

GA26-1592-4
File No. S360/370-07

Systems

**Reference Manual for
IBM 3830 Storage Control
Model 1 and
IBM 3330 Disk Storage**

IBM

Preface

The IBM 3830 Storage Control and IBM 3330 Disk Storage form a large capacity, high speed direct access storage facility for general purpose data storage and system residence. Attached to the central processing unit through a block multiplexer channel, the facility operates under direct program control of the CPU.

For experienced programmers, this manual provides readily accessible reference material related to channel command words, sense bytes, track format, track capacities, and error recovery.

Less experienced programmers will find sufficient information to create channel programs to best utilize the standard and special features of the 3830/3330 facility.

This manual is organized by the following topics:

- INTRODUCTION - describes the basic units and lists highlights and functions.
- INPUT/OUTPUT - describes operation between the channel and the Storage Control.
- CHANNEL COMMANDS - lists each command and describes its use.
- CHANNEL PROGRAM EXAMPLES - gives programming examples and command usage.
- STANDARD FEATURES - describes command retry, rotational position sensing (with formula for calculating sector positions), and other standard features.
- SPECIAL FEATURES - describes two channel switch and two channel switch additional operation.

Fifth Edition (October, 1974)

This is a merge reprint that replaces and makes Order No. GA26-1592-3 obsolete. All changes and additions contained in Technical Newsletters, GN26-0303 and GN26-0306, have been incorporated into this edition. Significant changes or additions to the specifications contained in this publication are continually being made. Before using this publication in connection with the operation of IBM equipment, contact the local IBM Branch Office for revisions.

The illustrations in this manual have a code number in the lower corner. This is a publishing control number and is not related to the subject matter.

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, Product Publications, Department G24, San Jose, California 95193.

© Copyright International Business Machines Corporation 1971, 1973

- ERROR RECOVERY PROCEDURES - explains the error correction function and lists error recovery techniques for various failures that may be encountered. The failures and recovery techniques are listed in table form.
- OPERATING INSTRUCTIONS - gives a complete description of switches, indicators, and procedures for loading and unloading disk packs for systems installation operators.

In addition, a console error message is illustrated and each field is explained in detail.

A complete summary of sense information is contained in Appendix A. Record/Track capacities for various record lengths are listed in Appendix B. The track capacity formula used to calculate these capacities is also included.

Programmers should be familiar with the information contained in IBM System/360 Principles of Operation, Order No. GA22-6821, and IBM System/370 Principles of Operation, Order No. GA22-7000. Operators should be familiar with the material presented in the system summary for the parent system. Order numbers for system summary and other related publications can be found in IBM System/360 and System/370 Bibliography, Order No. GA22-6822.

For definitions of terms used in connection with direct access storage devices, see Data Processing Glossary, Order No. GC20-1699.

CONTENTS

INTRODUCTION	1	Descriptions of Individual Sense Commands . . .	47
Highlights	1	Descriptions of Individual Write Commands . . .	51
General Description	1	CHANNEL PROGRAMS	57
Speed and Capacity	2	STANDARD FEATURES	64
IBM 3336 Disk Pack	2	Multiple Track (MT) Operation	64
Pack Initialization	2	Record Overflow	64
Record Format	2	Formatting Overflow Records	65
Count Area	2	Processing Overflow Records	66
Key Area	2	End-Of-File	66
Data Area	4	Rotational Position Sensing	67
Track Format	4	Multiple Requesting	69
Home Address	4	Command Retry	69
Track Descriptor Record	4	Statistical Usage/Error Recording	70
Data Records	4	Storage Control Diagnostics	70
Gaps	7	Usage Meter	70
Detecting and Correcting Errors	7	3830 Storage Control Meter	70
CPU Parity	7	SPECIAL FEATURES	71
Error Correction Code	7	Two Channel Switch and Two Channel	
Data Integrity	7	Switch Additional	71
INPUT/OUTPUT OPERATIONS	8	Channel Selection Switch	71
General Description	8	Device Status	72
Channel Operation	9	Addressing	72
Channel Address Word	9	Resets	72
Channel Command Word	9	ERROR RECOVERY PROCEDURES	73
Channel Status Word	9	Error Correction Function	73
Program Status Word	12	Example	77
Command Chaining	12	Construction of Restart CCWs	77
Data Chaining	13	OPERATING INSTRUCTIONS	79
Branching in Channel Programs	13	Loading a Disk Pack	79
Unit Selection and Device Addressing	14	Unloading a Disk Pack	79
CHANNEL COMMANDS	15	Changing a Drive Address	79
Control Commands	15	Console Error Message Analysis	82
Search Commands	15	APPENDIX A. SENSE BYTES	83
Read Commands	15	APPENDIX B. RECORD/TRACK	
Write Commands	15	CAPACITIES	91
Formatting Write Commands	15	INDEX	95
Update Write Commands	16		
Sense/Test I/O Commands	16		
Channel Command Descriptions	17		
Descriptions of Individual Control Commands .	18		
Descriptions of Individual Search Commands . .	30		
Descriptions of Individual Read Commands . . .	37		

The IBM 3830 Storage Control and the IBM 3330 Disk Storage combine to provide a high capacity direct access storage facility for medium-to-large scale IBM computers. Attached to a block multiplexer channel, each facility provides fast access to as many as 800,000,000 bytes of online storage. Standard checking and retry features increase system reliability and availability for batch processing and data base applications.

The 3830/3330 facility can be attached to the IBM 2880 Block Multiplexer Channel or to system channels with equivalent characteristics.

HIGHLIGHTS

- 30 milliseconds average access time.
- Online capacity of 200 million to 800 million bytes in 100-million-byte increments.
- Data rate of 806,000 bytes per second (1,612,000 decimal digits per second).
- Average latency (rotational delay) of 8.4 milliseconds.
- Powered drawers and frontal pack loading.
- Rotational position sensing permits the channel to disconnect during rotational delay.
- Multiple requesting enables multiple channel programs to be simultaneously active on a single facility.
- Command retry enables the facility to recover from most storage control and disk storage errors without the use of error recovery programs.
- Error correction circuitry in the storage control detects and corrects an error burst of up to 11 bits in length.
- Interchangeable address plugs permit online servicing of one 3330 drive while processing continues on other 3330 drives.

GENERAL DESCRIPTION

The 3830/3330 facility, consisting of an IBM 3830 Storage Control and one to four IBM 3330 Disk Storage modules, is attached to an IBM 2880 Block

Multiplexer Channel or integrated system channels with block multiplexing capability. Two versions of the 3330 are available: 3330 Model 1, containing two independent disk drives; 3330 Model 2, containing one disk drive. Each drive holds an IBM 3336 Disk Pack, which provides up to 100 million bytes of storage. On mixed 3330 model combinations, the physical address of a Model 2 drive will be the higher of the two available for the module. Example: A Model 2 drive in a module 4 position would have a physical address of H.

Wherever possible, the 3830/3330 facility has been made program compatible with other IBM direct access storage devices. Major areas of compatibility are the data format, channel commands, and permissible instruction sequences. Additional commands are provided for new features and increased serviceability. File scan commands (standard on the IBM 2314 Direct Access Storage Facility) are not usable with the 3830/3330 facility.

The following standard and special features are included or available with the facility:

ROTATIONAL POSITION SENSING (RPS): Allows the channel and storage control to be released during most of record search time, thus increasing channel and control unit availability for other operations.

MULTIPLE REQUESTING: Allows up to eight channel programs (one per disk drive) to be simultaneously active in the facility.

COMMAND RETRY: A channel-storage control procedure which, under certain conditions, causes a command to be retried without an I/O interruption. This procedure is initiated by the storage control and used to recover from correctable errors.

RECORD OVERFLOW: Provides a means of processing logical records which span track boundaries within a cylinder.

USAGE/ERROR RECORDING: The storage control maintains a statistical data record of usage and error information for each drive.

TWO CHANNEL SWITCH: A special feature that enables two channels to share the storage control and drives.

TWO CHANNEL SWITCH ADDITIONAL: A special feature. With the two channel switch, it enables four channels to share the storage control and drives.

Removable address plugs permit changing the logical device addresses of the drives within the facility. An additional service plug, provided with each facility, permits customer engineer servicing from the CE panel.

A usage meter is provided for the 3830 Storage Control Unit. There are no meters on the 3330 Disk Storage.

The functions of each unit in a basic system configuration are shown in Figure 1.

Speed and Capacity

- *Average access times:*

One cylinder	10 milliseconds.
Average number of cylinders	30 milliseconds.
Maximum number of cylinders	55 milliseconds.
- *Data rate:* 806 kilobytes per second.
- *Rotational delay:*

Minimum	0 milliseconds.
(min. of 250 μ s required for channel connection).	
Average	8.4 milliseconds.
Maximum	16.7 milliseconds.
- *Cylinders per pack:* 411 (including 7 alternates).
- *Tracks per cylinder:* 19.
- *Tracks per pack:* 7,809 (including 133 alternates).
- *Capacity:*

Per track	13,030 bytes.
Per cylinder	247,570 bytes.
Per pack	100,000,000 bytes.

IBM 3336 DISK PACK

The IBM 3330 Disk Storage uses the IBM 3336 Disk Pack (Figure 2). The pack is removable and interchangeable; information written on a pack by one 3330 drive can be read and updated by any other 3330 drive.

The 3336 is a compact disk assembly weighing approximately 20 pounds. Protective disks located at the top and bottom of the disk array minimize physical damage that could result from mishandling. In addition, the pack has a two-piece cover to prevent dust accumulation during storage. The bottom cover has a shock absorbing bumper strip for additional pack protection. For information relating to pack handling, see IBM Disk Pack and Cartridge Handling Procedures, Order No. GA26-5756.

Pack Initialization

All 3336 packs are initialized at the factory, with a home address and eight-byte track descriptor record (R0) written on all tracks. Any defective track is flagged and an alternate track is assigned.

An IBM utility program is available to flag defective tracks and assign alternate tracks if the data areas of the pack should become defective during normal operation.

Another IBM utility program is available to write the volume, volume table of contents (VTOC) and, initial program load (IPL) records. It also determines the number of flagged tracks for entry into the VTOC.

RECORD FORMAT

The basic unit of information recorded on the 3830/3330 is a byte, consisting of eight bits. A group of bytes separated by a special gap is called an area. Areas are combined to make a record, the logical unit of information.

A record consists of three areas: count area, key area (optional), and data area. The significance of the bytes within these areas is given in Figure 3.

Count Area

The count area contains the location of a data record on a specific track, and defines the size of the key and data areas of that record. The count area is written when the record is formatted and is not changed until the record is reformatted.

Key Area

Use of the key area is at the discretion of the programmer. When used, the key area of the record contains the primary identification of the data portion of the record (such as the social security number, man number, part number, or any other uniquely identifying information).

Key area length is defined by the KL byte in the count area. If the KL byte is zero, the key area and following gap are omitted from the record.

Once the key area is formatted, the contents -- but not the length -- may be altered. If the key area is altered, the data area of the record must also be rewritten.





Figure 1. Functional Description



11834

Figure 2. IBM 3336 Disk Pack

Data Area

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

Data information is organized and arranged by the programmer.

The length of the data area is defined by the DL bytes in the count area. If the DL bytes in the count area are zero, an end-of-file record is written. (See "End of File.")

Once the data area is formatted, the contents -- but not the length -- may be altered. The contents of the data area may be altered without affecting any other area in the record.

TRACK FORMAT

All tracks are formatted beginning at index and ending at the following index. Each track has the same basic format: home address, track descriptor record, and one or more data records. The records -- and areas within the records -- are separated by gaps.

Home Address

Each track contains one home address, which defines physical location of the track (track address) and condition of the track. Home address, the first recorded area following index, is separated from index by gap G1.

Special commands are used for writing and reading the home address area: write home address and read home address. Writing home addresses is normally done at the IBM plant.

Track Descriptor Record (R0)

This record is always the first record on the track following the home address area. Although R0 may be used as a normal data record, it is usually reserved by the operating system to store pertinent track information.


In IBM programming systems, the count field CCHH bytes of the defective track provide the address of the alternate track. If it is an alternate track, the CCHH bytes of the count area provide the address of the defective track. (The 3830 uses this information for internal error recovery procedures.) No key area is specified in the KL byte. An eight-byte data field is used to store the number of records on the track and the number of bytes remaining on the track.

Special commands, write R0 and read R0, are used for writing and reading the track descriptor record. Track descriptor records are normally written on the disk packs at the IBM plant.

Data Records

One or more data records may follow R0 on a track. The IBM 3330 uses self-formatting records in which the count area of the record specifies the format and length of the record. Record format is determined at the time the count, key, and data areas of the record are originally written by execution of a format write command. (See "Channel Commands.") The format of the record is not changed until the entire record is rewritten by another format write command.

Data records, as well as track descriptor records, can be formatted with or without keys. Generally, file organization determines whether keys are used. For example, if a sequential file is always processed sequentially, there is no point in formatting with keys. If, however, there is an appreciable amount of random processing, records should be formatted with keys for faster access.



HOME ADDRESS

Index: Indicates the physical beginning of each track. All tracks on the disk pack are synchronized by the same index marker.

G1(Gap1): Separates index and home address.

HOME ADDRESS

F (Flag): Defines the condition of the track and/or indicates a CE disk pack. This is the only flag byte transferred to or from the channel.
 Bits 0 through 4 - unused and written as 0's.
 Bit 5 - when on, this bit indicates a CE disk pack. This bit must be zero on customer packs, or diagnostic routines may destroy customer data.
 Bits 6 and 7 - 00 = normal track
 01 = alternate track
 10 = defective track

CC (Cylinder Number): Specifies the cylinder number (from 0 to 410).

HH (Head Number): Specifies the read/write head within the selected cylinder (from 0 to 18).

ECC (Error Correction Code): Generated by the storage control - used for error detection and correction.

G2 (Gap 2): Separates home address from R0 count field.

RECORD ZERO

R0 COUNT AREA

F (Flag): Defines the condition of the track.
 Bits 0 through 4 - unused and written as 0's.
 Bit 5 - always 0.
 Bits 6 and 7 - 00 = normal track
 *01 = alternate track
 *10 = defective track

CC (Cylinder Number): Specifies the cylinder number (from 0 to 410).

HH (Head Number): Specifies the read/write head number within the selected cylinder (from 0 to 18).

R (Record Number): Specifies the sequential number of the record on the track (zero in this case).

KL (Key Length): Specifies the number of bytes in the R0 key field (from 0 to 255 bytes).

DL (Data Length): Specifies the number of bytes in R0 data field (from 1 to track capacity).

ECC (Error Correction Code): Generated by the storage control - used for error detection and correction.

R0 KEY AREA

G2 (Gap 2): Precedes all key areas.

Key Field: Identifies the information in the data field.

ECC (Error Correction Code): Generated by the storage control - used for error detection and correction.

R0 DATA AREA

G2 (Gap 2): Precedes all data areas.

Data Field: Contains the information identified by the count and key areas.

ECC (Error Correction Code): Generated by the storage control - used for error detection and correction.

* See "Write R0", page 51.

DATA RECORD

DATA RECORD COUNT AREA

G3 (Gap 3): Precedes all count areas (except R0).

F (Flag): Defines the condition of the track, and indicates whether this is an overflow record.

Bits 0 through 3 - unused and written as 0's.

Bit 4 - When on, indicates that the logical record continues on the next track.

Bit 5 - Always 0.

Bits 6 and 7 - 00 = normal track

01 = alternate track

10 = defective track

CC (Cylinder Number): Specifies the cylinder number (from 0 to 410).

HH (Head Number): Specifies the read/write head within the selected cylinder (from 0 to 18).

R (Record Number): Specifies the sequential number of the record on the track.

KL (Key Length): Specifies the number of bytes in the key field (from 0 to 255 bytes).

DL (Data Length): Specifies the number of bytes in the data field (from 1 to track capacity).

ECC (Error Correction Code): Generated by the storage control - used for error detection and correction.

DATA RECORD KEY AREA

G2 (Gap 2): Precedes all key areas.

Key Field: Identifies information in the data field.

ECC (Error Correction Code): Generated by the storage control - used for error detection and correction.

DATA RECORD DATA AREA

G2 (Gap 2): Precedes all data areas.

Data field: Contains the information identified by the count and key areas.

ECC (Error Correction Code): Generated by the storage control - used for error detection and correction.

G4 (Gap 4): 0's are written from the end of the last data field to index.

Figure 3. Record and Track Format

Gaps

Gaps are written by the storage control to delimit records and areas within those records. Gaps generally include a unique combination of bits and recording areas to maintain orientation and synchronization between the storage control and disk storage. Gaps are not accessible to, nor under control of, the using system.

DETECTING AND CORRECTING ERRORS

CPU Parity

To check data accuracy, a parity bit is associated with each byte within the CPU and channel. When a byte is formed, the parity bit is set to either 1 or 0 to maintain an odd number of 1-bits within the byte (i. e. odd parity). Each byte of data to be written is checked for correct parity as it is received by the IBM 3830.

Error Correction Code

As data is transferred from the channel to disk storage (write operation), the storage control removes the parity bit associated with each byte. It then computes the error correction code bytes, which are written after each recorded area. The correction code bytes, coded to represent the data in the recorded area, are used for both error detection and correction.

As data is transferred from disk storage to the channel (read operation), each area is inspected by the storage control and the error correction code

bytes are recalculated for each area. The 3830 correction code corrects single bursts of 11 bits or less.

If a correctable data error is detected in the home address, count, or key areas, the storage control internally executes the error correction function through the use of command retry. (See "Command Retry.") If an uncorrectable data error, or a correctable data error in a data area, is detected, the correction function is determined by the system error recovery procedures. (See "Error Recovery Procedures.")

The correction code bytes are removed and proper parity is generated by the storage control before the data is transferred to the channel.

Data Integrity

Unless corrected immediately, soft write errors cause hard read errors. Therefore, where data integrity is required, verification should be incorporated within the program. Thus, in the event of soft errors, the record can be rewritten and verified before the original data is destroyed.

Either of two verification methods may be used: full readback check or correction code check.

FULL READ BACK CHECK: All of the data just written is read back into main storage and compared, byte-for-byte, with the original information.

CORRECTION CODE CHECK: A read operation is performed with the skip bit on. This method causes the storage control to check the validity of the record using the error correction code bytes.

Input/Output Operations

GENERAL DESCRIPTION

I/O operations, initiated by I/O instructions in the CPU program, are controlled by commands fetched from main storage by the channel. Arithmetical and logical decision operations are performed while the processing unit is in the problem state; for I/O operations, the processing unit must be in the supervisor state.

The processing unit is changed from problem to supervisor state when a supervisor call instruction is executed or when an I/O interrupt occurs. The status of the system existing at the time of the change is stored in the program status word. (See "Program Status Word.")

In the supervisor state, the CPU can execute the following I/O instructions:

1. Start I/O -- Initiates an I/O operation if the addressed channel, storage control, and disk drive are available.
2. Start I/O Fast Release -- Initiates an I/O operation if the addressed channel is available. The storage control and disk drive are assumed to be available. If not, an I/O interrupt occurs to indicate an unavailable condition.
3. Halt I/O -- Terminates the operation in progress at the channel, and the storage control is disconnected from the channel.
4. Halt Device -- Terminates the operation in progress at the storage control without interfering with other I/O operations at the channel. This instruction should be used instead of halt I/O to terminate an operation on a device attached to IBM block multiplexer channels.
5. Test I/O -- Sets the condition code in the program status word to indicate the status of the addressed channel, sub-channel, storage control, and disk drive.


After the specified instruction has been executed, the CPU can return to the problem state and continue the interrupted program by reloading the program status word originally stored when the program entered the supervisor state.

The format for I/O instructions is shown in Figure 4.

I/O Instruction Format



Bit Position	Field Designation	Function
0 - 7	Operation (OP) Code	Designates the operation to be performed.
8 - 14	Not Used	
15		Set to 1 for start I/O fast release and halt device.
16 - 19	Base Address Register Location (B ₁)	Designates the address of a general register in the CPU. The register is 32 bits in length, but only the low order 24 bits are used.
20 - 31	Displacement (D ₁)	Bits 16-31 of the sum obtained by the addition of the contents of the register at B ₁ and the contents of the D ₁ field identifies the channel and the device addressed by the instruction. The result has the following format:



Bit Position	Field Designation	Function
0 - 7	Operation (OP) Code	Designates the operation to be performed.
8 - 14	Not Used	
15		Set to 1 for start I/O fast release and halt device.
16 - 20	Must be Zero	
21 - 23	Channel Address	
24 - 28	Storage Control Address	
29 - 31	Device Address	

26569A

Figure 4. I/O Instruction Format

CHANNEL OPERATION

After successful execution of an I/O instruction, the channel independently selects and governs the storage control and drive addressed by the instruction. Reserved main storage locations contain information and instructions that enable the channel to perform those functions necessary to complete the operation.

Channel Address Word

Issuing a start I/O or start I/O fast release instruction causes the channel to fetch the channel address word from main storage location 72. Bits 0 through 3 of the channel address word form the protection key for all commands associated with the I/O instruction. The protection key establishes the right of access (that is, whether data can be stored or fetched) to the particular main storage locations.

The command address in bits 8 through 31 designates the address of the first channel command word. The three low order bits of the command address must be zero to specify the channel command word on doubleword boundaries.

Fetching of channel address words is a channel hardware function. The information must be set up in main storage location 72 prior to issuing the I/O instruction.

The format for the channel address word is shown in Figure 5.

Channel Command Word

The channel fetches the first channel command word (CCW) from the address specified in the channel address word. The CCW specifies the operation to be performed, the main storage locations to be used, and the action to be taken when the operation is completed.

The channel, if available when it receives the channel command word, attempts to select the device specified in the I/O instruction by sending the address to all attached control units. If the addressed device is attached to the channel and has power on, the command code portion of the channel command word is sent to the storage control, which responds with an initial status byte to the channel.

At this point, the start I/O instruction is finished, releasing the CPU to perform the next instruction. The results of the attempt to initiate execution of the command are indicated by the condition code in the program status word. If the I/O operation was not

started, new status information containing the reason for this condition is normally set in the channel status word.

The format for the channel command word is shown in Figure 6.

Channel Status Word

The channel status word (CSW), stored at main storage location 64, informs the program of I/O device status or the conditions under which an I/O operation was terminated. The CSW is formed or changed during I/O interruptions and instruction execution. Status stored in the CSW remains unchanged until a subsequent interrupt occurs or a new I/O instruction is processed.

The format for the channel status word is shown in Figure 7.


Status Presentation

Status is presented twice (initial status and ending status) for all commands except those seek commands that require access motion, set sector commands, and immediate commands not chained from write commands.

Seek and seek cylinder commands present initial status, channel end status (after transfer of the seek address), and device end (after the access mechanism is positioned).

The set sector command presents initial status, channel end (after transfer of sector number), and device end (when angular position is reached).

Channel Address Word (CAW)



CAW fields are allocated for the following purposes:


CAW Bit

Position	Field Designation	Function
0-3	Protection Key	Forms the storage protection key for all commands associated with start I/O. This key must match the storage key.
4-7		Always zero.
8-31	Command Address	Designates the location of the first CCW in main storage.

30069

Figure 5. Channel Address Word

Channel Command Word




CCW Bit Position	Field Designation	Function	CCW Bit Position	Field Designation	Function
0-7	Command Code	Specify the operation to be performed. The two low-order bits, or when these bits are 00, the four low-order bits of the command code identify the operation to the channel. The channel distinguishes the operations: write, control, read, sense, or transfer in channel. Commands that initiate I/O operations cause all eight bits to be transferred to the storage control.	34	Suppress Length Indicator (SLI)	When set to one, an incorrect length condition is suppressed (except when the CCW count is not exhausted, channel end is present and data chaining is indicated). Should be set to one for restore, recalibrate, no-op, and some space count commands.
			35	Skip Flag	When set to one, specifies suppression of a transfer of information to storage during a read or sense operation. Checking takes place as though the information had been placed in storage. When bit 35 is zero, normal transfer of data takes place.
8-31	Data Address	Specifies address of the area associated with data transfer operations.			
32	Chain Data	When set to one, specifies chaining of data. Make sure the data rate of the I/O device permits chaining by the particular system model before using. See "Data Chaining".	36	Program Control- Interruption	When set to one, causes the channel to generate an interruption condition upon fetching the CCW. When bit 36 is zero, normal operation takes place.
33	Chain Command (CC)Flag	When set to one, and when the CD flag is zero, specifies chaining of commands. It causes the operation specified by the command code in the next CCW to be initiated on normal completion of the current operation.	37-39		Bit positions 37-39 of every CCW other than one specifying transfer in channel must contain zeros. Violation of this restriction generates the program-check condition.
			40-47		Not used.
			48-63	Count	Specify the number of 8-bit byte locations in the storage area designated by the data address.

