

GA26-1664-1  
File No. S/370-07, 4300-07

**Systems**

**IBM 3380**  
**Direct Access Storage**  
**Description and User's Guide**

**IBM**

## 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 used in the United States of America.

### Second Edition (December 1981)

This is a major revision of and obsoletes *IBM 3380 Direct Access Storage Description and User's Guide*, Order No. GA26-1664-0, and Technical Newsletter TNL GN 26-0362. Technical changes to the original publication are indicated by a vertical line to the left of the changes.

Information in this publication is subject to change. Before using this publication in connection with the use of IBM equipment, contact the local IBM branch office for revisions.

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

Copies of this and other IBM publications can be obtained through IBM branch offices.

A form for reader's comments is provided at the back of this publication. If the form has been removed, send your comments to the address below.

This manual was prepared by the IBM General Products Division, Technical Publishing, Department G26/H15, 555 Bailey Ave., San Jose, California, 95150.

© Copyright International Business Machines Corporation 1981.

## Preface

This manual describes the 3380 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 *IBM 3380 Storage Control Description* is the manual for detailed reference information on channel commands, sense bit definitions, and 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.

- 3081 Processor support is added.
- 3380 Storage Control Model 13 is added.
- Information regarding fixed heads is removed.
- Information and guidelines are added on the use of the 3380 Speed Matching Buffer for the 3380 feature.
- Information is added on VM/SP and ACP/TPF support of 3380 model AA.
- Information on error handling is expanded with more precise guidelines on use of Device Support Facilities and assignment of physical IDs. Error handling includes the capability of skipping up to seven defects on a track before assigning an alternate track. The entire section, "Device Status and Error Handling" is moved to a more appropriate location next to the section, "Backup and Recovery."
- Information on addressing is modified to describe new techniques that more accurately define the reason an addressed device is not available to perform operations. Also, information is added on an addressing requirement of the operating system.
- Information on 3380 operator controls and indicators is completely changed.
- A new section is added on initialization.
- A correction is made to information that previously described blocking of VSAM records inaccurately.

### 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*, Order No. GC20-1699, defines terms used in IBM publications.

### Bibliography

- *Airline Control Program/Transaction Processing Facility General Information Manual*, Order No. GH20-2140.
- *ACP System Concepts and Facilities*, Order No. GH20-1473.
- *Data Facility Device Support, General Information*, Order No. GC26-3954.
- *Data Facility Device Support: User's Guide and Reference*, Order No. SC26-3952, gives detailed information on processing disk storage volumes with indexed VTOCs.
- *Data Facility Data Set Services: User's Guide and Reference*, Order No. SC26-3949, describes dump and restore functions.
- *Data Facility Extended Function: General Information*, Order No. GC26-3960, describes catalog programming.
- *Device Support Facilities*, Order No. GC35-0033, describes the commands to initialize and maintain disk storage volumes.
- *Direct Access Storage Device Migration Aid, General Information*, Order No. GC26-3972.

- *Environmental Recording, Editing, and Printing (EREP) Program*, Order No. GC28-0772.
- *IBM 3880 Storage Control Description*, gives detailed descriptions of channel commands, and status and sense information. For 3880 Models 2 and 3, use Order No. GA26-1661. For 3880 Model 13, use Order No. GA32-0062.
- *IBM Input/Output Equipment Installation Manual – Physical Planning for System/360, System/370, and 4300 Processors*, Order No. GC22-7064, gives information for planning floor space requirements and other physical planning considerations for the 3380.
- *MVS Initialization and Tuning Guide*, Order No. 6C28-1029.
- *Operator's Library: OS/VS2 MVS System Commands*, Order No. GC38-0229, describes the commands that can be issued from the system operator consoles.
- *OS/VS2 Access Method Services*, Order No. GC26-3841, describes the service program commands used to manipulate VSAM data sets.
- *OS/VS2 Message Library: VS2 System Messages*, Order No. GC38-1002, describes the messages issued on the system operator consoles.
- *OS/VS MVS Data Management Macro Instructions*, Order No. GC26-3873, describes the macro instructions used for the access method routines (except VSAM).
- *OS/VS2 MVS Data Management Services Guide*, Order No. GC26-3875, describes data management, including how to define a data set and use of the access methods (except VSAM).
- *OS/VS2 MVS Hierarchical Storage Manager: General Information*, Order No. GH35-0007.
- *OS/VS2 MVS JCL*, Order No. GC28-0692, contains information necessary to code Job Control Language (JCL), Job Entry Subsystem 2 (JES2) control statements, and Job Entry Subsystem 3 (JES3) control statements.
- *OS/VS2 MVS System Product General Information Manual*, Order No. GC28-1025, describes the functions of MVS/System Product. This manual should be referred to for a comprehensive list of publications needed for each release of MVS/System Product.
- *OS/VS2 MVS Utilities*, Order No. GC26-3902, describes how to use OS/VS utilities under OS/VS2 MVS. It includes information on IEHMOVE, IEBGENER, IEBISAM, and IEBCOPY, which can be used to convert data sets to the 3380.
- *OS/VS2 System Programming Library: Data Management*, Order No. GC26-3830, provides information on how to modify and extend the data management provisions of the system control programs. It is the source for information on issuing EXCP directly instead of using the access method routines.
- *OS/VS2 System Programming Library: System Generation Reference*, Order No. GC26-3792.
- *OS/VS2 Virtual Storage Access Method (VSAM) Programmer's Guide*, Order No. GC26-3838, describes how to create VSAM data sets and code the macro instructions required to process them.
- *VM/SP General Information Manual*, Order No. GC20-1838, describes the functions of VM/System Product. This manual should be referred to for a comprehensive list of publications needed for 3380 operations under VM/SP.
- *VM/SP Operators Guide*, Order No. SC19-6202.
- *VM/SP Planning and System Generation Manual*, Order No. SC19-6201.
- *VM/SP System Messages and Code*, Order No. SC19-6204.
- *VM/SP System Programmer's Guide*, Order No. SC19-6203.

# Contents

<b>Preface</b> .....	iii	<b>Emergency Power Off</b> .....	21
<b>Technical Changes and Additions</b> .....	iii	<b>AC Power Indicator</b> .....	21
<b>Terminology</b> .....	iii	<b>Path and Device On and Off</b> .....	21
<b>Bibliography</b> .....	iii	<b>Device Status</b> .....	21
<b>Figures</b> .....	vi	<b>Physical Identifier Labels</b> .....	21
<b>IBM 3380 Quick Reference</b> .....	vii	<b>Hardware Planning Considerations</b> .....	22
<b>Introduction</b> .....	1	<b>Processor Attachment Requirements</b> .....	22
<b>Models and Attachment Configurations</b> .....	3	<b>Physical Planning</b> .....	23
<b>Models</b> .....	3	<b>Power and Air Conditioning</b> .....	23
<b>Methods of Attachment to the 3880 Storage Control</b> ..	3	<b>Space Requirements</b> .....	23
<b>Model AA Attachment</b> .....	3	<b>Cable Requirements</b> .....	23
<b>Channel Attachment and Transfer Rate</b> .....	6	<b>Addressing</b> .....	24
<b>Data Streaming and Speed Matching Buffer</b> .....	6	<b>I/O Address</b> .....	24
<b>Channel Switches</b> .....	6	<b>Address Assignments</b> .....	24
<b>Addressed Selection Not Possible</b> .....		<b>Addressed Selection Not Possible</b> .....	25
<b>Seek Address</b> .....		<b>Seek Address</b> .....	25
<b>Seek Address Not Valid</b> .....		<b>Seek Address Not Valid</b> .....	25
<b>Device Characteristics and Formats</b> .....	8	<b>Program Planning Considerations</b> .....	27
<b>Head and Disk Assembly</b> .....	8	<b>MVS/SP Support</b> .....	27
<b>Characteristics Summary</b> .....	8	<b>System Product and Data Facility Device</b>	
<b>Record and Track Formats</b> .....	10	<b>Support Programs</b> .....	27
<b>Record Format</b> .....	10	<b>MVS/SP System Generation</b> .....	27
<b>Track Format</b> .....	10	<b>3380 MVS/SP Releases 1 and 3 Uses</b> .....	28
<b>Physical Records Per Track</b> .....	11	<b>Programs Using OS/VS and the 3380</b> .....	28
<b>Operations</b> .....	12	<b>VM/SP Support</b> .....	28
<b>Channel Commands</b> .....	12	<b>VM/SP System Generation</b> .....	28
<b>Control Commands</b> .....	12	<b>3380 VM/SP Uses</b> .....	28
<b>Search Commands</b> .....	13	<b>Airline Control Program/Transaction Processing</b>	
<b>Read Commands</b> .....	13	<b>Facility Support</b> .....	29
<b>Write Commands</b> .....	13	<b>ACP/TPF System Generation</b> .....	29
<b>Sense Commands</b> .....	14	<b>ACP/TPF 3380 Restrictions</b> .....	29
<b>Diagnostic Commands</b> .....	14	<b>Additional Programming Support</b> .....	29
<b>Command Sequence</b> .....	14	<b>Data Facility Extended Function</b> .....	29
<b>Commands for Adjusted Data Rate</b> .....	14	<b>Data Facility Data Set Services</b> .....	29
<b>Standard Features</b> .....	15	<b>Device Support Facilities</b> .....	30
<b>Speed Matching Buffer in 3880 Models 2 and 3</b> .....	16	<b>IBM Direct Access Storage Device</b>	
<b>Operations for 1.5 MB/sec</b> .....	16	<b>Migration Aid</b> .....	30
<b>Operations for 1.5 and 3.0 MB/sec, Same Processor</b>	17	<b>Hierarchical Storage Manager</b> .....	30
<b>Model AA Attached to 3880 Model 13</b> .....	17	<b>Sort/Merge</b> .....	30
<b>Overlapping Operations, Models A and AA</b> .....	18	<b>Environmental Recording, Editing, and</b>	
<b>Internal Paths and Device Access, Models A and AA</b> ..	18	<b>Printing</b> .....	30
<b>Dynamic Path Selection, Model AA</b> .....	18	<b>Volume Initialization and VTOC Index</b> .....	31
<b>Model AA with VM/SP or ACP/TPF</b> .....	20	<b>Data Management and Data Set Conversion</b> .....	33
<b>Dynamic Reconnection, Model AA</b> .....	20	<b>Job Control Language Uses</b> .....	33
<b>Device Reserve/Release Changes</b> .....	20	<b>Data Set Conversion Under MVS/SP</b> .....	34
<b>Standard Access Methods</b> .....		<b>User-Written Access</b> .....	34
<b>Operator Controls and Indicators</b> .....	21	<b>Data Set Conversion Under VM/SP</b> .....	34
<b>Power On and Off</b> .....	21		

<b>Effective Use of the 3380</b> .....	35
Criteria for Using the 3380 .....	35
Space and Performance Guidelines .....	35
Optimum Blocking for 3380 Records .....	35
Performance Considerations .....	38
Application Considerations .....	39
Page Data Sets, Special Considerations .....	39
<b>Device Status and Error Handling</b> .....	41
Control and Function Errors .....	41
Data Errors .....	41
Correction of Data Errors .....	42
Response to Data Errors .....	43
MVS/IEHATLAS .....	44
Data Verification by Program .....	44
Equipment Errors .....	44
Service Procedures .....	44
Advantages of Two Controllers .....	44
Statistical Information and Logging .....	45
Physical Identifiers .....	46

<b>Data Backup and Recovery</b> .....	47
<b>Conversion Checklist</b> .....	49
<b>Appendix A: Space Calculations</b> .....	51
Formulas .....	51
Equal Length Physical Records .....	51
Unequal Length Physical Records .....	52
Space Count Maintenance .....	52
Tables, Without and With Keys .....	53
Tables, Ranges of Key Lengths .....	56
<b>Appendix B: SAM-E Program Product</b> .....	60
<b>Appendix C: Addressing Reference Information</b> .....	61
<b>Abbreviations</b> .....	62
<b>Index</b> .....	63

## Figures

Figure 1. Example of Four 3380 Model A Strings Attached to 3880 Storage Directors .....	4
Figure 2. Example of Two 3380 Model AA Strings Attached to 3880 Storage Directors .....	4
Figure 3. Representative 3880/3380 Valid Configurations .....	5
Figure 4. Channel Switching for 3880 Storage Directors and Attached 3380 Disk Storage ..	7
Figure 5. 3380 Characteristics .....	9
Figure 6. Comparison of Characteristics of Different IBM Disk Storage Types .....	9
Figure 7. Ratio of Access Mechanisms for Approximate Equivalent Capacity .....	9
Figure 8. Track and Record Formats, Simplified Schematic .....	11
Figure 9. Four Internal Paths in a 3380 Model A with One Controller .....	19
Figure 10. Four Internal Paths and Dynamic Path Selection in a 3380 Model AA with Two Controllers .....	19
Figure 11. 3380 Control Panel .....	21
Figure 12. Comparisons of Approximate Floor Space Requirements Based on Full String Configurations .....	23
Figure 13. 3880 and 3380 Addressing .....	24
Figure 14. Sets of Valid 3880/3380 Addresses .....	25
Figure 15. Addressing Example: Two Model A Strings Attached to One 3880 Storage Director ...	26
Figure 16. Addressing Example: Two Model AA Strings Attached to the Same Two 3880 Storage Directors .....	26
Figure 17. Volume Layout of Initialization Information .....	32
Figure 18. Differences in Device Space Requirements between the 3350 and 3380 .....	36
Figure 19. Data Error Correction .....	42
Figure 20. Physical ID Assignment .....	46
Figure 21. Disk-to-Tape Dumping .....	48

# IBM 3380 Quick Reference

## Processor Attachments

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

## Storage Control Attachments

3880 Models 2, 3, and 13

## 3380 Models

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

A4 – one controller, two drives  
AA4 – two controllers, two drives, dynamic path selection  
B4 – two drives

## Program Support

OS/VS MVS – base system control program  
System Product (SP) – system support  
Data Facility Device Support – data management support, VTOC index  
Data Facility Extended Function – catalog  
Data Facility Data Set Services – dump and restore  
Direct Access Migration Aid – data set conversion  
Device Support Facilities – device initialization and maintenance  
Hierarchical Storage Manager (HSM) – space management, dump and restore  
VM/SP – base system control program and system support  
Airline Control Program/Transaction Processing Facility – control program



## Introduction

The IBM 3380 Direct Access Storage is a disk storage device with a storage capacity of 2.5 gigabytes (billion bytes) per unit, an increase of almost four times the capacity of the IBM 3350 Direct Access Storage. 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 new 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. Like the 3350, 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, or 13 to any of the following IBM processors:

- 3031, 3032, and 3033 Processors
- 3042 Attached Processor Model 2
- 3081 Processor
- 4341 Processors

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

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

The 3880 Storage Control has 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 megabyte-per-second (MB/sec) data transfer rate. A speed matching buffer in the 3880 accommodates attachment to channels with data transfer rates of 1.5 MB/sec and 3.0 MB/sec. Special channel commands have been added to optimize performance of write operations when the speed matching buffer is used.

