


Publication Number
GA26-1664-3

IBM 3380 Models A04, AA4,
and B04
Direct Access Storage

Description and User's Guide

Cross-System

IBM



IBM 3380 Models A04, AA4,
and B04
Direct Access Storage

Description and User's Guide

Cross-System

Publication Number
GA26-1664-3

File Number
S/370-07, 4300-07

e

.

.

Federal Communications Commission (FCC) Statement

Warning: This equipment generates, uses, and can radiate radio frequency energy and if not installed and used in accordance with the instructions manual, may cause interference to radio communications. It has been tested and found to comply with the limits for a Class A computing device pursuant to Subpart J of Part 15 of FCC Rules, which are designed to provide reasonable protection against such interference when operated in a commercial environment. Operation of this equipment in a residential area is likely to cause interference in which case the user at his own expense will be required to take whatever measures may be required to correct the interference.

In addition to the FCC statement above, the user of this manual should be aware that the referenced statement applies only to devices manufactured after January 1, 1981, and used in the United States.

Fourth Edition (December 1984)

This major revision obsoletes GA26-1664-2. Technical changes or additions to the text and illustrations are indicated by a vertical line to the left of the changes.

Periodically, changes or additions may be made to the information in this manual. Before using this publication in connection with the operation of IBM systems, consult your local IBM branch office for the editions that are current and applicable.

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

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

A form for readers' comments is provided at the back of this publication. If the form has been removed, comments may be addressed to IBM General Products Division, Device and Interactive Information, Department J62, 555 Bailey Avenue, San Jose, California, 95150. IBM may use or distribute whatever information you supply in any way it believes appropriate without incurring any obligation to you.

© Copyright International Business Machines Corporation 1981, 1983, 1984

Preface

This manual describes the 3380 Models A04, AA4, and B04 Direct Access Storage and gives installation and conversion information. It contains detailed information on 3380 models, attachment configurations, device characteristics, addressing, and space calculations. It also contains general information and guidelines on how to use this particular storage system, including the supporting programs. However, it does not cover specifics on how to prepare channel, system, and application programs. Such information is detailed in manuals that apply to other devices as well as the 3380. These manuals are listed in the bibliography.

The 3380 Models AD4, BD4, AE4, and BE4 are described in the following manuals:

- *IBM 3380 Models AD4, BD4, AE4, and BE4 Direct Access Storage General Information*, GC26-4193
- *IBM 3380 Direct Access Storage: Planning and Use*, GC26-4208-0
- *IBM 3380 Direct Access Storage: Migration*, GC26-4197-0

The *IBM 3380 Storage Control Description* is the manual for detailed reference information on channel commands, sense bit definitions, and system error recovery actions for all of the disk storage devices that attach to it, including the 3380.

It is assumed that the reader has a knowledge of disk storage, data management, and system operations, including the particular IBM operating system to be used in the installation.

Technical Changes and Additions

Technical changes and additions are indicated by a vertical line to the left of the revision in the body of the manual. They are summarized below.

- The 3380 can be attached (through a 3880) to a 4361 or a 4381 Processor.
- The 3380 is supported by VSE/Advanced Functions (VSE/AF).
- The 3380 is supported by OS/VS1 with DFDS (Data Facility Device Support).
- The IBM 3380 Model AA4 dynamic path selection function provides the capability to simultaneously access two devices (access mechanisms) from the two controllers in the Model AA4. However, in a VSE/Advanced Functions or OS/VS1 environment,

the related commands for enhanced performance for the speed matching buffer for 3380 feature and the dynamic path selection function are not supported.

- There is a new channel command, Diagnostic Control, for diagnostic purposes.

Terminology

Because this book contains both hardware and programming information, the term *physical record* is used to distinguish records as they are formatted on the tracks from the programmer's *logical records*. Both the terms *file* and *data set* are used to refer to a collection of logical records. The term *volume* corresponds to the amount of space accessible by one access mechanism.

A glossary of abbreviations used in this manual is given in the back of the book.

The *IBM Vocabulary for Data Processing, Telecommunications, and Office Systems*, GC20-1699, defines terms used in IBM publications.

Prerequisite Publications

IBM 3380 Storage Control Models 1, 2, 3, and 4 Description, GA26-1661, *IBM 3380 Storage Control Model 13 Description*, GA32-0067, and *IBM 3380 Storage Control Model 23 Description*, GA32-0083 give detailed descriptions of channel commands, status information, and sense information.

IBM Input/Output Equipment Installation Manual-Physical Planning For System/360, System/370, and 4300 Processors, GA22-7064, gives information for planning floor space requirements and other physical planning considerations for the 3380.

Pocket Reference

A summary of IBM 3380 information is contained in the *IBM 3380 Reference Summary*, GX26-1678, a pocket-size reference card.

References to Other Publications

All publications referred to in this manual are listed in the bibliography, which precedes the index. The bibliography lists other related publications, which may or may not be referred to in the text.

Contents

Preface	iii	Overlapping Operations	28
Technical Changes and Additions	iii	Internal Paths and Device Access	29
Terminology	iii	Dynamic Path Selection, Model AA4	29
Prerequisite Publications	iii	Device Reserve and Release Changes	31
Pocket Reference	iii	Reserve Resource on Shared Volume	31
References to Other Publications	iii	System Related Reserve, Model AA4	31
		Dynamic Reconnection, Model AA4	31
Figures	vii	Dynamic Path Selection Controls, Model AA4	31
		Path Group Identifier, Model AA4	32
IBM 3380 Quick Reference	viii	Model AA4 with OS/V51	32
		Model AA4 with VSE/Advanced Functions	32
Introduction	1	Model AA4 with VM/SP or ACP/TPF	32
		Hardware Planning Considerations	33
Models and Attachment Configurations	3	Processor Attachment Requirements	33
Models	3	Physical Planning	34
Methods of Attachment to the 3880 Storage Control	3	Power and Air Conditioning	34
Attachment to 3880 Model 2	3	Space Requirements	34
Attachment to 3880 Model 3	3	Cable Requirements	34
Attachment to 3880 Model 13 or 23	3	Addressing	35
Model AA4 Attachment	4	Device Address	35
Channel Attachment and Transfer Rate	8	Address Assignments	35
Data Streaming and Speed Matching Buffer	8	Defining Addresses at System Generation	36
Examples of 3380 Attachment to a Single Processor	9	Addressed Selection Not Possible	38
Examples of 3380 Attachment to Two Processors	14	Seek Address	38
Channel Switches	17	Seek Address Not Valid	38
		Operator Controls and Indicators	39
Device Characteristics and Formats	19	Power On and Off	39
Head and Disk Assembly	19	Emergency Power Off	39
Characteristics Summary	19	AC Power On Light	39
Record and Track Formats	21	Path and Device On and Off	39
Record Format	21	VARY Command and Path Group ID	40
Track Format	21	Device Status	43
Physical Records Per Track	22	Missing Interrupt	43
		Physical Identifier Labels	43
I/O Operations	23	Installation and Initialization	44
Channel Commands	23	Installation	44
Control Commands	23	Initialization	44
Search Commands	24	Volume Label and VTOC (Minimal Level)	44
Read Commands	24	HA and R0 Validation (Medial Level)	44
Write Commands	24	Surface Checking (Maximal Level)	44
Sense Commands	25	Initialization Under VSE/Advanced Functions	45
Diagnostic Commands	25	Index Creation and Support	45
Command Sequence	25	Program Planning Considerations	47
Commands for Speed Matching Buffer	25	System Support	47
Commands for Cache	26	OS/V51 Support	47
Standard Features	26	MVS/SP Version 1 Support	47
Speed Matching Buffer in 3880 Models 2 and 3	27	MVS/SP Version 2 (MVS/XA) Support	48
Operations for 1.5 MB/sec	27		
Operations for 1.5 and 3.0 MB/sec, Same Processor	28		
Model AA4 Attached to 3880 Model 13 or 23	28		

VM/SP Support	48	Device Status and Error Handling	62
VSE/SP with VSE/Advanced Functions Support	49	Program Functions	62
ACP/TPF Support	49	Error Handling Guidance	62
Device Support	49	Error Handling for Different Types of Errors	63
MVS/SP Version 1 with DFDS	49	Programming Errors	63
MVS/SP Version 1 with MVS/370 DFP or MVS/SP Version 2 with MVS/XA DFP	50	Control Errors	63
OS/VS1 with DFDS	50	Data Errors	63
VSE System IPO/E Release 1.4.2	50	Equipment Errors	64
VSE/Advanced Functions Release 1.3.5	50	Statistical Information and Logging	65
Additional Separate Programs	51	Physical Identifiers	65
Device Support Facilities	51	Data Backup and Recovery	67
Environmental Recording, Editing, and Printing	51	Conversion Checklist	69
IBM Direct Access Storage Device Migration Aid	51	Appendix A: Space Calculations	72
Data Facility Data Set Services	51	Formulas	72
Hierarchical Storage Manager	52	Equal Length Physical Records	72
Data Set Conversion	53	Unequal Length Physical Records	73
Data Set Conversion Under MVS/SP and OS/VS1	53	Space Count Maintenance	73
Approaches	53	An Example of Copying a Partitioned Data Set	73
Programming Changes	53	Tables, without and with Keys	74
Data Set Movement	54	Equal Length Physical Records without Keys	75
Use of Existing Programs	54	Equal Length Physical Records with Ranges of Key Lengths	77
User-Written Access	54	Appendix B: Addressing Reference Information	82
Data Set Conversion Under VM/SP	54	Abbreviations	83
Data Set Conversion Under VSE/Advanced Functions	55	Bibliography	84
Effective Use of the 3380	56	Index	87
Criteria for Using the 3380	56		
Space and Performance Guidelines	56		
Optimum Blocking for 3380 Records	56		
Performance Considerations	59		
Application Considerations	60		
Page Data Sets, Special Considerations	60		

Figures

Figure 1. Example of Two 3380 Model AA4 Strings Attached to 3880 Storage Directors	5	Figure 17. Ratio of Access Mechanisms for Approximate Equivalent Capacity	20
Figure 2. Example of Four 3380 Model A04 Strings Attached to 3880 Storage Directors	6	Figure 18. Track and Record Formats, Simplified Schematic	22
Figure 3. Representative 3880/3380 Valid Configurations	7	Figure 19. Four Internal Paths in a 3380 Model A04 with One Controller	30
Figure 4. 3380 Model A04 Strings Attached to 3880 Model 2	9	Figure 20. Four Internal Paths and Dynamic Path Selection in a 3380 Model AA4 with Two Controllers	30
Figure 5. 3380 Model A04 Strings Attached to 3880 Model 3	10	Figure 21. Comparisons of Approximate Floor Space Requirements Based on Full String Configurations	34
Figure 6. 3380 Model AA4 Strings Attached to 3880 Model 13 or 23	10	Figure 22. 3880 and 3380 Addressing	35
Figure 7. 3380 Model AA4 Strings Attached to Two 3880 Model 2s	11	Figure 23. Sets of Valid 3880/3380 Addresses.	36
Figure 8. 3380 Model AA4 Strings Attached to 3880 Model 3	11	Figure 24. Addressing Example: Two Model A04 Strings Attached to One 3880 Storage Director.	37
Figure 9. 3380 Strings Attached to Several 3880s.	13	Figure 25. Addressing Example: Two Model AA4 Strings Attached to the Same Two 3880 Storage Directors	37
Figure 10. 3380 Model A04 Strings Attached to 3880 Model 3	14	Figure 26. 3380 Control Panel.	39
Figure 11. 3380 Model AA4 Strings Attached to 3880 Model 13 or 23	14	Figure 27. Example of Standard Operating Procedures	42
Figure 12. 3380 Model AA4 Strings Attached to Two 3880 Model 2s	15	Figure 28. Volume Layout of Initialization Information.	46
Figure 13. 3380 Strings Attached to Several 3880s.	16	Figure 29. Differences in Device Space Requirements between the 3350 and 3380	57
Figure 14. Channel Switching for 3880 Storage Directors and Attached 3380 Disk Storage	18	Figure 30. Data Error Correction	63
Figure 15. 3380 Characteristics.	20	Figure 31. Possible Physical IDs for Model AA4 Controllers	66
Figure 16. Comparison of Characteristics of Different IBM Disk Storage Types.	20	Figure 32. Disk-to-Tape Dumping.	68

IBM 3380 Quick Reference

Processor Attachments

System/370 Models 158, 158-3, 168, and 168-3
3031, 3032, and 3033 Processors
3042 Processor Model 2
3081, 3083, and 3084 Processors
4341 Processors
4361 Processor
4381 Processor

Storage Control Attachments

3880 Models 2, 3, 13, and 23

3380 Models

(Each drive has a head and disk assembly with two access mechanisms.)

A04 – one controller, two drives
AA4 – two controllers, two drives, dynamic path selection
B04 – two drives

Operating Systems

OS/VS1 with DFDS
MVS/SP Version 1 with DFDS
MVS/SP Version 1 with MVS/370 DFP
MVS/SP Version 2 (MVS/XA) with MVS/XA DFP
VM/SP
VSE/Advanced Functions with Device Support Facilities (ICKDSF)
ACP/TPF

Introduction

The IBM 3380 Models A04, AA4, and B04 Direct Access Storage is a disk storage with a capacity of 2.5 gigabytes (billion bytes) per unit. An average seek time of 16 milliseconds and a data transfer rate of 3 megabytes-per-second contribute to the effective use of this increase in online storage.

A film head technology is used in the 3380 to achieve writing and reading of data recorded at higher densities than in previous disk storage devices. The read/write heads and disks are sealed in an enclosure to protect the disk surfaces from contaminants. The entire head and disk assembly is permanently mounted on the drive to maintain the critical head and track alignments.

Two movable access mechanisms (actuators) service each head and disk assembly. Each access mechanism has its own address, operates independently, and accesses 630.2 megabytes (MB) of data.

The 3380 attaches through an IBM 3880 Storage Control Model 2, 3, 13, or 23 to any of the following IBM processors:

3031, 3032, and 3033 Processors
3042 Attached Processor Model 2
3081, 3083, and 3084 Processors
4341 Processor
4361 Processor (3880 Models 2 and 3 only)
4381 Processor

Also, the 3380 attaches through an IBM 3880 Storage Control Model 2 or 3 to the following IBM processors:

System/370 Models 158 and 158-3
System/370 Models 168 and 168-3

The 3880 Storage Control Models 2, 3, 13, and 23 have two independent storage directors, each of which attaches to a block multiplexer channel. Data streaming in the processor accommodates attachment to channels with a 3.0 megabyte-per-second (MB/sec) data transfer rate. A speed matching buffer in the 3880 accommodates attachment to channels with data transfer rates less than 3.0 MB/sec. Special channel commands optimize performance of write operations when the speed matching buffer is used. The 3880 Model 13 and 23 contain electronic storage (cache) to retain frequently used data. This allows the processor to access the data faster.

There are model As and model Bs of the 3380. The model As contain the controller functions required to attach 3380s to a 3880 Storage Control. The 3380 can be connected in strings of one model A and up to three model Bs. A full string has 16 access mechanisms with a

total storage capacity of 10.08 gigabytes. Two 3380 strings can be attached to each 3880 storage director.

Model As are available with one or two controllers. When there are two controllers, both are in the same unit and have identical functions. The 3380 model As with two controllers attach to two 3880 storage directors, which in turn attach to two channels, thus achieving two paths for command and data transfers. The 3380 models with two controllers contain logic for a collection of functions called dynamic path selection. Implementation of the functions requires operating system support. The dynamic path selection functions include:

- Simultaneous data transfer over two paths
- Automatic use of the other controller path in the event one of the controllers is unavailable
- Extension of device reservation to a group of paths having a common system identifier
- Reconnection to the first available path in a group of paths having a common system identifier, when operating with System/370 extended architecture

Floor space and power consumption are becoming increasingly significant considerations. The 3380 provides increased storage capacity without requiring additional floor space, and provides important savings in power consumption on a per-byte basis over previous disk storage products.

The 3380 has essentially the same operating characteristics as other current disk storage. It uses the count, key, and data record format, and the channel commands that support this format. Status, sense, and usage information is supplied to the system by the 3880 Storage Control.

Several of functions that have proven successful for improving disk storage performance are incorporated in the 3380. These include such features as rotational position sensing and command retry.

Along with the advances in technology used in the 3380, there are improvements in programming support. Both new and modified programs are designed to improve online operations and to aid in storage management in environments that use a variety of storage media, including the 3380 disk storage. To this end, the programming support embraces a range of functions.

The 3380 is supported under the following operating systems:

- Operating System/Virtual Storage Version 1 (OS/VS1), which implements System/370 architecture
- Multiple Virtual Storage/System Product Version 1 (MVS/SP Version 1), which implements System/370 architecture
- Multiple Virtual Storage/System Product Version 2 (MVS/SP Version 2), which implements System/370 extended architecture
- Virtual Machine/System Product (VM/SP), used alone or with MVS/SP Version 1, and the VM Migration Aid for System/370 extended architecture
- Virtual Storage Extended/System Package (VSE/SP)
- Airline Control Program/Transaction Processing Facility program product (ACP/TPF)

Programming tools for device support and management are either components or utilities provided with the operating system, or they are separately orderable programs.

For MVS/SP operating systems, many of the needed data management and device support functions are consolidated in three program products:

- Data Facility Device Support (DFDS)
- MVS/370 Data Facility Product (MVS/370 DFP)
- MVS/XA Data Facility Product (MVS/XA DFP)

For OS/VS1, device support is provided with the DFDS program product. For VSE/SP, device support is provided with Device Support Facilities (ICKDSF), and with several programs contained in the VSE System

IPO/E Release 1.4.2 and VSE/Advanced Functions Release 1.3.5.

Programming support includes:

- Volume initialization, with creation of a volume table of contents (VTOC)
- Index for a volume table of contents to reduce VTOC search time
- Integrated catalog function that can replace, or in some cases, coexist with existing catalogs
- Space management including space allocation and utilization
- Access method routines and functions that support the access methods in carrying out input and output (I/O) operations
- Preparation for data conversion from an existing to a new device
- Copy or move data from one device to another
- Dump and restore full or partial volumes of data for backup purposes
- Problem determination and error handling

In addition to the tools for these functions, there is a storage management program, Hierarchical Storage Manager, that monitors the amount of use of data and automatically relocates it to make the best use of the various types of storage. The degree of automatic operation is specified by the user.

The 3380 can fill the need for increased disk storage capacity and improve performance over the 3350 disk storage. With the available programming support, the 3380 meets the requirements of most data storage applications.

Models and Attachment Configurations

Models

There are three models of the 3380 Direct Access Storage. All contain two head and disk assemblies, each with two access mechanisms.

Model A04 contains four access mechanisms and one controller. It attaches to one 3880 storage director.

Model AA4 contains four access mechanisms and two controllers. It attaches to two 3880 storage directors and contains the logic for dynamic path selection functions.

Model B04 contains four access mechanisms. Up to three Model B04s can attach to a 3380 Model A04 or AA4.

All units have the mechanical and electronic systems necessary to operate the disk and to write and read data. The controllers in the Model A04s perform the following functions, usually in cooperation with the 3880 storage director to which they are attached.

- Interpret and execute orders issued by the 3880
- Control the writing and interpretation of the track format on a field basis
- Clock and serialize or deserialize data as it transfers between the device and the 3880
- Provide data integrity through error detection and error correction
- Furnish string status to the 3880
- Perform diagnostic evaluation of the string

Model changes between a Model A04 and B04 are available at the time of manufacture. Model changes between Model A04 and AA4 can be made at the user's location.

Methods of Attachment to the 3880 Storage Control

The 3380 attaches to storage directors of the IBM 3880 Storage Control Models 2, 3, 13, and 23, in strings that consist of one 3380 Model A04 or one 3380 Model AA4, and up to three Model B04s.

Attachment to 3880 Model 2

Two 3380 strings (consisting of a 3380 Model A04 and up to three 3380 Model B04s, or a 3380 Model AA4 and up to three 3380 Model B04s) can attach to storage director 2 of a 3880 Model 2. Storage director 1 of a 3880 Model 2 is used for other types of direct access storage.

Figure 1 shows an attachment configuration of two 3380 Model AA4 strings attached to storage director 2 of two 3880 Model 2s. Figure 3 shows several ways 3380

strings can be attached to storage director 2 of a 3880 Model 2. The chart below shows how 3380 strings can be attached to 3880 Model 2 storage directors.

3380 Model	Attachment to a 3880 Model 2 Storage Director	Attachment to Another 3880 Model 2 or 3880 Model 3 Storage Director
A04	Storage director 2	(not possible)
AA4	Storage director 2	Model 2, storage director 2
AA4	Storage director 2	Model 3, either storage director

Attachment to 3880 Model 3

Two 3380 strings of the same type (two 3380 Model A04 strings or two 3380 Model AA4 strings) can attach to each storage director of a 3880 Model 3. (A 3380 string can consist of a 3380 Model A04 and up to three 3380 Model B04s, or a 3380 Model AA4 and up to three 3380 Model B04s.) Other types of direct access storage cannot be attached to a 3880 Model 3.

Figure 2 shows an attachment configuration of four 3380 Model A04 strings attached to the two storage directors of a 3880 Model 3. Figure 3 shows several ways 3380 strings can be attached to storage directors of a 3880 Model 3. The chart below shows how 3380 strings can be attached to 3880 Model 3 storage directors.

3380 Model	Attachment to a 3880 Model 3 Storage Director	Attachment to Another 3880 Model 3, or 3880 Model 13 or 23 Storage Director
A04	Storage director 1	(not possible)
A04	Storage director 2	(not possible)
AA4	Storage director 1 and Storage director 2	(not possible)
AA4	Storage director 1	Either model, storage director 1 or 2
AA4	Storage director 2	Either model, storage director 1 or 2

Attachment to 3880 Model 13 or 23

Two 3380 strings (consisting of a 3380 Model AA4 and up to three 3380 Model B04s) can attach to each storage

director of a 3880 Model 13 or 23. Other types of direct access storage cannot be attached to 3880 Models 13 and 23.

Figure 3 shows several ways 3380 strings can be attached to storage directors of 3880 Model 13 or 23. The chart below shows how 3380 Model AA4 strings can be attached to 3880 Model 13 or 23 storage directors.

3380 Model	Attachment to a 3880 Model 13 or 23 Storage Director	Attachment to Another 3880 Model 13 or 23 Storage Director
A04	(not possible)	(not possible)
AA4	Storage director 1 and Storage director 2	(not possible)
AA4	Storage director 1 (dual frame)	Storage director 2 (dual frame)

Model AA4 Attachment

When a 3380 string includes a Model AA4, the two controllers in the Model AA4 can attach to two 3880 storage directors. The two 3880 storage directors can be:

- Storage director 2 of a 3880 Model 2 and
 - Storage director 2 of another 3880 Model 2, or
 - Either storage director of a 3880 Model 3

- Both storage directors in the same 3880 Model 3
- Either storage director in two different 3880 Model 3s
- Both storage directors in the same 3880 Model 13
- Either storage director in two different dual-frame 3880 Model 13s
- Both storage directors in the same 3880 Model 23
- Either storage director in two different dual-frame 3880 Model 23s

A Model AA4 cannot attach to two storage directors that share the same channel. The two storage directors must have different channels as shown in Figure 1.

The channels may be on the same or different processors.

If there are two connected strings, each with Model AA4, both strings attach to the same two storage directors.

A Model AA4 cannot share the same storage director as a Model A04.

Although the Model AA4 can attach to two storage directors in the same 3880 Model 3, 13, or 23, the potential for even higher availability is achieved by attaching the controllers to storage directors in different 3880s. This way if one 3880 should fail completely, the 3380 data is still accessible through the other 3880.

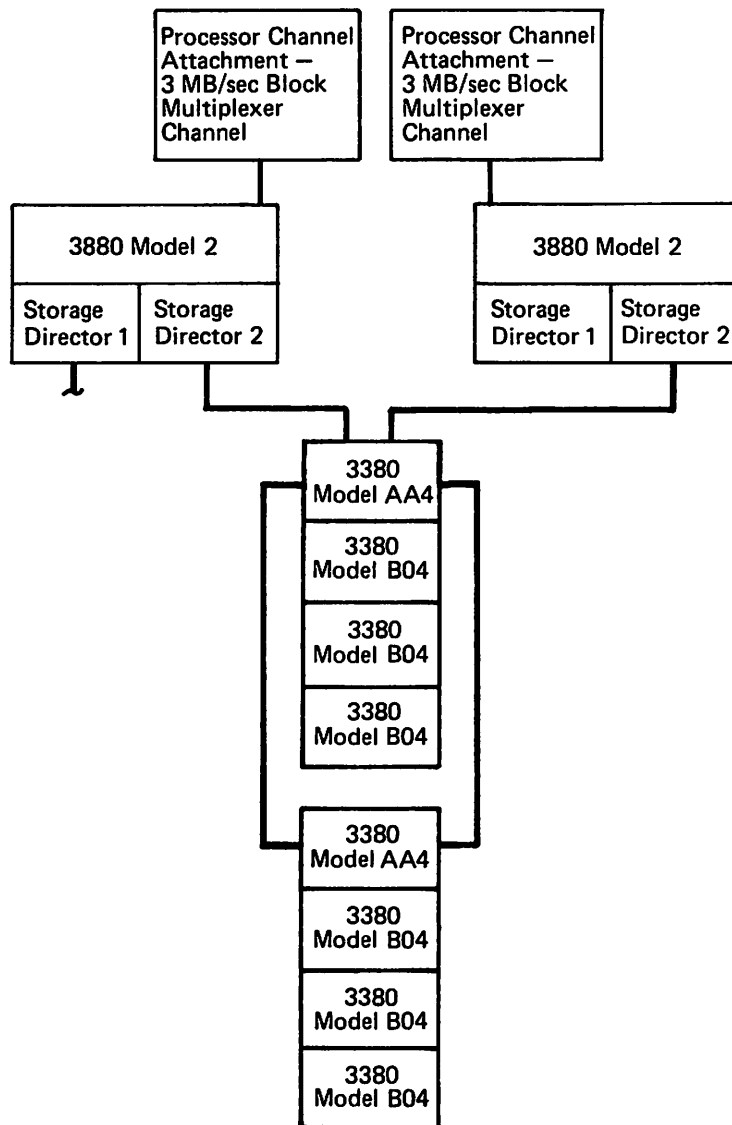


Figure 1. Example of Two 3380 Model AA4 Strings Attached to 3880 Storage Directors

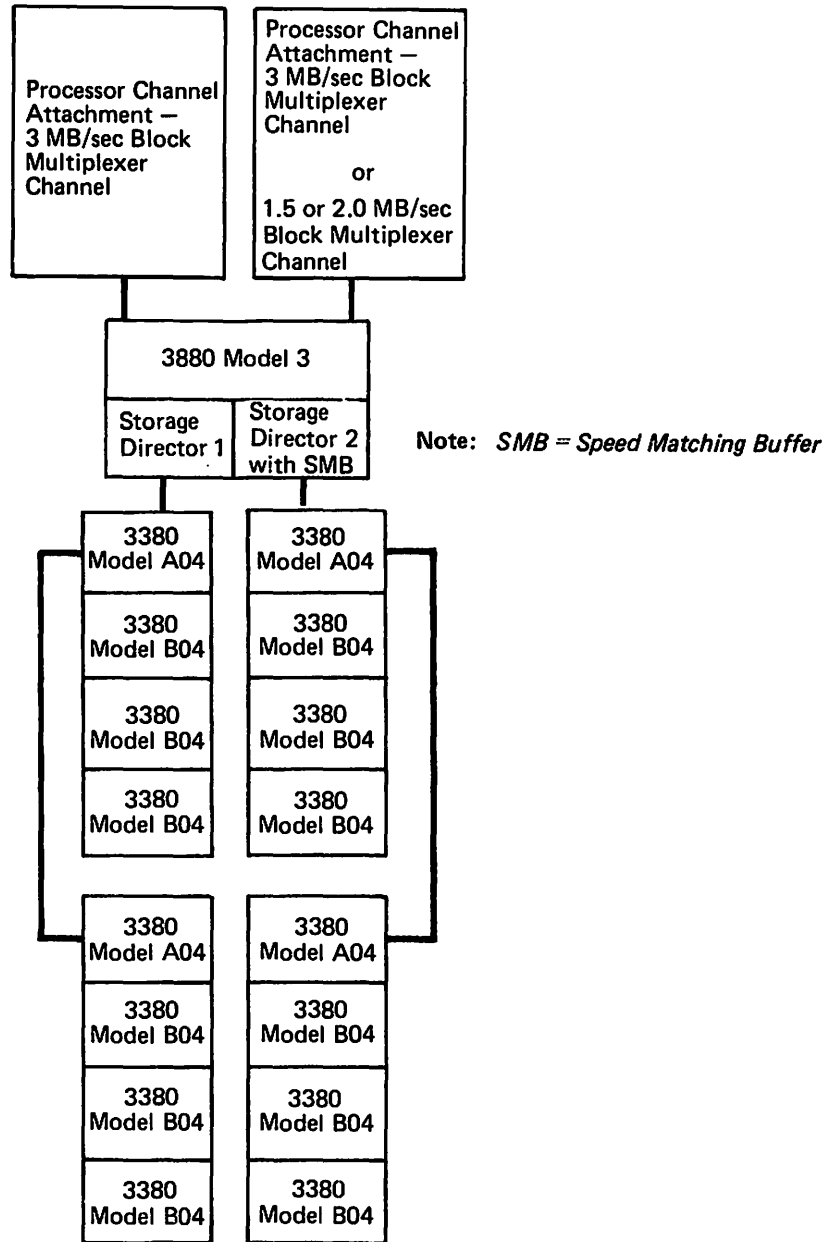


Figure 2. Example of Four 3380 Model A04 Strings Attached to 3880 Storage Directors

