

The IBM logo, consisting of the letters "IBM" in a bold, sans-serif font, is positioned inside a dark square.

Systems Reference Library

**IBM System/360 Component Descriptions –
2314 Direct Access Storage Facility and
2844 Auxiliary Storage Control**

The functional and operating characteristics of IBM 2314 model 1, A-series, and B-series Direct Access Storage Facility (DASF) and IBM 2844 Auxiliary Storage Control are described in this manual. These devices provide high speed, large capacity, direct access storage and control for IBM System/360 Models 30, 40, 50, 65, 67, 75, 85, 91, and 195 selector channels, or System/370 selector and block multiplexer channels.

This manual describes the instruction sets, command sequences, track formats, and operating procedures for the above devices.



PREFACE

The purpose of this manual is to provide reference material for programmers and operators of IBM 2314 Direct Access Storage Facilities. You should be familiar with IBM System/360 Principles of Operation, Order No. GA22-6821, and the system summary or functional characteristics manual applicable to the system to which the IBM 2314 is attached. Order numbers for system summary and functional characteristics manuals can be found in IBM System/360 Bibliography, Order No. GA22-6822.

Although there are three distinct versions of the IBM 2314 Direct Access Storage Facility (model 1, A-series, and B-series), the instruction sets, command sequences, and operating procedures described in this manual apply to all. All differences in performance,

special feature attachment, and component location are described in the appropriate sections of the manual.

Unless specifically noted otherwise, references to the 2314 encompass the facility as a whole. That is, the control unit and its attached disk storage modules.

The IBM 2844 Auxiliary Storage Control (not available with the 2314 B-series) performs essentially the same functions as the 2314 control unit. Therefore, the instruction set, command sequences, etc. described for the 2314 apply to the 2844. Operating procedures, addressing, and feature attachment peculiar to the 2844 are described in a separate section of this manual.

Seventh Edition (November, 1971)

This is a reprint of GA26-3599-5, incorporating changes released in Technical Newsletters GN26-0265 (dated January 15, 1971) and GN26-0274 (dated August 16, 1971).

Significant changes or additions to the specifications contained in this publication are continually being made. When using this publication in connection with the operation of IBM equipment, check the latest SRL Newsletter for revisions or contact the local IBM Branch Office.

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 Systems Development Division, Product Publications, Department G24, San Jose, California 95114.

IBM 2314 DIRECT ACCESS STORAGE FACILITY	1	END OF FILE	47
INTRODUCTION	1	MULTIPLE TRACK (M-T) OPERATION	48
IBM 2314 Direct Access Storage Facility, Model 1	1	RECORD OVERFLOW	48.1
IBM 2314 Direct Access Storage Facility, A-Series	1	Formatting Overflow Records	48.1
IBM 2314 Direct Access Storage Facility, B-Series	1	Processing Overflow Records	49
IBM 2844 Auxiliary Storage Control	1	CHANNEL PROGRAMS	50
FUNCTIONAL DESCRIPTION	2	TWO-CHANNEL SWITCH SPECIAL FEATURE	52
Comparison with the IBM 2311 Disk Storage Drive	2	Channel Selection Switch	52
IBM 2316 Disk Pack	3	Device Reservation	52
Access Mechanism and Disk Organization	3	Addressing with Two-Channel Switch	54
DATA FORMAT	3	Resets with Two-Channel Switch	54
Information Transfer	4	Untagged-Tagged Switch	54
TRACK FORMAT	5	Usage Meter with Two-Channel Switch	54
Index Marker	5	Power Control with Two-Channel Switch	54
Gaps	5	REMOTE SWITCH ATTACHMENT SPECIAL FEAT- TURE	54
Home Address Area	5	IBM 2844 AUXILIARY STORAGE CONTROL	55
Track Descriptor Record	7	Attachment and Addressing	55
Data Records	9	Program Control of 2314/2844	55
Data Checking	11	Status Presentation	56
RECORD CAPACITIES	12	Untagged-Tagged Toggle Switch	56
TIMING CONSIDERATIONS	12	Timing and Control	57
Access Times	12	Power Control	57
Rotational Delay	13	Configuration Control	57
IN-LINE MODE OF OPERATION	14	Usage Meter	57
RESETS	14	SUGGESTED ERROR RECOVERY PROCEDURES	58
2314 CHANNEL ATTACHMENT	14	OPERATOR INFORMATION	61
Enable/Disable Switch	14	DISK PACK HANDLING	61
Usage Meter	15	Disk Pack Labeling	61
INPUT/OUTPUT OPERATIONS	16	Disk Pack Cover	61
CHAINING	18	Disk Pack Storage	61
Disk I/O Operations	18	Error Free Disk Packs	61
TERMINATION OF I/O OPERATIONS	19	OPERATOR PROCEDURES	62
I/O INSTRUCTIONS	19	Disk Pack Loading	62
Addressing	21	Disk Pack Unloading	63
Channel Status Word	21	Changing a Module Address	63
CHANNEL COMMANDS	23	OPERATOR CONTROLS AND INDICATORS	63
TEST I/O	23	Module Identifier Plugs	63
Status Byte	23	Start/Stop Toggle Switches	63
SENSE COMMAND	23	Select Lock Indicators	63
Sense Bytes	25	Ready Indicator	64
BRANCHING IN 2314 CHANNEL PROGRAMS	27	Enable/Disable Switch	64
CONTROL COMMANDS	29	APPENDIX A. RECORD CAPACITIES	65
Seek Commands	31	APPENDIX B. COMMAND SUMMARY	66
SEARCH COMMANDS	32	APPENDIX C. TRACK ORIENTATION	67
Search Identifier	34	INDEX	68
Search Key	35		
Search Key and Data	36		
READ COMMANDS	39		
WRITE COMMANDS	43		
Formatting Write Commands	43		
Data Write Commands	47		