The Model 13 is a new model of the 3880 Storage Control. It contains electronic storage to retain frequently used 3380 data for faster access by the processor.

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.

For the first time, 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. This achieves two paths for command and data transfers. Both of the storage directors and channels have access to the data in a 3380 string at the same time, instead of one at a time as with string switching. The dynamic path selection function controls operation of the two controllers, including simultaneous data transfer over the two paths. When supported by the operating system, each controller can serve as an alternate path in the event the other controller is unavailable.

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 most of the same channel commands as the 3350 disk storage. Status, sense, and usage information are supplied by the 3880 Storage Control to the system in a similar manner as for the 3350.

An accumulation 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, and such programming techniques as chained scheduling and search direct for sequential data sets.

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 OS/VS2 MVS with the System Product (SP) and Data Facility Device Support program products. The operating system is referred to as MVS/SP. Data Facility Device Support includes provisions for adding an index to the volume table of contents (VTOC). Addition of an index reduces the time required to search the VTOC for a data set entry.

A catalog facility program product, Data Facility Extended Function, is an optional replacement of current VSAM and other catalogs, or it can coexist with them on the same system. It provides substantial improvements over existing catalog architecture and is available for MVS/SP.

Data Facility Data Set Services is a versatile program product that provides for dump and restore of full or partial volumes, and deletion of data sets. It supports current IBM disk storage and tape products, and the IBM 3850 Mass Storage System.

The Device Support Facilities program is provided to initialize the 3380 and as a tool for problem determination and recovery.

The Hierarchical Storage Manager program product provides space management in a storage hierarchy that includes disks, the IBM 3850 Mass Storage System, and tape devices. The program 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 OS/VS Sort/Merge program product supports the 3380.

A program product is available to help convert existing data sets to the 3380. This program is the IBM Direct Access Storage Device Migration Aid.

The 3380 also is supported by:

- Virtual Machine/System Product (VM/SP) program product, used alone or with MVS
- Airline Control Program/Transaction Processing Facility (ACP/TPF) program product

The 3380 can fill the need for increased disk storage capacity and, with the available programming support, meet 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 A4** contains four access mechanisms and one controller with control functions to attach to one 3880 storage director.

**Model AA4** contains four access mechanisms and two controllers. It includes control functions to attach to two storage directors and the logic required for dynamic path selection.

**Model B4** contains four access mechanisms and can attach to a 3380 Model A4 or AA4. Up to three Model B4s can attach to a model A or AA.

All units have the mechanical and electronic systems necessary to operate the disk and to write and read data. In addition to these functions, the controllers in the model As 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 A and model B are available at the time of manufacture only. Model changes between the model As can be made at the user's location.

## Methods of Attachment to the 3880 Storage Control

The 3380 attaches to the IBM 3880 Storage Control Models 2, 3, and 13. A 3380 Model A4 or AA4 can attach to a 3880 Storage Control Model 2 or 3. A 3880

Storage Control Model 13 attaches only 3380 Model AA4.

The 3380 attaches to a 3880 storage director in strings consisting of one model A or AA and up to three model Bs.

Two 3380 strings can attach to one of the storage directors of a 3880 Model 2, or two strings can attach to each storage director of a 3880 Model 3 or 13. See Figure 1 for an example.

If a 3880 Model 2 is used, one storage director can be used for 3380s and the other storage director for other types of disk storage. It is not possible to have other disk storage types on a 3880 Model 3 or 13.

### *Model AA Attachment*

A 3380 string with two controllers, Model AA4, attaches to two 3880 storage directors. See Figure 2.

The two storage directors can be in the same 3880 Model 3, or in different 3880 Models 2 or 3, or the two storage directors must be in the same 3880 Model 13.

A Model AA4 cannot attach to any two storage directors that share the same channel. The two storage directors must have different channels as shown in Figure 2.

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

Models AA4 cannot share the same storage director as Models A4.

Although the AA models can be attached to two storage directors in the same 3880 Model 3, the potential for even higher availability is achieved by attaching one controller to one of the storage directors in one 3880 and the other controller to a storage director in another 3880. With this attachment, if one 3880 should fail completely, the 3380 data is still accessible through the other 3880.

Figure 3 summarizes possible configurations.

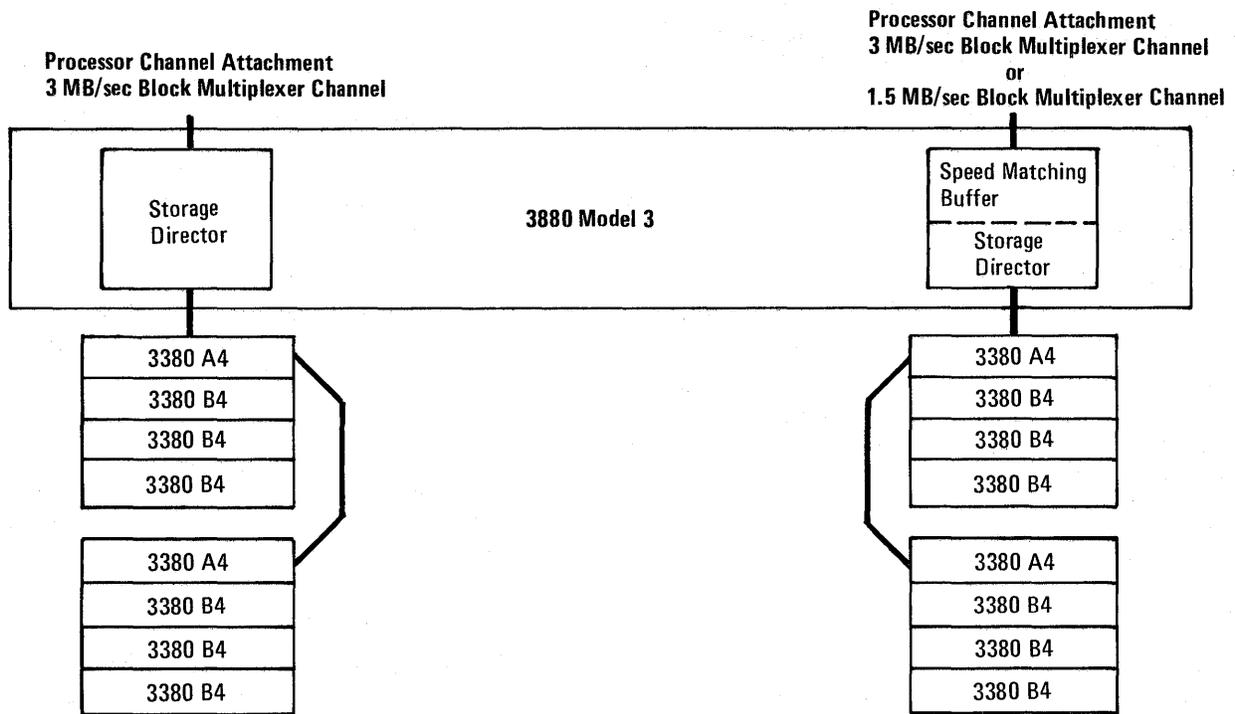


Figure 1. Example of Four 3380 Model A Strings Attached to 3880 Storage Directors

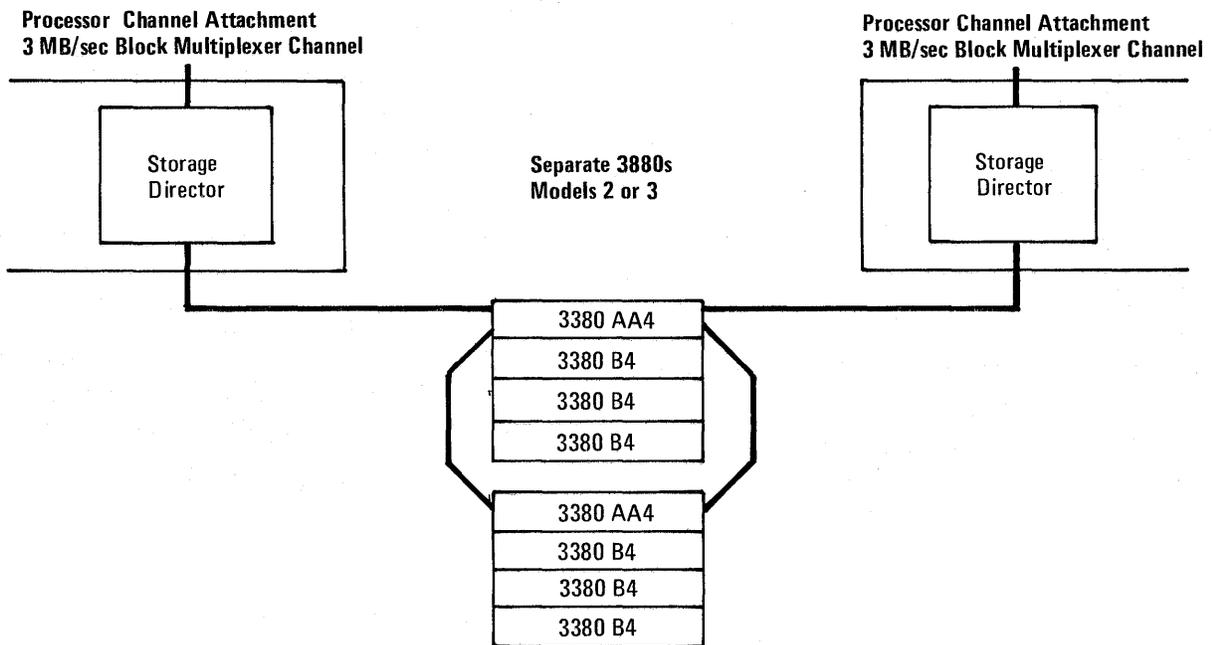
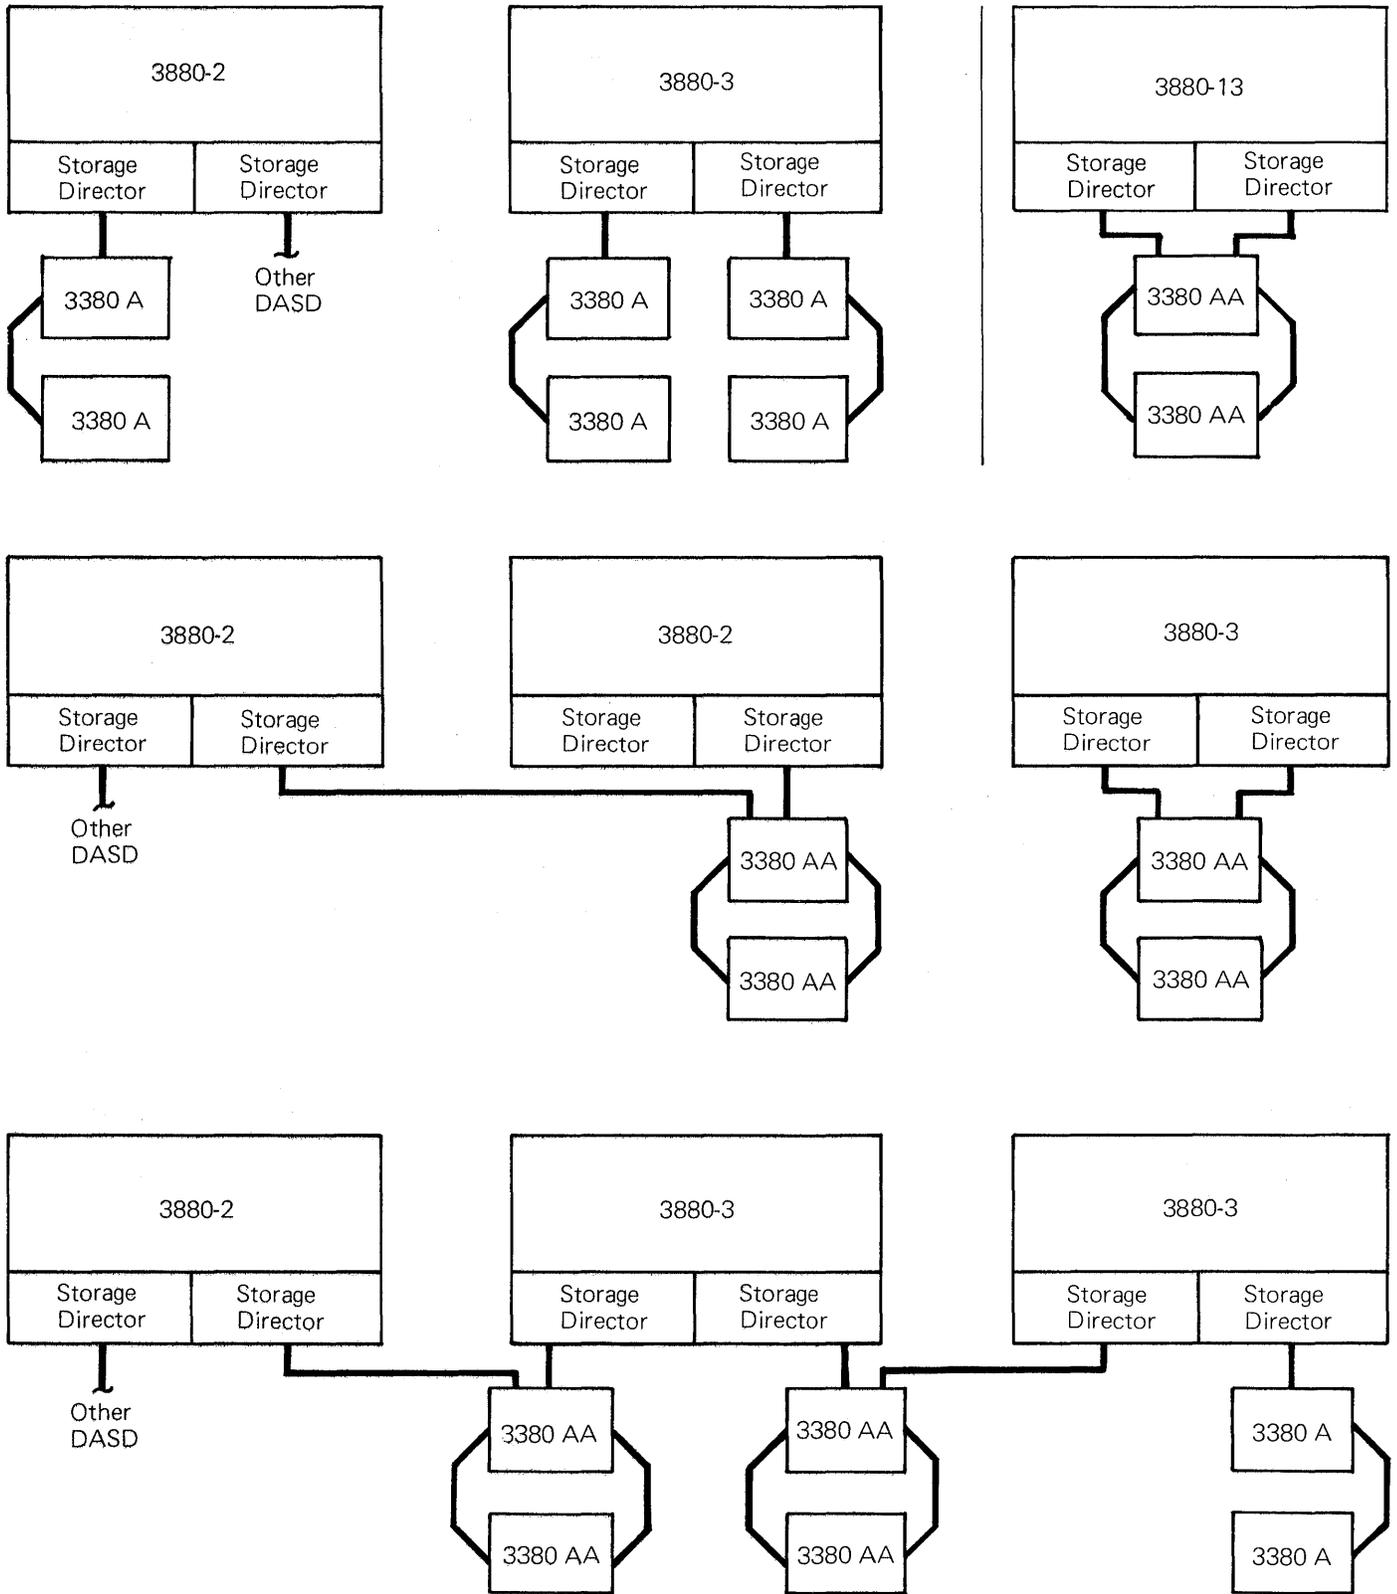