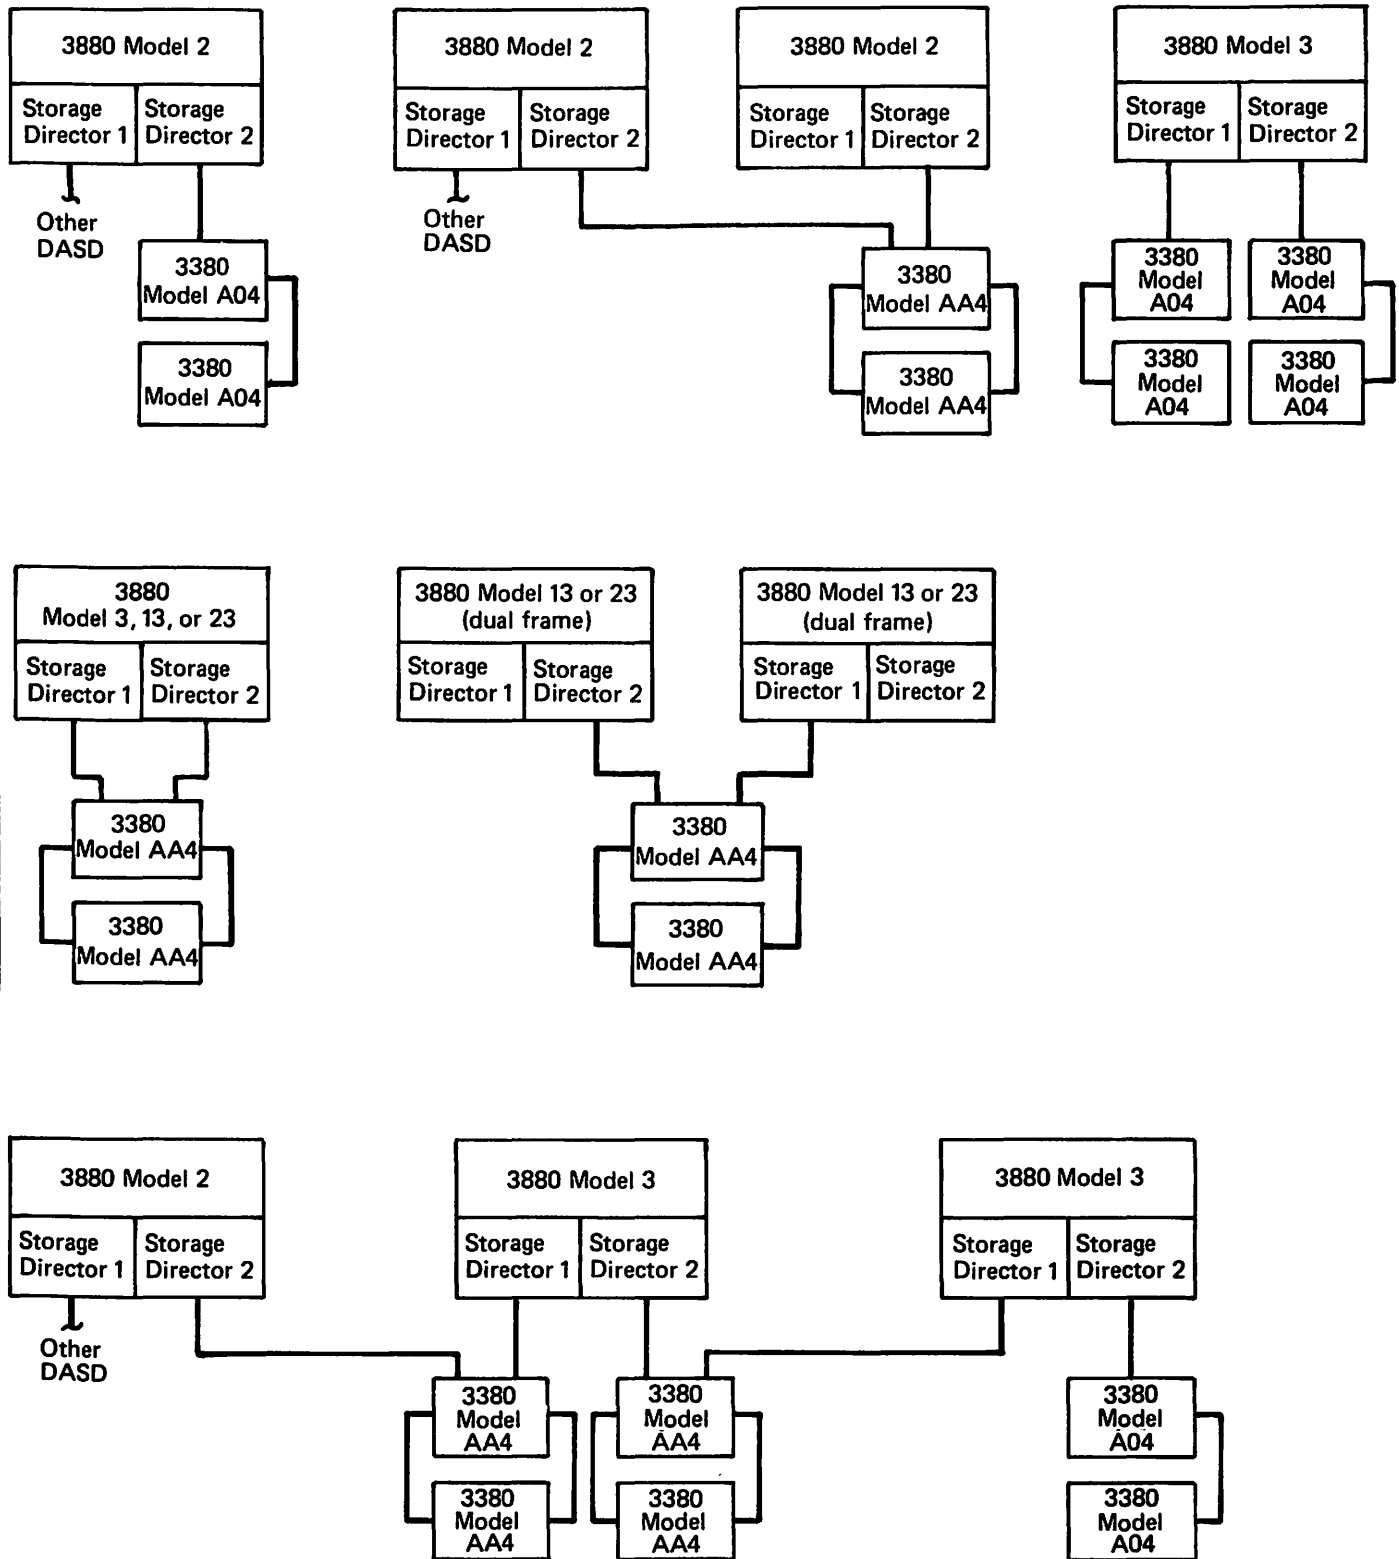


Figure 3. Representative 3880/3380 Valid Configurations

Channel Attachment and Transfer Rate

The 3380 transfers data at 3.0 megabytes-per-second (MB/sec) to and from the 3880 storage director. The 3880, in turn, transfers data to and from the channel at a rate of 3.0 or 1.5 megabytes-per-second (MB/sec) depending on the processor and the features used.

Data Streaming and Speed Matching Buffer

Data streaming provides for 3880 storage director attachment through a block multiplexer channel with a 3.0 MB/sec data transfer rate and, thus, achieves the full 3.0 MB/sec data transfer potential of the 3380.

Data streaming can be installed as a special feature for certain 3.0 MB/sec channels of 3031, 3032, and 3033 Processors, and the 3042 Attached Processor Model 2. Attachment to a 4341 Processor channel with a 3.0 MB/sec data transfer rate does not require a special feature for data streaming. (However, an engineering change is needed on 4341 models manufactured before March 1981, as described in the section Hardware Planning Considerations.) Data streaming is standard and can be specified for block multiplexer channels of the 3081, 3083, and 3084 Processors. Data streaming can also be specified for 3.0 MB/sec block multiplexer channels of 4381 Processors. Data streaming is standard on 4341 Processors manufactured after March 1981.

To also accommodate attachment to channels that do not have a 3.0 MB/sec data transfer rate, a speed matching buffer for 3380 feature can be installed in 3880 Storage Control Models 2 and 3 (but not Model 13 or 23). With this buffer between the 3380 and the channel, channels that do not have a 3.0 MB/sec channel can also be used. The speed matching buffer for 3380 feature is required for:

- Channels of System/370 Models 158, 158-3, 168, and 168-3
- Channels of 3031, 3032, 3033, 3042, and 4341 Processors, other than 3.0 MB/sec channels with data streaming

The speed matching buffer for 3380 feature, when used with 4361 or 4381 channels, must be run at a 3.0 MB/sec data transfer rate.

The speed matching buffer stores 32K bytes of data and operates at two speeds, 1.5 and 3.0 MB/sec. It supports 3.0 MB/sec channels at a 3.0 MB/sec data transfer rate and slower channels at a 1.5 MB/sec data transfer rate. Switches in the 3880 Storage Control are set for the speed of the particular channel. The switches are set by the customer engineer.

With data streaming in a 3031, 3032, 3033, 3042, or 4341 Processor and a speed matching buffer in the 3880 storage director, channels of different speeds (from the same processor) can access the same 3380 devices. For example, a 3380 attached to a storage director with a four-channel switch can operate at 3.0 MB/sec on two channels and 1.5 MB/sec on the two other channels.

The speed matching buffer for 3380 feature can be installed in one or both storage directors of a 3880 Model 3 or one storage director of a 3880 Model 2. Figure 2 shows a configuration where one storage director has the speed matching buffer for 3380 feature and the other does not. If a speed matching buffer is to be installed in only one storage director of a 3880 Model 3, it must be in the number 2 storage director.

When a 3380 Model AA4 is attached to two storage directors that are controlled by the same system, **either both storage directors must have a speed matching buffer or both storage directors should not have a speed matching buffer.** This is because both paths to the Model AA4 must support the special commands used when a speed matching buffer is installed. The related commands for enhanced performance for the speed matching buffer for 3380 feature is not supported by VSE/Advanced Functions or by the OS/VS1 operating system.

The following chart summarizes where the speed matching buffer for 3380 feature can be installed in storage directors of the various 3880 models.

3880 Model	Possible Location of Speed Matching Buffer for 3380 Feature
Model 2	Storage director 2
Model 3	Storage director 2
Model 3	Storage director 1 and 2
Model 13 or 23	(cannot be installed on the 3880 Model 13 or 23)

The following sections show valid attachment configurations with channels that include the data streaming feature and 3880 storage directors that have the speed matching buffer for 3380 feature.

Note: A 3031, 3032, 3033, 3042, or 4341 Processor may have the data streaming feature in one or more of the 3.0 MB/sec channels. When the attached 3880 storage director has the speed matching buffer for 3380 feature, the processor channel (whether or not it has data streaming) has access to an attached 3380 string whether the channel data transfer rate is 3.0 MB/sec, 2.0 MB/sec, or 1.5 MB/sec. A 4381 Processor may have the data streaming

feature in one or more of the 3.0 MB/sec channels. The processor channel (whether or not it has data streaming) has access to an attached 3380 string at a data transfer rate of 3.0 MB/sec.

Refer also to the section Hardware Planning Considerations.

Examples of 3380 Attachment to a Single Processor

Figures 4 through 9 show several ways 3380 Direct Access Storage strings can be attached, through 3880 Storage Controls, to a single processor.

Sharing 3380s Among Several Channels

In several configuration diagrams, two 3380 Model AA4 strings are attached to two storage directors of a 3880 (or to one of the storage directors of two 3880s). A channel is also attached to each 3880 storage director. It might appear that the two channels attached to a 3880 Storage Control can access every one of the 3380 access mechanisms attached to both storage directors of the 3880. This is not the case.

When a 3380 Model AA4 unit (a 3380 controller unit) is indicated in a configuration diagram, it is assumed that the controller unit is attached to a full 3380 Model AA4 string (three Model B04 units attached to a Model AA4 unit). Both controllers in the 3380 controller unit can address any of the 16 access mechanisms in the string (for details about the I/O address of each access mechanism, see the section Addressing). Each access mechanism I/O (unit) address identifies an access mechanism and its 3380 Model AA4 string. (Actually, the I/O address identifies a controller in one of the Model AA4 units. However, both controllers in each Model AA4 unit have the same address, either 0 or 1. Therefore, the controller portion of the I/O address identifies a string rather than a specific controller in the 3380 controller unit.)

A 3880 storage director can be attached to a maximum of two 3380 Model AA4 strings (at the most, thirty-two 3380 access mechanisms). Each of the thirty-two access mechanisms on the two 3380 Model AA4 strings has a unique 1-byte I/O address through a specific channel-storage director path. Therefore, each channel-storage director path is attached to its maximum number of 3380 access mechanisms. In these examples, the channel attached to a storage director can only access the 3380 access mechanisms in the strings attached to that storage director.

An access mechanism I/O address also specifies the address of the 3880 storage director currently accessing

it. Through a two-, four-, or eight-channel switch, a 3880 storage director can be attached to more than one channel. For this reason, the storage director portion of the I/O address can be a value from binary 000 to 111 to indicate the channel port of the storage director.

Therefore, by attaching to a two-, four-, or eight-channel switch, several channels can have access to a 3380 string through the same 3880 storage director. In this case, each access mechanism on the 3380 string attached to the storage director has several different 1-byte I/O addresses; one for each channel-storage director path. This is not the case, however, with the configuration examples shown in Figures 4 through 13. These figures illustrate only one channel (at the most) attached to a 3880 storage director.

Attachment to 3880 Model 2

The attachment configuration in Figure 4 shows a single processor attached, through one channel, to storage director 2 of a 3880 Model 2 and to two 3380 Model A04 strings. Storage director 2 is the only 3880 Model 2 storage director that a 3380 string can attach to.

If the channel has a data transfer rate of 3.0 MB/sec, the speed matching buffer for 3380 feature is not needed on storage director 2. Data streaming is either a standard feature of the channel, or an optional feature, depending on the processor.

If the channel has a data transfer rate of less than 3.0 MB/sec, storage director 2 requires the speed matching buffer for 3380 feature. The channel must be attached to storage director 2, and does not have the data streaming feature.




Figure 4. 3380 Model A04 Strings Attached to 3880 Model 2

Attachment to 3880 Model 3

The attachment configuration in Figure 5 shows a single processor attached, through two channels, to both storage directors of a 3880 Model 3. Each storage director is attached to two separate 3380 Model A04 strings.

If both channels have a data transfer rate of 3.0 MB/sec, the speed matching buffer for 3380 feature is not needed on either storage director. Data streaming is either a standard feature of each channel, or an optional feature, depending on the processor.

If the channel attached to storage director 2 has a data transfer rate of less than 3.0 MB/sec, storage director 2 requires the speed matching buffer for 3380 feature. The slower channel does not have the data streaming feature. If the channel attached to storage director 1 (or both the channels) has a data transfer rate of less than 3.0 MB/sec, both storage directors require the speed matching buffer for 3380 feature.




Figure 5. 3380 Model A04 Strings Attached to 3880 Model 3

Attachment to 3880 Model 13 or 23

The attachment configuration in Figure 6 shows a single processor attached, through two channels, to both storage directors of a 3880 Model 13 or 23. Each storage director is attached to two 3380 Model AA4 strings, which are shared between the storage directors.

Both channels have a data transfer rate of 3.0 MB/sec. The speed matching buffer for 3380 feature is not allowed on either storage director. Data streaming is either a standard feature of each channel, or an optional feature, depending on the processor.




Figure 6. 3380 Model AA4 Strings Attached to 3880 Model 13 or 23

Attachment to Two 3880 Model 2s

The more complex attachment configuration in Figure 7 shows a single processor attached, through two channels, to storage director 2 of two separate 3880 Model 2s. Each storage director is attached to two 3380 Model AA4 strings, which are shared between the storage directors. Storage director 2 is the only storage director of a 3880 Model 2 that a 3380 string can attach to.

If both channels have a data transfer rate of 3.0 MB/sec, the speed matching buffer for 3380 feature is not needed on either storage director. Data streaming is either a standard feature of both channels, or an optional feature, depending on the processor.

If one of the channels has a data transfer rate of less than 3.0 MB/sec, both storage directors require the speed matching buffer for 3380 feature. When a 3380 Model AA4 is attached to two storage directors that are controlled by the same system (processor and operating system), and one of the storage directors has the speed matching buffer for 3380 feature, the other storage director must also have the speed matching buffer for 3380 feature.

Because the speed matching buffer for 3380 feature can only be installed on storage director 2 (of a 3880 Model 2), a channel that intends to access the 3380 Model AA4 string must be attached to storage director 2. The slower channel does not have the data streaming feature.

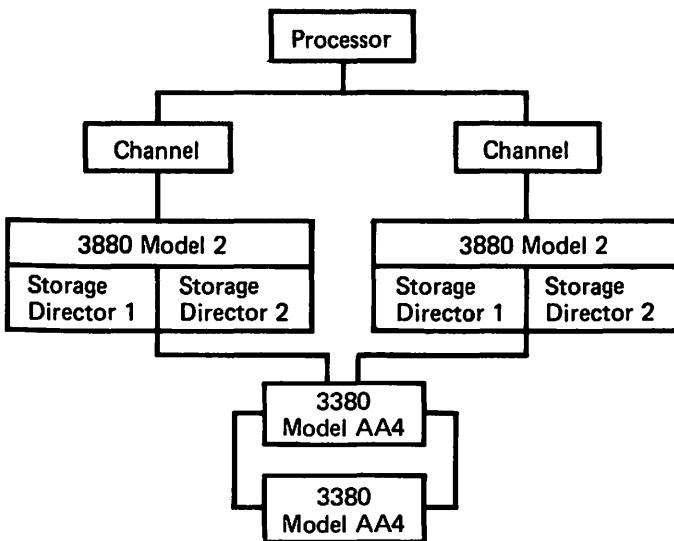


Figure 7. 3380 Model AA4 Strings Attached to Two 3880 Model 2s

Attachment to 3880 Model 3

The attachment configuration in Figure 8 shows a single processor attached, through two channels, to both storage directors of a 3880 Model 3. Each storage director is attached to two 3380 Model AA4 strings, which are shared between the storage directors.

If both channels have a data transfer rate of 3.0 MB/sec, the speed matching buffer for 3380 feature is not needed on either storage director. Data streaming is either a standard feature of both channels, or an optional feature, depending on the processor.

If one of the channels has a data transfer rate of less than 3.0 MB/sec, both storage directors require the speed matching buffer for 3380 feature. When a 3380 Model AA4 is attached to two storage directors that are controlled by the same system (processor and operating system), and one of the storage directors has the speed matching buffer for 3380 feature, the other storage director must also have the speed matching buffer for 3380 feature. The slower channel does not have the data streaming feature.




Figure 8. 3380 Model AA4 Strings Attached to 3880 Model 3

Attachment to Several 3880s

The complex attachment configuration in Figure 9 shows a single processor attached, through five channels, to three 3880 Storage Controls and several 3380 Direct Access Storage strings.

The Processor is Attached, through Channel 1, to storage director 2 of a 3880 Model 2. The processor is also attached, through channel 2, to storage director 1 of a 3880 Model 3. Storage director 2 of the 3880 Model 2 is attached to two 3380 Model AA4 strings (32 access mechanisms), which are shared with storage director 1 of the 3880 Model 3. In this case, the 3380s attached to storage director 2 of the 3880 Model 2 can be accessed through either of two channels (1 and 2).

If channel 1 and 2 each have a data transfer rate of 3.0 MB/sec, the speed matching buffer for 3380 feature is not needed on the storage director attached to either channel. Data streaming is either a standard feature of each channel, or an optional feature, depending on the processor.

If either channel 1 or 2 has a data transfer rate of less than 3.0 MB/sec, storage director 2 of the 3880 Model 2 requires the speed matching buffer for 3380 feature.

Because the speed matching buffer for 3380 feature can only be installed on storage director 2 (of a 3880 Model 2), a channel that intends to access the 3380 Model AA4 string must be attached to storage director 2. The slower channel does not have the data streaming feature.

Storage director 1 of the 3880 Model 3 also requires the speed matching buffer for 3380 feature. When a 3380 Model AA4 is attached to two storage directors that are controlled by the same system (processor and operating system) and one of the storage directors has the speed matching buffer for 3380 feature, the other storage director must also have the speed matching buffer for 3380 feature.

Storage director 2 of the 3880 Model 3 also requires the speed matching buffer for 3380 feature, because for storage director 1 to have the speed matching buffer for 3380 feature installed, storage director 2 must also have the speed matching buffer for 3380 feature.

The Processor is Attached, through Channels 2 and 3, to both storage directors of a 3880 Model 3. Each storage director is attached to two 3380 Model AA4 strings. The 3380 strings attached to storage director 1 are shared with a 3880 Model 2. In this case, the 3380s

attached to storage director 1 can be accessed through either of two channels (1 and 2).

The 3380 strings attached to storage director 2 are shared with another 3880 Model 3. In this case, the 3380s attached to storage director 2 can be accessed through either of two channels (3 and 4).

If channels 3 and 4 each have a data transfer rate of 3.0 MB/sec, the speed matching buffer for 3380 feature is not needed on either storage director. Data streaming is either a standard feature of both channels, or an optional feature, depending on the processor.

If either channel 3 or 4 has a data transfer rate of less than 3.0 MB/sec, both storage directors require the speed matching buffer for 3380 feature. When a 3380 Model AA4 is attached to two storage directors that are controlled by the same system (processor and operating system) and one of the storage directors has the speed matching buffer for 3380 feature, the other storage director must also have the speed matching buffer for 3380 feature. The slower channel does not have the data streaming feature.

Storage director 2 of the 3880 Model 3 also requires the speed matching buffer for 3380 feature, because for storage director 1 to have the speed matching buffer for 3380 feature installed, storage director 2 must also have the speed matching buffer for 3380 feature.

The Processor is Attached, through Channels 4 and 5, to both storage directors of a 3880 Model 3. Storage director 1 is attached to two 3380 Model AA4 strings. Storage director 2 is attached to two 3380 Model A04 strings. The 3380 strings attached to storage director 1 are shared with another 3880 Model 3. In this case, the 3380 access mechanisms attached to storage director 1 can be accessed through either of two channels (3 and 4). The 3380 access mechanisms attached to storage director 2 can be accessed through channel 5.

If channel 5 has a data transfer rate of 3.0 MB/sec, the speed matching buffer for 3380 feature is not needed. Data streaming is either a standard feature of this channel, or an optional feature, depending on the processor.

If channel 5 has a data transfer rate of less than 3.0 MB/sec, storage director 2 requires the speed matching buffer for 3380 feature. Channel 5 does not have the data streaming feature.

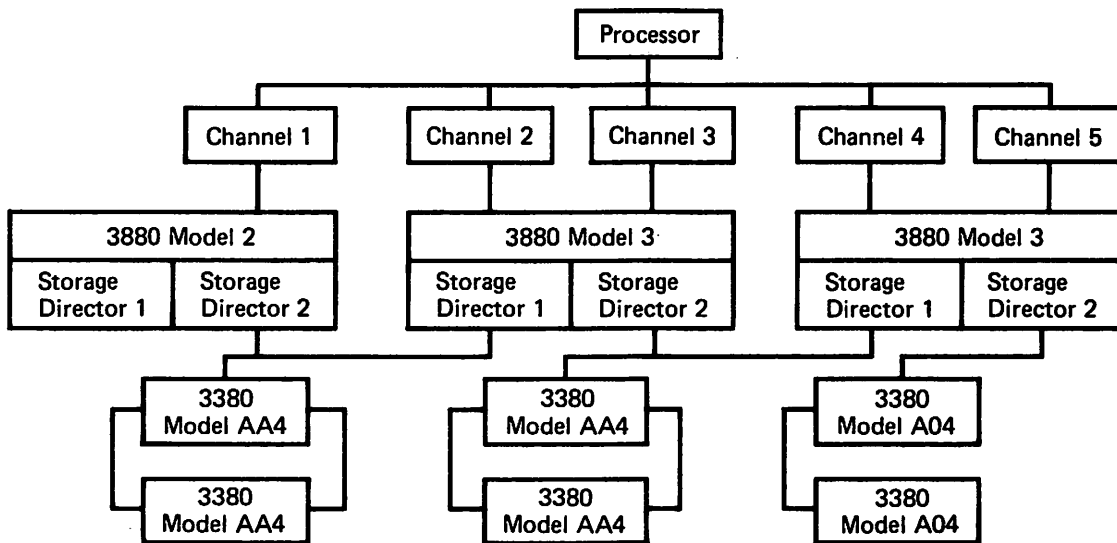


Figure 9. 3380 Strings Attached to Several 3880s

Examples of 3380 Attachment to Two Processors

Figures 10 through 13 show several ways 3380 Direct Access Storage strings can be attached, through 3880 Storage Controls, to two processors.

Attachment to 3880 Model 3

The attachment configuration in Figure 10 shows two processors attached, through one channel from each processor, to both storage directors of a 3880 Model 3. Each storage director is attached to two separate 3380 Model A04 strings.

If the channel from either processor has a data transfer rate of 3.0 MB/sec, the speed matching buffer for 3380 feature is not needed on the storage director attached to that channel. Data streaming is either a standard feature of the channel, or an optional feature, depending on the processor.

If the channel from either processor has a data transfer rate of less than 3.0 MB/sec, the storage director attached to that channel requires the speed matching buffer for 3380 feature. The slower channel does not have the data streaming feature. If the channel attached to storage director 1 (or both channels) has a data transfer rate of less than 3.0 MB/sec, both storage directors require the speed matching buffer for 3380 feature.

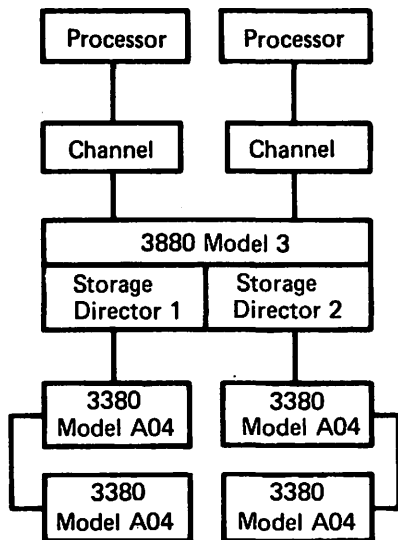


Figure 10. 3380 Model A04 Strings Attached to 3880 Model 3

Attachment to 3880 Model 13 or 23

The attachment configuration in Figure 11 shows two processors attached, through one channel from each processor, to both storage directors of a 3880 Model 13 or 23. Each storage director is attached to two 3380 Model AA4 strings, which are shared between the storage directors.

Both channels have a data transfer rate of 3.0 MB/sec. The speed matching buffer for 3380 feature is not allowed on either storage director. Data streaming is either a standard feature of each channel, or an optional feature, depending on the processor.




Figure 11. 3380 Model AA4 Strings Attached to 3880 Model 13 or 23

Attachment to Two 3880 Model 2s

The more complex attachment configuration in Figure 12 shows two processors attached, through one channel from each processor, to storage director 2 of two separate 3880 Model 2s. Each storage director is attached to two 3380 Model AA4 strings, which are shared between the storage directors. Storage director 2 is the only storage director of a 3880 Model 2 that a 3380 string can attach to.

If the channel from either processor has a data transfer rate of 3.0 MB/sec, the speed matching buffer for 3380 feature is not needed on the storage director attached to that channel. Data streaming is either a standard feature of the channel, or an optional feature, depending on the processor.

When a 3380 Model AA4 is attached to two storage directors that are controlled by the same system (processor and operating system), and one of the storage directors has the speed matching buffer for 3380 feature, the other storage director must also have the speed matching buffer for 3380 feature.

Because the speed matching buffer for 3380 feature can only be installed on storage director 2 (of a 3880 Model 2), a channel that intends to access the 3380 Model AA4 string must be attached to storage director 2. The slower channel does not have the data streaming feature.




Figure 12. 3380 Model AA4 Strings Attached to Two 3880 Model 2s

Attachment to Several 3880s

The complex attachment configuration in Figure 13 shows two processors, attached through two channels from the first processor and three channels from the second processor, to three 3880 Storage Controls and several 3380 Direct Access Storage strings.

The First Processor is Attached, through Channel 1, to storage director 2 of a 3880 Model 2. It is also attached, through channel 2, to storage director 1 of a 3880 Model 3. Storage director 2 of the 3880 Model 2 is attached to two 3380 Model AA4 strings (32 access mechanisms), which are shared with storage director 1 of the 3880 Model 3. In this case, the 3380s attached to storage director 2 of the 3880 Model 2 can be accessed through either of two channels (1 and 2).

If channel 1 and 2 each have a data transfer rate of 3.0 MB/sec, the speed matching buffer for 3380 feature is not needed on the storage director attached to either channel. Data streaming is either a standard feature of each channel, or an optional feature, depending on the processor.

If either channel 1 or 2 has a data transfer rate of less than 3.0 MB/sec, storage director 2 of the 3880 Model 2 requires the speed matching buffer for 3380 feature. Because the speed matching buffer for 3380 feature can only be installed on storage director 2 (of a 3880 Model 2), a channel that intends to access the 3380 Model AA4 string must be attached to storage director 2. The slower channel does not have the data streaming feature.

Storage director 1 of the 3880 Model 3 also requires the speed matching buffer for 3380 feature. When a 3380 Model AA4 is attached to two storage directors that are controlled by the same system (processor and operating system) and one of the storage directors has the speed matching buffer for 3380 feature, the other storage director must also have the speed matching buffer for 3380 feature.

Storage director 2 of the 3880 Model 3 also requires the speed matching buffer for 3380 feature, because for storage director 1 to have the speed matching buffer for 3380 feature installed, storage director 2 must also have the speed matching buffer for 3380 feature.

The First Processor is Attached, through Channel 2, to storage director 1 of a 3880 Model 3. The second processor is attached, through channel 3, to storage director 2 of a 3880 Model 3. Each storage director is attached to two 3380 Model AA4 strings.

The 3380 strings attached to storage director 1 are shared with a 3880 Model 2. In this case, the 3380 access mechanisms attached to storage director 1 can be accessed from both systems through either of two channels (1 and 2). The 3380 strings attached to storage director 2 are shared with another 3880 Model 3. In this case, the 3380s attached to storage director 2 can be accessed from both processors, through either of two channels (3 and 4).

If channels 3 and 4 each have a data transfer rate of 3.0 MB/sec, the speed matching buffer for 3380 feature is not needed on either storage director. Data streaming is either a standard feature of both channels, or an optional feature, depending on the processor.

If either channel 3 or 4 has a data transfer rate of less than 3.0 MB/sec, both storage directors require the speed matching buffer for 3380 feature. When a 3380 Model AA4 is attached to two storage directors that are controlled by the same system (processor and operating system) and one of the storage directors has the speed matching buffer for 3380 feature, the other storage director must also have the speed matching buffer for 3380 feature. The slower channel does not have the data streaming feature.

Storage director 2 of the 3880 Model 3 also requires the speed matching buffer for 3380 feature, because for storage director 1 to have the speed matching buffer for 3380 feature installed, storage director 2 must also have the speed matching buffer for 3380 feature.

The Second Processor is Attached, through Channels 4 and 5, to both storage directors of a 3880 Model 3. Storage director 1 is attached to two 3380 Model AA4 strings. Storage director 2 is attached to two 3380 Model A04 strings. The 3380 strings attached to storage director 1 are shared with another 3880 Model 3. In this case, the 3380 access mechanisms attached to storage director 1 can be accessed through either of two channels (3 and 4). The 3380 access mechanisms attached to storage director 2 can be accessed through channel 5.

If channel 5 has a data transfer rate of 3.0 MB/sec, the speed matching buffer for 3380 feature is not needed. Data streaming is either a standard feature of this channel, or an optional feature, depending on the processor.

If channel 5 has a data transfer rate of less than 3.0 MB/sec, storage director 2 requires the speed matching buffer for 3380 feature. Channel 5 does not have the data streaming feature.

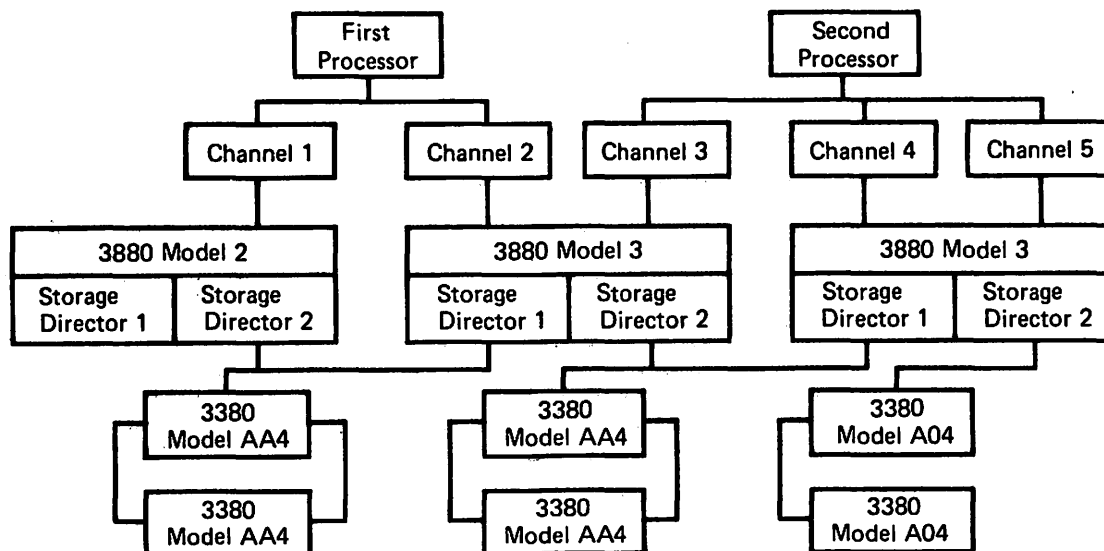


Figure 13. 3380 Strings Attached to Several 3880s

Channel Switches

Two- and four-channel switches are available as special features on the 3880 Storage Control Models 2, 3, and 13, and eight-channel switches are available on Models 2 and 3. The 3880 Model 23 provides for the attachment of two channels to each storage director. The two-channel switch pair feature is required on the 3880 Model 23, and the two-channel switch pair, additional and eight-channel switch are available as optional special features. The three features can be used as described as follows and in Figure 14.


