

IBM 3375
Direct Access Storage

Description and User's
Guide

Cross-System

IBM



**IBM 3375
Direct Access Storage**

**Description and User's
Guide**

Cross-System

Publication Number
GA26-1666-2

File Number
S/370-07, 4300-07

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 instruction manual, may cause interference to radio communications. It has been tested and found to comply with the limits for a Class A computing device pursuant to Subpart J of Part 15 of FCC Rules, which are designed to provide reasonable protection against such interference when operated in a commercial environment. Operation of this equipment in a residential area is likely to cause interference in which case the user at his own expense will be required to take whatever measures may be required to correct the interference.

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

Third Edition (April 1983)

This major revision obsoletes GA26-1666-1 and Technical Newsletters GN26-0367 and GN26-0368. Changes or additions to the text and illustrations are indicated by a vertical line to the left of the change.

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

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

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

A form for readers' comments is provided at the back of this publication. If the form has been removed, comments may be addressed to IBM General Products Division, Technical Publications, Department G26, 555 Bailey Avenue, San Jose, California 95150. IBM may use or distribute whatever information you supply in any way it believes appropriate without incurring any obligation to you.

© Copyright International Business Machines Corporation 1980, 1981, 1983.

Preface

This manual describes the 3375 Direct Access Storage and gives installation and conversion information. It contains detailed information on 3375 models, attachment configurations, device characteristics, addressing, and capacity formulas. These formulas, allow calculations for equal and unequal data lengths and are supplemented by tables listing capacities for commonly used data lengths.

This manual also contains general information and guidelines on how to use this particular storage system, including the supporting programs. A checklist aids the user with step-by-step conversion tips. This manual 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 3375. These manuals are listed in the bibliography, which precedes the Index of this book.

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.

The *IBM Vocabulary for Data Processing, Telecommunications, and Office Systems*, GC20-1699, defines terms used in IBM publications. A glossary of abbreviations used in this manual is given in the back of this book, preceding the Bibliography.

Prerequisite Publications

- *IBM 3880 Storage Control Models 1, 2, 3, and 4 Reference Manual*, Order No. GA26-1661, gives detailed descriptions of channel commands, and status and sense information.
- *IBM Disk Storage Management Guide*, Order No. GBOF 1205, describes data and equipment errors that can occur in disk storage operations. It also provides tutorial information on the Environmental Recording, Editing, and Printing (EREP) program System Exception reports.
- *IBM Reference Summary for Handling Data Checks*, Order No. GX26-1669, contains information extracted from the *IBM Disk Storage Management Guide, Error Handling* manual, Order No. GA26-1672.
- *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 3375.

References to Other Publications

All publications referred to in this book are listed in the Bibliography, which precedes the Index. The Bibliography also lists other related publications that are not referred to in text.

Contents


Introduction	1	Service Procedures	22
General Description	2	Alternate Access	22
3375 Models and Configurations	2	Statistical Information and Logging	23
String Switch	2		
Speed Matching Buffer for 3375	2	Hardware Planning Considerations	24
Programming Support	2	Physical Planning	24
		Power and Air Conditions	24
Data Organization and Capacity	5	Space Requirements	24
Head Disk Assembly	5	Processor Attachment Requirements	24
Track and Record Formats	5	Channel Switches	24
Track Format	5	Addressing	26
Record Format	6	Address and Attachment Configurations	27
Physical Records Per Track	6		
Track and Record Capacity Formulas	7	Program Planning Considerations	32
Equal Length Physical Records	7	DOS/VSE Support	32
Unequal Length Physical Records	7	VM/SP Support	32
Selection of Data Block Sizes	8	OS/VS1 Support	32
Track, Cylinder, and Access Mechanism Capacity		MVS/SP Support	32
Tables	9	MVS/XA DFP and MVS/SP Version 2	32
		MVS/370 DFP and SP 1.3.0	33
		Programs Using OS/VS and the 3375	33
Channel Commands	14	Additional Programming Support	33
Control Commands	14	Device Support Facilities	33
Search Commands	14	Program Products Not Supported By DOS	33
Read Commands	15	Data Facility Device Support	33
Write Commands	15	Data Facility Extended Function	33
Sense Commands	15	Data Facility Data Set Services	33
Diagnostic Commands	15	IBM Direct Access Storage Device Migration Aid ..	34
Command Sequence	16	Hierarchical Storage Manager	34
Speed Matching Buffer	16		
Speed Matching Buffer in 3880, Models 1 and 2	16	Data Management and Data Set Conversion	35
Operations for 1.5 MB/sec	16	Job Control Language Uses	35
Operations for 1.5 and 1.859 MB/sec,		Data Set Conversion	35
Same Processor	17	Standard Access Methods	35
Standard Features	18	User-Written Access	36
Command Retry	18	Conversion Under VM/SP	36
End-of-File	18		
Multiple Track Operations	18	Effective Use of the 3375	37
Format Write Release	18	Criteria for Using the 3375	37
Rotational Position Sensing	18	Space and Performance Guidelines	37
Surface Defect Skipping	18	Optimum Blocking for 3375 Records	37
Overlapping Operations	18	Background Information	37
		Performance Considerations	38
Device Status and Error Handling	20	Application Considerations	39
Program Resources	20	Page Data Sets, Special Considerations	39
Error Handling Guidance	20		
Programming Errors	21	Conversion Checklist	40
Control Errors	21		
Data Errors	21	Operator Controls and Indicators	43
Correction of Data Errors	21	Unit Operator Panel	43
Data Verification by Program	22	Ready Indicators	43
Equipment Errors	22		

Ready Switches.....	43
Read/Write or Read Only Switches	43
Controller Operator Panel.....	44
Delayed Power Off/On Switch	44
Power On Indicator	44
Data Path A and B Enable/Disable Switches	44

Glossary	45
Bibliography	46
Index	48

Figures

Figure 1. Comparison of Characteristics of Different IBM Disk Storage Types	1
Figure 2. Functional Overview.....	4
Figure 3. Track and Record Formats, Simplified Schematic.....	5
Figure 4. Differences in Device Space Requirements between the 3350 and 3375	8
Figure 5. Record Capacities, Without Keys	10
Figure 6. Record Capacities, With Keys.....	12
Figure 7. Comparison of Approximate Floor Space Requirements Based on Full String Configurations	24
Figure 8. Channel Switching for 3880 Storage Directors by 3375 Disk Storage Attachment.....	25
Figure 9. 3880 and 3375 Device Addressing.....	26
Figure 10. Sets of Valid 3880/3375 Addresses.....	26
Figure 11. Example of Four 3375 A1 Strings Attached to Storage Directors of 3880-1.....	28
Figure 12. Possible Configurations Between 3880 and 3375 Dual Controller Strings.....	29
Figure 13. Unsupported Processor/Channel Address Designation.....	30
Figure 14. Supported Processor/Channel Address Configuration.....	31
Figure 15. 3375 Unit Operator Panel	43
Figure 16. 3375 Controller Operator Panel.....	44



3375 Direct Access Storage (Design Model)

Introduction

The IBM 3375 Direct Access Storage, for intermediate and large systems, provides up to 3,729 megabytes of online storage per string. The 3375 uses the count, key, data format and command set that is compatible with previous IBM direct access storage devices.

The 3375 attaches through a 3880 Storage Control model 1, 2, or 4 to the 4331 Processor model group 2 and the 4341 Processor.

A prerequisite for the 3375 attachment is a channel that can support the 1.859 megabytes per second (MB/sec) data rate. However, the speed matching buffer (in 3880 models 1 and 2 only) provides for attachment to channels with data transfer rates as low as 1.5 MB/sec. This feature allows the 3375 attachment to System/370 models 145, 145-3, 148, 155-II, 158, 158-3, 165-II, 168, and 168-3.

The 3375 can also be attached through the 3880 (models 1 or 2) to the 3031, 3032, 3033, 3042-2, 3081, 3083, or 3084 processors. On 3031, 3032, 3033, and 3042 model 2 processors, the 3375 can attach via the data streaming or speed matching buffer features. Attachment to a 3081, 3083, or 3084 is through any block multiplexer channel.

The 3375 writes at higher densities and has significantly higher capacity than the IBM 3340, 3340/3344, or 3350. The read and write heads, along with the disks, are sealed in an enclosure called a head and disk assembly (HDA), which provides protection from contaminants. This head and disk assembly has two separately addressed access mechanisms (actuators). Each access mechanism services half of the magnetic disks in the head and disk assembly. Seeking with either access mechanism may be overlapped with seeking, reading, or writing by the other access mechanism.

The 3375 has a number of functions that result in improved disk storage performance. These include such features as rotational position sensing, command retry, surface defect skipping, and such programming techniques as chained scheduling and search direct for sequential data sets. Improved device performance characteristics include faster seek times and a significantly higher data transfer rate than that of the 3340, 3340/3344, and 3350. In operation, the 3375 requires less operating power and floor space than these other storage devices.

The efficiency and economy of the 3375 results from the following characteristics:

- Storage capacity: 409.8 megabytes per access mechanism, 819.7 megabytes per unit, more than three billion bytes per four-unit string.
- Storage cost per megabyte: below that of the 3340, 3340/3344, or 3350.
- Nominal data rate: 1.859 megabytes per second.
- Average access time: 19 milliseconds (ms).
- Average rotational delay time: 10.1 milliseconds (0.0101 seconds).
- Energy efficiency: savings in floor space, power and heat output can be attained when the 3375 replaces existing 3340, 3340/3344, and 3350 units.

Figure 1 compares the 3375 performance characteristics with previous IBM disk storage.

Performance and Physical Characteristics	3375	3350	3340	3344
Seek Time (ms) Average	19	25	25	25
Full Track Rotation Time (ms)	20.2	16.7	20.24	20.24
Data Transfer Rate (MB/sec)	1.859	1.198	0.885	0.885
MB Capacity: per Unit	819.7	635	140	560
per Drive	819.7	317.5	70	280
per Access Mechanism	409.8	317.5	70	280
MB Capacity, Fixed-Head Unit:				
per Access Mechanism	0	1.14	0.5	1.0
per Drive	0	2.28	1.0	2.0
Data Cylinders per Access Mechanism	959	555	696	4X696
Data Tracks per Cylinder	12	30	12	12
Data Bytes: per Track	35,616	19,069	8,368	8,368
per Cylinder	427,392	572,070	100,416	100,416

Figure 1. Comparison of Characteristics of Different IBM Disk Storage Types

General Description

The IBM 3375 Direct Access Storage operates under the control of the IBM 3880 Storage Control to which it is attached (see Figure 2). A 3375 string can consist of up to four disk storage units. Each unit contains a head and disk assembly that is not removable by the customer.

The first disk storage unit in the 3375 string contains a controller that communicates with the 3880 Storage Control. On dual controller strings, the last attached disk storage unit contains an additional controller. The controller decodes the orders received from the storage control, determines the addressed device, and monitors subsequent operations with the disk storage.

3375 Models and Configurations

The 3375 disk storage is available in the following three models:

Model A1 is a disk storage unit with one head and disk assembly and a controller that contains logic for attachment to the 3880 Storage Control. It also contains logic and power for attachment of up to three 3375 model B1s or two model B1s and one model D1.

Model B1 is a disk storage unit with one head and disk assembly. Up to three model B1s can be attached to one model A1.