Figure 2. Example of Two 3380 Model AA Strings Attached to 3880 Storage Directors



**Figure 3. Representative 3880/3380 Valid Configurations**

### ***Channel Attachment and Transfer Rate***

The 3380 transfers data at 3.0 megabytes per second 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 rate and, thus, achieves the full 3.0 MB/sec data transfer potential of the 3380.

Data streaming is a special feature for 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 rate does not require a special feature for data streaming. However, an engineering change is needed on models manufactured before March 1981, as described in the section, "Hardware Planning Considerations." Data streaming is standard with the 3081 processor.

To also accommodate attachment to channels that do not have a 3.0 MB/sec data rate, a Speed Matching Buffer for 3380 special feature can be installed in 3880 Storage Control Models 2 and 3. With this buffer between the 3380 and the channel, processors that do not have a 3.0 MB/sec channel also can use the 3380. The speed matching buffer stores 32K bytes of data and operates at two speeds, 1.5 and 3.0 MB/sec, depending on the speed of the channel to which it is attached.

The speed matching buffer for 3380 feature can be installed in one or both storage directors of a 3880. This applies to 3380 model As attaching to the same or different processors and to 3380 model AAs attaching to different processors. Figure 1 shows a configuration where one storage director has a speed matching buffer and the other does not.

If a 3380 model AA is attached to two storage directors that connect to the same processor, both storage directors must have a speed matching buffer or neither must have one. This is because both paths must support the special commands used when a speed matching buffer is installed.

With a speed matching buffer in both storage directors, two attaching channels can operate at different speeds of 3.0 and 1.5 MB/sec. This configuration permits channels of different speeds to access the same 3380 devices from a single 3031, 3032 or 3033 processor.

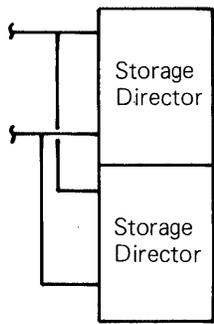
### ***Channel Switches***

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 three features can be used as described as follows and in Figure 4.

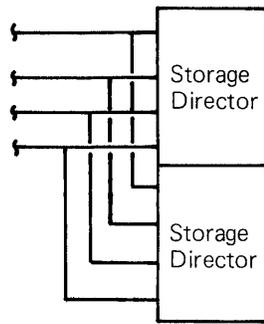
- With a two-channel switch pair feature, each storage director of a 3880 can connect to two different channels. Both storage directors can connect to the same or different channels.
- With a two-channel switch pair, additional feature, each storage director of a 3880 can connect to up to four different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

Channel switching is especially significant for 3380 model AAs with two controllers. In this configuration, 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.

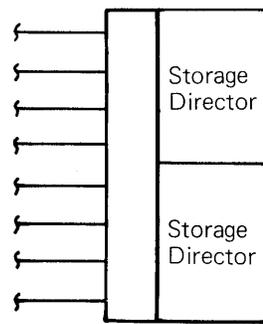
**Storage Directors Connected to the Same Channels**



**Two-Channel Switch Pair**

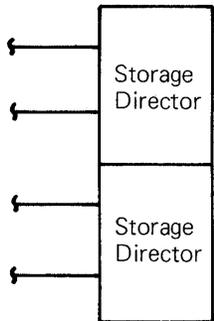


**Two-Channel Switch Pair, Additional**

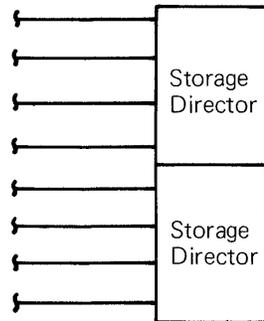


**Eight-Channel Switch**

**Storage Directors Connected to Different Channels**

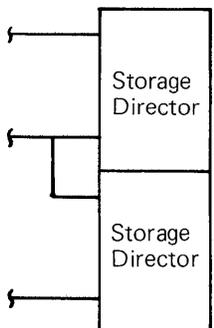


**Two-Channel Switch Pair (Four unique channels)**

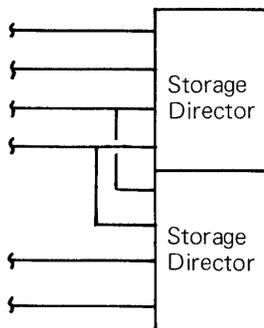


**Two-Channel Switch Pair, Additional (Eight unique channels)**

**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 4. 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 must be initialized with a track address and track descriptor record as described in the section, "Volume Initialization and VTOC Index."

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 5, 6, and 7 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 <sup>1</sup>	
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 <sup>2</sup>	8.3 ms
Alternate	1	Data Rate <sup>3</sup>	3.0 MB/sec
Customer Engineer	1	<b>Notes:</b>	
Data Tracks per Cylinder	15	1. <b>Seek time, or access motion time,</b> 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.	
Tracks per Access Mechanism		2. <b>Rotational delay time, or rotational latency time,</b> 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.	
Data	13,275	3. <b>Data rate</b> is the instantaneous speed at which bytes are transferred.	
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 5. 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

**Note:** The 3375 comparison has been added for this edition.

Figure 6. Comparison of Characteristics of Different IBM Disk Storage Types

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

Figure 7. 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. 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 8.

**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 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.

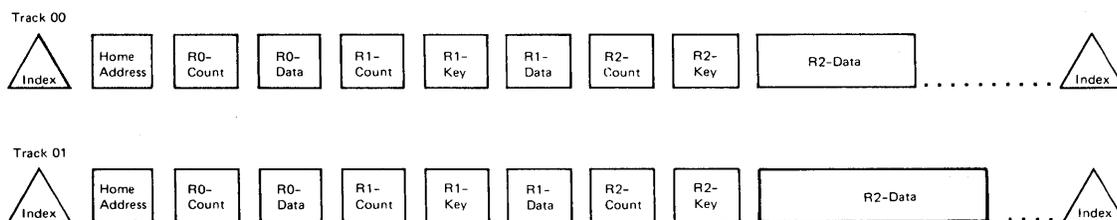
When preparing channel programs, the physical record size is the number of bytes in the data area, as defined in the count area as the data length (DL). This is specified with a channel Write Count, Key, and Data command or Write Count, Key, and Data Next Track command.

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/VS 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 5 and 6 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**



**Figure 8. Track and Record Formats, Simplified Schematic**

# 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 model 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.

Channel commands are executed by the storage director to which the 3380 is attached and are described in detail in the *IBM 3880 Storage Control Description* manual. The description of each command is accompanied by a description of the status conditions associated with the execution of the command. Also included are explanations of the channel address word (CAW), channel command word (CCW), channel status word (CSW), and program status word (PSW), including the formats.

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.

**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 execution 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 execution 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	**

\* 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. Execution of 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 with a single Start Input/Output (SIO) instruction. 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 Adjusted Data Rate***

New commands are used when the speed matching buffer for 3380 feature is used in the 3880 storage director. 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 executed 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 Sector and Search ID commands. The new write commands are for writing multiple equal length records. These commands are also supported in read operations.

**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.

### 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* manual.

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

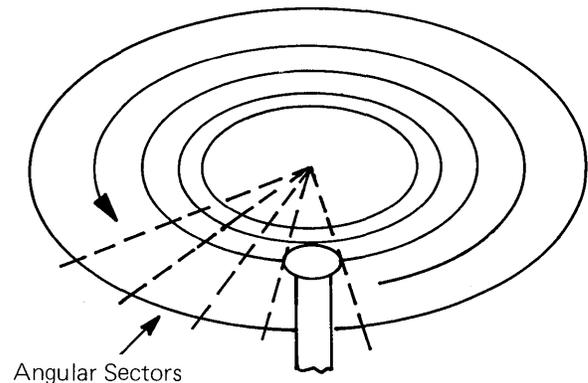
**Command retry** is a channel and storage director procedure that causes a channel command to be automatically retried. The retry does not cause an I/O interruption in the processor, and programmed error recovery procedures are not required. One of its uses is to correct certain data and seek errors 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. Execution of the 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.

### *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 new command extensions, but there is a difference in performance.

Without use of the Define Extent and Locate commands, execution of 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 command can be executed successfully when the correct record location is reached on the disk.

## Programming Support

The new commands for adjusted data rate are supported where possible in MVS/System Product, VM/System Product, and many of the supporting products. Specific system functions that are supported include:

- MVS paging and swapping
- VM paging
- JES2 and JES3
- VSAM (including catalog)
- SAM-E
- OSAM in DL/I
- Sort/Merge
- Indexed VTOC
- Media Manager
- Device Support Facilities (DSF)

Specifically not supported are ISAM and QSAM. Any programming that includes search key type channel commands prior to a write operation are not supported because of the difficulty giving prior notification for this type of operation.

With MVS, the I/O supervisor adds a Define Extent command to all read and write command sequences destined for a speed matching buffer. If the channel program contains a Locate Record command, the Define Extent and Locate Record commands replace the Seek Set Sector and Search command sequence.

For existing user channel programs, EXCP prefixes a Locate Record command to write operations if the program begins with a Set Sector command (optional) followed by a command sequence of Search ID Equal, TIC, and a write command. The write command may be Write Count, Key, and Data; Write Key and Data; or Write Data. However, the Locate Record command is not prefixed to the user channel program if any command of the sequence has a program-controlled interrupt (PCI) or suppress-length (SLI) flag bit set.

When constructing a Locate Record command to be prefixed to a user channel program, the count information in the Write command is used 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. 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.

If the user channel program begins with a Locate Record command, only a Define Extent command is prefixed for the write operation.

#### **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 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 Start I/O instructions 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 solve 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 MB/sec channel path offline so that it is truly only a backup path. (This is done at the operator console with the VARY PATH OFFLINE system command.) Under MVS, 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. (Refer to the Parameter Library Member IECIOS in the MVS Initialization and Tuning Guide listed in the bibliography.)

#### **Model AA Attached to 3880 Model 13**

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

When first requested for processing, data 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 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 Model 13 cache is not for paging.

Data also can be specifically placed in the 3880 cache by the programmer and retained there until released by the programmer.

The MVS/SP operating system supports the Model 13 for 3380 attachment.

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

### **Overlapping Operations, Models A and AA**

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 devices controlled by both models A and AA. In addition, command and data transfers can be overlapped with model AA; that is, two transfer operations can be in progress at the same time.

### **Internal Paths and Device Access, Models A and AA**

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 9, access mechanisms 0, 1, 8, and 9 are arranged on path 0.

If the 3380 is a model A, 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 9, 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, 12, or 13.

### **Dynamic Path Selection, Model AA**

If the 3380 is a model AA, 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 10, 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 10, 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 10 on path 1 through the other controller and storage director.

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 via another storage director and the other controller in the model AA.