25584A

Figure 6. Channel Command Word

Channel Status Word



CSW Bit Position	Field Designation	Function	CSW Bit Position	Field Designation	Function																		
0 -3	Protection Key	The storage protection key used in the chain of operations.	36	Channel End	Set at the end of each channel command.																		
4 -7	Not Used	Always zero.	37	Device End	Indicates that an access mechanism is free to be used.																		
8 -31	Command Address	An address eight positions higher than the address of the last CCW used.	38	Unit Check	Set whenever an unusual or error condition is detected. A sense I/O command may then be used to identify the condition.																		
32	Attention	Not used by 3830.	39	Unit Exception	Indicates an end-of-file has been detected during a read R0, read IPL, read CKD, read KD, read D, write KD, or a write D operation. It results from a data length of zero being detected in the count area of a record. When this condition is detected no data is transferred from the data area. If the key length is not zero, the key area is transferred.																		
33	Status Modifier	Set whenever a search high, search equal, or a search high or equal command has been executed and the condition satisfied. The status modifier is also set whenever the storage control is busy. This bit, in conjunction with the busy bit, signifies storage control busy.	40 -47	Channel Status	Indicate channel conditions as follows: <table border="1" style="font-size: small; width: 80%; margin-left: 20px;"> <tr><th>Bit</th><th>Designation</th></tr> <tr><td>40</td><td>Program-controlled interruption</td></tr> <tr><td>41</td><td>Incorrect length</td></tr> <tr><td>42</td><td>Program check</td></tr> <tr><td>43</td><td>Protection check</td></tr> <tr><td>44</td><td>Channel data check</td></tr> <tr><td>45</td><td>Channel control check</td></tr> <tr><td>46</td><td>Interface control check</td></tr> <tr><td>47</td><td>Chaining check</td></tr> </table>	Bit	Designation	40	Program-controlled interruption	41	Incorrect length	42	Program check	43	Protection check	44	Channel data check	45	Channel control check	46	Interface control check	47	Chaining check
Bit	Designation																						
40	Program-controlled interruption																						
41	Incorrect length																						
42	Program check																						
43	Protection check																						
44	Channel data check																						
45	Channel control check																						
46	Interface control check																						
47	Chaining check																						
34	Control Unit End	Set if a storage control busy status has been generated previously and the busy condition has been terminated. The device address associated with Control Unit End may be any device address recognized by the 3830.	48-63	Count	The residual count from the last CCW used.																		
35	Busy	Indicates that the selected device is busy. In conjunction with the status modifier bit, indicates the storage control is busy. It is set when a new command chain is initiated while the storage control is causing a track to be erased following a format write command.																					

25585A

Figure 7. Channel Status Word

INITIAL STATUS: The initial status byte is zero for test I/O and all non-immediate commands unless one or more of the following conditions exists:

- Storage control is busy.
- A status condition is pending. See "Pending Status."
- A unit check occurred.

Immediate commands (commands not requiring data transfer) present channel end and device end in initial status.

ENDING STATUS: In most cases, channel end and device end are presented as the normal ending sequence for an operation. The exceptions are noted in the individual command descriptions. See "Channel Commands."

If an error occurred during the operation, unit check will accompany the channel end-device end status.

PENDING STATUS: A pending status condition may exist for either the storage control or a disk drive. Status is pending for the storage control if:

- A disconnect was signaled after a command was issued, but before channel end status was accepted.
- Busy, channel end, or unit check status was stacked by the channel.
- Zero status, in response to a test I/O, was stacked by the channel.
- Control unit busy was presented to the channel.
- Unit check was detected for an operation after device end had been cleared.
- Device end status for a set sector command was stacked by the channel.

Status pending for the storage control causes the storage control to appear busy for all devices except the device for which the status condition exists. Unless it is busy, the storage control will request service to clear the pending status condition. Status is cleared when presented to, and accepted by, the channel.

Status is pending for a drive if:

- Channel end appears alone.
- Busy status is presented.
- The drive has gone from not ready to ready.

Status pending for a drive causes the storage control to request service when both the storage control and drive are not busy. The status is cleared when presented to, and accepted by, the channel.

CONTINGENT CONNECTION: A contingent connection is established in the storage control after the channel accepts a status byte containing unit check. The connection lasts until: (1) a command (other than test I/O or no-op) receives an initial status byte of zero for the storage control and device address that generated the unit check or, (2) a selective or system reset occurs.

During the contingent connection state, the storage control appears busy to all storage control and device addresses other than the address for which the contingent connection was established.

Program Status Word

Two program status words (PSW) are associated with 3830/3330 interrupt conditions: an "old" PSW which contains the status information of the system existing at the time on of the interrupt, and a current or "new" PSW which is used to control instruction sequencing and hold the status of the system in relation to the program being executed.


By storing the current PSW during an interruption, CPU status is preserved for subsequent inspection by the program. Loading a new PSW causes the state of the CPU to be initialized or changed to "branch to a new instruction sequence." If, at the conclusion of an interrupt routine, an instruction is executed that restores the old PSW as the current PSW, the system is restored to the state existing prior to the interruption, and the interrupted routine continues.

The format for the program status word is shown in Figure 8.

Command Chaining

The 3830/3330 has the ability to execute a series of channel commands as a result of a single start I/O

Program Status Word



PSW Bit Position	Field Designation		PSW Bit Position	Field Designation
0	Channel 0 mask	}	14	Wait state (W)
1	Channel 1 mask		15	Problem state (P)
2	Channel 2 mask		16-31	Interruption code
3	Channel 3 mask		32-33	Instruction length code (ILC)
4	Channel 4 mask		34-35	Condition code (CC)
5	Channel 5 mask			
6	Channel 6 mask		36	Fixed-point overflow mask
7	External mask		37	Decimal overflow mask
8-11	Protection key		38	Exponent underflow mask
12	Must be zero for System/370		39	Significance mask
13	Machine check mask (M)		40-63	Instruction address

25583A

Figure 8. Program Status Word

instruction; this method of operation is called command chaining. Command chaining is initiated by turning on bit 33 in the channel command word. The channel fetches a new CCW (specifying a new I/O operation) upon completion of the current CCW. The new I/O operation is automatically executed when the 3330 has completed the current operation and signaled device end to the channel.

The completion of the current CCW does not cause an I/O interrupt, and the count, indicating the amount of transferred data, is not available to the program.

Command chaining is normally used with all 3830/3330 channel programs. Time is available to execute command chaining functions in the gap area between record areas.

Certain restrictions regarding command sequence within a chain do exist. These restrictions, together with the individual command descriptions, are discussed in the "Channel Commands" section of this manual.

Data Chaining

Data transferred between main storage and the 3330 may be chained, which permits blocks of data to be transferred to or from noncontiguous areas of main storage.

Data chaining may be used to rearrange information as it is transferred between main storage and the 3330. It may also be used in conjunction with the skip flag to enable the program to place selected portions of a block of data in main storage.

When data chaining is specified (i. e. when bit 32 of the channel command word is on), the channel fetches a new CCW, which specifies a new storage location, upon completion of data transfer for the current channel command. Unless the command code specifies transfer-in-channel, the new CCW command code is ignored.

Data chaining occurs immediately after the last byte of data designated by the current CCW has been transferred to main storage or accepted by the 3330.

If both data chaining and command chaining are indicated in the channel command word, data chaining takes precedence and command chaining is ignored.

Note: Data chaining capabilities are dependent on several variable factors, including system type, I/O configuration, channel loading, etc. Because of these dependencies, read or write data chaining within record areas may cause unpredictable overruns or chaining checks. If these conditions are encountered or suspected, consult your IBM representative.

Branching in Channel Programs

Normally, the next CCW in a chain of channel commands is taken from an address eight positions higher than the address of the current CCW. This sequence can be modified in either of two ways:

1. If command chaining is specified in a search command, and execution of the command results

in a status modifier indication (search satisfied), the channel fetches the next CCW from a main storage location sixteen positions higher than the current channel command.

2. The transfer in channel command (TIC) may be used to modify the sequence of a chain of commands. The data address portion of the TIC CCW specifies the main storage location of the next channel command word. Therefore, the next CCW may be fetched from any valid main storage location.

These methods of modifying the sequence of a CCW chain provide branching capabilities within a channel program.

Unit Selection and Device Addressing

The I/O addresses of the 3830 and 3330 are designated by an eight bit binary number in an I/O instruction. These addresses consist of two parts: (1) the

storage control address (determined by the customer when the unit is initially installed) in the five high order bits and (2) the disk drive address (determined by the logical address plugs) specified in the three low order bits.

The storage control accepts any drive address from 000 through 111. If the specified drive is either not attached or off-line, the operation is terminated with unit check status. Multiple responses to an address, due to duplicate logical address plugs or hardware failures, also causes the operation to be terminated with unit check status.

Note: The addressing options provided in 3830s, coupled with addressing options provided by external switches, can cause difficulty in drive identification. For example, the same drive could be called 1A1, 2B1, 3C1, and 4C1 by system messages. This difficulty can be avoided by asking the CE installing the system to wire all interfaces identically. This causes addresses in the foregoing example to be the same; that is, 1A1, 2A1, 3A1, and 4A1.

Channel Commands

CONTROL COMMANDS

Control commands do not involve a transfer of data records between the storage control and main storage. However, in certain operations control bytes are transferred from main storage to the storage control. These bytes enable the operation to take place and are parity checked during transfer.

SEARCH COMMANDS

During the execution of search commands, the channel operates in write mode while the disk storage operates in read mode. The storage control compares the data coming from main storage against that coming from the drive. When the search criteria has been satisfied (for example, compared equal, high, etc.) the storage control returns a status modifier bit with channel end and device end. This bit causes the channel to skip the next CCW in the chain and fetch the next command from a storage location 16 positions higher than the current CCW.

Each search command operates on one record at a time. To search another record, the command must be reissued. This is normally done by chaining a TIC command to the search command, as follows:

```
Search Key Equal
TIC*-8
Read Data
```

If the search is unsuccessful, the TIC command following the search command causes the search to be repeated. When a search is successful, the status modifier causes the TIC command to be skipped and the read data command is executed.

At the end of every field searched, data validity is verified by the correction code bytes following the searched field. After the correction code check, the appropriate ending status is generated and presented to the channel.

If a data overrun or data check is detected, the storage control attempts recovery through use of command retry. If command retry is unsuccessful, channel end, device end, and unit check status are presented.

READ COMMANDS

A read command is used to transfer information from disk storage to the central processing unit. Read commands may operate in either single track or multiple track mode.

Note: Read IPL and read sector do not operate in multitrack mode.

On all read commands, the storage control checks the validity of each record area as it is transferred from the disk storage to the storage control. After the correction code bytes have been examined and data validity is established, the storage control sends an ending status byte of channel end and device end to the channel.

If a data overrun or data check is detected, the storage control normally attempts recovery through use of command retry. If command retry is unsuccessful or not used, channel end, device end, and unit check are presented to the channel.

WRITE COMMANDS

Formatting Write Commands

Formatting write commands are used to initialize tracks and records and establish the length of the areas within each record. Error correction code bytes are calculated and written after each area of a record.

The formatting write commands are:

- Write home address. (See Note in Write Home Address Channel Command Description).
- Write R0. (See Note in Write R0 Channel Command Description).
- Write count, key, and data.
- Write special count, key, and data.
- Erase.

The command prerequisites and file mask settings for these commands are explicit; any violation prevents command execution.

Format write commands may be chained together if each satisfies the required prerequisites. After the last format write command in a chain has been completed, the storage control causes the remaining portion of the track to be erased.

If a command (other than a format write command) is chained from a format write command, it is executed after the track has been padded. If the command is a control type command, the storage control utilizes the command retry function to free the channel while the track is being erased. If a new command chain is attempted before the end of the track is reached, a short control unit busy sequence (busy and status modifier bits) is presented to the channel. In this case, a control unit end signal is generated at the end of the track.

Update Write Commands

Update (non-formatting) write commands are used to update existing records and must operate on previously formatted tracks. Error correction code bytes are calculated and written after each key and/or data area in the record.

The update write commands are:

- Write data.
- Write key and data.

If a data overrun occurs during an update write operation (excluding the second and subse-

quent segments of an overflow record), the storage control attempts recovery through the use of command retry. If the retry is unsuccessful, channel end, device end, and unit check status are presented to the channel.

SENSE/TEST I/O COMMANDS

These commands are used to determine the status of the IBM 3830/3330 facility and identify the specific nature of errors or unusual conditions that have occurred.

Note: Since the test I/O "command" is not the result of the channel executing a CCW, its operation is explained at this time instead of with the other channel commands. A test I/O command (command code 0000 0000) is not written by the programmer. A command code of all 0's is considered invalid and causes a program check.

The test I/O command is generated automatically by the channel when the channel requires status information, or it is the result of processing a test I/O instruction. In either case it appears to the storage control as a command byte of all 0's and is treated as an immediate command. Test I/O requests the storage control to send all outstanding status information to the channel and, normally, presents an all-zero status byte. Stacked or pending status (if any) is presented in initial status.

CHANNEL COMMAND DESCRIPTIONS

COMMAND	COMMAND CODE				
	Multiple Track OFF		Multiple Track ON (if applicable)		
	Hexadecimal	Binary	Hexadecimal	Binary	
CONTROL	No Operation	03	0000 0011		
	Recalibrate	13	0001 0011		
	Seek	07	0000 0111		
	Seek Cylinder	0B	0000 1011		
	Seek Head	1B	0001 1011		
	Space Count	0F	0000 1111		
	Set File Mask	1F	0001 1111		
	Set Sector	23	0010 0011		
	Restore	17	0001 0111		
	Transfer in Channel	x8	xxxx 1000		
	Diagnostic Load	53	0101 0011		
	Diagnostic Write	73	0111 0011		
SEARCH	Home Address Equal	39	0011 1001	B9	1011 1001
	Identifier Equal	31	0011 0001	B1	1011 0001
	Identifier High	51	0101 0001	D1	1101 0001
	Identifier Equal or High	71	0111 0001	F1	1111 0001
	Key Equal	29	0010 1001	A9	1010 1001
	Key High	49	0100 1001	C9	1100 1001
Key Equal or High	69	0110 1001	E9	1110 1001	
READ	Home Address	1A	0001 1010	9A	1001 1010
	Count	12	0001 0010	92	1001 0010
	Record 0	16	0001 0110	96	1001 0110
	Data	06	0000 0110	86	1000 0110
	Key and Data	0E	0000 1110	8E	1000 1110
	Count, Key, and Data	1E	0001 1110	9E	1001 1110
	IPL	02	0000 0010		
Sector	22	0010 0010			
SENSE	Sense I/O	04	0000 0100		
	Read Reset Buffered Log	A4	1010 0100		
	Release	94	1001 0100		
	Reserve	B4	1011 0100		
	Read Diagnostic Status 1	44	0100 0100		
WRITE	Home Address	19	0001 1001		
	Record 0	15	0001 0101		
	Erase	11	0001 0001		
	Count, Key, and Data	1D	0001 1101		
	Special Count, Key, and Data	01	0000 0001		
	Data	05	0000 0101		
	Key and Data	0D	0000 1101		
Notes:					
x Not significant (data addresses should not exceed storage capacity).					

CONTROL

SEARCH

READ

SENSE

WRITE

Use of command codes other than those listed above presents unit check in initial status. A subsequent sense operation indicates command reject.

NO-OP

0 Command Code	7 8 Data Address	31 32 Flags	36 37 39 40 000	47 48	63 Count
Binary 0000 0011 Hex 03	Not checked for validity; should not exceed addressing capacity.	SLI flag (bit 34) should be on			Must be non-zero; zero count will cause a program check.

Chaining and Special Requirements: See following description.

NO-OP, an immediate command; causes no action at addressed device.

CHANNEL END is presented in initial status.

DEVICE END is presented in initial status.

INDISCRIMINATE USAGE must be avoided; a no-op resets orientation information causing all or part of records to be skipped.

EXAMPLE: a no-op inserted between read count and read data causes the following record's data to be read.

EXAMPLE: a no-op inserted between a command that reads the data field of record n-1 and a command that must process the count area of record n, may skip record n and process the count area of record n+1.

NO-OP CCW count field must not be zero.

SLI FLAG must be on to avoid incorrect length indication.

ZERO COUNT will set the program check bit (bit 42) in the CSW.

RECALIBRATE

0 Command Code	7 8 Data Address	31 32 Flags	36 37 39 000	40 47 48	63 Count
Binary 0001 0011 Hex 13	Not checked for validity, but should not exceed addressing capacity.	SLI flag (bit 34) should be on.			Must be non-zero. A zero count will cause a program check.

Chaining and Special Requirements: **Must be preceded by a Set File Mask permitting seek commands.**

RECALIBRATE causes addressed drive to seek to cylinder zero/head zero.

INITIAL STATUS byte normally zero; not processed as an immediate command.

CHANNEL END presented in ending status.

DEVICE END presented when drive positions access mechanism to cylinder zero/head zero.

FILE MASK must be set to allow seek commands.

SLI BIT must be on in recalibrate CCW to avoid incorrect length indication.

SEEK

0 Command Code	7 8 Data Address	31 Flags	32 36 000	37 39 40 47 48	63 Count
Binary 0000 0111 Hex 07	Specifies main storage location of the seek address.	Used at discretion of programmer.			Six

Chaining and Special Requirements: Must be preceded by a set file mask permitting seek commands.

SEEK transfers the six-byte seek address from channel to storage control.

INITIAL STATUS normally zero.

STORAGE CONTROL selects drive, moves access to proper cylinder and selects proper head.

ACCESS MOTION, if any, initiated after transfer of seek address

CCW COUNT > SIX: transfers six bytes of address information.

CCW COUNT < SIX: seek command is not executed; unit check, channel end, and device end are presented in ending status. A subsequent sense command indicates command reject.

VALID SEEK ADDRESS checked by storage control.

Bytes 0, 1, and 4 must be zero.

Bytes 2 and 3 must not exceed 410 (decimal).

Byte 5 must not exceed 18 (decimal).

INVALID SEEK ADDRESS: seek command is not executed; unit check, channel end, and device end are presented in ending status. A subsequent sense command indicates command reject.

PARITY ERROR detected in transfer of seek address: command is not executed; unit check, channel end, and device end presented in ending status. A subsequent sense command indicates bus-out parity error.

COMMAND EXECUTION does not require preceding CCW.

FILE MASK must be set to allow seeks, or unit check is presented in initial status.

CHANNEL END presented after transfer of seek address.

DEVICE END presented with channel end if no movement required.

DEVICE END presented after access is positioned if movement is required.

NOTE: Several successive seeks, without an intervening data read or write, may cause a seek incomplete condition in the storage control. The storage control uses its internal error recovery procedures to correct the failure. If it cannot correct the failure, unit check with equipment check and permanent error in the sense bytes is posted.

SEEK CYLINDER

0	7	8	31	32	36	37	39	40	47	48	63
Command Code	Data Address						Flags	000			Count
Binary 0000 1011 Hex 0B	Specifies main storage location of the seek address.						Used at discretion of programmer.				Six

Chaining and Special Requirements: Must be preceded by a set file mask permitting seek commands.

SEEK CYLINDER transfers the six-byte seek address from channel to storage control.

INITIAL STATUS normally zero.

STORAGE CONTROL selects drive, moves access to proper cylinder and selects proper head.

ACCESS MOTION, if any, initiated after transfer at seek address.

CCW COUNT > SIX transfers six bytes of address information.

CCW COUNT < SIX: seek cylinder command is not executed; unit check, channel end and device end are presented in ending status. A subsequent sense command indicates command reject.

VALID SEEK ADDRESS checked by storage control.

Bytes 0,1, and 4 must be zero.

Bytes 2 and 3 must not exceed 410 (decimal).

Byte 5 must not exceed 18 (decimal).

INVALID SEEK ADDRESS: seek cylinder command is not executed; unit check, channel end and device end are presented in ending status. A subsequent sense command indicates command reject.

PARITY ERROR detected in transfer of seek address: command is not executed; unit check, channel end and device end are presented in ending status. A subsequent sense command indicates bus-out parity error.

COMMAND EXECUTION does not require preceding CCW.

FILE MASK must be set to allow seeks, or unit check is presented in initial status.

CHANNEL END presented after transfer of seek address.

DEVICE END presented with channel end if no movement required.

DEVICE END presented after access is positioned if movement is required.

NOTE: Several successive seeks, without an intervening data read or write, may cause a seek incomplete condition in the storage control. The storage control uses its internal error recovery procedures to correct the failure. If it cannot correct the failure, unit check with equipment check and permanent error in the sense bytes is posted.

SEEK HEAD

0 Command Code	7 8 Data Address	31 Flags	36 000	37 39 40 47 48	63 Count
Binary 0001 1011 Hex 1B	Specifies main storage location of seek address.	Used at discretion of programmer.			Six

Chaining and Special Requirements: Must be preceded by a set file mask permitting seek head commands.

SEEK HEAD transfers seek address from channel to storage control.

INITIAL STATUS normally zero.

STORAGE CONTROL selects drive and proper head.

VALID SEEK ADDRESS required; however, only the head address specified in the sixth byte is significant (i.e. another cylinder address is ignored).
 Bytes 0,1 and 4 must be zero.
 Bytes 2 and 3 must not exceed 410 (decimal).
 Byte 5 must not exceed 18 (decimal).

INVALID SEEK ADDRESS: seek head command is not executed; unit check, channel end and device end are presented in ending status. A subsequent sense command indicates command reject.

PARITY ERROR detected in transfer of seek address; command is not executed; unit check, channel end and device end presented in ending status. A subsequent sense command indicates bus-out parity error.

COMMAND EXECUTION does not require preceding CCW.

FILE MASK must be set to allow head seeks, or unit check is presented in initial status.


CHANNEL END/DEVICE END presented after transfer of seek address.

SPACE COUNT

0	Command Code	7	8		31	32	36	37	39	40	47	48	63
				Data Address		Flags		000				Count	
	Binary 0000 1111 Hex 0F			Specifies main storage location of the key and data lengths of record to be recovered.		Used at discretion of programmer.						Three	

- Chaining and Special Requirements:
1. Cannot be chained from a format write or erase command.
 2. Must not be followed by a write, erase or set file mask command in the same chain.

SPACE COUNT allows bypassing of a defective count area on a track for recovering data in key and/or data areas following the defective area.



1. Searches for index.
2. Clocks thru gap 1, home address and gap 2.
3. Spaces over R0 count area.
4. Receives key and data length transfer from channel.
5. Sets an "end of count area" internal orientation state indicator.
6. Presents channel end and device end to channel.

Using the above:

- a. space count followed by a read key and data recovers or bypasses defective R0 count area.
- b. space count followed by a read CKD causes R1 to be read.

1. Orients at the beginning of next count area.
2. Spaces over the count area.
3. Receives key and data length transfer from channel.
4. Sets an "end of count area" internal orientation state indicator.
5. Presents channel end and device end to channel.

Using the above:

command chain (a) may be used to recover key and data areas of record N. (N≠0). Command chain (b) may be used to recover record N+1.

- | | |
|---------------------------------------------|---------------------------------------------|
| (a) Set Sector | (b) Set Sector |
| Search ID | Search ID |
| (record n-1) | (record n-1) |
| TIC*-8 | TIC*-8 |
| Space Count | Space Count |
| (must specify correct key and data lengths) | (must specify correct key and data lengths) |
| Read KD | Read CKD |

DATA TRANSFERRED FROM CHANNEL is used by the storage control as the key length (first byte) and data length (last two bytes) of the record to be recovered.

CCW COUNT > THREE: three bytes are transferred.

CCW COUNT < THREE: specified number of bytes is/are transferred.

NO BYTES TRANSFERRED: storage control assumes a value of zero. Read data and read key and data commands will receive unit exception status, and read CKD commands may detect data checks.

CHAINING REQUIREMENTS must be met; otherwise unit check is presented in initial status.

SET FILE MASK

0	Command Code	7	8	Data Address	31	32	36	37	39	40	47	48	Count	63
	Binary 0001 1111 Hex 1F			Specifies main storage location of mask byte.				Used at discretion of programmer.					One	

Chaining and Special Requirements: One set file mask command permitted in a CCW chain.