- With a two-channel switch pair feature, each storage director of a 3880 can connect to two different channels. Both storage directors can connect to the same or different channels.
- With a two-channel switch pair, additional feature, each storage director of a 3880 can connect to up to four different channels. Both storage directors can connect to the same or different channels.

- With an eight-channel switch feature, both storage directors of a 3880 connect to the same eight channels.


It is not possible for a 3380 Model AA4 to be attached to two storage directors in the same 3880 with the eight-channel switch because the two storage directors attaching a 3380 Model AA4 must have different channels. However, channel switching is especially significant for 3380 Model AA4s attached to two storage directors in different 3880 storage controls. Data in a 3380 string can be accessed by as many as 16 different channels. This is possible if one 3380 controller is attached to a storage director in one 3880 and the other controller is attached to a storage director in a different 3880, and both 3880s have the eight-channel switch feature.

Storage Directors Connected to the Same Channels


Two-Channel Switch Pair



Two-Channel Switch Pair, Additional




Eight-Channel Switch




Storage Directors Connected to Different Channels

**Two-Channel Switch Pair
(Four unique channels)**




**Two-Channel Switch Pair, Additional
(Eight unique channels)**



Storage Directors Connected to the Same and Different Channels

**Two-Channel Switch Pair
(Three unique channels)**



**Two-Channel Switch Pair, Additional
(Six unique channels)**




Figure 14. Channel Switching for 3880 Storage Directors and Attached 3380 Disk Storage

Device Characteristics and Formats

Head and Disk Assembly

The 3380 has a head and disk assembly on each drive that cannot be removed by the user. Although permanently mounted for normal operations, the assembly can be removed by a customer engineer.

Each head and disk assembly contains two access mechanisms. One access mechanism services one half of the disks and the other access mechanism services the other half of the disks.

Each access mechanism is addressable and accesses a volume. Volumes are initialized with track addresses and track descriptor records as described in the section **Installation and Initialization**.

Servo information is prerecorded at the time of manufacture. It contains information used by the device for seeking, track following, data clocking, index point signal generation, and rotational position sensing.

Characteristics Summary

Figures 15, 16, and 17 summarize the physical and performance characteristics of the 3380, and compare them with previous IBM disk storage products.

Physical Characteristics		Performance Characteristics	
Head and Disk Assemblies per Unit	2	Seek Time (Note 1)	
Access Mechanisms per Head and Disk Assembly	2	Minimum	3.0 ms
Heads per Access Mechanism	15 data and 1 servo	Average	16.0 ms
Cylinders per Access Mechanism		Maximum	30.0 ms
Data	885	Rotational Delay Time (Note 2)	8.3 ms
Alternate	1	Data Rate (Note 3)	3.0 MB/sec
Customer Engineer	1	Notes:	
Data Tracks per Cylinder	15	1. <i>Seek time, or access motion time, is the time required to position the access mechanism at the track (cylinder). The average seek time is for one third of the cylinders. If the mechanism is already at the correct track (cylinder), there is no access motion and seek time is zero.</i>	
Tracks per Access Mechanism		2. <i>Rotational delay time, or rotational latency time, is the average time required for the specified record to rotate to the read and write head so that the data transfer can begin. It is one-half the disk rotation time.</i>	
Data	13,275	3. <i>Data rate is the instantaneous speed at which bytes are transferred.</i>	
Alternates	15		
Maximum Data Capacity per Track	47,476 bytes		
per Cylinder	712,140 bytes		
per Access Mechanism	630.2 MB		
per Head and Disk Assembly	1.26 GB		
per Unit	2.52 GB		
per String	10.08 GB		

Figure 15. 3380 Characteristics

Performance and Physical Characteristics	3380	3375	3350	3340	3330
Seek Time (ms)					
Average	16	19	25	25	30
Full Track Rotation Time (ms)	16.7	20.2	16.7	20.24	16.7
Data Transfer Rate (MB/sec)	3.0	1.859	1.198	0.885	0.806
MB Capacity per Unit	2,521	819.7	635	140	200 or 400
MB Capacity per Drive	1,260	819.7	317.5	70	100 or 200
MB Capacity per Access Mechanism	630	409.8	317.5	70	100 or 200
MB Capacity per Fixed Head Unit	0	0	2.28	0.5	0
Data Cylinders per Access Mechanism	885	959	555	696	404 or 808
Tracks per Cylinder	15	12	30	12	19
Bytes per Track	47,476	35,616	19,069	8,368	13,030
Bytes per Cylinder	712,140	427,392	572,070	100,416	247,570

Figure 16. Comparison of Characteristics of Different IBM Disk Storage Types

<p>One 3380 access mechanism to three 3330-11 access mechanisms One 3380 access mechanism to nine 3340 access mechanisms One 3380 access mechanism to two 3350 access mechanisms</p>
--

Figure 17. Ratio of Access Mechanisms for Approximate Equivalent Capacity

Record and Track Formats

The 3380 uses the count, key, and data record format, which provides continuity for transition from existing IBM direct access storage products to the 3380.

Record Format

This section describes the format of physical records as they are written on the tracks. (One or more logical records or portions of logical records may be included in a physical record as shown in Figure 18. The format of logical records refers to whether they are of fixed, variable, or undefined length.)

A physical record contains three areas: count, key, and data. Each area is separated by gaps. Check bytes, which later are used for detecting and correcting data errors, are added to each area when it is written. Records are formatted on the tracks through the use of channel commands.

Count Area: The count area contains the location of a data area on a specific device. The location is specified by cylinder number, head number and record number (CCHHR) in five bytes, and is referred to as the ID. The record number customarily identifies the relative number of the record on the track.

The location information in the count area represents the physical, or actual, address of a record. (The CCHH bytes correspond to those in the seek channel commands, and the CCHHR bytes correspond to those in the search ID channel commands.)

In addition to addressing information, the count area contains certain control information.

The count area also defines the length in bytes of the key and data areas of the record, key length (KL) and data length (DL).

The count area is written when the record is formatted and is not changed until the record is reformatted.

Key Area: The key area is optional. The key is used by the programmer to identify the information in the data area of the record. Examples of key identifiers are social security numbers, part numbers, or employee numbers.

Once the key area is formatted, the contents may be rewritten, but the length of the key may not be changed without reformatting the record. If the key area of a record is rewritten, the data area of the record must also be rewritten.

Data Area: The data area contains the information identified by the count and key areas of the record. The

data area contains the user's logical records, which are organized and arranged by the programmer.

Once the data area is formatted, the contents may be rewritten, but the length of the data may not be changed without reformatting the record. The contents of the data area may be rewritten without changing any other area of the record.

Track Format

All tracks are formatted beginning at the index point and ending at the following index. Each track has the same basic format: track home address, track descriptor record, and one or more user data records. See Figure 18.

Home Address: Each track contains a home address (HA) recorded following the index point. It contains the track address defined by cylinder and head numbers (CCHH).

The home address area also contains information describing the condition of the track. This information is carried in the flag (F) byte, and indicates whether the track is usable, defective, or an alternate.

Special channel commands are used for writing and reading a home address area: Write Home Address and Read Home Address. Only certain portions of the area are transferred with these commands. The Diagnostic Read and Diagnostic Write Home Address commands transfer the entire home address area information.

Track Descriptor Record: Record 0 (R0) is always the first record on the track following the home address. In IBM programming systems, if the track is flagged defective in the home address, the count area of the following R0 contains the track address of the alternate track. The count area of the R0 on the alternate track contains the track address of the defective primary track.

A record 0 is referred to as standard or nonstandard. A standard R0 is reserved for the operating system for nonuser data. It always has a key length of zero and a data length of eight bytes, written with all zeros. A nonstandard R0 is used as a normal data record with the key length and data length determined by the user.

Special channel commands are used for writing and reading a track descriptor record: Write Record 0 and Read Record 0.

Data Records: One or more user data records may follow the track descriptor record (R0) on a track.

Data records are formatted using the channel Write Count, Key, and Data command or Write Count, Key and Data Next Track command. The format can be changed by subsequently executing another Write Count, Key, and Data command; however, the previously written records will be destroyed.

Data records, as well as track descriptor records, can be formatted with or without keys. Generally, file organization determines whether keys are used.

Physical Records Per Track

The number of physical records that can be placed on a track depends on the size of the records, the track capacity, and whether there are keys. The size of records is determined by the user, except for track descriptor records.

The way physical record size is specified depends on whether the programmer is preparing channel commands or defining data characteristics for the operating system.

When preparing channel programs, the physical record size is the number of bytes in the data area, as defined in the count area as the data length (DL). This is specified with a channel Write Count, Key, and Data command or Write Count, Key, and Data Next Track command.

The method for describing data characteristics depends on the operating system and access method used. For MVS non-VSAM records, physical record size is the number of bytes in the block, as defined in the BLKSIZE parameter of the OS/VIS data control block (DCB) macro. (LRECL specifies the length of logical records.)

If keys are used, the key length (KL of the count area and KEYLEN of the DCB macro) must be added to the calculation. (There are some exceptions in ISAM to the relationships DL=BLKSIZE and KL=KEYLEN.)

The maximum length of a physical record is the capacity remaining on one track after overhead has been accounted for. Track overhead is the space required on each track for home address, standard track descriptor record (R0), and skip defect allowance. Record overhead is the space required for gaps and the count area. Overhead has been accounted for in the data capacities given in Figures 15 and 16, and in Appendix A. Information is provided in Appendix A for calculating the number of records that can be placed on a 3380 track or cylinder. Block sizes are discussed in the section Effective Use of the 3380.

Cylinder 03

Head 00



Head 01




Figure 18. Track and Record Formats, Simplified Schematic

I/O Operations

Channel Commands

The channel commands used for 3380 operations support the count, key, and data record format. The commands are the same or similar to those used for the IBM 3330 and 3350 disk storage operations, with the following differences.

- There is no Write Special Count, Key, and Data command, because there is no record overflow function.
- Interpretation of the Device Reserve and Device Release commands can be extended with two new commands, Set Path Group ID and Sense Path Group ID, if the 3380 has dynamic path selection.
- Names of the Sense I/O and Sense I/O Type commands have been changed to Sense and Sense ID, but the functions remain the same.
- The Diagnostic Sense/Read command has been retained, and there are two new diagnostic commands, Diagnostic Write Home Address and Diagnostic Read Home Address.

Besides these changes to the basic command set, there are new commands for use with the 3880 speed matching buffer for 3380 feature, and for use of the cache in the 3880 Model 13 and 23.

The command information is carried in a channel command word (CCW). The channel command word is in one of two formats, depending on the system architecture. Format 0 is used for the 24-bit addressing of System/370 architecture and Format 1 is used for the 31-bit addressing of System/370 extended architecture. The channel commands are processed by the storage director to which the 3380 is attached and are described in detail in the *IBM 3880 Storage Control Description* manuals, listed in the bibliography. The description of each command is accompanied by a description of the status conditions that result when the processing of the command completes.

Users who need information for writing channel programs should refer to the *3880 Storage Control Description* manual, but for general information, the commands are briefly defined as follows.

Control Commands

Seek, or **Seek Cylinder** transfers the seek address, consisting of the cylinder and head numbers (00CCHH), from main storage to the storage director. It causes positioning of the 3380 access mechanism specified in the I/O address if access motion is required.

Seek Head transfers the seek address from main storage to the 3880 storage director. It is used to select a different head at the same cylinder where the access mechanism is positioned.

Set File Mask transfers the mask byte from main storage to the storage director. The mask byte defines the write and seek operations that can be used in the channel command word (CCW) chain and defines the interaction between command retry and program controlled interrupt (PCI). The provision for defining the commands that are permitted and those that are inhibited is a means of protecting data integrity.

Set Sector transfers the sector number (a relative angular track position) from main storage to the storage director. This command is used in conjunction with the Read Sector command to achieve rotational position sensing (RPS). The Read Sector command is used to obtain the angular position of a record. A subsequent Set Sector command causes the 3880 storage director to attempt reconnection to the channel at the approximate location of the desired record based on the information obtained with the read sector operation. This command does not guarantee record orientation. Search commands must still be used for this function.

Set Path Group ID transfers the path group ID from main storage to the storage director. The path group ID is used to identify the system control program governing a group of channels. It is used with the Device Reserve command to reserve a 3380 device (an access mechanism) for the channels encompassed by the path group ID.

Recalibrate causes the addressed device to seek to cylinder zero, head zero.

Space Count provides a means of recovering or bypassing a defective count area. It is used for data recovery.

No Operation (No-Op) causes no action at the addressed device. It is usually used as a place holder in a channel program that may be modified. It is sometimes used to force channel-end and device-end status to occur together at the end of a channel program by delaying the end of a channel program. For example, a No-Op after a write count, key, and data operation causes channel end and device end to be delayed until after the pad to index.

Restore causes no action at the addressed device. It is maintained for compatibility with other IBM disk storage devices.

Diagnostic Control transfers four bytes of diagnostic information from the channel to the 3880 storage director.

Transfer-in-Channel (TIC) provides chaining capabilities for CCWs that are not located in adjacent main storage locations. It does not initiate any I/O operation at the channel. Instead of a sequential processing of the CCWs in a chain, the TIC causes a branch (for example to repeat a search operation until it is satisfied) and then resumes processing of the chain.

Search Commands

Search Home Address Equal causes the storage director to compare the home address (CCHH only) from main storage with the home address on the track.

Search Identifier Equal causes the storage director to compare the record ID (CCHHR) from main storage with the record ID in the count area of the next record encountered on the track. If the command is repeated, the storage director continues to make comparisons until the search condition is satisfied or until the track index is reached twice. The search can be continued on multi-tracks, if so directed, until the end of the cylinder.

Search Identifier High operates the same as Search Identifier Equal except for the type of comparison.

Search Identifier Equal or High operates the same as Search Identifier Equal except for the type of comparison.

Search Key Equal operates the same as search identifier commands except the comparison is with the key area of the record.

Search Key High operates the same as Search Key Equal except for the type of comparison.

Search Key Equal or High operates the same as Search Key Equal except for the type of comparison.

Read Commands

Read Home Address transfers the home address area (FCCHH only) from the track to main storage.

Read Record Zero (R0) transfers the count, key, and data areas of record zero from the track to main storage.

Read Count transfers the count area of the record to main storage.

Read Data transfers the data area of the record from the track to main storage.

Read Key and Data transfers the key and data areas of the record from the track to main storage.

Read Count, Key, and Data transfers the count, key, and data areas of the record from the track to main storage.

Read Multiple Count, Key, and Data transfers the next record encountered (excluding record 0) and all remaining records on the track to main storage. It provides a means for reading all the records on a track in a single disk revolution. (This command cannot be used with the multiple track feature, described later in this section.)

Read Initial Program Load causes the addressed device to seek to cylinder zero, head zero, and to read the data area of record 1. The load is initiated from the system operator console.

Read Sector transfers the sector number of the last traversed count area to main storage.

Write Commands

The first three of the following commands are used to format records on the track. The next two commands are used to update a record that has already been formatted.

Write Home Address causes the home address area (FCCHH only) to be transferred from main storage and written on the track. Use of this command should be limited to flagging defective tracks and assigning alternate tracks, using the Device Support Facilities program.

Write Record Zero (R0) causes the count, key, and data areas of record zero to be transferred from main storage and written on the track. If it is the last command in a chain, a Write Record Zero command erases the remainder of the track.

Write Count, Key, and Data causes an entire record to be transferred from main storage and written on the track. This command formats a user data record by establishing the key length (KL) and data length (DL) that is allocated on the track for the record. The lengths are contained in the count area of the record. They are the maximum lengths that can be written with subsequent update write commands.

Write Data causes the specified data in main storage to be written in the data area of the selected record. This command is used to update a record.

Write Key and Data causes the specified data in main storage to be written in the key and data areas of the selected record. This command is used to update a record.

Erase writes zeros in the count, key, and data areas of the specified record, then pads the remainder of the track with zeros. The erased record can no longer be read by the system.

Sense Commands

Sense transfers 24 bytes of sense information describing the cause of a unit check status from the storage director to main storage. (Refer to the section, Device Status and Error Handling.)

Sense Identification (ID) transfers seven bytes of sense information describing the type and model of the storage control and device being addressed from the storage director to main storage.

Byte	0	1	2	3	4	5	6
	FF	38	80	*	33	80	**

FF is a hexadecimal notation.

** Indicates commands implemented and 3880 model number.*

*** Indicates if dynamic path selection is installed.*

Sense Path Group ID transfers 12 bytes of information describing the path group identification and its state from the storage director to main storage.

Read and Reset Buffered Log transfers 24 bytes of statistical information from the storage director to main storage.

Device Reserve reserves the addressed device (access mechanism) to the channel(s) issuing the command, and transfers 24 bytes of sense information from the storage director to main storage.

Unconditional Reserve performs the functions of the Device Reserve command and, in addition, reserves the device

to the issuing path even when the device was reserved or in use through another path. This command causes 24 bytes of sense information to be transferred from the storage director to main storage.

Device Release terminates the reservation of the addressed device, and transfers 24 bytes of sense information from the storage director to main storage.

Diagnostic Commands

Diagnostic Sense/Read transfers the contents of the trace/dump buffer from the storage director to the channel.

Diagnostic Read Home Address transfers the entire home address area from the track to main storage. This command is used by the Device Support Facilities program.

Diagnostic Write Home Address causes the entire home address area to be transferred from main storage and written on the track. This command is used by the Device Support Facilities program.

Command Sequence

Channel command words (CCWs) can be chained so that a series of operations in successive main storage locations can be performed as a single channel program. Command chaining is indicated with a flag bit setting in the CCW. Data chaining also can be implemented by the setting of a flag bit. Data chaining provides for reading or writing data that is not in contiguous main storage locations.

A channel program consists of a chain of commands in various combinations, such as the following general sequence.

Seek
Set File Mask
Set Sector
Search ID
TIC* – 8
Read or Write
Read Sector

Commands for Speed Matching Buffer

New commands are used when the speed matching buffer for 3380 feature is used in the 3880 Storage Control Models 2 and 3. The objective of the new commands is to give advance notification of a write operation so that data transfer can begin from main storage to the buffer. With a portion of the data resident in the buffer, a write command can be issued as soon as

the requested sector and the correct record are reached. This synchronization of data transfer averts the possibility of missing a revolution of the disks.

Use of the Define Extent and Locate Record commands supplants the use of the Seek, Set File Mask, Set Sectors, and Search ID commands. The new write commands are for writing multiple equal length records. These commands are also supported in read operations. The related commands for enhanced performance for the speed matching buffer for 3880 feature are not supported by VSE/Advanced Functions or by the OS/VS1 operating system.

Define Extent provides the file mask, record (block) size, and limits within which subsequent commands in the channel program can operate. (The limits are specified by the beginning and ending CCHH addresses of the extent.)

Locate Record specifies the address (CCHH) of the record on the volume, the number of records to be processed, and the type of operation to be performed. An optional transfer length factor may be used by programs dealing with variable length records. The transfer length factor replaces the block size of the Define Extent command or a transfer length factor retained from a previous Locate Record command.

Write Update Data causes data to be transferred from the channel and written in the data area of a previously formatted record.

Write Update Key and Data causes data to be transferred from the channel and written in the key and data areas of a previously formatted record.

Write Count, Key, and Data Next Track erases the remainder of the current track; switches to the head on the next track; verifies that there is a record following home address; skips over that record; and writes count, key, and data areas.

Commands for Cache

Special commands are used for controlling cache storage with the 3880 Storage Control Models 13 and 23. The control of cache storage is not supported in the VSE/Advanced Functions environment.

Define Extent defines limits on subsequent operations of the channel program and indicates global access intent information.

Set High-Performance Storage Limits establishes a high-performance storage area for the addressed device or terminates the high-performance storage areas in the subsystem.

Set Subsystem Mode activates the cache operations for the addressed device or makes the subsystem storage available or unavailable to the subsystem. When data is transferred with read or write commands, it may be transferred to the 3880 cache or to the device.

Sense Subsystem Status transfers the status of the subsystem to the channel.

Sense Subsystem Counts transfers the subsystem performance data from the storage director to the channel.

Standard Features

The 3380 supports the following standard features commonly used in disk storage operations to improve performance. The features are described in detail in the *3880 Storage Control Description* manuals.

The 3380 does not support record overflow, sometimes referred to as track overflow.

Command Retry is a channel and storage director procedure that causes a channel command to be automatically retried. The retry does not cause an I/O interruption in the processor, and programmed error recovery procedures are not required. One of its uses is to correct certain data and seek errors without involving system recovery procedures.

End-of-File provides a means to define the end of a logical group of records. An end-of-file record is written by executing a Write Count, Key, and Data (or a Write R0) command with a data length specified as zero. This command causes the 3380 to write a data area consisting of 20 bytes of zeros.

Multiple Track Operations allows the storage director on all search and most read commands to select the next sequentially numbered head. This eliminates the need for Seek Head commands in a chain of read or search commands.

Format Write Release (or Write Padding) frees the channel and storage director while the 3380 erases to the end of the track after a format write command (Write Count, Key, and Data; Write Record Zero or Write Home Address).

Rotational Position Sensing (RPS) allows a search command to be started just before the required record comes under the read/write head, instead of starting the search at a random location of the track. It is based on a division of the track into evenly spaced angular sectors. The channel and storage director can disconnect while the track rotates to a specified sector location, thus permitting certain I/O operations to be overlapped.

Speed Matching Buffer in 3880 Models 2 and 3

With a speed matching buffer in the 3880 storage director, data is transferred between the buffer and the 3380 at the 3380 transfer rate of 3.0 MB/sec, and between the storage director and channel at the speed supported by the channel. If the channel supports data streaming, the speed is 3.0 MB/sec; otherwise, it is 1.5 MB/sec. The related commands for enhanced performance for the speed matching buffer for 3380 feature is not supported by VSE/Advanced Functions or by the OS/VS1 operating system.

Operations for 1.5 MB/sec

The following information regarding the speed matching buffer applies when the channel speed is 1.5 MB/sec.

Read Operations

A read operation can be performed with the count, key, and data commands, with or without the new Define Extent and Locate Record commands. Data is read from the 3380 into the buffer at 3.0 MB/sec and out of the buffer onto the channel path at 1.5 MB/sec.

Write Operations

A write operation can be performed with count, key, and data commands, with or without the Define Extent and Locate Record commands, but there may be a difference in performance.

Without use of the Define Extent and Locate commands, a write operation CCW chain causes the device to orient so the read/write head is positioned before the record to be written. The Write command is then issued, but the storage director does not proceed with the write operation because the necessary amount of data is not yet available. The storage director disconnects from the channel and then reconnects with enough lead time to load the buffer with the necessary amount of data so the data can be written successfully. A disk revolution (16.7 milliseconds) is lost in this process and if several write commands are chained together, a revolution is lost on each write command.

With use of the Define Extent and Locate Record commands, this performance degradation is avoided. The new commands notify the storage director in advance of a write operation of the disk location to be used and the expected size of the records. Because the storage director is notified prior to the actual write operation, it can start the transfer of data to the buffer early enough to ensure that at least one-half the data has been transferred to the buffer prior to reaching the record position on the rotating disk. When at least

one-half the data is resident in the buffer, the write operation can be performed successfully when the correct record location is reached on the disk.

Programming Support

The Define Extent and Locate Record commands are included in the MVS/SP operating system and many of the supporting products. Specific functions that include these commands are:

- MVS paging and swapping
- VM paging
- JES2 and JES3
- VSAM (including catalog)
- SAM-E and SAM of Data Facility Product
- OSAM in DL/I
- Sort/Merge
- Indexed VTOC of DFDS and DFP
- Device Support Facilities
- IEBCOPY

All of these functions build a Locate Record command into channel programs. Those that are IOS drivers under MVS also build in a Define Extent command. The Define Extent and Locate Record functions are not supported in the VSE/Advanced Functions environment.

The access methods ISAM, BDAM, and SAM (prior to SAM with extended functions) are not specifically supported. However, the EXCP driver builds a Define Extent and Locate Record prefix for certain write command sequences. This is not done for any search key type channel commands prior to a write operation.

EXCP prefixes a Define Extent to all channel programs. EXCP also prefixes a Locate Record command if the channel program begins with a Set Sector command (optional) followed by a command sequence of Search ID Equal, TIC, and one or more write commands. The write command may be Write Count, Key, and Data; Write Key and Data; or Write Data. (The Locate Record is prefixed only to the first write command.) However, the Locate Record command is not prefixed to the channel program if any command of the sequence has a program-controlled interrupt (PCI) or suppress-length interrupt (SLI) flag bit set. If the channel program begins with a Locate Record command, EXCP adds only a Define Extent command for the write operation.

The I/O supervisor adds a Seek and Define Extent command to all channel programs destined for a speed matching buffer if the program does not already have a Define Extent command.

When constructing a Locate Record command to be prefixed to an existing channel program, EXCP uses the count information in the write command to establish the Locate Record transfer length factor. A mismatch between the transfer length factor and the lengths contained in the count area on the track (established with a Write Count, Key, and Data command) results in an intent violation that is mapped by EXCP into a command rejection. (An intent violation is indicated by the setting of sense byte 2, bit 4.) Command rejection of update write commands (Write Key and Data or Write Data) can be avoided by setting the SLI flag to deliberately inhibit the prefixing of the Locate Record command.

Possible Anomalies

Due to the asynchronous nature of speed matching buffer operation at 1.5 MB/sec, some anomalies can occur. An example of this is a read operation that uses program-controlled interrupt (PCI) fetching of load modules. The gap time between records when the data is transferred out of the buffer is less than half the gap time when the data actually is read from a 3380 track. With PCI and the use of very short records, an increase in the number of interruption conditions can result, because the shorter gap time does not allow the PCI process to extend the channel command chain. Relinking the load modules to a larger block size will relieve this particular problem.

Operations for 1.5 and 3.0 MB/sec, Same Processor

In a configuration where two channels with different speeds access the same 3380 strings from the same processor, performance is optimized by using the 3.0 MB/sec channel as often as possible. One way to do this is to keep the 1.5 or 2.0 MB/sec channel path offline so that it is truly only a backup path. For the MVS/SP Version 1 operating system, another way to cause the 3.0 MB/sec path to be used when it is available is to specify the sequential option for physical path selection. The 3.0 MB/sec path will always be tried first because it has a lower channel number designation. An exception is the 4341 Model Group 12 where the lower channel number is a 2.0 MB/sec channel.

Model AA4 Attached to 3880 Model 13 or 23

When attached to a 3880 Storage Control Model 13 or 23, frequently used data from a 3380 Model AA4 string can be retained in a cache in the 3880. The cache is electronic storage and therefore provides for faster data access.

When data is first requested from the 3380 for processing, it is transferred from the 3380 to main storage in the usual manner. A copy is then made for retention in the 3880 cache where it is available when next requested. The data is retained in the 3880 cache until space is needed for additional data. The least recently used data is then deleted from the cache to make room for a copy of data newly requested for processing. (As with all disk storage operations, once data is written on the disk it remains there until written over or erased.) System and application data sets can be placed in the cache, but the cache is not for optimized paging. Data also can be specifically placed in the 3880 cache by the programmer and retained there until released by the programmer.

Attachment of 3380 to 3880 Model 13 or 23 is supported by the MVS/SP Version 1 and MVS/SP Version 2 (MVS/XA) operating systems.

The 3380, when attached to a 3880 Model 13 or 23, can be used without dynamic path selection and provide a capability to simultaneously access two access mechanisms. The related commands for enhanced performance for the speed matching buffer for 3380 feature and the dynamic path selection function are not supported in the VSE/Advanced Functions environment.

The 3880 Model 13, Model 23, and Models 2 and 3, are described in different manuals, as listed in the bibliography of this manual.

Overlapping Operations

As with other disk storage, certain operations can be overlapped at the various access mechanisms in a 3380 string. Specifically, one access mechanism in a string can be using the path to the channel for a transfer operation, while at the same time, seek or set sector operations, already initiated by channel commands, continue at the other access mechanisms in the string. This type of overlapping is possible because the channel disconnects while seek and set sector operations are being performed by the device. With the 3380, this overlapping of operations can include access mechanisms at the same drive or different drives. For example, one access mechanism at a drive can be transferring data while another access mechanism is performing a set sector operation.

This overlapping of transfer and seek or set sector operations applies to 3380 access mechanisms controlled by both Models A04 and AA4. In addition, command and data transfers can be overlapped with Model AA4; that is, two transfer operations can be in progress at the same time.

Internal Paths and Device Access

There are four internal paths between a 3380 controller and the access mechanisms in a string. Up to two drives (four access mechanisms) are arranged on each of the four paths. For example, in Figure 19, access mechanisms 0, 1, 8, and 9 are arranged on path 0.

With a 3380 Model A04, with one controller, there is one external path between the controller and storage director, and only one of the four internal paths can be active. As shown by the broken line in Figure 19, internal path 2 is the only active internal path. It can be transferring data between the controller and one of its four access mechanisms: 4, 5, C, or D.

Dynamic Path Selection, Model AA4

With a 3380 Model AA4, with two controllers, the same four internal paths to the access mechanisms can be serviced by either controller.

Each of the two controllers attaches to a different storage director, and the two external paths between the controllers and storage directors can both be used to access the internal paths.

Only one controller at a time can have access to a particular internal path. For example, in Figure 20, if one controller is transferring with access mechanism 0, access mechanisms 1, 8, and 9 (which occupy the same internal path) are not accessible to the other controller until the transfer on the path is complete. Dynamic path selection governs access to the internal paths by the two controllers.

Two of the four internal paths to the access mechanisms can be actively transferring data. An access mechanism on one internal path can be transferring data, while at the same time, an access mechanism on any one of the other three paths can also transfer data. As shown by the broken lines in Figure 20, two simultaneous transfer operations are occurring. One transfer is with access mechanism 4 on internal path 2 through one controller

and storage director, and another transfer is with access mechanism A on path 1 through the other controller and storage director. While two access mechanisms are involved in transfer type operations, other access mechanisms in the string may be carrying out seek or set sector type operations.

If there are two Model AA4 strings, both sets of controllers are attached to the same pair of storage director-channel paths, therefore, there are still two paths for transfer to the two strings.

If the controller designated in the I/O address is busy or disabled, the dynamic path selection allows an alternate path to the device to be established through another storage director and the other controller in the Model AA4.

Logic for these two dynamic path selection functions, simultaneous transfer and use of an alternate path, is in the 3380 Model AA4. These functions are supported by the MVS/SP Version 1 and MVS/SP Version 2 (MVS/XA) operating systems. The related commands for enhanced performance for the dynamic path selection function of the 3880 is not supported by VSE/Advanced Functions or by the OS/VS1 operating system.

Dynamic path selection is different than the string switching of previous device types. String switching permits transfer of data over only one path at a time to a given string. However, one transfer can be made to one string and another transfer can be made to a second string at the same time over a different channel-storage control path when there is both channel and string switching. Dynamic path selection uses two entirely independent paths, and both can transfer data at the same time to a given string.

