IBM Disk Storage Management Guide

Background Reference Information

Cross-System

IBM

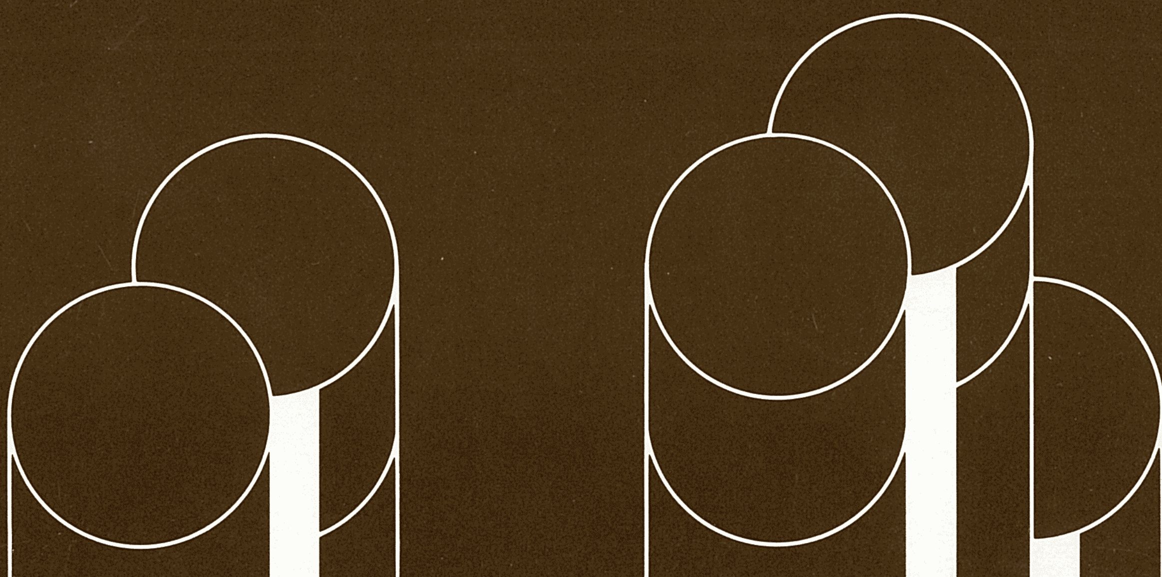

IBM Disk Storage Management Guide

Background Reference Information

Cross-System

Publication Number
GA26-1675-0

File Number
S370/4300-07

Any reference to an IBM program product in this document is not intended to state or imply that only IBM's program product may be used. Any functionally equivalent program may be used instead.

First Edition (December 1982)

This is the first edition of the *IBM Disk Storage Management Guide,* Publication No. GBOF 1205, *Background Reference Information,* Publication No. GA26-1675.

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

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

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

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

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

Preface

The *IBM Disk Storage Management Guide* provides tutorial information and guidance to help you with the physical management of your disk storage.

The information is intended for IBM disk storage customers. It is written especially for operations personnel, including operators and system programmers and others who may have responsibility for disk storage management tasks.

The material is in separate manuals, organized as a set. Each manual has a separate publication (order) number, but the manuals as a set have a common Bill of Form Publication (Order) Number, GBOF 1205. Manuals in the *IBM Disk Storage Management Guide* are:

Background Reference Information	GA26-1675
Error Handling	GA26-1672

The material in the guide applies to the following products.

Disk Storage: 3330, 3340, 3344, 3350, 3370, 3375, 3380

Storage Control: 3830-2, 3880-1, 2, 3, 11, 13

Operating Systems: OS/MVS, OS/VS1, DOS

IBM Disk Storage Management Guide, Background Reference Information

This manual contains reference information that may be needed as background for disk storage management. It is written at an introductory level.

Although much of the information appears in other forms in other publications, the information on errors will be new to most readers.

Table of Contents

Disk Storage Subsystem and Transfer Path

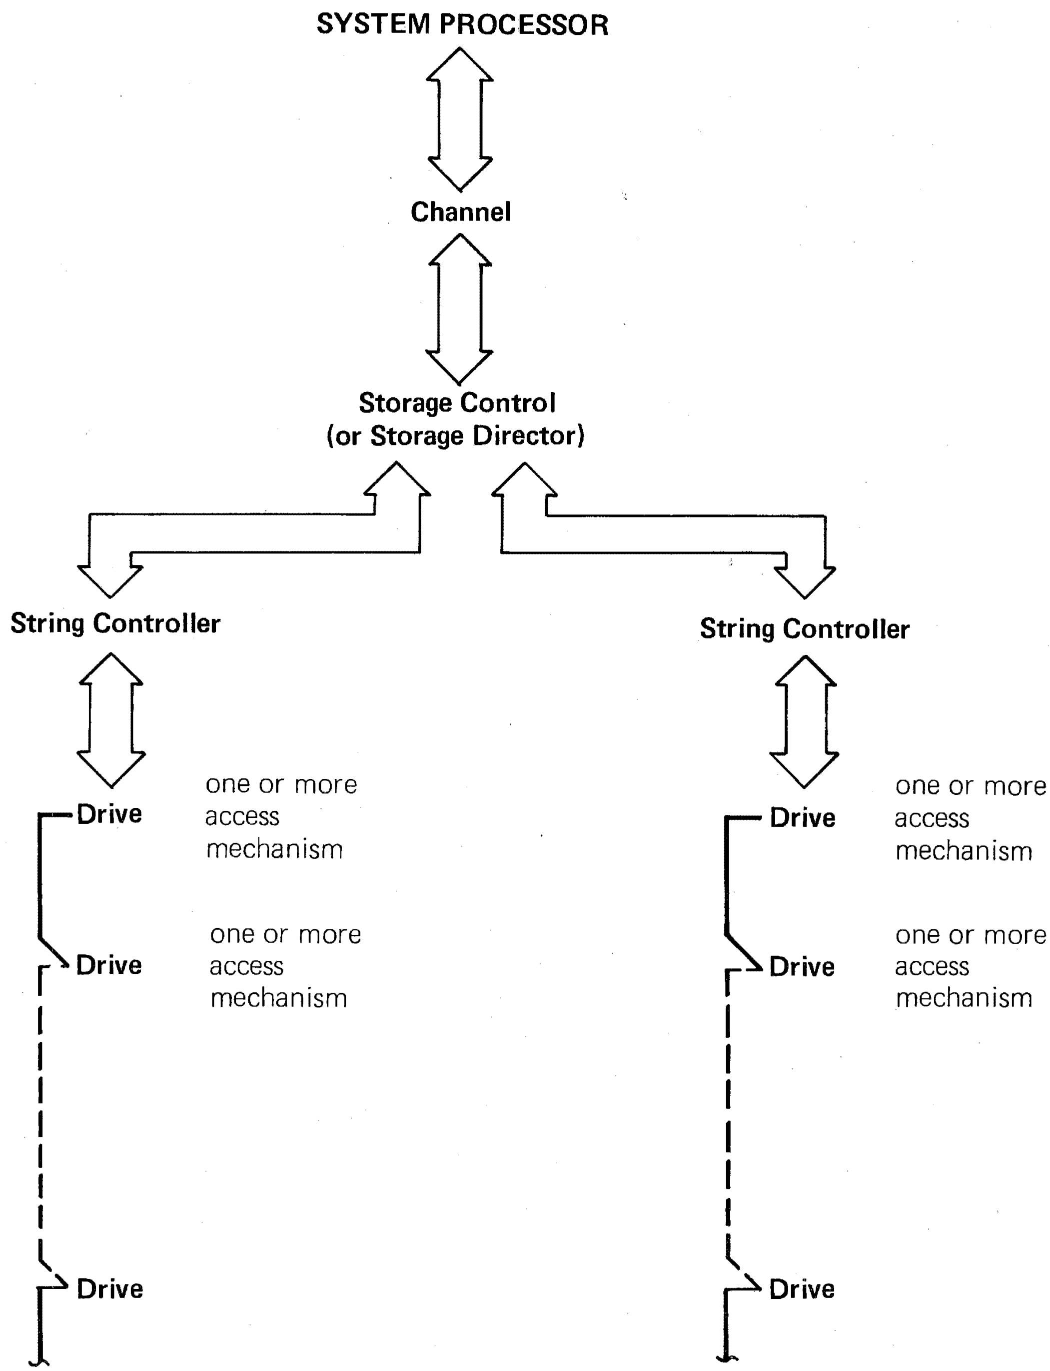

Disk Storage Subsystem

Although many readers are familiar with the general characteristics of disk storage, the characteristics are reviewed here to establish a common understanding of terminology and to reflect current technology. The subjects covered are common to all disk storage devices. Emphasis is on information that helps in understanding the tasks of disk storage management. Information is included on error detection and the recovery actions performed by the system and subsystem. The section, "Error Descriptions," includes definitions of errors that all readers will need for using the *Error Handling* manual, Publication (Order) Number GA26-1672.

Disk Storage Subsystem Components

Disk storage consists of a drive, spindle, and disks, with one or two access mechanisms. The drive controls rotation of the disks, which are stacked on the spindle. An access mechanism controls access to the tracks and carries read/write heads. A disk storage subsystem consists of disk storage, controllers, and storage control.

One or more disk storage devices can be connected in a string. The number of devices that can be in a string depends on the disk storage type, and a string may be made up of less than the maximum allowable devices.

A string of disk storage devices has one or two controllers. For example, a 3350 may have two controllers, where one operates as an alternate. A 3375 and 3380 may have two controllers, both of which are active controllers of the string of devices. The 3370 has one controller per string.