SET FILE MASK sets the write and seek masks which provide protection for 3330 data and defines command retry-PCI interaction.

Bit 0	Bit 1	Function	Bit 3	Bit 4	Function	Bit 5	Function	Bit 7	Function
0	0	Inhibit write home address and write R0.	0	0	Permit all seek commands.	0	Inhibit diagnostic write commands.	0	Not PCI fetch mode.
0	1	Inhibit all write commands.	0	1	Permit seek cylinder and seek head.	1	Permit diagnostic write commands.	1	PCI fetch mode. (The storage control presents unit check if command retry is used to recover from ECC uncorrectable data errors.)
1	0	Inhibit all format write commands.	1	0	Permit seek head.				
1	1	Permit all write commands.	1	1	Inhibit all seek commands and head switching.				

19475

Bits 2 and 6 must be zero, or unit check, channel end, and device end are presented in initial status.

COMMAND EXECUTION is allowable only once within a CCW chain. An attempt to issue more than one set file mask in a CCW chain causes a unit check in initial status.

COMMAND REJECT is indicated by a subsequent sense command.

FILE MASK RESET to 0's at end of CCW chain.

WRITE COMMANDS that violate file mask are not executed.

UNIT CHECK is presented in initial status.

COMMAND REJECT is indicated by a subsequent sense command.

SEEK COMMANDS that violate the file mask are not executed.

UNIT CHECK is presented in initial status.

FILE PROTECTED is indicated by a subsequent sense command.

MULTI-TRACK/OVERFLOW operations that violate the file mask indicate unit check and file protected.

CHANNEL END/DEVICE END are presented to the channel after transfer of mask byte.

SYSTEM OR SELECTIVE RESET resets the file mask to 0's.

START I/O executed after a reset without a set file mask CCW permits seek and write commands (except write home address and write R0).

SET SECTOR

0	7	8	31	32	36	37	39	40	47	48	63
Command Code		Data Address				Flags		000		Count	
Binary 0010 0011 Hex 23		Specifies main storage location of desired sector.				Used at discretion of programmer.				One	

Chaining and Special Requirements: None

SET SECTOR, used on block multiplexer channels, eliminates the necessity to maintain channel and storage control connection during rotational delay.

COMMAND EXECUTION transfers a sector number (128 possibilities) from main storage to storage control.

ANGULAR POSITIONS are checked for validity by the 3830.

***VALID ARGUMENT (0-127):**

1. Storage control presents channel end and disconnects.
2. Device end is signaled when angular position is reached and channel reconnects to continue chain.
3. If reconnection does not occur, the storage control attempts reconnection on subsequent revolutions.

ZERO ARGUMENT:

Storage control attempts reconnection just prior to index.

ARGUMENT > 127 < 255:

Channel end, device end and unit check presented in ending status. Command reject indicated in a subsequent sense command.

ARGUMENT = 255:

1. Command is treated as a no-op.
2. Channel end/device end presented in ending status.
3. Track orientation is destroyed.

*All valid arguments are adjusted by the storage control to compensate for channel reselection delay.

Programming Note:

1. The set sector command does not guarantee record orientation. The search commands must still be used for this function.
2. Indiscriminate use of set sector with multitrack search may result in missing the desired record. A set sector 0, read HA, search M/T sequence will avoid this exposure.

CONTROL

RESTORE

0 Command Code	7 8 Data Address	31 32 Flags	36 37 39 000	40 47 [Shaded]	48 63 Count
Binary 0001 0111 Hex 17	Not checked for validity; must not exceed addressing capacity.	SLI flag (bit 34) should be on.		[Shaded]	Must be non-zero. Zero count will cause a program check.

Chaining and Special Requirements: None

RESTORE is maintained primarily for compatibility with other IBM Direct Access Storage Devices and causes no action to be performed.

INITIAL STATUS normally zero.

CHANNEL END/DEVICE END immediately follows initial status.

SLI BIT must be on in the restore CCW to avoid incorrect length indication.

TRANSFER IN CHANNEL

0	7	8	31	32	36	37	39	40	47	48	63
Command Code		Data Address						Flags		Count	
Binary XXXX 1000 Hex X8		Specifies storage location from which next CCW will be taken.						Ignored		Ignored	

- Chaining and Special Requirements:
1. Cannot be first CCW designated by channel address word.
 2. One TIC command cannot transfer directly to another.

TRANSFER IN CHANNEL provides chaining capabilities for CCW's not located in adjacent main storage locations.

TIC DATA ADDRESS FIELD specifies next CCW to be fetched.

COMMAND EXECUTION does not initiate I/O operations or signal I/O device.

PROGRAM CHECK SIGNAL is generated when chaining requirements are not met or an invalid address is specified.
(TIC CCW data address field does not specify a double word boundary.)

ERROR DETECTION terminates chaining operations.

BIT POSITIONS 0-3 and 32-63 are ignored; bits 29-31 must be zero for double word boundary requirements.

NOTE: TIC is the only CCW that allows a zero count field; an incorrect length indication cannot occur since flags and count are ignored.

ASSEMBLER LANGUAGE notation TIC * -8 indicates an unconditional branch to the TIC storage address (*) minus a count of eight. TIC * -16 indicates an unconditional branch to the TIC storage address (*) minus a count of 16.

X = ignored

CONTROL

DIAGNOSTIC LOAD

0	7	8	31	32	36	37	39	40	47	48	63
Command Code		Data Address			Flags		000				Count
Binary 0101 0011 Hex 53		Specifies main storage location of control byte.			Used at discretion of programmer.						One

Chaining and Special Requirements: None

DIAGNOSTIC LOAD transfers a 512 byte block of data from storage control read-only storage to storage control buffer.

DATA BLOCK transferred is a functional microprogram diagnostic test.

INITIAL STATUS normally zero.

CONTROL BYTE specifying diagnostic microprogram ID number, is transferred from main storage to storage control.

***TRACK ADDRESS** (0-31) is specified by bits 0-4.

***SECTOR NUMBER** (0-7) is specified by bits 5-7.

COMMAND EXECUTION allows any drive address to be used with the storage control address.

READ DIAGNOSTIC STATUS 1 command transfers the diagnostic test from storage control buffer to main storage.

CHANNEL END/DEVICE END presented after transfer of diagnostic test to storage control buffer.

CAUTION

This command is intended for maintenance purposes only.
Any use other than that provided by IBM diagnostic programs may yield unpredictable results.

*Track address and sector number are references to the read only storage device attached to the 3830, not to a 3330 disk drive.

DIAGNOSTIC WRITE

0	7	8	31	32	36	37	39	40	47	48	63	
Command Code		Data Address						Flags	000			Count
Binary 0111 0011		Specifies main storage location of diagnostic test.						Used at discretion of programmer.				512
Hex 73												

Chaining and Special Requirements: File mask must be set to allow diagnostic write command (bit 5 = 1).

DIAGNOSTIC WRITE transfers a 512 byte diagnostic test from main storage to storage control.

INITIAL STATUS normally zero.

DATA TRANSFER COMPLETE: test execution begins.

TEST COMPLETE: 16 byte error code message is stored in storage control buffer.

COMPATIBILITY is verified by storage control comparing a key within the diagnostic test against the engineering level of the microprogram.

INVALID COMPARISON causes command termination; channel end, device end and unit check are presented in ending status.

CCW COUNT > 512: only 512 bytes are transferred.

CCW COUNT < 512: only the specified number of bytes is transferred, command is terminated and channel end, device end and unit check are presented in ending status.

ERROR CODE MESSAGE (16 bytes) is transferred from storage control buffer to main storage by a subsequent read diagnostic status 1 command.

CHANNEL END/DEVICE END presented after transfer of diagnostic test to the storage control and completion of test.

CAUTION

This command is intended for maintenance purposes only. Any use other than that provided by IBM diagnostic programs may yield unpredictable results.

SEARCH HOME ADDRESS EQUAL

0 Command Code	7 8 Data Address	31 32 Flags	36 37 39 000	40 47 48	63 Count
Binary 0011 1001 Hex 39 MT Binary 1011 1001 B9	Specifies main storage location of a cylinder number (CC) and head number (HH).	Used at discretion of programmer.			Four

Chaining and Special Requirements: None

SEARCH HOME ADDRESS EQUAL causes storage control to search for index.

INITIAL STATUS normally zero.

CYLINDER/HEAD NUMBERS from main storage and track home address area are compared by storage control when index is detected.

FLAG BYTE is not transferred or compared during command execution.

COMPARISON EQUAL: channel end/device end/status modifier are presented to the channel.

COMPARISON UNEQUAL: channel end/device end presented to the channel.

CCW COUNT > FOUR: only first four bytes used.

CHANNEL END/DEVICE END presented to terminate the command.

STATUS MODIFIER presented if comparison was equal.

CCW COUNT < FOUR: comparison of main storage and track data continues until CCW count is zero.

CHANNEL END/DEVICE END presented when home address and correction code bytes are read and checked.

STATUS MODIFIER presented if search is satisfied on short field.

MULTI-TRACK NOT USED: search is confined to one track; search continues (as long as channel repeats command) until search condition is satisfied or two index points are detected.

CHANNEL END/DEVICE END/UNIT CHECK presented to channel upon detection of second index.

MULTI-TRACK USED: causes search to continue (as long as channel repeats command); head number automatically increments at index until search condition is satisfied or end of cylinder is reached.

CHANNEL END/DEVICE END/UNIT CHECK presented to channel upon detection of end of cylinder.

SEARCH ID EQUAL

0	7	8	31	32	36	37	39	40	47	48	63
Command Code		Data Address				Flags		000		Count	
Binary 0011 0001 Hex 31 MT Binary 1011 0001 B1		Specifies main storage location of a five-byte record identifier (CC HH R).				Used at discretion of programmer.				Five	

Chaining and Special Requirements: None

SEARCH ID EQUAL compares the main storage ID and the count area ID. ID to be compared is next ID on the track (including R0).

INITIAL STATUS normally zero.

COMPARISON EQUAL: channel end/device end/status modifier presented to the channel.

COMPARISON UNEQUAL: channel end/device end presented to the channel.

CCW COUNT > FIVE: only first five bytes used.

CHANNEL END/DEVICE END presented to terminate command.

STATUS MODIFIER presented if comparison was equal.

CCW COUNT < FIVE: comparison of main storage and track data continues until CCW count is zero.

CHANNEL END/DEVICE END presented to channel when ID and correction code bytes are read and checked.

STATUS MODIFIER presented if search is satisfied on the short field.

MULTI-TRACK NOT USED: search is confined to one track; search continues (as long as channel repeats command) until search condition is satisfied or two index points are detected.

CHANNEL END/DEVICE END/UNIT CHECK presented to channel upon detection of second index.

MULTI-TRACK USED: causes search to continue (as long as channel repeats command); head number automatically increments at index until search condition is satisfied or end of cylinder is reached.

CHANNEL END/DEVICE END/UNIT CHECK presented to channel upon detection of end of cylinder.

SEARCH

SEARCH ID HIGH

0 Command Code	7 8 Data Address	31 32 Flags	36 37 39 000	40 47 [Shaded]	48 Count
Binary 0101 0001 Hex 51 MT Binary 1101 0001 D1	Specifies main storage location of a five byte record identifier (CC HH R).	Used at discretion of programmer.		[Shaded]	Five

Chaining and Special Requirements: None

SEARCH ID HIGH compares the main storage ID and the disk drive count area ID. ID to be compared is next ID on the track (including R0).

INITIAL STATUS normally zero.

COMPARISON HIGH: channel end/device end/status modifier presented to the channel. ID on drive is higher than ID in main storage.

COMPARISON NOT HIGH: channel end/device end presented to the channel.

CCW COUNT > FIVE: only first five bytes used.

CHANNEL END/DEVICE END presented to terminate command.

STATUS MODIFIER presented if comparison was equal.

CCW COUNT < FIVE: comparison of main storage and track data continues until CCW count is zero.

CHANNEL END/DEVICE END presented to channel when ID and correction code bytes are read and checked.

STATUS MODIFIER presented if search is satisfied on the short field.

MULTI-TRACK NOT USED: search is confined to one track; search continues (as long as channel repeats command) until search condition is satisfied or two index points are detected.

CHANNEL END/DEVICE END/UNIT CHECK presented to channel upon detection of second index.

MULTI-TRACK USED: causes search to continue (as long as channel repeats command); head number automatically increments at index until search condition is satisfied or end of cylinder is reached.

CHANNEL END/DEVICE END/UNIT CHECK presented to channel upon detection of end of cylinder.

SEARCH ID EQUAL OR HIGH

0	7	8	31	32	36	37	39	40	47	48	63	
Command Code		Data Address						Flags	000	Count		
Binary 0111 0001 Hex 71 MT Binary 1111 0001 F1		Specifies main storage location of five byte record identifier (CC HH R).						Used at discretion of programmer.		Five		

Chaining and Special Requirements: None

SEARCH ID EQUAL OR HIGH compares the main storage ID and the disk drive count area ID. ID to be compared is next ID on the track (Including R0).

INITIAL STATUS normally zero.

COMPARISON EQUAL OR HIGH: channel end/device end/status modifier presented to the channel. ID on drive is equal to or higher than ID in main storage.

COMPARISON NOT EQUAL OR HIGH: channel end/device end presented to the channel.

CCW COUNT > FIVE: only first five bytes used.

CHANNEL END/DEVICE END presented to terminate command.

STATUS MODIFIER presented if comparison was equal.

CCW COUNT < FIVE: comparison of main storage and track data continues until CCW count is zero.

CHANNEL END/DEVICE END presented to channel when ID and correction code bytes are read and checked.

STATUS MODIFIER presented if search is satisfied on the short field.

MULTI-TRACK NOT USED: search is confined to one track; search continues (as long as channel repeats command) until search condition is satisfied or two index points are detected.

CHANNEL END/DEVICE END/UNIT CHECK presented to channel upon detection of second index.

MULTI-TRACK USED: causes search to continue (as long as channel repeats command); head number automatically increments at index until search condition is satisfied or end of cylinder is reached.

CHANNEL END/DEVICE END/UNIT CHECK presented to channel upon detection of end of cylinder.

SEARCH KEY EQUAL

0	7	8	31	32	36	37	39	40	47	48	63
Command Code	Data Address			Flags	000				Count		
Binary 0010 1001 Hex 29 MT Binary 1010 1001 A9	Specifies main storage locations to which key is compared.			Used at discretion of programmer.					Equal to length of argument.		

Chaining and Special Requirements: None

SEARCH KEY EQUAL compares main storage key to key area read from track. Key to be compared is next key on track (excluding R0).

NOTE: When command is chained from search ID or read count, key compared is in same record as ID or count. Search key equal bypasses R0 unless chained from search ID command which searched R0 ID.

INITIAL STATUS normally zero.

COMPARISON EQUAL: channel end/device end/status modifier presented to the channel.

COMPARISON UNEQUAL: channel end/device end presented to the channel.

CCW COUNT > KL: search operation completed when key area is read.

CHANNEL END/DEVICE END terminates command.

STATUS MODIFIER presented if comparison was equal.

CCW COUNT < KL: track and main storage data comparison continues until CCW count is zero.

CHANNEL END/DEVICE END presented after key area and the following correction code bytes are read and checked.

STATUS MODIFIER presented if search was satisfied on the short field.

MULTI-TRACK NOT USED: search is confined to one track; search continues (as long as channel repeats command) until search condition is satisfied or two index points are detected.

CHANNEL END/DEVICE END/UNIT CHECK presented to channel upon detection of second index.

MULTI-TRACK USED: causes search to continue (as long as channel repeats command); head number automatically increments at index until search condition is satisfied or end of cylinder is reached.

CHANNEL END/DEVICE END/UNIT CHECK presented to channel upon detection of end of cylinder.

COMMAND EXECUTION on a record with zero KL does not set a status modifier. If followed by a chained read data command, the data area read is that of the next record.

SEARCH KEY HIGH

0	7	8	31	32	36	37	39	40	47	48	63
Command Code		Data Address				Flags		000		Count	
Binary 0100 1001 Hex 49 MT Binary 1100 1001 C9		Specifies main storage location to which key is compared.				Used at discretion of programmer.				Equal to length of argument.	

Chaining and Special Requirements: None

SEARCH KEY HIGH compares main storage key to key area read from track. Key to be compared is next key on track (excluding R0).

NOTE: When command is chained from search ID or read count, key compared is in same record as ID or count. Search key equal bypasses R0 unless chained from search ID command which searched R0 ID.

INITIAL STATUS normally zero.

COMPARISON HIGH: channel end/device end/status modifier presented to the channel. Key on drive is higher than main storage argument.

COMPARISON NOT HIGH: channel end/device end presented to the channel.

CCW COUNT > KL: search operation completed when key area is read.

CHANNEL END/DEVICE END terminates command.

STATUS MODIFIER presented if comparison was equal.

CCW COUNT < KL: track and main storage data comparison continues until CCW count is zero.

CHANNEL END/DEVICE END presented after key area and the following correction code bytes are read and checked.

STATUS MODIFIER presented if search was satisfied on the short field.

MULTI-TRACK NOT USED: search is confined to one track; search continues (as long as channel repeats command) until search condition is satisfied or two index points are detected.

CHANNEL END/DEVICE END/UNIT CHECK presented to channel upon detection of second index.

MULTI-TRACK USED: causes search to continue (as long as channel repeats command) head number automatically increments at index until search condition is satisfied or end of cylinder is reached.

CHANNEL END/DEVICE END/UNIT CHECK presented to channel upon detection of end of cylinder.

COMMAND EXECUTION on a record with zero KL does not set a status modifier. If followed by a chained read data command, the data area read is that of the next record.

SEARCH

SEARCH KEY EQUAL OR HIGH

0	7	8	31	32	36	37	39	40	47	48	63
Command Code		Data Address						Flags	000	Count	
Binary 0110 1001 Hex 69 MT Binary 1110 1001 E9		Specifies main storage locations to which key is compared.						Used at discretion of programmer.		Equal to length of argument.	

Chaining and Special Requirements: None

SEARCH KEY EQUAL OR HIGH compares main storage key to key area read from track. Key to be compared is next key on track (excluding R0).

NOTE: When command is chained from search ID or read count, key compared is in same record as ID or count. Search key equal bypasses R0 unless chained from search ID command which searched R0 ID.

INITIAL STATUS normally zero.

COMPARISON EQUAL OR HIGH: channel end/device end/status modifier presented to the channel. Key on drive is equal to or higher than main storage argument.

COMPARISON NOT EQUAL OR HIGH: channel end/device end presented to the channel.

CCW COUNT > KL: search operation completed when key area is read.

CHANNEL END/DEVICE END terminates command.

STATUS MODIFIER presented if comparison was equal.

CCW COUNT < KL: track and main storage data comparison continues until CCW count is zero.

CHANNEL/END DEVICE END presented after key area and the following correction code bytes are read and checked.

STATUS MODIFIER presented if search was satisfied on the short field.

MULTI-TRACK NOT USED: search is confined to one track; search continues (as long as channel repeats command) until search condition is satisfied or two index points are detected.

CHANNEL END/DEVICE END/UNIT CHECK presented to channel upon detection of second index.

MULTI-TRACK USED: causes search to continue (as long as channel repeats command); head number automatically increments at index until search condition is satisfied or end of cylinder is reached.

CHANNEL END/DEVICE END/UNIT CHECK presented to channel upon detection of end of cylinder.

COMMAND EXECUTION on a record with zero KL does not set a status modifier. If followed by a chained read data command, the data area read is that of the next record.

READ HOME ADDRESS

0 Command Code	7 8 Data Address	31 32 Flags	36 37 39 000	40 47 [Shaded]	48 Count
Binary 0001 1010 Hex 1A MT Binary 1001 1010 9A	Specifies main storage location where home address is to be stored.	Used at discretion of programmer.		[Shaded]	Five

Chaining and Special Requirements: None

READ HOME ADDRESS transfers the F CC HH bytes of the home address area to main storage.

INITIAL STATUS normally zero.

DATA VALIDITY is verified by correction code bytes following the home address area.

DATA OVERRUN/DATA CHECK, if detected, causes storage control to attempt recovery by command retry.

COMMAND RETRY, if unsuccessful, signals unit check to the channel.

PARITY BIT is added to each byte prior to transferring byte to the channel.

CHANNEL END/DEVICE END presented to the channel at completion of correction code check of home address.

READ

READ COUNT

0 Command Code	7 8 Data Address	31 Flags	32 36 000	37 39 40 47	48 63 Count
Binary 0001 0010 Hex 12 MT Binary 1001 0010 92	Specifies main storage location where first byte of count data is to be transferred.	Used at discretion of programmer.			Eight

Chaining and Special Requirements: None

READ COUNT transfers the eight bytes (CC HH R KL DL DL) of the next count area encountered on the track (excluding R0) from disk storage to main storage.

INITIAL STATUS normally zero.

DATA VALIDITY is verified by correction code bytes following the count area.

DATA OVERRUN/DATA CHECK, if detected, initiates a storage control recovery attempt by command retry.

COMMAND RETRY, if unsuccessful, signals unit check to the channel.

PARITY BIT is added to each byte prior to transferring byte to the channel.

CHANNEL END/DEVICE END are signaled to the channel at completion of the correction code check.

READ R0

0	7	8	31	32	36	37	39	40	47	48	63
Command Code		Data Address					Flags		000	Count	
Binary 0001 0110 Hex 16 MT Binary 1001 0110 96		Specifies main storage location where first byte of R0 count data is to be transferred.					Used at discretion of programmer.			Specifies number of count, key, and data bytes to be read.	

Chaining and Special Requirements: None

READ R0 transfers count, key and data areas of R0 from disk storage to the channel.

INITIAL STATUS normally zero.

STORAGE CONTROL searches for index, clocks through gap 1, home address, and gap 2.

DATA TRANSFER of the R0 count area is initiated by storage control.

DATA VALIDITY is verified by correction code bytes following each area.

DATA OVERRUN/DATA CHECK, if detected, initiates a storage control recovery attempt by command retry.

NOTE: If a correctable data error (error burst of 11 bits or less) is detected in the data area, unit check is signaled to the channel.

COMMAND RETRY, if unsuccessful, signals unit check to the channel at the end of the area in which the error occurred.

PARITY BIT is added to each byte prior to transferring byte to the channel.

COMMAND EXECUTION is accomplished immediately if read R0 is chained from a search home address or read home address command; the storage control will not search for index in these cases.

CHANNEL END/DEVICE END are presented to the channel at completion of the correction code check of the data area.

READ

READ DATA

0	Command Code	7	8	Data Address	31	32	36	37	39	40	47	48	Count	63	
	Binary 0000 0110 Hex 06 MT Binary 1000 0110 86	Specifies main storage location where first byte of data is to be transferred.					Used at discretion of programmer.						Specifies number of bytes to be read.		

Chaining and Special Requirements: None

READ DATA transfers the data area of a record from disk storage to main storage. The data read is:

1. data area of record read by search ID or search key command from which read command is chained.
2. data area of record read by read count command from which command is chained.
3. data area of record following next count area on the track (excluding R0).

INITIAL STATUS normally zero.

DATA VALIDITY is verified by correction code bytes following each area.

DATA OVERRUN/DATA CHECK, if detected, initiates a storage control recovery attempt by command retry.

NOTE: If a correctable data error (error burst of 11 bits or less) is detected in the data area, unit check is signaled to the channel.

COMMAND RETRY, if unsuccessful, signals unit check to the channel.

PARITY BIT is added to each byte prior to transferring byte to the channel.

CHANNEL END/DEVICE END are presented to the channel at completion of the correction code check of the data area.

READ KEY and DATA

0	7	8	Data Address	31	32	36	37	39	40	47	48	63		
Command Code									Flags			Count		
Binary 0000 1110 Hex 0E MT Binary 1000 1110 8E			Specifies main storage location where first byte of key data is to be transferred.						Used at discretion of programmer.			Specifies the number of key and data area bytes to be read.		

Chaining and Special Requirements: None

READ KEY AND DATA transfers key and data areas of a record from disk storage to main storage. The key and data are:

1. key and data area of record read by search ID command from which read key and data is chained.
2. key and data areas of record read by read count command from which read key and data is chained.
3. key and data areas of record following next count area on the track (excluding R0).

INITIAL STATUS normally zero.

DATA VALIDITY is verified by correction code bytes following each area.

DATA OVERRUN/DATA CHECK, if detected, initiates a storage control recovery attempt by command retry.