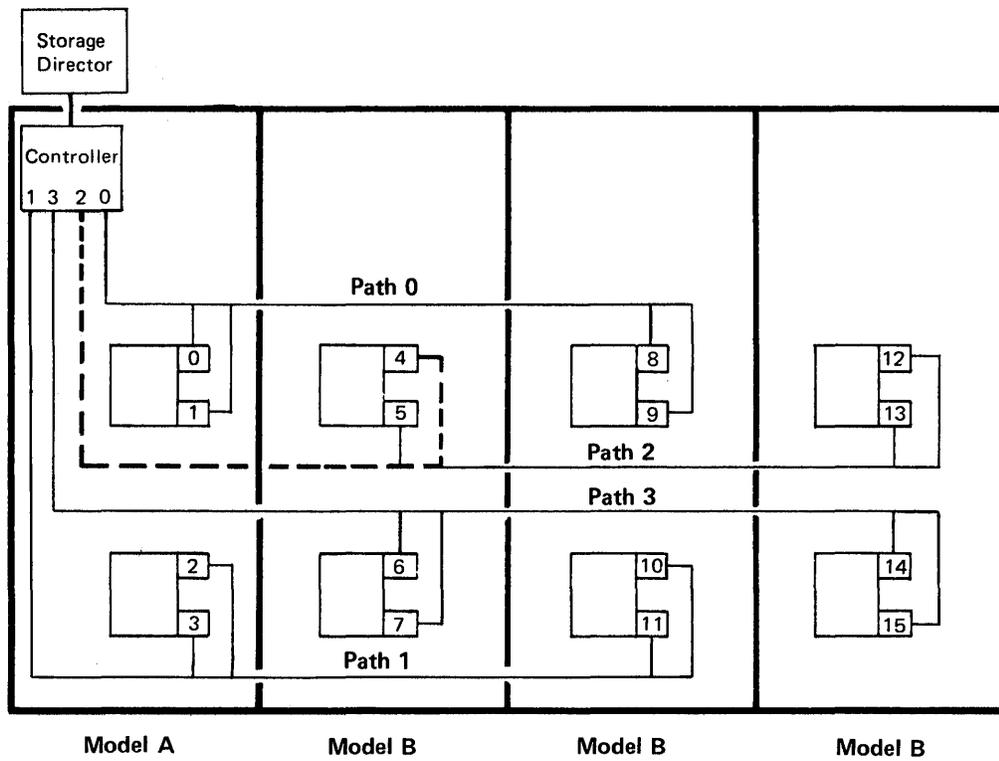


Figure 9. Four Internal Paths in a 3380 Model A with One Controller. The dashed line represents one path that can be active.

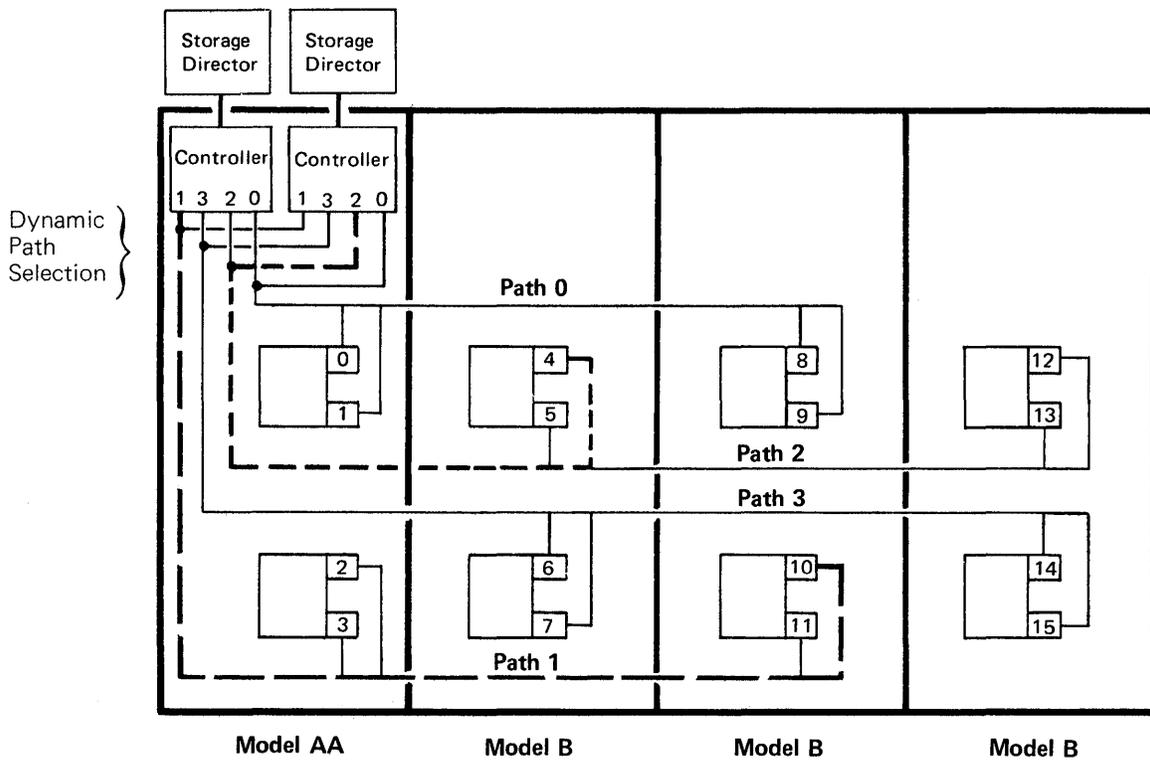


Figure 10. Four Internal Paths and Dynamic Path Selection in a 3380 Model AA with Two Controllers. The two dashed lines represent two of the paths that can be active simultaneously.

## Model AA With VM/SP or ACP/TPF

Both VM/SP and ACP/TPF operating systems support attachment of 3380 AA models. Both storage director/controller paths have access to all the access mechanisms in a string.

With VM/SP, a model AA 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.

## Dynamic Reconnection, Model AA

When connected to a 3081 Processor operating in System/370 Extended Architecture mode, a 3380 model AA can dynamically reconnect to the first available channel. In other situations the reconnection must be on the same, originating channel, instead of the first channel available.

## Device Reserve/Release Changes

The Device Reserve 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. Two functions have been added to improve flexibility of access to a shared volume in multiprocessing environments.

- With MVS/SP Release 2, 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.
- For a 3380 model AA, with full dynamic path selection support, there are new channel commands that reserve the volume for channels having a common path group identification. Channels in the group can access data sets on a reserved volume, but channels not encompassed in the path group identification are excluded from access to the entire volume. (This function is implemented by the system and does not require any specification by the user.) This extension to the device reserve/release function is supported by MVS/SP but not by VM/SP nor ACP/TPF operating systems.

## 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 A and AA units. See Figure 11. 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 obtain its operational status.

### Power On and Off

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

### Emergency Power Off

The Unit Emergency switch on the 3380 model A or AA 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 Indicator

The AC Power Indicator 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

The path to the 3380 is activated with the VARY PATH ONLINE system command and deactivated with the VARY PATH OFFLINE system command.

The data path switch on the 3380 control panel must be in the Enable position when the VARY PATH ONLINE command is issued. All panels have spaces for two Enable/Disable switches; however, if there is only one controller in the 3380 model A attached to one storage director, only one of the two switches is enabled.

Prior to setting the data path switch to the Disable position, a VARY PATH OFFLINE command should be issued. If the path is not offline, a false equipment check error message may be generated.

The VARY ONLINE command makes a volume ready for system use. Successful execution of this command indicates that the device is ready and can read the volume table of contents. A device is placed offline with the VARY OFFLINE command.

### Device Status

In response to a MVS DISPLAY U system console command, a message (IEE 4501) 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.

### Physical Identifier Labels

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

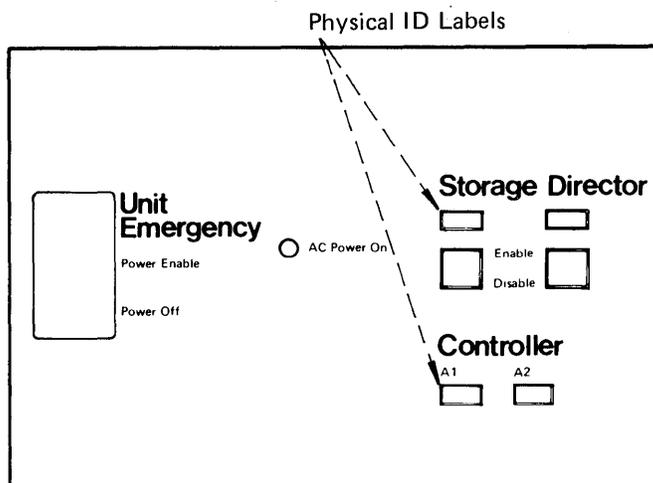


Figure 11. 3380 Control Panel

## 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, or 13. The 3880 Model 3 or 13 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 second storage director for other types of disk storage.

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 rate. The Speed Matching Buffer for 3380 feature cannot be installed in a 3880 Model 13.

Attachment of the 3380 to an IBM System/370 Model 158, 158-3, 168, or 168-3 requires installation of the:

- Speed Matching Buffer for 3380 special feature in the 3880 storage director.

Attachment of the 3380 to an IBM 3081 Processor is through a 3.0 MB/sec channel, with data streaming standard.

Attachment of the 3380 to an IBM 4341 processor channel that has a 3.0 MB/sec data 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 rate.

## 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*, listed in the bibliography.

### 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 only four percent more power than a 3350 unit, but has twice as many access mechanisms and 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 12 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 maximum cumulative length of signal cable that can be used to connect two 3380 controllers to a 3880 storage director is 61 meters (200 feet).

Model	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

\* 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 12. Comparisons of Approximate Floor Space Requirements Based on Full String Configurations

## Addressing

### I/O Address

A two-part I/O address is specified in the Start I/O instruction. The first part is the channel address. The second part is an 8-bit binary number used for 3880 and 3380 addressing, as described as follows and in Figure 13. There is supplementary reference information on addressing in Appendix C.

**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.

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

**Bit 3** is the 3380 controller (models A and AA) address. The controller addresses depend on the number of strings:

- 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 AA, 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 in the model AA, the controller bit addresses in one string must be 0s and the controller bits in the other string must be 1s. (Each string 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 programmer-assigned 3-hexadecimal digit unit address is even, this controller bit 3=0; if the second hexadecimal digit is odd, this controller bit 3=1.*

**Bits 4 through 7** are the access mechanism (device) address. The access mechanism addresses are 0 through 15. For a given string, access mechanism addresses start with 0 and are sequential.

Addresses of controllers and access mechanisms are set by the customer engineer when the 3380 is installed.

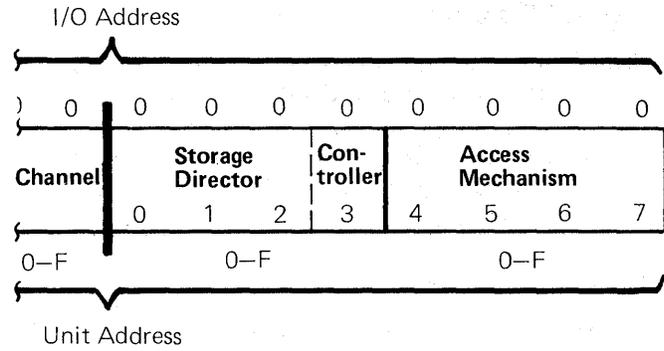


Figure 13. 3880 and 3380 Addressing

### Address Assignments

The I/O address is expressed as a 3-hexadecimal digit *unit address* when it is assigned at system generation. The first of the three digits refers to the channel. For the 3380, the second digit represents the storage director and controller address (bits 0 through 3), and the third digit represents the device (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 second and third digits of the unit address are B0, B1, B2 . . . BF.

Figure 14 lists the sets of 2-hexadecimal digits that are valid for the storage director/controller and device portions of the 3-hexadecimal digit unit address.

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 14. Sets of Valid 3880/3380 Addresses

The table in Figure 14 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 and 10, 11, 12 ... 1F, 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 15 and 16 are examples of address assignments.

If multiple storage directors attach to the same channel, each storage director must have a different address for use with that channel. For example: 8C0 and 8D0. If the same storage director attaches to different channels that are attached to the same processor, the storage director must have the same address for use with all of these channels. For example: 840, 940, and A40.

### 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.

If bits 0 through 2 of the I/O address refer to an unattached storage director, condition code 3 (not operational) is set in the program status word and a No Paths Available system message is issued at the operator console.

If bit 3 of the I/O address refers to a controller that has never been selected since the last 3880 initial microcode load (IML), condition code 3 (not operational) is set in the program status word and a No Paths Available system message is issued at the operator console. Bit 3 may refer to a controller that was previously selected but is now unavailable because of an error, the path switch is off or the power is off. If this occurs, unit check status is set in the channel status word, the equipment check sense bit is set, and a Permanent I/O Error system message is issued at the operator console.

If bits 4 through 7 of the I/O address refer to a device that is not installed, unit check status is set in the channel status word and an Intervention Required system message is issued at the operator console. (This condition also may occur when the customer engineer has set the device to service mode.) If bits 4 through 7 refer to a device that is physically attached but without power, unit check status is set in the channel status word, the equipment check sense bit is set, and a Permanent I/O Error system message is issued at the operator console.

### Seek Address

The seek address, transferred by the Seek channel command, specifies the record, cylinder, and head location 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, unit check status is set in the channel status word, the command reject sense bit is set, and a Permanent I/O Error system message is issued at the operator console.

