNEL SUBSYSTEM
- 3084 CONSOLE CONFIGURATIONS
- 3082 PROCESSOR CONTROLLER

IBM 3084 PROCESSOR COMPLEX

3084 OVERVIEW

- FOUR INTEGRATED CENTRAL PROCESSORS
  - 26NS CYCLE
- 32, 48, OR 64 MEGABYTES OF REAL STORAGE
- 48 CHANNELS
  - DATA STREAMING STANDARD ON ALL BLOCK CHANNELS
  - MAXIMUM 8 BYTE CHANNELS
- OPERATING MODES
  - SINGLE IMAGE
    - MVS/XA ONLY
  - PARTITIONED
    - MVS/XA, MVS/SP OR VM/SP

IBM 3084 PROCESSOR COMPLEX

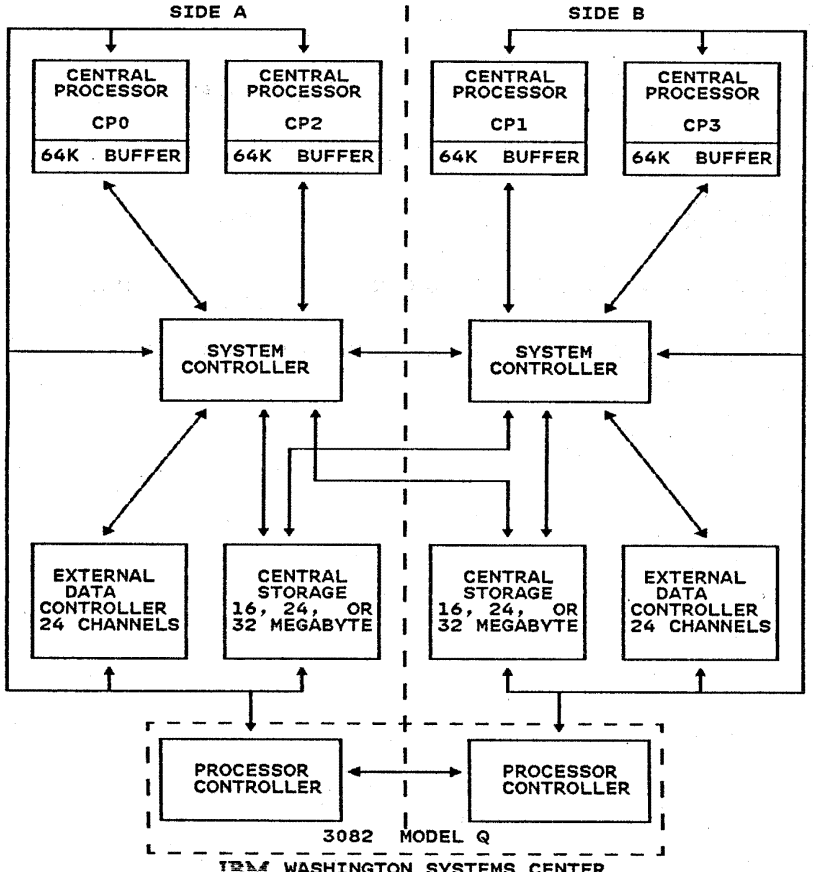

IBM 3084 PROCESSOR COMPLEX

3084 OVERVIEW - UPGRADEABILITY

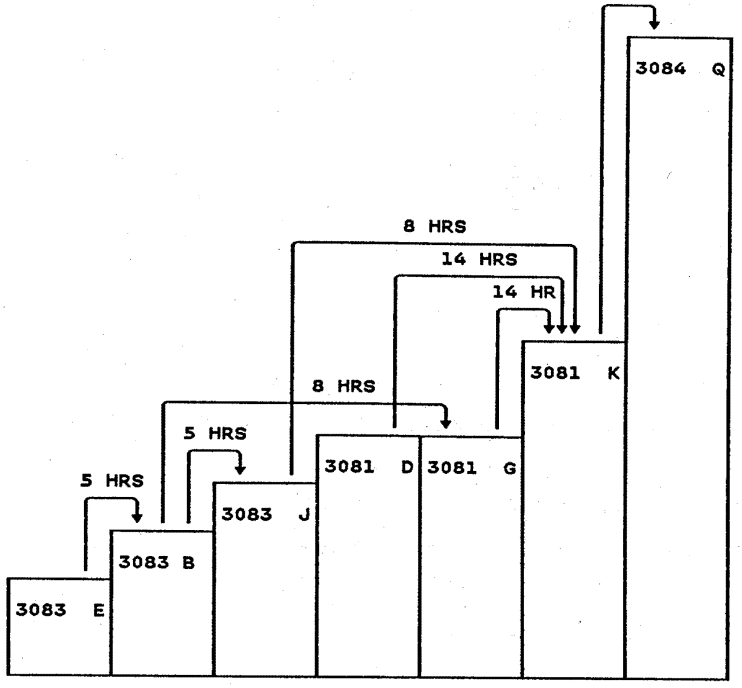

IBM 3084 PROCESSOR COMPLEX

3084 OVERIVEW

- STANDARD FEATURES
  - S/370 EXTENDED FACILITY
  - 3033 EXTENSION FEATURE
  - EXTENDED ADDRESSING
  - VIRTUAL MACHINE ASSIST
  - PREFERRED MACHINE ASSIST
- HARDWARE AND MICROCODE CONTROLLED
- WATER AND AIR COOLED

IBM 3084 PROCESSOR COMPLEX

3084 OVERVIEW - TECHNOLOGY

- THERMAL CONDUCTION MODULE (TCM)