NOTE: If a correctable data error (error burst of 11 bits or less) is detected in the data area, unit check is signaled to the channel.

COMMAND RETRY, if unsuccessful, signals unit check to the channel at the end of the area in which the error occurred.

KEY LENGTH = ZERO: command operates as a read data command.

PARITY BIT is added to each byte prior to transferring byte to the channel.

CHANNEL END/DEVICE END are presented to the channel at completion of the correction code check of the data area.

READ

READ COUNT, KEY, and DATA

0 Command Code 7	8 Data Address 31	32 Flags 36	37 000 39	40 47	48 Count 63
Binary 0001 1110 Hex 1E MT Binary 1001 1110 9E	Specifies main storage location where first byte of count data is to be transferred.	Used at discretion of programmer.			Specifies the number of count, key, and data bytes to be read.

Chaining and Special Requirements: None

READ COUNT, KEY, AND DATA transfers the next record encountered on the track from disk storage to main storage (excluding RO).

INITIAL STATUS normally zero.

DATA VALIDITY is verified by correction code bytes following each area.

DATA OVERRUN/DATA CHECK, if detected, initiates a storage control recovery attempt by command retry.

NOTE: If a correctable data error (error burst 11 bits or less) is detected in the data area, unit check is signaled to the channel.

COMMAND RETRY, if unsuccessful, signals unit check to the channel at the end of the area in which the error occurred.

PARITY BIT is added to each byte prior to transferring byte to the channel.

CHANNEL END/DEVICE END are signaled to the channel at completion of the correction code check of the data area.

READ IPL

0	7	8	31	32	36	37	39	40	47	48	63	
Command Code		Data Address						Flags 000				Count
Binary 0000 0010 Hex 02		Specifies main storage location where first byte of data is to be transferred.						Used at discretion of programmer.				Specifies number of bytes to be transferred.

Chaining and Special Requirements: Must not be preceded by a set file mask in the same chain.

READ INITIAL PROGRAM LOAD causes storage control to recalibrate to cylinder 0, head 0 of selected drive and search for index.

DATA AREA read, after index is detected, is the first record after R0.

COMMAND INITIATION is normally accomplished by setting the direct access storage device address in the load unit switches and pressing IPL key on console.

INITIAL STATUS normally zero.

DATA VALIDITY is verified by correction code bytes following the data area.

DATA OVERRUN/DATA CHECK, if detected, initiates a storage control recovery attempt by command retry.

NOTE: If a correctable data (error burst 11 bits or less) is detected in the data area, unit check is signaled to the channel.

COMMAND RETRY, if unsuccessful, signals unit check to the channel.

PARITY BIT is added to each byte prior to transferring byte to the channel.

CHANNEL END/DEVICE END are signaled to the channel at completion of the correction code check.

READ SECTOR

0 Command Code	7 8 Data Address	31 32 Flags	36 37 39 000	40 47 [Shaded]	48 63 Count
Binary 0010 0010 Hex 22	Specifies the main storage location where sector number is to be stored.	Used at discretion of programmer.		[Shaded]	One

Chaining and Special Requirements: None

READ SECTOR transfers one byte of data from storage control to main storage.

INITIAL STATUS normally zero.

BYTE TRANSFERRED contains sector number required to access the last record processed.

NOTE: If a drive power on sequence or system reset occurred, or a seek or set sector command was executed after a record was processed, this byte will be zero. If the last record processed was an overflow record, the angular position is that of the last segment.

COMMAND EXECUTION resets orientation information in the storage control.

CHANNEL END/DEVICE END presented after sector number is transferred.

SENSE I/O

0 Command Code	7 8 Data Address	31 32 Flags	36 37 39 000	40 47 [Shaded]	48 Count
Binary 0000 0100 Hex 04	Specifies storage location where bytes are to be transferred.	Used at discretion of programmer.		[Shaded]	Twenty-four

Chaining and Special Requirements: None

SENSE I/O transfers twenty-four bytes of sense information from the storage control to the channel.

INITIAL STATUS normally zero.

DESCRIBES:

UNIT CHECK STATUS

CURRENT STATUS of the device that performed operation, and

SYSTEM ERROR RECOVERY information.

UNIT CHECK should always be followed by a sense command, whether or not sense information is used; otherwise, expected future interrupts may not occur and some I/O access paths may be unavailable.

CHANNEL END/DEVICE END presented after sense bytes are transferred.

See "Sense Bytes" for a description of the sense information pertaining to 3830/3330 operations.

SENSE

READ AND RESET BUFFERED LOG

0 Command Code	7 8 Data Address	31 32 Flags	36 37 39 000	40 47 [Shaded]	48 63 Count
Binary 1010 0100 Hex A4	Specifies main storage location of first error byte or usage information.	Used at discretion of programmer.		[Shaded]	Twenty-four

Chaining and Special Requirements: None

READ AND RESET BUFFERED LOG transfers 24 bytes of usage or error information from storage control to the channel.

INITIAL STATUS normally zero.

USAGE/ERROR INFORMATION, generated and available when their respective counters overflow, pertains to the storage control addressed by start I/O and the disk storage drive identified in sense byte 4.

COUNTERS reset after data transfer.

CHANNEL END/DEVICE END presented after data transfer.

See "Statistical Usage/Error Recording"

DEVICE RELEASE

0 Command Code	7 8 Data Address	31 Flags	36 000	37 39 40	47 48 Count
Binary 1001 0100 Hex 94	Specifies main storage location where sense bytes are to be transferred.	Used at discretion of programmer.			Twenty-four

Chaining and Special Requirements: Must not be preceded by a set file mask in the same chain.
Two channel switch or two channel switch additional feature must be installed.

DEVICE RELEASE terminates reservation of the addressed drive.

INITIAL STATUS normally zero.

SENSE I/O command functions are performed by a device release command, i.e. 24 bytes of sense information are transferred to the channel.

NORMAL BUSY conditions cause command rejection; busy bit is set in the CSW.

ABNORMAL FILE status conditions (file unsafe, off-line, etc.) do not halt command execution.

CHANNEL END/DEVICE END presented after sense bytes are transferred.

UNIT CHECK, causing command rejection, is presented if:
Two channel switch or two channel switch additional feature is not installed in storage control.
Set file mask precedes command in the same chain.

DEVICE RESERVE

0 Command Code	7 8 Data Address	31 32 Flags	36 37 39 000	40 47 [Shaded]	48 Count
Binary 1011 0100 Hex B4	Specifies main storage location where sense bytes are to be transferred.	Used at discretion of programmer.		[Shaded]	Twenty-four

Chaining and Special Requirements: Must not be preceded by a set file mask in the same chain.
Two channel switch or two channel switch additional feature must be installed.

DEVICE RESERVE command reserves the addressed drive to the channel issuing the command.

INITIAL STATUS normally zero.

RESERVATION MAINTAINED until either a device release command or a system reset is performed by the channel.

SENSE INFORMATION (twenty-four bytes) is transferred to the channel.

NORMAL BUSY CONDITIONS cause a command reject; busy bit is set in the CSW.

ABNORMAL FILE STATUS conditions (e.g. file unsafe, off-line, etc.) do not halt command execution.

CHANNEL END/DEVICE END presented after sense byte transfer.

UNIT CHECK, causing command rejection, is presented if:

- Two channel switch or two channel switch additional feature is not installed in storage control.
- Set file mask command precedes command in same chain.

READ DIAGNOSTIC STATUS 1

0	7	8	31	32	36	37	39	40	47	48	63
Command Code		Data Address				Flags		000		Count	
Binary 0100 0100 Hex 44		Specifies main storage location where data accumulated during prior diagnostic load or diagnostic write is to be stored.				Used at discretion of programmer.				16 or 512	

Chaining and Special Requirements: None

READ DIAGNOSTIC STATUS 1 may perform either of two functions:

COMMAND FOLLOWS A DIAGNOSTIC WRITE COMMAND:

ERROR CODE MESSAGE (16 bytes) transferred from storage control buffer to main storage.

CCW COUNT FIELD should specify 16 bytes.

CHANNEL END/DEVICE END presented after transfer.

COMMAND FOLLOWS A DIAGNOSTIC LOAD COMMAND:

DIAGNOSTIC TEST (512 bytes) transferred from storage control buffer to main storage.

CCW COUNT FIELD should specify 512 bytes.

CHANNEL END/DEVICE END presented after transfer.

INITIAL STATUS normally zero.

DIAGNOSTIC LOAD/DIAGNOSTIC WRITE must precede the read diagnostic status 1 command, otherwise sixteen bytes of data are transferred from storage control buffer area which normally contains the error message.

CHANNEL END/DEVICE END are presented after data transfer.

CAUTION

This command is intended for maintenance purposes only. Any use other than that provided by IBM diagnostic programs may yield unpredictable results.

S M S

WRITE HOME ADDRESS

0 Command Code	7 8 Data Address	31 32 Flags	36 37 39 000	40 47 [Shaded]	48 Count
Binary 0001 1001 Hex 19	Specified main storage location of home address bytes (F CC HH).	Used at discretion of programmer.		[Shaded]	Five

Chaining and Special Requirements: Must be preceded by a set file mask permitting write home address commands.

WRITE HOME ADDRESS establishes track identity, a prerequisite for data operations on that track.

INITIAL STATUS normally zero.

STORAGE CONTROL orients on index, writes gap 1, home address, and ECC bytes.

FLAG BYTE, transferred from main storage. (Bit 5 must be zero.)

CCW COUNT < FIVE: 3830 records 0's until five bytes are written.

CCW COUNT > FIVE: First five bytes are written.

CHAINING REQUIREMENTS must be met; otherwise unit check is presented in initial status.

CHANNEL END/DEVICE END presented after ECC bytes are written.

NOTE: Home address is normally prewritten at the IBM plant. The use of this command should be limited to identifying defective tracks and assigning alternate tracks. Utility programs are available to perform this function.

WRITE R0

0	Command Code	7	8	Data Address	31	32	36	37	39	40	47	48	63
						Flags		000				Count	
	Binary 0001 0101 Hex 15			Specifies main storage location of R0 count, key and data bytes.		Used at discretion of programmer.						Specifies total number of bytes in R0 count, key and data areas.	

Chaining and Special Requirements: Must be chained from a successful write home address or search home address equal command.

WRITE R0 causes specified data in main storage to be written on selected drive.

INITIAL STATUS normally zero.

COUNT AREA is made up of the first eight bytes from main storage.

NOTE: The flag byte is generated by the storage control; the remaining data is written in the key and data areas as specified by the KL and DL bytes in the count area.

CORRECTION CODE BYTES are written by the storage control at the end of each record area.

CCW COUNT FIELD specifies the number of bytes (8 + KL + DL) to be transferred from main storage to drive.

CCW COUNT < 8 + KL + DL: storage control writes 0's in remainder of record.

CHAINING REQUIREMENTS must be met; otherwise unit check is presented in initial status.

CHANNEL END/DEVICE END is signaled after correction code bytes are written for the data area.

NOTE: Record zero is normally written on the disk pack at the IBM plant. The use of this command should be limited to identifying defective tracks and assigning alternate tracks. Alternate and defective tracks are identified in the R0 count area instead of the normal CCH bytes. Utility programs are available to perform these functions. Proper operation with Operating System (OS) requires an eight-byte field in R0.

ERASE

0 Command Code	7 8 Data Address	31 Flags	32	36	37 39 000	40	47	48 Count	63
Binary 0001 0001 Hex 11	Specifies main storage location where count, key, and data areas of the record are located.	Used at discretion of programmer.						Specifies number of bytes in count, key, and data areas of the record.	

Chaining and Special Requirements: Must be chained from either write R0, write CKD, *search ID equal or *search key equal.

ERASE writes count, key, and data areas on selected drive.

ZEROS are written in each area.

INITIAL STATUS normally zero.

CHANNEL END/DEVICE END are signaled at the end of the data area. Remainder of track is padded with 0's.

ERASED RECORD and all records that follow on the track are unrecoverable.

CHAINING REQUIREMENTS must be met; otherwise unit check is presented in initial status.

FORMAT WRITE command must not be chained from an erase command.

*Search commands must compare equal on all bytes of the searched field.

WRITE COUNT, KEY, and DATA

0	7	8	31	32	36	37	39	40	47	48	63		
Command Code		Data Address						Flags		000		Count	
Binary 0001 1101 Hex 1D		Specifies main storage location where count, key and data bytes of record are located.						Used at discretion of programmer.				Specifies total number of bytes in count, key, and data areas.	

Chaining and Special Requirements: Must be chained from either write R0, write CKD, *search ID equal or *search key equal.

WRITE COUNT, KEY, AND DATA causes specified data in main storage to be written on selected drive.

INITIAL STATUS normally zero.

COUNT AREA is made up of the first eight bytes from main storage.

FLAG BYTE is generated by storage control; the remaining data is written in the key and data areas as specified by the KL and DL bytes in the count area.

CORRECTION CODE BYTES are written by the storage control at the end of each record area.

CCW COUNT FIELD specifies number of bytes (8 + KL + DL) to be transferred from main storage to the drive.

CCW COUNT < 8 + KL + DL: storage control writes 0's in the remainder of the record.

READ DATA/READ KEY AND DATA may be inserted between search CCW and write CKD CCW.

CHAINING REQUIREMENTS must be met; otherwise unit check is presented in initial status.

CHANNEL END/DEVICE END are signaled to the channel after correction code bytes are written for the data area.

*Search commands must compare equal on all bytes of the searched field.

WRITE SPECIAL COUNT, KEY, and DATA

0	7	8	31	32	36	37	39	40	47	48	63		
Command Code		Data Address						Flags		000		Count	
Binary 0000 0001 Hex 01		Specifies main storage location where count, key and data areas of the record are located.						Used at discretion of programmer.				Specifies number of bytes in the count, key, and data areas of the record segment.	

Chaining and Special Requirements: Must be chained from a write RO, write CKD, *search ID equal or *search key equal command.

WRITE SPECIAL COUNT, KEY, AND DATA formats a segment of an overflow record; last segment is written by a normal write CKD command.

INITIAL STATUS normally zero.

COUNT AREA is made up of the first eight bytes from main storage.

FLAG BYTE contains a 1 in bit position 4; generated and written by the storage control, this bit indicates that another part of the record is located on the next track.

CORRECTION CODE BYTES are written by the storage control at the end of each record area.

CCW COUNT FIELD specifies number of bytes (8 + KL + DL) to be transferred from main storage to the drive.

CCW COUNT < 8 + KL + DL: storage control writes 0's in the remainder of the record.

READ DATA/READ KEY AND DATA may be inserted between search CCW and write special CKD CCW.

CHAINING REQUIREMENTS must be met; otherwise unit check is presented in initial status.

CHANNEL END/DEVICE END are signaled to the channel after correction code bytes are written for the data area.

*Search commands must compare equal on all bytes of the searched field.

WRITE DATA

0	Command Code	7	8	Data Address	31	32	36	37	39	40	47	48	Count	63
	Binary 000 0101 Hex 05	Specifies main storage location of data used to update record.			Used at discretion of programmer.						Specifies number of data bytes to be written.			

Chaining and Special Requirements: Must be chained from a *search ID equal or *search key equal command.

WRITE DATA performs normal record updating after track formatting.

INITIAL STATUS normally zero.

COMMAND EXECUTION causes specified data in main storage to be written in data area of selected record.

CORRECTION CODE BYTES are written by the storage control at the end of the data area.

NUMBER OF BYTES WRITTEN:

1. is specified in the count field of the write data CCW.
2. may be less than data length specified in formatted record.

CCW COUNT < COUNT AREA DL: Storage control writes 0's in remaining data area, writes ECC bytes and presents channel end/device end to channel.

CCW COUNT > COUNT AREA DL: Storage control writes only the number of bytes indicated in the count area DL, then writes ECC bytes.

CHAINING REQUIREMENTS must be met; otherwise unit check is presented in initial status.

CHANNEL END/DEVICE END are signaled to the channel after correction code bytes are written for the data area.

*Search commands must compare equal on all bytes at the searched field.

WRITE

WRITE KEY and DATA

0	7	8	31	32	36	37	39	40	47	48	63	
Command Code		Data Address						Flags	000			Count
Binary 0000 1101 Hex 0D		Specifies main storage location of data to be used to update record.						Used at discretion of programmer.				Specifies number of key and data bytes to be written.

Chaining and Special Requirements: Must be chained from a *search ID equal command.

WRITE KEY AND DATA is used for record updating after track formatting.

INITIAL STATUS normally zero.

COMMAND EXECUTION causes data from main storage to be written in key and data area of selected record.

CORRECTION CODE BYTES are written by the storage control at the end of each area.

NUMBER OF BYTES WRITTEN:

1. is specified in the count field of the write key and data CCW.
2. may be less than key and data length specified in formatted record.

CCW COUNT < KL/DL BYTE COUNT: Storage control writes 0's in the remaining areas, writes ECC bytes, and presents channel end/device end to channel.

CCW COUNT > KL/DL BYTE COUNT: Channel end/device end are presented after the number of bytes indicated in the count area KL/DL and ECC bytes are written.

CHAINING REQUIREMENTS must be met; otherwise unit check is presented in initial status.

CHANNEL END/DEVICE END presented after ECC bytes have been written for the data area.

*Search command must compare equal on all bytes of the searched field.

CHANNEL PROGRAMS

The following channel programs are typical examples of how CCW's are arranged to format, read, and write records on the 3830/3330 facility. The examples given do not include the CPU program, which would be used to initiate the channel program.

Unless otherwise noted, all numbers used are hexadecimal.

Example 1: Format track 6A on head 8 with home address, record 0, and records R1, R2 and R3 for customer records. Assuming R0 has a key length of zero and a data length of eight bytes, and that R1, R2, and R3 have a key length of 6 bytes and a data length of 03E8 (1000 bytes).

The channel program used is:

- Seek
- Set File Mask
- Set Sector
- Write Home Address
- Write Record Zero
- Write CKD
- Write CKD
- Write CKD

SEEK

0	7	8	31	32	36	37	39	40	47	48	63
Command Code		Data Address		Flags		000				Count	
Binary 0000 0111 Hex 07		C C H H 03E8 = 00 00 00 6A 00 08		01000		000				0006	
<p>Comments: The seek command is used to position the access at the desired cylinder and to select the proper head. All seek commands transfer six bytes of data from main storage to the storage control. (Thus the byte count of six.) The first two bytes of the seek address are always 0's, the cylinder number (6A) is specified in the third and fourth bytes, and bytes five and six specify the desired head (00 08 at 03EC and 03ED).</p>											

SET FILE MASK

0	7	8	31	32	36	37	39	40	47	48	63
Command Code		Data Address		Flags		000				Count	
Binary 0001 1111 Hex 1F		03EE = C0		01000		000				0001	
<p>Comments: The set file mask command is used to specify the types of operations that can be performed in this channel program. The mask byte in this case (1100 000 at address 03EE) permits all write and seek commands. The mask is reset to zero at the beginning of each chain of commands.</p>											

SET SECTOR

0	7	8	31	32	36	37	39	40	47	48	63
Command Code		Data Address				Flags		000		Count	
Binary 0010 0011 Hex 23		1390 = 00				01000		000		0001	
<p>Comments: Execution of a set sector command, with an argument at zero, orients the track to index. During the time that the 3830 is waiting for index, the channel is available to perform other operations on other drives.</p>											

WRITE HOME ADDRESS

0	7	8	31	32	36	37	39	40	47	48	63
Command Code		Data Address				Flags		000		Count	
Binary 0001 1001 Hex 19		F C C H H 03EF = 00 00 6A 00 08				01000		000		0005	
<p>Comments: The write home address command creates the home address area on the track. The home address area is five bytes long (F C C H H). When formatting tracks, the flag byte is normally zero. The cylinder number is in the CC bytes, and the head number is in the HH bytes.</p> <p>Write home address is the only write command in which the flag byte is transferred from main storage. The flag byte is generated automatically by the 3830 for other write commands. See note in Write Home Address.</p>											

WRITE R0

0	7	8	31	32	36	37	39	40	47	48	63
Command Code		Data Address				Flags		000		Count	
Binary 0001 0101 Hex 15		C C H H R K L D L D L 07D0 = 00 6A 00 08 00 00 00 08 07D8 = 00 00 00 00 00 00 00 00				01000		000		0010	
<p>Comments: Following the home address area is record 0. The write R0 command writes a count area, a key area (if the key length specified is not zero), and a data area whose length is dependent upon the value specified in the DL bytes of the count area. In this example, the data address is at 07D0 and at byte count of sixteen is specified.</p> <p>Since the key length specified is zero, address 07D5 is coded 00 and no key area is written. The data length is eight bytes so addresses 07D6 and 07D7 are coded 00 08, and the data in the following eight main storage locations is written in the data area.</p> <p>Note that the byte count in the write R0 command is sixteen and the 3830 requested sixteen bytes (eight for the count area and eight for the data area). Therefore no incorrect length error is generated.</p> <p>The flag byte preceding the count area is generated by the storage control and is not included in the CCW count. See note in Write R0.</p>											

WRITE CKD

0	7	8	31	32	36	37	39	40	47	48	63		
Command Code		Data Address						Flags	000			Count	
Binary	0001 1101	R1	C	C	H	H	R	KL	DL	DL			
Hex	1D	0BB8 = 00 6A 00 08 01 06 03 E8						01100	000			0008	
		R2	0FA0 = 00 6A 00 08 02 06 03 E8						01100	000			0008
		R3	1388 = 00 6A 00 08 03 06 03 E8						00100	000			0008

Comments: Execution of the write CKD commands causes a count area, key area (if the key length specified is not zero), and a data area whose length is dependent upon the value specified in the DL bytes of the count area, to be written on the disk.

The main storage locations specified in the data address are coded with the cylinder number, head number, record number, key length, and data length of each record. Since the key length specified is six, a key area of six bytes long will be created. The data length specified is 03E8 (1000 bytes). Although the CCW byte count is only eight, and the channel byte count will go to zero after eight bytes have been written, the 3830 is committed to writing a key area six bytes long and a data area 1000 bytes long. Therefore the 3830 inserts 0's in the applicable positions on the track until the 3830 byte count reaches zero.

The difference in the channel byte count and the 3830 byte count will cause an incorrect length indication. Therefore the SLI flag (bit 34) is on in the CCW's.

In this example, six bytes of 0's will be recorded in the key area followed by the error correction code bytes, a gap, 1000 bytes of 0's and more error correction code bytes. At a later time data can be recorded in the key and data areas with the following CCW sequence.

```

Set Sector
Search ID Equal (R1)
TIC * -8
Write Key and Data
Search ID Equal (R2)
etc.
    
```

Example 2: Update Frank Smith's payroll record. Assumed:

1. The disk is organized by key areas.
2. Each key area contains a man number.
3. Frank Smith's man number is 656151.
4. This man number is located on track 0C head-04.
5. Key areas are 6 bytes long and data areas 64 (100₁₀) bytes long.

The channel program used is:

Seek
 Search Key Equal
 TIC *-8
 Write Data

SEEK

0	7	8	31	32	36	37	39	40	47	48	63		
Command Code		Data Address						Flags		Count			
Binary 0000 0111 Hex 07		C C H H 03E8 = 00 00 00 0C 00 04						01000		000		0006	
<p>Comments: As explained in example 1, the seek command transfers the track address to the storage control, moves the access mechanism, and selects the specified head.</p>													

SEARCH KEY EQUAL

0	7	8	31	32	36	37	39	40	47	48	63		
Command Code		Data Address						Flags		Count			
Binary 0010 1001 Hex 29		(man number) 07D0 = F6F5F6F1F5F1						01000		000		0006	
<p>Comments: After locating the proper cylinder and track, it is necessary to find Frank Smith's record. Since the disk is organized by keys, a search key equal command is executed. Execution of this command causes the 3830 to search the key field of the next record encountered on the track. If the key is not equal to Frank Smith's man number, (main storage locations 07D0 to 07D5) the 3830 signals channel end and device end to the channel and the TIC command (back to search key equal) is executed. Subsequent key areas are searched until Frank Smith's record is found. The 3830 then signals channel end, device end, and status modifier to the channel. The status modifier bit in the ending status byte causes the channel to skip the next command (TIC) and execute the write data command.</p>													

TRANSFER IN CHANNEL (TIC)

0	7	8	31	32	36	37	39	40	47	48	63
Command Code		Data Address		Flags		000					Count
Binary XXXX 1000 Hex X8		Address of search key equal		XXXXX		XXX					XXXX
Comments: X = positions ignored.											