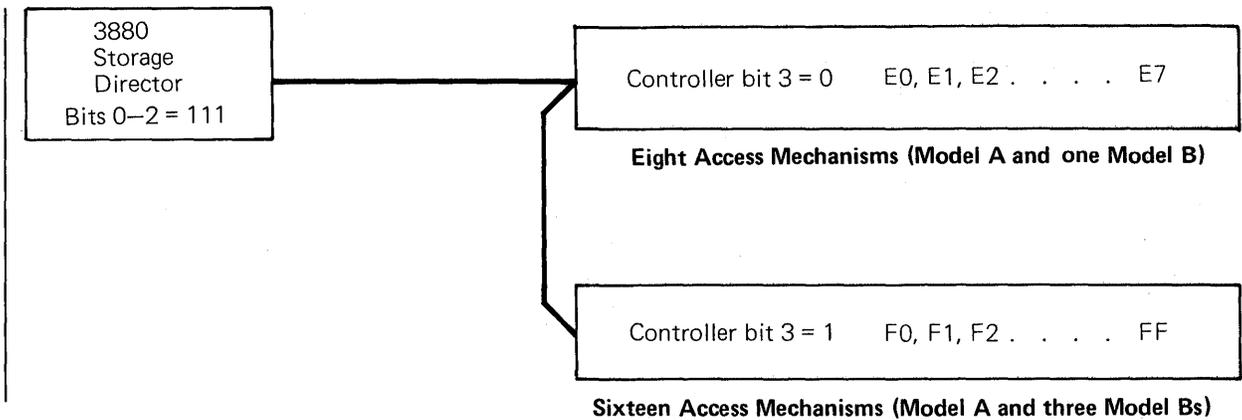


Figure 15. Addressing Example: Two Model A Strings Attached to one 3880 Storage Director

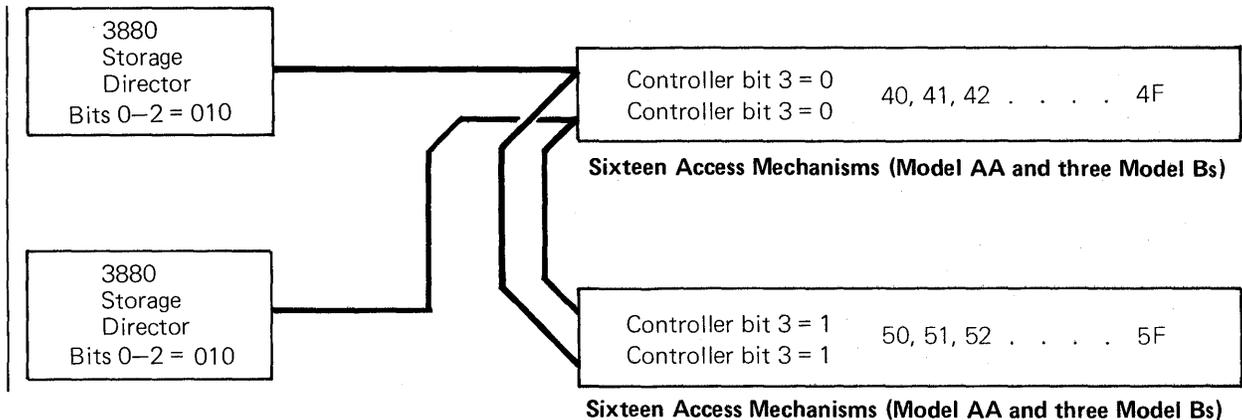


Figure 16. Addressing Example: Two Model AA Strings Attached to the Same Two 3880 Storage Directors

## Program Planning Considerations

The 3380 is supported under the following base system control programs: Operating System/Virtual Storage Version 2 Multiple Virtual Storage (OS/VS2 MVS), Virtual Machine/370, and the Airline Control Program. 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 base 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.

Documents that contain information on system generation and installation of the various program products that support the 3380 are listed in the front of the manual. Also there are program directories included with the program distribution.

### **MVS/SP Support**

The 3380 is supported under OS/VS2 MVS with a combination of system control programs and program products. The programs support all models of the 3380 and all of the processors that can connect 3380s. They support the full functional capabilities of the 3380 including dynamic path selection.

OS/VS2 MVS support is under Release 3.8. In addition to the base system control programs in Release 3.8, support for the 3380 requires the following:

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

### ***System Product and Data Facility Device Support Programs***

The System Product and Data Facility Device Support programs encompass modified system control modules and new functions for support of the 3380. System Product provides system program support and incorporates the earlier MVS system extensions. DFDS provides data management support.

DFDS includes provisions for supporting an index to the volume table of contents (VTOC). The index reduces the time required for searching the VTOC entries. Instead of sequentially searching through all of the information associated with data set entries, the index supplies the name of the data set and location of the entry in the VTOC. Creation of the index provision is optional. It 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. (The Device Support Facilities program, described later in this section, can be used to create an index for a new VTOC when a volume is initialized or to build an index for an existing VTOC.)

An alternate master catalog can be used during an initial program load (IPL) operation.

The location of data sets on a system residence volume can be identified by a special character string, which then may be used in place of the volume serial number and unit identifier in the catalog.

The Direct Access Device Space Management (DADSM) routines have been enhanced with preprocessing and postprocessing installation exits for certain functions.

### ***MVS/SP System Generation***

Users planning MVS support for the 3380 should have MVS Release 3.8 installed. System Product (SP) Release 1 or 2 must be installed prior to installation of Data Facility Device Support. The installation of SP requires at minimum a stage 1 system generation with TYPE = ALL. The 3380 support is provided with installation of Data Facility Device Support.

If the MVS/SP Installation Productivity Option (IPO) Release 3.8F is being used, only the Data Facility Device Support needs to be installed.

Because 3380s are being added to the system, both stage 1 and stage 2 system generations with TYPE = I/O must be done.

If system residence or system libraries are desired on the 3380, the data sets must be copied. Or a second system generation can be performed using the newly created system to perform the system generation with the 3380 as the target device.

The device-related tasks associated with system generation include adding the 3380 and providing addressing information using the appropriate macro instructions. The I/O address (see the section, "Hardware Planning Considerations") is assigned as a 3-hexadecimal digit *unit address* (OS/VS IODEVICE macro ADDRESS parameter). The ADDRESS parameter includes the unit address and the number of sequentially numbered units. A device type (the IBM-assigned generic name, 3380) is included. Information supplied by the IODEVICE macro is used to establish a unit control block (UCB) for each 3380 device (access mechanism).

The UNITNAME macro is used to assign a user-defined esoteric name to a group of devices in which the 3380 may be included.

Because there is no restriction on assigned but unused addresses, it is easiest to include in the system generation all the addresses for 3380 units that are anticipated to be installed.

The system generation process uses the information supplied in the IODEVICE and UNITNAME macros to build internal tables to be used for device allocation.

### **3380 MVS/SP Releases 1 and 3 Uses**

The 3380 can be used for the following storage needs under MVS/SP Releases 1 and 3.

- System residence
- Paging
- JES3 setup and scheduling
- System and user data sets

In addition, spooling and checkpointing are supported under SP 3.

### **Programs Using OS/VS and the 3380**

OS/VS program products using SAM, ISAM, PAM, DAM, or VSAM unmodified and not requiring specifications of device characteristics can operate on the 3380.

### **VM/SP Support**

The 3380 is supported by Virtual Machine/System Product (VM/SP) program product. 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 and all of the processors that can connect 3380s. It does not support the

alternate controller function of dynamic path selection nor use of the new commands associated with device reservation and path group identification. 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 also does not support CMS/VSAM data sets on the 3380.

### **VM/SP System Generation**

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. The ADDRESS operand includes the device address and number of sequentially numbered device addresses. DEVTYPE = 3380 must be specified.

### **3380 VM/SP Uses**

The 3380 can be used for the following storage needs under 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.

- Remote Spooling Communications Subsystem (RSCS) system residence.
- Dedicated device attached to a virtual machine.

VM/SP supports the 3380 for CP paging and spooling. Both paging and spooling areas can reside on the same volume. However, performance characteristics should be carefully considered before designating a 3380 as a paging or 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."

## **Airline Control Program/Transaction Processing Facility Support**

The 3380 is supported by the Airline Control Program/Transaction Processing Facility (ACP/TPF) for system residence and system and user data sets. ACP is a standalone operating system and requires no optional programs to support the 3380.

### ***ACP/TPF System Generation***

The System Initialization Program (SIP) is used to generate an ACP system. ACP/TPF support of 3380 is established during the SIP process. Existing disk storage support macros in the SIP process will be modified to accept 3380 device specifications as valid ACP/TPF devices. 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 3380 Restrictions***

The following restrictions apply to ACP with regard to the 3380 support:

- Dynamic path selection, and extensions to the 3380 command set are not supported.
- The 3380 is supported by ACP/TPF only in conjunction with data streaming.
- Up to four types of mixed disk storage can be included in an ACP/TPF system.

### **Additional Programming Support**

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

### ***Data Facility Extended Function***

The Data Facility Extended Function (DFEF) program is a catalog facility that can be used under OS/VS2 (MVS) with the SP and DFDS installed. Data Facility Extended Function has a new architecture that enhances reliability, recoverability, and performance. It is an optional replacement for current VSAM master, user, or OS catalogs. DFEF catalogs, VSAM catalogs, and OS catalogs can coexist. The new catalog supports both VSAM and non-VSAM data sets. It uses the Access Method Services (IDCAMS) command syntax.

## ***Data Facility Data Set Services***

The Data Facility Data Set Services program provides dump and restore functions for the 3380. Data Set Services has both system and standalone versions and supports all current IBM storage products except the 2314 and 2319 disk storage types. 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 VSAM, 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.

### ***Device Support Facilities***

The Device Support Facilities (DSF) 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.
- Skip up to seven defective areas on a track and assign alternate tracks, 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.

### ***IBM Direct Access Storage Device Migration Aid***

The IBM Direct Access Storage Device Migration Aid can help in converting data sets from other direct access storage devices to the 3380. It operates with OS/VS and supports all data sets created by the standard access methods. It 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.

### ***Hierarchical Storage Manager***

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.

### ***Sort/Merge***

The OS/VS Sort/Merge program product Release 4 provides support for the 3380 for use as SORTIN, SORTOUT, and SORT/WORK devices.

### ***Environmental Recording, Editing, and Printing***

A new release of the Environmental Recording, Editing, and Printing (EREP) program supports editing and printing of 3380 error records. The new release is a complete replacement of previous revisions.

## Volume Initialization and VTOC Index

The 3380 volume is initialized before shipment with a home address (HA) and a track descriptor in record zero (R0) on all tracks. The home address contains the track address (CCHH) and a flag indicator if the track has been found defective. The record zero of a defective primary track contains the address (CCHH) of an assigned alternate, and record zero of the assigned alternate contains the address of the associated primary track. The 3380 is shipped from the plant of manufacture with no tracks flagged as defective.

At installation the volume must be initialized for use by the operating system. The Device Support Facilities (DSF) program is provided for this initialization. See Figure 17.

With the DSF INIT command, a 3380 volume can be initialized to one of three levels. Higher levels provide additional functions. The following summary highlights the differences in the three levels.

All levels of initialization include:

- Writing initial program load (IPL) bootstrap records as records one and two, and a volume label as record three on cylinder 0, track 0.
- Specifying the starting location and the size of the volume table of contents. The default in OS/VS for the location is cylinder 0, track 1, and the default for the size is one track.

The second 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 address and record zero, which causes the rest of the track to be erased. This level is required if for any reason a head and disk assembly is replaced. The Validate parameter should be used to initialize both volumes of the head and disk assembly by writing a home address and record zero on all tracks.

The third level of initialization adds surface checking for the condition of all tracks. With this check, special patterns are written and read. 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, DSF automatically bypasses the defect. Up to seven defective areas are skipped, after which an alternate track is assigned. If tracks that were previously flagged defective are found to be usable, they can be reclaimed. Data checks are discussed later in the section, "Device Status and Error Handling."

An index to the volume table of contents also can be created with the 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 DSF BUILDIX command changes an existing OS format VTOC to the indexed format VTOC.)

Cylinder 0:0 Track

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

Cylinder 0:1 Track

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

Cylinder 443:4 Track

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 a 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, track 0, records 1, 2, and 3.
4. The VTOC address (CCHH) is specified with Device Support Facilities INIT command. Track 1 is the OS/VS default for the starting location of the VTOC.
5. An index to the VTOC is specified with the Device Support Facilities INIT command. There is no location default.

Figure 17. Volume Layout of Initialization Information

## Data Management and Data Set Conversion

This section discusses data management and conversion for users of MVS. Both VM and ACP/TPF have special data structures and management techniques.

All of the access methods that are supported by the operating system can be used for data sets that are stored on the 3380. The IBM-supplied access methods are:

- Sequential Access Method (SAM),
  - Basic and Queued (BSAM, QSAM)
- Indexed Sequential Access Method (ISAM),
  - Basic and Queued (BISAM, QISAM)
- Partitioned Access Method, Basic (BPAM)
- Direct Access Method, Basic (BDAM)
- Virtual Storage Access Method (VSAM)

The Sequential Access Method - Extended (SAM-E) is a program product that can be used in place of SAM. It is briefly described in Appendix B.

The access method routines code the channel programs, construct the control blocks necessary for execution of the channel programs, and issue the execute channel program (EXCP) macro instruction. The user supplies information for the data control block (DCB) with the system DCB macro. Instead of using the available access methods, the user can perform the functions provided by the access methods and issue the EXCP macro instruction directly. (VSAM uses access method control blocks and issues SIO instructions.)

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

### Job Control Language Uses

Job Control Language (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 I/O Addressing in the section, "Hardware Planning Considerations" and System Generation in the section, "Program Planning Considerations.") The UNIT parameter may specify:

- Unit Address: Three-character address representing a channel, control unit, and unit number, such as UNIT = 16E
- Device Type: IBM-defined generic name, such as UNIT = 3380
- Group Name: User-defined esoteric name, such as UNIT = DISK. At system generation, this name is equated by the system generation UNITNAME macro to a list of device addresses.

UNIT parameters specifying a unit address 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 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.

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. 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.

The IBM Direct Access Storage Migration Aid program product reduces the manual effort involved in making JCL changes. This program can be used to generate customized JCL, identify data set JCL statements in procedure libraries affected by the conversion, and create a jobstream to update the JCL statements.

## Data Set Conversion Under MVS/SP

The programming changes required to convert existing data sets to the 3380 are primarily those that reflect changes in the 3380 device characteristics, such as track capacity and number of tracks per cylinder. The IBM Direct Access Storage Device Migration Aid program can be used to identify the JCL changes needed and generate the required statements, as described in the previous section.