11441A

IBM 2314 Direct Access Storage Facility



11446B

IBM 2844 Auxiliary Storage Control with IBM 2314 Direct Access Storage Facility, Model 1 or A-Series



24487

IBM 2314 Direct Access Storage Facility, B-Series

INTRODUCTION

Direct access storage enables a record retrieval without sequential searching through a file. Thus, when only part of the records of a given run are active, a direct access storage device provides significantly better throughput than a device that handles each record sequentially. Direct access storage also enables random inquiries during record processing. If desired, a direct access device can also process records sequentially. Therefore, with a direct access storage device, records can be stored sequentially or randomly.

There are three versions of the IBM 2314 Direct Access Storage Facility (DASF).

IBM 2314 Direct Access Storage Facility - Model 1

The 2314 model 1 contains eight on-line disk storage modules plus one spare. Each module can store 29.17 million bytes or 58.35 million packed decimal digits on an IBM 2316 Disk Pack. The total on-line storage capacity is 233.4 million bytes. The average access time is 75 milliseconds.

IBM 2314 Direct Access Storage Facility - A-Series

The A-series contains a maximum of eight on-line disk storage modules and one spare. The available configurations of the A-series provide modular growth of the facility by attaching various combinations of IBM 2312, 2313, and 2318 disk storage units to an IBM 2314 Storage Control. Each module can store 29.17 million bytes, or 58.35 million packed decimal digits, providing a total on-line capacity of 233.4 million bytes.

Average access time of the 2314 A-series is 60 milliseconds.

IBM 2312 Disk Storage: The 2312 adds one disk storage module to the facility, providing 29.17 million bytes of storage.

IBM 2313 Disk Storage: The 2313 adds four disk storage modules to the facility, providing 116.7 million bytes of storage.

IBM 2318 Disk Storage: The 2318 adds two disk storage modules to the facility, providing 58.35 million bytes of storage.

IBM 2314 Direct Access Storage Facility - B-Series

The B-series includes the IBM 2314 Storage Control, Model B1 and the IBM 2319 Disk Storage, Models B1 and B2. Each 2319 contains three modules. The 2319-B1 attaches to a 2314-B1; the 2319-B2 attaches

to a 2319-B1 or another 2319-B2. The B-series facility can be configured with three, six, or nine (eight on-line and one spare) modules. Each module stores up to 29.17 million bytes, providing a total on-line capacity of 233.4 million bytes. The average access time of the B-series is 60 milliseconds.

The IBM 2844 Auxiliary Storage Control, available with the 2314 model 1 and the 2314 A-series, is not available with the B-series.

Other than the differences noted above, all versions of the IBM 2314 are functionally the same. Therefore, the instruction sets, command sequences, metering, etc. described in this manual apply to all versions of 2314.