One or more strings of disk storage, each with one or two controllers, can be connected to a storage control. (A 3880 Storage Control has two storage directors, each of which has all of the functions of an independent storage control.) These components of a disk storage subsystem are shown on the facing page. They are packaged in different ways in unit enclosures.

Transfer Path

A channel, storage control, and string controller form a transfer path between the processor storage and disk storage. A transfer path is used to transfer:

- Control information from the system to the subsystem
- User data to and from disk storage
- Status and sense information from the subsystem and channel to the system

In most cases, a disk storage device can be accessed over different possible transfer paths, made up of different components. Different paths are possible because:

- Different operating systems and different processors can share the same disk storage resources, and a processor can have one or more channels for possible access to the same disk storage.
- Multiple channels of the same or different processors can connect to the same storage control. Channel switches in the storage control are used to control which of the connected channels are allowed to select the storage control. The number of channel switches in a storage control depends on the storage control type.

- Two storage controls can connect to the same string of disk storage. String switches in the string controller are used to control which of the connected storage controls are allowed to select the string. When there are two strings, each can have a string switch. Not all disk storage types have string switching.

A channel, storage control, and string controller can be active with only one transfer at a time.

Figure 1 shows two processors with paths to the same two strings of disk storage:

> Processor I has two possible paths. It can access either string a or b through channel 2 and storage control A or through channel 4 and storage control B.
>
> Processor II has only one path. It can access either string a or b through channel 3 and storage control A.

In this configuration, transfers can be in progress to each of the two strings at the same time in the following ways.

> Processor I, channel 2, storage control A and Processor I, channel 4, storage control B
>
> Or
>
> Processor I, channel 4, storage control B and Processor II, channel 3, storage control A

In addition to the connections described, the storage controls in the illustration have other channel switches that can be connected to other channels.

This flexibility of configurations permits access from different requestor locations, and may provide alternate components in the event the originally requested component is not available. A requested component in a path may not be available because it is busy, disabled by an equipment malfunction, or deliberately removed from availability for system selection.

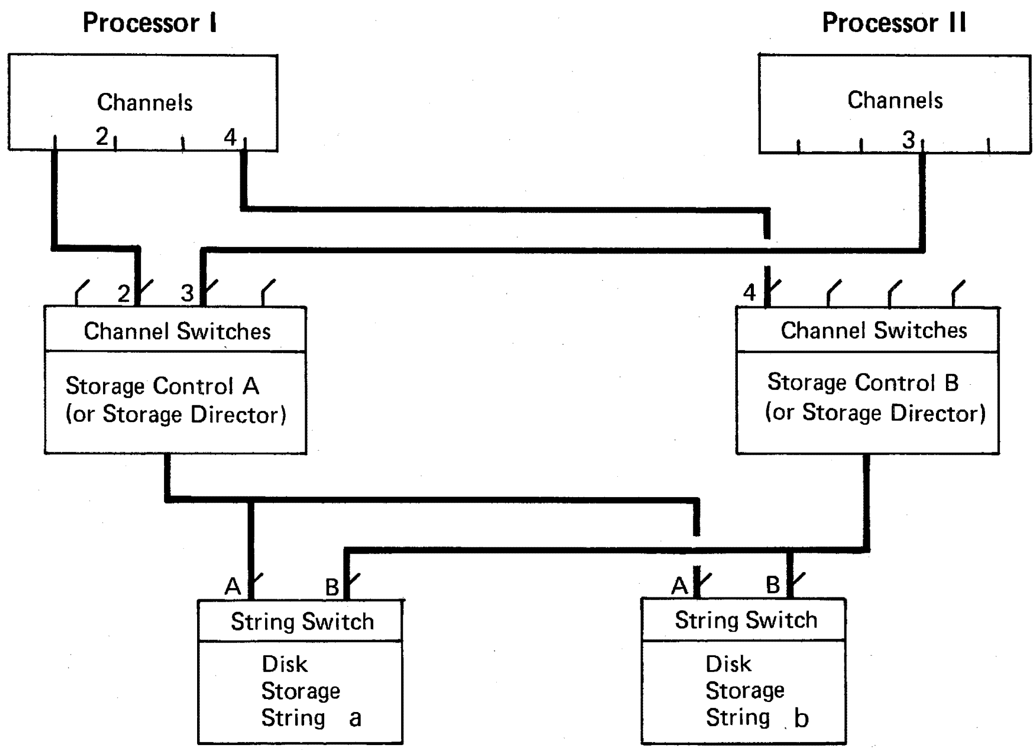

Figure 1. Example of Possible Channel and String Switches

Control of Availability

Disk storage and the path components are made available for system selection or removed from availability by system operator commands, or in some cases, by switches at a separate controlling console. Commands for OS/VS are the VARY Online or Offline commands; for DOS they are the DEVICE Up or Down commands.

When channel or string switches are installed, there are enable and disable switches on the storage control and disk storage unit control panels or at a remote control panel. There is a switch for each channel that can potentially be switched to the storage controls or for each storage control that can potentially be switched to the string. A switch must be set in the enable position to allow the storage control or string controller to be made available for selection by system console commands. More than one switch can be in the enable position but only one at a time can be selected and active. Use of system operator commands to control availability of a disk storage device also depends on the manual setting of switches at unit control panels.

Request and Response

Traditionally the same path (through the same components) is used to send a response from the disk storage as was used to send the request to disk storage. Recently it has become possible to respond over the first available channel of the processor, if there is the required system support. (For example, this has been announced as dynamic reconnection for a 3380 Model AA attached to a 3081 processor, using System 370/Extended Architecture.)

Device Reserve

A disk storage device can be reserved for the exclusive use of one channel or a group of channels. Reservation and release are controlled by channel commands.

Simultaneous Transfer

Some of the newer subsystems permit transfer over two paths to the same string at the same time, as long as the transfer is to different devices in the string. This is possible for 3375 and 3380 disk storage that have two controllers (attached to two different 3880 storage directors and channels). In these cases, two paths can simultaneously perform transfer type operations.

Overlap of Transfer with Control Operation

For most of the disk storage types covered in this manual, a transfer over a path may overlap with certain control operations at a device, such as locating a position on a track. In this case, only one transfer path is involved. Once the control operation is requested over the path, the device carries out the requested process independently, leaving the path free for a transfer operation to or from a different disk storage in the string. The control operations that can be carried out independent of the channel and storage control are a seek to a track address and a set sector for a particular track position (this is accomplished with a technique known as rotational position sensing).

References to Bibliography

For information on possible attachments and configurations, refer to the disk storage and storage control manuals. These manuals include information on channel switches, string switches, and dual controllers (if a disk storage has two controllers).

For information on system operator commands, refer to the manuals on commands for operations control for the applicable operating system.

Identifiers

Two types of identifiers are needed for disk storage:

- One type for I/O selection and response. These identifiers are the means used by the system for accessing a disk storage device and components that make up a transfer path to the device. They are assigned based on the conventions of the particular disk storage type and the operating system, and physically set in the hardware.
- One type for service. These identifiers are the means used by service representatives and operations personnel to identify a device or path component that needs attention. The particular type of service identifier and method of assignment depends on the device type.

Besides these two types of identifiers, some systems provide a path group identifier for multiple channels used by a single system.

I/O Selection

The disk storage and each of the components that can make up a transfer path are assigned identifiers so they can be selected for an input or output (I/O) operation.

With System 370 architecture, identifiers for I/O selection are carried as a 2-byte address in a start I/O instruction.

Unit Address

An installation assigns identifiers that can be used by its operating system for I/O selection. Each possible configuration of identifiers makes up a *unit address*. Unit addresses are established for system use at system generation.

A unit address is 2 bytes, expressed as 4 hexadecimal digits. The first of the four digits is always zero. The second digit identifies a channel. The next two digits identify a storage control, controller, and device. For example, with a unit address of 0217, the 2 refers to a channel; the 1 and 7 refer to a storage control, controller, and device. The device portion of the address refers to the ultimate destination – the drive and access mechanism. If there are two access mechanisms per drive, each has a device address.

A disk storage device may be accessed through many different channels and storage controls; consequently, a given disk storage may be selected by different unit addresses.

Different operating systems may impose certain constraints on addressing although they are not required by the hardware. For example, under OS/MVS, a single operating system must use the same identifiers to select a disk storage, except that the channel identifiers can differ. When selection of a disk storage is controlled from different operating systems, each different system may have different channel and storage control identifiers or they can be the same.

Although a given disk storage device may be selected by unit addresses with different channel and storage control identifiers, the device portion of the address is always the same for a given device. This is because each access mechanism in a string of disk storage devices has a sequentially numbered address that is physically set in the hardware. These are referred to as *physical addresses* and they cannot be changed by the user.

There are three exceptions to addressing as described. A 3330 address can be moved by an operator, and the moved address is called a logical address. A 3350, when in 3330 emulation mode, has two addresses per access mechanism because it stores the equivalent amount of data of two 3330s. These are called logical addresses. A 3344 has four addresses per access mechanism for storing the equivalent amount of data of four 3340s. These are called logical addresses. In these cases, a *logical address* gives access to a volume of data.

Other than the exceptions described, physical and logical addresses are the same, and the address gives access to a volume of data.

For each disk storage type, there are ranges of addresses that are acceptable for assignment to storage controls, controllers, and devices for one or two strings.

Addresses used for I/O selection are reported in system console messages to the operator and are included in various usage and error reports.

Meaning of Addressing Bits

The bits of the 2-byte address have the following meaning.

Left-most Byte: the first 4 bits are always zeros. The next 4 bits identify a channel.

Right-most Byte: the eight bits identify a storage control, controller, and device (drive and access mechanism). Bits 0 and 1 are always for a storage control and bits 5, 6, and 7 are always for a device. The meaning of bits 2, 3, and 4 may differ depending on the device type and interpretations made by the storage control.