Model D1 is a disk storage unit with one head and disk assembly and control logic for attachment to the 3880 Storage Control. The model D1 can only be used in a complete string of four units and replaces the last model B1. It provides a dual controller function and a second path to each HDA. Model D1 may be attached to the same system as the model A1 or to a different system. It cannot be attached to the same 3880 storage director nor through the same channel on the processor. See the Hardware Planning Considerations section for attachment and addressing configurations.

The three models, A1, B1, and D1 can be attached in the following storage configurations:

One model A1

Storage capacity - 819.7 megabytes

One model A1 and one model B1

Storage capacity - 1,639 megabytes

One model A1 and two model B1s

Storage capacity - 2,459 megabytes

One model A1 and three model B1s

or

One model A1, two model B1s, and one model D1

Storage capacity - 3,279 megabytes

String Switch

The string switch is an optional feature that allows a 3375 model A1 or model D1, or both, and all model B1s to be attached to a second storage director. The two storage directors can be attached to the same processor or different processors. Switching between the two storage directors is under program control. The 3375 may be dedicated to a single attachment by manual Enable/Disable switches at the controller.

Speed Matching Buffer For 3375

The speed matching buffer for 3375 allows the 3375 to attach to block multiplexer channels that operate at data transfer rates as low as 1.5 MB/sec. Special commands have been added to maintain performance when the channel operates at a slower data rate than the 3375. See the Channel Commands section for a description of these commands.

The speed matching buffer can be installed in one storage director of a 3880 model 2 and both storage directors of a 3880 model 1. If a 3375 is attached to two storage directors that connect to the same processor, either both storage directors must have the speed matching buffer or neither can have one. This is because both paths must support the special commands used with the speed matching buffer.

The speed matching buffer is required to attach 3375s to System/370 models 145, 148, 155-II, 158, 165-II, 168, and to block multiplexer channels without the data streaming feature. If a storage director is attached to a channel with a data rate faster than the 3375 data rate (1.859 MB/sec) and a channel slower than the 3375 data rate, the speed matching buffer supports the faster channel at the 3375 data rate of 1.859 MB/sec and the slower channel at 1.5 MB/sec data rate.

Programming Support


The 3375 is supported by the following system control programs (SCPs):

- Disk Operating System/Virtual Storage Extended (DOS/VSE) with Advanced Functions (AF3 and above)
- Virtual Machine/370 (VM/370) Release 6 with VM System Product


- **Operating System/Virtual Storage 1 (OS/VS1) Release 7 with Data Facility Device Support**
- **Operating System/Virtual Storage 2 Multiple Virtual Storage (OS/VS2 MVS) with MVS/System Product and Data Facility Device Support**
- **Multiple Virtual Storage/Extended Architecture (MVS/XA) Data Facility Product and MVS/SP Version 2**

- **Multiple Virtual Storage/370 Data Facility Product and System Product**

See the Program Planning Considerations section for additional details.



Console




Processor

- Executes I/O instructions
- Stores data
- Stores status
- Stores channel program

Channel



- Fetches channel commands
- Controls transfer of data between the processor and the 3880

3375 Disk Storage



* Model 4 has only one storage director.

Other Disk Storage



Disk Storage Device

- Responds to orders from the 3880.
- Positions access mechanism
- Selects head
- Reads and writes data on a count, key, and data basis.
- Serializes and deserializes data
- Performs error detection and correction

3880 Storage Control

- Interprets and executes commands from channel
- Starts and controls disk storage operations with orders.
- Controls channel and disk storage interfaces
- Furnishes status to system
- Performs diagnostic evaluation of the storage control and attached disk storage.

Figure 2. Functional Overview

Data Organization and Capacity

This section describes the initial data recorded on the head and disk assembly and the track and record formats of the data. It also includes record-per-track formulas and track and cylinder capacity charts for various record lengths.

Head and Disk Assembly

The 3375 HDA is initialized at time of manufacture with a home address and a standard track descriptor record (R0) on all tracks. (Addressing is described in the Hardware Planning Considerations section.)

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

The 3375 is shipped from the plant of manufacture with no flagged defective tracks. If data errors that cannot be corrected are encountered during use, the Device Support Facilities program can be used to generate skip displacements or assign an alternate track.

Track and Record Formats

The 3375 has a count, key, and data recording format compatible with existing IBM direct access storage products.

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 3.

Home Address

Each track contains a home address (HA) that identifies the track. It is the track address and is defined by cylinder and head number (CCHH). In addition, the home address area contains information describing the condition of the track. Home address is the first recorded area following the index point.

Special channel commands are used for writing and reading a home address area: Write Home Address and Read Home Address.

Track Descriptor Record


Record 0 (R0) is always the first record on the track following the home address. The R0 can be standard or non-standard. Standard R0s contain a count area, no key area, and an eight-byte data area. Usually, R0 is reserved by the operating system for nonuser functions. However, non-standard R0s may be written with all zeros in the data field or treated as a normal user data record.

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

In IBM programming systems, the count area of R0 on a defective track provides the address of the alternate track. The count area of R0 on an alternate track provides the address of the defective track it replaced.

Cylinder 03

Head 00



Head 01




Figure 3. Track and Record Formats, Simplified Schematic

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. 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.

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 of the key and data areas of the record.

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.

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 the records is determined by the user, except for track descriptor records.

Physical record size is specified depending 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 Write Count, Key, and Data channel command.

The method for describing data characteristics depends on the operating system and access method. The following section, Track and Record Capacity Formulas, provides information for determining the number of records that can be placed on the HDA tracks and cylinders.

Track and Record Capacity Formulas

The maximum length (block size) of a physical record is the capacity remaining on one track after device overhead has been accounted for. Device overhead refers to the space required on each track for hardware information, such as the home address, gaps, and count areas. Information and guidelines for determining optimum block sizes are given later in this section under Selection of Data Block Sizes.

All data recorded on the HDA is written in 32-byte segments. The formulas for making 3375 capacity calculations are based on this 32-byte segment and differ from those for the 3340, 3344, and 3350 disk storage devices.

The following formulas can be used to determine the number of equal length records that can be placed on a track. Equal length records are assumed, unless stated otherwise.

Equal Length Physical Records

The number of equal length physical records per track can be calculated as follows:

$$\text{Equal length records per track} = \frac{1125}{C + K + D}$$

1125 is the number of 32-byte segments per track available for user data records.

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

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

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

KL is the key length in bytes.

DL is the data length in bytes.

where C = 12 if KL = 0

C = 17 if KL ≠ 0

K = 0 if KL = 0

$$K = \frac{KL^*}{32} \quad \text{if } KL \neq 0$$

$$D = \frac{DL^*}{32}$$

* KL and DL each should be rounded up to a multiple of 32.

With a partitioned data set directory where KL = 8 and DL = 256:

$$K = \frac{8}{32} = \frac{8}{32} \text{ rounded up to a multiple of 32}$$

$$= \frac{256}{32} = 8 \text{ segment}$$

$$D = \frac{256}{32} = 8 \text{ segments}$$

$$\text{Then: } \frac{1125}{17+1+8} = \frac{1125}{26} = 43.27 \text{ records}$$

(rounded down to 43 equal length records per track)

A tabular summary of this formula for records with and without keys appears in Figures 5 and 6.

Note: The following technique is an easy way to maintain the count of the amount of space available on an HDA track. It assumes standard R0.

1. Start with 36,000 bytes
2. For each record written without a key, subtract: (data length)* + 384 bytes
3. For each record written with a key, subtract: (data length)* + (key length)* + 544 bytes

* Denotes rounding to a multiple of 32.

Unequal Length Physical Records

In the case of unequal length records, the number of records (n) per track must satisfy the following relationship:

<p>Standard R0 (KL = 0, DL = 8)</p> $1125 - \sum_{i=1}^n (C_i + K_i + D_i) \geq 0$
--

<p>Non-standard R0 (R0 used for data record)</p> $1137 - \sum_{i=1}^n (C_i + K_i + D_i) \geq 0$

C, K, and D are the same as defined for equal length records for the ith record.

Selection of Data 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 (for example, track formats, capacities, and blocking factors) of the particular device type must be considered.

Figure 4 illustrates differences in space requirements between the 3350 and 3375. For comparison, assume a data set of 400,000 logical records of 80 bytes each. The larger block size uses fewer tracks or cylinders on both the 3350 and the 3375.

When converting data sets where space has been allocated in track and 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 3375, one record with a data length of 35,616 bytes can be placed on a track. However, the largest block size supported by the IBM-supplied access method is 32,760 bytes. EXCP must be used to obtain block sizes larger than 32,760 bytes.

The next best use of space is obtained with two records of 17,600 bytes, which achieve approximately 99% utilization. A block size larger than one-half 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 128 bytes were chosen, about 25% of the capacity would be available for user data. Block sizes that correspond to the 80-byte IBM punched card use only 17% of the full data capacity of a 3375.

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 3,400 and 3,040 bytes accommodates 10 physical records, without keys, on a 3375 track. The 3,400 data block size utilizes approximately 95% while the 3,040 data block size utilizes about 85% of the available capacity. The track capacity charts, Figures 5 and 6, list block sizes that achieve the best use of space for a given number of records.

In general, block sizes between 3,000 and 8,000 bytes are most practical for the 3375. 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 87% space utilization on the 3340, 3344, 3350, and 3375 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.

Blocking Factor	Block Size	Blocks Required	Blocks per Track		Tracks Required		Blocks per Cylinder		Cylinders Required	
			3350	3375	3350	3375	3350	3375	3350	3375
10	800	40,000	19	30	2,105	1,334	570	360	71	112
100	8,000	4,000	2	4	2,000	1,000	60	48	67	84

Figure 4. Differences in Device Space Requirements between the 3350 and 3375

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 3375 is so large (35,616 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 11,616 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 logic records in the data set. For example, if the data set contains 80-byte logical records, a block size of 6,800 might be chosen. This allows eighty-five 80-byte logical records to be blocked into the block, therefore, permitting five blocks to be placed on a track. This results in 96% utilization.

Blocking Aid

The *Direct Access Storage Device Migration Aid, General Information* manual can help to determine which existing data sets need to be reblocked and the block sizes to be used. (This applies to OS/VS only and does not apply to VSAM.) This program provides a means for determining which block sizes meet the user's criterion for percentage of utilization.

Track, Cylinder, and Access Mechanism Capacity Tables

The following tables give the number of equal length records that can be placed on a track and cylinder. Figure 5 gives the number of records if there are no keys (KL = 0); Figure 6 gives the number of records if there are keys: Part 1 is used for key lengths 1 through 32 bytes; Part 2 is used for key lengths greater than 32 bytes up to a maximum of 255 bytes.

The number of 32-byte segments that can be placed on a track depends on the data length (DL). A range of data lengths results in a given number of segments. For example, for records without keys, data lengths from:

5,601 to 6,816 bytes allow 5 records per track, 60 records per cylinder.

6,817 to 8,608 bytes allow 4 records per track, 48 records per cylinder.

8,609 to 11,616 bytes allow 3 records per track, 36 records per cylinder.

The tables show the maximum data lengths that can be used for the given number of records. Key and data areas are assumed to occupy the full 32-byte segments. These two tables include the total number of bytes used for the given number of records on a track and cylinder.