WRITE DATA

0	7	8	31	32	36	37	39	40	47	48	63
Command Code		Data Address		Flags		000					Count
Binary 0000 0101 Hex 05		(data to update record) 0BB8 = XX XX XX to 0C1C		00000		000					0064
<p>Comments: The write data command transfers the data to update Frank Smith's payroll record from main storage locations 0BB8 to 0C1C to the disk.</p> <p>Note: If Frank Smith's payroll record had not been on track 0C head 04, the program would loop between the search key equal and TIC until every key on the track had been searched. The 3830 would then signal unit check to the channel. A subsequent sense I/O command would indicate no record found.</p> <p>The data just written could be verified by chaining the following CCW's to the write data command:</p> <ul style="list-style-type: none"> Read Sector (store sector address) Set Sector (locate sector) Search Key Equal (locate record) TIC *-8 Read Data (verify data) 											

Example 3: Find and read Joe Brown's insurance policy number. Assume:

1. The disk is organized by ID - no keys.
2. Joe Brown's employee serial number is 12341.
3. The data length of each record is 00AA (170 bytes).
4. His policy number is in the data area.
5. The data set begins on cylinder 0A track 00.

Using the record capacity chart in Appendix B, it is known that 43 - 170 byte records can be written on a 3330 track. Since the disk is organized by ID's (Joe Brown's = 12341) the track and record location can be determined by dividing the ID by the number of records per track. In this case:

$$\frac{12341}{43} = 287 \text{ Note: Add 1 to the remainder to establish the address of the specific record.}$$

Thus Joe Brown's ID is 287 tracks from the beginning of the data set. There is no remainder so the first record on the track will be Joe Brown's.

The CC HH R for the seek command is then determined by converting the 287 tracks to cylinders and adding the results to the beginning of the data set.

	<u>Cylinder</u>	<u>Track</u>	<u>Record</u>	<u>C</u>	<u>C</u>	<u>H</u>	<u>H</u>	<u>R</u>
Starting Address:	10	00	0	00	0A	00	00	00
Displacement:*	15	02	1	00	0F	00	02	01
Result:	25	02	1	00	19	00	02	01

* = Determined by dividing 287 by 19.

The channel program used is:

Seek
 Search ID Equal
 TIC* -8
 Read Data

SEEK

0	7	8	31	32	36	37	39	40	47	48	63
Command Code		Data Address				Flags		000		Count	
Binary 0000 0111 Hex 07		C C H H 03E8 = 00 00 00 19 00 02				01000		000		0006	
Comments: The seek command is executed to position the access mechanism at cylinder 19 (decimal 25) and select head 02.											

SEARCH ID EQUAL

0	7	8	31	32	36	37	39	40	47	48	63
Command Code	Data Address						Flags	000			Count
Binary 0011 0001 Hex 31	C C H H R 05DC = 00 19 00 02 01						01000	000			0005

Comments: The search ID equal command causes the first ID encountered on the track to be compared with Joe Brown's ID. All unequal comparisons of ID's cause the 3830 to signal channel end - device end to the channel, and the TIC command (back to the search ID equal) is executed. When an equal comparison is encountered (ID of record 1) the 3830 signals channel end, device end, and status modifier to the channel. Status modifier causes the next command (TIC) to be skipped and the read data command is executed.

If the search ID equal is not satisfied and index is passed twice, unit check is sent in the status byte. A subsequent sense I/O command would indicate no record found. The course of action would then be determined by the error recovery procedures.

TRANSFER IN CHANNEL (TIC)

0	7	8	31	32	36	37	39	40	47	48	63
Command Code	Data Address						Flags	000			Count
Binary XXXX 1000 Hex X8	Address of search ID equal command.						XXXXXX	XXX			XXXX

Comments: X = positions ignored.

READ DATA

0	7	8	31	32	36	37	39	40	47	48	63
Command Code	Data Address						Flags	000			Count
Binary 0000 0110 Hex 06	(insurance policy number) 0BB8 = XX XX XX to 0C62						00000	000			00AA

Comments: Execution of the read data command causes the data area, containing Joe Brown's insurance policy number, to be read into main storage locations 0BB8 to 0C62.

MULTIPLE TRACK (MT) OPERATION

On all search and most read commands, the storage control can automatically select the next sequentially numbered head on the disk drive under control of bit 0 of the command code. If bit 0 is a 1 and data transfer of the command has not been initiated, the next sequentially numbered head is selected at index. Thus, the need for seek head commands in a chain of read or search commands is eliminated.

Note: Channel end, device end, and unit check are signaled to the channel if the head switching operation crosses a file-protected boundary or exceeds the limits of the cylinder.

Discretion must be used when using the MT bit. For example, assume that during a multi-track search operation the desired record is on the first

track searched and the search commences after that record is passed. The head number, therefore, is advanced to the next track without comparing on the desired record. Also, should a set sector command with a sector value of zero precede a multi-track command, head switching could occur before the desired record is reached. To avoid these conditions, a single track read home address or read R0 should be placed before the search, thus ensuring that the search commences at R0 or R1 of the track. (See "Figure 9.")

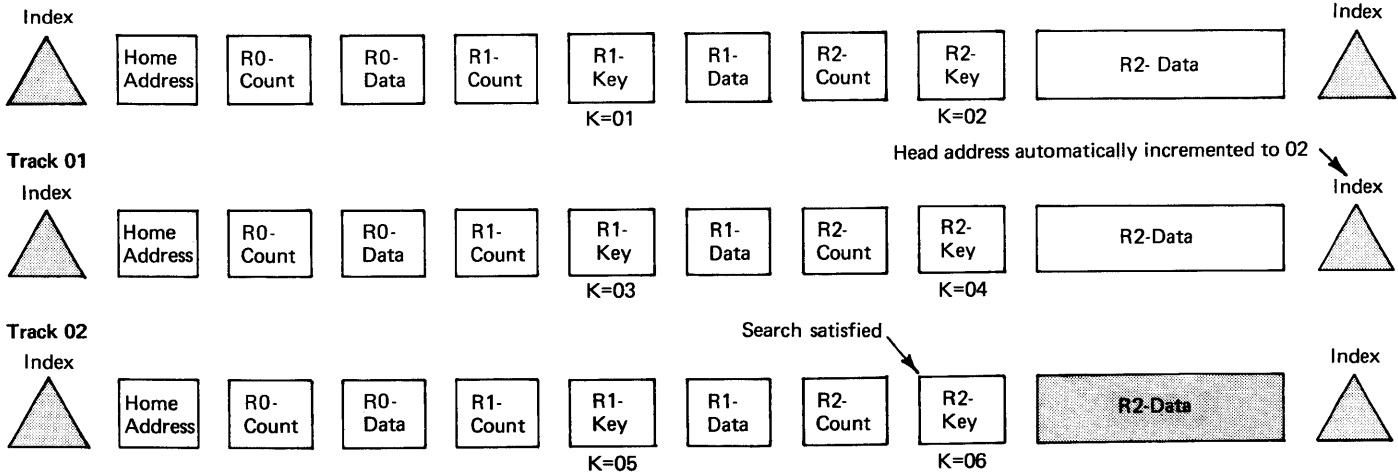
Multi-track operations are not used on read IPL, read sector, or read diagnostic status 1 commands.

RECORD OVERFLOW

The record overflow function provides a means of processing logical records which exceed the capacity

MULTI-TRACK OPERATION

Cylinder 02
Track 00



Channel program using multiple track search.

Object: Update John Doe's payroll record.
Assume: The disk is organized by keys, and the physical address of the record is unknown.
 Set File Mask (allow write and seek commands).
 Seek (cylinder 02, head 00).
 Read Home Address (make sure all records are read).
 Search Key Equal (MT bit on, argument = 06).
 TIC * -8.
 Write Data (updates shaded area).

25563A

Figure 9. Multiple Track Operation

of a track. When using overflow records, the factor limiting the size of the record is the cylinder boundary.

Formatting Overflow Records

That portion of an overflow record written on (or read from) one track is called a record segment. Each segment contains a count field, key field (optional), and a data field. The key and data lengths specified in the KL and DL bytes of the count field pertain only to that segment, not the entire overflow record. Since only the key field of the first segment has significance, overflow records are usually formatted without key fields (KL = 0).

Write special count, key, and data commands are used to format all segments of an overflow

record except the last segment. As shown in Figure 10, the last segment is formatted with a normal write count, key, and data command.

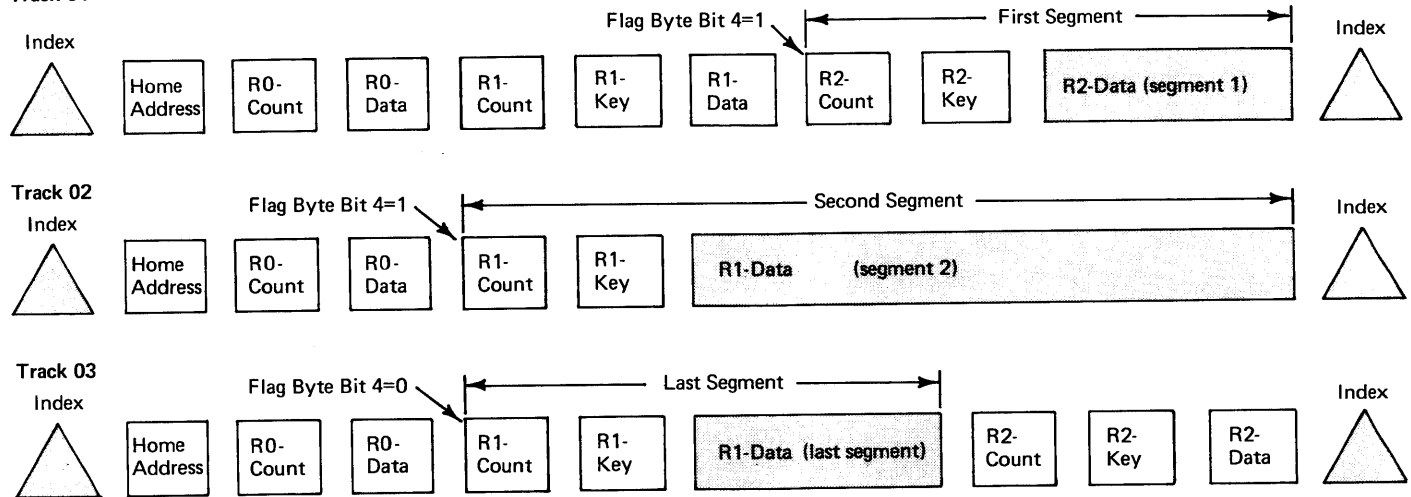
Write special CKD commands cause a 1 to be written in flag byte bit position four of the record segment being written. This bit, which identifies the record as an overflow segment, indicates to subsequent record processing commands that the logical record continues on the following track.

No internally generated head switching is associated with formatting overflow records; all head seeking must be done by the formatting program. (Figure 10). Head switching will not occur:

- in violation of the file mask,
- past the end of the cylinder,
- to a defective track,
- to an alternate track.

OVERFLOW RECORD

Cylinder 02
Track 01



Typical channel programs for formatting, updating, and reading overflow records.

Formatting:

- Set sector
- Search ID R1 (track 1)
- TIC* -8
- Write special CKD (segment 1)
- Seek head (next track)
- Search ID R0 (track 2)
- TIC* -8
- Write special CKD (segment 2)
- Seek head (next track)
- Search ID R0 (Track 3)
- TIC* -8
- Write CKD (last segment)

Updating:

- Set sector
- Search ID R2 (segment 1)
- TIC* -8
- Write data (updates shaded areas)

Reading:

- Set sector
- Search ID R2 (segment 1)
- TIC* -8
- Read data (reads shaded areas)

25564A

Figure 10. Record Overflow

All segments of an overflow record—except the first—must be written immediately following R0; all segments—except the last—must be the last physical record on their respective tracks.

Processing Overflow Records

The following commands may be used to read or update previously formatted overflow records.

- Read count, key, and data.
- Read key and data.
- Read data.
- Write key and data.
- Write data.

When any of the above are used to process an overflow record, the operation will not terminate at the end of a record segment when the segment is flagged with bit four (on) in the flag byte. Instead, the head address is incremented by 1 at index and the operation continues in the data field of record one on the next track. If this record segment is also flagged with bit four (on) in the flag byte, the operation continues on the next track. When a segment is found that is not flagged, the operation terminates at the end of the data field. The net effect of this procedure is that the data fields of all the record segments appear as a single logical data field.

Should a data overrun occur during the first segment, the storage control attempts recovery through use of command retry. If a data overrun occurs during an operation involving the second (or subsequent) segments, unit check is signaled immediately during a read operation, or at the end of the associated segment during write operations.

If a data check or bus out parity error occurs, unit check is signaled at the end of the associated area.

Note: If a write operation was in progress, unit check is signaled at the end of the record segment.

If the CCW count is less than the number of bytes in the logical record, the operation continues to the end of the logical record before presenting ending status.

Spacing over overflow records does not occur automatically. The channel program must be written so that the entire logical record is spaced over, not just the first segment. For example, in the sequence:

```
Set sector
Search ID (first segment)
TIC* -8
Read CKD (multi-track)
```

the read CKD does not read the next logical record on the cylinder. It commences reading the overflow record at the count field of the second segment.

The sequence:

```
Set sector
Search ID (first segment)
TIC* -8
Read Key and data (skip and SLI flags on)
Read CKD (multi-track)
```

reads the count, key, and data of the next logical record.

Multiple track operations should not be confused with overflow record operations. Head switching - when processing overflow records - occurs regardless of whether the MT bit is on or off.

END-OF-FILE

An end-of-file record, used to define the end of a logical group of records, is written by executing a write count, key, and data command with the DL bytes in the count area set to zero. Execution of a write CKD with a data length of zero causes the storage control to write a data area consisting of one byte of zeros followed by the error correction code bytes. (Figure 11.)

The KL portion of the count area can be either zero or non-zero. If KL equal zero, the end-of-file record contains only the contents of the count area and data area. If the key length is not zero, the key area is written as specified by the KL byte.

Detection of a zero data length causes unit exception status to be generated. No data from the data area is transferred to the channel. A read R0, read CKD, or read KD will transfer the key area (if any) to the channel.

The unit exception is generated during execution of read IPL, read R0, read CKD, read KD, read data, write KD, and write data commands.

END-OF-FILE

Cylinder 02
Track 00

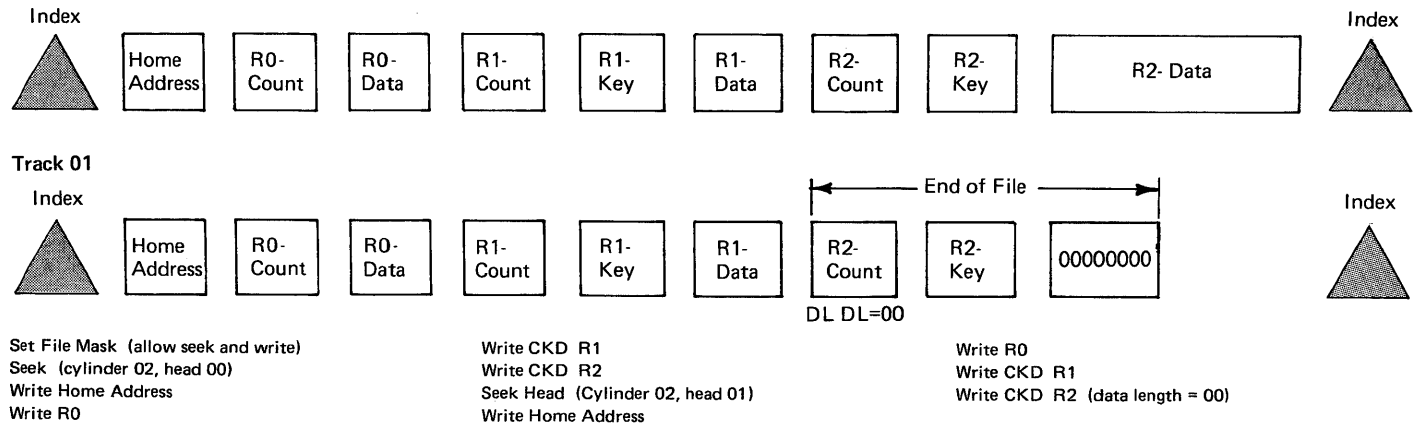


Figure 11. End of File

ROTATIONAL POSITION SENSING

Rotational position sensing reduces the time the channel is busy searching for a record. This procedure permits a search command to be initiated just before the desired record is positioned under the read/write heads.

To accomplish this, a "sector" concept is employed. The tracks in each cylinder of a disk storage drive are divided into 128 equally spaced sectors; each record on the track has a sector location as well as a record address. Although the sector location is not physically indicated on the tracks, the sector number is stored at the beginning of all read, write, and search commands. When chained to a read, write, or search CCW, the read sector command provides the sector number required to access the record processed by the previous command. A subsequent set sector command can be used to fetch the sector number from main storage to reposition the track at that record. This type of operation is particularly useful in write verification (Figure 12) and sequential disk processing operations.

The sector in which a record is recorded is a function of the length of all records that precede it and its sequential position on the track. Therefore, the sector location can be calculated with the following formula. IBM access methods that support 3330 calculate the proper sector number.

If:

n = Record number of desired record
 $S(n)$ = Sector number
 KL = Length of key field (in bytes)
 DL = Length of data field (in bytes)

Then:

$$n = 0: S(n) = 0$$

$$n = 1: S(n) = \frac{128}{13440} [237]$$

$$n > 1: S(n) = \frac{128}{13440} \left[237 + \sum_{i=1}^{n-1} (135 + KL_i + DL_i + C) \right]$$

where

$C = 0$ if KL_i is zero.

$C = 56$ if KL_i is not zero.


A standard R0 key area ($KL = 0$, $DL = 8$) is assumed.

The following example shows some of the advantages of using rotational position sensing to locate and retrieve records.

Without RPS

Channel program 1.

Command	Selector Channel and Storage Control Status
Seek	Available as soon as the storage control accepts the seek address.



Channel program for write verification of record n.

Seek

- A** Search ID Rn
TIC* -8
- B** Write data Rn
- C** Read Sector (82)
- D** Set Sector (82)
After channel reselection:
- E** Search ID Rn
TIC* -8
- F** Read data Rn

25561A

Figure 12. Rotational Position Sensing

<u>Command</u>	<u>Selector Channel and Storage Control Status</u>
Channel program 2.	
Search ID Equal	Busy (average 12.5 ms on the 2314).
TIC *-8	
Read Data	Busy

With RPS

When the sector address is known or can be calculated, the following channel program can be used:

<u>Command</u>	<u>Block Multiplexer Channel and Storage Control Status</u>
Seek	Available during access movement.
Set Sector	Available until sector is located.
Search ID Equal	Busy (average 250 μ s on the 3330).
TIC *-8	Normally the first ID read is that of the desired record and the TIC is not executed.
Read Data	Busy

Note that with RPS only one channel program is required to locate the record and transfer the data. This eliminates a seek I/O interrupt and the I/O processing required to schedule a data transfer channel program.

Also, the channel and disk storage are available during access motion and rotational positioning, allowing seek and set sector operations to be overlapped with other I/O operations on the storage control and channel.

MULTIPLE REQUESTING

Use of block multiplexer channels and rotational position sensing enables the IBM 3830/3330 facility to disconnect from the channel during mechanical delays resulting from execution of arm positioning seek or set sector commands. Reconnection is attempted when the access mechanism is positioned at the desired track or when the specified rotational position has been reached.

During the time the channel and storage control are disconnected, the CPU is free to initiate I/O

operations on other drives attached to the 3830 although the disconnected channel program is not completed. Thus, separate channel programs may be operating simultaneously on each 3330 attached to the storage control.

The storage control stores the file mask, seek, or set sector arguments required to successfully complete the disconnected chains.

COMMAND RETRY

Command retry is a channel/storage control procedure that causes an improperly executed command in a channel program to be automatically retried. Retry is requested when the storage control sends unit check with channel or device end and status modifier. The re-execution does not cause an I/O interrupt, therefore no CSW is stored and programmed error recovery procedures are not required.

Command retry is used:

1. To recover from correctable data errors (error burst 11 bits or less) that occur during a search or read operation on a home address, count, or key area.

During a search or read operation, the home address, count, or key read from the disk is placed in a buffer in the storage control. When a correctable data error occurs, the storage control corrects the data in the buffer and requests the channel to reissue the command which originally caused the error. During reorientation to the record, the storage control disconnects and frees the channel. When the failing search or read command is re-executed, the corrected data in the buffer is used, instead of the actual data from the track.

2. When an uncorrectable data error (an error burst longer than 11 bits) is detected on any position of record during a read or search operation.

The failing command is reissued by the storage control. If retry is successful, the channel program continues normally. If retry is unsuccessful, the storage control retries the operation again.

If after any retry the error becomes correctable, the procedure outlined in 1 applies. If the error does not become correctable, the operation is terminated and the program is interrupted.

3. When a seek malfunction is detected.

The storage control retries the command in an attempt to position the access mechanism correctly.

4. When an alternate or defective track condition is detected before data transfer begins.

The storage control determines the location of the alternate or defective track (from R0 on the track), initiates a seek to this track, orients on index, and reissues the original command.

5. When a command overrun (or late command chaining) condition occurs because of interference from another channel or the CPU.

The storage control initiates a retry of the command that was late.

6. When a data overrun occurs except:
 - a. A data overrun occurring during a record overflow operation in the second or subsequent segments.
 - b. A data overrun occurring during a format write.

Execution of command retry may cause the following conditions to be detected by the initiating program:

1. A CCW containing a PCI may, if retried because of command retry, cause multiple PCI interruptions to occur.
2. A channel program consisting of a single, unchained CCW specifying an immediate command may cause a condition code of zero rather than one to be set. This setting of the condition code occurs if the control unit signals command retry at the time initial status is presented to the command. The channel program then causes a later interruption upon completion of the operation.
3. If premature termination of the execution of a channel program occurs during the retry of a command, the residual count and command address field in the CSW may not necessarily indicate the extent of main storage used.
4. If a CCW used in an operation is changed before that operation has been successfully completed, the results are unpredictable.

STATISTICAL USAGE/ERROR RECORDING

The 3830 maintains a statistical data record of usage and error information for each logical device in the facility. The usage information provides an accumulated count of the total number of access motions, and the total number of data bytes read. The error information provides an accumulated count of the total number of seek errors, correctable data errors, and uncorrectable data errors which were recovered by the storage control retry procedure. Also included in the error information is the total number of command and data overrun conditions which were retried by the storage control.

Any time the number of errors exceeds a predetermined level, or the number of seeks or data bytes processed exceeds a predetermined level, the storage control generates a unit check signal. A unit check is presented to the channel in response to the next start I/O instruction addressed to the storage control. The following sense information is asso-

ciated with the unit check: sense byte 2, bit 3 indicates environmental data present, and sense byte 7 indicates usage/error statistics. Usage/error information is reset after it is transferred to the channel by the sense I/O command.

The read and reset buffered log command is used to offload the usage/error information after a pack change or at the end of day.

A system reset will reset usage/error statistics for only those devices which have a pack change device end outstanding.

STORAGE CONTROL DIAGNOSTICS

To provide maximum facility availability, the 3830 can execute diagnostic tests on a drive, concurrent with normal system operations on the remaining drives. This mode of operation allows the customer engineer to diagnose and repair most drive failures while the facility continues to operate other attached drives. The 3830 provides a transient block of 512 bytes (128 words) of control storage to allow temporary residence for a specific diagnostic test.

The transient area is loaded by the system under control on the On-Line Test Executive Program (OLTEP). A special command -- diagnostic write -- loads a selected test into control storage and instructs the storage control to execute the test. This loading and execution may also be initiated from the CE panel.

After the test, error message information or test results are transferred from the 3830 to main storage by a read diagnostic status 1 command. If the CE panel is used, the test results are displayed on the CE panel indicators.

USAGE METER

3830 Storage Control Meter

If the Enable/Disable switch is in the Enable position when a power on sequence occurs, meter time will be recorded as long as the CPU meter is recording or until the usage meter and 3830 are disabled from the channel.

The usage meter and 3830 are disabled when the following conditions exist simultaneously:

- The Enable/Disable switch is in Disable.
- The CPU is in a stop or wait state.
- Command chaining is not in effect.
- The 3830 channel selection switch is not selected to a channel. (See "Special Features.")
- The 3830 is not performing an operation.
- There is not any status pending. See "Pending Status."