The 2314 model 1 and the B-series with nine modules retain one disk storage module as a service module. (The service module is optional on 2314 A-series.) The service module, which can be any one of the attached modules, is available for immediate use if servicing or routine maintenance is required on any one of the other modules. A switch allows the customer engineer to maintain the service module with in-line diagnostics while the customer program operates the remaining on-line modules. (See "In-Line Mode of Operation" in this manual.)

The IBM 2316 Disk Pack, used on the 2312, 2313, 2318, and 2319, can be quickly removed from or mounted in any of the module locations. Interchangeability of disk packs provides a high degree of program flexibility and virtually unlimited off-line storage capacity.

The large capacity, high speed, direct access storage characterized by the 2314 is ideally suited for real-time data processing and teleprocessing applications. The IBM 2314 is designed for attachment to IBM System/360 Models 30, 40, 50, 65, 67, 75, 85, 91, and 195 selector channels, or System/370 selector and block multiplexer channels. As in other System/360 equipment, the 2314 circuitry is packaged using solid logic technology (SLT).

IBM 2844 Auxiliary Storage Control

An IBM 2844 Auxiliary Storage Control can be ordered with the 2314 model 1 or an A-series with 8 on-line modules and a spare. The IBM 2844 is a second control unit (in addition to the basic 2314 control) for the disk modules in the 2314. It attaches to the 2314 facility as an integral unit. Any on-line 2314 disk module can be controlled through either the 2844 or the basic 2314 control. Switching of any module to operate with either control unit is effected by programming.

This second control unit, when attached to two System/360 Model 30's or 40's or to a Model 50 and

up (where a second channel is available), potentially doubles 2314 throughput and significantly increases system throughput. A 2314/2844 complex is ideally suited for very high activity applications and for real-time and teleprocessing applications that require maximum disk storage availability.

The 2314/2844 complex provides for:

1. Simultaneous operation (such as for reading and writing operations) of any two 2314 on-line disk modules with two selector channels.
2. Systems availability to the 2314 modules if either control unit (i. e. , the basic 2314 or the 2844) should require preventive or unscheduled maintenance.

Descriptions of program control of the 2314 also apply to the 2844. See IBM 2844 Auxiliary Storage Control for additional considerations that apply to the 2314/2844 complex.

FUNCTIONAL DESCRIPTION

The IBM 2314 model 1 DASF consists of a control unit, eight on-line disk storage modules, and a service module. The 2314 A-series consists of a 2314 control unit, from one to eight on-line disk storage modules, and a service module (optional). An access mechanism for twenty read/write heads is provided at each module location. This arrangement permits separate access to each of the on-line IBM 2316 Disk Packs.

The control unit in the 2314 DASF:

- Interprets and executes commands sent from the CPU via the selector channel to which the 2314 is attached.
- Provides a path for and translates data transferred between the CPU and the disk packs.
- Furnishes, to the CPU, information (status and sense) concerning the operational conditions of the IBM 2314.
- Checks for accurate transfer of data.

Comparison with the IBM 2311 Disk Storage Drive

Characteristics of the IBM 2314 listed here are more fully described in other sections of this publication. If you are unfamiliar with 2311 disk storage operations, you may omit this section.

Information formats and control unit functions of the IBM 2314 are identical to those used by the IBM 2841 Storage Control Unit with attached IBM 2311 Disk Storage Drives. When compared with the 2841 with attached 2311's, the 2314:

- Has the same track format as the 2311.
- Has the same addressing structure as the 2311, except that the 2314 can store more information per track (7294 bytes) and has more tracks per cylinder (20).
- Uses the same error recovery procedures as the 2841/2311. (Sense bytes 0, 1, 2, and 5 are identical. Sense bytes 3 and 4 have special usage.)
- Uses file protection in the same way as the 2841/2311.
- Uses the same IBM programming system volume and file labels as the 2311.
- Provides the same defective surface flagging and related recovery functions as those used in the 2841/2311.
- Has three additional tracks per surface for use as alternate tracks should any of tracks 000 through 199 become defective. (As in the 2311, the alternate tracks are numbered 200, 201, and 202.)
- Allows the channel as much or more time for performance of command chaining functions than does the 2841/2311.
- Has the same disk rotation speed as the 2311 (25 milliseconds per revolution).
- Implements functions of the two-channel switch (a special feature) in the same way as the 2841/2311.
- Operates on all commands in the same way as the 2841/2311, except that addresses referenced by the control-*seek* commands are extended to include the additional 2314 read/write heads.
- Has a data rate of 312,000 bytes (624,000 packed decimal digits) per second. Note that this rate approximately doubles (over the 2841/2311) the data load impressed on the selector channel to which the 2314 is attached.
- Has the ability to have any on-line disk storage module switched with a service module for maintenance purposes. (Changing the module identifier plug places the affected module in the same physical and logical state as if a disk pack had been changed.)
- Has file scan and record overflow as standard features.