  - UP TO 133 CHIPS PER TCM
    - SINGLE 4 7/8 BY 5 1/2 BY 1 3/8 INCH MODULE HAS ABOUT AS MANY CIRCUITS AS A S/370 MODEL 148
  - FILLED WITH HELIUM AND WATER COOLED
  - MOUNTED ON MULTI-LAYERED BOARD
  - OVER 90% OF THE CIRCUITS USE TCM/BOARD TECHNOLOGY
- LSI CARD ON BOARD
  - CHANNEL INTERFACE
- OTHER CARD ON BOARD
  - PROCESSOR CONTROLLER
- HIGHER DENSITY STORAGE

IBM 3084 PROCESSOR COMPLEX

3084 OVERVIEW - FLOOR PLAN

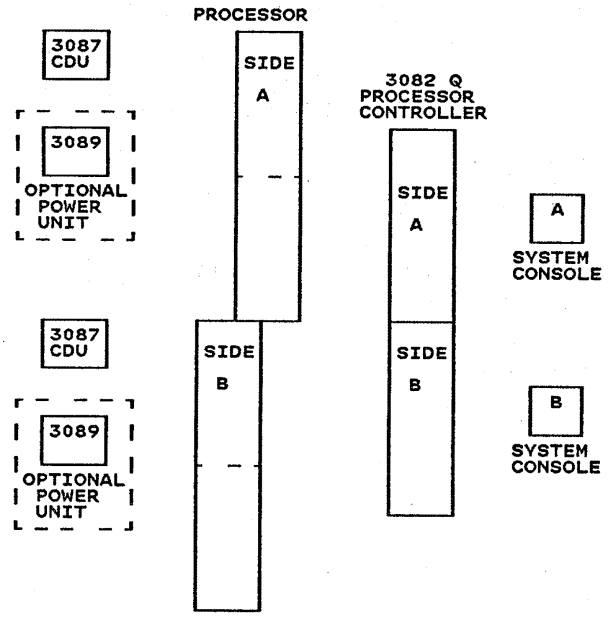

IBM 3084 PROCESSOR COMPLEX

3084 OPERATING MODES

- TWO MODES OF OPERATION
  - SINGLE IMAGE MODE
  - PARTITIONED MODE
- DYNAMIC RECONFIGURATION

IBM 3084 PROCESSOR COMPLEX

3084 SINGLE IMAGE MODE

- ENTIRE 3084 PROCESSOR COMPLEX OPERATING UNDER A SINGLE COPY OF MVS/XA
  - DYNAMIC CHANNEL SUBSYSTEM
    - ALL PROCESSORS HAVE ACCESS TO ALL DEVICES
  - EACH SIDE HAS DIRECT ACCESS TO ALL OF CENTRAL STORAGE
  - SYSTEM AVAILABILITY
  - BOTH SIDES MUST BE AT SAME EC LEVEL