References to Bibliography

For the specific meaning of each bit of the right-most 2-byte address, refer to the disk storage manuals. These manuals also include the ranges of addresses that are valid for assignment for one and two strings.

Service

An identifier used for service is either a *physical identifier* (ID) or a physical address, depending on the device type. A physical ID is 3-bytes long. A physical address is 2-bytes long, as described in the previous section.

The assignment of physical identifiers (IDs) is fairly new and not all devices have them. They are used to identify subsystem components involved in an error situation. The 3 bytes are expressed as 6-hexadecimal digits, where 2 digits identify the storage control, 2 digits the controller, and 2 digits the access mechanism. (The channel is not included in a physical ID.) When physical identifiers are assigned, they are set with switches at installation, and a label with the identifier is affixed to the outside of the unit, where it can be seen. They are included in sense information and reported in system console messages to the operator and in certain error reports.

(In addition to the use in error handling, a physical ID may be helpful in configuration management. For instance, the physical ID may be used to communicate the identification of a particular string that is to be disabled at a controlling console.)

When a physical identifier is not assigned, the physical I/O address is used as a service identifier. The device portion of the address is included in sense information, but addresses of the channel and storage control used to access the device are not included. However, these are all stored in the system log and are given in certain error reports.

Physical Identifiers

For 3880 storage controls and 3375 and 3380 disk storage devices, physical IDs are assigned.

A 2-hexadecimal digit physical identifier (ID) is assigned to each storage director and controller. These identifiers, unlike the I/O address, are unique within the installation. In addition to the assigned storage director and controller physical IDs, the device (access mechanism) also has a physical ID. The device physical ID is the same as the access mechanism portion of the unit address. The first of the two hexadecimal digits of the device physical ID is always zero.

Assignment of physical IDs for a 3380 is shown in Figure 2. The 3375 assignment is made in the same way, but the device address does not use bit 4.

When a 3330, 3340, 3350, or 3370 disk storage is attached to a 3880 Storage Control, the storage director to which the device is attached has an assigned physical ID, although there are none for the controller and device.

Identifiers for Other Storage Controls and Disk Storage

When no physical ID is available, the physical address is used to identify a unit needing service or attention.

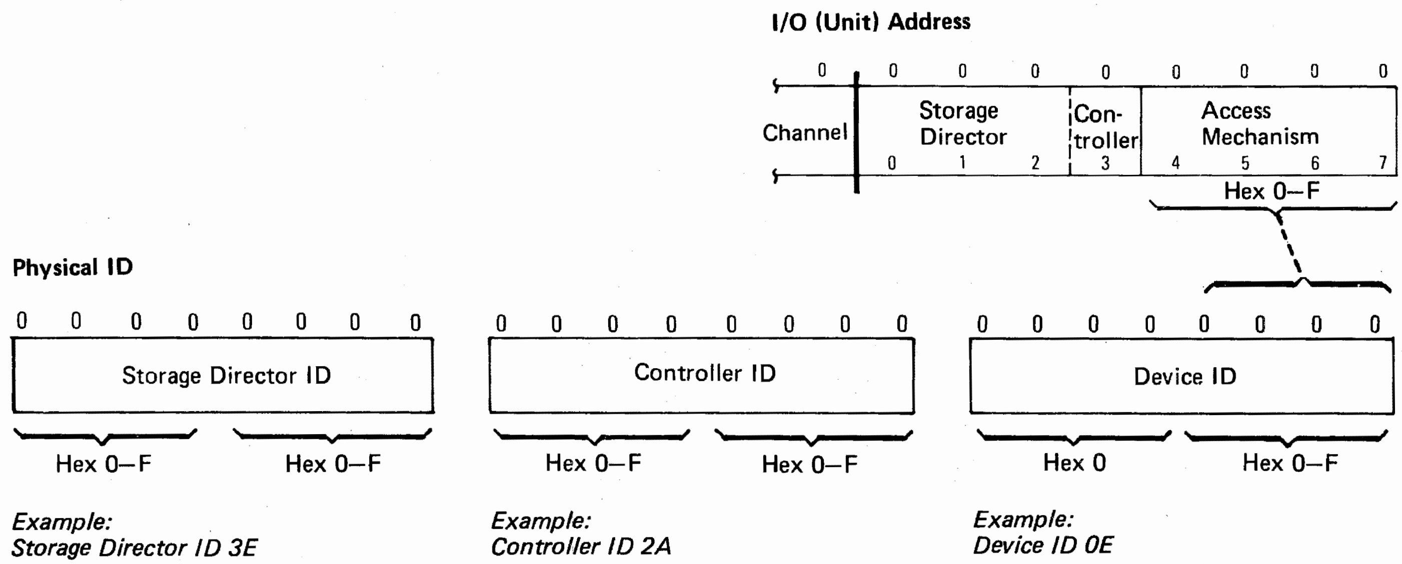

Figure 2. Physical ID Assignment for 3380

References to Bibliography

For information on assignment of physical IDs, refer to the disk storage manuals of devices that support them.

Disk Storage Physical Characteristics

Disks are stacked on a spindle (or shaft) with space for access arms to move back and forth between the disks. Heads, which are carried on the access arms, write and read data (bit patterns) on tracks around the disk surface. See Figure 3. During operations the heads do not contact the disk surface but, in effect, fly above it. Access arms with read/write heads can be positioned to each track. Both sides of a disk are used.

Tracks

Tracks are separate, concentric circles on the disk surface. The same amount of data (number of bytes) can be written on each track of a given device type. On the outermost tracks, the data is more spread out with more space between each bit, whereas on the innermost tracks, data is more tightly packed with less space between each bit. Different disk storage device types have different numbers of tracks per disk surface, and the amount of data that can be stored per track differs for each device type.

Access Mechanism

Access arms that service a set of addressable disk surfaces are attached to an access mechanism that moves all of the heads at the same time and for the same distance. Therefore, when the access mechanism is moved to a track on one disk surface, a track on each of the other surfaces is accessible without another movement of the access mechanism. Although there are multiple read/write heads on the access mechanism, only one head at a time transfers data.

Some types of disk storage have one access mechanism per drive and some have two access mechanisms per drive. When there are two access mechanisms per drive, each is addressable and operates independently at a different set of disk surfaces.

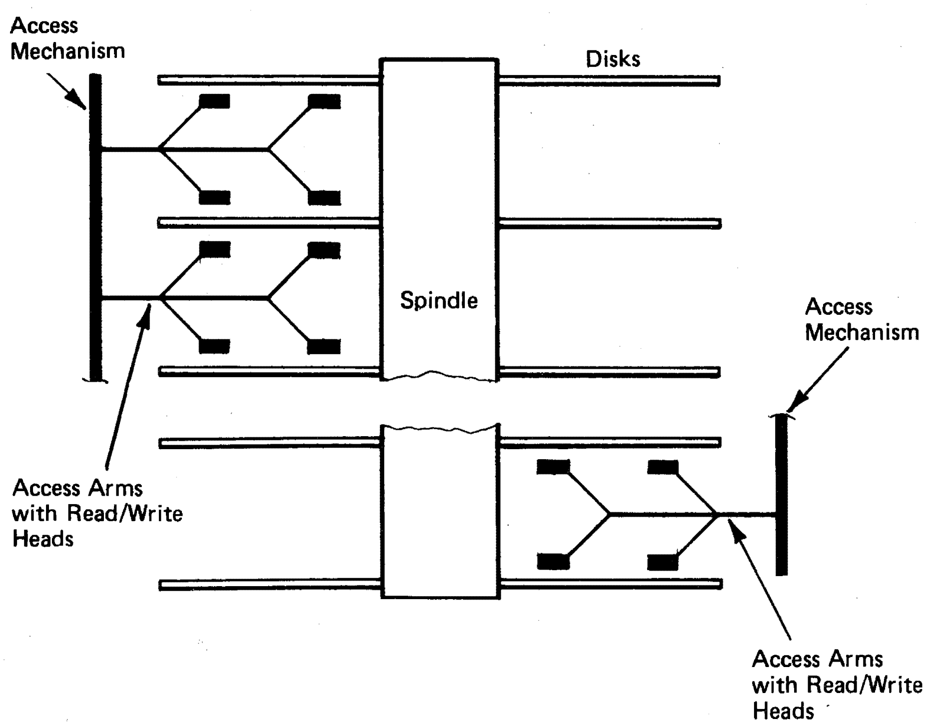

Figure 3. Disk and Access Mechanisms with Two Read/Write Heads per Surface

Cylinder

All tracks accessible by a particular positioning of the access mechanism are referred to as a cylinder. It is faster to switch electronically from one head to another to access tracks on different surfaces in the same cylinder than to mechanically move a head to each track on the same disk surface. When reading or writing in a continuous sequence, all of the data in a cylinder is read or written by switching from one head to the next.

Heads Per Disk Surface

The 3330 has one head per disk surface and one track per surface in a cylinder. The 3340, 3344, 3350, 3370, 3375, and 3380 have two heads per disk surface and two tracks, per surface, in a cylinder. The heads are spaced so that one head writes and reads the outer band of the tracks, and the other head writes and reads the inner band of the tracks, as shown in Figure 4.

Fixed Heads and 3880 Cache for Faster Access

Instead of heads that are moved by the access mechanism, some disk storage types have tracks with permanently positioned heads (with one head per track). These fixed heads provide faster access than movable heads. Fixed heads are available for certain tracks on models of the 3340, 3344, and 3350 disk storage.