Programs written for the 2841/2311 will operate with the 2314. However, if file scan and record

overflow are not used in existing 2841/2311 programs, it will be desirable, in most cases, to include these functions because file scan and record overflow are standard features in the 2314. Also, because the 2314 has a greater track capacity and more read/write heads per cylinder than the 2311, files should be reorganized to take full advantage of 2314 storage capabilities.

IBM 2316 Disk Pack

The IBM 2316 Disk Pack (Figure 1) is a compact disk assembly, 15 inches in diameter (with cover), and weighs about 13 pounds. The disk pack is composed of 11 disks (each 14 inches in diameter) mounted one-half inch apart on a vertical shaft. Circular plates are located above the top disk and below the bottom disk to protect the assembly. The disks provide 20 surfaces on which data can be recorded. The entire assembly (disks, vertical shaft, and protective plates) rotates at a speed of 2,400 revolutions per minute in the 2314. (25 milliseconds is required to complete one revolution.)

Data written on a 2316 disk pack can be read by any 2314 drive.

The various storage categories and capacities of 2316's in a 2314 are shown in Figure 2. (See Record Capacities for further information.)

Care and handling procedures for 2316 disk packs are described in IBM Disk Pack and Cartridge Handling Procedures, Order No. GA26-5756.




Figure 1. IBM 2316 Disk Pack

Access Mechanism and Disk Organization

Information is written on and read from disk surfaces of the 2316 disk packs by read/write heads in the 2314. The 20 read/write heads at each disk-module location are positioned by a movable comb-like access mechanism. (Each disk module has its own access mechanism.) Two read/write heads are attached to each arm of an access mechanism. For example, read/write heads 0 and 1 are both attached to the uppermost access arm in a disk module. (The read/write heads are numbered 0 to 19, top to bottom.) When operating, the heads float over the disk surfaces on thin films of air.

Because all of the 20 read/write heads are always positioned in the same vertical plane, 20 tracks are available for reading or writing purposes without movement of the access mechanism. That is, when read/write head 0 is positioned at track 5 of its disk surface, then read/write heads 1 through 19 are also positioned at track 5 of their respective disk surfaces. Therefore, the positioning of the read/write heads at any track location provides access to a cylinder (or drum) of 20 tracks.

Figure 3 illustrates how the entire disk pack consists of 203 concentric cylinders of information. The numbering is from 000 (outermost cylinder) to 202 (innermost cylinder). Tracks in cylinders 200 through 202 are alternate tracks that can be used if any of tracks 000 through 199 should become defective.

The address of an individual track in a specific disk module location consists of the cylinder number and the read/write head number.

DATA FORMAT

The basic unit of information representation in System/360 is the byte. A byte is made up of eight information bits and, in many devices, a parity bit for checking purposes.

The eight bits of each byte can each be set on (value of 1) or off (value of 0) to represent any one of 256 combinations. Figure 4 illustrates the decimal values represented by the setting on of the various bits in a byte.

The bit pattern in a byte can represent:


1. A single alphabetic, special, or numeric character in zoned format. (Bits 0 through 3 of the low order byte in a zoned-format numeric field carry the sign of the field.)

	Per Track	Per Cylinder	Per Disk Storage Module (All Models)	Per IBM 2314 Model 1	Per IBM 2314 Model A-Series
Disk Storage Drives				8	1 to 8
Cylinders			200	1,600	200 to 1,600
Tracks		20	4,000	32,000	4,000 to 32,000
Bytes (Alphameric Characters)	7,294	145,880	29,176,000	233,408,000	29,176,000 to 233,408,000
Packed Decimal Digits (Numeric Only)	14,588	291,760	58,352,000	466,816,000	58,352,000 to 466,816,000

Note: All figures are based on one record per track.

11322D

Figure 2. IBM 2314 Storage Capacity



10342B

Figure 3. Cylinder Concept

- Two packed-decimal digits. (The sign of a packed-decimal field is carried in bits 4 through 7 of the low-order byte.)
- Part of a fixed-point or floating-point number.
- Special information (such as status), where the value of each bit in the byte specifies the presence or absence of a particular condition.