The usage meter can then be enabled, provided:

- The CPU is in the stop or wait state.
- The Enable/Disable switch is in Enable.

TWO CHANNEL SWITCH AND TWO CHANNEL SWITCH ADDITIONAL

The two channel switch special feature provides the ability for the IBM 3830 Storage Control to be shared by two channels. The combination of two special features--two channel switch and two channel switch additional--permits the 3830 to be shared by four channels. The channels may be attached to either the same or different central processing units. With appropriate programming or operator action, individual drives attached to the storage control may be reserved for the exclusive use of any of the channels. Channel switching and device reservation are controlled by the channel program. Two special commands are associated with the features: device reserve and device release. (See "Channel Commands.")

Channel Selection Switch

Channel selection is determined by a three or five position program-controlled switch in the 3830. When the switch is in neutral, the 3830 can be selected by any channel. The channel A position indicates that the storage control has been selected by channel A; the channel B position indicates that the storage control has been selected by channel B; and so on.

Once the 3830 has been selected by a channel, it is switched to that channel until the channel disconnects. The channel selection switch will then return to neutral unless:

- Chaining is indicated and device end is included in the status.
- Chaining is indicated without device end in the status, and the channel does not disconnect.
- Chaining is indicated without device end in the status, the channel disconnects, and the storage control becomes busy to allow:
 1. Execution of a storage control error recovery procedure.
 2. Execution of a diagnostic load or diagnostic write command.
 3. Completion of a format write operation.

- Chaining is indicated and a format write operation is in progress.
- The last status byte was part of a channel-initiated signal sequence and was stacked by the channel.
- A contingent connection is established.
- Ending status associated with an interface disconnect has not been accepted by the channel.

Device Status

Multi-tagged status: presented to all interfaces not partitioned from the storage control. Multi-tagged status conditions cause status to be generated for each of the attached channels. The status must be accepted by a channel for that channel to use the device.

Untagged status: not associated with any particular interface and is presented to only one channel -- the first channel to accept the status from the device. Other channels may be presented a status byte of all zeros. This type of status transfer is accomplished by considering the status as multi-tag until one channel accepts the status; at that time the status condition is cleared for other channels.

Tagged status: associated with a particular interface and made available solely to that interface. The status remains pending until accepted over the interface identified by the tag.

When a device is busy for any reason (including reservation to channel A), any command from channel B, C, or D addressed to that device will be rejected with a busy status. This, in turn, causes the storage control to attempt to present to channel B, C, or D a status byte containing device end after the busy condition has been terminated. The address byte associated with this status byte will be the same as that associated with the busy status byte.

Device end status resulting from any channel command will be presented to the channel that issued the command.

Device end status resulting from a not-ready to ready transition will be presented under control of the multi-tagged/untagged switch.

Addressing

The base address (five high-order bits) of the storage control on one channel is independent of the base address on the other channel. However, the three low-order address bits for any attached device must be the same on all channels.

Resets

A system reset may be initiated by any channel at any time. A system reset resets all reservations

and status conditions stored in the storage control for the resetting channel, terminates all block multiplex command chains in progress on the resetting channel, and resets all device interrupts not associated with the other channels. Reservations, status, and device interrupts for the other channels, as well as block multiplex chains in progress on the other channels, are not affected. If a channel initiates a system reset while the selection switch is connected to the other channels, a machine reset is performed when the selection switch goes to neutral. A selective reset has no effect on device reservations or status.

Error Recovery Procedures

The error condition table (Figure 13) identifies all unique configurations of sense bits in sense bytes 0, 1, and 2, posted by the storage control. In addition, it maps each of these configurations into a specific recovery action to be invoked by the system. The recovery action table (Figure 13) specifies the action to be taken for each error condition.

Error Correction Function

The recovery action table uses an error correction function as a step in recovering from data errors. The error correction function is used when the storage control posts the data check and correctable sense bits in the sense information. These bits are posted if a correctable data error is detected in any data area.

Correctable data errors in home address, count, and key areas are corrected internally by the storage control by using command retry. Data check and correctable sense bits are not posted for these errors, and do not cause a system interrupt.

When the correctable and data check sense bits are included in the sense information, sense bytes 18 through 22 provide the error pattern and displacement.

Error correction is accomplished by aligning the error pattern provided in sense bytes 20 through 22 with the erroneous data in main storage and exclusively ORing the error pattern and main storage bytes.

The location of the erroneous data in main storage is determined by using displacement information provided in the sense bytes, and the counts provided in the interrupted CCW chain. The storage control specifies the location of the error bytes, relative to the first byte transferred in the operation which incurred the error. The displacement between the first byte transferred and the first byte in error is calculated by subtracting the error displacement provided in sense bytes 18 and 19 from the restart displacement provided in sense bytes 15 through 17. The result constitutes the forward error displacement and is used, in conjunction with the count specified in the interrupt CCW, to locate the erroneous main storage data.

If data chaining was indicated in the operation which posted the correctable error, the forward displacement may reference data from the second (or subsequent) CCW in the data chain.

Prior to applying the error correction function, it must be determined whether any error bytes were not transferred, due to the skip bit being on, due to a short count in the CCW, or if the error bytes are not contiguous in main storage due to data chaining between CCW's.

- If any of the error bytes are contained in data specified by a CCW which has the skip bit on, the error correction function must be bypassed for those bytes which were not transferred to main storage.
- If any of the error bytes are contained in data not transferred to main storage due to a short count in the CCW, the error correction function must be bypassed for those bytes which were not transferred to main storage.
- If no short count in the CCW is detected and bit 7 of sense byte 23 indicates channel truncation, the error correction function cannot be applied correctly.
- If the error pattern spans non-contiguous main storage boundaries due to data chaining, the error correction function must be selectively applied to the non-contiguous storage locations.
- If the error displacement in sense bytes 18 and 19 is less than 3, the error is partially or totally contained in the correction code bytes. In this case, the error pattern in sense bytes 20-22 is constructed as follows:
 1. If the error displacement is zero, the error pattern must be set to zero by the error recovery programs (ERPs).
 2. If the error displacement is one, the two low-order error pattern bytes (bytes 21 and 22) must be set to zero by the ERPs. The high-order bytes contain the correction syndrome.
 3. If the error displacement is two, the low-order pattern byte must be set to zero by the ERPs. The high-order bytes contain the correction syndrome.

Note: Case (1) also occurs if the error is totally contained in the gap byte which immediately precedes the data area.

Error Correction Table					
Byte	Bit	Name	General Description	Action	Logged
0	0	Command Reject	Programming error	2	No
0	1	Intervention Required	Drive off-line	3	No
0	2	Bus Out Parity	Bus Out Parity Error	3	Yes
0	3	Equipment Check	Equipment Malfunction	4	Yes
0	3	Equipment Check	Equipment malfunction	1	Yes
1	0	Permanent Error	Storage Control retry exhausted or undesirable		
0	4	Data Check	Uncorrectable data check. Storage Control retry exhausted.	1	Yes
1	0	Permanent Error			
0	4	Data Check	Correctable data check in data area or data area of last overflow segment.	5	No
2	1	Correctable			
0	4	Data Check	Correctable Data Check in data area of overflow segment, not last segment.	6	No
2	1	Correctable			
1	7	Operation Incomplete			
0	4	Data Check	Data check in second or subsequent overflow segment other than a data field correctable error.	6A	No
1	7	Operation Incomplete			
0	5	Overrun	Control unit retry exhausted on a service overrun	1	Yes
1	0	Permanent Error			
0	5	Overrun	Service overrun in second or subsequent overflow segment or during a format write.	4	Yes
1	1	Invalid Track Format	Track capacity exceeded	2	No
1	2	End of Cylinder	Cylinder boundary detected during a basic multitrack operation.	8	No
1	2	End of Cylinder	Cylinder boundary detected during a basic overflow operation.	9	No
1	7	Operation Incomplete			
1	4	No Record Found	Record not found during basic command sequence.	2	No
1	5	File Protected	The seek command or read/search multitrack operation violated file mask.	10	No
1	5	File Protected	A read or write overflow operation violated file mask.	11	No
1	7	Operation Incomplete			
1	7	Operation Incomplete	One of the following was detected after initiation of data transfer during an overflow operation: a. A defective or alternate track condition. b. A seek error in the second or subsequent segment.	7	No
2	3	Environmental Data Present	Statistical usage/error log information present.	4	Yes
0	0	Command Reject	A write command received with write inhibit switch in read only position.	1	No
1	6	Write Inhibit			

Recovery Action Table	
Action	Explanation
1	Print console error message for operator and/or customer engineer notification. (see Figure 14)
2	Exit with programming error or unusual condition indication.
3	a. Repeat the operation one time. b. If error condition persists, do action 1.
4	a. Repeat the operation. b. If the error condition persists after ten retries, do action 1.

11828.0B

Figure 13. Error Recovery Procedures (Part 1 of 3)

Recovery Action Table (continued)	
Action	Explanation
5	<p>a. Perform error correction function.</p> <p>b. Examine bit 7 of the file mask (PCI). If this bit is off, go to step (c). If this bit is on, return to user with indication that data has been corrected. (User is operation in PCI fetch mode and must, therefore, supply his own restart recovery action.)</p> <p>c. If the user's chain has not been completed, examine the next non-TIC command in the user's chain. If bit 3 of this command is on (count area), go to step (d). If bit 3 is off, do action 5A.</p> <p>d. Continue the user's chain by executing the following CCW chain:</p> <p>Seek * Set File Mask (same as original) Set Sector (Sector data provided in sense byte 13) Search Equal ID (CCHHR provided in sense bytes 8-12) TIC * -8 TIC (channel status word)</p>
5A	<p>Continue the user's chain by executing the following command chain:</p> <p>Seek * Set File Mask (same as original) Set Sector (Sector data provided in sense byte 13) Search Equal ID (CCHHR provided in sense bytes 8-12) TIC * -8 Read Count (skip bit on) TIC (channel status word)</p>
6	<p>a. Perform error correction function.</p> <p>b. Examine bit 7 of the file mask (PCI). If this bit is off, go to step (c). If this bit is on, return to user with indication that data has been corrected. (User is operating in PCI fetch mode and must supply his own restart recovery action.)</p> <p>c. Construct restart CCW 2.</p> <p>d. Complete the interrupted operation and continue the user's chain (if appropriate) by executing the following command chain.</p> <p>Seek * (increment seek argument by one) Set File Mask (same as original) Set Sector (argument 0) Search ID Equal (record 1) TIC * -8 Restart CCW 2 TIC (channel status word)</p> <p>Note: If the modified seek argument is not within the user's extent then IOS must supply the correct seek argument before issuing the seek. If that is impossible, then IOS must do action 2.</p>
6A	<p>a. Examine bit 7 of the file mask (PCI). If this bit is off, go to step (b). If this bit is on, return to user with indication that data been corrected. (User is operating in PCI fetch mode and must supply his own restart recovery action.)</p> <p>b. Construct restart CCW 2.</p> <p>c. Complete the interrupted operation and continue the user's chain (if appropriate) by executing the following command chain.</p> <p>Seek * Set File Mask (same as original) Set Sector (argument 0) Search ID Equal (record 1) TIC * -8 Restart CCW 2 TIC (channel status word)</p>
7	<p>a. Construct restart CCW 1.</p> <p>b. Continue the user's chain by executing the following command chain:</p> <p>Seek * Set file Mask (same as original) Set Sector (argument 0) Search ID Equal (record 1) TIC * -8 Restart CCW 1 TIC (channel status word)</p>

* Cylinder bytes and high order head byte obtained from user.
Low order head byte obtained from bits 3 thru 7 of sense byte 6.

11828.1B

Figure 13. Error Recovery Procedures (Part 2 of 3)

Recovery Action Table (continued)	
Action	Explanation
8	<p>a. Increment the cylinder address of the user's seek argument by one. Reset the head address</p> <p>b. Continue the operation by executing the following command chain:</p> <p>Seek (argument from step a) Set File Mask (same as original) TIC (channel status word -8)</p> <p>Note: If the modified seek argument is not within the user's extent then IOS must supply the correct seek argument before issuing the seek. If that is impossible, then IOS must do action 2.</p>
9	<p>a. Increment the cylinder address of the user's seek argument by one. Reset the head address.</p> <p>b. Construct restart CCW1.</p> <p>c. Complete the interrupted operation and continue the user's chain (if appropriate) by executing the following command chain:</p> <p>Seek (argument from step a) Set File Mask (same as original) Set Sector (argument 0) Search ID Equal (record 1) TIC * -8 Restart CCW 1 TIC (channel status word)</p> <p>Note: If the modified seek argument is not within the user's extent, then IOS must supply the correct argument before issuing the seek. If that is impossible, this IOS must do action 2.</p>
10	<p>a. Determine if the interrupted command is a seek. If yes, go to step b. If no, do action 10A.</p> <p>b. Continue the operation by executing the following command chain:</p> <p>Seek (user's argument) Set File Mask (same as original) TIC (channel status word)</p> <p>Note: If seek argument is not within the user's extent, then IOS must supply the correct seek argument before issuing the seek. If that is impossible, then IOS must do action 2.</p>
10A	<p>a. This is a multi-track operation. Increment the user's seek argument by one.</p> <p>b. Continue the operation by executing the following command chain:</p> <p>Seek (argument from step a) Set File Mask (same as original) TIC (channel status word -8)</p> <p>Note: If the modified seek argument is not within the user's extent, then IOS must supply the correct seek argument before issuing the seek. If that is impossible, then IOS must do action 2.</p>
11	<p>a. Increment the user's seek argument by one.</p> <p>b. Construct restart CCW 1.</p> <p>c. Complete the interrupted operation and continue the user's chain (if appropriate) by executing the following command chain:</p> <p>Seek (argument from step a) Set File Mask (same as original) Set Sector (argument 0) Search ID Equal (record 1) TIC * -8 Restart CCW 1 TIC (channel status word)</p> <p>Note: If the modified seek argument is not within the user's extent, then IOS must supply the correct seek argument before issuing the seek. If that is impossible, then IOS must do action 2.</p>

11828.2A

Figure 13. Error Recovery Procedures (Part 3 of 3)

Example


Assume the following:

Key length = 2

Data length = 10

The CSW-8 points to CCW 1 in the following chain:

CCW	Commands	Address	Count	Flags
1	Read key and data	A	2	data chaining
2	TIC	CCW 3	--	----
3	----	B	4	data chaining, skip
4	----	C	1	suppress incorrect length



Suppose the error affected bytes 6, 7, and 8 as follows:

Byte 6 -----XX

Byte 7 XXX-----

Byte 8 X-----

where (-) corresponds to correct bit
(X) corresponds to incorrect bit

25587

The illustrated condition generates a restart displacement of 12 and an error displacement of 7.

The error pattern would be generated as follows:

```

Pattern byte 1 (sense byte 20)  0 0 0 0 0 0 1 1
Pattern byte 2 (sense byte 21)  1 1 1 0 0 0 0 0
Pattern byte 3 (sense byte 22)  1 0 0 0 0 0 0 0
    
```

Application of the error correction algorithm, as outlined in the preceding sections, would result in the following system recovery action.

1. Pattern byte 1 would not be applied to data byte number six, since this byte was not transferred to main storage due to the skip flag in the CCW 3.
2. Pattern byte 2 would be exclusively ORed to main storage location B, where data byte 7 resides.
3. Pattern byte 3 would not be applied to data byte 8, since this byte was not transferred to main storage due to a short count in CCW 4.

Construction of Restart CCW's

If operation incomplete (byte 1 - bit 7) is set in the sense information, it indicates that an error or unusual condition occurred during a logical operation after data transfer had been initiated. By constructing restart channel command words, the error recovery procedures are able to correct the unusual condition and continue the operation in progress from the point of interruption to the normal ending point.

Restart CCW 1

Restart CCW 1 is constructed as follows:

1. The command code byte is provided in sense byte 3.
2. The data address is that of the interrupted CCW, plus the count of that CCW, minus the residual count in the channel status word.
3. The flags (except PCI) are those of the interrupted CCW.
4. The count is the residual count in the CSW. If the residual count is zero, a count of one must be used. If a write command was in progress, the data address should specify a byte containing 00. If a read command was in progress, the skip bit should be on.

Restart CCW 2

Restart CCW 2 is constructed as follows:

1. The command code is provided in sense byte 3.
2. The count is constructed as follows:
 - a. Fetch the count of the CCW designated by CSW-8, and set a pointer to this CCW.
 - b. Subtract the restart displacement from the count obtained in (a). If this result is positive, go to step (f); otherwise go to step (c).
 - c. Check the chain data flag of the CCW designated by the pointer. If the flag is not set go to step (e); otherwise go to step (d).
 - d. Advance the pointer to the next non-TIC CCW in the data chain and add the count of this CCW to the counts of all preceding non-TIC CCW's in the data chain. Return to step (b).

- e. Truncation occurred. Set restart CCW 2 count equal to one. Go to Step 3 and include the skip bit in the restart CCW flags.
 - f. Set restart CCW 2 count equal to the result of the subtraction in step (b). Go to Step 3.
3. The flags (except PCI) are those of the CCW designated by the pointer in Step 2. The skip bit is also set if Step 2e was executed.
 4. The data address is that of the CCW designated by the pointer in Step 2, plus the count

of that CCW, minus the restart CCW count generated in Step 2.

If another "operation incomplete" occurs while executing the restart CCW, a new restart CCW may be generated from the old restart CCW.

Note: Be sure to avoid destroying the old restart CCW before generating the new one.

Operating Instructions

Loading a Disk Pack

1. Place the start/stop switch on the 3330 operator panel in the stop position.
2. Place the open/close switch on the 3330 operator panel in the open position.
3. Remove the bottom cover of the disk pack by pressing the two handles on the bottom cover together.
4. Place the disk pack (in its top cover) on the drive spindle.
5. Turn the top cover in a clockwise direction until it comes to a full stop.
6. Lift the top cover from the disk pack.
7. Place the open/close switch in the close position.
8. Place the start/stop switch in the start position to return the drive to normal operation.
9. Reassemble the top and bottom covers.

With the pack identification label facing forward, place the reassembled cover in the recessed "well" on top of the 3330. The cover for the pack in the upper drive should be placed in the well on the left, and the cover for the pack in the lower drive in the well on the right. When stored in this manner, the pack identification is over the logical address plug associated with the drive in which the pack is mounted.

Do not store disk packs on top of the disk drives.

Unloading a Disk Pack

1. Place the start/stop switch on the 3330 operator panel in the stop position.
2. Place the open/close switch on the 3330 operator panel in the open position.
3. Remove the bottom cover of the disk pack by pressing the two handles on the bottom cover together.
4. Place the disk pack (in its top cover) on the drive spindle.
5. Turn the top cover in a clockwise direction until it comes to a full stop.
6. Lift the top cover from the disk pack.
7. Place the open/close switch in the close position.
8. Place the start/stop switch in the start position to return the drive to normal operation.
9. Reassemble the top and bottom covers.

2. Place the open/close switch on the 3330 operator panel in the open position.
3. Place the top cover on the disk pack and turn the cover in a counter-clockwise direction for two full turns.
4. Lift the top cover, now containing the disk pack, from the spindle.
5. Immediately attach the bottom cover.
6. Unless another pack is being loaded, place the open/close switch in the close position.
7. Store the removed disk pack in a clean cabinet or on a clean shelf.


Changing a Drive Address

To change the address of a drive:

1. Make sure that the program controlling the 3830 is in a wait state, or that the existing conditions allow removal of the logical address plug.
2. Remove the logical address plug from the affected 3330 operator panel and perform any necessary pack changes.
3. Place the desired address plug in the socket on the operator panel.

The drive is now ready to resume normal (or CE) operation.

3830 STORAGE CONTROL PANEL



A Toggle switch that must be in the enable position before the 3830 Storage Control is available to the channel. If the two channel switch feature (and possibly the two channel switch additional feature) is installed, a separate switch is provided for each channel.

B Toggle switch that determines how the device end generated by the drive, in a not-ready-to-ready sequence, is presented to the channel.

Multitag Position: A drive is available to a channel after it clears the device end generated by the drive in a not-ready-to-ready sequence. Before any other channel can use the drive, it must also accept the not-ready-to-ready sequence device end.

Off Position: A drive is made available to all channels after one of the channels clears the device end generated by the drive in a not-ready-to-ready sequence.

C **Power Off:** A momentary pushbutton that can be used to remove ac power from the 3330 facility.

If system power is on when the pushbutton is pressed, ac power is removed from the 3330 facility. If system power is later turned off, then on, ac power is reapplied to the 3330 facility; operation of the power on pushbutton is not required.

Power On: A momentary pushbutton that can be used to reverse the effect of the power off switch. If system power is on, and the power off switch is pressed to remove ac power from the 3330 facility, then pressing the power on switch will restore ac power to the 3330 facility.


Whenever system power is brought up, ac power is applied to the 3330 facility, regardless of what was previously done to the two pushbuttons.

See "3830 Storage Control Meter" for usage meter operation.

3330 DISK STORAGE PANEL

There is one operator panel for each pair of disk drives attached to the 3330 facility.

This panel contains switches and indicators associated with individual drives.



A Opens and closes the drawer of one disk drive to permit operator access.

B A logical address plug with one unique address (0-7) must be inserted in the socket associated with each drive. The plugs are interchangeable among drives; simply remove the plug and insert the desired one in its place.

C Starts or stops one disk drive. When the switch is on START, the drive motor starts, a brush cycle is taken, and the read/write heads load. When the switch is on STOP, the heads unload and the drive motor stops.

D Ready indicator. On when the drive is running track following, and ready for use.

E This indicator comes on if a read/write malfunction occurs in the drive.

F Write inhibit. In READ position, only read operations can be performed on the disk.

NOTE: Upper elements are for upper drive, lower elements are for lower drive.

11833D

CONSOLE ERROR MESSAGE ANALYSIS

BASIC CONSOLE ERROR MESSAGE FORMAT

- The end of each field, except the last, is marked by a comma.
- Fields are sometimes omitted. In this case it is delimited by its comma.
- The message may consist of one or two lines. Each such line uses the basic format but prints different fields.

Notes On Specific Fields: (Refer to Figure 14)

- 1** CUA = Channel/Unit Address. Three characters.
- 2** Error descriptions applicable to 3830/3330. Twelve characters.
- 3** CM = Command code of last CCW executed. Two characters.
- 4** CSW Status. First two hex characters = unit status; Second two hex characters = channel status. Four characters.
- 5** Sense bytes printed in hex character pairs. Variable length.
- 6** Seek address = BBCCHH. Twelve characters.
- 7** Serial number of the volume mounted on the device. Six characters.
- 8** Job name. Eight characters.

Error Message Content

Error Description	Fields Printed	Sense Bytes in Field 5
Intv Req'd	Line 1: 1,2,3,4,5,,7,8 Line 2: None	0-4
Equip Check Data Check	Line 1: 1,2,3,4,,7,8 Line 2: 1,,,5	0-23
Program Errors	Line 1: 1,2,3,4,,7,8 Line 2: 1,,,5,6	0-7

Typical Error Message Response

An error message has been printed on the console. Refer to message format to interpret the meaning. Read field **1** to verify that CUA is for 3830/3330 installed on the system.

Read the error description from field **2** and follow corresponding directions.

Intervention Required: One-line message. Device requiring intervention is specified in CUA field **1**

Command Reject: Two-line message code = '09YZ', where 'YZ' are hex digits of sense byte 7 in field **5**

Bus Out Check: Two-line message. Look for details in sense bytes.

Overrun: Two-line message. Look for details in sense bytes.

Equipment Check: Two-line message. The symptom code is the four hex digits of sense bytes 22 and 23 at the end of field **5**

Interface Control Check: No error message. An abnormal condition detected.

Write Inhibited: Check to ensure that write inhibit switch on the drive operator panel is set to R/W position.

Data Check: Two-line message. Data check is written for permanent, ECC-uncorrectable read errors described under Format 4. Check EREP history printouts and look for an invalid track format indication for this disk pack and seek address. If this condition occurred, suspect program problem with count field too large (writing into index).

Verify that the error is not due to a disk pack surface by moving the failing disk pack to another physical drive. If a surface defect is suspected, the operator can use utilities ATLAS or DASDI/DASDR to attempt data recovery and assign alternate surface for the defective track.

The symptom code is the 4 hex digits of sense bytes 22 and 23 at the end of field **5**.