Record Length (* Total Bytes)	Track Capacity		Cylinder Capacity		Access Mechanism Capacity	
	Records	Bytes	Records	Bytes	Records	Bytes
35,616	1	35,616	12	427,392	11,508	409,868,928
17,600	2	35,200	24	422,400	23,016	405,081,600
11,616	3	34,848	36	418,176	34,524	401,030,784
8,608	4	34,432	48	413,184	46,032	396,243,456
6,816	5	34,080	60	408,960	57,540	392,192,640
5,600	6	33,600	72	403,200	69,048	386,668,800
4,736	7	33,152	84	397,824	80,556	381,513,216
4,096	8	32,768	96	393,216	92,064	377,094,144
3,616	9	32,544	108	390,528	103,572	374,516,352
3,200	10	32,000	120	384,000	115,080	368,256,000
2,880	11	31,680	132	380,160	126,588	364,573,440
2,592	12	31,104	144	373,248	138,096	357,944,832
2,368	13	30,784	156	369,408	149,604	354,262,272
2,176	14	30,464	168	365,568	161,112	350,579,712
2,016	15	30,240	180	362,880	172,620	348,001,920
1,856	16	29,696	192	356,352	184,128	341,741,568
1,728	17	29,376	204	352,512	195,636	338,059,008
1,600	18	28,800	216	345,600	207,144	331,430,400
1,504	19	28,576	228	342,912	218,652	328,852,608
1,408	20	28,160	240	337,920	230,160	324,065,280
1,312	21	27,552	252	330,624	241,668	317,068,416
1,248	22	27,456	264	329,472	253,176	315,963,648
1,152	23	26,496	276	317,952	264,684	304,915,968
1,088	24	26,112	288	313,344	276,192	300,496,896
1,056	25	26,400	300	316,800	287,700	303,811,200
992	26	25,792	312	309,504	299,208	296,814,336
928	27	25,056	324	300,672	310,716	288,344,448
896	28	25,088	336	301,056	322,224	288,712,704
832	29	24,128	348	289,536	333,732	277,665,024
800	30	24,000	360	288,000	345,240	276,192,000
768	31	23,808	372	285,696	356,748	273,982,464
736	32	23,552	384	282,624	368,256	271,036,416
704	33	23,232	396	278,784	379,764	267,353,856
672	34	22,848	408	274,176	391,272	262,934,784
640	35	22,400	420	268,800	402,780	257,779,200
608	36	21,888	432	262,656	414,288	251,887,104
576	37	21,312	444	255,744	425,796	245,258,496
544	38	20,672	456	248,064	437,304	237,893,376
512	40	20,480	480	245,760	460,320	235,683,840

*Includes DL and unused bytes when rounding up to 32-byte segments.

Figure 5. Record Capacities, Without Keys (Part 1 of 2)

Record Length (*Total Bytes)	Track Capacity		Cylinder Capacity		Access Mechanism Capacity	
	Records	Bytes	Records	Bytes	Records	Bytes
480	41	19,680	492	236,160	471,828	226,477,440
448	43	19,264	516	231,168	494,844	221,690,112
416	45	18,720	540	224,640	517,860	215,429,760
384	46	17,664	552	211,968	529,368	203,277,312
352	48	16,896	576	202,752	552,384	194,439,168
320	51	16,320	612	195,840	586,908	187,810,560
288	53	15,264	636	183,168	609,924	175,658,112
256	56	14,336	672	172,032	644,448	164,978,688
224	59	13,216	708	158,592	678,972	152,089,728
192	62	11,904	744	142,848	713,496	136,991,232
160	66	10,560	792	126,720	759,528	121,524,480
128	70	8,960	840	107,520	805,560	103,111,680
96	75	7,200	900	86,400	863,100	82,857,600
64	80	5,120	960	61,440	920,640	58,920,960
32	86	2,752	1,032	33,024	989,688	31,670,016

*Includes DL and unused bytes when rounding up to 32-byte segments.

Figure 5. Record Capacities, Without Keys (Part 2 of 2)

Data Length (*Total Bytes)	Track Capacity		Cylinder Capacity		Access Mechanism Capacity	
	Records	Bytes	Records	Bytes	Records	Bytes
35,424	1	35,424	12	425,088	11,508	407,659,392
17,408	2	34,816	24	417,792	23,016	400,662,528
11,424	3	34,272	36	411,264	34,524	394,402,176
8,416	4	33,664	48	403,968	46,032	387,405,312
6,624	5	33,120	60	397,440	57,540	381,144,960
5,408	6	32,448	72	389,376	69,048	373,411,584
4,544	7	31,808	84	381,696	80,556	366,046,464
3,904	8	31,232	96	374,784	92,064	359,417,856
3,424	9	30,816	108	369,792	103,572	354,630,528
3,008	10	30,080	120	360,960	115,080	346,160,640
2,688	11	29,568	132	354,816	126,588	340,268,544
2,400	12	28,800	144	345,600	138,096	331,430,400
2,176	13	28,288	156	339,456	149,604	325,538,304
1,984	14	27,776	168	333,312	161,112	319,646,208
1,824	15	27,360	180	328,320	172,620	314,858,880
1,664	16	26,624	192	319,488	184,128	306,388,992
1,536	17	26,112	204	313,344	195,636	300,496,896
1,408	18	25,344	216	304,128	207,144	291,658,752
1,312	19	24,928	228	299,136	218,652	286,871,424
1,216	20	24,320	240	291,840	230,160	279,874,560
1,120	21	23,520	252	282,240	241,668	270,668,160
1,056	22	23,232	264	278,784	253,176	267,353,856
960	23	22,080	276	264,960	264,684	254,096,640
896	24	21,504	288	258,048	276,192	247,468,032
864	25	21,600	300	259,200	287,700	248,572,800
800	26	20,800	312	249,600	299,208	239,366,400
736	27	19,872	324	238,464	310,716	228,686,976
704	28	19,712	336	236,544	322,224	226,845,696
640	29	18,560	348	222,720	333,732	213,588,480
608	30	18,240	360	218,880	345,240	209,905,920
576	31	17,856	372	214,272	356,748	205,486,848
544	32	17,408	384	208,896	368,256	200,331,264
512	33	16,896	396	202,752	379,764	194,439,168
480	34	16,320	408	195,840	391,272	187,810,560
448	35	15,680	420	188,160	402,780	180,445,440
416	36	14,976	432	179,712	414,288	172,343,808
384	37	14,208	444	170,496	425,796	163,505,664
352	38	13,376	456	160,512	437,304	153,931,008
320	40	12,800	480	153,600	460,320	147,302,400
288	41	11,808	492	141,696	471,828	135,886,464
256	43	11,008	516	132,096	494,844	126,680,064
224	45	10,080	540	120,960	517,860	116,000,640
192	46	8,832	552	105,984	529,368	101,638,656
160	48	7,680	576	92,160	552,384	88,381,440
128	51	6,528	612	78,336	586,908	75,124,224
96	53	5,088	636	61,056	609,924	58,552,704
64	56	3,584	672	43,008	644,448	41,244,672
32	59	1,888	708	22,656	678,972	21,727,104

*Data length rounded up to 32-byte segments. This table also adjusted for key lengths of 1 through 32 bytes. See Figure 6 (Part 2 of 2) for key lengths of 32 through 255 bytes.

Figure 6. Record Capacities, With Keys (Part 1 of 2)

Record Length (*Total Bytes)	Track Capacity		Cylinder Capacity		Access Mechanism Capacity	
	Records	Bytes	Records	Bytes	Records	Bytes
35,456	1	35,456	12	425,472	11,508	408,027,648
17,440	2	34,880	24	418,560	23,016	401,399,040
11,456	3	34,368	36	412,416	34,524	395,506,944
8,448	4	33,792	48	405,504	46,032	388,878,336
6,656	5	33,280	60	399,360	57,540	382,986,240
5,440	6	32,640	72	391,680	69,048	375,621,120
4,576	7	32,032	84	384,384	80,556	368,624,256
3,936	8	31,488	96	374,856	92,064	362,363,904
3,456	9	31,104	108	373,248	103,572	357,944,832
3,040	10	30,400	120	364,800	115,080	349,843,200
2,720	11	29,920	132	359,040	126,588	344,319,360
2,432	12	29,184	144	350,208	138,096	335,849,472
2,208	13	28,704	156	344,448	149,604	330,325,632
2,016	14	28,224	168	338,688	161,112	324,801,792
1,856	15	27,840	180	334,080	172,620	320,382,720
1,696	16	27,136	192	325,632	184,128	312,281,088
1,568	17	26,656	204	319,872	195,636	306,757,248
1,440	18	25,920	216	311,040	207,144	298,287,360
1,344	19	25,536	228	306,342	218,652	293,868,288
1,248	20	24,960	240	299,520	230,160	287,239,680
1,152	21	24,192	252	290,304	241,668	278,401,536
1,088	22	23,936	264	287,232	253,176	275,455,488
992	23	22,816	276	273,792	264,684	262,566,528
928	24	22,272	288	267,264	276,192	256,306,176
896	25	22,400	300	268,800	287,700	257,779,200
832	26	21,632	312	259,584	299,208	248,941,056
768	27	20,736	324	248,832	310,716	238,629,888
736	28	20,608	336	247,296	322,224	237,156,864
672	29	19,488	348	233,856	333,732	224,267,904
640	30	19,200	360	230,400	345,240	220,953,600
608	31	18,848	372	226,176	356,748	216,902,784
576	32	18,432	384	221,184	368,256	212,115,456
544	33	17,952	396	215,424	379,764	206,591,616
512	34	17,408	408	208,896	391,272	200,331,264
480	35	16,800	420	201,600	402,780	193,334,400
448	36	16,128	432	193,536	414,288	185,601,024
416	37	15,392	444	184,704	425,796	177,131,136
384	38	14,592	456	175,104	437,304	167,924,736
352	40	14,080	480	168,960	460,320	162,032,640
320	41	13,120	492	157,440	471,828	150,984,960
288	43	12,384	516	148,608	494,844	142,515,072
256	45	11,520	540	138,240	517,860	132,572,160
224	46	10,304	552	123,648	529,368	118,578,432
192	48	9,216	576	110,592	552,384	106,057,728
160	51	8,160	612	97,920	586,908	93,905,280
128	53	6,784	636	81,408	609,924	78,070,272
96	56	5,376	672	64,512	644,448	61,867,008
64	59	3,776	708	45,312	678,972	43,454,208

*Includes KL, DL, and unused bytes when rounding up to 32-byte segments.

How To Use This Table

Assume DL=820, KL=48, then:




Figure 6. Record Capacities, With Keys (Part 2 of 2)

Channel Commands

The channel commands used for 3375 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.
- The names of the Sense I/O and Sense I/O Type commands have been changed to Sense and Sense ID, respectively, but the functions remain the same.
- The following changes have been made to the diagnostic commands:

Diagnostic Sense has been changed to Diagnostic Sense/Read.

Diagnostic Write Home Address and Diagnostic Read Home Address have been added.

Diagnostic Write and Diagnostic Load are not used.

Channel commands are executed by the storage director to which the 3375 is attached and are described in detail in the *IBM 3880 Storage Control Models 1, 2, 3, and 4 Reference* manual. The description of each command is accompanied by a description of the status conditions associated with the execution of the command. Also included is an explanation of the channel command word (CCW) including the format used. Users who need information for writing channel programs should refer to the 3880 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 access mechanism specified by the device address to the desired cylinder 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).

This 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.

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. It is normally initiated by selecting the address of the device in the Load Unit switches on the system console and pressing the IPL pushbutton.

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

Write Commands

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.

Write Data causes the specified data in main storage to be written in the data area of the selected 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.