A group of related bytes is a field; a series of related fields is a record; and a series of similar records is a logical file (Figure 5). The application

Decimal Values of Byte Positions

128	64	32	16	8	4	2	1
0	1	2	3	4	5	6	7

Binary Representation of Decimal Number 19

0	0	0	1	0	0	1	1
0	1	2	3	4	5	6	7

11250B

Figure 4. Bit Positions of a Byte


determines which of the many possible ways records and logical files are to be organized. Further descriptions about System/360 information for has are in IBM System/360 Principles of Operation, Order No. GA22-6821.

Information Transfer

Information is transferred one byte at a time between the CPU and the control unit in the 2314. A ninth bit, the odd parity or check bit, is added when needed. If the bit structure of a byte is such that an even number of information bits are on (set to 1), then the parity bit is set on. If an odd number of information bits are on, then the parity bit is off. Therefore, an odd number of 1 bits is maintained for each byte sent to or from the CPU.

Nine bit lines are provided to transfer a byte between the channel and the control unit. Nine bits (each set to a value of either 0 or 1) are transferred simultaneously (in parallel) between the channel and the control unit. This transfer method is called parallel-by-byte.

Information is transferred between the 2314 control unit and a disk module one bit at a time



Records and Fields within a File 11251A

Figure 5. Records and Fields within a File

(in serial). This transfer method is called serial-by-bit. The parity bit is not transferred to the disk pack. Instead, cyclic and bit count checking are used. Cyclic and bit count check bytes are written with the data and checked when the data is read. (See Data Checking.)

The 2314 control unit converts information from parallel-by-byte to serial-by-bit or from serial-by-bit to parallel-by-byte to provide information movement between the selector channel of the System/360 and disk packs in the 2314 (Figure 6).

TRACK FORMAT

The same track format is used for all tracks on a 2316 Disk Pack. Each track has a home address, a track descriptor record (R0), and one or more data records (Figure 7).

Each record has a count area, a key area (use is optional), and a data area. A gap (G) separates each area.




Figure 6. Data Transfer Methods 11546 A

Index Marker

The beginning of a track is signalled when the index marker (index point) is detected. All tracks on a disk pack are synchronized by the same index marker. No index indication appears on individual records.

Many illustrations in this publication show an index marker at the beginning of a track. Note, however, that there is only one index marker per disk pack. The marker is automatically recognized by a special sensing device and its detection is not dependent upon any operation of the read/write head.

Gaps (G)

A gap (Figure 8) is a string of bytes written by the 2314 to delimit the recorded areas of a record. The gap contains no data and is used only by the control unit for orientation and synchronization between the control unit and drive. Each gap contains one byte that indicates to the control unit the beginning of the next recorded area. This byte contains different information for each type gap. It indicates by the information in it what type recorded area follows.

There are five types of gaps. G1, G2, G4, and G5 are fixed length. G3 varies with the combined length of the key and data areas.

The byte of information in the G3 gap is referred to as an address marker. The address marker is used by the control unit to locate the beginning of a logical record.


A G4 gap precedes home address area, a G5 gap precedes record zero count area, a G3 gap precedes R1-Rn count area, a G1 gap precedes R0-Rn key area, and a G2 gap precedes R0-Rn data areas.

Gap (G4)

This is a fixed length gap of 73 bytes. (778 bytes in the case where the home address area is defective.) It is generated by the 2314 control unit to separate the home address from the index. It contains one byte which indicates to the control unit that the next recorded area is home address (Figure 8).

Home Address Area (Nine Bytes)

There is one home address area on each track (Figure 9). It follows the index marker on the track and



* G3 gap contains an address marker
 ** Key area may not be present

20587

Figure 7. Track Format

is nine bytes in length. The first five bytes of information describe the condition of the track and the physical location of the track within the storage device. The last 4 bytes contain the error detection bytes for the home address area.

The purpose of the home address is to give each track a unique identity that is not affected by normal programming operations. Each track within a storage device can be located directly by cylinder number and head number. Normal programming operations may use the home address but do not affect its contents. Home addresses can only be transferred from the CPU to the storage device by a write home address command, and from the storage device to the CPU only by a read home address command. The writing of home addresses is usually accomplished by utility programs.

The cylinder and head numbers of the home address must be identical with the last three bytes of the seek address for that track to insure against spurious seek check indications. Portions of these bytes are checked against information in the control unit during some operations, to help insure against faulty seeking or head advancing.