Figure 14. Console Error Message Format

SENSE BYTE 0	
Bit 0 Command Reject	<ol style="list-style-type: none"> Invalid command code. Invalid command sequence. Invalid or incomplete argument transferred by a control command. Track formatted without home address. Write portion of file mask violated. Improper alternate/defective track pointer. A write command has been issued with the write inhibit switch in the read only position (Sense byte 1, bit 6 also one).
Bit 1 Intervention Required	<ol style="list-style-type: none"> Addressed device not physically attached to system. Addressed device not ready. Diagnostic write or diagnostic load command issued and microdiagnostic is resident in 3B30 control storage.
Bit 2 Bus Out Parity	The 3B30 has detected bad parity in data transferred from the channel.
Bit 3 Equipment Check	An unusual hardware condition originated in the channel, storage control, or drive. (Condition further defined in sense bytes 7 thru 23.)
Bit 4 Data Check	<ol style="list-style-type: none"> A correctable data error has been detected in information received from a disk drive. (Byte 2, bit 1 on, and correction information is provided in sense bytes 15 thru 22.) An uncorrectable data error has been detected in information received from a disk drive. (Condition further defined in sense byte 7.)
Bit 5 Overrun See Note 1.	<ol style="list-style-type: none"> The storage control received a byte from a drive before the last byte read was accepted by the channel. A data byte was received too late from the channel during a write operation.
Bit 6	Not used – set to 0.
Bit 7	Not used – set to 0.
<p>Note 1: The storage control posts overrun only if the condition occurs: (1) more than ten times in a CCW chain, (2) in the second or subsequent segments of an overflow record, or (3) during a format write operation.</p> <p>Detection of an overrun immediately stops data transmission. When writing, the remaining portion of the record area is padded out with 0's. With the following two exceptions, all data overrun conditions are retried by the storage control.</p> <ol style="list-style-type: none"> Data overruns that occur on the second or subsequent segments of an overflow record. Data overruns that occur during format write operations. <p>If the overrun condition exists after retry is exhausted, byte 1 bit 0 (permanent error) is posted with overrun.</p>	

SENSE BYTE 1	
Bit 0 Permanent Error	<ol style="list-style-type: none"> Storage control retry has been attempted and was unsuccessful. A drive unsafe condition has been detected and retry should not be attempted.
Bit 1 Invalid Track Format	An attempt has been made to write data exceeding track capacity.
Bit 2 End of Cylinder	<ol style="list-style-type: none"> A multi-track read or search operation has attempted to continue beyond the addressable cylinder boundary. An overflow operation has attempted to continue beyond the addressable cylinder boundary. (Byte 1 bit 7, operation incomplete, also set.)
Bit 3	Not used – set to 0.
Bit 4 No Record Found	<ol style="list-style-type: none"> Two index points sensed in the same command chain without an intervening read operation in the home address area or in a data area. Two index points sensed in the same command chain without an intervening write, sense, or control command.
Bit 5 File Protected	<ol style="list-style-type: none"> A seek command has violated the file mask. A multi-track read or search operation has violated the file mask. An overflow operation has violated the seek portion of the file mask. (Byte 1 bit 7, operation incomplete, also set.)
Bit 6	Write Inhibited.
Bit 7 Operation Incomplete	<p>One of the following conditions occurred during the processing of an overflow record:</p> <ol style="list-style-type: none"> Overflow to a file protected boundary. (Byte 1 bit 5, file protected, also set.) Overflow past the cylinder boundary. (Byte 1 bit 2, end of cylinder, also set.) A correctable data error was detected in a data field other than the last segment. (Byte 2 bit 1, correctable, also set.) A correctable data check was detected in a home address or count area associated with a segment other than the first segment. An uncorrectable data check was detected in any area associated with a segment other than the first segment. A defective or alternate track condition was detected after initiation of data transfer. A seek error was detected in the second or subsequent segment.
SENSE BYTE 2	
Bit 0	Not used – set to 0.
Bit 1 Correctable	Indicates that the data check posted in sense byte 0 bit 4 is correctable. Sense bytes 15 thru 22 identify the error pattern and error pattern displacement.
Bit 2	Not used – set to 0.
Bit 3 Environmental Data Present	Indicates that the sense bytes 8 thru 23 contain either usage/error statistics or error log information. Sense byte 7 identifies the format of bytes 8 thru 23.
Bits 4 thru 7	Not used – set to 0.

SENSE BYTE 3										
Bits 0 thru 7 Restart Command	When byte 1 bit 7 (operation incomplete) is set, this byte identifies the operation in progress when the interrupt occurred. 0000 0110 = A read operation was in progress. 0000 0101 = A write operation was in progress. When byte 1 bit 7 is zero, this byte is zero.									
SENSE BYTE 4										
Bits 0 & 1 Storage Control Identification	Provides the physical identification of the storage control as specified by the customer engineer on the storage control/drive interface card.									
Bits 2 thru 7 Drive Identification	Provides the physical address of each disk drive as follows:									
<table border="1"> <tr> <td>Drive G = 001110</td> <td>Drive E = 011100</td> <td>Drive C = 101010</td> <td>Drive A = 111000</td> <td rowspan="2">3830</td> </tr> <tr> <td>Drive H = 000111</td> <td>Drive F = 010101</td> <td>Drive D = 100011</td> <td>Drive B = 110001</td> </tr> </table>		Drive G = 001110	Drive E = 011100	Drive C = 101010	Drive A = 111000	3830	Drive H = 000111	Drive F = 010101	Drive D = 100011	Drive B = 110001
Drive G = 001110	Drive E = 011100	Drive C = 101010	Drive A = 111000	3830						
Drive H = 000111	Drive F = 010101	Drive D = 100011	Drive B = 110001							
SENSE BYTE 5										
Bits 0 thru 7 Cylinder-low	Identifies the low-order cylinder address of the most recent seek argument from the channel.									
SENSE BYTE 6										
Bit 0 Reverse	Last seek (excluding retry seeks) was in reverse direction – towards track 00.									
Bit 1 Cylinder – High	High order bit of cylinder address in sense byte 5.									
Bit 2 Difference	High order bit of difference count in sense byte 16 format 1.									
Bits 3 thru 7 Head Address See Note 2	Identifies head address of last seek (excluding retry seeks). Head address is updated during multi-track and overflow operations.									
<p>Note 2: If an alternate track condition is detected and operation incomplete is posted during an overflow operation, byte 6 is set to the head address of the defective track plus 1. This information is used by the ERP's to construct the seek argument to continue the operation.</p>										

SENSE BYTE 7	
Bits 0 thru 3 Format	<p>A Specifies the format of sense bytes 8 thru 23 as follows:</p> <p>0000 = Format 0 – Programming or system check.</p> <p>0001 = Format 1 – Disk drive equipment check.</p> <p>0010 = Format 2 – Storage control equipment check.</p> <p>0011 = Format 3 – Storage control control check.</p> <p>0100 = Format 4 – Data checks not providing displacement information.</p> <p>0101 = Format 5 – Data checks providing displacement information.</p> <p>0110 = Format 6 – Usage/error statistics</p>
Bits 4 thru 7 Message	<p>M Describes the specific nature of error conditions for each of the above formats. The "Message Table" that accompanies each format description specifies the function of the message bits for that format..</p>

A FORMAT 0 – PROGRAMMING OR SYSTEM CHECK	
SENSE BYTES 8–21 Not Used (Set to 0)	
SENSE BYTES 22 and 23 – Error Symptom Code	
M MESSAGE TABLE – FORMAT 0	
Sense byte 7- bits 4 thru 7 =	
0000	No message.
0001	Invalid command.
0010	Invalid sequence.
0011	CCW count less than required.
0100	Data value not as required.
0101	Diagnostic write not permitted by file mask.
0110	Channel discontinued retry operation.
0111	Channel returned with incorrect retry CCW.
1000	23FD – not ready.
1001	23FD – hard seek check.
1010	23FD hard read check.
1011	Improper alternate/defective track pointer.
1100	SERDES malfunction - no ST 4's.
1101	Diagnostic write control code mismatch.
1110	Control storage busy with microdiagnostic.
1111	Retry byte count/sector value incorrect.
B FORMAT 1 – DISK DRIVE EQUIPMENT CHECK	
SENSE BYTE 8 – MODULE STATUS	
Bit 0	Index error.
Bit 1	Offset active.
Bit 2	Seek incomplete.
Bit 3	Seek complete.
Bit 4	Online.
Bit 5	Attention.
Bit 6	Busy.
Bit 7	Record ready.
SENSE BYTE 9 – MONITOR MODE	
Bit 0	Not used.
Bit 1	Diagnostic 4.
Bit 2	Diagnostic 2.
Bit 3	Diagnostic 1.
Bit 4	Not used.
Bit 5	Mode 4.
Bit 6	Mode 2.
Bit 7	Mode 1.

SENSE BYTE 10 – MONITOR STATE			
Bit 0	Monitor state 8.		
Bit 1	Monitor state 7.		
Bit 2	Monitor state 6.		
Bit 3	Monitor state 5.		
Bit 4	Monitor state 4.		
Bit 5	Monitor state 3.		
Bit 6	Monitor state 2.		
Bit 7	Monitor state 1.		
SENSE BYTE 11 – CHECK STATUS			
Bit 0	CE program stop.		
Bit 1 thru 3	Not used.		
Bit 4	CUDI bus-out parity.		
Bit 5	Monitor Check.		
Bit 6	Not used.		
Bit 7	Drive command reject.		
SENSE BYTE 12 – SAFETY			
Bit 0	Data safety.	Bit 6	Not heads loaded.
Bit 1	Servo safety.	Bit 7	Even
Bit 2	Not used.		
Bit 3	Not used.		
Bit 4	Power on reset.		
Bit 5	Not used.		
SENSE BYTE 13 – TA REG/EXPECTED			
Expected data for messages 1, 6, 7, 8, 9 otherwise TA register.			
SENSE BYTE 14 – ND REG/RECEIVED			
Bits 0 thru 7	Drive status for message 9 or ND reg.		
SENSE BYTE 15 – TAG BUS/TD REG			
Bits 0 thru 7	Contents of TD register.		
SENSE BYTE 16 – 20 Not Used			
Bits 0 thru 7	Unused		
SENSE BYTE 21– CUDI CHECK			
Bit 0	Drive selection error.		
Bit 1	Tag invalid.		
Bit 2	Device check.		
Bit 3	TA register check.		
Bit 4	CUDI register check.		
Bit 5	TD register check.		
Bit 6	Not used.		
Bit 7	Not used.		
SENSE BYTES 22 AND 23 ERROR SYMPTOM CODE			
M MESSAGE TABLE – FORMAT 1			
Sense byte 7- bits 4 thru 7 =			
0000	No message.		
0001	Set target error.		
0010	Not used.		
0011	No write gate at drive.		
0100	No write current sense.		

0101	Not used.	
0110	Transmit cylinder error.	
0111	Transmit head error.	
1000	Transmit difference error.	
1001	File status not as expected.	
1010	Seek error.	
1011	Seek incomplete on retry.	
1100	No interrupt from drive.	
1101 - 1111	Not used.	
C FORMAT 2 – STORAGE CONTROL EQUIPMENT CHECK		
SENSE BYTE 8 – CONTROL CHECK		
TWO CHANNEL SWITCH		TWO CHANNEL SWITCH ADDITIONAL
Bit 0	Channel buffer read error	Channel buffer read error
Bit 1	Channel A check	Channel A or C check
Bit 2	Channel B check	Channel B or D check
Bit 3	Data transfer error	Data transfer error
Bit 4	SERDES, control unit/device interface check or ECC check. (Further defined in sense bytes 9 thru 11.)	SERDES, control unit/device interface check or ECC check. (Further defined in sense bytes 9 thru 11.)
Bit 5	PLO check	PLO check
Bit 6	Sector count check	Sector count check
Bit 7	Not used	Multi-connect, Channel C or D check
SENSE BYTE 9 – SERDES CHECK		
Bit 0	CUDI check.	
Bit 1	Write parity check.	
Bit 2	Read parity check.	
Bit 3	Bit ring check.	
Bit 4	Write compensation check.	
Bit 5	ECC check.	
Bit 6	Missing PLO.	
Bit 7	VFO phase.	
SENSE BYTE 10 – ECC CHECK		
Bit 0	No input data received.	
Bit 1	P0 or write error.	
Bit 2	P1 or P3 error.	
Bit 3	P2 error.	
Bits 4-7	Zero	
SENSE BYTE 11 – Not Used		
SENSE BYTE 12 Not Used – Set to 0		
SENSE BYTE 13 Contents of TA register.		
SENSE BYTE 14 Contents of ND register		
SENSE BYTE 15 Contents of TD register.		
SENSE BYTES 16-20 Not Used – Set to 0		
SENSE BYTE 21-CUDI CHECK		
Bit 0	Drive selection error.	
Bit 1	Tag invalid.	
Bit 2	Device check.	
Bit 3	TA register check.	
Bit 4	CUDI register check.	
Bit 5	TD register check.	
Bit 6	Not used.	
Bits 6 & 7	Zero	
SENSE BYTES 22 AND 23 ERROR SYMPTOM CODE		

SENSE BYTE 7	
Bits 0 thru 3 Format	Specifies the format of sense bytes 8 thru 23 as follows: 0000 = Format 0 – Programming or system check. ----- CE information ----- 0001 = Format 1 – Disk drive equipment check. 0010 = Format 2 – Storage control equipment check. D 0011 = Format 3 – Storage control control check. E 0100 = Format 4 – Data checks not providing displacement information. F 0101 = Format 5 – Data checks providing displacement information. 0110 = Format 6 – Usage/error statistics
Bits 4 thru 7 Message	M Describes the specific nature of error conditions for each of the above formats. The "Message Table" that accompanies each format description specifies the function of the message bits for that format.

MESSAGE TABLE – FORMAT 2	
Sense byte 7- bits 4 thru 7 = 0000 0001 0010 0011-1111	No message. ECC P1 or P3 compare failure. ECC P2 compare failure. Unused.
D	FORMAT 3 – STORAGE CONTROL – CONTROL CHECK
SENSE BYTE 8 – FAILING INSTRUCTION ADDRESS (1)	
Bits 0 thru 7	High order address byte of control storage word addressed when error was detected.
SENSE BYTE 9 – FAILING INSTRUCTION ADDRESS (2)	
Bits 0 thru 7	Low order address byte of control storage word addressed when error was detected.
SENSE BYTE 10 – ERROR LATCHES (1)	
Bit assignment is dependent upon the state of bit 0 as follows:	
Bit 0 (on) One Bit 0 (off) Zero	
Bit 1 Clock.	Clock
Bit 2 CA decode even.	CS decode.
Bit 3 CA decode odd.	Zero.
Bit 4 CB decode even.	A register.
Bit 5 CB decode odd.	B register.
Bit 6 Branch status.	ALU.
Bit 7 Special operation.	23FD parity.
SENSE BYTE 11 – ERROR LATCHES (2)	
Bit assignment is dependent upon the state of byte 10 – bit 0 as follows:	
Bit 0 (on) Not used Bit 0 (off) Storage address bus 1 – 7.	
Bit 1 Storage read multiple 0/1.	Storage address bus 8 – 13.
Bit 2 Storage ECC multiple 2/3.	Storage write bus 0/2.
Bit 3 Not used.	Storage write bus 1 – 3
Bit 4 Cycle control.	Address bus 1 – 13 low.
Bit 5 CD decode.	Address bus 1 – 13 high.
Bit 6 Not used.	23FD not ready.
Bit 7 Not used.	Zero
SENSE BYTE 12 – STORAGE ERROR PATTERN	

Bits 0 thru 7	Identifies the failing bits of a control storage cycle.
SENSE BYTE 13 – T REGISTER (1)	
Bits 0 thru 7	Contains the contents of the TC register after an unsolicited selective reset. The TC register is reset if selective reset is in response to disconnect in from storage control.
SENSE BYTE 14 – T REGISTER (2)	
Bits 0 thru 7	Contains the contents of the TG register after an unsolicited selective reset. The TG register is reset if selective reset is in response to disconnect in from the storage control.
SENSE BYTES 15–21 Not Used (Set to 0) SENSE BYTES 22 and 23 – Error Symptom Code	
M	MESSAGE TABLE FORMAT 3
Sense byte 7- bits 4 thru 7 = 0000 0001 thru 1111	No message. Not used.
E	FORMAT 4 – DATA CHECKS NOT PROVIDING DISPLACEMENT INFORMATION (See Note)
SENSE BYTE 8 – CYLINDER (1)	
Bits 0 thru 7	High order cylinder byte of last seek address.
SENSE BYTE 9 – CYLINDER (2)	
Bits 0 thru 7	Low order cylinder byte of last seek address.
SENSE BYTE 10 – HEAD (1)	
Bits 0 thru 7	High order head byte of last seek address.
SENSE BYTE 11 – HEAD (2)	
Bits 0 thru 7	Low order head byte of last seek address.
SENSE BYTE 12 – RECORD	
Bits 0 thru 7	Record number of record in error.
SENSE BYTE 13 – SECTOR	
Bits 0 thru 7	Sector number of record in error.
SENSE BYTE 14 OFFSET	
Bits 0 thru 7	Amount of offset used to recover from error.
SENSE BYTE 15-RETRIES	
Bits 0 thru 7	Number of retries required to recover from error.
SENSE BYTE 16 – SOURCE DRIVE IDENTIFICATION	

Bits 0 and 1	Identifies the storage control that was used to record the data in which the error occurred.									
Bits 2 thru 7	Identifies the disk drive that was used to record the data in which the error occurred. Drive ID is as follows:									
	<table border="1"> <tr> <td>Drive G = 001110</td> <td>Drive E = 011100</td> <td>Drive C = 101010</td> <td>Drive A = 111000</td> <td rowspan="2">3830</td> </tr> <tr> <td>Drive H = 000111</td> <td>Drive F = 010101</td> <td>Drive D = 100011</td> <td>Drive B = 110001</td> </tr> </table>	Drive G = 001110	Drive E = 011100	Drive C = 101010	Drive A = 111000	3830	Drive H = 000111	Drive F = 010101	Drive D = 100011	Drive B = 110001
Drive G = 001110	Drive E = 011100	Drive C = 101010	Drive A = 111000	3830						
Drive H = 000111	Drive F = 010101	Drive D = 100011	Drive B = 110001							
SENSE BYTE 17 thru 21 NOT USED – SET TO 0										
SENSE BYTES 22AND 23 ERROR SYMPTOM CODE										
M	MESSAGE TABLE – FORMAT 4									
Sense byte 7- bits 4 thru 7 = 0000 0001 0010 0011 0100 0101 0110 0111 1000 1001 1010-1111	HA field ECC uncorrectable. Count field ECC uncorrectable. Key field ECC uncorrectable. Data field ECC uncorrectable. HA field no sync byte found. Count field no sync byte found. Key field no sync byte found. Data field no sync byte found. Unused AM detection failure on retry. Unused									
F	FORMAT 5 – DATA CHECKS PROVIDING DISPLACEMENT INFORMATION (See Note)									
SENSE BYTE 8 – CYLINDER (1)										
Bits 0 thru 7	High order cylinder byte of last seek address.									
SENSE BYTE 9 – CYLINDER (2)										
Bits 0 thru 7	Low order cylinder byte of last seek address.									
SENSE BYTE 10 – HEAD (1)										
Bits 0 thru 7	High order head byte of last seek address.									
SENSE BYTE 11 – HEAD (2)										
Bits 0 thru 7	Low order head byte of last seek address.									
SENSE BYTE 12 – RECORD										
Bits 0 thru 7	Record number of record in error.									
SENSE BYTE 13 – SECTOR										
Bits 0 thru 7	Sector number of record in error.									
SENSE BYTE 14 – OFFSET										
Bits 0 thru 7	Amount of offset used to recover from error									
<i>Note: Contents of bytes 8-12 are unreliable if message code in byte 7 is either 0 or 4 (error occurred in HA), 1 or 5 (error occurred in count field), or 9 (AM detection on retry).</i>										

SENSE BYTE 7	
Bits 0 thru 3 Format	Specifies the format of sense bytes 8 thru 23 as follows: 0000 = Format 0 – Programming or system check. CE information 0001 = Format 1 – Disk drive equipment check. 0010 = Format 2 – Storage control equipment check. 0011 = Format 3 – Storage control control check. 0100 = Format 4 – Data checks not providing displacement information. 0101 = Format 5 – Data checks providing displacement information. G 0110 = Format 6 – Usage/error statistics
Bits 4 thru 7 Message	M Describes the specific nature of error conditions for each of the above formats. The "Message Table" that accompanies each format description specifies the function of the message bits for that format..

BYTES 15 thru 17 – RESTART DISPLACEMENT	
Specifies the number of bytes processed by the storage control to end of data field in error.	
BYTES 18 and 19 – ERROR DISPLACEMENT	
Displacement of first byte in error relative to end of the data field where error occurred.	
BYTES 20 thru 22 – ERROR PATTERN	
Contain error pattern used for error correction function. See "Error Correction Function."	
SENSE BYTE 23	
Bits 0-6	Not used - set to zero
Bit 7	Channel truncation

MESSAGE TABLE – FORMAT 5	
Sense byte 7- bits 4 thru 7	0000 HA field correctable 0001 Count field correctable 0010 Key field correctable 0011 Data field correctable 0100-1111 Unused
FORMAT 6 – USAGE/ERROR STATISTICS	
SENSE BYTES 8 thru 11 – BYTES READ/SEARCHED	
These four bytes provide an accumulated count of the number of bytes processed by the storage control in read or search operations. Bytes processed during retry operations are not included in this count. Only key and data field counts are accumulated.	
SENSE BYTES 12 and 13 – CORRECTABLE DATA CHECKS	
These two bytes provide an accumulated count of the number of ECC correctable data checks which were detected by the storage control.	
SENSE BYTES 14 and 15 – RETRY DATA CHECKS	
These two bytes identify the number of ECC uncorrectable data checks which were successfully retried by the storage control.	
SENSE BYTES 16 and 17 – SEEKS	
These two bytes provide a count of the number of access motions initiated by the channel.	
SENSE BYTE 18	
Bit 0 set to zero = bytes 20-23 contain information for interfaces A and B. Bit 0 set to one = bytes 20-23 contain information for interfaces C and D.	