Data sets can be copied, or moved, to the 3380 using IBM-supplied utilities. The utility used depends on the characteristics of the data set and the functions desired. The *OS/VS2 MVS Utilities* manual, listed in the bibliography, describes the various utilities and gives information on the control specifications needed.

Data sets can be converted directly from existing disk storage, such as the 3350, to the 3380. Or, in many cases, they can be unloaded in a sequential organization to tape (for instance, to create a backup copy or to hold a copy when an existing device is removed before the new disk storage is installed). When the unloaded data sets are loaded to the 3380 from tape, the data is reconstructed, if necessary, to its original organization.

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.

The migration aid program can be used to generate the necessary utility statements.

### *Standard Access Methods*

If the data organization is not changed, the following procedures can be used to convert data sets from an existing device to the 3380.

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

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

ISAM data sets can be converted to the 3380 using the IEBISAM utility. Probably ISAM data sets will not require new programs because the user's existing reorganization programs can be used. Minor changes to the load programs will be required because of the 3380 device characteristics. Changes such as space allocations can be made with JCL data definition statements.

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 the 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 copied using the Access Method Services REPRO command.

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

Besides the utilities described above, the Hierarchical Storage Manager program can be used as a conversion tool. Data sets with a physical sequential organization and fixed-block, variable-block or variable-block-spanned format can be reblocked automatically when recalled to a 3380 from a migrated location.

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 or SIO instruction 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* manual listed in the bibliography.

## 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. The VM/SP DASD dump/restore utility can be used to move data on a cylinder basis.

## 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.

While performance needs may be the primary criterion for placing data sets on the early 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.

Besides taking advantage of the physical characteristics of the device, certain data sets may be placed on the 3380 to take advantage of programming enhancements. For example, the indexed VTOC and the new catalog facility should prove especially beneficial for data sets used for interactive programs, such as the time sharing option (TSO) where access traffic is heavy.

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 deblocked 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/V5 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 7.) 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 18 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. However, the largest block size supported by the IBM-supplied access methods is 32,760 bytes. 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 3,859 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 18. 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 relative track address.

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*

The IBM Direct Access Storage Device Migration Aid 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.

## Performance Considerations

Performance of disk storage operations is affected by factors such as 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 AA with two controllers, placement of the data sets relative to the internal paths should be taken into consideration.
- When a 3380 storage director has a speed matching buffer and a channel with a 1.5 MB/sec transfer rate, performance benefits are obtained by using the new Define Extent and Locate 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.
  - Access multiple blocks for each access by specifying chained scheduling for sequential data sets (OPTCD=C). Chained scheduling provides for a succession of read or write operations by synchronizing I/O operations and use of buffer space. There are certain restrictions on its use. With MVS, chained scheduling is used wherever possible whether or not it is requested. (Chained scheduling in SAM-E is described in Appendix B.)

When using chained scheduling, be sure the number of chained buffers multiplied by the block size can be accommodated by main storage.

2. Minimize arm movement when accessing.
  - Place the most active permanent data sets near the middle cylinder (cylinder 443). In general, this placement reduces the maximum length of seeks outside the data set. If absolute addressing is not used, however, MVS 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 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 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 accelerates operations by reading the requested record and the count field for the second record. This is instead of searching the ID of the previous record to read the requested record as normally done. For fixed standard records (FS or FBS), the access methods use a form of search direct technique. Search direct cannot be used when chained scheduling is used, and chained scheduling takes priority. (Search Direct in SAM-E is described in Appendix B.)

- Use a smaller block size for data sets where direct access predominates.
  - Use a 3.0 MB/sec channel, where possible.
  - Use commands for adjusted data rate when using a 1.5 MB/sec channel.
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 also can be implemented.
- Use the search direct option or fixed-length records, standard format, as described in item 3.
  - 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 Start I/O (SIO) 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, but 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, unless a fixed-head device is being used.

The PLPA may be structured, through the PAKLST option, so that the most active pages are on the front of the PLPA data set. Then these most active pages can be allocated under the fixed heads.

- 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 reasonable 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 the combination of status conditions sent to the channel. Depending on the other conditions in the combination, the channel initiates procedures to correct the error internally 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. Error recovery procedures (ERPs) are provided with the IBM operating systems.

Further explanation of the cause of a unit check status is obtained 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 device 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* manual.

The 3880 Storage Control manual also includes 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 automatically implemented by the system error recovery procedures.

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

The system control program makes status and sense information available to the problem program through error analysis routines that the user specifies through operating system macro instructions (MVS SYNAD routine). 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.

Many errors can be corrected by the 3880 storage director with internal command retry procedures or by the system with error recovery procedures. In the event that neither of these procedures is possible or successful, a system message is issued to bring the problem to the attention of the operator. The system messages include status and certain sense information.

A Permanent I/O Error system message may indicate a problem that requires a customer engineer, but in some cases, action can be taken to prepare for service or to handle the problem without requiring a service call.

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

### Control and Function 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.

Programming errors such as invalid track format and no record found cause a command rejection, and are passed to the user program for recovery action. An error message is issued that the job has been terminated.

### 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. In many cases, even though the first attempt to read the data is not successful, with subsequent attempts it can be read correctly.

### Correction of Data Errors

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 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 19.

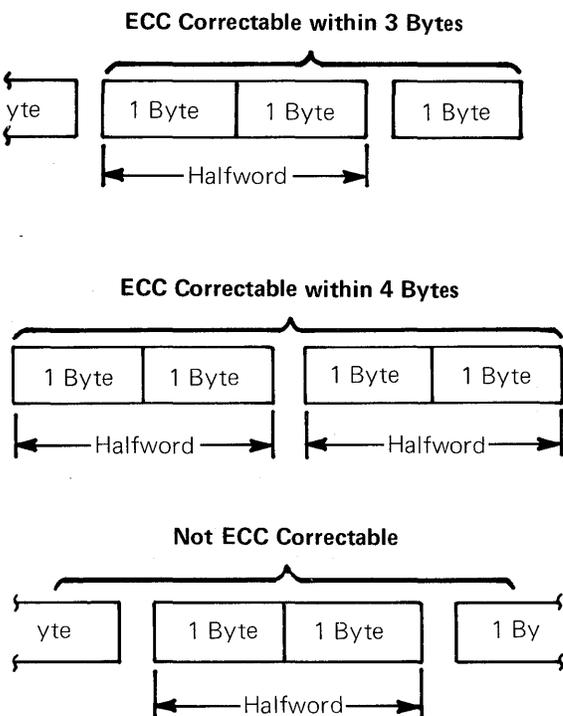


Figure 19. Data Error Correction

Errors that can be corrected with the error 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.

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.)

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. A Permanent I/O Error message is issued at the operator console with a description of the data check.

### **Response to Data Errors**

The cause of data checks may need to be investigated if the errors are permanent, or if they are temporary yet suggest a possible problem (for example, if errors occur repeatedly).

There are several sources of information regarding the nature of a data check and its location. The most immediate source is the system console message. In addition, the 3380 EREP has two reports that provide information related to data checks, the event history report and the outboard record (OBR) edit report. The basic source of detailed information regarding all errors is the sense information. Sense information pertaining to a particular error can be obtained from console messages or the two EREP reports. Both the console message and EREP reports contain the cylinder and head numbers at which the data check occurred.

The Device Support Facilities (DSF) program is provided as a tool to aid in problem determination and resolution of data check problems. The following information gives general guidelines, but for a complete description of use of the DSF program, refer to the Device Support Facilities manual listed in the bibliography.

#### **Hardware and Data Verification Tests**

Sometimes it is not evident whether a read or write problem is caused by a disk surface problem or an equipment problem (for example with the access arm). The ANALYZE command of the DSF can be used to help determine the cause. Two types of tests can be invoked by the operator. The results are reported in diagnostic messages.

One type of tests exercises the various hardware components by trying Seek, Read and Write (in the CE cylinder) commands. If a hardware problem is indicated, a customer engineer should be called. The diagnostic output should be saved for use by the customer engineer.

If no problem is indicated by the hardware tests, data verification tests can be performed to test the readability of data. The user's records are read from the cylinders specified or from all of the tracks on the volume. (The data is not transferred to main storage.) If a data check occurs during the data verification test, the device address and cylinder and head numbers are included in the diagnostic message.

#### **Surface Checking**

The INSPECT command of DSF provides for surface checking of specific tracks suspected of having a problem.

With this check, special patterns are written and read. When specific tracks are specified, DSF uses a set of advanced patterns to thoroughly analyze the condition of the specified tracks. The Preserve parameter is the default and causes the data that was on the tracks to be saved. However, if the No Preserve option is specified, the chances for successful track inspection are improved.

It also is possible to perform surface checking of all of the tracks on the volume by specifying INSPECT with the All Tracks and Check parameters. The Preserve parameter is the default and is required to inspect all tracks. This specification provides a sifting function as with the INIT command (third level). 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 the advanced, and more time-consuming, set of patterns.

#### **Bypassing Defects**

If defects are confirmed while surface checking with the set of advanced patterns, the 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 displacements 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.

If tracks that were previously flagged defective are found to be usable, they are reclaimed, if so specified in the INSPECT command.

#### **Data Checks in Home Address and Record Zero**

Special attention needs to be given to data checks that may occur in the home address or record zero of a track. A data check in the home address area can be determined from sense information or from diagnostic messages generated as a result of execution of the DSF ANALYZE command.

If a data check occurs in the home address of a single track, issue the INSPECT command to surface check the specific track. Both INSPECT and INIT surface checking (advanced patterns) handle data checks in the home address automatically. Attempts are made to read

enough of the home address area to rewrite it in the same location or to a skip displacement location. If the home address cannot be reconstructed, a message is issued at the operator console, and a customer engineer should be called.

If data checks occur in the home address of more than one track, there may be a potential problem reading home addresses of other tracks also. The problem can be corrected by reinitializing all of the home addresses. The data from the entire volume should be dumped to a backup device and the INIT command with the Validate parameter (second level of initialization) used to rewrite the home address and record zero of all tracks. The data that was dumped can then be restored from the backup.

### **MVS IEHATLAS**

An operating system utility (IEHATLAS) also can be used in certain cases to assign alternate tracks and transfer data from an error track to an alternate track.

### **Backup Data**

Data on a 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, 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 3330 and 3350 disk storage. If data checks occur in the home address or record zero of any track, the data from the

entire volume should be dumped to a backup device and the DSF INIT command with the validate option used to rewrite the HA and R0 of all tracks. The data can then be restored for the backup.

### **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. The problem may be intermittent or it may be continuous.

If the error is intermittent, system error recovery procedures may be able to correct the problem by retrying the operation. With the 3380, the system retries the operation up to 10 times. If the error is not intermittent, or if the system error recovery procedures are not successful, a Permanent I/O Error message is issued describing an equipment check. A customer engineer should be called.

*Note: Information about an equipment check may suggest a possible data error, as described in "Response to Data Errors."*

### **Service Procedures**

The customer engineer (CE) 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 AAs 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 AA 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 megabyte 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

New identifiers in 3380 sense information pinpoint the location of a problem more precisely than in the past. 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 3-hexadecimal digit 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 and 3380 controllers are each assigned a 1-byte (2-hexadecimal characters) physical ID.

The storage director identifiers must be assigned in pairs in each 3880 Storage Control. The two numbers must be consecutive, with the even number having the smaller value. For example: 02 and 03, 30 and 31 are valid pairs; 05 and 06, 23 and 24 are not valid pairs because they begin with odd numbers. Numbering should not begin with 00 and 01.

In 3380s with only one controller (model A), only one physical ID is needed, but in models with two controllers (model AA), two physical IDs are assigned, one for each controller. As with the storage directors, the two numbers must be consecutive, with the even numbers having the smaller value. Physical IDs for controller 0 must be between 02 and 7F; those for controller 1 must be between 80 and FD.

In addition to the assigned storage director and controller physical IDs, the device also has a physical ID. The device physical ID is the same as the device address. The first of the two hexadecimal digits of the physical ID is always zero.

Physical IDs are illustrated in Figure 20.

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 A and AA 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.

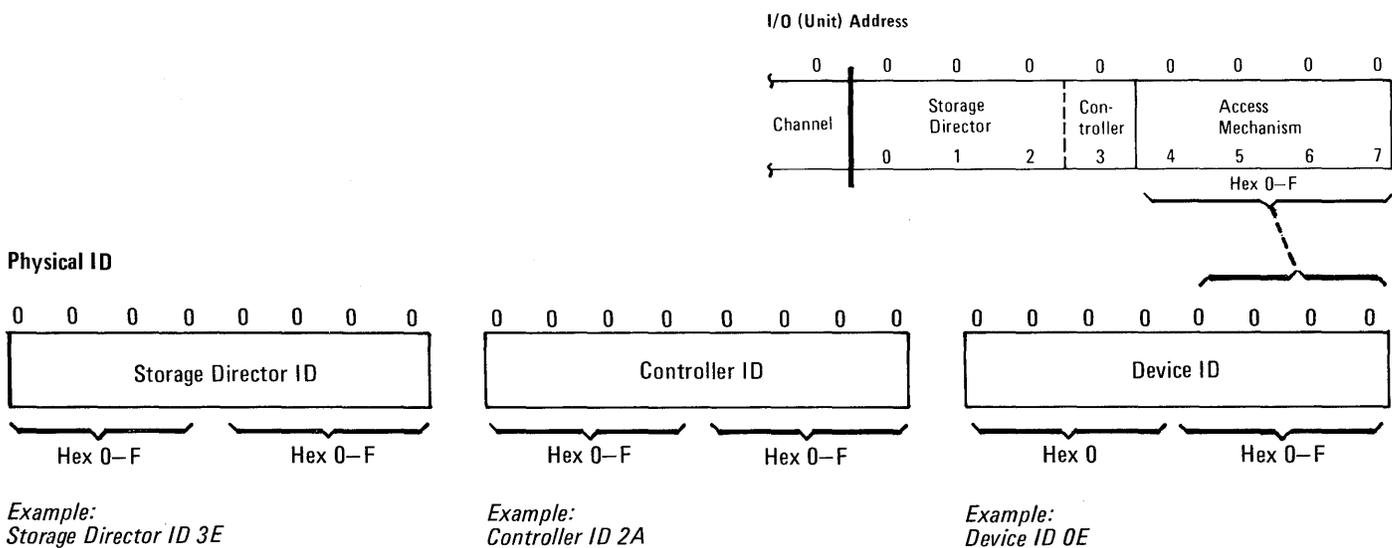


Figure 20. Physical ID Assignment

## 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.

The Data Facility Data Set Services program, described in the section, "Program Planning Consideration," 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. Once they have been dumped, data sets can be deleted.