Write Count, Key, and Data causes an entire record to be transferred from main storage and written on the track.

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 all 24 bytes of sense information from the storage director to main storage.

Sense ID transfers the seven bytes of sense information that describe the type and model of the storage control and device being addressed 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.

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.

Device Release terminates the reservation of the addressed device.

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 can be performed as a single channel program. 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

Speed Matching Buffer

Five commands (Define Extent; Locate Record; Write Update Data; Write Update Key and Data; and Write Count, Key, and Data Next Track) are used to maintain performance when the channel operates at a slower data rate than the 3375. 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.

Two commands, Define Extent and Locate Record, are used to notify the channel of the types of data transfer commands that are going to be executed and the limits in which they may operate. These commands replace the Seek, Set Sector, Set File Mask, and Search commands used in other channel programs.

The data transfer commands, Write Update Data, Write Update Key and Data, and Write CKD Next Track, are provided for write operations on consecutive records following execution of the Define Extent and Locate Record commands.

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 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 the programs dealing with variable length records. The transfer length factor replaces the block size in the Define Extent command or the transfer length 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.

Speed Matching Buffer in 3880, Models 1 and 2

With a speed matching buffer in the 3880 storage director, data is transferred between the buffer and the 3375 at the 3375 transfer rate of 1.859 MB/sec, and between the storage director and channel at the speed supported by the channel. If the channel speed is greater than 1.859 MB/sec, data is transferred at 1.859 MB/sec; otherwise, data is transferred at 1.5 MB/sec. A speed control switch must be set in the 3880 to run at high speed.

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 3375 into the buffer at 1.859 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 (20.2 milliseconds) is lost in this process and if several write commands are chained together, a revolution is lost on each write command.

With the 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 the appropriate amount of data to avoid overrun has been transferred to the buffer prior to reaching the record position on the rotating disk. When the appropriate amount of 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, VS1/Basic Program Extension, 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

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 (execute channel program) 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 indicator (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 Variations

Due to the asynchronous nature of speed matching buffer operation at 1.5 MB/sec, some variations 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 the gap time when the data actually is read from a 3375 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 1.859 MB/sec, Same Processor

In a configuration where two channels with different speeds access the same 3375 strings from the same processor, performance is optimized by using the faster 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 1.859 MB/sec path to be used when it is available is to specify the sequential option for physical path selection.

Note: *Two storage director paths that connect a 3375 to the same processor cannot be mixed. Either both paths*

must have the speed matching buffer or neither must have one.

Standard Features

The 3375 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 Reference* manual.

The 3375 does not support record overflow, sometimes referred to as track overflow.

Command Retry

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 condition 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

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 3375 to write a data area consisting of 32 bytes of zeros.

Multiple Track Operations

Multiple track operations allow 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.

Caution must be used in chaining multitrack CCW operations to a format write because head switching will not always occur due to loss of track orientation. (See explanation under the Format Write Release section.)

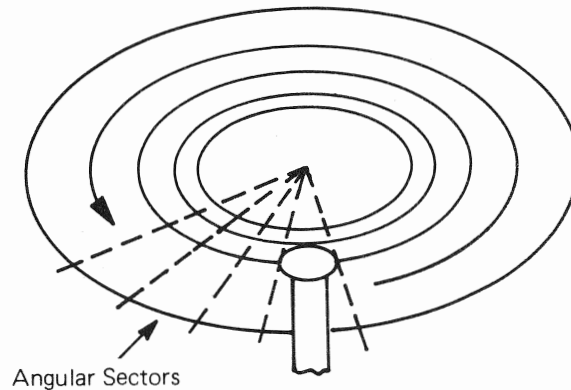
Format Write Release

Format write release (or write padding) frees the channel and storage director while the 3375 erases to the end of the track after a format write command (Write Count, Key, and Data; Write Record Zero; or Write Home Address).

Loss of track orientation may result when a channel fails to reconnect immediately following the pad complete device end interrupt for write padding completion.

Rotational Position Sensing

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.



Surface Defect Skipping

Surface defect skipping allows the 3375 to handle surface defects. The track remains intact for writing and reading data except for small areas. Surface defect skipping allows data to be stored before and after a surface defect. This averts having to assign an alternate track and the delays associated with repositioning of the read/write heads.

Data integrity is not affected by surface defect skipping, nor is the total access mechanism capacity.

Overlapping Operations

As with other disk storage, certain operations can be overlapped at the various access mechanisms in the 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 3375, 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 performing a set sector operation using rotational position sensing while the other access mechanism at the same drive is reading data.

In 3375 strings that have two controllers, the second controller provides a dual path to any disk storage attached. The D1 and A1 controller units can operate at the same time using two separate storage directors, but only one controller at a time can address a given head and

disk assembly. When one of the controllers is transferring data with one unit, the other can transfer data with any of the other three units in the four-unit string.

Device Status and Error Handling

Status conditions provide information to the system regarding the progress of an I/O operation at the 3375. 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 provide further information for system error recovery procedures and for diagnosing and isolating the cause of an error condition.

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 successful, a system message is issued to the operator. The system message includes status and sense information.

The *3880 Storage Control Reference* manual contains detailed descriptions of sense information and procedures that can be used for error recovery for the 3375. It 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 implemented by the system error recovery procedures.

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

Program Resources

Various program functions are available to interpret error 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 parameters of the appropriate define-the-file (DTF) macro for DOS/VSE and data control block (DCB) macro for VS1 and MVS. When a permanent 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 3375 includes a report on failure types to help the customer engineer in diagnosing a problem.

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

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

Error Handling Guidance

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

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

Programming Errors

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

Control Errors

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

Data Errors

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

Correction of Data Errors

Error detection and correction information is added by the 3375 to the home address of each track as well as to each count, key, and data area of a record when it is written. Later, when the home address or 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. The error correction code (ECC) generated by the 3375 can detect and correct single burst errors contained within two contiguous bytes.

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 that is not correctable with the ECC, 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 repeated attempts are made to read the data. 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, it is regarded as a temporary error. If a data error is not within the ECC correction range and the specified number of command retries prove unsuccessful, a permanent (or uncorrectable I/O) error condition exists. The 3375 then provides sense byte information that indicates the error condition and identifies the area (HA, count, key, or data field) of the record affected by the error.

Temporary and Permanent Data Error

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

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

Response to Data Checks

If a data check is caused by a defect on the disk surface, the Device Support Facilities program can be used to bypass the area. Defects are bypassed by skipping the defective area to use a displacement location on the same track. Up to seven defective areas on a track are skipped. Space for the defect 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 959 is reserved for assignment of up to 12 alternate tracks.

The Device Support Facilities program can also be used to rewrite home addresses, which eliminates the cause of certain data checks.

Backup Data

Data on an area or track that is to be bypassed may have to be restored from a backup copy. As described in the Backup and Recovery section, 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 reduce the effects of undetected errors which may have occurred during processing, but requires time and, therefore, may not be considered practical for routine use.

A read operation with the CCW skip bit on causes the data to be transferred back to the storage director instead of to main storage, and the storage director checks the data using the error correction code (ECC). With DOS/VSE, this validity check (or write check) can be done by specifying the CHECK macro and coding the appropriate DTF macro of the I/O control system (IOCS). With VS1 and MVS, this is done by specifying OPTCD=W in the DCB macro.

Equipment Errors

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

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

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

Service Procedures

The customer engineer performs diagnostic and problem determination procedures to identify the source of the failure and replaces the failing components, which are referred to as field-replaceable units (FRUs). Indicators in the sense information identify the 3880 storage director and the 3375 controller, thereby pinpointing the location of the hardware problem more precisely.

During installation of the 3375, the customer engineer must select a unique identification (ID) for the model A1 or model D1 controller. This controller ID, which must be different for each controller, is set with switches within the controller logic and is not associated with system addressing. The unique controller ID assists the customer engineer in identifying the unit at fault when errors occur by being presented in the sense information.

The customer engineer can service portions of the logic and controls of one unit without interrupting user operations on the remaining units in the 3375 string.

If a unit requires service, the customer engineer uses the IBM Maintenance Device (MD) to quickly isolate the cause of the problem. The MD is a portable microcomputer with a hand-held keyboard display, a diskette reader, and storage. It connects to the 3375 model A1 or D1 and performs diagnostic routines to help the customer engineer isolate problems and verifies that the unit is operating normally. The MD replaces the permanent CE panel used on previous IBM direct access storage devices.

Alternate Access

The 3375 model D1 provides an alternate path 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 3375, the dual controllers back up each other. Manual intervention is not required.

Statistical Information and Logging

The 3880 storage director maintains counters for its attached 3375s. 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 and made available in the EREP report.

Temporary seek and data checks are counted by the 3880 storage director error rate counters. (Error rate for seek checks 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. Subsequent checks

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 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. Information about the nature and location of the error is issued in a message to the operator and documented in the EREP report.

Hardware Planning Considerations

Careful planning will help to ensure the successful installation of a 3375. An installation planning representative can assist with the planning, and should be notified early in the planning stage of the installation.

Physical Planning

The following information highlights some of the significant physical planning considerations for installing the 3375. Detailed physical planning and signal cable 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 3375 string requires considerably less power, in kilo Volt Amp (kVA), than a 3350 string. A reduced power requirement translates into a reduced air conditioning load, which in turn may also reduce power consumption.

Space Requirements

For equivalent floor space, a 3375 requires approximately 45 percent less space than the 3350. That is, for about 10 gigabytes of storage, the 3375 requires approximately 16 square meters (172 square feet), whereas the 3350 requires approximately 30 square meters (300 square feet). Figure 7 illustrates the floor space saving of the 3375 over 3340, 3340/3344, and 3350 disk storage. The calculations are based on full string configurations, including allowance for one-half of the service clearance dimensions.

Processor Attachment Requirements

The 3375 is attached by the 3880 Storage Control model 1, 2, or 4 to various processors as follows:

- Attachment by 3880 model 1, 2, or 4 to a 4331 model group 2 or 4341 requires a block multiplexer that

supports the 3375 data rate.

- Attachment by 3880 model 1 or 2 to S/370 model 145, 145-3, 148, 155-II, 158, 158-3, 165-II, 168, or 168-3 requires the speed matching buffer feature.
- Attachment by 3880 model 1 or 2 to 3031, 3032, 3033, or 3042 model 2 requires data streaming or a speed matching buffer feature.
- Attachment by 3880 model 1 or 2 to 3081, 3083, or 3084 is via any block multiplexer channel.

Channel Switches

Channel switch features for attaching up to four channels are supported by the 3375 attachment.

Two- and four-channel switches are available as special features on the 3880 Storage Control. These features can be used as described and as shown in Figure 8.

- With the two-channel switch feature, the storage director of the 3880-4 can connect to the same or different channels.
- With the 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 the 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.

Channel switching is important for the 3375 with the dual controllers. In this configuration, data in the 3375 can be accessed by as many as eight different channels. This is possible if one 3375 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 two-channel switch pair, additional (four-channel switch) feature.


Model	Capacity (GB)	Units Required	Floor Space Required *		MB of Storage	
			Square Meter	Square Feet	Per Square Meter	Per Square Feet
3340	10.08	72	128.6	1,384.0	78.4	7.28
3344	9.1	20	41.0	445.6	219.8	20.4
3350	10.16	16	27.6	297.1	366.1	34.2
3375	9.84	12	15.7	163.4	626.2	60.2

*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.