Another means for improving access time is the cache storage contained in Models 11 and 13 of the 3880 Storage Control. This is electronic storage where a copy of frequently needed data can be stored for faster access and transfer to main storage. (The 3880 Model 11 attaches 3330 and 3350 disk storage. The 3880 Model 13 attaches 3380 disk storage. The Models 11 and 13 are used for different purposes.)

Maintenance Information

Certain areas on a device are reserved for maintenance information for use by the subsystem, hardware service representative, and maintenance programs. Most types of disk storage contain information that can be used to reconstruct an original track home address if it becomes unreadable. (Refer to the "Count, Key, and Data Format" section for an explanation of the home address.)

Packaging

Disks and read/write heads are packaged in different ways. The disk pack used on a 3330 and the data module used on the 3340 can be moved from one disk drive to another by the user. With the head and disk assemblies used on the 3344, 3350, 3370, 3375, and 3380, the disks and access mechanisms are sealed in an enclosure and permanently mounted on the disk drive. They are not removable, except by a service representative. These non-removable, sealed assemblies help prevent problems that can sometimes cause data errors.

References to Bibliography

For detailed information on the characteristics of a disk storage type, refer to the disk storage and storage control manuals in the bibliography.

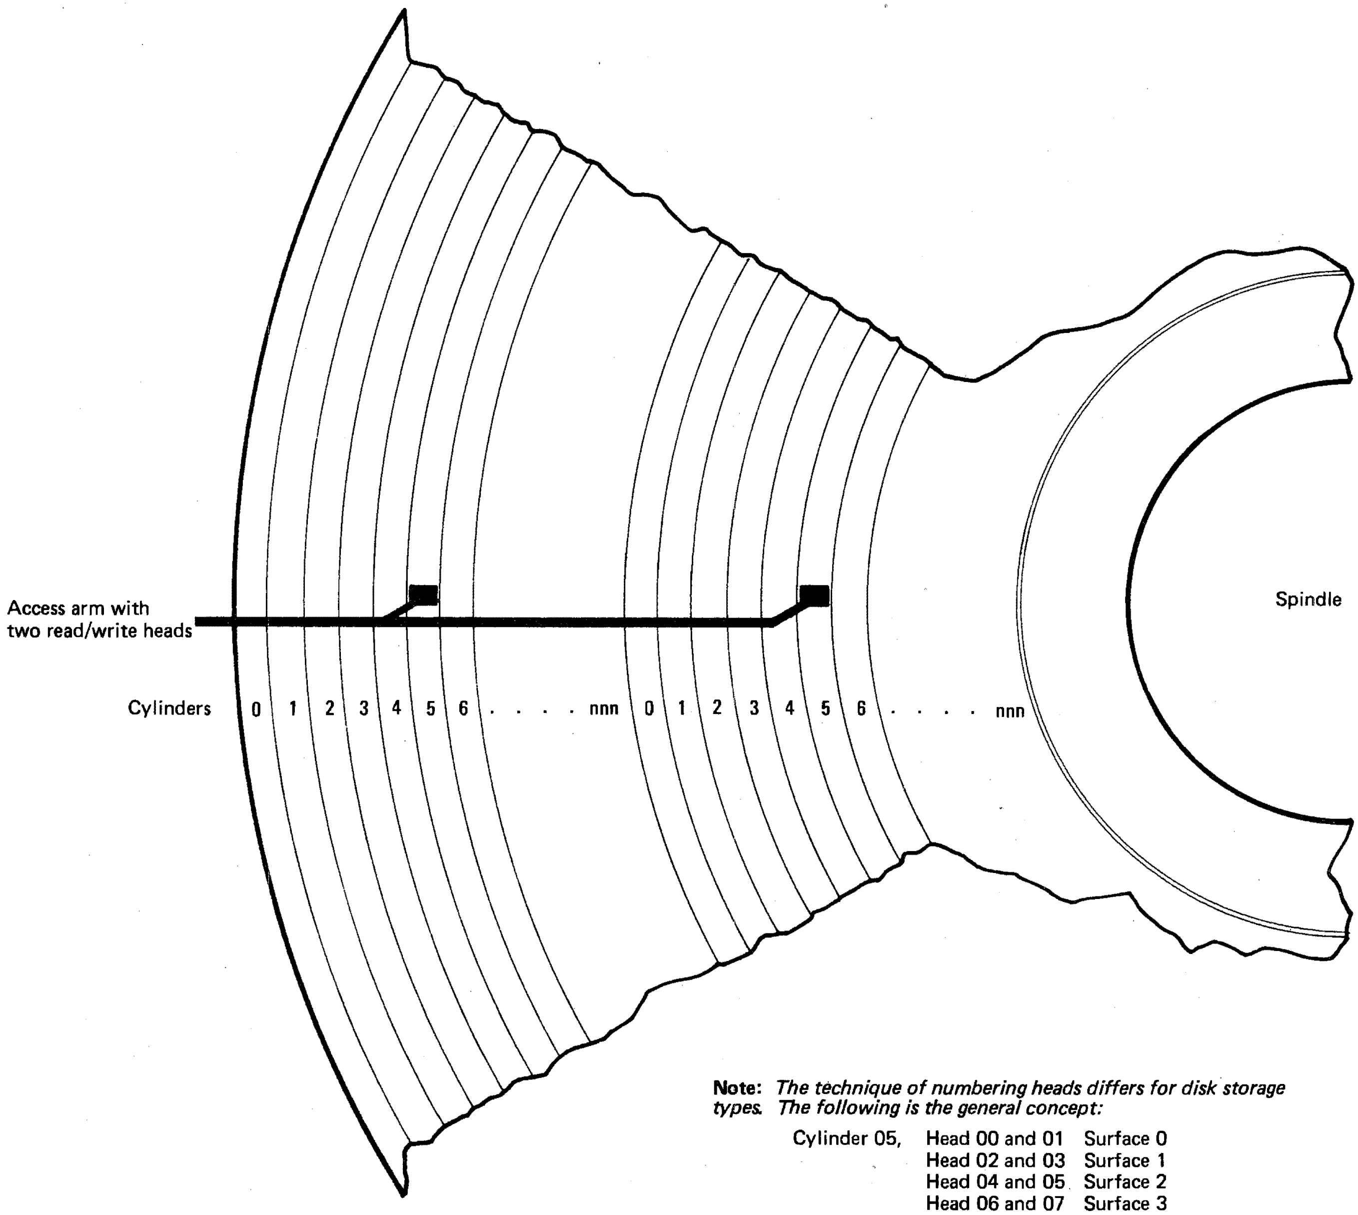

Note: *The technique of numbering heads differs for disk storage types. The following is the general concept:*

Cylinder 05,	Head 00 and 01	Surface 0
	Head 02 and 03	Surface 1
	Head 04 and 05	Surface 2
	Head 06 and 07	Surface 3
	Head nn and nn	Surface n

Figure 4. Disk Surface with Two Read/Write Heads

Tracks and Records

For the tasks associated with the physical use of disk storage, it helps to be familiar with the methods and terminology used to describe:

- The unit of data stored on the disk
- The location on the device where data is stored and how it is selected
- The unit of data transferred between disk storage and processor storage

In addition, you need to know certain details concerning the way tracks and records are formatted; that is, the pattern, or format, in which data is arranged on the device.

Physical Records

The unit of data stored on a disk is a *physical record*. A physical record contains a data area, plus control information needed to describe it and information needed to check it for correctness when it is read.

A physical record is written in two possible formats, and a disk storage type supports one format or the other. The controls for accessing the data are different for the two formats. Of the devices covered in this manual, all use the *count, key, and data format and controls*, except the 3370 disk storage type. The 3370 uses a *fixed-block architecture with a fixed-block format and controls.*

For both count, key, data, and fixed-block architecture formats, physical records are separated by gaps.

With a count, key, and data device type, the programmer decides the length, in bytes, of each physical record; that is, the number of bytes allocated for the data area of the physical record. The channel programmer formats (writes) certain control information for each record in the pattern that is predefined for the count, key, and data format.

With fixed-block architecture devices, the length of the data area of the physical record is predetermined for the device type, and all are the same length. Therefore, the user does not need to specify the number of bytes for the data area. The records are formatted (that is, certain control information is written) at the time of manufacture in a pattern that is pre-established on the track.

Track Address

A physical record has a *track address* that identifies its location on the device. A track address is made up of a cylinder number and a head number of four bytes (CCHH). Up to four hexadecimal digits are used for the cylinder and up to four hexadecimal digits for the head. Both are numbered sequentially beginning with zeros.

On devices with two heads per surface, the two tracks that are accessed when the access mechanism is positioned have the same cylinder number. The cylinder address number has two possible head address numbers per surface. For example, cylinder address 02 may have head addresses 02 or 03 on one surface, 04 or 05 on another surface, and so on depending on the number of disk surfaces. See Figure 4.

Record Selection

Count, key, and data formatted records are specified for selection by a 5-byte record identifier. The 5 bytes consist of the track address (referred to as a seek address), and customarily a record number.

Fixed-block architecture records are specified for selection by relative block numbers. The entire device is formatted into a continuous sequence of numbered blocks, arranged at evenly spaced sector locations. The subsystem converts the block number to the track address and sector location.

Unit of Transfer

The unit of data transfer between disk storage and processor storage is the physical record.

With count, key, and data devices, the transfer may be the entire record including the control areas, or the data area only, or the control information only.

With fixed-block devices, the transfer may be the entire record including the control information, the data area only, or the control information only. Or the operating system may combine multiple records for transfer as a control interval. (A control interval comprises an integral number of physical records.)