IBM 3084 PROCESSOR COMPLEX

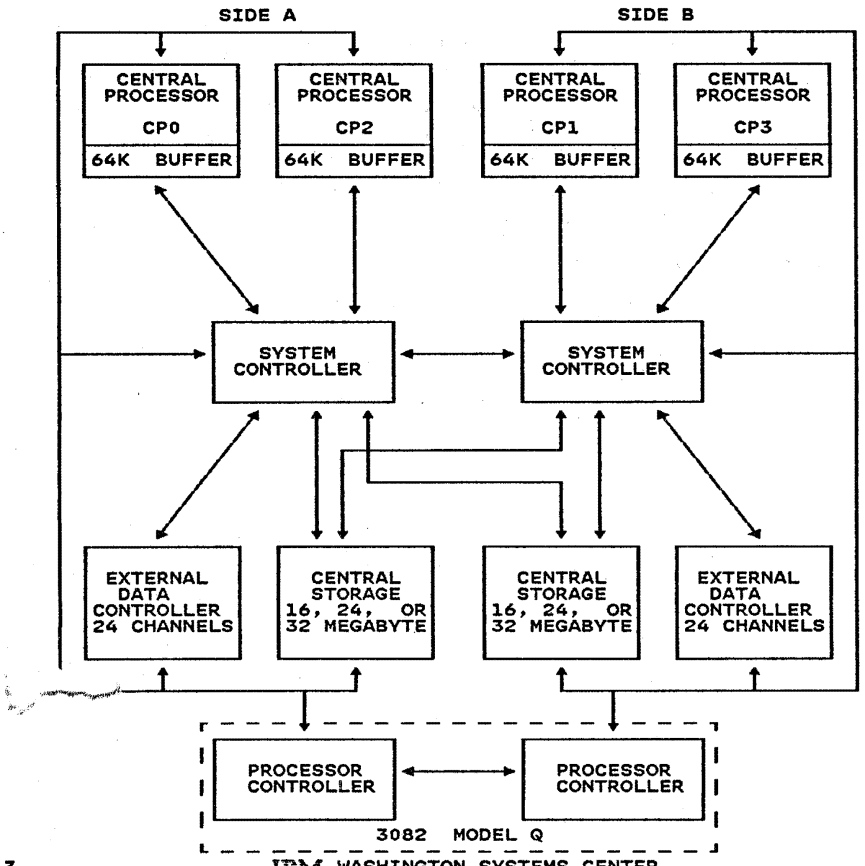

IBM 3084 PROCESSOR COMPLEX

3084 PARTITIONED MODE

- 3084 SPLIT INTO TWO INDEPENDENT SYSTEMS
  - EACH SIDE CAN RUN EITHER MVS/XA, MVS/SP, OR VM/SP
  - MUST BE SYMMETRIC
    - CENTRAL PROCESSORS
    - CHANNELS
    - STORAGE