Home Address Flag Byte (One Byte)

The flag byte in the home address indicates track condition. Normally all eight bits of the flag are

set to zero when the home address is first written. The flag byte in the home address is transferred automatically to the using system only by performing a read home address operation on the track. The bit significance is:

Flag Byte Bit	Function or Setting
0	Zero
1	Zero
2	Zero
3	Zero
4	Zero
5	Zero
6	Track Condition: 0 indicates operative track. 1 indicates defective track.
7	Track Use: 0 indicates primary track. 1 indicates alternate track.

The home address flag byte, bits 1 - 7, is propagated automatically by the 2314 control unit into the flag byte of each record on the track after the home address, (except in cases where bit 1 is set to a 1 in formatting overflow record segments). Bits 2 - 7 of all flag bytes on the track are identical.



* G3 gap contains an address marker
 ** Key area may not be present

20588

Figure 8. Gaps

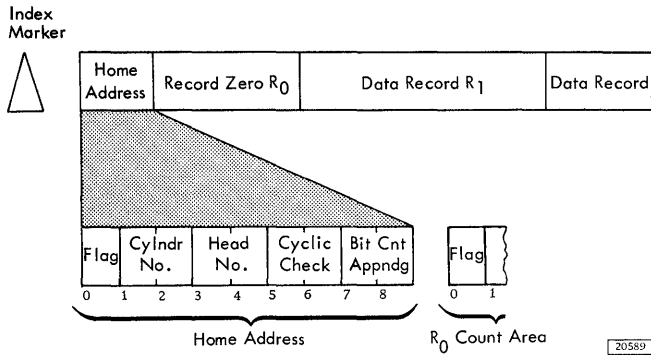


Figure 9. Home Address

The only flag byte on any track that is ever transferred to or from the channel is the one in the home address area. This flag byte is transferred only by using the write home address or read home address commands. When a track is found to be defective, the normal procedure is to assign an alternate track.

The address of the alternate track is written by the programmer in the track descriptor record (R0) of the defective track.

If a search or read operation (other than HA or R0) is performed on the defective track, a check condition is signalled to the channel. The program can then read the track descriptor record (R0) of the defective track into main storage. The address of the alternate track is obtained from this record and a seek to the alternate track can be performed. For further information, see Suggested Error Recovery Procedures - byte 0, bit 6, Track Condition Check.

Cylinder Number (Two Bytes)

The group of tracks available to the read/write heads at each access mechanism position is called a cylinder. The cylinder number (two bytes) identifies the cylinder within which data is stored. All bits in byte 1 of the home address (high order byte of cylinder number) must be set to zero. Byte 2 is set to one of the values in the range of 0 to 202.

Read/Write Head Number (Two Bytes)

The read/write head number (two bytes) specifies a particular read/write head within the selected cylinder. All bits in byte 3 of the home address (high order byte of head number) must be set to zero. Byte 4 is set to one of the values in the range 0 to 19.

The combination of cylinder and read/write head numbers is used to locate a specific track.

Note that no disk module address is used in any record. The module address is derived from the program instruction used to initiate the I/O operation. (See Input/Output Operations.)

Data Check (Four Bytes)

Four bytes of data check information are appended to the home address. These bytes are used for error detection. (See Data Checking.)

Gap (G5)

This is a fixed length gap of 43 bytes. It is generated by the 2314 control unit to separate the home address from the next recorded area. It contains one byte which indicates, to the control unit, that the next recorded area is record zero count area (Figure 8).

Track Descriptor Record (R0)

The first record (Figure 10) following the home address on each track is R0, the track descriptor record. Although R0 can be used to store data, this practice is discouraged, because R0 is designed to enable the entire content of a track to be moved to alternate tracks if a portion of the primary track becomes defective. A primary track is the original track on which data is stored. An alternate track contains data that has been repositioned from a defective primary track. This repositioning is independent of the file organization scheme used. R0 is not preceded by an address marker as are other data records on the track. Special read and write commands are provided that operate only on R0.




Figure 10. Track Descriptor Record (R0)

Count Area R0 (Thirteen Bytes)

The 13-byte count area is composed of the flag, cylinder number, read/write head number, record number, key length, data length, and data check bytes.