Relation of Logical Record to Physical Record

A logical record is the user's record and is the unit of data used by the processing program. The user specifies the format and record length. A single logical record can be the same length as a physical record or multiple logical records may be grouped, that is, *blocked* into one physical record. Logical records are usually relatively small, and combining them into physical records has two advantages. The larger physical records conserve space on the track for user data by reducing the amount of space required for addressing information and gaps. Blocking also may increase processing efficiency by reducing the number of I/O operations required to process the data.

The operating system handles the *blocking* of logical records into physical records for storage on disk and *deblocking* them for use by the processing program. A logical record sometimes extends beyond one physical record into another physical record, especially if the physical records are of a predetermined fixed length that is relatively small, such as on the 3370 disk storage. In this case, the operating system separates the portions of the logical record for storage on the disk and then recombines them for use by the processing program. If the basic access methods are used, the blocking and deblocking functions are not performed automatically by the operating system. The methods used for describing logical records and their relation to physical records are fully described in the data management manuals listed in the bibliography.

References to Bibliography

Refer to the disk storage manuals for further details on tracks and record formats and for record addressing for the particular device type.

Refer to the data management manuals for additional information on records from a programmers point of view.

Refer to *Effective DASD Utilization* for guidelines on establishing block sizes.

Count, Key, and Data Format

For count, key, and data devices, a track address and, by convention, a track descriptor record are written at the beginning of the track. The track address written at the beginning of the track is called the home address (HA). The descriptor record is record zero (R0). The home address and record zero are written at the time of manufacture. Later it may be necessary to rewrite them. Home addresses and record zeros are rewritten using the Device Support Facilities program. When R0 is written, any data on the rest of the track is erased by the writing of a fixed pattern, if the operation is the last command in the chain.

Normally, the track address is also written at the beginning of each record. This is done when the record is formatted. Once a record has been formatted, data can be written and read in the length of space allocated for the data area.

Track Format

All count, key, and data tracks are formatted beginning at the index point and ending at the following index. Each track has the same basic format: track home address and track descriptor record, followed by user records. Records are numbered sequentially, normally beginning with record 0. See Figure 5.

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

The home address area also contains information describing the condition of the track. This information, carried in the flag (F) byte, indicates if the entire track is defective, or if the track is an alternate for a defective track. Additional control information provides for skipping defective areas on the track.

Track Descriptor Record: The track descriptor is always the first record (record zero) on the track following the home address. In IBM programming systems, if the flag in the home address area indicates the track is defective, the count area of the following record zero (R0) contains the track address of the track that is to be used as an alternate. The count area of the R0 on the alternate track contains the track address of the defective, primary track.

The data area of a standard R0 is usually written with a fixed pattern.

Data Records: One or more user records, with count, key, and data areas, can be written following the track descriptor record (R0) on a track.

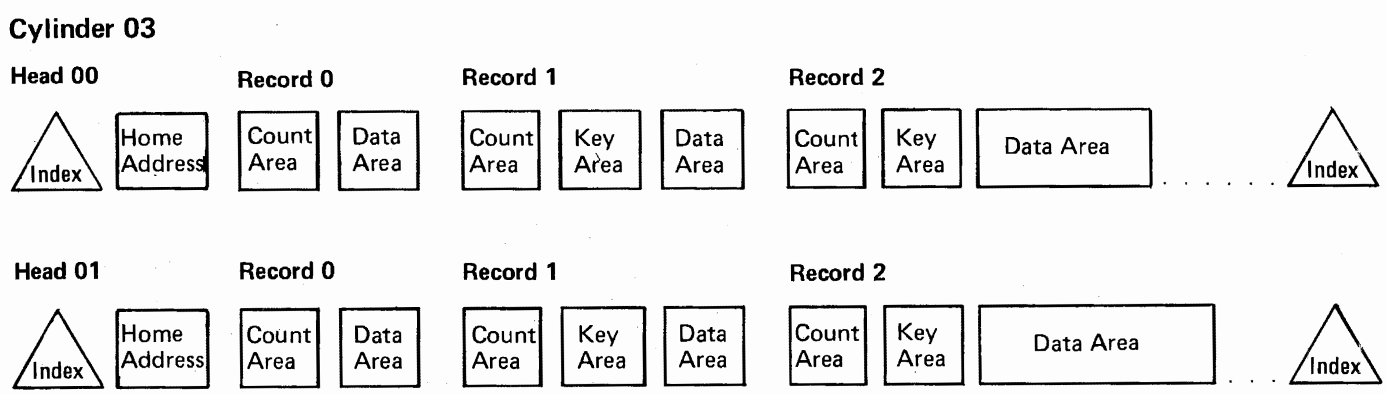

Figure 5. Track and Record Formats for Count, Key, and Data Devices

Record Format

When a disk storage device type supports count, key, and data format, a physical record contains three areas: count, key, and data. Each area is separated by gaps. Checking information is added to each area when it is written that is used later for detecting and correcting data errors. See Figure 5.

Count Area: The count area contains the location of the data area that follows. The location is specified by the track address (cylinder and head numbers) and record number. The record number customarily identifies the number of the record on the track, 0 - n. The location is expressed as five bytes (CCHHR).

In addition to addressing information, the count area specifies the length, in bytes, of the key area and data area that are part of the record.

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. The standard record 0 does not have a key area.

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

Records per Track

The number of records that can be placed on a track depends on the length of the data areas of the records and the overhead required by the device. The records may be of equal or of unequal lengths. For the 3375 and 3380, the length of a single physical record may not exceed the capacity of the track available for user data. For 3330, 3340, and 3350 disk storage, there is hardware support for extending a physical record that exceeds the track capacity to another track. This is referred to as record overflow. For each overflow segment to another track, there is separate control information (that is, a separate count area).

Note that the terms *block* and *physical block* are sometimes used to encompass the entire physical record and other times to refer only to the data area of the record. When block sizes are specified by the user, the size refers to the data area, not to the entire physical record. When calculating the number of records that can be placed on a track, the data length refers to the data area of a physical record. Therefore, the number of bytes specified for a block size cannot exceed the data length specified in the count area for the data area of the physical record.

References to Bibliography

The description manual for each disk storage type includes formulas and tables for calculating the number of records of given block sizes that can be placed on a track for the device type.

Refer to the storage control manuals for further information on record overflow.

Fixed-Block Architecture Format

For fixed-block architecture device types (3370), tracks and records are formatted at the time of manufacture. There is no track home address as in count, key, and data device types.

Track and Record Formats

Each fixed-block record has an ID area and a data area of a fixed number of bytes, as shown in Figure 6.

The ID area contains the block number and cylinder and head numbers. If the block is skipped because it is defective, information pointing to the alternate block is contained in the ID area of the defective block. The ID area of the alternate block contains information pointing to the defective area it is replacing.

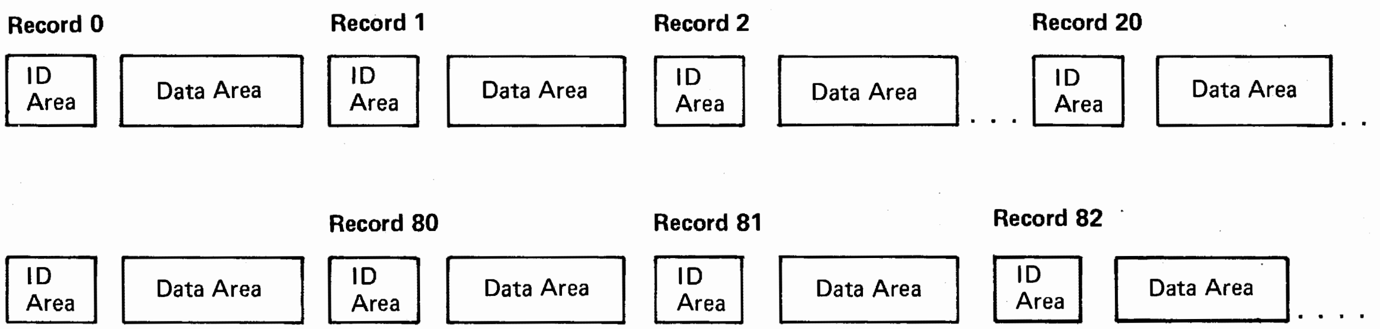

Figure 6. Track and Record Formats for Fixed-Block Architecture Devices

References to Bibliography

For additional information regarding fixed-block architecture records, refer to the 3370 disk storage manual.

Volume

A disk storage volume is a collection of data sets and unused space that can be selected through an I/O selection identifier (address) and accessed by a particular access mechanism.

Volume Initialization

Certain identifying information is written on disk storage so it can be used as a volume by an operating system. This preparation is referred to as *initialization.*

Initialization by the user requires writing a volume label and creating a volume table of contents. The home address and track descriptor, originally written at the plant of manufacture, can optionally be rewritten at this time. Initial program load information also may be written on the volume when it is initialized.

Initialization is accomplished by the user with the Device Support Facilities program.

Volume Label

A volume label (written at initialization) is written at the beginning of the volume. The label includes:

- Volume owner identification (ID)
- Volume serial number
- Volume table of contents address (VTOC)

For count, key, and data devices, the volume label is written as record 3 at cylinder 0, head 0. For fixed-block architecture devices, the volume label is written at block 1.

Volume Table of Contents

Each volume in OS/VS and DOS environments contains a table of contents of the data sets and unused space on the volume.

The VTOC is used to:

- Control the allocation of space on a volume.
- Determine where a data set is located on a volume. The location information is then used by access method routines and device support programs to construct the proper channel commands for accessing a disk location.

You may need a listing or printout of the VTOC for space management or data recovery tasks.

A location for the volume table of contents (VTOC) is assigned and space reserved when the volume is initialized. The track address (CCHH) or block number assigned for the beginning of the VTOC is recorded in the volume label.