IBM 3084 PROCESSOR COMPLEX

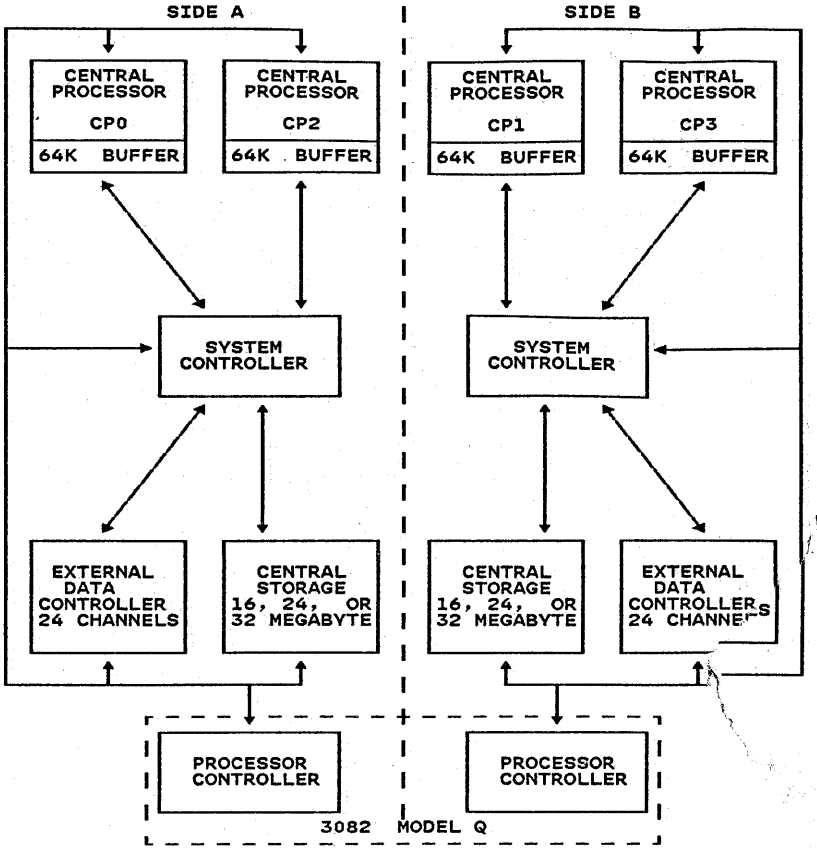

IBM 3084 PROCESSOR COMPLEX

3084 DYNAMIC RECONFIGURATION

- ABILITY TO DYNAMICALLY SWITCH OPERATING MODES WHILE RUNNING UNDER MVS/XA
  - MVS/XA CAN OPERATE, WITHOUT INTERRUPTION, WHILE SWITCHING BETWEEN SINGLE IMAGE AND PARTITIONED MODE