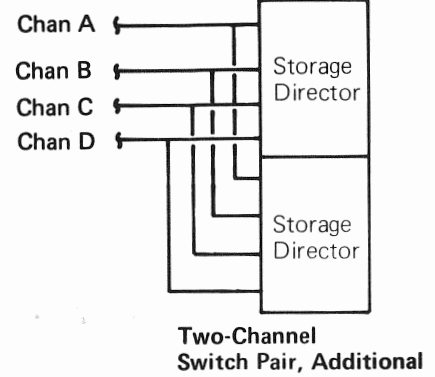
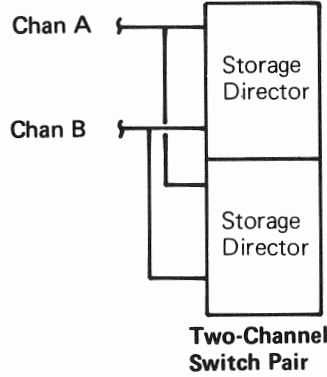
Figure 7. Comparisons of Approximate Floor Space Requirements Based on Full String Configurations

Storage Directors Connected to the Same Channels

3880 Model 4




3880 Models 1 and 2

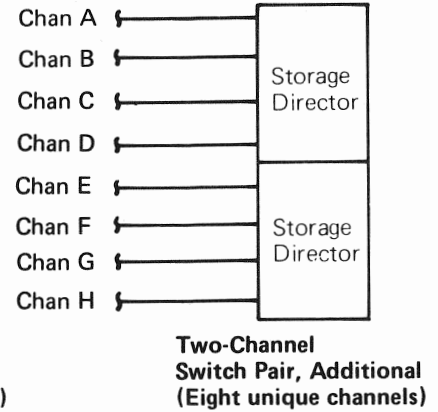
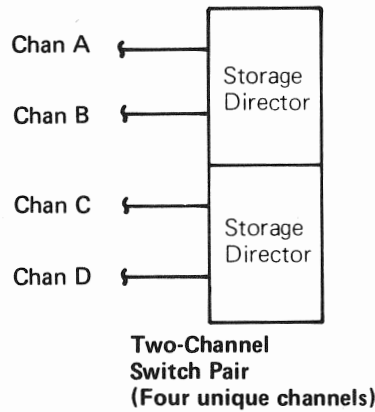


Storage Directors Connected to Different Channels

3880 Model 4



3880 Models 1 and 2



Storage Directors Connected to the Same and Different Channels

3880 Models 1 and 2





Figure 8. Channel Switching for 3880 Storage Directors Supported by 3375 Disk Storage Attachment

Addressing

Device Address

The device address used for the 3880 and 3375 addressing is an 8-bit binary number and is described as follows and in Figure 9.

Bits 0, 1, and 2 are the 3880 storage director address bits. If these three bits refer to an unattached storage director, select out is propagated.

A 3880 storage director can have up to eight different addresses, one for each of the channels to which it is connected. The addresses are set manually by the customer engineer.

Bits 3 and 4 are the 3375 controller (model A1 or D1) address. If these bits refer to a controller that is unattached or without power, select out is propagated.

Addresses of controllers are set with switches by the customer engineer when the 3375 is installed.

Bits 5 through 7 are the access mechanism address. If an addressed mechanism is unattached, without power, or in service mode, unit check is presented in initial status. The access mechanism addresses are 0 through 7. For a given controller, access mechanism addresses start with 0 or 8 and are sequential.

The table in Figure 10 lists the address ranges for 8, 16, and 32 access mechanisms connected to the same storage director. Addresses must begin at the first address of an address range. For example, with two strings attached to the same 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. Two strings, of eight devices each and connected to different storage directors, can be assigned non-adjacent address ranges like 20 through 27 and 40 through 47.

If less than the 8 or 16 access mechanisms are used, all of the addresses within the range are still reserved. For example, if a string of four devices is attached to a storage director and assigned addresses E0 through E3, addresses E4 through E7 are not available for assignment to another set of devices on the same channel. Nor can the unused addresses be used for another string of devices on the same storage director.




Figure 9. 3880 and 3375 Device Addressing

Address Ranges			
One String of 8 Access Mechanisms or Less	Two Strings of 16 Access Mechanisms or Less	Four Strings of 32 Access Mechanisms or Less	
00-07	80-87	00-0F	00-1F
08-0F	88-8F	10-1F	20-3F
10-17	90-97	20-2F	40-5F
18-1F	98-9F	30-3F	60-7F
20-27	A0-A7	40-4F	80-9F
28-2F	A8-AF	50-5F	A0-BF
30-37	B0-B7	60-6F	C0-DF
38-3F	B8-BF	70-7F	E0-FF
40-47	C0-C7	80-8F	
48-4F	C8-CF	90-9F	
50-57	D0-D7	A0-AF	
58-5F	D8-DF	B0-BF	
60-67	E0-E7	C0-CF	
68-6F	E8-EF	D0-DF	
70-77	F0-F7	E0-EF	
78-7F	F8-FF	F0-FF	

Figure 10. Sets of Valid 3880/3375 Addresses

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 in current disk storage devices.) If the seek address is not valid, unit check is presented in ending status.

The cylinder addresses are:

Data cylinders	0-958
Alternate cylinder	959

The head addresses are: 0-11

Address and Attachment Configurations

The 3375 attaches to the 3880 in strings consisting of one model A1 and up to three model B1s. With the dual controller, the model D1 replaces the last attached model B1 of a full string to provide a dual path. A full string provides eight separately addressable access mechanisms. Up to four strings can attach to each storage director of a 3880 Model 1 or to one storage director of a 3880 model 2.

Figure 11 illustrates an example attachment configuration of four 3375 strings and typical address assignments. Figure 12 shows the possible configurations between 3880 and 3375 dual controller strings. The configurations shown are typical ways to attach the devices. Consult an installation planning representative for details.




Figure 11. Example of Four 3375 Model A1 Strings Attached to Storage Directors of 3880-1




Figure 12. Possible Configurations Between 3880 and 3375 Dual Controller Strings

Alternate Path Address Restrictions

There are restrictions on the addresses that can be assigned when processors have alternate paths to the access mechanisms of the 3375. These restrictions are due to attachment hardware limitations. Another restriction is imposed by the operating system other than 370-XA.

The attachment hardware does not allow the 3880 and 3375 address combination to be used more than once per physical channel. The operating system restricts the alternate path address so the only part of the I/O address that can vary between alternate paths to a 3375 is the

channel part of the I/O address. Figure 13 illustrates an unsupported configuration. If, as shown, one of the paths is designated address A28 (where A is the channel address), the address of the other path cannot be A28 because of the attachment hardware restriction, but must be X28 (where X is any address but A) because of restrictions of operating systems other than 3370-XA.

Figure 14 shows a supported alternate path configuration where path 1/A is processor 1, channel A; path 1/B is processor 1, channel B; path 2/A is processor 2, channel 1; and path 2/B is processor 2, channel B.

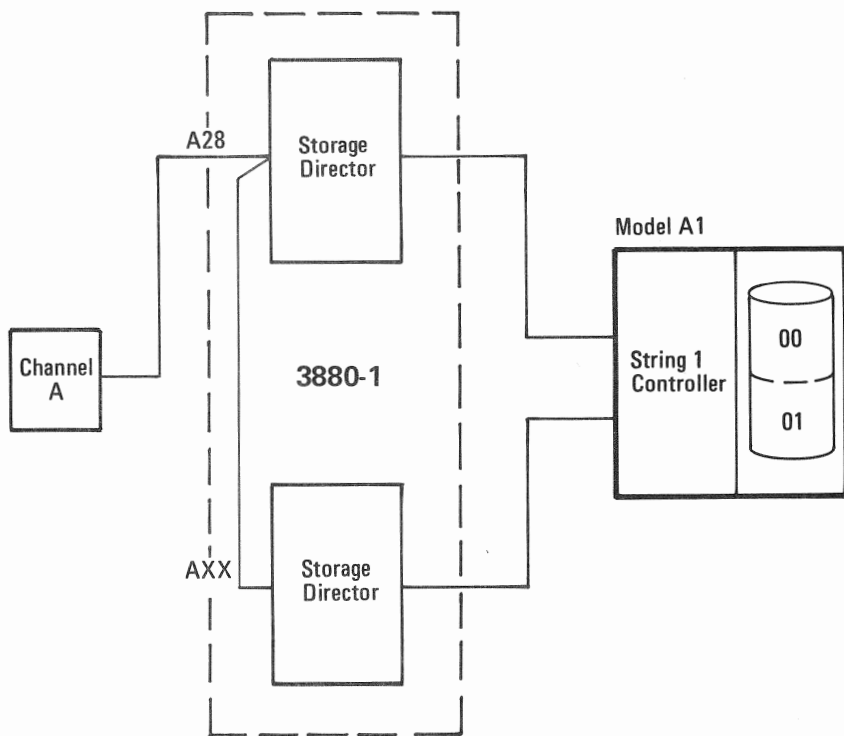


Figure 13. Unsupported Processor/Channel Address Designation

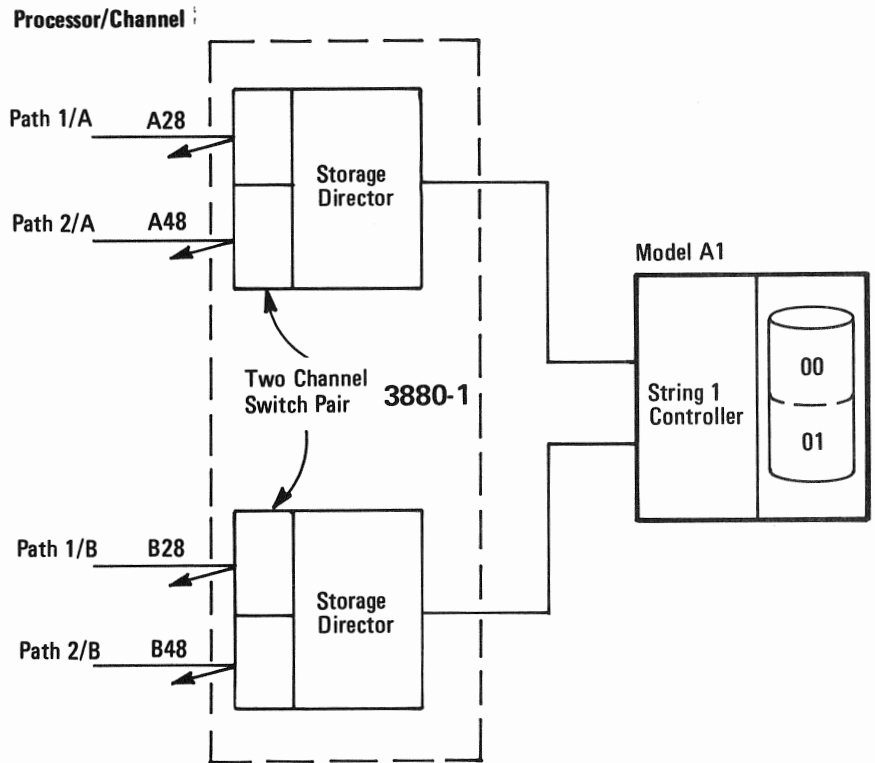


Figure 14. Supported Processor/Channel Address Configuration

Program Planning Considerations

The 3375 is supported under the following base system control programs: DOS/VSE, VM/SP, OS/VS1, OS/VS2 MVS, and MVS/XA. A collection of required and optional programs are available to support 3375 operations in these environments. The optional programs, some of which have standalone versions, are described in this section.