VTOC Creation with OS/VS

With OS/VS, information is recorded in the VTOC in records called data set control blocks (DSCBs). The DSCBs have different meanings depending on the format. For each data set on the volume, a format 1 DSCB is constructed for entry in the VTOC that names and describes the data set. The data set description includes the cylinder and head location and the number of extents. (An extent is an addition to the originally allocated space at another location.)

When a volume is initialized, the Device Support Facilities builds the volume table of contents at the location specified by the user and creates the number of DSCB entries permitted by the amount of space allocated for the VTOC by the user. This process includes creating entries that describe the device, the volume, and the volume table of contents (referred to as format 4 DSCBs), and entries that list the unassigned space (referred to as format 5 DSCBs).

The Device Support Facilities program formats records that are free for entry of data set descriptions (referred to as format 0 DSCBs). These free space records have key and data areas of all zeros.

When space on a volume is requested, the Direct Access Device Space Management (DADSM) routines allocate the space on the volume and enter a data set description in one of the free space records of the VTOC. The format 0 DSCB is thereby changed to a format 1 DSCB in the VTOC, and format 5 entries are updated to reflect subtraction of the free space from the list.

VTOC Index for OS/VS

For many operations, the DSCBs in an OS/VS volume table of contents are searched sequentially. To reduce search time, an index to the VTOC can be built. An index is built with the Device Support Facilities program, but requires the Data Facility Device Support program to support it.

References to Bibliography

For additional information on initializing a disk storage volume under OS/VS, refer to the Device Support Facilities manual.

For additional information on the VTOC under OS/VS, refer to the DADSM manuals.

For information on listing or printing a VTOC, refer to the utility manuals and to the manual on the Data Facility Data Set Services program.

VTOC Creation with DOS/VSE

With DOS/VSE, a VTOC is made up of DASD label information established with DLBL and its associated EXTENT control statements. These are supplied by the user. The DASD labels are called File Labels and are of one of four different formats that differ in use as follows:

Format 1 — is the primary label for a file and defines the area(s) of the volume where the records are written. This label can define up to three separate extents for a single logical file. This label is written on each volume upon which a multi-volume file resides.

Format 2 — is required and maintained by ISAM and carries updated information from one use of the file to the next.

Format 3 — is written to contain the additional extents whenever a logical file contains more than three extents and is pointed to by the Format-1 label. For a multi-volume this label is written only on the volumes to which the additional extents apply.

Format 4 — defines the VTOC itself and is written when the volume is initialized.

File labels are written and checked by DOS/VSE OPEN and CLOSE routines.

References to Bibliography

For additional information on initializing a disk storage volume under DOS/VSE, refer to the Device Support Facilities manual.

For additional information on the VTOC under DOS/VSE, refer to the manual DOS *VSE/Advanced Functions Serviceability Aids and Debugging Procedures* and the DOS utilities manuals.

Input and Output Operations

The following brief summary contains information about input and output operations required for most disk storage management tasks.

Input and output (I/O) operations are initiated by the system. A disk storage device is selected and a transfer path established. I/O operations are then directed by channel commands. The code for the command and other information regarding the operation are carried in a channel command word (CCW). Certain CCWs can be chained so that system participation is not required except to begin the chain. A series of channel commands is a channel program.

Channel programs are prepared for access method routines and device support programs. (Or the user may perform the channel programming functions without an access method.) Two different sets of channel commands are used for disk storage. One set directs operations for devices with the count, key, and data record format, and the other set directs operations for devices with the fixed-block architecture record format.

There are commands for the following common functions:

- Formatting records for count, key, and data devices.
- Locating records to be written or read.

  For count, key, and data records, this process uses seek and search commands. Different commands are used for locating records when a speed matching buffer is used for the 3375 and 3380 and when these devices are attached to a 3880 Model 11 or 13 with buffer storage. In these cases, define extent and locate commands are used.

  For fixed-block records, the process uses define extent and locate commands.
- Writing or reading records
- Obtaining usage and error information regarding the disk storage operation
- Reserving and releasing a disk storage device for use by one or a group of channels

In addition, some disk storage types have special commands.

With count, key, and data devices, when a path and device have been selected, the channel sends the cylinder and head numbers to start a *seek* operation and the subsystem causes the access mechanism to be positioned to that track. When the access mechanism has been positioned at a cylinder, and a head has been activated, a write or read operation can be performed. However, the record on the track must still be located. This is accomplished with a command that causes the subsystem to search for the record identifier. The desired record identifier is compared to record numbers in the count area of records on the track until a match is achieved. The function of rotational position sensing allows the search operation to be started just before the desired record location rotates to the read/write head.

With fixed-block architecture devices, when a path and device have been selected, the channel sends block numbers to start define extent and locate operations. The subsystem converts the block numbers to the proper cylinder, head, and sector location on the track and causes the disk storage to position directly to that location.

References to Bibliography
Channel commands are described in the storage control manuals. Rotational position sensing and other functions directed by channel commands also are included.

Device Status and Maintenance

Maintenance procedures are built into the subsystem and system to ensure reliable disk storage operations. Maintenance functions include:

- Counting operations, such as number of bytes read
- Detecting errors that occur and automatically recovering from the problem where possible
- Giving notification of an error situation that needs attention
- Generating records regarding operations

Information regarding operations is communicated by the subsystem to the system in status and sense information.

Status information pertains to the availability of subsystem components and the progress of operations. A particular status indicates a particular condition, and usually a combination of status indicators is sent to more fully describe the circumstance.

If an error is detected during execution of a channel command, the subsystem includes a unit check status in the combination of status conditions sent to the channel. This may cause error recovery procedures to be performed by the subsystem, or cause an I/O interruption to get the attention of the system. An I/O interruption usually causes the system to request sense information for additional information regarding the error. Based on this information, the system may perform error recovery procedures.

Subsystem Role in Maintenance

The subsystem, including the storage control, controller, and disk storage, participates in maintenance of disk storage operation.

Although the exact functions differ depending on the disk storage device type, the subsystem performs the following general roles:

- Adds error detecting and correcting information to records when they are written.
- Detects errors in reading data, in performing control operations, in functioning of the hardware, and in programming.
- Recovers from data and control errors where possible.
- Keeps count of the use of the disk storage (such as number of bytes read), and for some disk storage, keeps count of errors. (The subsystem log is described in a separate section, "Counting and Logging.")
- Assembles usage and error information in the form of sense information.

Sense information is usually obtained from the subsystem after it has interrupted the system with a unit check status. Four conditions may cause the interruption:

- A permanent error condition exists.
- An error condition existed that was successfully recovered and should be logged by the system.
- An error condition exists that the system is to handle.
- Counters of usage or errors (within the subsystem) are full and need to be sent to the system for logging.

For the first three of these four conditions, the sense bytes identify the conditions that caused the unit check status, indicate the type of error, and provide information as to where the error was detected. The sense bytes may further provide information for system error recovery procedures and for diagnosing and isolating the cause of an error condition.

There are 24 bytes of sense information. Sense bytes 0 through 2 describe the error. The meaning of each bit of these first three sense bytes is similar for all disk storage types, although there are a few special definitions. For example, an error detected when data is being read causes a bit defined as a data check to be set. (Indications of data and equipment type errors are referred to as data *checks* and equipment *checks* in sense information.) The setting of a particular combination of bits in the first 3 bytes define an *error condition.* For example, the data check bit and the permanent error bit are an error condition. For each error condition, there is a recovery action recommended for the system error recovery procedures to perform. Sense bytes 3 through 23 provide additional explanation of the error, including its location. Interpretation of bytes 8 through 23 depend on the format used.

Error messages and reports give sense bit settings in hexadecimal form. The 24 bytes are listed, 0 through 23. The two last bytes usually contain a symptom code. The service representative interprets the symptom code based on the definition for the particular device type.

System Role in Maintenance

System error recovery procedures for disk storage are included as part of the operating system.

(The system error recovery procedures for disk storage are used by most programs. However, some programs have procedures that differ from those implemented by the system, and our descriptions regarding error procedures may not apply in these situations.)

The system error recovery procedures perform the following functions:

- Implement recovery actions where possible
- Issue system messages at the operator console
- Create and log records of usage and error information. (The system log is described in a separate section, "Counting and Logging.")

A specific recovery action is based on a particular error condition. The error condition is defined by the combination of bits set in the first three sense bytes sent from the subsystem. (The 3380 also uses a bit of byte 4.) System recovery actions include, for example, retrying an operation when an equipment check is reported in the sense information. The type of system recovery action for data checks depends on the disk storage type and where the error occurs.

References to Bibliography

Sense bit definitions are described in the storage control manuals.

The manuals also include definitions of error conditions and the recommended system recovery actions.

Error Description

This section explains the various dimensions of an error situation, and then elaborates on the specific techniques used by the subsystem and system for automatically handling data type errors. This background information will be helpful in your understanding and handling of error situations. Your role in error handling and the resources available are described in the manual *Error Handling,* Publication Number GA26-1672.

Dimensions of an Error

There are four dimensions to consider when evaluating an error situation.

- Type of error
- Recoverability of error
- Duration of error
- Source of error

Types of Error

Errors can be categorized as *programming errors, data errors, equipment errors,* and *seek and overrun errors.* The type of error is given in sense information, and in error messages and reports. The type of error is similar to a symptom; it is evidence of a problem, but does not necessarily reveal the source nor the cause.

Programming Error

A programming error, such as an invalid track format or incorrect record specification, causes a command rejection. This type of error is passed to the user program for recovery action and, therefore, is not described in this manual.