IBM 3084 PROCESSOR COMPLEX
3084 CONSOLES - SINGLE IMAGE MODE

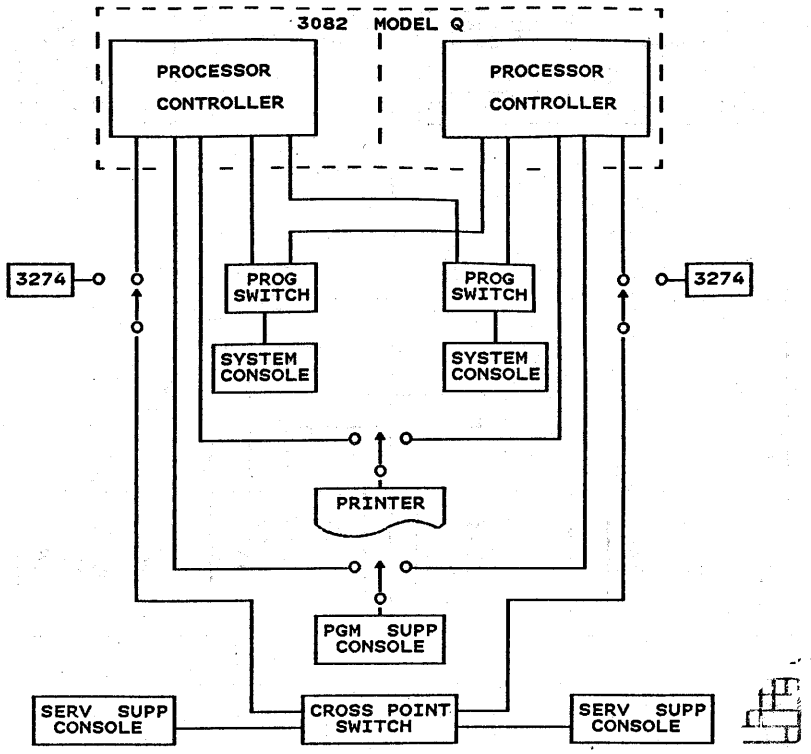

IBM 3084 PROCESSOR COMPLEX
3084 CONSOLES - PARTITIONED MODE

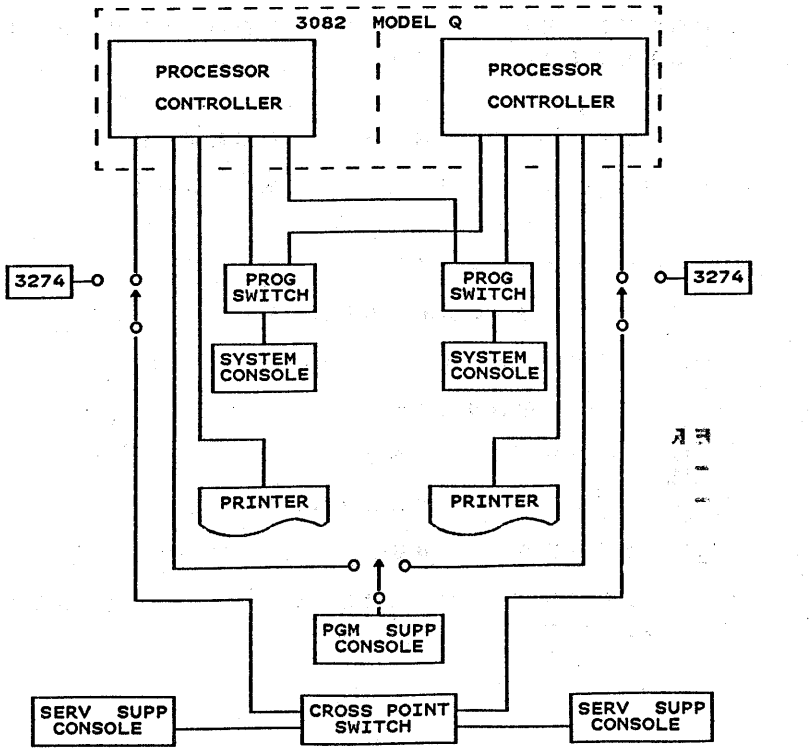

IBM 3084 PROCESSOR COMPLEX

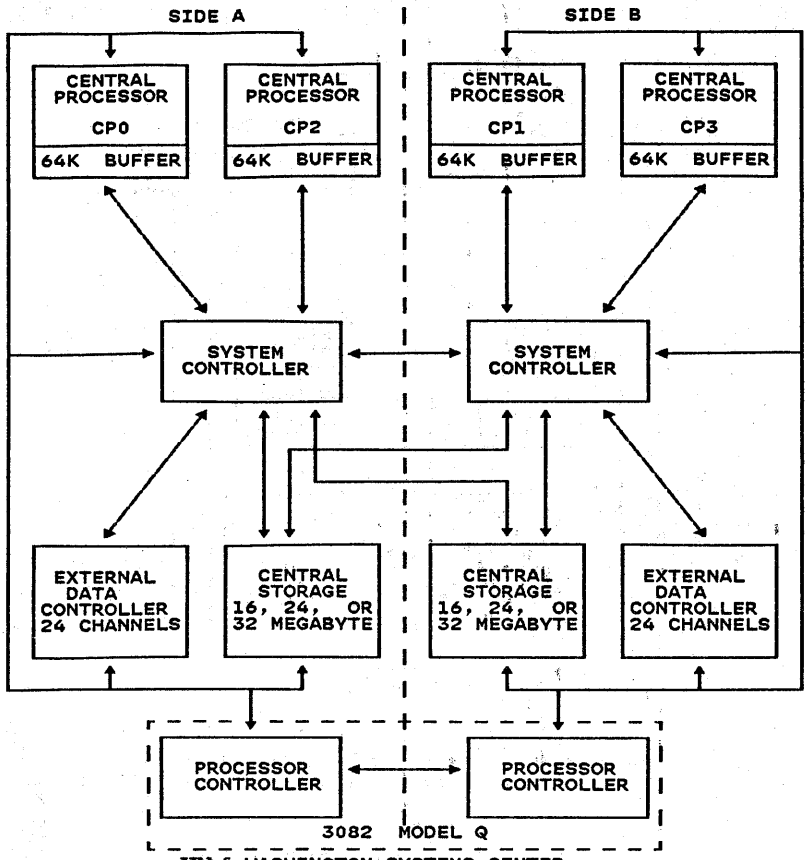

IBM 3084 PROCESSOR COMPLEX

3082 MODEL Q

- MONITORING AND SYSTEM SUPPORT FACILITY (MSSF)
  - POWER SEQUENCING, MONITORING AND CONTROL
  - MICROCODE LOADING AND INITIALIZATION
  - CONFIGURATION CONTROL
    - LOGICAL AND PHYSICAL
  - ERROR HANDLING
    - RECOVERY
    - RECORDING
    - DATA ANALYSIS
  - INTERACTIVE CE INTERFACE
- ENHANCED 3082 RECOVERY IN SINGLE IMAGE MODE