- 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.

All VSAM data sets can be dumped and restored as full or partial volumes with the Data Set Services program. Selected data sets can be dumped and restored using the EXPORT and IMPORT commands of Access Method Services. If the new catalog facility, DFEF, is installed, volumes containing VSAM catalogs or data sets can be dumped and restored without introducing catalog/data set synchronization problems.

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

	<b>3420 Model 8</b>	<b>3420 Model 6</b>
Data Recording Density	6250 bpi	6250 bpi
Data Transfer Rate	1.25 MB/sec	0.78 MB/sec
3420 Tape Reels per 3380 Volume <sup>1</sup>	2	2
Approximate time elapsed <sup>2</sup>	5 min	6.6 min

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.

**Figure 21. 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 probably 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.

- Data processing (DP) 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 needed publications through an IBM branch office.

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: Checkout of 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 and Type II and Type III programs currently in use 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 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 to initialize volumes.

Examine performance considerations.

- Consider the performance guidelines as described in “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/1)? 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, such as the fixed-length standard format, search direct option, and chained scheduling for sequential data sets.

- 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 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 indicator on the operator panel. Procedures should be implemented to ascertain device status from the system console.

## Appendix A: Space Calculation

Space calculations are determined for physical records. The data length (DL) of a physical record is the block size. With MVS, 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 (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

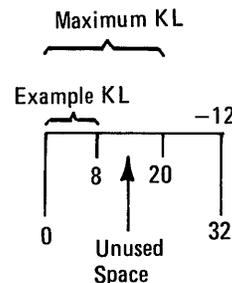
$$K^\dagger = \frac{KL + 12}{32} \text{ if } KL \neq 0$$

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

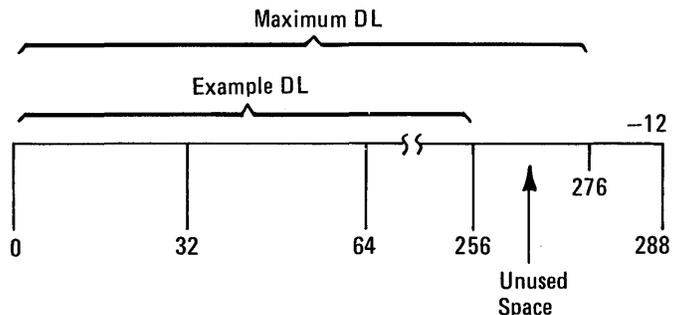
† These values are rounded up to the next integer. KL and DL each must have 12 added and each must be rounded up to a multiple of 32. The 12 bytes are for error checking. For example, with a partitioned data set directory where KL = 8 and DL = 256: 12 bytes must be added to the 8, and 12 bytes must be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

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  $32 \times 1 = 32$  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  $32 \times 9 = 288$  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 with a standard track descriptor record zero (R0).

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

When the track descriptor record (R0) is not standard, the following formula applies:

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

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

Capacity calculations for the 3380 can be performed by using the TRKCALC macro in OS/VS. 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 as follows:
  - 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 MVS Release 3.8 with PTFs UZ90059 and UZ28937. It is described in *OS/VS2 System Programming Library; Data Management*, Order No. GC26-3830.

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).

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

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 = (Key Length + 12) rounded up to a multiple of 32 if KL  $\neq$  0
  - D = (Data Length + 12) rounded up to a multiple of 32.

## Tables, Without and With Keys

The following tables give the numbers of equal length records that can be placed on a track and cylinder. One table gives the number of records if there are no keys, and the other table gives the number of records if there are keys.

The number of records that can be placed on a track depends on the data length and whether or not there are keys. A range of data lengths results in a given number of records. For example, for records without keys, data lengths from 7,477 to 9,076 bytes allow 5 records per track, 75 records per cylinder. The tables show the maximum data lengths that can be used for the given number of records.

The key lengths (KL) and data lengths (DL) in the tables are the number of bytes available. The table without keys gives the data length available after overhead and 12 bytes for error checking have been subtracted. When

there are keys, there is additional overhead and another 12 bytes for error checking (236 bytes). The table with keys gives the number of bytes available for the the key and data lengths (KL + DL) after both data and key overhead and error checking bytes have been subtracted. Key and data areas are assumed to occupy the full 32-byte increments. If less space is specified by the programmer, the storage control pads to the full 32 bytes.

These two tables also include the total number of bytes used on a track and cylinder for the given number of records.

To determine the data length when there is a given key length, subtract the key length (including any padding to round up to a 32-byte increment) from total KL and DL bytes available as given in the table with keys. The next section provides these data lengths for ranges of key lengths.

### Equal Length Physical Records Without Keys

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
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

DL Bytes (Max)	Track Capacity		Cylinder Capacity	
	Records	Bytes	Records	Bytes
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

**| Equal Length Physical Records With Keys\***

KL+DL Bytes (Max)	Track Capacity		Cylinder Capacity	
	Records	Bytes	Records	Bytes
47,240	1	47,240	15	708,600
23,240	2	46,480	30	697,200
15,240	3	45,720	45	685,800
11,240	4	44,960	60	674,400
8,840	5	44,200	75	663,000
7,240	6	43,440	90	651,600
6,120	7	42,840	105	642,600
5,256	8	42,048	120	630,720
4,584	9	41,256	135	618,840
4,040	10	40,400	150	606,000
3,624	11	39,864	165	597,960
3,240	12	38,880	180	583,200
2,952	13	38,376	195	575,640
2,696	14	37,744	210	566,160
2,440	15	36,600	225	549,000
2,248	16	35,968	240	539,520
2,088	17	35,496	255	532,440
1,928	18	34,704	270	520,560
1,768	19	33,592	285	503,880
1,640	20	32,800	300	492,000
1,544	21	32,424	315	486,360
1,448	22	31,856	330	477,840
1,352	23	31,096	345	466,440
1,256	24	30,144	360	452,160
1,160	25	29,000	375	435,000
1,096	26	28,496	390	427,440
1,032	27	27,864	405	417,960
968	28	27,104	420	406,560
904	29	26,216	435	393,240
840	30	25,200	450	378,000
808	31	25,048	465	375,720
744	32	23,808	480	357,120
712	33	23,496	495	352,440
680	34	23,120	510	346,800
616	35	21,560	525	323,400
584	36	21,024	540	315,360
552	37	20,424	555	306,360

KL+DL Bytes (Max)	Track Capacity		Cylinder Capacity	
	Records	Bytes	Records	Bytes
520	38	19,760	570	296,400
488	39	19,032	585	285,480
456	40	18,240	600	273,600
424	41	17,384	615	260,760
392	42	16,464	630	246,960
360	44	15,840	660	237,600
328	45	14,760	675	221,400
296	46	13,616	690	204,240
264	48	12,672	720	190,080
232	49	11,368	735	170,520
200	51	10,200	765	153,000
168	53	8,904	795	133,560
136	55	7,480	825	112,200
104	57	5,928	855	88,920
72	59	4,248	885	63,720
40	62	2,480	930	37,200

\* KL and DL both are rounded up to multiples of 32— 12 as described earlier.

### Tables, Ranges of Key Lengths

The following table gives the data length (DL) available and the number of equal length records that can be placed on a track for records with given key lengths.

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

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				

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				

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				

## Appendix B: SAM-E Program Product

Sequential Access Method - Extended (SAM-E) is a program product designed to replace the standard SAM. It provides BPAM, BSAM, and QSAM with the capability of using the EXCPVR macro instruction instead of EXCP for most direct access operations. It saves processor time by providing, as a default, a technique similar to chained scheduling of channel commands.

In addition, SAM-E uses the search-direct technique for practically all disk accesses and, thus, performs the input/output operation for the best performance possible.

For further information about SAM-E, refer to *OS/VS2 MVS Sequential Access Method - Extended (SAM-E) General Information Manual*, Order No. GH20-9122.



## Abbreviations

<b>ABSTR</b>	absolute track	<b>KEYLEN</b>	key length (DCB parameter)
<b>ATLAS</b>	An OS/VS macro	<b>KL</b>	key length
<b>BDAM</b>	basic direct access method	<b>LRECL</b>	logical record length (DCB parameter)
<b>BISAM</b>	basic indexed sequential access method	<b>MB</b>	megabytes
<b>BLKSIZE</b>	block size (DCB parameter)	<b>ms</b>	milliseconds
<b>BPAM</b>	basic partitioned access method	<b>MSS</b>	Mass Storage System
<b>BSAM</b>	basic sequential access method	<b>MVS</b>	Multiple Virtual Storage
<b>CCW</b>	channel command word	<b>OPTCD</b>	optional services code (DCB parameter)
<b>CMS</b>	conversational monitor system	<b>OS/VS</b>	operating system/virtual storage
<b>CP</b>	command processor	<b>PCI</b>	program-controlled interruption
<b>CSW</b>	channel status word	<b>PDS</b>	partitioned data set
<b>DCB</b>	data control block	<b>PLPA</b>	Pageable Length Pack Area
<b>DD</b>	data definition	<b>PROCLIB</b>	Procedures Library
<b>DISP</b>	data set disposition (DCB parameter)	<b>PS</b>	physical sequential
<b>DL</b>	data length	<b>QISAM</b>	queued indexed sequential access method
<b>DLIB</b>	distribution library	<b>QSAM</b>	queued sequential access methods
<b>DSF</b>	Device Support Facilities	<b>R</b>	record
<b>ECC</b>	error correction code	<b>RECFM</b>	record format (DCB parameter)
<b>EREP</b>	environmental recording, editing, and printing program	<b>RPS</b>	rotational position sensing
<b>ERP</b>	error recovery procedures	<b>SIO</b>	Start Input/Output
<b>EXCP</b>	execute channel program	<b>SP</b>	System Product
<b>F</b>	fixed-length records	<b>SYNAD</b>	synchronous error routine address (DCB parameter)
<b>FB</b>	fixed-length, blocked records	<b>SYSIN</b>	system input stream
<b>FBS</b>	fixed-length, blocked, standard records	<b>SYSOUT</b>	system output stream
<b>FS</b>	fixed-length, standard records	<b>TIC</b>	transfer in channel
<b>GB</b>	gigabytes	<b>TSO</b>	time sharing option
<b>HA</b>	home address	<b>VM</b>	Virtual Memory
<b>HSM</b>	Hierarchical Storage Manager	<b>VSAM</b>	virtual storage access method
<b>ID</b>	identification	<b>VTOC</b>	volume table of contents
<b>I/O</b>	input/output		
<b>IPL</b>	initial program load		
<b>ISAM</b>	indexed sequential access method		
<b>JCL</b>	job control language		
<b>JES</b>	job entry system		

# Index

## A

access mechanisms  
  address 26  
  capacity 9  
  data placement 38, 39  
  internal paths 18  
  number of 8, 9  
  related to volume iii, 8

access methods  
  block size considerations 37  
  data set conversion 34  
  3380 support 33

Access Method Services (IDCAMS)  
  Data Facility Data Set Services Commands 29  
  EXPORT command 47  
  IMPORT command 47  
  REPRO command 34  
  VSAM 33

actuator (*see* access mechanisms)

address, absolute 33, 38

address, controller  
  assignment 24-26  
  JCL change 33  
  not selectable 25  
  odd or even 24  
  system generation 28

address, cylinder  
  home address 10  
  number 25  
  record 10  
  seek operations 12  
  sense information 41

address, device  
  assignment 24-26  
  JCL change 33  
  not selectable 25  
  physical ID 46  
  sense information 41  
  system generation 28

address, head  
  home address 10  
  number 25  
  record 10  
  seek operation 12  
  sense information 41

address, home  
  cylinder and head 10  
  data check 43  
  initialization 31

address, I/O  
  assignment 24-26  
  JCL change 33  
  not selectable 25  
  sense information 41  
  system generation 28

address, record  
  command 12  
  count area 10  
  cylinder and head 10

address, seek  
  command 12  
  cylinder and head 25  
  not valid 25

address, storage director  
  assignment 24-26  
  JCL change 33  
  sense information 41  
  system generation 28

address, unit  
  (*see* address, I/O)

Airline Control Program 20, 29

alternate  
  controller 18-20, 44  
  path 18-20, 44  
  track assignment 9, 31, 44

attachment 1, 3-7, 22

## B

backup  
  frequency 47  
  program support 47

block  
  access method control 33  
  data control 33  
  physical record 11, 35-37

block size  
  access method maximum 36  
  blocking aid 37  
  considerations 35-37, 39  
  fixed length 35  
  fixed standard 39  
  main storage consideration 38, 43  
  system data sets 37  
  tables 54-59  
  track capacity 36, 46  
  3380 optimum 37

buffer  
  size, main storage 37  
  speed matching 6, 16

## C

- cable length 23
- capacity
  - calculations 51
  - comparisons 9
  - cylinder 9
  - track 9
- cache 17
- catalog 29, 47
- chain
  - channel commands 12
  - data 12
  - scheduling 38, 60
- channel
  - attachment 6, 22
  - data rate 6, 16
  - program 12-15
  - switching 6
- channel commands
  - Define Extent 14
  - Device Release 14
  - Device Reserve 14
  - Diagnostic Read Home Address 14
  - Diagnostic Sense/Read 14
  - Diagnostic Write Home Address 14
  - Erase 13
  - Locate Record 14
  - No-Operation 12
  - Read and Reset Buffered Log 13
  - Read Count 13
  - Read Count, Key, and Data 13
  - Read Data 13
  - Read Home Address 13
  - Read Initial Program Load 13
  - Read Key and Data 13
  - Read Multiple Count, Key, and Data 13
  - Read Record 0 13
  - Read Sector 13
  - Recalibrate 12
  - Search Home Address Equal 13
  - Search Identifier Equal 13
  - Search Identifier Equal or High 13
  - Search Identifier High 13
  - Search Key Equal 13
  - Search Key Equal or High 13
  - Search Key High 13
  - Seek 12
  - Seek Cylinder 12
  - Seek Head 12