Data Error

A data error is an error detected in the bit pattern read from the disk. A problem may have interfered with writing the data correctly so that later, when read, an error is detected in the pattern previously recorded on the disk. Or the data may have been written correctly, but a subsequent problem interferes with reading it correctly. (A data error cannot be detected when data is written. Sometimes data is written and immediately read back to the storage control, where it is checked to verify that it was correctly written. If a data error is detected at this time, the data can then be rewritten immediately.)

Seek and Overrun Errors

There are two control type errors: an overrun error that results from timing problems and a seek error that results from failure to position the access mechanism to the correct cylinder and head address.

Equipment Error

An equipment error is an error detected in mechanical or electrical operations.

Recoverability - Temporary or Permanent Errors

When a data, control, or equipment error is detected by the subsystem, either the subsystem or system may attempt recovery, depending on the nature of the error. The terms *temporary* and *permanent* define whether or not the automatic recovery is successful.

(In addition to subsystem and system error recovery procedures, some programs have procedures for error recovery that are called on when recovery attempts by the subsystem and system are not successful.)

Temporary Error

A temporary error is one that is recovered by correcting the data or by retrying the operation. In both cases the operation in progress is completed. If the subsystem performs the recovery procedures, the action does not cause an I/O interruption. If the system performs the recovery procedure, an I/O interruption is always involved. A temporary error is sometimes referred to as a soft error.

Permanent Error

A permanent error is one that is not correctable and cannot be recovered by retrying the operation. It always causes an I/O interruption at the system and causes the subsystem to send unit check status to the system. A permanent error usually terminates a job step and may terminate the job. Both actions depend on the error recovery procedures at the application level. (A permanent error is sometimes referred to as a hard error and sometimes as an I/O uncorrectable error.)

Recoverability by Error Type

Temporary or Permanent Data Error. With a data error, information may be available to actually correct the data. In other cases, even though the first attempt to read the data is not successful, it may be read correctly with subsequent attempts. If neither of these is successful, the data error is regarded as permanent. The techniques used to handle data checks and the role played by the subsystem and system differ for different disk storage device types. A fuller explanation of data errors is given later in this section.

Temporary or Permanent Control Error. When a control (overrun and seek) error is detected, the subsystem retries the operation a specified number of times. If the retry is not successful within the specified number of tries, the error is reported to the system as a permanent equipment type error. If the error is an overrun, the system may then retry the operation. (Seek checks for the 3340 are always retried by the system.)

Temporary or Permanent Equipment Error. Equipment errors always cause the subsystem to interrupt the system, and the system retries the operation a specified number of times. If the operation is then successful, the error is recorded as temporary; if the operation is still not successful, the error is recorded as permanent.

Duration – Intermittent or Continuous Errors

An error may be *intermittent* or *continuous*, regardless of whether it is temporary or permanent in terms of recoverability.

Intermittent Error

An intermittent error is one that changes with time. For example, an error classified as permanent (because it was not recovered automatically) may be caused by a transient condition, such as electrical noise due to lightning. Although the operation could not be completed at the time, a subsequent attempt to perform the same operation may be completely successful. Therefore, some permanent errors may not require any action other than restarting the job. (Some application environments provide a level of error recovery above the system, and the applications may be able to recover, automatically, if the permanent error symptom is due to an intermittent cause.)

Continuous Error

A continuous error is one that does not change with time. It requires outside intervention to restore normal function. For example, a data error caused by a small scratch on the disk surface may always be correctable, therefore, the data check is temporary. However, if the error occurs every time the data is read, it is continuous. This is usually referred to as a repeatable temporary error. A continuous permanent error is sometimes called a solid failure.

To determine if an error is intermittent or continuous, it is necessary to review error reports over a period of time or attempt to reproduce the situation to see if the error is repeated.

Source of Error

The source of an error is important in evaluating an error situation. A data error can be evidence of a hardware problem in one of the subsystem components, such as the controller, but it is more likely that the source of a data error is at the recording media. A data error can be caused by such things as a slight misalignment of the head with the center of the track, a transient electrical interference, or an imperfection on the disk surface.

Equipment errors are evidence of a problem in the controller, the drive, or the head and disk assembly, or in a component that makes up the path to the disk storage. The error may be caused by a mechanical feature or an electronic component.

Determining the source of an error, and then the exact cause, usually requires analysis of sense information and other diagnostics. These problem determination functions are performed with program aids or by a service representative.

Automatic Handling of Data Errors

When a data check occurs, there are two methods that may be attempted for automatic recovery:

- The data actually may be corrected.
- The data may be read correctly when the read operation is retried.

All disk storage types described in this manual add error checking information to each count, key, and data area of a record when it is written; or to each ID and data area of a fixed-block architecture device. Later, when the record is read, the information is used to detect errors that may be present and to correct the data where possible.

Error sensing and detection codes have been provided in disk storage and disk storage controls from the beginning. Techniques for performing these error recognition functions are continuously being improved to keep pace with the higher density with which data is recorded and the higher speed with which it is written and read.

Error detection and correction information is a pattern of bits, coded to define a validity check for the full data in the recorded area. It is referred to as error checking and correction code (ECC). By comparing the code information generated from the data read from the track with the code information that was recorded with the data, any discrepancies are detected. Detected errors, within a certain range of bytes, can be corrected with information from the code.

Corrected Data Errors

When *ECC correctable,* the data is reconstructed to be the same as the data originally transferred to disk storage. The data correction may be done by the subsystem as data is transferred to processor storage, or it may be done by the system after the data is in processor storage. Although the data in processor storage is correct for use in processing, the data on the disk is not changed. The next attempt to read the data from the disk may again result in a data error.

Retried Data Errors

The other technique used for recovery from data errors is to retry the read operation by reissuing the command. The *retry* is done repeatedly in attempts to read the data correctly. If retry attempts are made a certain number of times without successfully reading the data, some disk storage types then physically adjust the access mechanism. This causes the head to move to different positions across the track in attempts to better read the data signals. This technique is referred to as retry with head *offset.* Retry is sometimes done by the subsystem and sometimes by the system, depending on the device type.

Success of Recovery

When one of the automatic recovery methods is successful, the data error is temporary.

If none of the automatic recovery methods succeeds, the error is permanent.

Recovery Methods by Device Type

The method used for data error recovery depends on the disk storage type and the area that contains the error. The recovery methods for different device types and areas are described as follows and summarized in Figure 7 and 8.

3330, 3375, and 3380

For 3330, 3375, and 3380 devices, when the data is ECC correctable, the correction is made by the subsystem or system, depending on where the error occurred on the track.

- If the error is in the home address area of the track or the count or key area of a record, the subsystem makes the correction. It uses the ECC correction information to reconstruct the data. The data correction and resumption of the command chain are done without interrupting the system.
- If the error is in the data area of a record, the system error recovery procedures make the correction. It uses pattern and displacement information derived from the ECC. This information is supplied to the system in sense information.

When the data is not ECC correctable, the subsystem retries the operation. If this is not successful, head offset is invoked, and attempts to read the data are repeated. These retry attempts are made in any area of the record in which an error is detected. 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.)

3340 and 3344

All data error recovery is done by the system.

- If the error is in a data area of a record and is ECC correctable, the system error recovery procedures make the correction. It uses pattern and displacement information derived from the ECC. This information is supplied in sense information.
- For data errors in a data area that are not ECC correctable and for all data errors in other areas, the system error recovery procedures attempt recovery by retrying the operation.

3350

Data error recovery for the 3350 is done by the subsystem or system as follows:

- If the error is in the data area of a record and is ECC correctable, the system error recovery procedures make the correction. It uses pattern and displacement information derived from the ECC. This information is supplied in sense information.
- For data errors in a data area that are not ECC correctable and for all data errors in other areas, the subsystem attempts recovery by retrying the operation. It does not retry with the head offset.

3370

For the 3370, when the data is ECC correctable, the correction is made by the subsystem or system, depending on where the error occurred on the track.

- If the error is in the ID area of the record, the subsystem makes the correction. It uses the ECC correction information to reconstruct the data. The data correction and resumption of the command chain are done without interrupting the system.
- If the error is in the data area of a record, the system error recovery procedures make the correction. It uses pattern and displacement information derived from the ECC. This information is supplied to the system in sense information.

When the data is not ECC correctable, the subsystem retries the operation. If this is not successful, head offset is invoked, and attempts to read the data are repeated. These retry attempts are made in any area of the record in which an error is detected. The command retries are made without interrupting the system.

References to Bibliography

For further information on the differences in error handling for each device, refer to the disk storage manuals listed in the bibliography.

Counting and Logging

There are two logs that may contain information about disk storage errors. These are the subsystem log and the system log.

Subsystem Log

A buffered log is kept in the storage control for each disk storage that attaches to it.

The log contains counts of the seeks made, the bytes read, and overruns.

For some devices, counts of data checks and seek checks are also kept. For the 3330 and 3370, counts are kept of all data checks and seek checks. For the 3350, counts are kept of retried data checks and seek checks. For the 3340, no count is kept of data nor seek checks. For the 3375 and 3380, seek checks and data error rates are calculated but the error counter values are not sent to the system. The subsystem detects the occurrence of an excessive number of temporary errors and causes logging of complete sense information for a specified number of subsequent errors.

When a counter is filled, the subsystem interrupts the system and sends the contents of *all* the counters for that device to the system. The counters are then reset to zeros. The content of the counters may also be obtained directly by the system by issuing a channel command to read the buffered log. This also causes all of the counters to be reset to zero.

System Log

Records of error and error usage information are stored in the system log as error recording data sets (ERDS). The information is stored in SYS1.LOGREC by OS/VS and in SYSREC by DOS. Much of the information stored in the log is supplied by the storage control as sense information.