SENSE BYTE 19 – SEEK ERRORS	
Bits 0 thru 7	Identifies the total number of seek errors which were successfully retried by the storage control.
SENSE BYTE 20 – COMMAND OVERRUN A or C.	
Bits 0 thru 7	Provides a count of the number of command overruns which were retried by the storage control for channel A or C.
SENSE BYTE 21 – DATA OVERRUN A or C.	
Bits 0 thru 7	Provides a count of the number of data overruns which were retried by the storage control for channel A or C.
SENSE BYTE 22 – COMMAND OVERRUN B or D.	
Bits 0 thru 7	Provides a count of the number of command overruns which were retried by the storage control for channel B or D.
SENSE BYTE 23 – DATA OVERRUN B or D.	
Bits 0 thru 7	Provides a count of the number of data overruns which were retried by the storage control for channel B or D.
MESSAGE TABLE – FORMAT 6	
Sense byte 7- bits 4 thru 7	0000 No message. 0001 thru 1111 Not used.

| APPENDIX B. RECORD/TRACK CAPACITIES

CAPACITIES WITH KEYS							
BYTES PER RECORD (KL+DL)		RECORDS PER				BYTES PER PACK (KL+DL)	
MINIMUM	MAXIMUM	TRACK	CYLINDER	PACK	8 DRIVES	MINIMUM	MAXIMUM
2	2	68	1292	521968	4175744	1043936	1043936
3	5	67	1273	514292	4114336	1542876	2571460
6	8	66	1254	506616	4052928	3039696	4052928
9	11	65	1235	498940	3991520	4490460	5488340
12	14	64	1216	491264	3930112	5895168	6877696
15	17	63	1197	483588	3868704	7253820	8220996
18	21	62	1178	475912	3807296	8566416	9994152
22	24	61	1159	468236	3745888	10301192	11237664
25	28	60	1140	460560	3684480	11514000	12895680
29	32	59	1121	452884	3623072	13133636	14492288
33	35	58	1102	445208	3561664	14691864	15582280
36	39	57	1083	437532	3500256	15751152	17063748
40	44	56	1064	429856	3438848	17194240	18913664
45	48	55	1045	422180	3377440	18998100	20264640
49	52	54	1026	414504	3316032	20310696	21554208
53	57	53	1007	406828	3254624	21561884	23189196
58	62	52	988	399152	3193216	23150816	24747424
63	67	51	969	391476	3131808	24662988	26228892
68	72	50	950	383800	3070400	26098400	27633600
73	77	49	931	376124	3008992	27457052	28961548
78	83	48	912	368448	2947584	28738944	30581184
84	89	47	893	360772	2886176	30304848	32108708
90	95	46	874	353096	2824768	31778640	33544120
96	101	45	855	345420	2763360	33160320	34887420
102	108	44	836	337744	2701952	34449888	36476352
109	115	43	817	330068	2640544	35977412	37957820
116	122	42	798	322392	2579136	37397472	39331824
123	130	41	779	314716	2517728	38710068	40913080
131	138	40	760	307040	2456320	40222240	42371520
139	146	39	741	299364	2394912	41611596	43707144
147	155	38	722	291688	2333504	42878136	45211640
156	164	37	703	284012	2272096	44305872	46577968
165	174	36	684	276336	2210688	45595440	48082464
175	185	35	665	268660	2149280	47015500	49702100
186	196	34	646	260984	2087872	48543024	51152864
197	207	33	627	253308	2026464	49901676	52434756
208	220	32	608	245632	1965056	51091456	54039040
221	233	31	589	237956	1903648	52588276	55443748
234	247	30	570	230280	1842240	53885520	56879160
248	262	29	551	222604	1780832	55205792	58322248
263	279	28	532	214928	1719424	56526064	59964912
280	296	27	513	207252	1658016	58030560	61346592
297	315	26	494	199576	1596608	59274072	62866440
316	335	25	475	191900	1535200	60640400	64286500
336	357	24	456	184224	1473792	61899264	65767968
358	381	23	437	176548	1412384	63204184	67264788
382	407	22	418	168872	1350976	64509104	68730904
408	435	21	399	161196	1289568	65767968	70120260
436	467	20	380	153520	1228160	66934720	71693840
468	501	19	361	145844	1166752	68254992	73067844
502	540	18	342	138168	1105344	69360336	74610720
541	583	17	323	130492	1043936	70596172	76076836
584	631	16	304	122816	982528	71724544	77496896
632	686	15	285	115140	921120	72768480	78986040
687	749	14	266	107464	859712	73827768	80490536
750	821	13	247	99788	798304	74841000	81925948
822	906	12	228	92112	736896	75716064	83453472
907	1005	11	209	84436	675488	76583452	84858180
1006	1125	10	190	76760	614080	77220560	86355000
1126	1271	9	171	69084	552672	77788584	87805764
1272	1454	8	152	61408	491264	78110976	89287232
1455	1689	7	133	53732	429856	78180060	90753348
1690	2003	6	114	46056	368448	77834640	92250168
2004	2442	5	95	38380	307040	76913520	93723960
2443	3100	4	76	30704	245632	75009872	95182400
3101	4197	3	57	23028	184224	71409828	96648516
4198	6391	2	38	15352	122816	64447696	98114632
6392	12974	1	19	7676	61408	49064992	99588424

CAPACITIES WITHOUT KEYS							
BYTES PER RECORD (KL+DL)		RECORDS PER				BYTES PER PACK (KL+DL)	
MINIMUM	MAXIMUM	TRACK	CYLINDER	PACK	8 DRIVES	MINIMUM	MAXIMUM
1	2	96	1824	736896	5895168	736896	1473792
3	3	95	1805	729220	5833760	2187660	2187660
4	5	94	1786	721544	5772352	2886176	3607720
6	6	93	1767	713868	5710944	4283208	4283208
7	8	92	1748	706192	5649536	4943344	5649536
9	9	91	1729	698516	5588128	6286644	6286644
10	11	90	1710	690840	5526720	6908400	7599240
12	12	89	1691	683164	5465312	8197968	8197968
13	14	88	1672	675488	5403904	8781344	9456832
15	16	87	1653	667812	5342496	10017180	10684992
17	18	86	1634	660136	5281088	11222312	11882448
19	19	85	1615	652460	5219680	12396740	12396740
20	21	84	1596	644784	5158272	12895680	13540464
22	23	83	1577	637108	5096864	14016376	14653484
24	25	82	1558	629432	5035456	15106368	15735800
26	27	81	1539	621756	4974048	16165656	16787412
28	29	80	1520	614080	4912640	17194240	17808320
30	31	79	1501	606404	4851232	18192120	18798524
32	33	78	1482	598728	4789824	19159296	19758024
34	35	77	1463	591052	4728416	20095768	20686820
36	38	76	1444	583376	4667008	21001536	22168288
39	40	75	1425	575700	4605600	22452300	23028000
41	42	74	1406	568024	4544192	23288984	23857008
43	45	73	1387	560348	4482784	24094964	25215660
46	47	72	1368	552672	4421376	25422912	25975584
48	50	71	1349	544996	4359968	26159808	27249800
51	53	70	1330	537320	4298560	27403320	28477960
54	55	69	1311	529644	4237152	28600776	29130420
56	58	68	1292	521968	4175744	29230208	30274144
59	61	67	1273	514292	4114336	30343228	31371812
62	64	66	1254	506616	4052928	31410192	32423424
65	67	65	1235	498940	3991520	32431100	33428980
68	70	64	1216	491264	3930112	33405952	34388480
71	73	63	1197	483588	3868704	34334748	35301924
74	77	62	1178	475912	3807296	35217488	36645224
78	80	61	1159	468236	3745888	36522408	37458880
81	84	60	1140	460560	3684480	37305360	38687040
85	88	59	1121	452884	3623072	38495140	39853792
89	91	58	1102	445208	3561664	39623512	40513928
92	95	57	1083	437532	3500256	40252944	41565540
96	100	56	1064	429856	3438848	41266176	42985600
101	104	55	1045	422180	3377440	42640180	43906720
105	108	54	1026	414504	3316032	43522920	44766432
109	113	53	1007	406828	3254624	44344252	45971564
114	118	52	988	399152	3193216	45503328	47099936
119	123	51	969	391476	3131808	46585644	48151548
124	128	50	950	383800	3070400	47591200	49126400
129	133	49	931	376124	3008992	48519996	50024492
134	139	48	912	368448	2947584	49372032	51214272
140	145	47	893	360772	2886176	50508080	52311940
146	151	46	874	353096	2824768	51552016	53317496
152	157	45	855	345420	2763360	52503840	54230940
158	164	44	836	337744	2701952	53363552	55390016
165	171	43	817	330068	2640544	54461220	56441628
172	178	42	798	322392	2579136	55451424	57385776
179	186	41	779	314716	2517728	56334164	58537176
187	194	40	760	307040	2456320	57416480	59565760
195	202	39	741	299364	2394912	58375980	60471528
203	211	38	722	291688	2333504	59212664	61546168
212	220	37	703	284012	2272096	60210544	62482640
221	230	36	684	276336	2210688	61070256	63557280
231	241	35	665	268660	2149280	62060460	64747060
242	252	34	646	260984	2087872	63158128	65767968
253	263	33	627	253308	2026464	64086924	66620004

CAPACITIES WITHOUT KEYS							
BYTES PER RECORD (KL+DL)		RECORDS PER				BYTES PER PACK (KL+DL)	
MINIMUM	MAXIMUM	TRACK	CYLINDER	PACK	8 DRIVES	MINIMUM	MAXIMUM
264	276	32	608	245632	1965056	64846848	67794432
277	289	31	589	237956	1903648	65913812	68769284
290	303	30	570	230280	1842240	66781200	69774840
304	318	29	551	222604	1780832	67671616	70788072
319	335	28	532	214928	1719424	68562032	72000880
336	352	27	513	207252	1658016	69636672	72952704
353	371	26	494	199576	1596608	70450328	74042696
372	391	25	475	191900	1535200	71386800	75032900
392	413	24	456	184224	1473792	72215808	76084512
414	437	23	437	176548	1412384	73090872	77151476
438	463	22	418	168872	1350976	73965936	78187736
464	491	21	399	161196	1289568	74794944	79147236
492	523	20	380	153520	1228160	75531840	80290960
524	557	19	361	145844	1166752	76422256	81235108
558	596	18	342	138168	1105344	77097744	82348128
597	639	17	323	130492	1043936	77903724	83384388
640	687	16	304	122816	982528	78602240	84374592
688	742	15	285	115140	921120	79216320	85433880
743	805	14	266	107464	859712	79845752	86508520
806	877	13	247	99788	798304	80429128	87514076
878	962	12	228	92112	736896	80874336	88611744
963	1061	11	209	84436	675488	81311868	89586596
1062	1181	10	190	76760	614080	81519120	90653560
1182	1327	9	171	69084	552672	81657288	91674468
1328	1510	8	152	61408	491264	81549824	92726080
1511	1745	7	133	53732	429856	81189052	93762340
1746	2059	6	114	46056	368448	80413776	94829304
2060	2498	5	95	38380	307040	79062800	95873240
2499	3156	4	76	30704	245632	76729296	96901824
3157	4253	3	57	23028	184224	72699396	97938084
4254	6447	2	38	15352	122816	65307408	98974344
6448	13030	1	19	7676	61408	49494848	100018280

Track Capacity

The number of records that can be recorded on a track depends on the record size. The following equation is used to determine the number of equal length records per track. Home address and standard RO space are accounted for.

$$\text{Number of equal length records per track} = \frac{13,165}{135 + C + KL + DL} \quad \begin{matrix} \text{(track capacity)} \\ \text{(bytes per record)} \end{matrix}$$

where

$$C \text{ (constant)} = \begin{matrix} 0 & \text{if } KL = 0 \\ 56 & \text{if } KL \neq 0 \end{matrix}$$

KL = key length
DL = data length

25581C

(This page intentionally left blank.)

A

- Access Time 2
- Address Plugs 1, 81
- Address Word, Channel 9
- Addressing
 - Device 14
 - Precautions 14
 - With Two Channel Switch 72

B

- Block Multiplexer Channel 1
- Bus Out Parity 83
- Busy Indicator 11

C

- Capacity
 - Record 91
 - Track 91
- Chain Command Flag 10
- Chaining, Command 12
- Chaining, Data 13
- Changing a Drive Address 79
- Channel Address Word
 - Description of 9
 - Format 9
- Channel, Block Multiplexer 1
- Channel Address Word (CAW) 9
- Channel Command Word
 - Description of 9
 - Format 10
- Channel Commands 15
- Channel End Indicator 11
- Channel Operation 9
- Channel Programs
 - Branching in 13
 - Example 1 57
 - Example 2 60
 - Example 3 62
- Channel Selection Switch 71
- Channel Status Word
 - Description of 9
 - Format 11
 - Status Presentation 9
- Command Address 11
- Command Chaining
 - Description of 12
 - Flag 10
 - Restrictions 13
- Command Codes 10, 17
- Command Overrun 89
- Command Reject 83
- Command Retry
 - Description of 69
 - Introduction to 1
 - Program Indications 70
- Console Error Message 82
- Construction of Restart CCW's 77
 - Restart CCW 1 77
 - Restart CCW 2 77
- Contingent Connection 12
- Control Check 85
- Control Commands
 - Command Codes 17
 - Diagnostic Load 28

Control Commands (Continued)

- Diagnostic Write 29
- General Description 15
- No-Op 18
- Recalibrate 19
- Restore 26
- Seek 20
- Seek Cylinder 21
- Seek Head 22
- Set File Mask 24
- Set Sector 25
- Space Count 23
- Transfer in Channel 27
- Control Unit End 11
- Correctable (Sense Bit) 83
- Correction Code Check 7
- Correction, Error 7
- Count Area 2, 5
- CPU Parity 7
- CUDI Check 85
- Cylinder Number 5

D

- Data Address 10
- Data Area 4, 5
- Data Chaining, Description of 13
- Data Check 83
- Data Integrity 7
- Data Length 5
- Data Overrun 89
- Data Rate 2
- Data Records 4
 - Count Area Format 5
 - Key Area Format 5
- Detection, Error 7
- Device Addressing (see Addressing)
- Device End Indicator 11
- Device Release 47
- Device Reserve 48
- Device Status 71
- Diagnostic Load 28
- Diagnostic Write 29
- Diagnostics, Storage Control 70
- Disk Pack
 - Cover Storage 79
 - Description of 2
 - Initialization 2
 - Loading 79
 - Unloading 79
- Disk Storage
 - Functional Description 3
 - General Description 1
 - Highlights 1
 - Special Features of 1
 - Standard Features of 1
- Disk Storage Panel, 3330 81
- Drive Address, Changing 79
 - Initial Assignment 14
- Drive Identification 83

E

- ECC Check 85
- Enable/Disable Switch 70
- End of Cylinder 83

- End-of-File
 - Description of 66
 - Example of 67
- Ending Status 12
- Environmental Data Present 83
- Equipment Check 83
- Erase 52
- Error Correction Code 7
- Error Correction Function
 - Construction of Restart CCW's 77
 - Error Displacement 73
 - Error Pattern 73
 - Example of 77
 - Restart CCW 1 77
 - Restart CCW 2 77
- Error Correction Table 74
- Error Detection and Correction
 - Correction Code Check 7
 - CPU Parity 7
 - Data Integrity 7
 - Error Correction Code 7
 - Error Correction Function 73
 - Error Displacement 89
 - Error Recovery Procedures 73
 - Full Read-Back Check 7
 - Latches 87
- Error Recovery Procedures
 - Error Correction Function 73
 - Error Correction Table 74
 - Error Message Format 82
 - Recovery Action Table 74
- Equipment Check 83, 85

F

- File Mask 24
- File Protected 83
- Flag 5
- Format, Record 2
- Formula
 - Track Capacity 93
 - Sector Location 67
- Full Read-Back Check 7

G

- Gap 1, 2, 3 5
- Gaps, Purpose of 7

H

- Halt Device 8
- Halt I/O 8
- Head Number 5
- Home Address
 - Cylinder Number 5
 - Data Length 4
 - Error Correction Code 5
 - Flag 5
 - Gap 5
 - Head Number 5
 - Read 37
 - Search Equal 30
 - Write 30

I

- I/O Instructions
 - Format of 8
 - Halt Device 8
 - Halt I/O 8

- I/O Instructions (Continued)
 - Start I/O 8
 - Start I/O Fast Release 8
 - Test I/O 8
- Index Marker 5
- Indicator, Select Lock 81
- Initial Program Load 43
- Initial Status 12
- Initialization Pack 2
- Input/Output Operations
 - Channel Operation 9
 - General Description 8
 - Instruction Format 8
- Interrupt Code 13
- Interrupt, Program Control 10
- Intervention Required 83
- Invalid Track Format 83

K

- Key Area 2, 5
- Key Field 5
- Key Length 5

L

- Loading Disk Packs 79
- Logical Address Plug 1, 81

M

- Message, Console Error 82
- Message Table
 - Description of Formats 0-6 85-89
- Meter, 3830 Storage Control 70
- Module Status 85
- Monitor Mode 85
- Monitor State 85
- Multi-tag Switch 80
- Multi-tagged Status 71
- Multiple Requesting
 - Description of 69
 - Introduction to 1
- Multiple Track (MT) Operation
 - Description of 64
 - Example of 64

N

- ND Register 85
- No-Op 18
- No Record Found 83

O

- Open/Close Switch 81
- Operating Instructions
 - Changing a Drive Address 79
 - Loading, Disk Pack 79
 - Logical Address Plug 81
 - Multi-tag Switch 80
 - Open/Close Switch 81
 - Power Off Switch 80
 - Power On Switch 80
 - Ready Indicator 81
 - Select Lock Indicator 81

Operating Instructions (Continued)

- Start/Stop Switch 81
- Unloading, Disk Pack 79
- Usage Meter 80
- 3330 Disk Storage Panel 81
- 3830 Storage Control Panel 80
- Operation Code 8
- Operation Incomplete 83
- Overflow Records (see also Record Overflow)
 - Formatting 65
 - Processing 66
 - Reading 65
 - Spacing Over 66
 - Updating 65
- Overrun 83

P

- Pending Status 12
- Permanent Error 83
- Power Off Switch 80
- Power On Switch 80
- Problem State 8
- Program Control Interruption 10
- Program Mask 13
- Program Status Word
 - Description of 12
 - Format 13
 - New 12
 - Old 12
- Protection Key 11

R

- R0 Record (see Track Descriptor Record (R0))
- Read and Reset Buffered Log 46
- Read Commands
 - Command Codes 17
 - General Description 15
 - Read Count 38
 - Read Count, Key, and Data 42
 - Read Data 40
 - Read Home Address 37
 - Read IPL 43
 - Read Key and Data 41
 - Read R0 39
 - Read Sector 44
- Read Count 38
- Read Count, Key, and Data 42
- Read Data 40
- Read Diagnostic Status 1 49
- Read Home Address 37
- Read IPL 43
- Read Key and Data 41
- Read R0 39
- Read Sector 44
- Ready Indicator 81
- Recalibrate 19
- Record Capacities
 - Maximum 2
 - With Keys 91
 - Without Keys 92
- Record Format
 - Count Area 2
 - Data Area 4
 - Key Area 2
- Record Number 5

- Record Overflow (see also Overflow Records)
 - Description of 64
 - Example of 65
 - Head Switching Restrictions 65
 - Introduction to 1
- Record Zero (see Track Descriptor Record (R0)) 5
- Recovery Action Table 74
- Resets 71
- Restart CCW 1 77
- Restart CCW 2 77
- Restart Command 83
- Restart Displacement 89
- Restore 26
- Rotational Delay 2
- Rotational Position Sensing RPS
 - Description of 67
 - Example of 68
 - Introduction to 1
- Retry Data Checks 89

S

- Search Commands
 - Command Codes 17
 - General Description 15
 - Search Home Address 30
 - Search ID Equal 31
 - Search ID Equal or High 33
 - Search ID High 32
 - Search Key Equal 34
 - Search Key Equal or High 36
 - Search Key High 35
- Search Home Address Equal 30
- Search ID Equal 31
- Search ID Equal or High 33
- Search ID High 32
- Search Key Equal 34
- Search Key Equal or High 36
- Search Key High 35
- Sector
 - Calculating 67
 - Concept 67
 - Set 25
 - Read 44
- Seek 20
- Seek Cylinder 21
- Seek Errors 89
- Seek Head 22
- Select Lock Indicator 81
- Sense Bytes, Summary of 83-90
- Sense Command
 - Command Codes 17
 - Device Release 47
 - Device Reserve 48
 - General Description 16
 - Read and Reset Buffered Log 46
 - Read Diagnostic Status 1 49
 - Sense I/O 45
 - Sense Bytes, Summary of 83-90
- Sense I/O 45
- Set File Mask 24
- Set Sector 25
- Skip Flag 10, 13
- Space Count 23
- Special Features 71-72
- Speed 2
- Standard Features 64-70

- Start I/O 8
- Start I/O Fast Release 8
- Start/Stop Switch 81
- Statistical Usage/Error Recording (see Usage/Error Recording)
- Status Presentation
 - Channel Status Word 9
 - Ending Status 12
 - Initial Status 12
 - Pending Status 12
- Storage Control
 - Functional Description 3
 - General Description 1
 - Highlights 1
 - Special Features of 1
 - Standard Features of 1
- Storage Control Diagnostics 70
- Storage Control Identification 83
- Storage Control Panel, 3830 80
- Supervisor State 8
- Switches
 - Enable/Disable 80
 - Multi-tag 80
 - Open/Close 81
 - Power 80
 - R/W - Read 81
 - Start-Stop 81
- Suppress Length Indicator (SLI) 10
- System Mask 13

T

- T Register 87
- Tagged Status 71
- Test I/O 8
- Test I/O Command 16
- TIC (see Transfer In Channel)
- Track Capacities 91
 - Formula for 93
 - Maximum 2
 - With Keys 91
 - Without Keys 92
- Track Descriptor Record (R0)
 - Count Area
 - Cylinder Number 5
 - Data Length 5
 - Error Correction Code 5
 - Flag 5
 - Head Number 5
 - Key length 5
 - Record Number 5
 - Data Area
 - Data Field 5
 - Error Correction Code 5
 - Description of 4
 - Key Area
 - Gap 2 5
 - Key Field 5
- Track Format
 - Data Records 4
 - Home Address 4

- Track Format (Continued)
 - Track Descriptor Record (R0) 4
- Transfer In Channel (TIC) 27
- Two Channel Switch and Two Channel Switch Additional Addressing 71
 - Channel Selection Switch 71
 - Description of 71
 - Device Release 47
 - Device Reserve 48
 - Device Status 72
 - Introduction to 1
 - Multi-tagged Status 71
 - Resets 71
 - Tagged Status 71
 - Untagged Status 71

U

- Unit Check Indicator 11
- Unit Exception Indicator 11
- Unit Selection 14
- Unloading Disk Packs 79
- Untagged Status 71
- Usage/Error Recording
 - Description of 70
 - Introduction to 1
- Usage Meter 70, 80

W

- Write Commands
 - Command Codes 17
 - Erase 52
 - Formatting 15
 - General Description 15
 - Update 16
 - Write Count, Key, and Data 53
 - Write Data 55
 - Write Home Address 50
 - Write Key and Data 56
 - Write R0 51
 - Write Special Count, Key, and Data 54
- Write Count, Key, and Data 53
- Write Data 55
- Write Home Address 50
- Write Inhibit 81
- Write Key and Data 56
- Write R0 51
- Write Special Count, Key, and Data 54

- 2880 Block Multiplexer Channel 1
- 3330 Disk Storage Panel 81
- 3336 Disk Pack 2
- 3830 Storage Control Meter 70
- 3830 Storage Control Panel 80

READER'S COMMENT FORM

3830 Storage Control Model 1
and 3330 Disk Storage

Form GA26-1592-4

- Your comments, accompanied by answers to the following questions, help us produce better publications for your use. If your answer to a question is "no" or requires qualification, please explain in the space provided. All comments and suggestions become the property of IBM.

	Yes	No	Explain
• Does this publication meet your needs?	<input type="checkbox"/>	<input type="checkbox"/>	_____
• Did you find the material:			_____
Easy to read and understand?	<input type="checkbox"/>	<input type="checkbox"/>	_____
Organized for convenient use?	<input type="checkbox"/>	<input type="checkbox"/>	_____
Complete?	<input type="checkbox"/>	<input type="checkbox"/>	_____
Well illustrated?	<input type="checkbox"/>	<input type="checkbox"/>	_____
Well indexed?	<input type="checkbox"/>	<input type="checkbox"/>	_____
Written for your technical level?	<input type="checkbox"/>	<input type="checkbox"/>	_____

- What is your occupation? _____
- How do you use this publication?

	Yes	No		Yes	No
As an introduction to the subject?	<input type="checkbox"/>	<input type="checkbox"/>	As an instructor in a class?	<input type="checkbox"/>	<input type="checkbox"/>
For advanced knowledge of the subject?	<input type="checkbox"/>	<input type="checkbox"/>	As a student in a class?	<input type="checkbox"/>	<input type="checkbox"/>
For information about operating procedures?	<input type="checkbox"/>	<input type="checkbox"/>	As a reference manual?	<input type="checkbox"/>	<input type="checkbox"/>
Other _____					

- We would appreciate your other comments; please give specific page and line references where appropriate. If you wish a reply, be sure to include your name and address.

- Reply Requested Yes No Name _____
 Address _____
Zip **▶** _____

- Thank you for your cooperation. No postage stamp necessary if mailed in the U.S.A.

YOUR COMMENTS PLEASE . . .

Your answers to the questions on the back of this form, together with your comments, will help us produce better publications for your use. Each reply will be carefully reviewed by the persons responsible for writing and publishing this material. All comments and suggestions become the property of IBM.

Note: Please direct any request for copies of publications, or for assistance in using your IBM system, to your IBM representative or to the IBM sales office serving your locality.

fold

fold

BUSINESS REPLY MAIL
NO POSTAGE STAMP NECESSARY IF MAILED IN U. S. A.

FIRST CLASS
PERMIT NO. 2078
SAN JOSE, CALIF.



POSTAGE WILL BE PAID BY . . .

IBM Corporation
Monterey & Cottle Rds.
San Jose, California
95193

Attention: Product Publications, Dept. G24

fold

fold



International Business Machines Corporation
Data Processing Division
1133 Westchester Avenue, White Plains, New York 10604
(U.S.A. only)

IBM World Trade Corporation
821 United Nations Plaza, New York, New York 10017
(International)



International Business Machines Corporation
Data Processing Division
1133 Westchester Avenue, White Plains, New York 10604
(U.S.A. only)

IBM World Trade Corporation
821 United Nations Plaza, New York, New York 10017
(International)