Some previous device types also have two controllers. But with these other types only one of the controllers is active at a time and use of the alternate controller requires manual intervention. With the 3380 Model AA4, both controllers are active, and if one should become unavailable, the other controller is used as an alternate automatically.

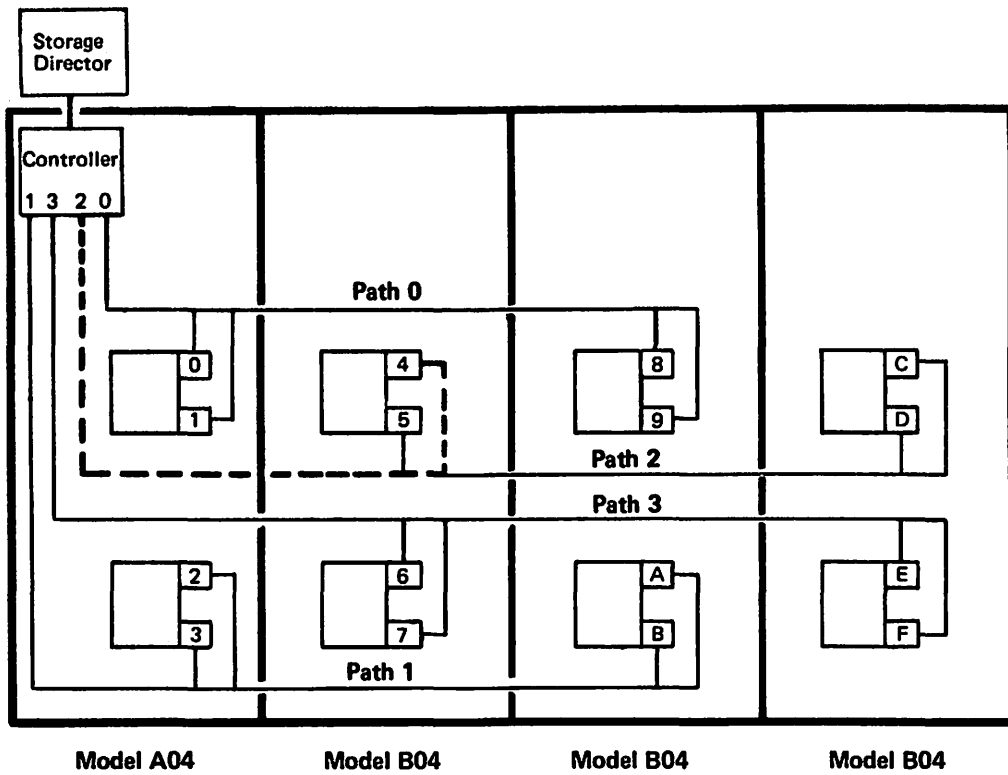


Figure 19. Four Internal Paths in a 3380 Model A04 with One Controller (dashed line represents one path that can be active)

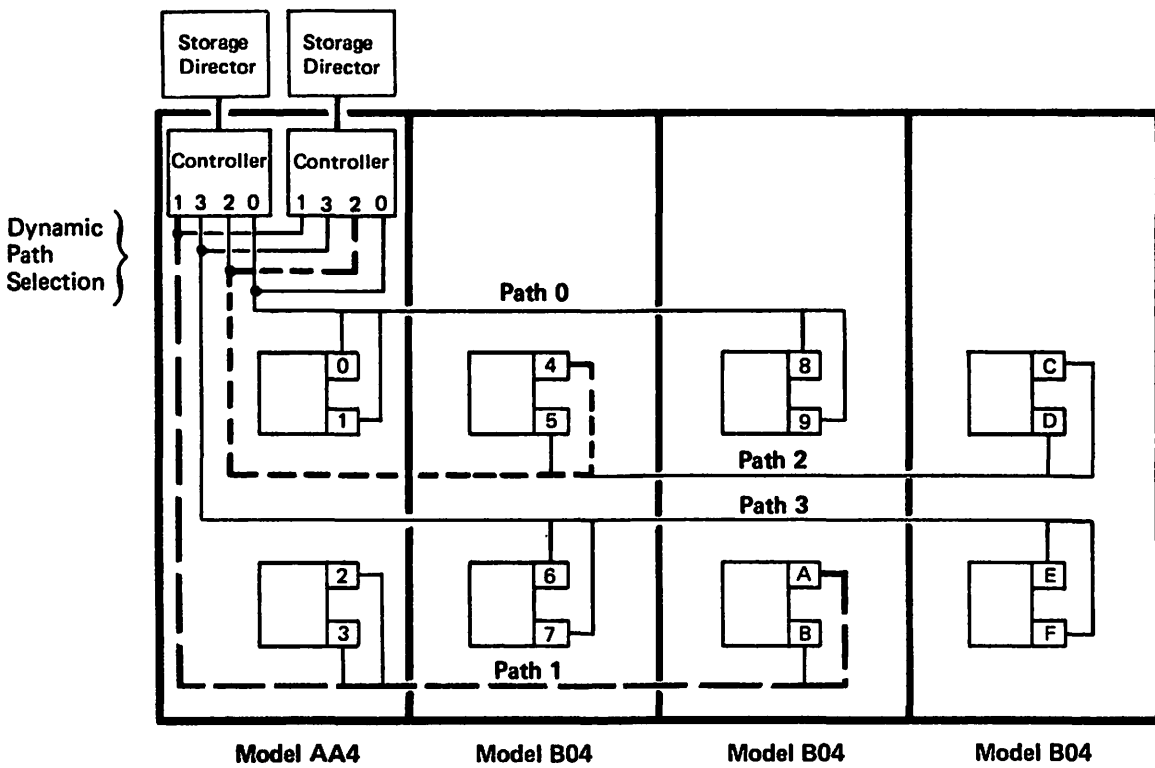


Figure 20. Four Internal Paths and Dynamic Path Selection in a 3380 Model AA4 with Two Controllers (dashed lines represent two of the paths that can be active simultaneously)

Device Reserve and Release Changes

The Device Reserve channel command causes an entire volume (space accessible by one access mechanism) to be reserved. This prevents access by any other requester to the volume until released with a Device Release command.

The reserve function is intended to prevent simultaneous attempts to update data on a volume. For example, a Device Reserve channel command is typically used when a VTOC or catalog is to be updated. With a Device Reserve command appended to the front of a channel command chain, the device is reserved for use only over the single, selected path. An Unconditional Reserve channel command allows suspending of the reserve on the original reserving path and applying the reserve to another path. This allows access to the device when there is no access over the original reserving path.

Two functions have been added to improve access to a shared volume in a multisystem environment. Neither function is supported by the OS/VS1 operating system.

Reserve Resource on Shared Volume

With the MVS/SP Version 1 (Release 2) and MVS/SP Version 2 (MVS/XA) operating systems, a global enqueue (ENQ) macro request allows processors that share a volume, in effect, to agree that a resource on the shared volume is reserved by one processor, while the rest of the volume is still available to the other processors. A typical resource might be a data set. This type of communication between processors requires channel-to-channel adapters. Global resource serialization permits data set level serialization and avoids volume level serialization of the reserve/release function.

System Related Reserve, Model AA4

A device in a 3380 Model AA4 string can now be reserved to a group of paths belonging to a system, instead of only to a single path. The paths in a group have a common, system identifier. Any path in the group can access data sets on a reserved volume, but paths not encompassed in the path group are excluded from access to the entire volume. (This function is implemented by the system and does not require any specification by the user.)

Logic for system related reservation is associated with dynamic path selection in the 3380 Model AA4. System related reserve is supported by the MVS/SP Version 1 and MVS/SP Version 2 (MVS/XA) operating systems. Because system related reserve is associated with dynamic path selection, this function is not supported by VSE/Advanced Functions or by the OS/VS1 operating system.

The Device Reserve and Device Release channel commands are used as before, but now, if one path is not available for the operation, another path in the group is used automatically. Any path having the same path group identifier as the system that issued the Device Reserve channel command is eligible. An attempt to use a path that is not in the group results in a device busy until a Device Release command is issued by the same system that issued the Device Reserve.

With a group of possible paths to a reserved device, use of the unconditional reserve function is significantly reduced.

Dynamic Reconnection, Model AA4

A device in a 3380 Model AA4 string can be dynamically reconnected to the first available path in a group of paths belonging to a system, instead of only to the path that originated the operation. The paths in a group have a common, system identifier. Dynamic reconnection is possible for devices operating under control of a system using System/370 extended architecture.

Logic for dynamic reconnection is associated with dynamic path selection in the 3380 Model AA4. Dynamic reconnection is supported by the MVS/SP Version 2 (MVS/XA) operating system. Because dynamic reconnection is associated with dynamic path selection, this function is not supported in the VSE/Advanced Functions environment.

Dynamic reconnection can be made after the device has completed an independent operation, such as a seek, during which it is disconnected from the system. The reconnection can be made to the first available path in a group of paths that have a common system identifier.

Dynamic reconnection is possible both for devices that are reserved and devices that are not reserved.

This increase in opportunities for reconnection represents an improvement in performance.

Dynamic Path Selection Controls, Model AA4

No specifications are required by the user for dynamic path selection.

Dynamic path selection is installed in all 3380 Models AA4. In response to a Sense ID channel command, the subsystem sends to the requesting system seven bytes of information that describe the subsystem. If byte 6, bit 6 is set, dynamic path selection is installed, and the supporting systems implement dynamic path selection functions.

Path Group Identifier, Model AA4

Each system sharing a 3380 Model AA4 has a unique path group identifier (ID) for dynamic path selection functions. The path group ID is a common, system identifier. The same path group ID is assigned to all of the paths a system has to a device. The principle use of the path group ID is for the dynamic path selection functions of system related reserve and dynamic reconnection.

A system path-group ID is associated with each path to each device and is stored in an array, in the 3380 Model AA4. The subsystem uses information in the array for determining which paths have access to a reserved device and which paths can be used to dynamically reconnect a device that has been disconnected.

To enter information in the 3380 array, a system issues a Set Path Group ID channel command. This command carries an 11-byte system identifier and a flag byte to indicate whether a member of a group is to be established or removed.

A group of paths is established by a system with system initial program load (IPL) and maintained with the VARY system operator commands. The VARY (device), VARY PATH, and VARY CHANNEL forms of the VARY command include a new function, when used for a 3380 with dynamic path selection, to invoke the Set Path Group ID command. Because this function is associated with dynamic path selection, it is not supported by VSE/Advanced Functions or the OS/VS1 operating system.

In response to a Sense Path Group ID channel command, the subsystem sends to the requesting system the path group ID that was sent in the Set Path Group ID command subsequent to the last system initial program load (IPL) or 3380 initial microcode load (IML).

Synchronization between information stored in the array and information used by the operating system must be maintained. This synchronization is maintained by following the operating procedures recommended in the section, Operator Controls and Commands.

Model AA4 with OS/VS1

The OS/VS1 operating system supports the 3380 Model AA4. However, OS/VS1 does not support dynamic path selection. Both storage director-controller paths of the Model AA4 can address all the access mechanisms in a string, but only one of the paths is used.

Model AA4 with VSE/Advanced Functions

In a VSE environment, the IBM 3380 Model AA4 can be used without dynamic path selection and provide the capability to simultaneously access two access mechanisms. The related commands for enhanced performance for the speed matching buffer for 3380 feature and the dynamic path selection function are not supported in the VSE/Advanced Functions environment.

Model AA4 with VM/SP or ACP/TPF

Both VM/SP and ACP/TPF operating systems support attachment of 3380 Model AA4 but not the full functions of dynamic path selection. With Model AA4 attachment, both storage director-controller paths have access to all the access mechanisms in a string.

With VM/SP, a Model AA4 configuration can be generated that defines a primary path through one storage director-controller path for some devices in the string and a primary path through the other storage director-controller for the rest of the devices in the string. For example:

```
RDEVICE ADDRESS=(350,8),ALTCU=450
RDEVICE ADDRESS=(458,8),ALTCU=350
RCTLUNIT ADDRESS=350,FEATURE=16-DEVICE
RCTLUNIT ADDRESS=450,FEATURE=16-DEVICE
RCHANNEL ADDRESS=3
RCHANNEL ADDRESS=4
```

Because VM/SP begins path selection for an I/O request with the primary path, and the devices in the string do not all have the same primary path, some load balancing can be achieved. All devices in the string are accessible by either storage director/controller, when the addressed controller is available, and both external paths can transfer at the same time. (Addressing considerations are described in the section Hardware Planning Considerations.)

When controlled by VM/SP or ACP/TPF operating systems, there is no switching to the other controller if the addressed controller is not available.

Hardware Planning Considerations

Careful planning will help to ensure the successful installation of a 3380. A customer engineer (installation planning representative) can assist with the planning, and should be notified early in the planning stage of the installation.

Processor Attachment Requirements

Attachment of the 3380 to any processor requires installation of an IBM 3880 Storage Control Model 2, 3, 13, or 23. The 3880 Model 3, 13, or 23 is required if both storage directors are used for 3380s. The Model 2 is required if one storage director is used for 3380s and the other storage director for other types of disk storage.

Data streaming or a speed matching buffer is required.

The speed matching buffer for 3380 feature can be installed in one storage director of a 3880 Model 2 and both storage directors of a 3880 Model 3, but a speed matching buffer cannot be installed in a 3880 Model 13 or 23. (If only one speed matching buffer is to be installed in a 3880, it must be in the number 2 storage director.)

Attachment of the 3380 to an IBM System/370 Model 158, 158-3, 168, or 168-3 always requires installation of the speed matching buffer for 3380 feature in the 3880 storage director.

Attachment of the 3380 to an IBM 3031, 3032, 3033 Processor, or 3042 Attached Processor Model 2 requires installation of one of the following special features:

- Data streaming feature in the processor for attachment to a channel with a 3.0 MB/sec data transfer rate.
- Speed matching buffer for 3380 feature in the 3880 Storage Control Model 2 or 3, for attachment to a

channel with a 1.5 MB/sec data transfer rate. The speed matching buffer in a 3880 storage director also supports channels with a 3.0 MB/sec data transfer rate.

Attachment of the 3380 to an IBM 4341 processor channel that has a 3.0 MB/sec data transfer rate does not require a special feature for data streaming. However, 4341 models manufactured before March 1981 require an engineering change to achieve data streaming. The customer engineer assisting with installation should arrange for this change. Data streaming is standard on 4341 models manufactured after March 1981. The speed matching buffer for 3380 feature can be used for 4341 channels that have less than the 3.0 MB/sec data transfer rate.

The 3380 can be attached (through a 3880) to an IBM 4381 Processor channel with a 3.0 MB/sec data transfer rate. This does not require a special feature for data streaming. Data streaming is standard on 4361 and 4381 Processors.

There are requirements as to which channels of a processor may have data streaming and which channels can connect to 3880 storage directors with the speed matching buffer. The requirements depend on the type of processor. Consult your market representative for the requirements that apply in your installation.

Attachment of the 3380 to IBM 3081, 3083, and 3084 Processors is always through a 3.0 MB/sec channel. Data streaming is standard on all block multiplexer channels of these processors.

Additional information may be found in the sections Models and Attachment Configurations, and Operations.

Physical Planning

The following information highlights some of the significant physical planning considerations for installing the 3380. Detailed physical planning information is in the *IBM Input/Output Equipment Installation Manual – Physical Planning for System/360, System/370, and 4300 Processors*.

Power and Air Conditioning

A 3380 string requires approximately 71 percent less power, in kilo Volt Amp (kVA), than 3350s for equivalent storage capacity. A single 3380 model A unit uses four percent more power than a 3350 unit, but has almost four times as much storage capacity. A reduced power requirement translates into a reduced air conditioning load, which in turn may also reduce power consumption.

Space Requirements

A 3380 string requires approximately 66 percent less space than 3350s for equivalent storage capacity. That is, for 10 gigabytes of storage, the 3380 requires approximately 10 square meters (100 square feet); whereas the 3350 requires approximately 30 square meters (300 square feet). Figure 21 illustrates the space saving of the 3380 over 3330 Model 1, 3330 Model 11, and 3350 disk storage. The calculations are based on full string configurations, including allowance for one-half of the service clearance dimensions.

Cable Requirements

The total length of signal cable that can be used to connect two 3380 controllers to a 3880 storage director cannot exceed 61 meters (200 feet).

DASD Type	Capacity (GB)	Units Required	Floor Space Required		MB of Storage	
			Square Meter	Square Feet	Per Square Meter	Per Square Feet
3330-1	9.6	48	139.70	1503.48	68.72	6.39
3330-11	9.6	24	69.85	751.74	137.44	12.78
3350	10.16	16	27.60	297.08	366.13	34.20
3380	10.08	4	9.24	99.44	1091.14	101.37

Note: These values are based on space occupied by the equipment plus half of the service clearance dimensions area. Half the service clearance is used because adjacent strings may overlap service areas.

Figure 21. Comparisons of Approximate Floor Space Requirements Based on Full String Configurations

Addressing

Device Address

Eight bits are used for 3880 and 3380 device physical addresses, as described as follows and in Figure 22. There is supplementary reference information on addressing in Appendix B. Bit assignments are as follows:

Bits 0, 1, and 2 are the 3880 storage director address. A 3880 storage control (two storage directors) can have up to eight different addresses, one per channel interface.

The storage director addresses are set with switches by the customer engineer.

Bit 3 is the 3380 controller (Models A04 and AA4) address. The controller addresses depend on the number of strings and whether it is a 3380 Model A04 or AA4.

- If there is only one string attached to a 3880 storage director, the 3380 controller bit may be 0 or 1.
- If there are two strings attached to a 3880 storage director, the controller bit for one string must be 0 and the controller bit for the other string must be 1.
- If there is one string with two controllers in the Model AA4, both controllers must have the same bit address, either 0s or 1s. (Each controller attaches to different storage directors and each storage director attaches to different channels.)
- If there are two strings, each with two controllers Models AA4, the controller bit addresses in one string must be 0s and the controller bits in the other string must be 1s. (Each of the two strings attaches to the same two storage directors, and each of the two storage directors attaches to different channels.)

Note: If the second hexadecimal digit of the 3-hexadecimal digit address assigned by the system programmer at system generation is even, this controller bit 3=0; if the second hexadecimal digit is odd, this controller bit 3=1. See Appendix B.

Bits 4 through 7 are the access mechanism (device) address. The access mechanism addresses are 0 through F. For a given string, access mechanism addresses must start with 0 and must be sequential from the beginning of the string.

Addresses of controllers are established by the customer engineer when the 3380 is installed. Device (access mechanism) addresses are pre-established and individual access mechanism addresses cannot be changed.




Figure 22. 3880 and 3380 Addressing

Address Assignments

The 8-bit device physical address is expressed as two hexadecimal digits.

The high-order digit represents the storage director and controller address (bits 0 through 3).

The low-order digit represents the access mechanism (bits 4 through 7).

For example, a storage director and controller are addressed as B, with 16 access mechanisms sequentially addressed as 0 through F. Thus the addresses are B0, B1, B2 . . . BF.

Figure 23 lists the sets of two hexadecimal digits that are valid for the storage director-controller and access mechanism.

Address Ranges	
One String, 16 Access Mechanisms or Less	Two Strings, 32 Access Mechanisms or Less
00-0F	00-1F
10-1F	20-3F
20-2F	40-5F
30-3F	60-7F
40-4F	80-9F
50-5F	A0-BF
60-6F	C0-DF
70-7F	E0-FF
80-8F	
90-9F	
A0-AF	
B0-BF	
C0-CF	
D0-DF	
E0-EF	
F0-FF	

Figure 23. Sets of Valid 3880/3380 Addresses

The table in Figure 23 lists the address ranges for 16 and 32 access mechanisms. Addresses must begin at the beginning of the address range. For example, with two strings attached to a storage director, addresses can be 00, 01, 02 ... 0F for the first string and 10, 11, 12 ... 1F for the second string, but *not* 10, 11, 12 ... 1F and 20, 21, 22 ... 2F.

If less than the 16 or 32 access mechanisms are used, all of the addresses within the range are still exhausted for 3380 use. However, the unused addresses may be used for other device types on a different storage director attached to the same channel as the 3380s. For example, if a string of eight 3380 devices is assigned addresses E0 through E7, addresses E8 through EF are not available for assignment to another set of 3380 devices on the same channel. However, E8 through EF can be assigned to another device type, such as the 3350 or 3375, attached to a different storage director.

Figures 24 and 25 are examples of address assignments.

Defining Addresses at System Generation

At system generation, the physical address is defined for use by the software. For System/370 architecture, the device physical address corresponds to the second and low-order digits of a 3-hexadecimal digit I/O (unit) address. For System/370 extended architecture, the device physical address and channel-path identifier (CHPID) are associated with a device number. Refer to the applicable system generation manual for information on making these specifications.

Operating systems may impose certain constraints on addressing although they are not required by the hardware. For example, under MVS/SP Version 1 a single system must use the same address to select a disk storage, except that the channel portion may differ.

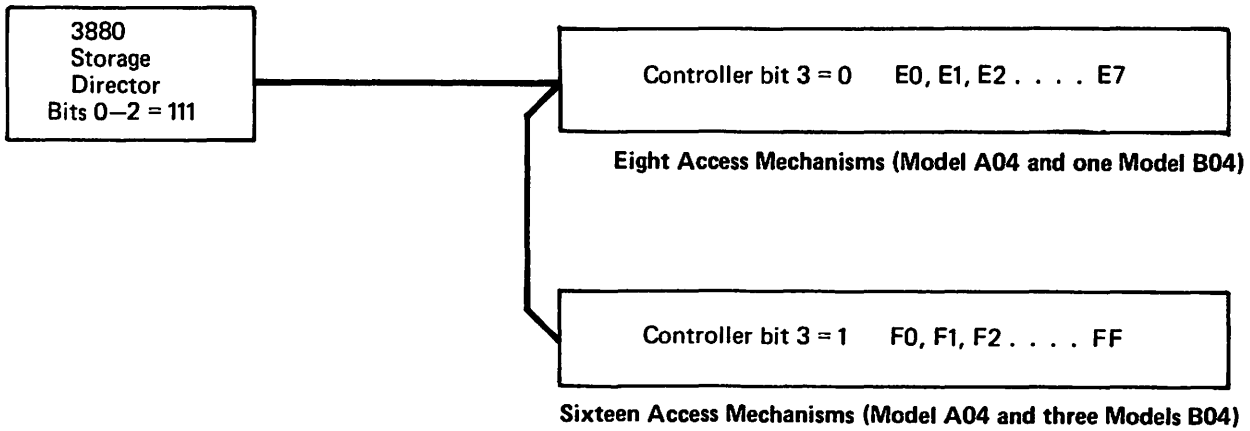


Figure 24. Addressing Example: Two Model A04 Strings Attached to One 3880 Storage Director




Figure 25. Addressing Example: Two Model AA4 Strings Attached to the Same Two 3880 Storage Directors

Addressed Selection Not Possible

An addressed storage director, controller, or device may not be selectable for different reasons. For the 3380, the different reasons are distinguishable in the system message issued at the operator console.

Message Issued	Reason Not Selectable
No Paths Available	Address bits 0-2 refer to a storage director that is not attached.
No Paths Available	Address bit 3 refers to a controller that cannot be selected by the 3880 during its checkout testing which immediately follows an initial micro-code load (IML).
Permanent I/O Error	Address bit 3 refers to a controller that was previously selected but now is not available because of a path communication error, the controller Enable/Disable switch on the 3380 is set to Disable, or power is off. Unit check status is set in the channel status word and the equipment check sense bit (byte 0, bit 3) is set on.
Intervention Required	Address bits 4-7 refer to a device that is not installed. Unit check status is set in the channel status word and the intervention required sense bit (byte 0, bit 1) is set on.
Permanent I/O Error	Address bits 4-7 refer to a device that is physically installed but without power. Unit check status is set in the channel status word and the equipment check sense bit (byte 0, bit 3) is set on.

Seek Address

The seek address, transferred by the Seek channel command, specifies the cylinder and head location of the record on the addressed device. It is a 6-byte address, 00CCHH, where CC is the cylinder number and HH is the head number. (The first two bytes are always zero.)

The cylinder addresses are:

Data cylinders	0-884
Alternate cylinder	885

The head (track) addresses are: 0-14

Seek Address Not Valid

If a seek address is not valid, a permanent I/O error system message is issued at the operator console. Unit check status is set in the channel status word and the command reject sense bit (byte 0, bit 0) is set on.

Operator Controls and Commands

Most 3380 operations are controlled from a system operator console instead of from operator panels on the 3380.

There is one control panel on the end cover of the 3380 Model A04 and AA4 units. See Figure 26. The control panel is used primarily by a customer engineer, but it is also used by the operator to turn off power in an emergency.

System commands are issued from the system operator console to control 3380 operations and to obtain its operational status.

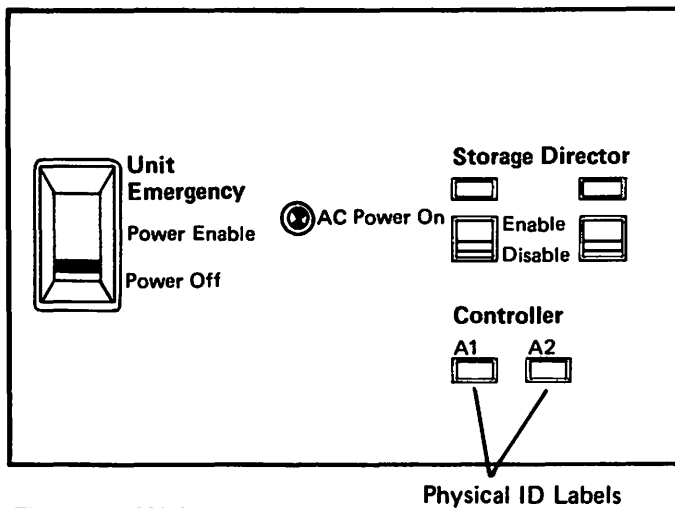


Figure 26. 3380 Control Panel

Power On and Off

Power for 3380 operations is usually controlled remotely from the system operator console when the system is powered on and off. The necessary physical power connections in the 3380 are made by the customer engineer and are transparent to the user.

There are two ways the 3380 can be powered on: by initiating a power-on sequence at the system operator console, or by turning on the power at each IBM 3380 Storage Control.

To initiate the power-on sequence at the central processor, the processor turns on power for each 3380 in sequence. Power sequencing for all units attached to a 3380 must complete before power sequencing for the next 3380 can begin.

If power is manually turned on at each 3380, power sequencing for each 3380 on the system can proceed independent of other 3380s on the system.

Warning: When the 3380 Power Sequence Complete light comes on, power sequencing for all 3380s attached to the 3380 is *not* complete. Wait three more minutes after the 3380 Power Sequence Complete light is on before assuming that all 3380s attached to that 3380 have completed the power-on sequence. If an attempt is made to IPL the system (to perform an initial program load of the operating system) before all 3380s have completed the power-on sequence, no IPL will occur. The resulting system messages are unpredictable.

The 3380-3880 subsystem is usually powered off at the 3380 operator panel. To allow for a complete power-off sequence, wait at least 90 seconds before turning the power back on.

Emergency Power Off

The Unit Emergency switch on the 3380 Model A04 or AA4 control panel should be used to turn off power in an emergency. When the switch is in the Power Off position, power is turned off to the entire string.

The Unit Emergency switch should not be used to turn power back on. An attempt to do this could result in incorrect power sequencing. Call a customer engineer to restore power. The Unit Emergency switch is not for use for normal power off and on by the user.

AC Power On Light

The AC Power On light on the 3380 control panel is lit when voltage is present. This does not necessarily mean the devices are operating.

Path and Device On and Off

Enable/Disable switches control physical availability of paths and strings of devices. The Enable/Disable switches on the 3380 control panel apply to the controller-storage director connection. These switches are set to the Enable position by the customer engineer. All panels have spaces for two Enable/Disable switches; but if there is only one controller in the 3380 (Model A04), only one of the switches is installed. Most of the controls performed with manual switches are done with switches at the 3380 storage director or a remote configuration controller, not at the 3380.

The VARY (device) and VARY PATH commands are used to control whether an access mechanism or path is available to the system. VARY (device) and VARY PATH commands are system operator commands. These commands can be issued through an OS/VS1, MVS/SP Version 1, or MVS/SP Version 2 system control program. As a general rule, all of the applicable Enable/Disable switches must be in the Enable position before the vary online versions of the command are issued, and the vary offline versions of the command should be issued before the applicable switches are disabled.

If the VARY PATH or VARY (device) Offline system command is not issued prior to setting the Enable/Disable switch on the 3380 to the Disable position, a false equipment check error message may be generated. (The system message is a permanent I/O error message indicating an equipment check error condition.)

A VARY PATH command applies to all of the components of a path including a channel, a 3880 storage director, and a 3380 controller to the specified 3380 devices (access mechanisms). A VARY PATH Online allows the system to attempt to use the specified 3380 controller and a VARY PATH Offline prevents the system from attempting to use the specified 3380 controller as a path to the specified device.

If one controller of a 3380 Model AA4 is not usable because of an equipment problem, the other controller will be used by systems that support dynamic path selection. However, the problem controller should be varied offline to avoid initial attempts by the system to use the problem controller.

A VARY (device) command applies to an access mechanism. A VARY (device) Online allows the system to use the device, and a VARY (device) Offline prevents the system from using the device from any path of the issuing system.

VARY Command and Path Group ID

For the 3380 Model AA4, the VARY command also maintains path group IDs in the array used for dynamic path selection functions. When a path is removed from the array with the VARY command, corresponding information is available to the operating system. However, a path can also be removed from the array by actions that cause hardware resets. If information used by the system is not changed prior to the reset, the system may repeatedly attempt to access a reserved device over a path that no longer has an eligible path group ID in the array. In this case, the reserved device will appear busy

even when it is not in use. In an MVS/XA operating system environment, the subsystem also will not be able to implement dynamic reconnection. This means the array in the Model AA4 that contains the path group IDs is out of synchronization with information used by the operating system.

Synchronization between the array and the system is maintained with the following procedures. See also Figure 27.

- Before the Enable/Disable switch for a storage director-controller on the 3380 control panel is set to Disable, all paths from all systems through the controller to the string of devices should be varied offline with the VARY PATH Offline command. The command should specify addresses of all devices in the string. After the switch is returned to the Enable position, issue the VARY PATH Online command.
- Before an Enable/Disable switch for a channel-storage director on a 3880 Storage Control is set to Disable, the path should be varied offline with the VARY PATH Offline command. (If the storage director is being disabled for servicing, every path through the storage director should be varied offline before the switches are set to Disable.) The command should specify addresses of all devices in the string. After the switch is returned to the Enable position, issue the VARY PATH Online command.
- Before there is an initial microcode load (IML) of a 3880 storage director, all paths from all systems through the storage director should be varied offline with the VARY PATH Offline command. The command should specify addresses of all devices in the string. After an IML, issue the VARY PATH Online command. (Setting the Subsystem Power switch on the 3880 Storage Control *off* and *on* causes an IML and resets.)

In the event the required procedure is not used and an out-of-synchronization problem occurs, action can be taken to identify the problem and restore synchronization.

For system related reserve, messages help determine that a problem exists. For dynamic reconnection, when the device is not reserved, there is no explicit indication that the array is out of synchronization. The first indication of a possible problem will probably be a Device End Missing system console message (for System/370) or a Start Pending system console message (for System/370-XA). Then, if messages issued as a result of the DISPLAY UNITS system operator command are A-BSY-R or O-BSY-R, an out-of-synchronization problem should

be suspected. If an out-of-synchronization condition does exist, these messages will be persistent.