For logging purposes, the system processes the sense information to produce the set of log records and adds information regarding the results of recovery actions.

Two types of records contain information about disk storage. Both are based on sense information supplied by the subsystem.

- Outboard Record (OBR) - contains error description (including record address) for each error situation
- Miscellaneous Data Record (MDR) - contains usage and error counts

The I/O address used for selection also is stored in the log.

These records stored in the system log are available to various programs that use the information for analyzing error situations and usage.

The Environmental Recording, Editing, and Printing (EREP) program prints error reports based on information it obtains from the system error log.

When an error is entered in the system log, the cylinder and head numbers at which the error occurred are available, because these locations are contained in the sense information.

Summary Tables by Device Type

Temporary Data Errors - Count, Key, and Data

	3330 and 3333	3340 and 3344	3350	3375 and 3380
HA, R0, Count Area	Data checks corrected with ECC or recovered with retry, with and without offset, by subsystem in HA, count, and key areas. Logged when in logging mode.[1]	Data checks recovered with retry by system in HA, count, and key area. Logged.	Data checks recovered with retry by subsystem in HA, count, and key area. Logged when in logging mode.[1]	Data checks corrected with ECC or recovered with retry, with and without offset, by subsystem in HA, count, and key areas. Logged when in logging mode.[2]
Data Area	Data checks corrected with ECC by system or recovered with retry, with and without head offset, by subsystem in data area. Logged when in logging mode.[1]	Data checks corrected with ECC or recovered with retry by system in data area. Logged.	Data checks corrected with ECC by system or recovered with retry, without offset, by subsystem in data area. Data checks in data area corrected with ECC are logged. Data checks recovered with retry are logged when in logging mode.[1]	Data checks corrected with ECC by system or recovered with retry, with and without offset, by subsystem in data area. Logged when in logging mode.[2]

Notes:

1. Logging begins when the data error counter for the volume overflows or when logging is forced. Logging continues for the next four data checks for the drive. Use of head offset with retry is not reported.
2. Logging begins when the data error rate threshold for the volume is exceeded. Logging continues after this first log for another 23 data checks (without offset) for the string. Use of offset with retry is always reported.

Permanent Data Errors - Count, Key, and Data

	3330 and 3333	3340 and 3344	3350	3375 and 3380
HA, R0, Count Area	Data checks in HA, count, and key area not corrected with ECC nor recovered with retry by subsystem. Logged.	Data checks in HA, count, and key areas not recovered with retry by system. Logged.	Data checks in HA, count, and key areas not recovered with retry by subsystem. Logged.	Data checks in HA, count, and key area not corrected with ECC nor recovered with retry by subsystem. Logged.
Data Area	Data checks in data area not corrected with ECC by system nor recovered with retry by subsystem. Logged.	Data checks in data area not corrected with ECC nor recovered with retry by system. Logged.	Data checks in data area not corrected with ECC by system nor recovered with retry by subsystem. Logged.	Data checks in data area not corrected with ECC nor recovered with retry by subsystem. Logged.

Note: When executing a Read Multiple Count, Key, and Data command, the above described procedures apply to the first record read. If errors are detected when reading subsequent records, unit check status is sent. The system retries the operation 10 times, attempting to read the data correctly. If this fails, a system console message is issued and a permanent data error is logged.

Figure 7. Summary of Subsystem and System Handling of Data Checks - By Disk Storage Type for Count, Key, and Data Devices

Temporary Data Errors – Fixed Block

	3370
ID Area	Data checks corrected with ECC or recovered with retry, with and without offset, by subsystem in ID areas. Logged when in logging mode.
Data Area	Data checks corrected with ECC by system or recovered with retry, with and without head offset, by subsystem in data area. Logged when in logging mode.
	Note: Logging begins when the data error counter for the volume overflows or when logging is forced. Logging continues for the next four data checks for the drive. Use of head offset with retry is not reported.

Permanent Data Errors – Fixed Block

	3370
ID Area	Data checks in ID area not corrected with ECC nor recovered with retry by subsystem. Logged.
Data Area	Data checks in data area not corrected with ECC by system nor recovered with retry by subsystem. Logged.

Figure 8. Summary of Subsystem and System Handling of Data Checks for 3370

Bibliography

Terminology	Order Number
IBM Vocabulary for Data Processing Telecommunications, and Office Systems	GC20-1699

Disk Storage

Reference Manual for IBM 3330 Series Disk Storage	GA26-1615
Reference Manual for IBM 3340/3344 Disk Storage	GA26-1619
Reference Manual for IBM 3350 Direct Access Storage	GA26-1638
IBM 3370 Direct Access Storage Description	GA26-1657
IBM 3375 Direct Access Storage Description and User's Guide	GA26-1666
IBM 3380 Direct Access Storage Description and User's Guide	GA26-1664

Storage Control

IBM 3830-2 Storage Control gives detailed descriptions of channel commands, and status and sense information.	GA26-1617
IBM Integrated Storage Control gives detailed descriptions of channel commands, and status and sense information.	GA26-1620
IBM 3880 Storage Control Description gives detailed descriptions of channel commands, and status and sense information.	
For 3880 Models 1, 2, and 3	GA26-1661
For 3880 Model 11	GA32-0060
For 3880 Model 13	GA32-0062

Programs

Data Facility Data Set Services: User's Guide and Reference describes dump and restore functions.	SC26-3949
Device Support Facilities describes the commands to initialize and maintain disk storage volumes for system and standalone users.	GC35-0033
DOS VSE/Advanced Functions Serviceability Aids and Debugging	SC33-6099
DOS VSE/Advanced Functions Utilities	SC33-6100
DOS/VSE Data Management 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.	GC24-5138
DOS/VSE Messages lists and interprets the messages that the system issues .	GC33-5379
DOS/VSE System Control Statements describes commands for operations control.	GC33-5376
DOS/VSE System Generation is a reference source for installing, generation, availability, and servicing of the DOS/VSE program. It also provides information for estimating the size of supervisor, libraries, and gives formulas to estimate disk storage space requirements.	GC33-5377
DOS/VSE System Utilities describes the utility programs.	GC33-5381
Effective DASD Utilization is a student text with guidance on block sizes and space allocation.	SR20-4748

Environmental Recording, Editing, and Printing (EREP) Program describes the functions of this program. GC28-0772

IBM Input/Output Equipment Installation Manual – Physical Planning for System/360, System/370, and 4300 Processors gives information for planning floor space requirements and other physical planning considerations for disk storage subsystems. GC22-7064

OS/VS1 Access Method Serv. es GC26-3840

OS/VS1 Direct Access Device Space Management describes these routines to automatically allocate space on a volume. SY26-3837

OS/VS1 Data Management Services GC26-3894

Operator's Library, VS1 Reference describes commands for operations control. GC38-0110

OS/VS Message Library: VS1 System Messages GC38-1001

OS/VS1 System Generation GC26-3791

OS/VS1 Utilities Manual describes the capabilities of the OS/VS1 utility programs and the control statements used with each. GC26-3901

OS/VS2 Access Method Services GC26-3841

OS/VS2 DADSM Logic describes the use of these routines to automatically allocate space on a volume. SY26-3828

OS/VS2 MVS Data Management Services Guide describes data management, including how to define a data set and use of the access methods (except VSAM). GC26-3875

OS/VS2 Message Library: VS2 System Messages describes the messages issued on the system operator consoles. GC38-1002

OS/VS2 System Programming Library: System Generation Reference GC26-3792

Operator's Library: OS/VS2 MVS System Commands describes commands for operations control. GC38-0229

OS/VS2 MVS Utilities describes how to use OS/VS utilities under OS/VS2 MVS. GC26-3902

Index

IBM Disk Storage Management Guide
Background Reference Information

Order No. GA26-1675-0

READER'S COMMENT FORM

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

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

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

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

Possible topics for comment are:

Clarity Accuracy Completeness Organization Coding Retrieval Legibility

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

____

____

____

What is your occupation ? ________________________

Number of latest Newsletter associated with this publication: ________________________

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

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

Cut Along This Line

Reader's Comment Form

Fold and Tape | **Please Do Not Staple** | Fold and Tape

NO POSTAGE NECESSARY IF MAILED IN THE UNITED STATES

BUSINESS REPLY MAIL

FIRST CLASS PERMIT NO. 40 ARMONK, N.Y.

POSTAGE WILL BE PAID BY ADDRESSEE:

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

Fold and Tape | **Please Do Not Staple** | Fold and Tape

IBM®

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

Cut or Fold Along Line

IBM Disk Storage Management Guide — Background Reference Information (File No S370/4300-07) Printed in U.S.A. GA26-1675-0

IBM Disk Storage Management Guide
Background Reference Information

READER'S COMMENT FORM

Order No. GA26-1675-0

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

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

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

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

Possible topics for comment are:

Clarity Accuracy Completeness Organization Coding Retrieval Legibility

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

____

____

____

What is your occupation ? ______________________________

Number of latest Newsletter associated with this publication: ______________________________

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

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

Cut Along This Line

GA26-1675-0

Reader's Comment Form

Fold and Tape — **Please Do Not Staple** — Fold and Tape

NO POSTAGE NECESSARY IF MAILED IN THE UNITED STATES

BUSINESS REPLY MAIL

FIRST CLASS PERMIT NO. 40 ARMONK, N.Y.

POSTAGE WILL BE PAID BY ADDRESSEE:

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

Fold and Tape — **Please Do Not Staple** — Fold and Tape

IBM®

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

Cut or Fold Along Line

Publication Number
GA26-1675-0

File Number
S370/4300-07

Printed in
USA

IBM®