- Sense 14
- Sense ID 14
- Sense Path Group ID 14
- Set File Mask 12
- Set Path Group ID 12
- Set Sector 12
- Space Count 12
- Transfer-In-Channel 13
- Unconditional Reserve 14
- Write Count, Key, and Data 13
- Write Count, Key, and Data Next Track 15
- Write Data 13
- Write Home Address 13
- Write Key and Data 13
- Write Record Zero 13
- Write Update Data 15
- Write Update Key and Data 15
- channel commands for speed matching buffer 14-17
- channel program 12
- commands
  - Access Method Services 29, 33
  - channel iii, 12-17
  - Device Support Facilities 31, 43
  - Hierarchical Storage Manager 30
  - system console 21
  - VM/SP 34
- command retry
  - error handling 42, 44
  - standard feature 15
- control panel 21
- controllers, 3380
  - alternate 18, 20
  - functions 3
  - internal paths 18
  - models A and AA 3, 18-20
  - path activation 21
- conversion, data sets 34

## D

- data
  - length 35, 51-59
  - record area 10
  - write/read commands 13-15
- data chaining 14
- data control block
  - JCL 33, 35
  - macro 33, 35
- data conversion
  - by access method 34
  - how to 33, 34
- data definition (DD) 33
- data errors 41-44

- Data Facility Data Set Services
    - backup/recovery 47
    - description 29
  - Data Facility Extended Function
    - description 29
    - dump restore advantage 47
  - data length
    - blocksize 33
    - effective use 35
    - space calculations 51-59
    - where specified 10, 11
  - data management 33
  - data rate 9, 16
  - data sets
    - access methods 33
    - conversion 34
    - logical record length 11, 35
    - management 33
    - physical record length 11, 35, 51
    - placement 38-40
    - space allocation 33, 51-59
  - data streaming
    - ACP/TPF 29
    - attachment requirements 22
    - channel 6
    - 3380 model 13 22
  - data verification 43-44
  - defect, surface 31, 42-44
  - device
    - address (*see* address, device)
    - physical ID 46
    - reserve release 14, 20
    - status 21, 41
  - Device Support Facilities
    - allocate VTOC space 31
    - Analyze command 43
    - assign alternate track 31, 43
    - BUILDIX command 31
    - create VTOC index 31
    - description 30
    - error handling 43
    - INIT command 31, 43
    - initialize volume 43
    - INSPECT command 43
    - place IPL bootstrap record 31
    - skip surface defects 43
    - uses 30, 31, 43
    - write volume serial 31
  - direct access device space management (DADSM) 27, 38
  - direct access method (DAM)
    - block size 37
    - data set conversion 34
    - dump/restore 47
    - space allocation 37
    - 3380 support 33
  - directories 39
  - disable 21
  - DISPLAY U 21
  - disk 8
  - drive 8
  - dump
    - frequency 47
    - number tape reels 47
    - program support 47
    - time 47
  - dynamic path selection
    - ACP support 29
    - control program support 18-20
    - controllers 18-20
    - data placement 18-20
    - device reserve 20
    - internal paths 18
    - overlapping operations 18
    - VM/SP support 20
- ## E
- emergency 21
  - Enable/Disable switch 21
  - Environmental Recording, Editing, and Printing (EREP) 30, 41, 42
  - equipment, errors 44
  - error
    - control 41
    - correction 41-44
    - data 42-44
    - defective track 31, 43
    - equipment 44
    - function 41
    - logging 45
    - reports 41, 45
  - error recovery procedures (ERP) 41
  - extended
    - architecture 20
    - commands 14-16, 39
  - floor space 23
- ## G
- group name 28, 33
  - group path ID
    - commands 14, 20
    - device reserve extension 20

## H

- hardware (*see* equipment)
- head (*see* address, head)
  - number of 9
- head and disk assembly
  - initialization 31
  - number of access mechanisms 9
  - number per unit 9
  - removal 9
  - replacement 31
- heat 23
- Hierarchical Storage Manager
  - backup 47
  - data set conversion 34
  - description 30
- home address
  - commands 13, 14
  - data errors 43-44
  - initialization 31
  - track format 10

## I

- IBM Direct Access Storage Migration Aid
  - blocking aid 37
  - description 30
  - JCL changes 33
  - utility statements 34
- index
  - track 11
  - VTOC 27, 30, 31
- indexed sequential access method (ISAM)
  - block size 37
  - data management 33
  - data set conversion 34
  - dump/restore 47
- initialization
  - levels of 31
  - volume 31, 32
- installation 22
- internal paths 18, 19
- I/O address (*see* address)

## J

- Job Control Language (JCL) 33
- Job Entry System (JES) 27

## K

- key
  - length 35, 51-59
  - record area 12
  - write/read commands 13

## L

- length
  - data 10, 35
  - key 10, 35
  - logical record 35
  - physical record 11, 35
- logging 45
- logical record
  - effect on block size 37
  - length 35
  - related to physical record 35

## M

- macros
  - ADDRESS 28
  - DATASET 33
  - DCB 33, 35
  - ENQ 20
  - EXCP 33, 34, 36
  - I/ODEVICE 28
  - RDEVICE 28
  - SYSRES 28
  - SYSTYPE 28
  - TRKCALC 52
  - UNITNAME 28, 33
- main storage
  - buffers 37
  - requirements 49
- migration, *see also* conversion
  - aid program 30
  - Hierarchical Storage Manager 30
- models AA
  - attachment to 3880 3-6
  - device reserve/release 20
  - dynamic path selection 18
  - dynamic reconnection 20
  - operating system support 20
  - overlapping operations 18
- models, 3380
  - attachment to 3880 3-5
  - changes 3
  - configurations 3-5
  - description 3-5
  - functions 3
  - two controllers 3

- models, 3880
  - processor attachment 1
  - 3380 attachment 3-5
  - with cache 17
- multiple count, key, and data operations
  - command 13
  - error handling 42
- multiple track operations 15
- Multiple Virtual Storage (MVS)
  - data set conversion 34
  - device reserve release 20
  - dynamic path selection 18
  - processors supported 27
  - system control program 27
  - system generation 27
  - uses 28

## O

- operating system 27-29
- operating system/virtual storage 27
- operator controls 21

## P

- page data set
  - effective 3380 use 39, 40
  - 3880 model 13 17
- paging
  - MVS/SP 28
  - VM/SP 28
- parameters and subparameters
  - BLKSIZE, JCL 33
  - DCB, JCL 33
  - KEYLEN, JCL 35
  - OPTCD 38, 39
  - RECFM 39
  - SPACE, JCL 33
  - UNIT, JCL 33
  - with DSF commands 31, 43
- partitioned access method
  - block size consideration 37
  - data management 33
  - data set conversion 34
  - dump/restore 47
- paths
  - enable/disable 21
  - VARY commands 21
  - external 18-20
  - internal 18
- path, group ID 14, 20
- physical identifiers 21, 46

- physical planning 22
- physical record
  - equal length 51
  - number per cylinder 36, 37
  - number per track 36, 37
  - related to block size 11, 35
  - related to logical record 11, 35
  - unequal length 52
- power on/off 21
- processors
  - operating system support 27-29
  - requirements 22
  - 3380 attachment 1
- programs
  - how to order 27
  - 3380 support 27-30

## Q

- queued access method 33

## R

- ratio, volume capacity 9
- read
  - commands 13
  - error handling 41-43
- record
  - address 10, 25
  - format 10
  - length 11
  - logical 11, 12
  - physical 11, 12
  - track descriptor (*see* record 0)
- record 0
  - track format 10
  - commands 13
  - initialization 31, 32
  - error handling 42
- recovery
  - error 41-45
  - backup 47
- reserve/release 14, 20
- restore 47
- rotational delay 15
- rotational position sensing
  - commands 12, 13
  - standard feature 15
  - user-written access 34
- rotational time 9

## S

- search
  - record format 10
  - command 13
  - direct technique 39, 60
  - previous technique 39
- seek
  - address 25
  - command 12
  - data placement 38
  - record format 10
  - time 9
- sense information 41
- sequential access method (SAM)
  - block size 37
  - data management 33
  - extended (SAM-E) 60
  - dump/restore 47
- sequential organization 37
- sort/merge 30
- space
  - calculation formulas 51
  - data set allocation 33, 35, 39
  - floor requirements 23
- speed matching buffer
  - channel speed 6
  - commands 14-15
  - operations 16
  - program support 16
  - processor requirement 22
- spooling
  - MVS/SP 28
  - VM/SP 28
- Start I/O Instruction (SIO)
  - address 24
  - data access 33, 34
- status
  - maintained by 3880 41
  - reported 41
  - unit check 25, 41
  - unit status 21
  - 3380 status 21
- storage directors, 3880
  - address 25
  - channel switching 6
  - speed matching buffer 6, 22
  - 3380 attachment 3-5
- strings, 3380 3
- switch
  - channel feature 6
  - 3380 control panel 21
- system console commands 21
- system control programs 27-29

- system error log record 41, 45
- system error recovery procedures 41-45
- system generation
  - ACP 29
  - MVS/SP 27
  - unit address 24
  - VM/SP 28
- System Product 27, 28
- System 370
  - attachment requirements 22
  - channel data rates 6, 22
  - data streaming 22
  - extended architecture 20
  - speed matching buffer 22
  - support 22
  - operating system support 27-29

## T

- time sharing option (TSO) 27
- track
  - alternate 31, 43
  - capacity 8, 51-59
  - condition information 31, 43
  - defective 31, 43
  - descriptor record 10, 13, 31
  - format 10
  - home address 10, 13-14, 31
  - index 10
  - initialization 31
  - number of 9
- transfer length factor 14, 16

## U

- units 3
- unit address
  - in error message 41
  - in sense information 41
  - system generation 24, 28
  - valid I/O 25
- unit check 41-44
- unit emergency switch 21
- unit status 41
- utilities
  - IEBISAM 34
  - IEBGENER 34
  - IEHATLAS 44
  - IEHCOPY 34
  - IEHMOVE 34
  - VM/SP DASD dump/restore 34

## V

- VARY commands 21
- virtual machine (VM)
  - data set conversion 34
  - limitations 28
  - minidisk size 28
  - processors supported 28
  - system control program 28
  - system generation 28
  - uses 28
  - 3380 model AA support 20
- virtual storage access method (VSAM)
  - access method service 33
  - block size consideration 37
  - data set conversion 34
  - data management 33
  - dump/restore 47
- VM/SP commands
  - COPYFILE 34
  - MOVEFILE 34
  - TAPE 34
  - TAPPDS 34
- volume
  - access mechanism 8
  - address 24
  - capacity 9
  - definition iii, 8
  - initialization 31
- volume table of contents (VTOC)
  - default 31
  - Device Support Facilities 30-31
  - placement 31
  - space allocation 31
- VTOC index
  - creation 31
  - DFDS support 29
  - invalid format 5 34
  - placement 31
  - use 29

## W

- write commands 13, 15
- write pad 15, 51

## Numeric

- 3031, 3032, 3033 Processors
  - attachment requirements 22
  - channel data rate 6, 22
  - data streaming 22
  - speed matching buffer 22
  - operating system support 27-29
- 3081 Processor
  - attachment requirements 22
  - channel data rate 6, 22
  - channel reconnection 20
  - data streaming 22
  - speed matching buffer 22
  - operating system support 27-29
- 32-byte increment 35, 51
- 3420 Magnetic Tape 47
- 3880 Storage Control
  - attachment of 3380 3-5
  - attachment to channels 6
  - cache 17
  - channel commands 12
  - command retry 15, 41-44
  - status and sense information 41
  - (*see also*, speed matching buffer)
- 4341 Processor
  - attachment requirements 22
  - channel data rate 6, 22
  - data streaming 22
  - speed matching buffer 22
  - operating system support 27-29



IBM 3380 Direct Access Storage Description  
and User's Guide

READER'S  
COMMENT  
FORM

Order No. GA26-1664-1

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.)

Please use pressure sensitive or other gummed tape to seal this form.

Cut Along This Line

Reader's Comment Form

Cut or Fold Along Line

IBM 3380 Direct Access Storage Description and User's Guide (File No. S/370-07, 4300-07) Printed in U.S.A. GA26-1664-1

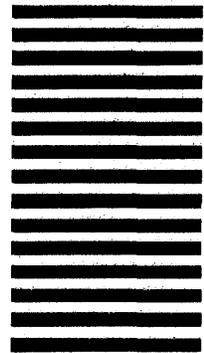
Fold and Tape

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.

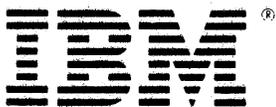
POSTAGE WILL BE PAID BY ADDRESSEE:

International Business Machines Corporation  
Department G26  
555 Bailey Avenue  
San Jose, California 95150

Fold and Tape

Please Do Not Staple

Fold and Tape



IBM 3380 Direct Access Storage Description  
and User's Guide

READER'S  
COMMENT  
FORM

Order No. GA26-1664-1

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.)

Please use pressure sensitive or other gummed tape to seal this form.  
Cut Along This Line

Reader's Comment Form

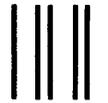
Cut or Fold Along Line

IBM 3380 Direct Access Storage Storage Description and User's Guide (File No. S/370-07, 4300-07) Printed in U.S.A. GA26-1664-1

Fold and Tape

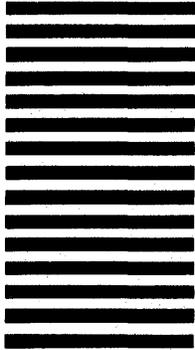
Please Do Not Staple

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 G26  
555 Bailey Avenue  
San Jose, California 95150

Fold and Tape

Please Do Not Staple

Fold and Tape



IBM 3380 Direct Access Storage Description and User's Guide (File S/370-07, 4300-07) Printed in U.S.A. GA26-1664-1

GA26-1664-1