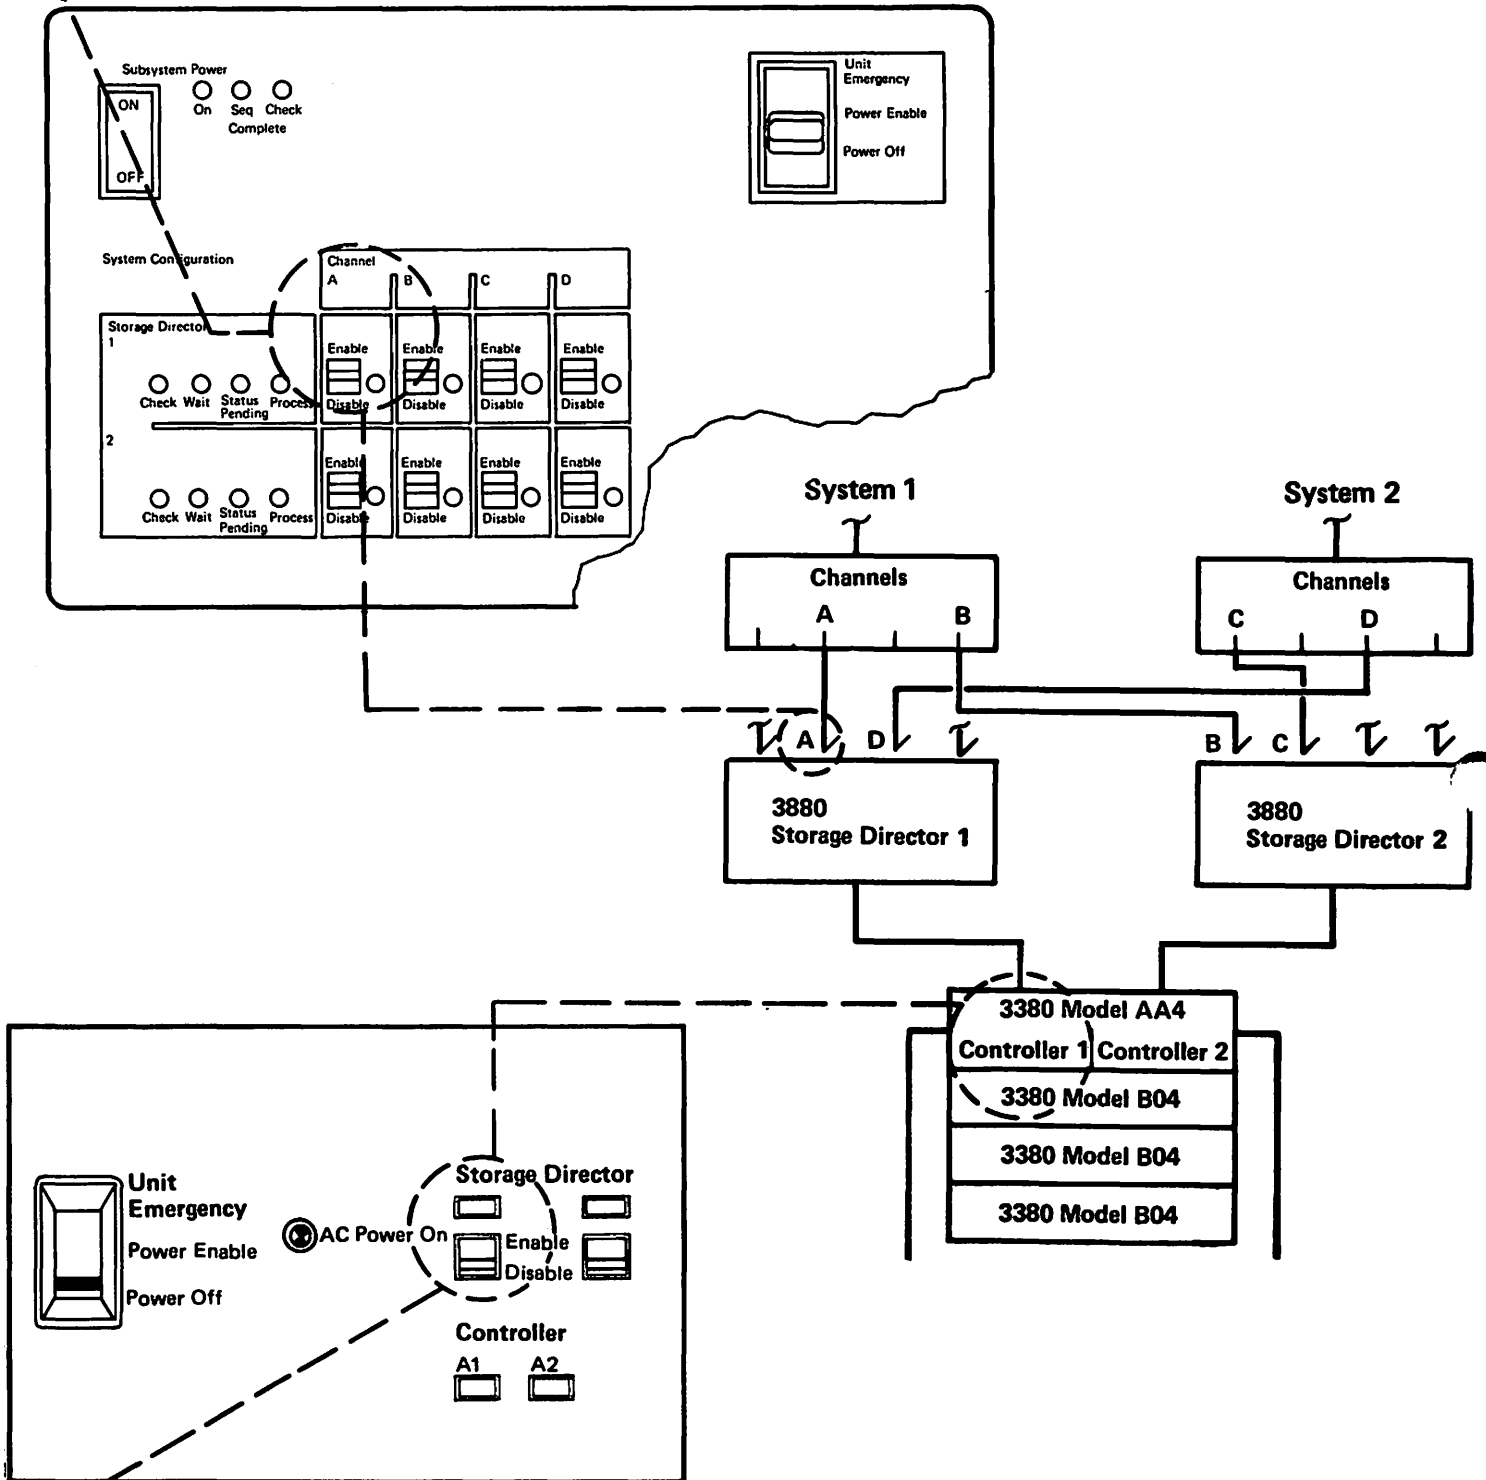
If information indicates an out-of-synchronization problem, use a VARY (device) Online command to vary online every device in the string. This action will refresh the array for all online paths to all devices and thus restore synchronization with all of the systems. A path varied offline with the VARY PATH command cannot be brought online with the VARY (device) Online command. When a path is varied offline, it can be brought back online only with the VARY PATH Online command. Successful restoration of synchronization can

be assumed if the Device End Missing or Start Pending system messages stop appearing.

An initial program load (IPL) by a system recreates the array and therefore maintains synchronization between the array and the particular system that is IPLed.

If information in the array is out of synchronization because of a hardware problem, a customer engineer is required. After the hardware is repaired, the array can be recreated with the VARY (device) Online command as described above or by a system IPL.

- Before setting the Channel A and Storage Director 1 switch on the 3880 to Disable, issue a VARY PATH Offline for path A-1.



- Before setting the Storage Director 1 and Controller 1 switch on the 3880 to Disable, issue a VARY PATH Offline for both paths A-1 and D-1.

Figure 27. Example of Standard Operating Procedures

Device Status

In response to an MVS/SP DISPLAY U system console command, a message is issued giving *unit* status. Messages may indicate the following status conditions: online, allocated, busy, and not ready. Other status conditions indicate whether the device is reserved and whether the volume contains system resident data sets. For VM operating systems, the Query DASD command gives device status.

There is no Ready light on the 3380.

For VSE/Advanced Functions, the VOLUME command provides the operator with a summary of the volumes on a DASD unit, together with an indication of which volumes are in use.

Missing Interrupt

There is no attention button on the 3380; therefore, missing interrupts cannot be resolved by generating an attention interrupt. With both MVS/SP and VM operating systems, most missing interrupt conditions are now handled by the operating system without operator intervention.

With VSE/Advanced Functions, however, operator intervention might be required, depending on the condition of the interrupt.

Physical Identifier Labels

Spaces are allocated on the 3380 control panel to affix labels of the physical identifiers, which are assigned to storage directors, controllers, and devices. Physical identifiers and their assignment are described in the section Device Status and Error Handling.

Installation and Initialization

At the plant of manufacture, the 3380 volume is initialized with a home address (HA) and a track descriptor in record zero (R0) on all tracks. (The HA and R0 are described earlier in the section Device Characteristics and Formats.)

All disks are analyzed for surface conditions. The 3380 is shipped from the plant of manufacture with no tracks flagged as defective. (A track flagged as defective requires use of an alternate track.) The 3380 may have bypassed defective areas on a track where data will not be recorded. Instead, data will be recorded at another location on the same track. Space for such defect displacements does not subtract from that available for user space calculations.

Special data is recorded on all disk surfaces to be used at installation for subsystem integration testing.

Installation

After installing the 3380, the customer engineer performs a subsystem integration test to ensure that data can be successfully transferred.

Initialization

After a new device is successfully installed and tested, the volumes must be initialized for use by the operating system. The Device Support Facilities (ICKDSF) program is provided for this initialization.

With the Device Support Facilities INIT command, a 3380 volume can be initialized to one of three levels. The three levels are referred to in the *Device Support Facilities User's Guide and Reference* as minimal, medial, and maximal. The first (minimal) level fulfills the requirements for volume use by the operating system. The second and third levels include the functions of the first level and add other functions. The following summary highlights the functions accomplished with the three levels.

Volume Label and VTOC (Minimal Level)

Each volume requires a volume label and volume table of contents (VTOC). See Figure 28. The first level of initialization includes:

- Writing initial program load (IPL) bootstrap records as records one and two, and a volume label as record three on cylinder 0, head 0. The volume label consists of a volume ID, volume serial, and the address of the starting location of the VTOC.

- Specifying the starting location and the size of the VTOC. The default for the VTOC location is cylinder 0, head 1, and the default for the VTOC size is one track.

Because it is frequently accessed, placement of the VTOC near the middle cylinder (cylinder 443) may have performance advantages.

As a general guide in calculating the size needed for the VTOC, determine the maximum number of data sets that will reside on the volume and divide this number by the number of data set control blocks (DSCBs) that can be placed on a track. For the 3380, 53 DSCBs can be placed on a track. If it is assumed that twice the number of data sets will reside on a 3380 volume as on an existing volume, approximately three times as many cylinders will be needed for the VTOC. For additional information on calculating the size of a VTOC, refer to the *Device Support Facilities User's Guide and Reference*.

During this first level of initialization, the Device Support Facilities program creates (formats) the VTOC.

HA and R0 Validation (Medial Level)

The second, or medial, level of initialization adds validation to ensure that the proper track addresses (CCHH) appear in all of the home address areas and record zeros. Validation rewrites the home addresses and record zeros, which causes the rest of the track to be erased.

This second level is required if for any reason a head and disk assembly is replaced. The Validate parameter of the INIT command should be used to initialize both volumes of the head and disk assembly to reformat the home address and record zero on all tracks.

Surface Checking (Maximal Level)

The third, or maximal, level of initialization adds surface checking for the condition of all tracks. With this check, special patterns are written and read to detect any surface defects that cause errors in reading data. Checking is first done with a set of primary patterns that accomplish a high-level sifting to expose possible errors. The suspected errors are then subjected to an advanced, and more time-consuming, set of patterns. Checking erases the contents of each track during initialization. If disk surface defects are confirmed during the advanced checking, Device Support Facilities automatically bypasses the defect. Areas causing both ECC correct-

able (temporary) and ECC uncorrectable (either temporary or permanent) errors are bypassed.

Up to seven defective areas are skipped, after which an alternate track is assigned. Space at other locations on the track is already reserved so skip displacements by ICKDSF do not subtract from space available for user records when making space calculations. If the space reserved for skip displacements is exhausted, the track is examined again. If there are too many defects, the track is flagged as defective and an alternate track is assigned.

Results When Checking All Tracks

If the maximal level of initialization is used, the following results should be understood.

When checking all tracks on a volume (using the INIT command), the temporary data errors detected may not appear to be consistent if the command is issued more than one time. Defects may not be detected when the function is performed the first time; but, if the command is issued again, defects may be detected and this inconsistency may persist on subsequent runs. This happens because the program first uses primary test patterns, which save time but do not permit as refined an analysis as the more advanced patterns. Any defects detected with these primary patterns are then subjected to the advanced tests and defects that are confirmed are bypassed. A defect not bypassed but detected on a subsequent run is associated with a very small surface irregularity. Such irregularities may be detected under slightly varying conditions. A data check resulting from this type of defect is correctable with the ECC and, therefore, is a temporary error. These small defects are not likely to become the source of permanent errors. If the temporary errors appear excessive, consult a hardware service representative (customer engineer).

Initialization Under VSE/Advanced Functions

The Device Support Facilities (ICKDSF) program, Release 1.6.0 and 1.7.0, is used to initialize and maintain the 3380. For details, see the *Detail Support Facilities User's Guide and Reference*.

Index Creation and Support

An index to the volume table of contents supplies the name of the data set and location of the entry in the VTOC. These location pointers are used instead of the usual method of sequentially searching through all of the information associated with data set entries in the VTOC, thus reducing the time for searching the VTOC entries. An index is most beneficial for volumes where the VTOC exceeds two tracks. An example is in a time-sharing option (TSO) environment, where a large number of data sets are on a single volume. An index is optional.

An index to the volume table of contents is created with the Device Support Facilities INIT command. There is no default for placement of the index on the volume but more effective performance may be obtained by locating it adjacent to the VTOC. The BUILDIX command changes an existing non-indexed format VTOC to the indexed format VTOC.

The *Device Support Facilities User's Guide and Reference* gives detailed information on how to calculate the size of a VTOC index.

Use of a VTOC index requires either the Data Facility Device Support program or the Data Facility Product program.

The concept of a VTOC index is not supported in the VSE/Advanced Functions environment.

Cylinder 0:0 Head

HA		R0			R1 and R2			R3		
F	Address	Count	Key	Data	Count	Key	Data	Count	Key	Data
	CCHH	CCHHR			CCHHR		IPL Bootstrap	CCHHR		Volume Label: Volume ID Volume Serial No. VTOC Address

Cylinder 0:1 Head

HA		R0			R1		
F	Address	Count	Key	Data	Count	Key	Data
	CCHH	CCHHR			CCHHR		Start of VTOC

Cylinder 443:4 Head

HA		R0			R1		
F	Address	Count	Key	Data	Count	Key	Data
	CCHH	CCHHR			CCHHR		Start of VTOC Index

Notes:

1. The home address (HA) contains the address of the track (CCHH) and a flag (F) indicating the condition of the track.
2. The record 0 (R0) count area contains primary and alternate track associations. For standard R0, there is no key, and the data area is eight bytes long.
3. Initialization information is specified with the Device Support Facilities INIT command for cylinder 0, head 0, records 1, 2, and 3.
4. The VTOC address (CCHH) is specified with Device Support Facilities INIT command. Cylinder 0, head 1 is the default for the starting location of the VTOC under OS/VS and MVS.
5. An index to the VTOC is specified with the Device Support Facilities INIT command. There is no location default.

Figure 28. Volume Layout of Initialization Information

Program Planning Considerations

The 3380 is supported by the following operating systems:

- Operating System/Virtual Storage Version 1 (OS/VS1), which implements System/370 architecture
- Multiple Virtual Storage/System Product Version 1 (MVS/SP Version 1), which implements System/370 architecture
- Multiple Virtual Storage/System Product Version 2 (MVS/SP Version 2), which implements System/370 extended architecture [system control program for the MVS/Extended Architecture (MVS/XA) environment]
- Virtual Machine/System Product (VM/SP), used alone or with MVS/SP Version 1, and the VM Migration Aid for System/370 extended architecture
- Virtual Storage Extended/System Package (VSE/SP)
- Airline Control Program/Transaction Processing Facility program product (ACP/TPF)

A collection of required and optional programs are available to support 3380 operations in these environments. The optional programs, some of which have standalone versions, are described later in this section.

Before the 3380 is installed, the correct version and release of the system control program should be in place and all of the associated 3380 required programs obtained. Programming support is available from the IBM Programming Information Department. It is usually distributed on tape.

The following section gives a summary of programming support for the 3380. For complete information on these programs and information about other available programming support, refer to the applicable programming manuals or consult a representative from an IBM branch office.

Documents that contain information on system generation and the various programs that support the 3380 are listed in the bibliography. Program directories are included with the program distribution.

System Support

OS/VS1 Support

The 3380 is supported by OS/VS1 with Data Facility Device Support (DFDS). OS/VS1 implements System/370 architecture, supports all models of the 3380, and supports many of the 3380 functions. However, OS/VS1 does not support dynamic path selection. To support the 3380, the OS/VS1 operating system requires DFDS program product (Release 1.2 or later). OS/VS1 provides system program support. DFDS provides device support. OS/VS1 together with DFDS is called an OS/VS1 environment.

OS/VS1 System Generation for the 3380

DFDS can be installed by updating the distribution library (DLIB) and performing a complete system generation (GENTYPE=ALL). To add 3380s to the system, they must be included during the system generation (SYSGEN) process. As an alternative, DFDS can be installed by performing a SMP APPLY operation. In this case, a system generation (GENTYPE=I/O) must be performed to add 3380s to the system. For more details about system generation, see the *OS/VS1 System Generation Reference*.

OS/VS1 Uses for the 3380

The 3380 can be used for the following storage needs by the OS/VS1 operating system:

- System residence
- Paging
- System and user data sets
- Writing to spool and checkpoint data sets

MVS/SP Version 1 Support

The 3380 is supported by MVS/SP Version 1. MVS/SP Version 1 is based on OS/VS2 MVS Release 3.8 and implements System/370 architecture. This system supports all models of the 3380 and the full functional capabilities of the 3380 including dynamic path selection.

MVS/SP Version 1 requires:

- System Product - JES 2 or JES 3 Release 1, or System Product - JES2 or JES3 Release 3 (MVS/SP-JES2, MVS/SP-JES3)
- DFDS program product or MVS/370 Data Facility Product (DFP) program product

MVS/SP Version 1 provides system program support. DFDS or MVS/370 Data Facility Product (DFP) provides device support. MVS/SP Version 1 together with DFDS or MVS/370 DFP is called an MVS/SP environment.

MVS/SP Version 2 (MVS/XA) Support

The 3380 is supported by MVS/SP Version 2. MVS/SP Version 2 is based on OS/VS2 Release 3.8 and implements System/370 extended architecture. This system supports all models of the 3380 and the full functional capabilities of the 3380, including dynamic path selection.

MVS/SP Version 2 requires:

- System Product - JES2 or JES3
- MVS/XA Data Facility Product (DFP) program product

MVS/SP Version 2 provides system program support. MVS/XA DFP provides device support. MVS/SP Version 2 together with MVS/XA DFP is called an MVS/XA environment.

MVS/SP System Generation for 3380

If the data facility product (MVS/XA DFP, MVS/370 DFP or DFDS) is installed by updating the DLIBs and performing a complete system generation (GENTYPE = ALL), 3380s may be added to the system by including them in the SYSGEN. If the data facility product (MVS/XA DFP, MVS/370 DFP, or DFDS) is installed by performing a SMP APPLY, a system generation (GENTYPE=I/O) must be performed to add 3380s to the system.

Refer to the applicable system generation manual listed in the bibliography.

MVS/SP Uses for 3380

The 3380 can be used for the following storage needs under MVS/SP Version 1, Release 3 and MVS/SP Version 2 (MVS/XA):

- System residence
- Paging
- JES3 setup and scheduling

- System and user data sets
- Spooling and checkpointing

MVS/SP Version 1 Release 1 does not support spooling and checkpointing.

VM/SP Support

The 3380 is supported by Virtual Machine/System Product (VM/SP) program product, which implements System/370 architecture. VM/SP incorporates earlier system extensions. VM/370 Release 6 is a prerequisite for VM/SP.

MVS/SP is supported as a guest (within a virtual machine) operating system.

VM/SP supports all models of the 3380, but it does not support the alternate controller function of dynamic path selection nor use of the commands associated with path group identification for device reservation. These new commands will result in unit check status with command reject in the sense data. This will be interpreted by the guest to mean that the dynamic path selection capability is not available.

VM/SP System Generation for 3380

A starter system for VM/SP is available for system generation. The 3380 is recognized at system generation as a target system residence device, when 3380 is specified on the SYSTYPE operand on the SYSRES macro. The 3380 must be added to the VM/370 real I/O configuration by including an RDEVICE macro in the DMKRIO file.

VM/SP Uses for 3380

The 3380 can be used for the following storage needs by VM/SP:

- Control Program (CP) system residence.
- Control Program (CP) paging and spooling.
- Conversational Monitor System (CMS) system residence device and user data. There are no limitations on minidisk size. All 885 cylinders of the 3380 can be allocated. For CMS minidisks on a 3380, optimum space utilization and performance is obtained with a 4K-byte (4,096 bytes) physical record size. The CMS FORMAT parameter default is 1K-byte (1,024 bytes) physical record size.
- Remote Spooling Communications Subsystem (RSCS) system residence.
- Dedicated device attached to a virtual machine.

Both paging and spooling areas can reside on the same

volume. However, performance characteristics should be carefully considered before designating a 3380 as both a paging and spooling volume. The large capacity of the 3380 permits one 3380 volume to contain VM/SP paging and spooling areas and yet to have space left for other uses. The user should consider whether or not the use of this remaining space introduces significant access mechanism contention on the 3380 before assigning paging and spooling areas to a 3380. Placement of page data sets is described in the section, Effective Use of the 3380.

VSE/SP with VSE/Advanced Functions Support

The 3380 is supported by Virtual Storage Extended/System Package (VSE/SP) with VSE/Advanced Functions, which implements System/370 architecture. Support for the 3380 under VSE applies to all program products that can be ordered and installed with the three categories of VSE System IPO/E Release 1.4.2 products:

- The products comprising VSE/System Package (VSE/SP) Release 1.1.2
- Optional program products, with the exception of VSE/Interactive Problem Control System (VSE/IPCS), which does not support the 3380
- Documented support program products, with the exception of Entry Level Interactive Application System-One (ELIAS-1), which does not support the 3380

VS/FORTRAN Release 1.3.1 supports the 3380 (VS/FORTRAN cannot be ordered through VSE System IPO/E). For a complete list of the program products that support the 3380, see Program Announcement Letter 284-031. See also the *VSE System Support for IBM 3380 and 3370 Models A02/B02* manual.

Current users of VSE/Advanced Functions Release 1.3.5 who are not ordering through VSE System IPO/E can receive programming support for the 3380 through standard distribution procedures. New orders of VSE/Advanced Functions Release 1.3.5 will contain programming support for the 3380.

ACP/TPF Support

The 3380 is supported by the Airline Control Program/Transaction Processing Facility (ACP/TPF). This program supports all models of the 3380. Although dynamic path selection is not supported, the two paths of the 3380 Model AA4 can be used effectively to access all devices in attached strings. Only data streaming

channels can be used, therefore, the commands used with a speed matching buffer are not supported.

ACP is a standalone operating system and requires no optional programs to support the 3380.

ACP/TPF System Generation for 3380

The System Initialization Program (SIP) is used to generate an ACP system. ACP/TPF support of 3380 is established during the SIP process. When 3380 devices are specified, the associated device tables will be included in the generated system. Since disk storage I/O operations in ACP are table-driven, inclusion of the tables establishes 3380 support.

ACP/TPF Uses for 3380

The ACP/TPF can be used for system and user data sets.

Device Support

MVS/SP Version 1 with DFDS

Data Facility Device Support (DFDS) supports I/O functions associated with the access methods. DFDS also provides support for an index to the VTOC. In addition, data management and device support are provided with either the components of the system control program or with separate programs.

Standard Access Method Routines and Data Management are provided with MVS/SP Version 1 with DFDS. The following are MVS/SP Version 1 access methods:

- Sequential access method (SAM)
- Partitioned access method (PAM)
- Direct access method (DAM)
- Indexed sequential access method (ISAM)
- Virtual storage access method (VSAM)

The extended functions of SAM-E use chained scheduling and search direct techniques. SAM-E is available as a separate program product.

Direct Access Device Space Management (DADSM) Routines are provided with MVS/SP Version 1. DADSM routines assign space on a volume and maintain the VTOC. DADSM routines have preprocessing and postprocessing exits. This gives an installation greater control of space management.

Utilities to Perform a Variety of Functions are provided with MVS/SP Version 1. To move data sets, the IEBCOPY, IEHMOVE, IEBGENER, and IEBISAM utilities are used. To obtain lists of information from the VTOC, the IEHLIST program is used. In addition, the Access

Method Services (IDCAMS) program has commands used to move VSAM data sets.

Catalog Support is provided with one or more of the following: MVS/SP Version 1 (CVOL), the IDCAMS program, or Data Facility Extended Function (DFEF) program product. The integrated catalog of DFEF supports both VSAM and non-VSAM data sets.

MVS/SP Version 1 with MVS/370 DFP or MVS/SP Version 2 with MVS/XA DFP

Data Facility Product (DFP) supports I/O functions associated with the access methods. DFP also provides support for an index to the VTOC. (These functions are functionally equivalent to those of DFDS.) In addition, DFP provides data management and device support.

Standard Access Method Routines and Data Management are provided with MVS/370 DFP or MVS/SP Version 2 with MVS/XA DFP. The following are MVS/SP access methods:

- Sequential access method (SAM)
- Partitioned access method (PAM)
- Direct access method (DAM)
- Indexed sequential access method (ISAM)
- Virtual storage access method (VSAM)

The extended functions of SAM-E use chained scheduling and search direct techniques. With DFP, the extended functions of SAM-E are part of SAM.

Direct Access Device Space Management (DADSM) Routines are provided with MVS/370 DFP or MVS/SP Version 2 with MVS/XA DFP. DADSM assigns space on a volume and maintains the VTOC. DADSM routines have preprocessing and postprocessing exits. This gives an installation greater control of space management.

Utilities to Perform a Variety of Functions are provided with MVS/370 DFP or MVS/SP Version 2 with MVS/XA DFP. To move data sets, the IEBCOPY, IEHMOVE, IEBGENER, and IEBISAM utilities are used. Improvements have been made to IEBCOPY and IEHMOVE. To obtain lists of information from the VTOC, the IEHLIST program is used. In addition, the Access Method Services (IDCAMS) program has commands used to move VSAM data sets.

An Integrated Catalog Facility (Part of DFP) supports both VSAM and non-VSAM data sets. It can replace or coexist with CVOL and VSAM catalogs.

OS/VS1 with DFDS

Data Facility Device Support (DFDS) supports I/O functions associated with the access methods. DFDS also provides support for an index to the VTOC. In addition, data management and device support are provided with either the components of the system control program or with separate programs.

Standard Access Method Routines and Data Management are provided with OS/VS1 and DFDS. The following are OS/VS1 operating system access methods:

- Sequential access method (SAM)
- Partitioned access method (PAM)
- Direct access method (DAM)
- Indexed sequential access method (ISAM)
- Virtual storage access method (VSAM)

Direct Access Device Space Management (DADSM) Routines are provided with the OS/VS1 operating system. DADSM routines assign space on a volume and maintain the VTOC. DADSM routines have preprocessing and postprocessing exits. This gives an installation greater control of space management.

Utilities to Perform a Variety of Functions are provided with the OS/VS1 operating system. To move data sets, the IEBCOPY, IEHMOVE, IEBGENER, and IEBISAM utilities can be used. To obtain lists of information from the VTOC, the IEHLIST program can be used. In addition, the Access Method Services (IDCAMS) program has commands used to move VSAM data sets.

Catalog Support is provided with OS/VS1 (CVOL) or with the IDCAMS program.

VSE System IPO/E Release 1.4.2

The 3380 is not supported as a system residence device in a VSE/SP or VSE System IPO/E Release 1.4.2 environment. Libraries and system files can be defined on the 3380 without the help of the Interactive Productivity Facility. The 3380 can also be used for data files, in addition to the minimum DASD requirements for VSE System IPO/E. For details on these requirements, see the *VSE System IPO/E Planning Guide*.

VSE/Advanced Functions Release 1.3.5

The 3380 is supported as SYSRES for system files and for data files under SAM and DAM, but not under ISAM. In addition, VSE/VSAM (including the space management and backup/restore functions) fully supports the 3380.

For more details about using the 3380 in a VSE/Advanced Functions environment, see *VSE System Support for IBM 3380 and 3370 Models A02/B02*.

IPL and Job Control

Device type 3380 can be used in the initial program load (IPL) ADD command:

```
ADD cuu [ (S) ] , 3380 [ , SHR ]
```

Device type 3380 can also be used in the ASSGN job control command or statement:

```
[ / ] ASSGN SYSxxx , 3380 , . . .
```

Additional Separate Programs

Operations under MVS/SP depend on the following two programs:

- Device Support Facilities
- Environmental Recording, Editing, and Printing (EREP) program

In addition to the programs required for 3380 support, the following programs aid in its use:

- IBM Direct Access Storage Migration Aid
- Data Facility Data Set Services
- Hierarchical Storage Manager

Device Support Facilities

The Device Support Facilities (ICKDSF) program is provided for initialization and media maintenance of the 3380. It has both system and standalone versions.

The Device Support Facilities is used to:

- Initialize 3380 volumes for use by the operating system.
- Build an index for a new VTOC or change the format of an existing VTOC.
- Reformat a previously initialized volume.
- Perform hardware and data verification tests to help determine the source of a data problem and, as a consequence, decide on appropriate action.
- Check the surface of tracks on the volume for possible defects.
- Bypass defective areas on a track and use an alternate track, if necessary.
- Reclaim tracks that were previously flagged as defective when surface checking finds them to be usable.

Support for the 3380 is not provided by the IEHDASDR and IBCDASDI utilities nor the Analysis Program-1 (AP-1) because the functions are incorporated in Device Support Facilities.

Use of this program is described in the sections Installation and Initialization, and Device Status and Error Handling.

Environmental Recording, Editing, and Printing

The Environmental Recording, Editing, and Printing (EREP) program supports editing and printing of 3380 error records.

For MVS/SP users, the EREP program prints a set of reports, the System Exception reports, that identify the source of errors. These reports are important for error handling and are described in the section Device Status and Error Handling.

IBM Direct Access Storage Device Migration Aid

The IBM Direct Access Storage Device Migration Aid program can help in converting data sets from other direct access storage devices to the 3380. It operates with MVS/SP and supports most data sets created by the access methods. The migration aid program can be used to:

- Collect information about current disk data sets.
- Calculate data set space requirements for the 3380.
- Identify data sets that may be candidates for reblocking and help to evaluate the appropriateness of block sizes for the new device (not for VSAM).
- Generate, and place in the procedures library, customized JCL and utility statements for migrating the data sets to the 3380. (The migration aid program does not actually move the data sets. This is accomplished with utilities.)
- Identify affected data set JCL statements in the procedures library and create a job stream to update the statements.