Before the 3375 is installed, the correct version and release of the base system control program should be in place and all of the associated 3375 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 3375 are listed in the front of the manual. Also there are program directories included with the program distribution.

DOS/VSE Support

The Disk Operating System/Virtual Storage Extended (DOS/VSE) supports all capabilities (except SMB) of the 3375. For initialization and maintenance of the 3375, Device Support Facilities is also required.

VM/SP Support

The 3375 is supported by Virtual Machine/System Product (VS/SP) program product. VM/SP incorporates earlier system extensions. VM/370 Release 6 is a prerequisite for VM/SP.

VM/SP supports all models of the 3375 and all of the processors that can connect to 3375s. It does not support use of the new commands associated with device reservation.

Device Support Facilities, standalone version, can be used for initial program loading (IPL) in a virtual machine and used for initialization and maintenance of the 3375.

OS/VS1 Support

The 3375 is supported under OS/VS1 with the following system control program and program products:

- OS/VS1, Release 7 Basic Program Extensions
- Data Facility Device Support, Release 1
- Device Support Facilities, Release 4

The programs support all models of the 3375 and all of the processors that can connect the 3375. They support the full functional capabilities of the 3375.

The Data Facility Device Support program allows the 3375 to take advantage of the improved volume table of contents (VTOC) structure for data. The 3375 is also supported by the DASD Migration Aid, which helps ease the process of conversion from current DASD. See Additional Programming Support for details.

MVS/SP Support

The 3375 is supported under OS/VS2 MVS with a combination of system control programs and program products. The programs support all models of the 3375 and all of the processors that can connect to 3375s. They support the full functional capabilities of the 3375.

OS/VS2 MVS support is under Release 3.8. In addition to the base system control programs in Release 3.8, support for the 3375 requires the following:

- System Product – JES 2 or JES 3 Release 1, or System Product – JES 2 or JES 3 Release 2, or System Product – JES 2 or JES 3 Release 3 (MVS/SP-JES2, MVS/SP-JES 3).
- Data Facility Device Support (DFDS) program product.
- Device Support Facilities, Release 4.

MVS/XA DFP and MVS/SP Version 2

MVS/XA Data Facility Product and MVS/System Product Version 2 is based on Release 3.8 of OS/VS2 MVS. Version 2 of MVS/SP requires the following:

- An IBM 3081 Processor Complex operating in extended architecture mode.
- MVS/XA Data Facility Product (5665-284).
- MVS Assembler H Version 2 (5668-962).

The MVS/XA Data Facility Product is a program product that contains data management support, device support, program library management support, and utility functions. Assembler H is required for installing and maintaining MVS/SP Version 2.

MVS/370 DFP and SP 1.3.0

The MVS/370 Data Facility Product and System Product 1.3.0 are program products that contain data management support, device support, program library management support, and utility functions for the 3375.

Programs Using OS/VS and the 3375

OS/VS program products using SAM, ISAM, PAM, DAM, or VSAM unmodified and not requiring specifications of device characteristics also operate on the 3375.

Additional Programming Support

Device Support Facilities

The Device Support Facilities program, used for initialization and maintenance of direct access storage devices, now includes support for the 3375. It has both system and standalone versions.

The device Support Facilities is the vehicle for performing the following functions:

- Initialize 3375 volumes for use by the operating system. Initialization includes creating the contents of cylinder 0, track 0 by writing IPL bootstrap and volume label records; and allocating space for a VTOC. If so specified, all tracks on the volume can be checked for defects, and skip displacements or alternates can be assigned if needed.
- Build an index for the VTOC, if Data Facility Device Support is installed.
- Reformat a previously initialized volume.
- Inspect selected tracks of a previously initialized volume for possible surface defects, and skip displacements or alternates can be assigned if needed.
- Perform tests to help determine the source of a problem, and, as a consequence, decide on appropriate action. The tests are invoked with the ANALYZE command as discussed in the Device Status and Error Handling section.

Support for the 3375 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.

Program Products Not Supported By DOS

The following program products support and aid 3375 use depending on the installation. These programs do not run under DOS/VSE.

Data Facility Device Support

The Data Facility Device Support (DFDS) program product provides data management support for the 3375 operating under VS1 or MVS.

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 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.)

Data Facility Extended Function

The Data Facility Extended Function (DFEF) program is a catalog facility that can be used under 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 DFEF 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 (DFDSS) program provides dump and restore functions for the 3375. 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 data.
- 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, DFDSS 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 allocated by absolute track.

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

Support for 3375 is not provided by IEHDASDR, IBCDMPRS, DRWDASDR, and DRUDMPRS utilities, because the functions are incorporated in Data Set Services and other programs.

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 3375. It operates with OS/VS and supports all data sets created by the standard access methods. It can:

- Collect information about current disk data sets.

- Calculate data set space requirements for the 3375.
- Identify non-VSAM data sets that may be candidates for reblocking and help to evaluate the appropriateness of block sizes for the new device.
- Generate, and place in the procedures library, customized JCL and utility statements for migrating the data sets to the 3375. (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 3375.

Hierarchical Storage Manager

The Hierarchical Storage Manager (HSM) is a space management tool under MVS 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.

Data Management and Data Set Conversion

This section discusses data management and conversion for users of DOS/VSE and OS/VS. Most of the access methods that are supported by the operating systems can be used for data sets stored on the 3375.

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 file definition record with the appropriate parameters of the DTF or DCB macros.

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 I/O instructions.)

Most of the programming required to store data sets on the 3375 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) 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 statements instead of specified with system macros (DTF or DCB). The 3375 is fully supported by all applicable JCL statements.

Data Set Conversion

The programming changes required to convert existing data sets to the 3375 are primarily those that reflect changes in the 3375 device characteristics, such as track capacity and number of tracks per cylinder.

Data sets can be copied, or moved, to the 3375 using IBM supplied utilities. The utility used depends on the characteristics of the data set and the functions desired. The DOS/VSE, OS/VS1, and MVS utilities manuals, listed in the bibliography, describe the various utilities and provide information on the control specifications needed.

Data sets can be converted directly from existing disk storage, such as the 3340 or 3350, to the 3375. Or, 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 3375 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.

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 3375.

SAM data sets can be converted to the 3375 using:

- COPY or FAST COPY DISK for DOS/VSE
- IEHMOVE or IEBGENER for OS/VS1 or MVS

PAM data sets can be converted by using IEHMOVE or IEBCOPY for OS/VS1 or MVS. PAM is not supported by DOS/VSE.

ISAM data sets can be converted by using IEBISAM for OS/VS1 or MVS. For DOS/VSE, ISAM data sets can be converted to the 3375 by first converting them to VSAM. Then the ISAM data set can be accessed with the ISAM/VSAM interface. Probably ISAM data sets will not require new programs, because the user's existing reorganization programs suffice. Minor changes to the load programs will be required because of the 3375 device characteristics. Changes such as space allocations can be made with JCL statements.

DAM data sets can be converted by using:

- COPY for DOS/VSE
- IEHMOVE for OS/VS1 or MVS

If more records are to be placed on a 3375 track than are on a source device track, the existing reorganization program can transfer data to the 3375. It may be necessary to re-program 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 (IDCAMS) REPRO command.

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

Subject to the restrictions on device dependencies and those indicated for DAM data sets, existing 3330, 3340/3344, or 3350 programs can be used without change to process data sets transferred to the 3375.

User-Written Access

User-written programs that issue the EXCP macro or I/O instructions directly require modifications to reflect the different 3375 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 3340, 3340/3344, and 3350 channel command word (CCW) lists operate on 3375 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 3375 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.

Conversion Under VM/SP

Current CMS command (COPYFILE, MOVEFILE, TAPE, TAPPDS) are available to move data to 3375 devices from other disk storage types. System areas on CP OWNED volumes must be reconstructed on the 3375. The VM/SP DASD dump/restore utility can be used to move data on a cylinder basis.

Effective Use of the 3375

This section discusses some of the factors involved in deciding which data sets to place on the 3375 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 3375 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 3375.

Criteria for Using the 3375

Assuming the requirement for fast, direct-access storage, one of the criteria for deciding which data sets to place on the 3375 will be performance needs. The 3375 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 3375.

While performance needs may be the primary criterion for placing data sets on the early installations, as more 3375s 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 3375 where a balance is needed for those data sets requiring high performance.

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 considered.

Optimum Blocking for 3375 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 overhead. 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 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 READKEY parameters of the DOS/VSE DTFDA macro (or BLKSIZE AND KEYLEN parameters of the OS/VS DCB macro).

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 3375, data is recorded in 32-byte segments. This is a consideration when deciding on block sizes because small block sizes are vulnerable to an accumulation of residual bytes. For example, a 993-byte block occupies as much space as a 1,024-byte block.

See the section on Data Organization and Capacity for more specific information.

Performance Considerations

For the most part, the principles that provide good performance for other DASD installations also apply to the 3375. There are two other factors that benefit 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.

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. Chained scheduling provides for a succession of read or write operations by synchronizing I/O operations and use of buffer space.

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 480). In general, this placement reduces the maximum length of seeks outside the data set.
 - 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 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.
3. Minimize the time required for a given access.
 - 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 dual controllers are available.
 - 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 or fixed length record, standard format 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 record as normally done. For fixed standard records, 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.
 - Use a smaller block size for data sets where direct access predominates.
 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 of fixed-length records, standard format, as described in item 3.
 - Where applicable, use an indexed VTOC (if Data Facility Device Support is installed) 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.

5. Reduce processor interference for I/O operations.

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

Application Considerations

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

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

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

Therefore, block sizes that result in satisfactory performance in a batch processing environment may not be desirable in an online, response-oriented system, where it may be preferable to sacrifice device space for a gain in

performance. The large storage capacity available with the 3375 helps to make this a reasonable trade-off, but in addition, the fast 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 3375 access mechanism. This reduces the possible contention for use of the access mechanism that could be caused by the paging subsystem.
- Place the pageable link pack area (PLPA) modules on the fastest available device.
- Place most of the local page data sets on moderate speed devices to avoid overloading the high-speed devices if the number of the faster devices (such as the 3375) is limited.
- Set up a relatively slow-speed device for data sets that 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 detected, and to permit some reasonable increase in the number of concurrent virtual I/O (VIO) data sets.

Conversion Checklist

Conversion from existing types of storage to the 3375 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 3375. 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 3375 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 3375 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 3375 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 3375 support.

- Reorder if necessary.
- Define alternatives if the 3375 is not supported.

Review application programs.

- EXCP programs using 3340s or 3350s may have to be modified (including CCWs, sense checks, and others).
- 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 is required for 3375 support. For main storage requirements for individual functional components, see the *Storage Estimates Manual*, listed in the bibliography.
- Additional main storage is required if block sizes are increased to take advantage of the 3375 higher capacity tracks.

Review device addressing requirements.

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

Review system parameter and BLDL lists.

- Parameter library and build lists should be reviewed and revised, because many modules may have been split or recommended for addition to residency lists in the new release. See *Storage Estimates Manual*, listed in the bibliography.

Initialize 3375 volumes.

- The 3375 volumes are initialized with a home address and track descriptor record before shipment. At installation, the Device Support Facilities program INIT command can be used to initialize the volume at a *minimal level*. This initialization includes placing three labels as R1, R2, and R3. The first label in the VTOC contains a format 4 -DASD label describing the VTOC itself. The second label is reserved, but is not used by DOS/VSE. The third label is the volume label at cylinder 0, track 0.