IBM 3084 PROCESSOR COMPLEX

3082 MODEL Q COMPONENTS PER SIDE

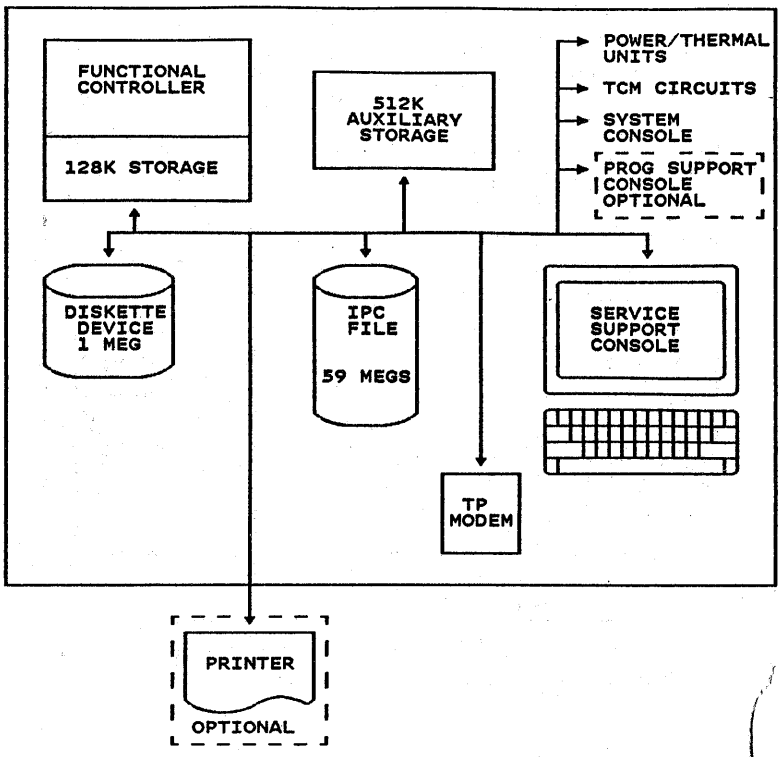

IBM 3084 PROCESSOR COMPLEX

3082 Q - SINGLE IMAGE MODE

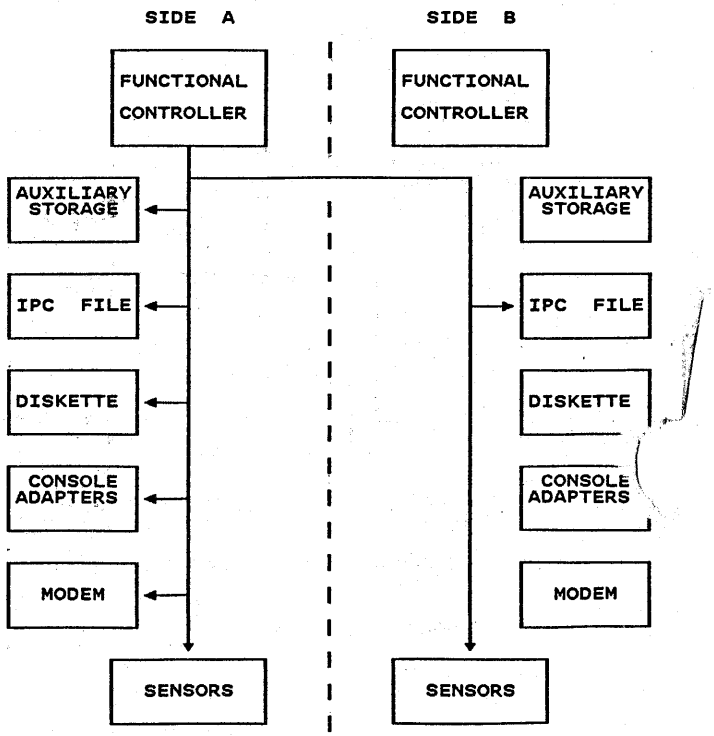

IBM 3084 PROCESSOR COMPLEX
3082 Q - PARTITIONED MODE

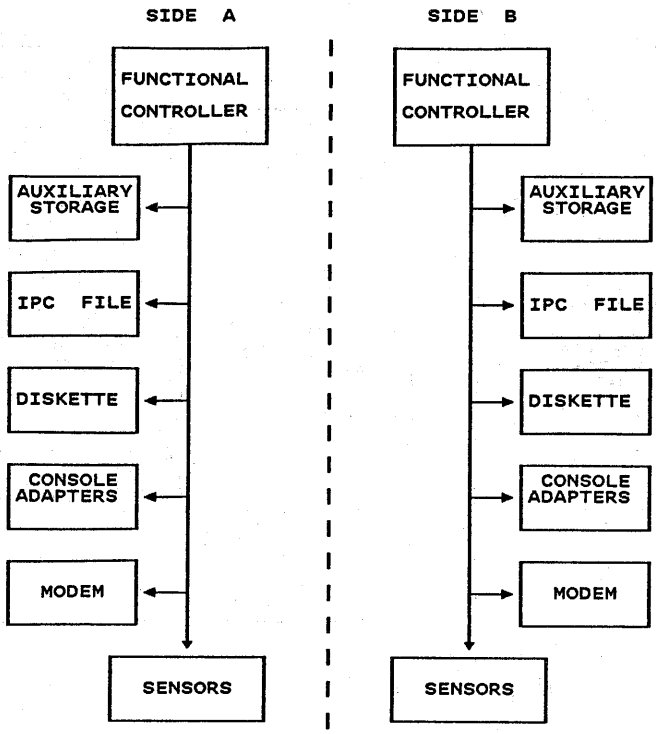

IBM 3084 PROCESSOR COMPLEX
082 Q SINGLE IMAGE ERROR RECOVERY

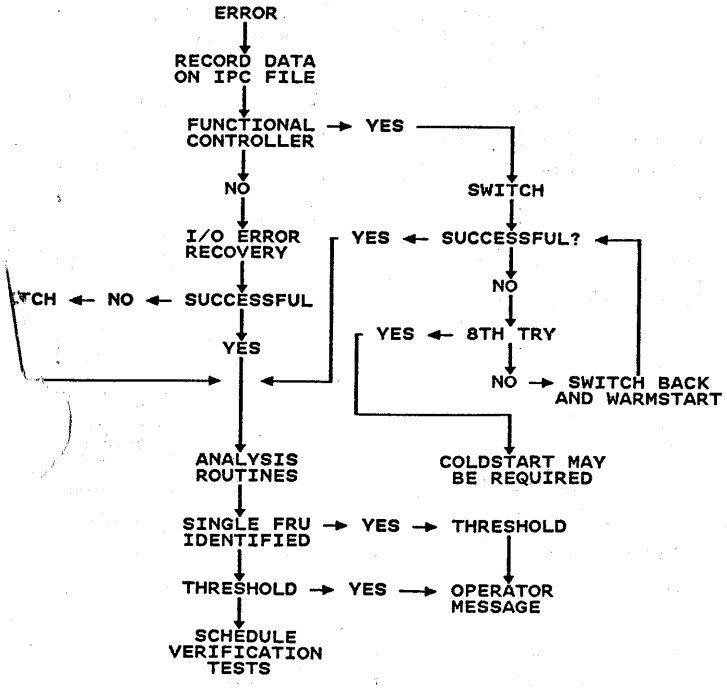

IBM 3084 PROCESSOR COMPLEX
3082 Q - SINGLE IMAGE MODE

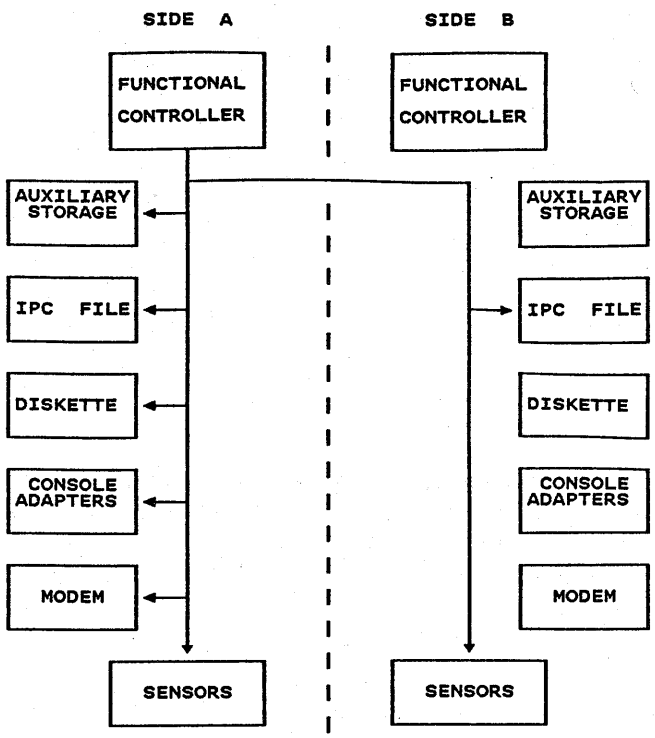

PLE
O

READER'S COMMENT FORM

Title: An Introduction To The
IBM 3084 Processor Complex
Structure And Capabilities
Washington Systems Center
Technical Bulletin ZZ05-0402-00

You may use this form to communicate your comments abou
cation, its organization, or subject matter, with the
that IBM may use or distribute whatever information you
way it believes appropriate without incurring any oblig

Please state your occupation: _____

Comments:

Please mail to: R. M. Schardt
IBM Washington Systems Center
Field Support
18100 Frederick Pike
Gaithersburg, Md. 20879

ZZ05–0402–00

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:

R. M. Schardt
IBM Washington Systems Center
Field Support
18100 Frederick Pike
Gaithersburg, Md. 20879

Fold and tape | Please Do Not Staple | Fold and tape

IBM®

An Introduction To The IBM 3084 Processor Complex Structure And Capabilities

Printed in U.S.A.

ZZ05–0402–00

An Introduction To The IBM 3084 Processor Complex Structure And Capabliities Printed in U.S.A. ZZ05–0402–00

NK

ZZ05–0402–00