Besides its use during migration, this programming aid includes a simple technique that helps determine the most efficient blocking factors for new data sets being placed on the 3380.

Use of this program is described in the section Data Set Conversion.

Data Facility Data Set Services

The Data Facility Data Set Services program product provides dump and restore functions for the 3380. Data

Set Services has both system and standalone versions and supports IBM disk storage products except the 2314 and 2319. The standalone version supports restore of full and partial (range of tracks) volumes. The system version supports both dump and restore of full and partial (range of tracks) volumes. The system version also supports dump and restore by data set name.

With Data Set Services:

- Data sets can be selected for dump or restore based on attributes such as creation date, last referenced date, and expiration date.
- Data sets can be dumped from multiple disk source volumes in a single operation; the output will be written to one sequential file.
- Once data has been dumped to a tape, it can be copied to one or more additional tapes to provide multiple copies.
- Certain data sets can be relocated on a volume to reduce fragmentation.
- Data may optionally be compressed to save space on the dump tape. The compressed data can be restored from the tape to disk with the system version of Data Set Services but not with the standalone version.

Besides its functions for dump and restore, the Data Set Services program can be used to copy data from one disk storage volume to another of the same device type. It also can be used to print a data set, range of tracks, or all or part of the VTOC.

Dump and restore by data set is not possible for VSAM data sets, nor is the print function. The defragmentation function cannot be used for data sets that are ISAM, unmovable, or addressed by absolute track.

Data Set Services uses a command language that follows the Access Method Services (IDCAMS) syntax.

Support for 3380 is not provided by IEHDASDR, IBCDMPRS, DRWDASDR, and DRUDMPRS utilities, because the functions are incorporated in Data Set Services and other programs.

Use of this program is described in the section Data Backup and Recovery.

Hierarchical Storage Manager (MVS/SP Only)

The Hierarchical Storage Manager (HSM) is a space management tool that helps to take advantage of the data handling and storage attributes of different types of storage, such as disk, tape, and mass storage cartridges. The program monitors data set use and automatically keeps active data sets on the faster access devices and migrates less active data sets to other devices, as specified by the user.

The HSM program also can be used to dump and restore data sets for data backup and recovery, and as a data conversion tool.

Use of this program is described in the sections Data Set Conversion and Data Backup and Recovery.

Data Set Conversion

Data set conversion involves changing programs to take into account differences in device characteristics, and moving data from the existing device to the newly installed 3380. This section points to the major programming tools available to perform these functions.

Data Set Conversion Under MVS/SP and OS/VS1

In planning for data conversion, the following factors need to be considered in terms of the effect they have on how and when the data is moved.

- The type of data set: system, application, or data base
- The type of processing: batch, online, or interactive
- The access method used: SAM, PAM, DAM, ISAM, or VSAM
- Any special placement of data sets, such as data sets currently on multiple volumes and those that are not movable

In addition, attention is needed for any special requirements for moving catalog data sets and for updating catalogs to reflect data sets that have been moved. A new volume table of contents is created when a volume is initialized and is automatically updated by space management services to reflect data set placement on the device.

New backup copies of data sets can be made with Data Facility Data Set Services (DFDSS). (Backup tapes that may have been produced by the DRWDASDR and IEHDASDR utilities cannot be restored by Data Facility Data Set Services.)

Approaches

There are two principal approaches to data set conversion. One approach is to use the Direct Access Storage Migration Aid program with Job Control Language (JCL) and utility specifications. Another approach (under the MVS/SP operating system only) is to use the Hierarchical Storage Manager program. Because each approach has certain advantages, it is sometimes beneficial to use both, specifically by using the migration aid program to augment the Hierarchical Storage Manager.

The Direct Access Migration Aid program produces reports on data set characteristics. The reports give information that can be used to determine the block sizes that are most effective in utilizing space and the amount of space allocation needed for given non-VSAM data sets. The reports identify the JCL changes needed.

With specifications supplied by the user, the migration aid can then be used to generate JCL statements regarding such items as block sizes and space allocation. The program also generates statements for utilities or Access Method Services programs to move the data. All access methods are supported. Full volumes, partial volumes, and specific data sets can be specified for movement and multiple volumes can be moved simultaneously.

The Hierarchical Storage Manager program (under the MVS/SP operating system only) can be used to move data sets to the 3380. Data set movement can be performed either as a function of storage management or as an independent function. In both cases, the data is, in effect, moved by way of an intermediate volume. The SAM, PAM, and VSAM access methods are supported. Certain SAM data sets can be reblocked automatically when moved to the new device. Space on the new device is allocated by the program.

Programming Changes

Most of the programming required to store data sets on the 3380 can be done with JCL statements; or where VSAM is used, with the Access Method Services (IDCAMS) program.

The JCL data definition (DD) statements specify the I/O requirements for a job step. Device-related information for both system and user data sets can be defined with JCL data definition statements instead of specified with system macros (DATASET and DCB).

Three DD statement parameters are likely to be affected by 3380 installation and data set conversion. These are UNIT; SPACE; and DCB, subparameter BLKSIZE.

The UNIT parameter contains the information needed for device allocation for a data set. (See also Addressing in the section Hardware Planning Considerations and System Generation in the section Program Planning Considerations.)

UNIT parameters specifying a unit address or device number probably have to be changed to establish a correspondence between the unit number and a specific 3380. UNIT parameters specifying a generic name (device type) have to be changed to the 3380. UNIT parameters specifying a group (esoteric) name have to be changed unless all devices in the group are changed to 3380, or the 3380 addresses are added to the group at system generation.

The JCL that requests mounting or demounting (volume count greater than unit count, deferred mounting) must be changed for the 3380.

The SPACE parameter designates the amount of space to be allocated for a data set. Space can be allocated in tracks, cylinders, or blocks, and in tracks at an absolute address. If space is allocated by number of tracks or cylinders or with an absolute address, the SPACE parameter has to be changed for the 3380 characteristics. As a general guideline for planning, approximately one-half as many tracks will be needed on a 3380 as on a 3350 for fixed-length records. The same number of cylinders will be needed since the 3380 has 15 tracks per cylinder compared with the 3350 which has 30 tracks per cylinder. If space is allocated by blocks, the SPACE parameter need not be changed unless the block sizes are changed.

The BLKSIZE subparameter of the DCB parameter specifies the maximum length of a block in bytes. Block sizes and therefore the BLKSIZE parameter should be changed to take full advantage of the increased track and cylinder capacity of the 3380. However, unless current block sizes are less than 1098 bytes, changing block sizes may be postponed to reduce the conversion effort. Blocking considerations are described in the next section, Effective Use of the 3380.

Data Set Movement

Data sets are moved from one device to another with utilities, or for VSAM data sets with Access Method Services.

Data sets can be copied to the new device and the original retained on the source volume, or they can be moved to the receiving volume and deleted from the source volume. All of the utilities provide a copy function. The IEHMOVE utility provides both a copy and a move option.

SAM data sets can be converted to the 3380 using the IEHMOVE or IEBGENER utility.

PAM data sets can be converted to the 3380 using the IEHMOVE or IEBCOPY utility. IEHMOVE reorders the members alphabetically while IEBCOPY does not change the order. IEBCOPY is generally faster.

ISAM data sets can be converted to the 3380 using the IEBISAM utility.

DAM data sets can be converted to the 3380 using the IEHMOVE utility. If more records are to be placed on a 3380 track than are on a source device track, the existing

reorganization program can be used to transfer data to the 3380. It may be necessary to reprogram the randomizing algorithm if a relative track address is used to calculate an actual address and if fewer or more tracks are allocated to the data set than before.

VSAM data sets can be converted using the Access Method Services REPRO command or the EXPORT/IMPORT commands.

In all of the above cases, the catalog must be changed for cataloged data sets.

Data bases under Information Management System (IMS) can be converted with the IEHMOVE utility invoked by the Direct Access Storage Migration Aid. However, to ensure backup, the IMS utilities IMAGECOPY/RECOVERY may be considered.

Use of Existing Programs

Subject to the restrictions on device dependencies and those indicated for DAM data sets, existing 3330, 3340, or 3350 programs can be used without change to process data sets transferred to the 3380.

User-Written Access

User-written programs that issue the EXCP macro directly require modifications to reflect the different 3380 device characteristics, if the programs have been coded with device dependencies. If the programs use relative track addresses, they may not require any modifications. The 3330, 3340, and 3350 channel command word (CCW) lists operate on the 3380 except those that are device or channel time dependent and those that support the file scan feature or the Write Special Count, Key, Data command. (The 3380 does not have track overflow.) Channel commands to obtain rotational position sensing (read and set sector commands) must be added by the user if this support is desired.

If an index is used for the VTOC, the format 5 data set control blocks are invalid and dependent user programs will have to be modified. Refer to the *Data Facility Device Support* or *Data Facility Product* manuals.

Data Set Conversion Under VM/SP

Current CMS commands (COPYFILE, MOVEFILE, TAPE, TAPPDS) are available to move data to 3380 devices from other disk storage types. System areas on CP OWNED volumes must be reconstructed on 3380.

Data Set Conversion Under VSE/Advanced Functions

Data sets can be converted from other DASD volumes to the 3380 by using VSE/Advanced Functions utility programs:

- To transfer system libraries and private libraries to the 3380, use the backup and restore utilities, or use

the CORGZ program of VSE/Advanced Functions.

- To transfer files to the 3380, use VSE/DITTO Release 1.1.0 (VSE/Data Interfile Transfer, Testing, and Operations utility program, a program product).
- To copy files to the 3380, use VSE/Fast Copy Data Set Release 1.2.0.

Effective Use of the 3380

This section discusses some of the factors involved in deciding which data sets to place on the 3380 and how to place the records on the device to achieve the best use of the space and performance. It also describes specific means for accomplishing the tasks involved. These factors can make important differences in the benefits realized from the increased 3380 storage capacity and in performance at both the device and system levels.

The points made in this section should be viewed as guidelines, rather than as explicit recommendations. The information is intended as a general aid to achieving effective use of the 3380.

Criteria for Using the 3380

Assuming the requirement for fast, direct-access storage, one of the criteria for deciding which data sets to place on the 3380 will be performance needs. The 3380 has the potential for higher performance in terms of throughput than most other storage devices; therefore, data sets with critical time requirements generally should have priority for placement on the 3380.

Performance needs will probably be the primary criterion for placing data sets on the early 3380 installations. As more 3380s are added, the exceptional large capacity may make it desirable to move, or add, data sets not justified previously for placement on a fast, direct-access device. Data sets such as these also may be candidates for placement on the 3380 where a balance is needed for those data sets requiring high performance.

Effective use of storage, of course, must be viewed in terms of effective use of the various storage media in the total storage system. Under these circumstances, the distribution of data sets to particular devices may change depending, for instance, on frequency of use. The Hierarchical Storage Manager, described in the previous section, provides a programming method to help make this distribution of data sets a more dynamic process.

Space and Performance Guidelines

For effective use of storage, the objective is to achieve a balance between space utilization and performance. Choices made regarding block sizes and data placement determine the amount of track space utilized for data, the number and distance of arm movements required to

access data, and the length of time required for data transfer. When making block size and data placement choices at the device level, their dependencies and potential effects on other components in the system need to be taken into account.

Optimum Blocking for 3380 Records

The amount of space on a track or cylinder actually utilized for data depends to a large extent on the block sizes selected. Grouping records into blocks conserves space because it reduces the amount of space required for address information and for separator gaps. Blocking also may increase processing efficiency by reducing the number of I/O operations required to process the data.

Background Information

A block is the unit of data that is accessed and transferred. It has a distinct location and a unique physical address on a device. A block may include one or more *logical records*, the unit of data used by the processing program. Logical records can be blocked and unblocked by the access method routines or the user's program at processor speeds. The term *blocking factor* refers to the number of logical records grouped in a block.

Block sizes correspond to the size of *physical records*, specifically to the data length (DL) of the physical record. If keys are used, the key length (KL) must be added to the data length. Therefore, the number of physical records that can be placed on a track depends on block sizes and whether or not there are keys. Block sizes and key lengths are specified by the user with the BLKSIZE and KEYLEN parameters of the OS/VSE DCB macro, or the JCL DD statement DCB parameter.

A certain amount of space on every track is required for overhead. The amount required depends on the number of physical records per track.

With the 3380, data is recorded in 32-byte increments. This is a consideration when deciding on block sizes because small block sizes are vulnerable to an accumulation of residual bytes. For example, a 981-byte block occupies as much space as a 1,012-byte block.

Capacity Calculations

Appendix A provides formulas for calculating the number of physical records (blocks) of equal and of unequal lengths that can be placed on a track. For equal length records, Appendix A shows the maximum block sizes for the given number of records, calculated with and without keys. In the following discussion, equal length records without keys are assumed, unless stated otherwise.

Selection of Block Sizes

The selection of block sizes depends on a number of factors including:

- Device space characteristics
- Size of logical records to be blocked
- Organization and accessing techniques used
- Amount of main storage for I/O buffer areas
- Potential effects on system resources

Device Space Characteristics: When data sets are being created or converted from one device to another, space characteristics of the particular device type must be considered. (For a comparison of the space characteristics of different devices, see Figure 15.) Block sizes that yield satisfactory use of space for existing device types, including the 3350, may not prove as satisfactory for the 3380.

The example in Figure 29 illustrates differences in space requirements between the 3350 and 3380 that are attributable to the difference in space characteristics of the two device types. For the comparison, assume a data set of 400,000 logical records of 80 bytes each. The smaller block size actually requires more cylinders on the 3380 than on the 3350. The larger block size uses fewer tracks or cylinders on both the 3350 and the 3380, but the savings in space is even more significant with the 3380. When converting data sets where space has been allocated in track or cylinder units, it is necessary to

recalculate the number of units required based on the differences in the two devices.

To achieve maximum use of the device space, it would be necessary to assign a block size of the full capacity of a track. With the 3380, one record with a data length of 47,476 bytes can be placed on a track. **The largest block size supported by IBM-supplied access methods is 32,760 bytes.** The block size that provides optimum space utilization on a 3380 for most types of data sets is 4,096 bytes (4K). The EXCP macro must be used to obtain block sizes larger than 32,760 bytes.

The next best use of space is obtained with two records of 23,476 bytes, which achieve approximately 99% utilization. A block size larger than one-half of a track, however, prevents placement of another equal length record on the track, and introduces the possibility that the residual space will be lost.

As block sizes decrease, the amount of data that can be stored on a track also decreases. This is because smaller block sizes increase the amount of space required for overhead (there are more address areas and gaps). For example, if physical records with block sizes of 116 bytes were chosen, about 19% of the capacity would be available for user data. Block sizes that correspond to the 80-byte IBM punched card utilize only 14% of the full data capacity of a 3380.

A span of block sizes allows a given number of physical records to be accommodated, and the largest block sizes in the range achieve the best use of space. For example, any block size between 4,276 and 3,861 bytes accommodates 10 physical records, without keys, on a 3380 track. A block size at the highest end of this range utilizes approximately 90%, while block sizes at the lowest end of the range utilize only about 82% of the available capacity. (The table in Appendix A gives the block size that achieves the best use of space for a given number of records.)

Blocking Factor	Block Size	Blocks Required	Blocks per Track		Tracks Required		Blocks per Cylinder		Cylinders Required	
			3350	3380	3350	3380	3350	3380	3350	3380
10	800	40,000	19	36	2105	1112	570	540	71	75
100	8,000	4,000	2	5	2000	800	60	75	67	54

Figure 29. Differences in Device Space Requirements between the 3350 and 3380

In general, block sizes between 3,000 and 8,000 bytes are most practical for the 3380. These sizes are likely to provide a better balance of space utilization and performance than blocks of smaller or larger sizes. Larger block sizes may give better space utilization, but may require too much main storage buffer space and may tie up other resources with the long data transfers.

A block size of 6,233 bytes has a distinct advantage. This size provides the potential for over 90% space utilization on the 3330, 3350, and 3380 and, thereby, offers device independence for many data sets. While this exact block size may rarely suit the needs of the application, it can be used as an optimum figure to arrive at a suitable block size.

For previous disk storage devices, it is recommended that system data sets be blocked in the maximum, full-track size — that is, one record per track. Because the maximum block size for the 3380 is so large (47,476 bytes), system data sets probably should be blocked with a size of one third the track length. The maximum utilization for three equal length records per track can be attained with a block size of 15,476 bytes.

Logical Record Size: The sizes of logical records to be blocked are another consideration of block sizes. A block size must be a multiple of the length of fixed-length logical records. One procedure for arriving at a suitable block size is to start by selecting a likely candidate as an ideal block size and then adjusting it to accommodate the logical records in the data set. For example, if the data set contains 80-byte logical records, an ideal block size of 6,233 might be chosen. The closest to this ideal block size can be attained with 77 logical records blocked in 6,160-byte sizes. This permits 7 blocks to be placed on a track and achieves 90% utilization.

Organization and Accessing Techniques: In general, data sets that are sequentially accessed can be blocked in larger sizes than data sets directly accessed. However, this general guideline takes on a new perspective with the 3-MB/sec data transfer rate of the 3380. Larger data transfers can be justified, even for direct data sets, than were reasonable at the slower rate of earlier devices.

Records that are organized and accessed in physical sequential order can make the best use of space because large block sizes can be used. However, there are two constraints to consider when establishing block sizes for these records. The first is that the size is limited by the amount of main storage available for I/O buffer areas. The second constraint applies when there is a mixture of data sets, where some are sequential and some require

direct access for interactive applications. Excessively long data transfers, such as occur with large block sizes, may monopolize the path, thereby interfering with response times for interactive users.

With an index sequential organization, the access method controls the index entries, but the user has control of blocking for the prime data records, usually over 70% of the space involved. When a data set can be processed either sequentially or randomly by keys, as with ISAM, both processing methods must be considered when establishing block sizes. For sequential processing, there are some space and performance advantages obtained with large block sizes, as already described. However, for random processing by keys, large block sizes can cause a degradation in data transfer time. Therefore, the ideal block size for ISAM records is the one that provides an optimum balance between sequential and random processing, limited by the amount of main storage available for I/O buffers.

Direct data organization may present some special considerations depending on the accessing techniques used. In the most commonly used technique, records are stored randomly with *indirect addressing*. Other techniques usually tend to be a variation of ISAM, that is, indexed by keys, and therefore involve the same considerations. For randomly stored records, better space utilization might be attained by randomizing to a relative block address rather than to a relative track address. ISAM support for the 3380 is not included in the VSE/Advanced Functions environment.

For VSAM data sets, the unit of data specified by the user is a control interval (CI), which is the unit of data transfer. CI sizes are multiples of 512 bytes until 8,192 and then multiples of 2,048 bytes. A control interval is made up of block sizes of 512, 1K, 2K, or 4K bytes. A commonly used block size of 4,096 bytes provides approximately 85% utilization of a 3380 track. Smaller sizes yield less effective space utilization. Users should specify a CI size that can be made up of the most effective block sizes.

Blocking Aid: For MVS/SP operating systems, the Direct Access Storage Device Migration Aid program can help to determine which existing data sets need to be reblocked and the block sizes to be used. (This does not apply to VSAM.) This program provides a means for determining which block sizes meet the user's criterion for percentage of utilization.

The IBM Hierarchical Storage Manager program optionally reblocks data sets when they are recalled during conversion to 3380.

Performance Considerations

Performance of disk storage operations is affected by factors such as the number of units in a string and placement of data sets. In addition, there are new characteristics of 3380 operations that affect performance.

- When the 3380 string is controlled by a Model AA4 with two controllers, placement of the data sets relative to the internal paths should be taken into consideration.
- When a 3880 storage director has a speed matching buffer and a channel with a 1.5 MB/sec data transfer rate, performance benefits are obtained by using the new Define Extent and Locate Record channel commands for write operations, and by using the programs that support these commands.

Both internal paths and speed matching buffer are discussed in the section Operations.

- Besides these specific considerations, there are two other factors that are reflected in 3380 performance. More data is traversed with each disk rotation or arm movement, and more data is transferred in a given length of time.

Techniques usually effective in obtaining performance benefits are as follows. In many cases, the techniques result in more than one type of performance benefit. (Parameters of the MVS DCB macro are in parentheses.)

1. Reduce the number of accesses required for a given workload.
 - Make frequently accessed data resident in main storage by placing it there at system or job step initialization time. Frequently accessed data typically includes program modules, indexes, directory lists, and access method routines.
 - Optimize the block size of sequential data sets to obtain large data transfers with a single access, but be careful not to interfere with interactive users who may be trying to access data sets over the same path.
 - If not using SAM-E or SAM of DFP: access multiple blocks for each access by specifying chained scheduling for sequential data sets (OPTCD=C). With MVS/SP operating systems, chained scheduling is used wherever possible whether or not it is requested. SAM-E is not supported by the OS/VS1 operating system.

2. Minimize arm movement when accessing.

- Place the most active permanent data sets near the middle cylinder (cylinder 443). Placement at a middle cylinder has potential performance advantages for data sets that are frequently accessed. This is because the access arm begins a seek operation from the position where it completed the last operation; therefore, the distance the arm must travel will, on the average, be less to reach a cylinder in the middle of the disk surface. If absolute addressing is not used, however, DADSM allocates new data sets from the lowest numbered track of the first space large enough to accommodate the request. On a sparsely filled volume, this could result in having permanent active data sets in the middle and new data sets close to the beginning of the volume.

- Place the smallest data sets of about equal activity nearer the center.
- Place concurrently active data sets adjacent to one another when they must be accessed by the same access mechanism.
- Place small concurrently active data sets on the same cylinder if possible (for example, the VTOC index, VTOC, and master catalog).
- Cluster the most active members of partitioned data sets near the directory within the data set.
- Allocate a realistic amount of space for data sets to eliminate time required to pass over excessive unused space.
- Use the Device Support Facilities program to skip up to seven defects on a track before assigning an alternate track, which requires a change of arm position.

3. Minimize the time required for a given access and optimize transfer speed.

- Place concurrently active data sets under different access mechanisms when possible to avoid arm contention.
- Place concurrently active data sets on different internal paths if dynamic path selection is in use.
- Distribute data sets to obtain a balance between those requiring frequent access and those that are seldom accessed.
- Use the techniques described in item 2 to reduce arm movement.

- Use a smaller block size for data sets where direct access predominates.
- Use a 3.0 MB/sec data transfer rate channel, where possible.
- Use the Define Extent and Locate Record commands for adjusted data rate when using a 1.5 MB/sec data transfer rate channel.
- If not using SAM-E or SAM of DFP: use the search direct option (OPTCD=Z) or fixed length record, standard format (RECFM=FS or FBS) for physical sequential data sets where possible. Search direct cannot be used when chained scheduling is used, and chained scheduling takes priority.

4. Reduce channel busy time. The most significant techniques for reducing channel busy time are already done automatically, such as disconnecting during seek operations and using rotational position sensing (RPS) where possible. The following techniques can also be implemented.

- Use an indexed VTOC when there are many data sets on a volume to save VTOC search time.
- Use a smaller block size for data sets where direct accesses predominate.
- Use small directories for partitioned data sets to avoid long searches.

5. Reduce processor interference for I/O operations.

- If using data streaming and standard commands, allocate space in tracks or cylinder units depending on the size of the data sets and on performance requirements. When space is allocated in tracks, the control program checks for record extents at the end of each track. To eliminate the time required for this extent checking, allocate space in cylinders if the data set is large enough to use a practical amount of the cylinder space. For small data sets, it is better to allocate in tracks.

If using the 3880 speed matching buffer and commands for adjusted data rate, the above does not apply. The Define Extent command specifies the beginning and ending physical addresses within which operations can be performed.

- If using the 3880 speed matching buffer and PCI fetching of load modules, use larger block sizes to avoid an increase in the number of Start I/O instructions.

- Use the techniques described in item 1 to reduce the number of access operations required.

Application Considerations

Many of these known performance factors involve trade-offs that can only be evaluated for a particular application.

Small block sizes decrease the time per access and permit more concurrent operations on a channel, but they reduce the effective data rate, result in more processor interference for sequential data sets, and use more device space for a given quantity of data.

Large block sizes allow a high effective data rate for sequentially accessed data sets and reduce processor overhead for I/O instruction processing, but they monopolize the path for longer periods of time.

Therefore, block sizes that result in satisfactory performance in a batch processing environment may not be desirable in an online, response-oriented system, where it may be preferable to sacrifice device space for a gain in performance. The large storage capacity available with the 3380 helps to make this a reasonable trade-off. In addition, the high data transfer rate allows block sizes to be increased while maintaining acceptable response times.

Page Data Sets, Special Considerations

Page data sets are critical for system operation. The following recommendations should improve system performance through the careful use of page data sets and devices.

- Avoid placing more than one page data set on any single 3380 volume to reduce the contention among multiple page data sets for the use of the access mechanism.
- Specify a data set to exclusively contain the pageable link pack area (PLPA). Place the PLPA modules on the fastest available device, because needed real-storage pages must be obtained from the PLPA data set. Avoid placing other page data sets on the same device.
- Place most of the page data sets on moderate speed devices to avoid overloading the fast device if the number of fast devices (such as the 3380) is limited.

- Set up a relatively slow device for data sets that will contain secondary copies of duplexed common areas. Secondary copies are read only some of the time for error recovery, and therefore should not occupy space on faster devices.
- Specify overspace (excess) for all page data sets to allow the creation of additional address spaces before current ones are deleted, and to permit some reason-

able increase in the number of concurrent virtual I/O (VIO) data sets.

- Protect page data sets with a password. Use the **DEFINE SPACE** command to specify password protection when preallocating the data set, or use the **ALTER** command at a later time to add password protection.

Device Status and Error Handling

Status conditions provide information to the system regarding the progress of an I/O operation at the 3380. As with previous disk storage, device status is supplied to the channel by the storage control. If an error is detected during execution of a channel command, the 3880 storage director includes a unit check status in a combination of status conditions sent to the channel. When an error occurs, the channel initiates procedures for the subsystem to recover from the error or, if this is not possible, causes an I/O interruption. If an I/O interruption occurs, the system can perform various steps to try to recover from the problem.

The system obtains further explanation of the cause of a unit check status by issuing a Sense channel command to obtain the 24 bytes of sense information developed by the 3880 storage director. Sense bytes identify the conditions that caused the unit check status, indicate the type of error, and specify the access mechanism and seek addresses where the error was detected. The sense bytes further provide information for system error recovery procedures and for diagnosing and isolating the cause of an error condition.

A description of sense information for the 3380 is contained in the *3880 Storage Control Description* manuals.

The 3880 Storage Control manuals also include a description of the *error condition* and *recovery action tables* for system error recovery. The error condition table gives the meanings of the various combinations of sense bit settings that can occur and pointers to the appropriate action in the recovery action table. Each unique configuration of sense bit settings has a recommended recovery action. These actions are implemented by the system error recovery procedures (ERP). Error recovery procedures are provided with the IBM operating systems.

Many errors are corrected by the subsystem or the system with error recovery procedures. In the event that these procedures are not possible or successful, a system message is issued to bring the problem to the attention of the operator. The system messages contain the address and type of error and also include status and sense information.

When an error situation cannot be recovered by the subsystem or the system, there are program functions available to treat the cause of certain types of problems. For other problems, the services of a customer engineer may be required to treat the cause of the error.

Program Functions

Various program functions are available to interpret error information and assist in problem determination and resolution.

The system makes status and sense information available to the problem program through error analysis routines. When an uncorrectable I/O error occurs, error analysis routines analyze the status and sense information and produce error messages for the user.

The system error recovery procedures include logging of certain error conditions. All 24 bytes of sense information are placed in the system log record. Data reduction and report generation based on information in the system log record are accomplished with the Environmental Recording, Editing, and Printing (EREP) program. EREP support for the 3380 is a complete replacement of previous EREP support. This new release includes a report on failure types to help the customer engineer in diagnosing a problem.

In addition to the EREP reports used by the customer engineer, the System Exception reports help to determine when and where error recovery action should be applied. The reports describe the type and location of an error and identify the probable source of the error. With this information, operations personnel can determine whether to call for a customer engineer or attempt recovery with program functions.

The Device Support Facilities program has error handling functions for use by operations personnel. This program exercises hardware to determine if the equipment is functioning properly, checks disk surfaces to determine if there are defects on the media, bypasses defects and uses alternate locations, and rewrites home addresses and record zeros.

Error Handling Guidance

Guidance on error handling for system programmers and operations personnel can be found in the *IBM Disk Storage Management Guide*.

Error Handling for Different Types of Errors

Error handling will be described for four general categories: programming errors, control errors, data errors, and equipment errors.

Programming Errors

Programming errors, such as invalid track format and no record found, cause a command reject sense bit to be set. The information of the I/O error is passed to the user program for recovery action. Error handling for programming type errors is described in programming manuals. An error message is issued that the job has been terminated.

Control Errors

Recovery from seek errors and timing checks, including command and data overruns, often can be accomplished by the 3880 storage director with the command retry function. The command is retried without interrupting the system. If successful, the check is regarded as a temporary error. If unsuccessful, a permanent I/O error message is issued. Seek and overrun problems that are not correctable with retries are equipment problems and require a customer engineer.

Data Errors

A data error is an error in the bit pattern read from the disk. These data errors can be caused by such things as a slight misalignment of the head with the track, a transient electronic noise situation, or a defect on the disk surface.

Error detection and correction information is added by the 3380 to each count, key, and data area of a record when it is written. Later, when the record is read, the information is used to detect data errors that may be present and to correct the data where possible.

The error information is a pattern of bits, coded to define a validity check for the full data in the recorded area. (This is similar to the common use of a hash check digit with an account number for number validity checking.)

The error correction code (ECC) generated by the 3380 is used to detect single burst errors contained within five contiguous bytes and to correct single burst errors contained within three contiguous bytes. Also the code can be used to correct single burst errors contained within four bytes if the errors are in two halfwords (byte pairs) adjacent to each other. That is, the two halfwords must have a common boundary relative to the beginning of the area (count, key, or data area). See Figure 30.

Errors that can be corrected with the error-checking and correction code are referred to as ECC correctable

errors. Sense information identifies whether the data check is ECC correctable or ECC uncorrectable.

The method used to correct the data when it is ECC correctable depends on which area of the record contains the error.

- If the error is in the home address, count, or key area, the 3880 storage director makes the correction. It uses the ECC correction information to correct the data to be transferred to main storage. The data correction and resumption of the command chain are done without interrupting the system.
- If the error is in the data area, the system error recovery procedures make the correction. The system error recovery procedures include an *error correction function* to correct the data in main storage. It uses pattern and displacement information derived from the ECC. This information is supplied to the system in sense information.

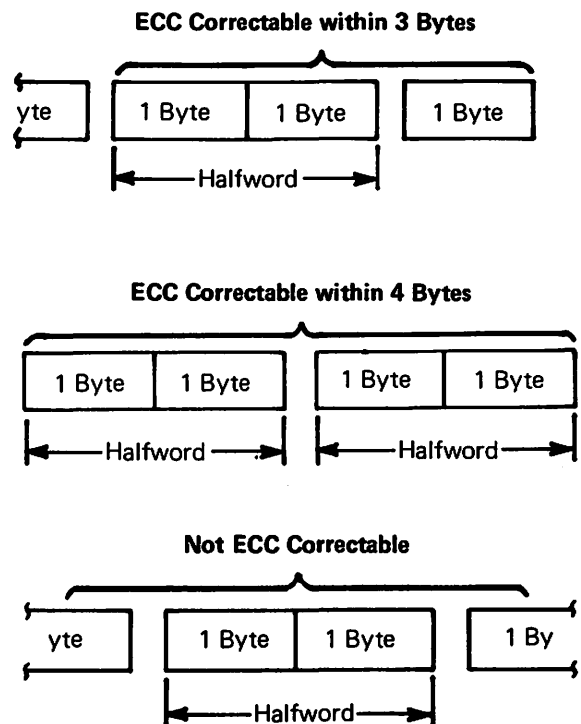


Figure 30. Data Error Correction

In many cases, even though the first attempt to read the data is not successful, with subsequent attempts it can be read correctly. If a data error is detected and the data is not ECC correctable, the 3880 uses command retry in repeated attempts to read it correctly. If this is not successful, offset adjustments are made with the access mechanism and attempts to read the data repeated. These attempts are made in any area of the record in which an error is detected. Within the specified number of retry attempts, the data may be read correctly or become ECC correctable. The command retries are made without interrupting the system.