Examine performance considerations.

- Consider the performance guidelines as described in the Effective Use of the 3375 section.
- Consider the increment of the capacity. The 3375 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 DTF or DCB parameters. Are record format and blocksize parameters specified in user programs (Assembler, COBOL, PL/1)? If so, the specifications may need to be changed.
- Consider the logical record format. All fixed block (QSAM and BSAM) data sets should be examined to see if the fixed length standard 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 3375 device parameters during system generation will aid greatly here.

- Review the procedures library to ensure optimum space allocation, block sizes, and unit separation. 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 3375 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 Effective Use of the 3375 section, 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 and after processing, store it on the 3375 devices. On the second day, use data from the 3375 and after processing, store the updated data on the 3375.

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, (DUMP for DOS/VSE, or DFDSS for OS/VS).

Update operating procedures.

- Modify machine room procedures to reflect installation of the 3375 and programming changes and additions.
- Plan to familiarize operator with new machine room procedures.

Operator Controls and Indicators

Two types of controls are found on a 3375 string: individual controls and indicators for each unit, and power controls and indicators that affect the entire string of up to four units.

Unit Operator Panel

Each unit has an operator panel (Figure 15) mounted on the front of the top surface. The controls and indicators are on this panel. The six square indentations on each half of the panel are used for unit address labels.

Ready Indicators

The green Ready lamps indicate when each associated access mechanism in the unit is on, track-following, and ready to operate.

Ready Switches

The Ready switch prepares the drive for use. When the Ready switch is pressed and released, the selected drive read/write heads are moved to track zero, the head and cylinder address registers are reset, and an attention signal is sent to the storage control. However, no action occurs on either access mechanism of the drive unless both are inactive.

Read/Write or Read Only Switches

The Read/Write or Read Only switches inhibit writing by the associated access mechanism when set to Read Only. When set to Read/Write, any legitimate write operation may be performed on the associated access mechanism.

The switch position can be changed at any time, but the state of the device does not change until the current operation is completed. This prevents accidental destruction of the data. When the switch is set to Read Only, any write command is rejected. The sense information sent to the storage director indicates that the writing was inhibited and the command was rejected.

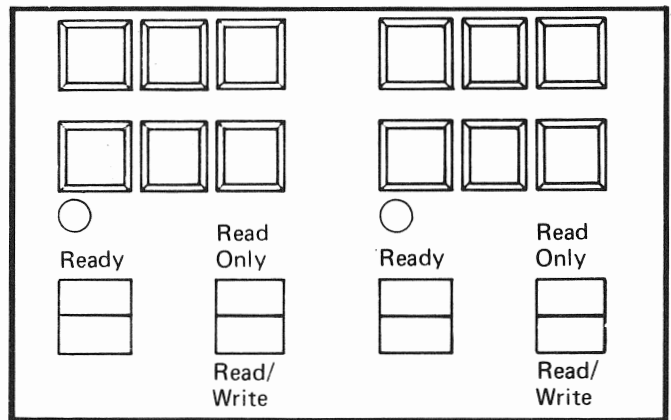


Figure 15. 3375 Unit Operator Panel

Controller Operator Panel

The controller operator panel shown in Figure 16 is mounted on the front top surface of the model A1 and model D1 controllers.

Delayed Power Off/Enable Switch

The Delayed Power Off/Enable switch is used to power on and power off the string. When the switch is set to Delayed Power Off, there is about a 12-second delay before power is removed from the string.

Check Indicator

The Check indicator comes on when the controller detects a check condition.

Power On Indicator

The Power On indicator is on when power is present in the controllers. It is activated by setting the Delayed Power Off/Enable switch to Enable. The string power sequencing operation then powers up each unit in turn.

Data Path A and B Enable/Disable Switches

If a 3375 string has the optional string switch feature, the A and B Enable/Disable switches are also located on the controller operator panel. These switches allow either storage director path interface to be enabled or disabled from the string. An internal interlock prevents the switch status from changing until the current operation is complete.

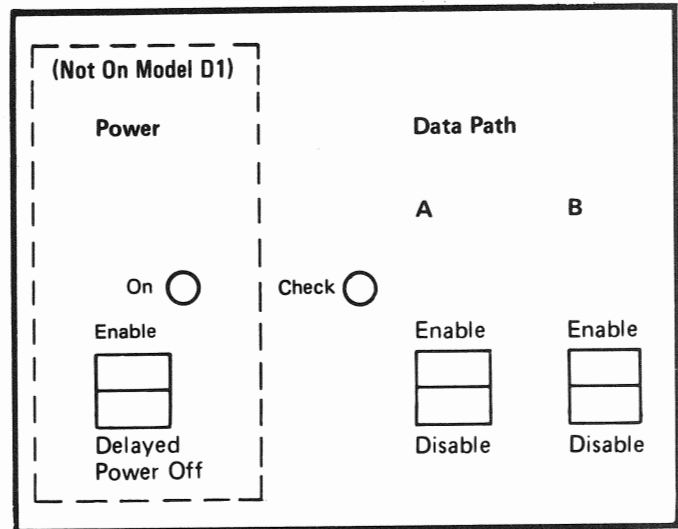


Figure 16. 3375 Controller Operator Panel

Glossary

AP-1	Analysis Program-1	KEYLEN	key length (DCB parameter)
BDAM	basic direct access method	KL	key length
BLKSIZE	block size (DCB parameter)	kVA	kilo volt amp
BSAM	basic sequential access method	LRECL	logical record length (DCB parameter)
CAW	channel address word	MB	megabyte
CCW	channel command word	MD	IBM Maintenance Device
chan	channel	MVS	multiple virtual storage
CKD	count, key, and data	OBR	outboard recorder
CSW	channel status word	OPTCD	optional services code (DCB parameter)
DAM	direct access method	OS/VS	operating system/virtual storage
DASD	direct access storage device	PAM	partitioned access method
DCB	data control block	PCI	program-controlled interrupt
DFDS	Data Facility Device Support program	PDS	partitioned data set
DFDSS	Data Facility Data Set Services program	PLPA	pageable link pack area
DFEF	Data Facility Extended Function program	PROCLIB	procedures library
DISP	data set disposition (DCB parameter)	PSW	program status word
DL	data length	PTF	program temporary fix
DOS	disk operating system	QISAM	queued indexed sequential access method
DP	data processing	QSAM	queued sequential access methods
DTF	define the file	R0	record zero
ECC	error correction code	RECFM	record format (DCB parameter)
EREP	Environmental Recording, Editing, and Printing program	RPS	rotational position sensing
ERP	error recovery procedures	SAM	sequential access method
EXCP	execute channel program	SIP	system initialization program
F	fixed-length records	SIO	start input/output
FBS	fixed-length, blocked, standard records	SLI	suppress-length indicator
FRU	field-replaceable unit	SP	system product
GB	gigabytes	SYNAD	synchronous error routine address (DCB parameter)
HA	home address	TIC	transfer-in-channel
HDA	head and disk assembly	TSO	time-sharing option
HSM	hierarchical storage manager	UCB	unit control block
ID	identification	VIO	virtual input/output
INIT	initialize	VM	virtual memory
I/O	input/output	VS	virtual storage
IOCS	input/output control system	VSAM	virtual storage access method
IPL	initial program load	VSE	virtual storage extended
ISAM	indexed sequential access method	VTOC	volume table of contents
JCL	job control language		
JES	job entry subsystem		

Bibliography

This bibliography includes all publications used in compiling this book, prerequisite publications, and other related and useful publications not referred to in text.

- *IBM 3880 Storage Control Models 1, 2, 3, and 4 Reference Manual*, Order No. GA-1661, gives detailed descriptions of channel commands, and status and sense information.
- *IBM Disk Storage Management Guide*, Order No. GBOF 1205, describes data and equipment errors that can occur in disk storage operations. It also provides tutorial information on the Environmental Recording, Editing, and Printing (EREP) program System Exception reports.
- *IBM Reference Summary for Handling Data Checks*, Order No. GX26-1669, contains information extracted from the *IBM Disk Storage Management Guide, Error Handling* manual, Order No. GA26-1672.
- *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 3375.
- *Introduction to DOS/VSE*, Order No. GC33-5370, contains a general description of DOS/VSE and the basic computing services available.
- *DOS/VSE Macro User's Guide*, Order No. GC24-5139, is for programmers using DOS/VSE macro instructions.
- *DOS/VSE Macro Reference*, Order No. GC24-5140, is a quick reference on data management and system control macros for experienced programmers.
- *DOS/VSE Data Management*, Order No. GC24-5138, discusses the data management concepts basic to any DOS/VSE access method. It explains the concepts of how information (data) can be defined, organized, and identified by the DOS/VSE user for processing by the system.
- *DOS/VSE System Utilities*, Order No. GC3305381, provides a full description of the use of the DOS/VSE System Utility Programs for application programmers.
- *OS/VS1 Access Method Services*, Order No. GC26-3840, describes the service program commands used to manipulate VSAM data sets.
- *OS/VS1 Data Management For System Product*, Order No. GC26-3837, 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/VS1 Direct Access Device Space Management*, Order No. SY26-3837, describes in detail each DADSM routine and its function.
- *OS/VS1 Storage Estimates Manual*, Order No. GC28-0604, gives information for estimating main storage requirements.
- *OS/VS2 Access Method Services*, Order No. GC26-3841, describes the service program commands used to create and manage VSAM data sets.
- *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 System Product General Information Manual*, Order No. GC28-1025.
- *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 3375.
- *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 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.
- *MVS/Extended Architecture Data Facility Product: General Information*, Order No. GC26-4007, contains general information about MVS/Extended Architecture (MVS/XA) Data Facility Product.

- *MVS/System Product Version 2 General Information Manual*, Order No. GC28-1118, gives an overview and planning information for Version 2 of either MVS/SP-JES2 (Program No. 5740-XC6) or MVS/SP-JES3 (Program No. 5665-2910). It also describes machine and programming requirements, storage requirements, and compatibility considerations.
- *IBM Virtual Machine Facility/370 Introduction*, Order No. GC20-1800, contains a general description of VM/370 and the basic programs available. It also references the associated publications.
- *VM/SP General Information Manual*, Order No. GC20-1838.
- *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, restore, and defragmentation functions.
- *Data Facility Extended Function: General Information*, Order No. GC26-3960.
- *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.