(For read multiple count, key, and data operations, the above internal correction procedures are used only on the first record read. If an error occurs in subsequent records, unit check status is presented so that system error recovery can be initiated.)

Temporary and Permanent Data Errors

If a data error is correctable with the ECC or readable with command retry, it is regarded as a temporary error. If a data error is not within the scope of ECC correction and the specified number of command retries prove unsuccessful, the error is regarded as a permanent data check.

Permanent data checks and excessive temporary data checks cause an I/O error message to be issued at the operator console. Such errors will be described in the EREP System Exception reports.

Response to Data Checks

If a data check is caused by a defect on the disk surface, the Device Support Facilities program can be used to bypass the area.

Defects are bypassed by skipping the defective area to use a displacement location on the same track. Up to seven defective areas on a track are skipped. Space for the defect displacement is already reserved so it does not subtract from that available for space calculations. If the defect displacements are exhausted, the track is flagged in the home address area, and an alternate track is assigned. (Record 0 carries the alternate and primary address associations.) Cylinder address 885 is reserved for assignment of up to 15 alternate tracks.

The Device Support Facilities program can also be used to rewrite home addresses. Rewriting home addresses can eliminate the cause of some types of data checks.

Backup Data

Data on an area or track that is to be bypassed may have to be restored from a backup copy. As described in the section Backup and Recovery, the Data Facility Data Set Services has the program functions to do this. The same program has a print function for printing the data from tracks where an error occurred.

Data Verification By Program

There is a programming method for immediately verifying that data is correctly written at the device. The verification helps to eliminate the chances for future data errors during processing, but requires time and, therefore, may not be considered practical for routine use.

A read operation with the CCW skip flag bit on causes the data to be transferred back as far as the storage director, and the storage director checks the data using the error correction code (ECC). With MVS/SP operating systems, this validity check (or write check) can be done by specifying the option OPTCD=W in the data control block (DCB) macro or the DCB parameter of the JCL DD statement. This procedure is the same as that used for the 3350 disk storage.

Equipment Errors

An equipment, or hardware, error can occur when there is a problem in the controller, the drive, or the head and disk assembly.

System error recovery procedures attempt recovery from an equipment error by retrying the operation. With the 3380, the system retries the operation 10 times.

Temporary and Permanent Equipment Checks

If the system retry procedures are successful, the error is regarded as temporary. If the retry procedures are not successful, the error is regarded as a permanent equipment check.

Permanent equipment checks cause an I/O error message to be issued at the operator console. The error will be described in the EREP System Exception reports.

Response to Equipment Checks

Equipment checks caused by a hardware problem require the service of a customer engineer to restore use.

Service Procedures

The customer engineer performs diagnostic and problem determination procedures to identify the source of the failure and repairs or replaces the failing components, which are referred to as field replaceable units (FRUs).

The customer engineer can service portions of the logic and controls of one access mechanism without interrupting user operations on the remaining access mechanisms on the 3380 string.

Advantages of Two Controllers

The 3380 model AA4s provide an alternate access to all access mechanisms in the string. If access is lost because of a controller, storage director, or channel malfunction, the data is still available through the other controller and its associated storage director and channel. With the 3380, the two controllers in the Model AA4 back up each other. This occurs without manual intervention. The dynamic path selection function allows the backup to be implemented.

Statistical Information and Logging

The 3880 storage director maintains counters for its attached 3380s. Statistical counters count the number of seeks, bytes read and searched, and command and data overruns where command retry was used. Accumulated values are entered in the system error log record and made available in the EREP report.

Temporary seek errors and data checks are counted by the 3880 storage director error rate counters. (Error rate for seek errors is the number of seek checks per seeks attempted. Error rate for data checks is the number of data checks per megabytes read or searched.) If the error rate threshold of a counter is exceeded, the threshold condition is entered in the system error log record. Subsequent errors continue to be logged until there is sufficient information to assist in corrective action. Information about the nature and location of the threshold condition is issued in a message to the operator and documented in the EREP report.

Temporary equipment checks are entered in the system error log record, and information about the nature and location of the error is documented in the EREP report.

Permanent errors are entered in the system error log record. Information about the nature and location of the error is issued in a message to the operator and documented in the EREP report.

Physical Identifiers

Identifiers in 3380 sense information pinpoint the location of a problem more precisely than in previous disk storage products. This is a response to the difficulty of trying to identify a failing device by its I/O (unit) address. In complex installations, a device may be accessed by one or more processors, by multiple channels, and, with installations using dynamic path selection, by different storage directors. As a result, a single device can be addressed by the same or many different unit addresses.

To overcome this identification problem, the customer and customer engineer now assign a unique physical identifier (ID) to each storage director and controller. These identifiers, unlike the unit address, are unique within the installation.

The 3880 storage directors, 3380 controllers, and 3380 devices are each assigned a 1-byte (two hexadecimal characters) physical ID.

The storage director identifiers must be assigned in pairs in each 3880 Storage Control. The numbers for the two physical IDs must be consecutive, with the even number having the smaller value. For example: 02 and 03, 30 and 31 are valid pairs because the first physical ID of each pair is even numbered. But 05 and 06, 23 and 24 are not valid pairs because they begin with odd numbers.

In 3380s with only one controller (Model A04), only one physical ID is needed and the number may be even or odd. In models with two controllers (Model AA4), two physical IDs are assigned, one for each controller. As with the storage directors, the numbers for the two physical IDs must be consecutive, with the even numbers having the smaller value. For example: controller A1 = 02 and controller A2 = 03 or controller A1 = 80 and controller A2 = 81. The even number is always assigned to controller A1. (The customer engineer will know which controller is A1 and which is A2.)

A string with a controller address of 0 must be assigned physical IDs between 02 and 7F. A string with a controller address of 1 must be assigned physical IDs between 80 and FD. See Figure 31 for a summary of controller physical ID assignment. For information on controller addresses, refer to the earlier section Addressing, Bit 3.

For storage directors and controllers, numbering should not begin with 00 and 01 and should not include FE and FF.

The device physical ID is the same as the device address. The first of the two hexadecimal digits of the device physical ID is always zero.

The physical identifiers are set with switches at installation and are in addition to I/O (unit) addresses that also are set by the customer engineer at installation. Spaces are provided on the control panel on Models A04 and AA4 to affix physical ID labels.

Physical IDs are included in the 24 bytes of sense information sent in response to a Sense channel command. They also are included in system messages to the operator and in EREP reports.

Besides its use to identify a failing storage director-controller path, the physical ID may be helpful in configuration management. For instance, it may be used to communicate the identification of a particular 3380 string that is to be disabled with a switch on the operator panel.

Controller A1 Address Bit 3=0	Controller A2 Address Bit 3=0
02	03
04	05
.	.
7C	7D
7E	7F

Controller A1 Address Bit 3=1	Controller A2 Address Bit 3=1
80	81
82	83
.	.
.	.
FA	FB
FC	FD

Figure 31. Possible Physical IDs for Model AA4 Controllers

Data Backup and Recovery

It is increasingly important to develop measures to protect against the loss of data as larger amounts are stored on a single device and multiple users become more dependent on its availability. Therefore there are new demands in terms of:

- Recognizing when a failure occurs and determining its cause
- Maintaining copies of data and keeping track of changes so a file can be restored from a backup if necessary
- Establishing an alternate for disabled storage
- Preparing for recovery of operations that are interrupted

To ensure that data is not permanently lost in the event of a failure, backup copies can be created, usually on a different storage medium such as tape. The data that is dumped to the backup device includes all of the associated address information, so it can be restored to the same device type from which it was originally dumped. Data sets vary considerably in terms of how frequently and extensively they are changed. Therefore, the amount of data and the frequency with which dumps are made vary with the application.

Routines for creating backup copies and strategies for recovery depend both on the frequency of change to the data and on the potential effects were the data temporarily unavailable or not up to date.

In OS/VS1 and MVS/SP operating systems, the Data Facility Data Set Services (DFDSS) program, described in the section Program Planning Considerations, provides flexibility for dump and restore operations to meet a variety of needs. (There are differences in functions available with the system and standalone versions.) With the system version:

- Data can be dumped from a 3380 to a backup device (including tape, the IBM 3850 Mass Storage System, or another disk storage) and later restored to the source 3380 or a different 3380.
- A full volume, one or more data sets, or selected tracks can be dumped. With SAM and PAM data sets, only the used portion of the allocated spaces are dumped.
- Data sets can be restored to multiple volumes in a single operation.

- Specific tracks or data sets can be restored to the 3380, even though the dump was taken for the full volume or multiple data sets.
- Data sets may be renamed during restore operations.
- Once they have been dumped, data sets can be deleted.

All VSAM data sets can be dumped and restored as full or partial volumes with the DFDSS program. Selected data sets can be dumped and restored using the EXPORT and IMPORT commands of Access Method Services.

In MVS/SP operating systems, the Hierarchical Storage Manager program can be used to control automatic backup of designated volumes. The user can specify the frequency of creating dumps, the number of backup versions to be made, and when the dumps are to be made. This function makes backup copies of only those data sets that have been changed since the last backup. There are some constraints on the types of data sets that can be backed up with this program; for instance, it cannot be used for ISAM data sets.

Recovery from a failure involves restoring the data from a backup copy and reconstructing it to the current level of processing. Files are dependable only to the last checkpoint or dump. They can be updated by applying a journal of intervening transactions that were made against the data set or by rerunning the jobs that accessed the file since the last checkpoint or dump.

In a batch environment, data can be restored from a backup, journals applied if they have been kept, and the job being processed when the interruption occurred can be restarted. The procedures are essentially the same in an online environment except there are many users of a single volume of data sets and, therefore, no journal made of transactions. Data can be restored from the last checkpoint, and the individual users notified to resubmit their jobs.

Figure 32 describes the performance that can be expected when using the Data Facility Data Set Services to dump 3380 volumes to IBM 3420 and 3480 Magnetic Tape units.

	3420 Model 8	3420 Model 6	3480
Data Recording Density	6250 bpi	6250 bpi	38,000 bpi
Data Transfer Rate	1.25 MB/sec	0.78 MB/sec	3.0 MB/sec
Tape Reels per 3380 Volume ¹	2	2	— ³
Approximate time elapsed ²	5 min	6.6 min	— ³

1. The number of reels required depends on the number of megabytes of data stored on a disk volume and the data set type mix. A volume of approximately 50% capacity is assumed.
2. The time to dump depends on the number of tracks. The figures are based on an assumption that 60% of the tracks in the volume are used (out of approximately 13,275 tracks per access mechanism). The time results also are based on the assumption that two tape drives are used, so that no rewind time is included.
3. No data available.

Figure 32. Disk-to-Tape Dumping

Conversion Checklist

Conversion from existing types of storage to the 3380 involves tasks that range from physical installation of the equipment to conversion of the data sets. Typically, this effort entails a series of stages from initial planning to full production. Checkout and testing are integral elements of the process. The objective is to make the conversion, while causing as little impact as possible on normal data processing operations. Performance goals are important to an overall successful conversion. However, with the many factors involved in conversion, extensive tuning of the system may not be practical at the beginning. Efforts in the initial conversion should be to apply basic performance guidelines and concentrate on those aspects of conversion that impact operations while the conversion is actually taking place.

The following checklist contains items that should be considered when converting to the 3380. The list does not contain all of the tasks associated with the effort, but it calls attention to significant areas.

Identify IBM support sources and establish communication.

- National Accounts Division (NAD) or National Marketing Division (NMD) account representative
- System engineering support in the IBM branch office
- Customer engineer hardware specialist
- Program support representative

Order the appropriate operating system release and the required and desired optional programs from the IBM Program Information Department.

- Order the most current level of Device Support Facilities program so it will be available prior to planned installation of 3380s. This program is used for 3380 initialization and error handling.
- Order the EREP program with the System Exception reports needed for error handling.

Order the needed publications through an IBM branch office.

- Arrange with your IBM branch office for System Library Subscription Service (SLSS), where applicable, for publications needed for 3380 use. Under the SLSS, revisions to existing publications and any new publications in a bill of form (BOF) series will be sent automatically.

- In addition to publications associated with the operating system and the required Data Facility Device Support or Data Facility Product, the following publications are useful:

- *Device Support Facilities User's Guide and Reference*
- *Environmental Recording, Editing, and Printing (EREP) Program User's Guide and Reference*
- *IBM Disk Storage Management Guide*

Plan for physical installation.

- Consider the physical characteristics of the 3380 in terms of heat dissipation, electrical requirements, cable lengths, weight, and size of the units.
- Consider the physical characteristics during the period when both the 3380 and the units it replaces are on the system.

Arrange for data center use. Testing that must take place prior to installation may have to be scheduled at an appropriate IBM data center.

- Review system configurations at the data center to ensure that backup and test requirements can be met.
- Determine IBM data center hardware compatibility, and communicate special requirements, if any.
- Schedule data center time.

Arrange for personnel training in the following areas.

- System programming
- Application programming
- Operations

Establish a detailed conversion plan and schedule that includes:

- Program and data set conversion plans
- Operational plans to minimize impact
- Application of performance guidelines
- Methods for controlling and monitoring progress
- Interim (parallel) operating procedures
- Checkout and acceptance criteria for release of old facilities

Note: Checking user programs and data sets on the 3380 may require a full cycle of data processing (depending on the users procedures). Procedures for verification of tests and parallel runs should be established with extra attention paid to online or communications systems.

Review the status of installed program products for 3380 support.

- Reorder if necessary.
- Define alternatives if the 3380 is not supported.

Review application programs.

- EXCP programs using 3330s or 3350s may have to be modified (including CCWs, sense checking).
- BDAM randomizing routines may require modification or complete reprogramming.
- Accounting routines may need to be modified if based on the number of tracks or cylinders used. (A billing by block allocation algorithm would not need to be changed.)

Plan for system generation.

Arrange for main storage requirements.

- A slight increase in resident storage (100 bytes in MVS) is required for 3380 support.
- Additional main storage is required if block sizes are increased to take advantage of the 3380 higher capacity tracks.

Review addressing requirements.

- Check that the addresses conform to the requirements described in the sections Hardware Planning Considerations and Program Planning Considerations.

Review system parameter and build lists.

- PARMLIB and BLDL lists should be reviewed and revised, because many modules may have been split or recommended for addition to residency list in the new release.

Initialize 3380 volumes.

- Use the Device Support Facilities program at installation to initialize volumes, as described in the section Installation and Initialization.

Examine performance considerations.

- Consider the performance guidelines as described in the section Effective Use of the 3380.
- Consider the increment of the capacity. The 3380 provides more storage per access mechanism than 3330, 3340, or 3350 disk storage.
- Consider how more data sets on the same volume can affect performance.
- Check DCB parameters. Are DCB parameters (BLKSIZE, RECFM) specified in user programs (Assembler, COBOL, PL/I)? If so, the specification may need to be changed.
- Consider the logical record format (RECFM). All fixed blocked QSAM and BSAM data sets should be examined to see if the fixed length standard (FS=FBS) can be used.

Set priorities for areas to address. A sequence of areas to address, in order of priority, might be the following:

- Attempt to optimize placement of system data sets, including page data sets.
- Distribute other data sets such as sort work areas, and compiler and linkage editor work data sets. Careful definition of the UNITNAME parameter during system generation will aid greatly here.
- Review procedures libraries (PROCLIB) to ensure optimum space allocation, block sizes, and unit separation via UNITNAME. Because procedures are used by many jobs, the performance return for careful design is great.
- Review all access method usage to take advantage of performance options not now being used.
- Evaluate block sizes of user data sets to take advantage of 3380 track size.
- Attempt to optimize user data set placement across volumes.
- Attempt to optimize user data set placement within volumes.

Prepare for data set conversion.

- Determine data set conversion priorities.
- Schedule machine time for one-time conversion runs.

- Check carefully to make certain all the data sets are converted prior to the running of the program for programs that use concatenated data sets.
- Consider the guidelines in the section Effective Use of the 3380, including use of the migration aid program and utilities.
- Determine if data conversion can be done as part of the regular processing cycle. On the first day, use data from the existing devices (such as 3330 or 3350) and after processing, store it on the new 3380 devices. On the second day, use data from the 3380 and after processing, store the updated data on the 3380.

Prepare backup and recovery procedures.

- Develop two strategies – one for conversion and parallel operation and one for full production. Conversion and parallel operation backup is more complicated than full production backup, because old and new devices are different. This may require keeping the old catalog, procedures, JCL, and duplicate data sets. Close control over volume serial number usage and catalog updating is required during parallel runs, to avoid confusion and errors due to duplicate volume IDs.

- Consider the accessibility and availability of backup systems or devices, such as distance, availability, and priorities.
- Arrange means for dumping and restoring data, including possible use of Data Facility Data Set Services.

Update operating procedures.

- Modify machine room procedures to reflect installation of the 3380 and new programs.
- Plan to familiarize the operator with new machine room procedures.
- Take into account that there is no Ready light on the operator panel. Procedures should be implemented to ascertain device status from the system console, as described in the section Operator Controls and Indicators.

Appendix A: Space Calculations

Space calculations are determined for physical records. The data length (DL) of a physical record is the block size. With MVS/SP operating systems, this is specified in the BLKSIZE parameter of the DCB macro. Information and guidelines for determining optimum block sizes are in the section, Effective Use of the 3380.

Formulas

All data on a 3380 is written in 32-byte increments. Therefore, the formulas for making space calculations differ from those used for the 3330 and 3350 disk storage devices.

Equal Length Physical Records

The number of equal length physical records per track can be calculated as follows. Track overhead for the home address, track descriptor record zero (R0), and skip defect allowance has already been accounted for.

$$\text{Equal length records per track} = \frac{1499}{C + K + D}$$

1499 is the number of 32-byte increments per track available for user data records.

C is the number of 32-byte increments used by the record overhead including gaps and the count area.

K is the number of 32-byte increments used by the key area.

D is the number of 32-byte increments used by the data area.

KL is the key length in bytes.

DL is the data length in bytes.

where C = 15 if KL = 0

C = 22 if KL ≠ 0

K = 0 if KL = 0


$$K^1 = \frac{KL + 12}{32} \text{ if } KL \neq 0$$

$$D^1 = \frac{DL + 12}{32}$$

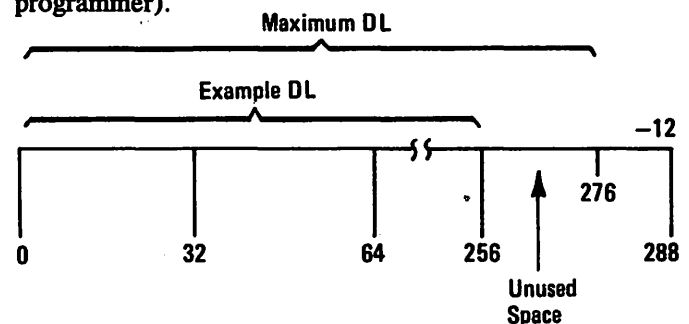
¹ These values are rounded up to the next integer. KL and DL each must have 12 added and each must be rounded up to a multiple of 32. The 12 bytes are for error checking. For example, with a partitioned data set directory where KL = 8 and DL = 256: 12 bytes must be added to the 8, and 12 bytes must be added to the 256, and each rounded up to the necessary 32-byte increments.

The following is a more extensive explanation of the use of the space calculation formula. The same partitioned data set directory is used as an example. The explanation shows the maximum data length and key length actually available, as well as the space required. The available space is given in the space calculation tables that follow.

In the directory example, the KL is 8. To the 8, add 12, yielding 20, which is then rounded up to 32 x 1 = 32 bytes of required space. Of the 32 bytes required, subtract 12, yielding 20 bytes available for the key, although only 8 are needed (as specified by the programmer).



In the directory example, the DL is 256. To the 256, add 12, yielding 268, which is then rounded up to 32 x 9 = 288 bytes of required space. Of the 288 bytes required, subtract 12, yielding 276 bytes available for data, although only 256 are needed (as specified by the programmer).



Therefore, KL + DL = 20 + 276 = 296 bytes record length available.

The following calculations show the number of 296-byte records (with keys) that can be placed on a track:

$$\frac{1499}{C + K + D} = \frac{1499}{22 + 1 + 9} = 46 \text{ records of } 296 \text{ bytes}$$

If less than the available space is specified by the programmer, the storage control pads to the full 32 bytes, that is, to the full available space.

Unequal Length Physical Records

In the case of unequal length records, the number of records (n) per track must satisfy the following relationship, if a standard track descriptor record zero (R0) has been written on the track.

$$\text{Unequal length records per track} = \sum_{i=1}^n (C_i + K_i + D_i) \leq 1499$$

When the track descriptor record (R0) is not standard, or has not been written on the track, the following relationship must be satisfied:

$$\text{Unequal length records per track} = \sum_{i=1}^n (C_i + K_i + D_i) \leq 1515$$

C, K, and D are the same as defined for equal length records for the i^{th} record.

Capacity calculations for the 3380 can be performed by using the TRKCALC macro. By using this macro, the user can do the following:

- Determine the number of records of a given size that can be written on a full track or the remainder of a track.
- Perform track balance calculations.
 - Determine if a given record size can be written in the space remaining on the track and return the new track balance.
 - Determine the maximum size record that can be written on the track if the given record does not fit.
 - Determine the track balance if the last physical record is removed from the track.

TRKCALC is available in OS/VS2 Release 3.8 with PTFs UZ90059 and UZ28937. TRKCALC is integral in MVS/SP Version 2 (MVS/XA). Use of the TRKCALC macro is described in the *System Programming Library: Data Management* manuals.

A generalized program that uses the TRKCALC macro can be assembled. This program can be called from a TSO CLIST and the various data parameters passed to the program from the terminal. In this fashion the TRKCALC macro can be used interactively. It can be a useful tool for planning and conversion.

Space Count Maintenance

The following formula can be used to maintain count of the amount of available space remaining on a 3380 track, assuming a standard track descriptor record zero (R0) has been written on the track.

$$\text{Number of 32-byte increments} = 1499 - \sum_{i=1}^n (C_i + K_i + D_i)$$

where C_i , K_i , and D_i are as previously defined.

This is equivalent to:

$$\text{Available space} = 47,968 - \sum_{i=1}^n (C_i + K_i + D_i)$$

where $C = 480$ bytes if $KL = 0$
 $C = 704$ bytes if $KL \neq 0$
 $K = 0$ if $KL = 0$
 $K = (\text{Key Length} + 12)$ rounded up to a multiple of 32 if $KL \neq 0$
 $D = (\text{Data Length} + 12)$ rounded up to a multiple of 32

An Example of Copying a Partitioned Data Set

One of the principal reasons for performing space calculations is to determine the amount of 3380 space to allocate when copying a data set from another device to a 3380 volume. It is fairly easy to determine the amount of space the original data occupied, but determining the amount of 3380 space needed to hold the same amount of data is less obvious.

For example, assume a partitioned data set (PDS) stored in fixed-length blocks (FB) on a 3350 has the following characteristics:

- Logical record length = 80 bytes (LRECL=80).
- Physical record length = 3,120 bytes (BLKSIZE = 3120). The logical records are blocked 39 to 1.

For this example, 7,200 blocks have been allocated for the directory and 5,424 blocks have been used. For the data portion of the PDS members, 50 cylinders of 3350 space have been allocated with 41 cylinders used. If all space is occupied by 3120-byte blocks, it would seem that the PDS contains 6,150 blocks. However, this assumption is not correct for the following reasons:

- The PDS contains 37,966 members, each of which requires an end-of-file mark (a 0-byte data record). The amount of 3350 space needed for the end-of-file marks is 12 cylinders and 6 tracks.

- The remaining space, 28 cylinders and 24 tracks, can be occupied by 3120-byte blocks. Assuming an optimum distribution of data blocks and end-of-file marks so no 3350 space is wasted, the PDS can contain 4,900 blocks.

To determine the correct amount of 3380 space to be allocated when the PDS is copied, use the space calculation tables in this appendix and perform the following steps:

1. Determine how much space is needed for the PDS directory. Figure A2 (Part 1 of 4) indicates that 7,200 directory blocks, each of which has an 8-byte key area and a 256-byte data area, require 11 cylinders.

46 Directory blocks per track
15 Tracks per cylinder
690 Directory blocks per cylinder

Therefore, 7,200 directory blocks fit in 11 cylinders (although the actual requirement is 10 cylinders, 7 tracks).

Note: Even though only 5,424 directory blocks have been used, space must be set aside for the number of directory blocks that have been allocated.

2. Determine how much space is needed for the PDS data blocks. Figure A1 indicates that 4,900 blocks at 3120 bytes per block require 26 cylinders.

13 3120-Byte data records per track
195 3120-Byte data records per cylinder

Therefore, 4,900 3120-byte data blocks fit in 26 cylinders (although the actual requirement is 25 cylinders, 2 tracks).

3. Determine how much space is needed for the PDS end-of-file marks. The PDS contains 37,966 members, each of which requires an end-of-file mark (a 0-byte data record). Figure A1 indicates that 37,966 end-of-file records require 28 cylinders.

93 0-Byte data records per track
1,395 0-Byte data records per cylinder

Therefore, 37,966 0-byte data records fit in 28 cylinders (although the actual requirement is 27 cylinders, 4 tracks).

4. Add the space requirements for each part of the PDS to determine the amount of space to allocate.

11 Cylinders for the directory
26 Cylinders for the PDS members
28 Cylinders for the end-of-file marks
+ _____
65 Cylinders for the PDS

The minimum amount of 3380 space to allocate for the PDS is 65 cylinders. However, this assumes that all end-of-file marks (0-byte data records) are clumped together in a 28-cylinder area. A more random (and not necessarily optimum) distribution would be expected, resulting in a higher space requirement.

It should also be noted that if the data set copying utility (IEBCOPY) needs to write a record, but there is not enough space left on the track to contain the record, the record is written on the next track. The space that is not used on the previous track is lost and not available to store anything associated with the PDS.

So, for this example, an initial primary allocation of 75 cylinders and a secondary allocation of 10 cylinders would be enough to contain the copy of the PDS.

Tables, without and with Keys

The following tables give the number of equal length records that can be placed on a 3380 track and cylinder, based on the length of the data area (DL) and key area (KL), if any. The number of records that can be placed on a track depends on the record data area length and whether or not the record includes a key area. The actual length of a record can be determined by using the formula for equal length records, discussed earlier.

Because partial records cannot be written on a track, the number of equal length records that can fit on a track is the same over a range of data lengths. For example, Figure A1 indicates that if a record data area length is between 3,861 bytes and 4,276 bytes, 10 records can fit on a track and 150 records can fit on a cylinder.

The table also helps determine the largest data area length for each equal length record if a track is to contain a given number of records. For example, Figure A1 indicates that 20 equal length records can fit on a track, as long as each record data area length is between 1,781 and 1,876 bytes.

The table in Figure A1 gives the number of equal length records that can be written on a track when the records do not include key areas. The tables in Figure A2 (parts 1 through 4) give the number of equal length records that can be written on a track when the records include a key area.

The data capacity values in the tables specify the number of bytes available for user data, assuming that the maximum number of equal length records have been written on the track or cylinder and that the record data area is the largest possible for that number of data records. For example, if ten 4,276-byte records are written on a track, the amount of user data contained on the track is 42,760 bytes (as shown in Figure A1). However, if fewer records are written or if shorter records are written on the track, the amount of user data actually on the track is less.

Equal Length Physical Records without Keys

To determine the number of equal length data records of a specific size that can be written on a track, see Figure A1. Using the left-most column, select the smallest DL

value that is equal to, or greater than, the desired data length. The number of records that can fit on a track and a cylinder can be read from the other columns.

For example, a data record of 256 bytes is greater than 244 and less than the next higher value, 276. Using the data in the row for records with a maximum length of 276, observe that, at the most, 62 equal length records can be written per track, and 930 records can be written per cylinder.

The other columns list the maximum track and cylinder data capacity. If a track is filled with 276-byte data records (that is, 62 records have been written), the maximum capacity of the track is 17,112 bytes of user data. The maximum capacity of the cylinder, 256,680 bytes of user data, is approximately 251K bytes of user data (1K byte = 1024 bytes).

Smaller records in the 245 to 276 range would yield a track user data capacity that is less than the value listed in the table. However, the track or cylinder would be able to contain the same number of equal length physical records.

DL Bytes (Max)	Track Capacity		Cylinder Capacity	
	Records	Bytes	Records	Bytes
47,476	1	47,476	15	712,140
23,476	2	46,952	30	704,280
15,476	3	46,428	45	696,420
11,476	4	45,904	60	688,560
9,076	5	45,380	75	680,700
7,476	6	44,856	90	672,840
6,356	7	44,492	105	667,380
5,492	8	43,936	120	659,040
4,820	9	43,380	135	650,700
4,276	10	42,760	150	641,400
3,860	11	42,460	165	636,900
3,476	12	41,712	180	625,680
3,188	13	41,444	195	621,660
2,932	14	41,048	210	615,720
2,676	15	40,140	225	602,100
2,484	16	39,744	240	596,160
2,324	17	39,508	255	592,620
2,164	18	38,952	270	584,280
2,004	19	38,076	285	571,140
1,876	20	37,520	300	562,800
1,780	21	37,380	315	560,700
1,684	22	37,048	330	555,720
1,588	23	36,524	345	547,860
1,492	24	35,808	360	537,120
1,396	25	34,900	375	523,500
1,332	26	34,632	390	519,480
1,268	27	34,236	405	513,540
1,204	28	33,712	420	505,680
1,140	29	33,060	435	495,900
1,076	30	32,280	450	484,200
1,044	31	32,364	465	485,460

DL Bytes (Max)	Track Capacity		Cylinder Capacity	
	Records	Bytes	Records	Bytes
980	32	31,360	480	470,400
948	33	31,284	495	469,260
916	34	31,144	510	467,160
852	35	29,820	525	447,300
820	36	29,520	540	442,800
788	37	29,156	555	437,340
756	38	28,728	570	430,920
724	39	28,236	585	423,540
692	40	27,680	600	415,200
660	41	27,060	615	405,900
628	42	26,376	630	395,640
596	44	26,224	660	393,360
564	45	25,380	675	380,700
532	46	24,472	690	367,080
500	48	24,000	720	360,000
468	49	22,932	735	343,980
436	51	22,236	765	333,540
404	53	21,412	795	321,180
372	55	20,460	825	306,900
340	57	19,380	855	290,700
308	59	18,172	885	272,580
276	62	17,112	930	256,680
244	65	15,860	975	237,900
212	68	14,416	1,020	216,240
180	71	12,780	1,065	191,700
148	74	10,952	1,110	164,280
116	78	9,048	1,170	135,720
84	83	6,972	1,245	104,580
52	88	4,576	1,320	68,640
20	93	1,860	1,395	27,900

Figure A1. Records and Bytes per Track and per Cylinder

Equal Length Physical Records with Ranges of Key Lengths

To determine the number of equal length keyed data records of a specific size that can be written on a track, see Figure A2 (parts 1 to 4) and perform the following steps:

1. Select the table appropriate for the range of key length values that includes the desired key length.
2. Select the smallest DL value that is equal to, or greater than, the desired data length in the left-most column of the selected table.
3. Read from the next column the number of records that can fit on a track.
4. Multiply the Records per Track value by 15 to determine the number of records that can fit on a cylinder.
5. Multiply the Records per Track value by the DL Bytes (Max) value to determine the maximum data capacity for a track.
6. Multiply the Records per Cylinder value by the DL Bytes (Max) value to determine the maximum data capacity for a cylinder.

For example, a keyed data record with a key length of 8 bytes [see Figure A2 (Part 2 of 4)] has a data length of 256 bytes. From the table, the DL Bytes (Max) value to select is 276. The number of keyed data records that can fit on a track is 46. Therefore, the number of records that can fit on a cylinder is 690 (46 x 15).

To determine the track and cylinder data capacity for this example, assume a track is filled with 276-byte keyed data records (that is, 46 records have been written) and determine the values as follows.

- Track capacity : $46 \times 276 = 12,696$ bytes
- Cylinder capacity : $690 \times 276 = 190,440$ bytes, or approximately 186K bytes of user data (1K byte = 1024 bytes)

Smaller records in the 245 to 276 range would yield a track user data capacity that is less than the maximum value listed in the table. However, the track or cylinder would be able to contain the same number of equal length physical records.

KL = 1 to 20 DL Bytes (Max)	Records per Track	KL = 21 to 52 DL Bytes (Max)	Records per Track	KL = 53 to 84 DL Bytes (Max)	Records per Track
47220	1	47188	1	47156	1
23220	2	23188	2	23156	2
15220	3	15188	3	15156	3
11220	4	11188	4	11156	4
8820	5	8788	5	8756	5
7220	6	7188	6	7156	6
6100	7	6068	7	6036	7
5236	8	5204	8	5172	8
4564	9	4532	9	4500	9
4020	10	3988	10	3956	10
3604	11	3572	11	3540	11
3220	12	3188	12	3156	12
2932	13	2900	13	2868	13
2676	14	2644	14	2612	14
2420	15	2388	15	2356	15
2228	16	2196	16	2164	16
2068	17	2036	17	2004	17
1908	18	1876	18	1844	18
1748	19	1716	19	1684	19
1620	20	1588	20	1556	20
1524	21	1492	21	1460	21
1428	22	1396	22	1364	22
1332	23	1300	23	1268	23
1236	24	1204	24	1172	24
1140	25	1108	25	1076	25
1076	26	1044	26	1012	26
1012	27	980	27	948	27
948	28	916	28	884	28
884	29	852	29	820	29
820	30	788	30	756	30
788	31	756	31	724	31
724	32	692	32	660	32
692	33	660	33	628	33
660	34	628	34	596	34
596	35	564	35	532	35
564	36	532	36	500	36
532	37	500	37	468	37
500	38	468	38	436	38
468	39	436	39	404	39
436	40	404	40	372	40
404	41	372	41	340	41
372	42	340	42	308	42
340	44	308	44	276	44

Figure A2 (Part 1 of 4). Records and Bytes per Track for Ranges of Key Lengths

The left-most column gives a value for a data length for a range of key lengths. To determine the number of records per cylinder, multiply the Records per Track value by 15.

KL = 1 to 20 DL Bytes (Max)	Records per Track	KL = 21 to 52 DL Bytes (Max)	Records per Track	KL = 53 to 84 DL Bytes (Max)	Records per Track
308	45	276	45	244	45
276	46	244	46	212	46
244	48	212	48	180	48
212	49	180	49	148	49
180	51	148	51	116	51
148	53	116	53	84	53
116	55	84	55	52	55
84	57	52	57	20	57
52	59	20	59		
20	62				

Figure A2 (Part 2 of 4). Records and Bytes per Track for Ranges of Key Lengths

The left-most column gives a value for a data length for a range of key lengths. To determine the number of records per cylinder, multiply the Records per Track value by 15.

KL = 85 to 116 DL Bytes (Max)	Records per Track	KL = 117 to 148 DL Bytes (Max)	Records per Track	KL = 149 to 180 DL Bytes (Max)	Records per Track
47124	1	47092	1	47060	1
23124	2	23092	2	23060	2
15124	3	15092	3	15060	3
11124	4	11092	4	11060	4
8724	5	8692	5	8660	5
7124	6	7092	6	7060	6
6004	7	5972	7	5940	7
5140	8	5108	8	5076	8
4468	9	4436	9	4404	9
3924	10	3892	10	3860	10
3508	11	3476	11	3444	11
3124	12	3092	12	3060	12
2836	13	2804	13	2772	13
2580	14	2548	14	2516	14
2324	15	2292	15	2260	15
2132	16	2100	16	2068	16
1972	17	1940	17	1908	17
1812	18	1780	18	1748	18
1652	19	1620	19	1588	19
1524	20	1492	20	1460	20
1428	21	1396	21	1364	21
1332	22	1300	22	1268	22
1236	23	1204	23	1172	23
1140	24	1108	24	1076	24
1044	25	1012	25	980	25
980	26	948	26	916	26
916	27	884	27	852	27
852	28	820	28	788	28
788	29	756	29	724	29
724	30	692	30	660	30
692	31	660	31	628	31
628	32	596	32	564	32
596	33	564	33	532	33
564	34	532	34	500	34
500	35	468	35	436	35
468	36	436	36	404	36
436	37	404	37	372	37
404	38	372	38	340	38
372	39	340	39	308	39
340	40	308	40	276	40
308	41	276	41	244	41
276	42	244	42	212	42
244	44	212	44	180	44
212	45	180	45	148	45
180	46	148	46	116	46
148	48	116	48	84	48
116	49	84	49	52	49
84	51	52	51	20	51
52	53	20	53		
20	55				

Figure A2 (Part 3 of 4). Records and Bytes per Track for Ranges of Key Lengths

The left-most column gives a value for a data length for a range of key lengths. To determine the number of records per cylinder, multiply the Records per Track value by 15.

KL = 181 to 212 DL Bytes (Max)	Records per Track	KL = 213 to 244 DL Bytes (Max)	Records per Track	KL = 245 to 255 DL Bytes (Max)	Records per Track
47028	1	46996	1	46964	1
23028	2	22996	2	22964	2
15028	3	14996	3	14964	3
11028	4	10996	4	10964	4
8628	5	8596	5	8564	5
7028	6	6996	6	6964	6
5908	7	5876	7	5844	7
5044	8	5012	8	4980	8
4372	9	4340	9	4308	9
3828	10	3796	10	3764	10
3412	11	3380	11	3348	11
3028	12	2996	12	2964	12
2740	13	2708	13	2676	13
2484	14	2452	14	2420	14
2228	15	2196	15	2164	15
2036	16	2004	16	1972	16
1876	17	1844	17	1812	17
1716	18	1684	18	1652	18
1556	19	1524	19	1492	19
1428	20	1396	20	1364	20
1332	21	1300	21	1268	21
1236	22	1204	22	1172	22
1140	23	1108	23	1076	23
1044	24	1012	24	980	24
948	25	916	25	884	25
884	26	852	26	820	26
820	27	788	27	756	27
756	28	724	28	692	28
692	29	660	29	628	29
628	30	596	30	564	30
596	31	564	31	532	31
532	32	500	32	468	32
500	33	468	33	436	33
468	34	436	34	404	34
404	35	372	35	340	35
372	36	340	36	308	36
340	37	308	37	276	37
308	38	276	38	244	38
276	39	244	39	212	39
244	40	212	40	180	40
212	41	180	41	148	41
180	42	148	42	116	42
148	44	116	44	84	44
116	45	84	45	52	45
84	46	52	46	20	46
52	48	20	48		
20	49				

Figure A2 (Part 4 of 4). Records and Bytes per Track for Ranges of Key Lengths

The left-most column gives a value for a data length for a range of key lengths. To determine the number of records per cylinder, multiply the Records per Track value by 15.

Appendix B: Addressing Reference Information

This appendix supplements the explanation of addressing that is in the section Hardware Planning Considerations. The 2-byte device address contains the address for the 3880 storage director and 3380 controller in bits 0 through 3, and for the device (access mechanism) in bits 4 through 7.

Storage Director				Con- troller	Device or Access Mechanism				Hexa- decimal	Access Mechanism Address (full string)				
Bits	0	1	2	3	Bits	4	5	6			7			
	0	0	0	0	=	0		0	0	0	0	=	0	1
	0	0	0	1	=	1		0	0	0	1	=	1	2
	0	0	1	0	=	2		0	0	1	0	=	2	3
	0	0	1	1	=	3		0	0	1	1	=	3	4
	0	1	0	0	=	4		0	1	0	0	=	4	5
	0	1	0	1	=	5		0	1	0	1	=	5	6
	0	1	1	0	=	6		0	1	1	0	=	6	7
	0	1	1	1	=	7		0	1	1	1	=	7	8
	1	0	0	0	=	8		1	0	0	0	=	8	9
	1	0	0	1	=	9		1	0	0	1	=	9	10
	1	0	1	0	=	A		1	0	1	0	=	A	11
	1	0	1	1	=	B		1	0	1	1	=	B	12
	1	1	0	0	=	C		1	1	0	0	=	C	13
	1	1	0	1	=	D		1	1	0	1	=	D	14
	1	1	1	0	=	E		1	1	1	0	=	E	15
	1	1	1	1	=	F		1	1	1	1	=	F	16
0000-1111						0-F		0000-1111				0-F		

Examples of 3880 and 3380 addressing:

0001 1000 = 18
 0110 0101 = 65
 1000 1110 = 8E
 1111 0111 = F7

Abbreviations

ACP	Airlines Control Program	KEYLEN	key length (DCB parameter)
AF	Advanced Functions	KL	key length
BDAM	basic direct access method	LRECL	logical record length (DCB parameter)
BISAM	basic indexed sequential access method	MB	megabytes
BLKSIZE	block size (DCB parameter)	MB/sec	megabytes per second
BOF	bill of forms	ms	milliseconds
BPAM	basic partitioned access method	MSS	mass storage system
BSAM	basic sequential access method	MVS	multiple virtual storage
CCW	channel command word	NAD	IBM National Accounts Division
CMS	conversational monitor system	NMD	IBM National Marketing Division
CP	command processor	OPTCD	optional services code (DCB parameter)
CSW	channel status word	OS/VS	operating system/virtual storage
DAM	direct access method	PAM	partitioned access method
DCB	data control block	PDS	partitioned data set
DD	data definition	PLPA	pageable length pack area
DFP	data facility product	PROCLIB	procedures library
DFDS	Data Facility Device Support program	PS	physical sequential
DFDSS	Data Facility Data Set Services program	QISAM	queued indexed sequential access method
DISP	data set disposition (DCB parameter)	QSAM	queued sequential access method
DL	data length	R	record
DLIB	distribution library	RECFM	record format (DCB parameter)
ECC	error correction code	RPS	rotational position sensing
EREP	Environmental Recording, Editing, and Printing program	SAM	sequential access method
ERP	error recovery procedures	SIO	start input/output
EXCP	execute channel program	SLSS	System Library Subscription Service
GB	gigabytes	SP	system product
HA	home address	SYSIN	system input stream
HDA	head and disk assembly	SYSOUT	system output stream
HSM	Hierarchical Storage Manager program	TIC	transfer-in-channel
ICKDSF	Device Support Facilities program	TSO	time-sharing option
ID	identification	VM	virtual machine
I/O	input/output	VSAM	virtual storage access method
IPL	initial program load	VSE	virtual storage extended
ISAM	indexed sequential access method	VTOC	volume table of contents
JCL	job control language	XA	extended architecture
JES	job entry system		

Bibliography

Note: Manuals may have different order numbers for different releases of a program. To verify the order number you need, refer to System/370 and 4300 Bibliography, GC20-0001, or contact an IBM Branch Office. Program directories contain order numbers for the manuals needed for the program. A program directory is delivered with the program. A program general information manual also lists manuals related to the program.

	Order Number		Order Number
MVS/SP Version 1 with DFDS		MVS/SP with MVS/370 DFP	
<i>IBM System/370 Principles of Operation</i>	GA22-7000	<i>IBM System/370 Principles of Operation</i>	GA22-7000
<i>MVS/System Product General Information</i>	GC28-1025	<i>MVS/System Product General Information</i>	GC28-1025
<i>OS/VS2 System Generation Reference</i>	GC26-3792	<i>MVS/370 System Generation Reference</i>	GC26-4063
<i>OS/VS Message Library: VS2 System Messages</i>	GC38-1002	<i>OS/VS Message Library: VS2 System Messages</i>	GC38-1002
<i>Operator's Library: OS/VS2 MVS System Commands</i>	GC38-1031	<i>Operator's Library: OS/VS2 MVS System Commands</i>	GC38-1031
<i>OS/VS2 MVS Utilities</i>	GC26-3902	<i>MVS/370 Utilities</i>	GC26-4065
<i>OS/VS2 MVS JCL</i>	GC28-0692	<i>MVS/370 Data Facility Product: General Information</i>	GC26-4050
<i>Data Facility Device Support, General Information</i>	GC26-3954	<i>MVS/370 System Programming Library: Data Management</i>	GC26-4056
<i>Data Facility Extended Function: General Information</i>	GC26-3960	<i>MVS/370 Data Management Macro Instructions</i>	GC26-4057
<i>OS/VS2 MVS Sequential Access Method - Extended (SAM-E) General Information</i>	GH20-9122	<i>MVS/370 Data Management Services</i>	GC26-4058
<i>OS/VS2 System Programming Library: Data Management</i>	GC26-3830	<i>MVS/370 VSAM Reference</i>	GC26-4074
<i>OS/VS2 MVS Data Management Macro Instructions</i>	GC26-3873	MVS/SP Version 2 (MVS/XA) with MVS/XA DFP	
<i>OS/VS2 MVS Data Management Services Guide</i>	GC26-3875	<i>IBM System/370 Extended Architecture Principles of Operation</i>	SA22-7085
<i>OS/VS Virtual Storage Access Method (VSAM) Programmer's Guide</i>	GC26-3838	<i>MVS/System Product Version 2 General Information</i>	GC28-1118
		<i>MVS/Extended Architecture Installation: System Generation</i>	GC26-4009
		<i>MVS/Extended Architecture Message Library: System Messages (2 volumes)</i>	GC28-1376 and GC28-1377
		<i>MVS/Extended Architecture Operations: System Commands</i>	GC28-1206
		<i>MVS/Extended Architecture Data Administration: Utilities</i>	GC26-4018

	Order Number	VSE/Advanced Functions and DSF	Order Number
<i>MVS/Extended Architecture Job Control Language (JCL)</i>	GC28-1148	<i>The VSE System, General Information</i>	GC33-6108
<i>MVS/Extended Architecture Data Facility Product: General Information</i>	GC26-4007	<i>VSE/Advanced Functions Planning and Installation</i>	SC33-6096
<i>MVS/Extended Architecture System-Data Administration</i>	GC26-4010	<i>VSE/Advanced Functions System Control Statements</i>	SC33-6095
<i>MVS/Extended Architecture Data Administration: Macro Instruction Reference</i>	GC26-4014	<i>VSE/Advanced Functions Operating Procedures</i>	SC33-6097
<i>MVS/Extended Architecture Data Administration Guide</i>	GC26-4013	<i>VSE/Advanced Functions System Utilities</i>	SC33-6100
<i>MVS/Extended Architecture VSAM Administration: Macro Instruction Reference</i>	GC26-4016	<i>Device Support Facilities User's Guide and Reference</i>	GC35-0033
OS/VS1 with DFDS		<i>VSE System IPO/E Planning Guide</i>	GC20-2003
<i>IBM System/370 Principles of Operation</i>	GA22-7000	<i>VSE/Fast Copy Data Set Installation Reference</i>	SC33-6082
<i>Data Facility Device Support: General Information</i>	GC26-3954	<i>DL/I DOS/VS Resource Definition and Utilities</i>	SH24-5021
<i>OS/VS1 System Generation Reference</i>	GC26-3791	<i>DOS/VS RPG II Language</i>	SC33-6031
VM/SP		<i>DOS/VS RPG II Installation Reference</i>	SC33-6032
<i>VM/SP General Information</i>	GC20-1838	<i>VSE System Support for IBM 3380 and 3370 Models A02/B02</i>	SC33-6170
<i>VM/SP Operators Guide</i>	SC19-6202	Airline Control Program/Transaction Processing Facility	
<i>VM/SP Planning and System Generation Guide</i>	SC19-6201	<i>Airline Control Program/Transaction Processing Facility General Information Manual</i>	GH20-2140
<i>VM/SP System Messages and Codes</i>	SC19-6204	<i>ACP System Concepts and Facilities</i>	GH20-1473
<i>VM/SP System Programmer's Guide</i>	SC19-6203	Additional Programming Support	
<i>Virtual Machine/System Product High Performance Option General Information</i>	GC19-6221	<i>Data Facility Data Set Services: User's Guide and Reference</i>	SC26-3949
<i>Virtual Machine/System Product High Performance Option: Planning and System Generation Guide</i>	SC19-6223	<i>Device Support Facilities User's Guide and Reference</i>	GC35-0033
<i>Virtual Machine/System Product High Performance Option: System Programmer's Guide</i>	SC19-6224	<i>Direct Access Storage Device Migration Aid, General Information</i>	GC26-3972

	Order Number
<i>3880 Models 1, 2, 3, and 4 Description</i>	GA26-1661
<i>3880 Model 13 Description</i>	GA32-0062
<i>3880 Model 23 Description</i>	GA32-0083
<i>Environmental Recording, Editing, and Printing (EREP) Program</i>	GC28-1378
<i>IBM Disk Storage Management Guide, Background Reference Information, contains introductory-level reference information that may be needed as background for disk storage management tasks.</i>	GA26-1675
<i>IBM Disk Storage Management Guide, Error Handling, contains guidance on interpreting EREP System Exception Reports and using the Device Support Facilities program for data errors. The IBM Disk Storage Management Guide manuals can be ordered as a set, by specifying order number GB0F-1205.</i>	GA26-1672
<i>OS/VS2 MVS Hierarchical Storage Manager: General Information</i>	GH35-0007

Index

A

- access mechanisms
 - address 35
 - capacity 20
 - data placement 59, 60
 - internal paths 29
 - number of 19
 - related to volume iii, 19
- access methods
 - block size considerations 56
 - data set conversion 53
 - 3380 support 49
- Access Method Services (IDCAMS)
 - EXPORT command 54
 - IMPORT command 54
 - REPRO command 54
 - VSAM 53
- actuator (*see* access mechanisms)
- address, absolute
- address, controller
 - assignment 35-37
 - JCL change 53
 - not selectable 38
 - odd or even 35
 - system generation 36
- address, cylinder
 - count area 21
 - home address 21
 - record 21
 - seek command 38
- address, device
 - assignment 35
 - JCL change 53
 - not selectable 38
 - physical ID 65
 - sense information 62
 - system generation 36
- address, head
 - count area 21
 - home address 21
 - number 38
 - record 21
 - seek command 38
 - sense information 62
- address, home
 - cylinder and head 21
 - data check 63
 - initialization 44
 - rewrite 44
 - track address 21
- address, I/O
 - assignment 35, 82
 - JCL change 53
 - not selectable 38
 - sense information 62
 - system generation 36
- address, record
 - command 25
 - count area 21
 - cylinder and head 21, 38
- address, seek
 - command 25
 - cylinder and head 21, 38
 - not valid 38

- address, storage director
 - assignment 35
 - JCL change 53
 - sense information 62
 - system generation 36
- address, unit
 - (*see* address, I/O)
- air condition 34
- Airline Control Program 32, 49
- alternate
 - controller 3, 29
 - path 3, 29
 - track assignment 44, 63
- attachment 3-12, 33

B

- backup
 - conversion 53
 - frequency 67
 - program support 51, 67
- block
 - data control 53
 - physical record 22, 56, 72
- block size
 - access method maximum 57
 - blocking aid 58
 - changing 53
 - considerations 56
 - main storage consideration 59
 - system data sets 58
 - track capacity 72
 - 3380 optimum 58
- buffer
 - size, main storage 58
 - speed matching 8, 27, 33

C

- cable length 34
- capacity
 - access mechanism (volume) 20
 - calculations 72
 - comparisons 20
 - cylinder 20
 - track 20
- cache
 - attachment 3
 - commands 26
- catalog
 - data conversion 53
 - support 49
- chain
 - channel commands 25
 - data 25
- channel
 - attachment 8, 33
 - data rate 8
 - program 25
 - switching 17
- channel commands
 - Define Extent 26, 28
 - Device Release 25
 - Device Reserve 25
 - Diagnostic Read Home Address 25
 - Diagnostic Sense/Read 25

- Diagnostic Write Home Address 25
- Erase 25
- Locate Record 26
- No-Operation 24
- Read and Reset Buffered Log 25
- Read Count 24
- Read Count, Key, and Data 24
- Read Data 24
- Read Home Address 24
- Read Initial Program Load 24
- Read Key and Data 24
- Read Multiple Count, Key, and Data 24
- Read Record 0 24
- Read Sector 24
- Recalibrate 23
- Search Home Address Equal 24
- Search Identifier Equal 24
- Search Identifier Equal or High 24
- Search Identifier High 24
- Search Key Equal 24
- Search Key Equal or High 24
- Search Key High 24
- Seek 23
- Seek Cylinder 23
- Seek Head 23
- Sense 25
- Sense ID 25
- Sense Path Group ID 25
- Set File Mask 23
- Set Path Group ID 23
- Set Sector 23
- Space Count 23
- Transfer-In-Channel 24
- Unconditional Reserve 25
- Write Count, Key, and Data 24
- Write Count, Key, and Data Next Track 26
- Write Data 24
- Write Home Address 24
- Write Key and Data 24
- Write Record Zero 24
- Write Update Data 26
- Write Update Key and Data 26
- channel commands for speed matching buffer 25, 27
- channel commands for cache 26
- channel program 25
- command overrun 62
- commands
 - cache 26
 - channel 23-26
 - speed matching buffer 25, 27
 - system operator console 39
- command retry
 - error handling 63
 - standard feature 26
- control panel 39
- controllers, 3380
 - addresses 35
 - alternate 29, 63
 - functions 3
 - internal paths 29
 - Models A4 and AA4 3, 29
 - path activation 40
- conversion, data sets 53

D

- data
 - length 21, 57, 73
 - record area 21
- data chaining 25
- data control block
 - JCL 53
 - macro 53
- data conversion
 - by access method 53
 - how to 53
- data definition (DD) 53
- data errors 63
- Data Facility Data Set Services
 - backup/recovery 67
 - data conversion 53
 - description 51
- Data Facility Data Set Services 51
- Data Facility Product 51
- data length
 - blocksize 53
 - effective use 55
 - space calculations 72
- data overrun 63
- data rate 8
- data sets
 - conversion 53
 - logical record length 22
 - physical record length 22, 72
 - placement 58
 - space allocation 53
- data streaming
 - ACP/TPF 49
 - attachment requirements 33
 - channel 8
 - performance 28, 59
 - 3880 Model 13 14
 - 3880 Model 23 14
- data verification 63
- defect, surface 44, 62
- device
 - physical ID 63
 - reserve/release 25
 - status 22, 39
- device, address
 - assignment 35
 - JCL change 53
 - not selectable 38
 - sense information 62
 - system generation 36
- Device support, programming 49
- Device Support Facilities 44
- direct access device space management (DADSM) 49
- direct access method (DAM) 53
- disable 39
- DISPLAY operator commands 43
- dump 53, 67
- dynamic path selection
 - ACP support 32
 - control program support 31
 - controls 28
 - controllers 3, 28
 - device reserve 31
 - dynamic reconnection 31
 - internal paths 29
 - MVS/SP support 29
 - overlapping operations 29
 - path group ID 31
 - VM/SP support 32
- dynamic reconnection 31

E

emergency 39
Enable/Disable switch 39
Environmental Recording, Editing, and Printing (EREP) 51, 55, 58
equipment, errors 64
errors 63, 64
error recovery procedures (ERP) 62

F

floor space 34

G

group, path ID
channel commands 23
description 31
dynamic path selection 31
operator controls 39

H

head (*see* address, head)
number of 20
head and disk assembly
initialization 39
number of access mechanisms 20
Hierarchical Storage Manager
backup 67
data set conversion 53
description 52
home address
commands 24
data errors 63
initialization 44
track format 21

I

IBM Direct Access Storage Migration Aid
conversion tool 53
description 51
JCL changes 53
index
track 21
VTOC 45
indexed sequential access method (ISAM)
block size 57
data management 50
data set conversion 54
dump/restore 67
initial program load 39
initial microprogram load 39
initialization 44
installation 44
I/O address (*see* address)

J

Job Control Language (JCL) 53
Job Entry System (JES) 48

K

key
length 72
record area 21

L

logging 65
logical records 56

M

main storage
buffers 57
requirements 69
migration, (*see* conversion)
models AA
attachment to 3880 4
device reserve/release 31
dynamic path selection 29
dynamic reconnection 31
operating system support 47
overlapping operations 29
models, 3380 3-6
models, 3880
processor attachment 33
3380 attachment 3
multiple track operations 26
Multiple Virtual Storage (MVS)
data set conversion 47
dynamic path selection 29
system generation 47
System/370 47
System/370 XA 47
uses 47
versions 49

O

operating system support 47
operator controls 39
operator commands 40
overruns 63

P

page data set
effective 3380 use 60
MVS/SP support 48
VM/SP 48
parameters 53
partitioned access method
block size consideration 56
data set conversion 53
dump/restore 67
paths
enable/disable 39
VARY commands 39
external 29
internal 29
path group ID
channel commands 23
description 32
dynamic path selection 31
operator controls 40

- performance
 - comparisons 1, 20
 - considerations 59
 - dump/restore 67
 - dynamic path selection 31
- physical identifiers
 - assignment 44
 - control panel 39
 - use 44
- physical planning 34
- physical record
 - equal length 72
 - number per cylinder 57
 - number per track 57
 - related to block size 21, 56
 - related to logical record 21, 56
 - unequal length 73
- power on/off 39
- processors
 - requirements 33
 - 3380 attachment 1, 33
- programs
 - 3380 support 47

R

- read
 - commands 24
 - error handling 63
- record format 21
- record 0
 - track format 21
 - commands 24
 - initialization 44
 - error handling 62
- recovery
 - error 62
 - backup 67
- reserve/release 31
- reset 40
- rotational position sensing
 - commands 23
 - standard feature 26
- rotational time 20

S

- search commands 24
- seek
 - address 38
 - command 23
 - time 20
- sense
 - commands 25
 - dynamic path selection 31
 - error handling 62
- sequential access method (SAM)
 - data management 49
 - extended (SAM-E) 27, 49
 - dump/restore 67
- sequential organization
 - data management 49
 - data migration 53
- space
 - calculation formulas 72
 - data set allocation 53, 59
 - floor requirements 34
- speed matching buffer
 - channel speed 8
 - commands 25

- operations 27
- performance 27, 59
- program support 27
- processor requirement 33
- spooling 47
- status
 - maintained by 3380 62
 - reported 62
 - unit check status 38, 62
 - unit status 40
 - 3380 status 40
- storage directors, 3380
 - address 35
 - channel switching 17
 - speed matching buffer 8
 - 3380 attachment 3
- switch
 - channel feature 17
 - 3380 control panel 39
 - 3380 control panel 39
- system control programs 47
- system error log record 62, 65
- system error recovery procedures 62
- system generation
 - ACP 49
 - address assignment 35
 - MVS/SP 47
 - VM/SP 47
- System Product 47
- system support 47
- System 370
 - attachment requirements 33
 - channel data rates 8, 33
 - data streaming 33
 - extended architecture 47
 - speed matching buffer 33
 - operating system support 47

T

- time sharing option (TSO) 45
- track
 - address 21
 - alternate 44, 63
 - capacity 19, 72
 - defective 63
 - descriptor record 21, 44, 63
 - format 21
 - home address 21, 44, 63
 - index 21
 - initialization 44
 - number of 21
- TRKCALC 73

U

- units 3
- unit address
 - in error message 62
 - in sense information 62
 - system generation 36
 - valid I/O 37
- unit check 37, 62
- unit emergency switch 39
- unit status 43, 62
- utilities 54

V

VARY commands 32, 39

virtual machine (VM)

data set conversion 53

dynamic path selection 31

system control program 47

system generation 47

uses 47

3380 model AA support 31

virtual storage access method (VSAM)

block size consideration 56

data set conversion 53

data management 49

dump/restore 67

VM/SP commands 54

volume

access mechanism 19

capacity 20

conversion 53

definition 19

initialization 44

label 44

serial 44

volume table of contents (VTOC) 44

VTOC index 45

W

write commands

description CCW 24

speed matching buffer 25

cache 26

IBM 3380 Models A04, AA4, and B04
Direct Access Storage
Description and User's Guide

READER'S
COMMENT
FORM

Publication Number
GA26-1664-3

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

You may use this form to communicate your comments about this publication, its organization, or subject matter, with the understanding that IBM may use or distribute whatever information you supply in any way it believes appropriate without incurring any obligation to you.

Your comments will be sent to the author's department for whatever review and action, if any, is deemed appropriate. Comments may be written in your own language; use of English is not required.

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

Possible topics for comment are:

Clarity Accuracy Completeness Organization Coding Retrieval Legibility

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

What is your occupation ? _____

Number of latest Newsletter associated with this publication: _____

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

Note: Staples can cause problems with automated mail sorting equipment.
Please use pressure sensitive or other gummed tape to seal this form.

Cut Along This Line

Reader's Comment Form

IBM 3380 Models A04, AA4, and B04 Direct Access Storage Description and User's Guide (File No. S/370-07 4300-07) Printed in U.S.A. GA26-1664-3

Cut or Fold Along Line


Fold and Tape

Please Do Not Staple

Fold and Tape



NO POSTAGE
NECESSARY
IF MAILED
IN THE
UNITED STATES



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

POSTAGE WILL BE PAID BY ADDRESSEE:

International Business Machines Corporation
Department J62
555 Bailey Avenue
San Jose, California 95150

Fold and Tape

Please Do Not Staple

Fold and Tape



IBM 3380 Models A04, AA4, and B04
Direct Access Storage
Description and User's Guide

READER'S
COMMENT
FORM

Publication Number
GA26-1664-3

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

You may use this form to communicate your comments about this publication, its organization, or subject matter, with the understanding that IBM may use or distribute whatever information you supply in any way it believes appropriate without incurring any obligation to you.

Your comments will be sent to the author's department for whatever review and action, if any, is deemed appropriate. Comments may be written in your own language; use of English is not required.

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

Possible topics for comment are:

Clarity Accuracy Completeness Organization Coding Retrieval Legibility

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

Note: Staples can cause problems with automated mail sorting equipment.
Please use pressure sensitive or other gummed tape to seal this form.

Cut Along This Line

What is your occupation ? _____

Number of latest Newsletter associated with this publication: _____

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

Reader's Comment Form

Cut or Fold Along Line

IBM 3380 Models A04, AA4, and B04 Direct Access Storage Description and User's Guide (File No. S/370-07, 4300-07) Printed in U.S.A. GA26-1664-3

Fold and Tape

Please Do Not Staple

Fold and Tape



NO POSTAGE
NECESSARY
IF MAILED
IN THE
UNITED STATES



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

POSTAGE WILL BE PAID BY ADDRESSEE:

International Business Machines Corporation
Department J62
555 Bailey Avenue
San Jose, California 95150

Fold and Tape

Please Do Not Staple

Fold and Tape



Publication Number
GA26-1664-3

File Number
S/370-07, 4300-07

Printed in
USA

IBM

GA26-1664-03