Index

A

access
 capacity 1
 mechanism 1
 time 1
 user-written 36
address
 (*see also*, seek address) 27
 device 26
 ranges 26
 restrictions 30
 seek 27
 table 26
addressing
 (*see also*, seek address) 27
 access mechanism 26
 configurations 27
 device 26
 for the 3880 and 3375 26
 of controllers 26
 ranges 26
 restrictions 30
 seek 27
 strings
 four 26
 one 26
 two 26
 table 26
alternate path 30
attachment
 by 3880 model 4 24
 by 3880 models 1 and 2 24
 by 3880 models 1, 2, and 4 24
 channel 1
 processor requirements 24
 requirements 1
 systems 1

B

block sizes 8
blocking aid 9

C

capacity
 access mechanism 1
 cylinder 1
 device overhead 7
 drive 1
 formulas
 record 7
 track 7
 physical records
 equal length 7
 unequal length 7
 tables
 with keys 12
 without keys 10
 track 1
 unit 1
CE panel 22
chained scheduling 1
chaining
 command 16

 data 16
channel
 attachment 1
 commands 14
 switching
 two-channel 24
 two-channel switch pair 24
 two-channel switch pair, additional 24
channel command words (CCWs) 16
channel commands (*see* individual commands) 14
checklist, conversion 40
command retry 18
command sequence 16
commands, channel
 control 14
 count area 6
 data area 6
 description 6
 diagnostic 15
 key area 6
 read 15
 search 14
 sense 15
 write 15
configurations
 dual controller strings 2
control errors 21
controller 2
controller operator panel
 (*see* operator panel) 43
controls
 (*see* individual controls) 44
 (*see* operator panel) 43
conversion
 checklist 40
 data set 35
 from other direct access storage 34
 under VM/SP 36
cylinder and head number (CCHH) 5

D

DAM data sets 35
data
 backup 22
 block sizes 8
 compressed 34
 errors 21
 formats
 record 5
 track 5
 management 35
 organization 5
 verification by program 22
data errors
 ECC correctable 21
 ECC uncorrectable 21
 permanent 21
 response 21
 temporary 21
Data Facility Data Set Services (DFDSS) 33
Data Facility Device Support (DFDS) 33
Data Facility Extended Function (DFEF) 33
data management 35
data set conversion 35
data transfer rate 1
Define Extent 16

Device Release 15
Device Reserve 15
Device Support Facilities 33
diagnostic
 Read Home Address 15
 Sense/Read 15
 Write Home Address 15
Diagnostic Load 14
Diagnostic Sense 14
Diagnostic Write 14
Direct Access Storage Device Migration Aid 34
DL data length 7
DOS/VSE support 32
dual controller 2

E

end-of-file 18
Environmental Recording, Editing, and Printing (EREP) 20
equipment checks 22
error
 alternate path 22
 control 21
 correction of 21
 data 21
 detection 21
 ECC correctable 21
 ECC uncorrectable 21
 equipment 22
 handling 20
 logging 23
 programming 21
 statistical information 23
error correction code (ECC) 21
error handling guidance 20
error recover procedures (ERPs) 20

F

features 2,18
 channel switch 24
 speed matching buffer for 3375 2
 standard
 command retry 18
 end-of-file 18
 format write release 18
 multiple track operations 18
 overlapping operations 18
 rotation position sensing 18
 surface defect skipping 18
 string switch 2
field-replaceable units (FRUs) 22
floor space required
format write release 18

H

head and disk assembly (HDA)
 addressing 26
 alternate path 30
 description 5
 initial data recorded 5
 servo information 5
Hierarchical Storage Manager (HSM) 34
home address (HA)
 home 5
 reading 5
 writing 5

I

index point 5
indicators
 (see individual indicators) 44
 (see operator panel) 43
ISAM data sets 35

J

JCL statements 34
Job Control Language (JCL) 35

K

KL key length 7

L

Locate Record 16

M

Maintenance Device (MD)
 diskette reader 22
 keyboard display 22
 uses 22
models, 3375
 A1 2
 B1 2
 D1 2
multiple track operations 18
MVS/SP support 32
MVS/XA DFP and MVS/SP Version 2 32
MVS/370 DFP and SP 1.3.0 33

N

No Operation (No-Op) 14

O

operator panel, unit
 controls
 Data Path A and B Enable/Disable switches 44
 Delayed Power Off/Enable switch 44
 Read Only switches 43
 Read/Write switches 43
 Ready switches 43
 indicators
 Check 44
 Power On 44
 Ready 43
OS/VS1 support 32

P

page data sets 39
PAM data sets 35
performance
 considerations 38
 guidelines 37
physical planning 24
planning considerations
 air conditioning 24

- application 39
- hardware 24
- page data sets 39
- performance 38
- power 24
- program 32
- space requirements 24
- power requirement 24
- program-controlled interrupt (PCI) 17
- programming errors 21
- programming support 2,17
 - DOS/VSE 32
 - MVS/XA 32
 - OS/VS1 32
 - OS/VS2 MVS 32
 - VM/SP 32

R

- Read and Reset Buffered Log 15
- Read commands
 - Count 15
 - Count, Key, and Data 15
 - Data 15
 - Home Address 15
 - Initial Program Load 15
 - Key and Data 15
 - Multiple Count, Key, and Data 15
 - Record Zero (R0) 15
 - Sector 15
- Recalibrate 14
- record format 6
- record overflow 18
- Restore 14
- rotation position sensing (RPS) 14,18
 - description 18
 - Read Sector 15
 - Set Sector 14
- rotational delay 1

S

- SAM data 35
- Search commands
 - Home Address Equal 14
 - Identifier Equal 14
 - Identifier Equal or High 15
 - Identifier High 15
 - Key Equal 15
 - Key Equal or High 15
 - Key High 15
- search direct 1
- Seek 14
- seek address
 - cylinder 27
 - head 27
- Seek Cylinder 14
- Seek Head 14
- sense bytes 20
- Sense commands
 - Device Release 15
 - Device Reserve 15
 - ID 15
 - Read and Reset Buffered Log 15
 - Unconditional Reserve 15
- Sense I/O 14
- Sense I/O Type 14
- Sense ID 15
- service clearance 24
- Set File Mask 14

- Set Sector 14
- Space Count 14
- speed matching buffer for 3375 2
- speed matching buffer in 3880
 - description 2
 - read operations 16
 - write operations 16
- standard access methods
 - DAM 35
 - ISAM 35
 - PAM 35
 - SAM 35
 - VSAM 35
- statistical counters
 - bytes read 23
 - bytes searched 23
 - command overruns 23
 - data overruns 23
 - seeks 23
- storage, (see capacity) 1
- string switch
 - description 2
- suppress-length indicator (SLI) flag bit set 17
- system control programs (SCPs) 2

T

- track descriptor record
 - IBM programming 5
 - Record 0 (R0)
 - reading 5
 - writing 5
- track format
 - description 5
 - home address 5
 - index point 5
 - physical records
 - equal length 7
 - unequal length 7
 - track descriptor record 5
- track overflow 18
- Transfer-in-Channel 14
- two-channel switch feature 24
- two-channel switch pair feature 24
- two-channel switch pair, additional feature 24

U

- Unconditional Reserve 15
- unit operator panel 43
- user-written programs 36

V

- VM/SP support 32
- VSAM data 34
- VSAM data sets 35

W

- Write commands
 - Count, Key, and Data 15
 - Count, Key, and Data Next Track 16
 - Data 15
 - Erase 15
 - Key and Data 15
 - Record Zero (R0) 15
 - Update Data 16

Update Key and Data 16
Write Special Count, Key, and Data 14

Numeric

3031, 3032, 3033 1
3042-2 1
3081, 3083, or 3084 1
3340 1
3340/3344 1
3350 1
3375

access mechanism capacity 9
attachment 1
channel commands 14
characteristics
 performance 1
 physical 1
configurations 2
data management 35
data organization 5

data set conversion 35
device status 20
effective use of
 space and performance guidelines 37
 space on a track 37
 unit of data 37
error handling 20
floor space required 24
models 2
performance considerations 38
planning considerations
 hardware 24
 physical planning 24
 program 32
speed matching buffer 16
standard features 18
statistical information 23
3880 1
4331 1
4341 1

Order No. GA26-1666-2

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

You may use this form to communicate your comments about this publication, its organization, or subject matter, with the understanding that IBM may use or distribute whatever information you supply in any way it believes appropriate without incurring any obligation to you.

Your comments will be sent to the author's department for whatever review and action, if any, is deemed appropriate. Comments may be written in your own language; use of English is not required.

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

Possible topics for comment are:

Clarity Accuracy Completeness Organization Coding Retrieval Legibility

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

Note: Staples can cause problems with automated mail sorting equipment.
Please use pressure sensitive or other gummed tape to seal this form.

Cut Along This Line

What is your occupation ? _____

Number of latest Newsletter associated with this publication: _____

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

Reader's Comment Form

Cut or Fold Along Line

IBM 3375 Direct Access Storage Description and User's Guide (File No. S/370-07, 4300-07) Printed in U.S.A. GA26-1666-2


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/H14
555 Bailey Avenue
San Jose, California 95150

Fold and Tape

Please Do Not Staple

Fold and Tape



International Business Machines Corporation
Data Processing Division
1133 Westchester Avenue, White Plains, N.Y. 10604

IBM World Trade Americas/Far East Corporation
Town of Mount Pleasant, Route 9, North Tarrytown, N.Y., U.S.A. 10591

IBM World Trade Europe/Middle East/Africa Corporation
360 Hamilton Avenue, White Plains, N.Y., U.S.A. 10601

Order No. GA26-1666-2

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

You may use this form to communicate your comments about this publication, its organization, or subject matter, with the understanding that IBM may use or distribute whatever information you supply in any way it believes appropriate without incurring any obligation to you.

Your comments will be sent to the author's department for whatever review and action, if any, is deemed appropriate. Comments may be written in your own language; use of English is not required.

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

Possible topics for comment are:

Clarity Accuracy Completeness Organization Coding Retrieval Legibility

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

Note: Staples can cause problems with automated mail sorting equipment.
Please use pressure sensitive or other gummed tape to seal this form.

Cut Along This Line

What is your occupation ? _____

Number of latest Newsletter associated with this publication: _____

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

Reader's Comment Form

IBM 3375 Direct Access Storage Description and User's Guide (File No. S/370-07, 4300-07) Printed in U.S.A. GA26-1666-2

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/H14
555 Bailey Avenue
San Jose, California 95150

Fold and Tape

Please Do Not Staple

Fold and Tape



International Business Machines Corporation
Data Processing Division
1133 Westchester Avenue, White Plains, N.Y. 10604

IBM World Trade Americas/Far East Corporation
Town of Mount Pleasant, Route 9, North Tarrytown, N.Y., U.S.A. 10591

IBM World Trade Europe/Middle East/Africa Corporation
360 Hamilton Avenue, White Plains, N.Y., U.S.A. 10601

Order No. GA26-1666-2

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

You may use this form to communicate your comments about this publication, its organization, or subject matter, with the understanding that IBM may use or distribute whatever information you supply in any way it believes appropriate without incurring any obligation to you.

Your comments will be sent to the author's department for whatever review and action, if any, is deemed appropriate. Comments may be written in your own language; use of English is not required.

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

Possible topics for comment are:

Clarity Accuracy Completeness Organization Coding Retrieval Legibility

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

Note: Staples can cause problems with automated mail sorting equipment. Please use pressure sensitive or other gummed tape to seal this form.

Cut Along This Line

What is your occupation ? _____

Number of latest Newsletter associated with this publication: _____

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


Reader's Comment Form

IBM 3375 Direct Access Storage Description and User's Guide (File No. S/370-07, 4300-07) Printed in U.S.A. GA26-1666-2


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/H14
555 Bailey Avenue
San Jose, California 95150

Fold and Tape

Please Do Not Staple

Fold and Tape



International Business Machines Corporation
Data Processing Division
1133 Westchester Avenue, White Plains, N.Y. 10604

IBM World Trade Americas/Far East Corporation
Town of Mount Pleasant, Route 9, North Tarrytown, N.Y., U.S.A. 10591

IBM World Trade Europe/Middle East/Africa Corporation
360 Hamilton Avenue, White Plains, N.Y., U.S.A. 10601

Publication Number
GA26-1666-2

File Number
S/370-07, 4300-07

Printed in
USA



GA26-1666-2