Flag (One Byte): Byte 0 of the count area is generated by the 2314 control unit as R0 is written. Bit significance is:

Flag Byte Bit	Function or Setting
0	Zero: For missing address marker detection. Always a zero in record zero and alternates one to zero in subsequent records.
1	Zero
2	Zero
3	Zero
4	Zero
5	Zero
6	Track Condition: 0 indicates operative track. 1 indicates defective track.
7	Track Use: 0 indicates primary track. 1 indicates alternate track.

The 2314 control unit causes bits 6 and 7 of flag bytes for all records written on a track to be set to the values of the corresponding bits in the home address flag byte.

Record Identifier (Five Bytes): The ID (record identifier) is composed of the cylinder number, read/write head number, and record number. In search ID operations, the ID of any record on a track (including R0) may be examined quickly enough to read the following key and data areas of that record without using a disk revolution.

Cylinder Number (Two Bytes): Bytes 1 and 2 of R0 contain the cylinder number of the track on which R0 is written. Byte 1 is set to zero; byte 2 is set to one of the values in the range 0 to 202. In IBM programming systems, if the track is defective, the cylinder number of the alternate track is placed in these two bytes of R0 of the defective track. The cylinder number of the defective track is placed in the cylinder number bytes of R0 of the alternate track. This arrangement provides a reference between the defective and alternate tracks.

Read/Write Head Number (Two Bytes): Bytes 3 and 4 contain the read/write head number for the disk surface on which the record is stored. Byte 3 is set to zero, and byte 4 can have a value in the range 0 to 19. In IBM programming systems, if

the track is defective, the head number of the alternate track is placed in these two bytes of R0 of the defective track. The head number of the defective track is placed in the head number bytes of R0 of the alternate track.

Record Number (One Byte): Byte 5 of the count area designates the sequential number of the record on the track. For R0, the record number is zero.

Key Length (One Byte): Byte 6 specifies the number of bytes in the key area of the record (excluding data check bytes). If the record has no key, byte 6 is set to all zeros. Byte 6 can indicate a key length from 0 to 225. Because of its intended use with alternate track procedures, R0 normally has no key area.

Data Length (Two Bytes): Bytes 7 and 8 specify the number of bytes in the data area of R0 (excluding data check bytes).

A data length of zero indicates the end of a logical file. The 2314 control unit sends special indications to the CPU when an end-of-file record is read or written.

In IBM programming systems, the data length in R0 is set to eight. Hence, end-of-file does not occur when R0 is read.

Data Check (Four Bytes): Bytes 9, 10, 11, and 12 are used for error detection as described in Data Checking.

Gap (G1)

This is a fixed length gap of 41 bytes, generated by the 2314 control unit, to separate recorded areas. It contains one byte which indicates to the control unit, that the next recorded area is a key area. This is the same G1 gap that precedes all R0-Rn key areas (Figure 8).

Key Area (R0 - 1 to 255 Bytes)

The length of the key area is recorded in the key-length byte of the R0 count area. Standard use of R0 by IBM programming systems does not include a key area.

A more detailed description of key area can be found in the Key Area subsection of the Data Records (R1-Rn) section in this manual.

Gap (G2)

This is a fixed length gap of 41 bytes, generated by the 2314 control unit, to separate recorded areas. It contains one byte which indicates, to the control unit, that the next recorded area is a data area. This is the same G2 gap that precedes all R0-Rn data areas (Figure 8).

Data Area (R0)

IBM programming systems use this area of R0 to record the identifier of the last record on the track and the number of bytes still available for writing new records. In IBM programming systems, the data area of R0 contains eight bytes (plus four data check bytes).

If the data length is zero, no data is transferred to the channel from the data area when this record is read, but the unit exception status is set to indicate end of file.

Data Records (R1-Rn)

One or more data records (Figure 11) can follow R0 on a track. Each record has a count area that is used to designate the track address, record number, and length of the key and data areas of the record. Count areas make each record self-formatting. Because of this identification method, file organization can make use of variable length records.

Gap (G3) Address Marker

This gap is a variable length gap (see below) generated by the 2314 control unit to separate recorded areas (Figure 8). It contains an address marker which indicates, to the control unit, that the next recorded area is a R1-Rn count area. Also, this address marker indicates to the control unit that this is the beginning of any logical record (count, key, and data or count and data, except R0).

$$G3 = ((K_L + D_L) \div 23 \frac{65}{256})^* + 43$$

Gap G3 can be computed approximately from:

$$G3 \text{ (Bytes)} = 43 + .043 (K_L + D_L)$$

*Truncate fraction

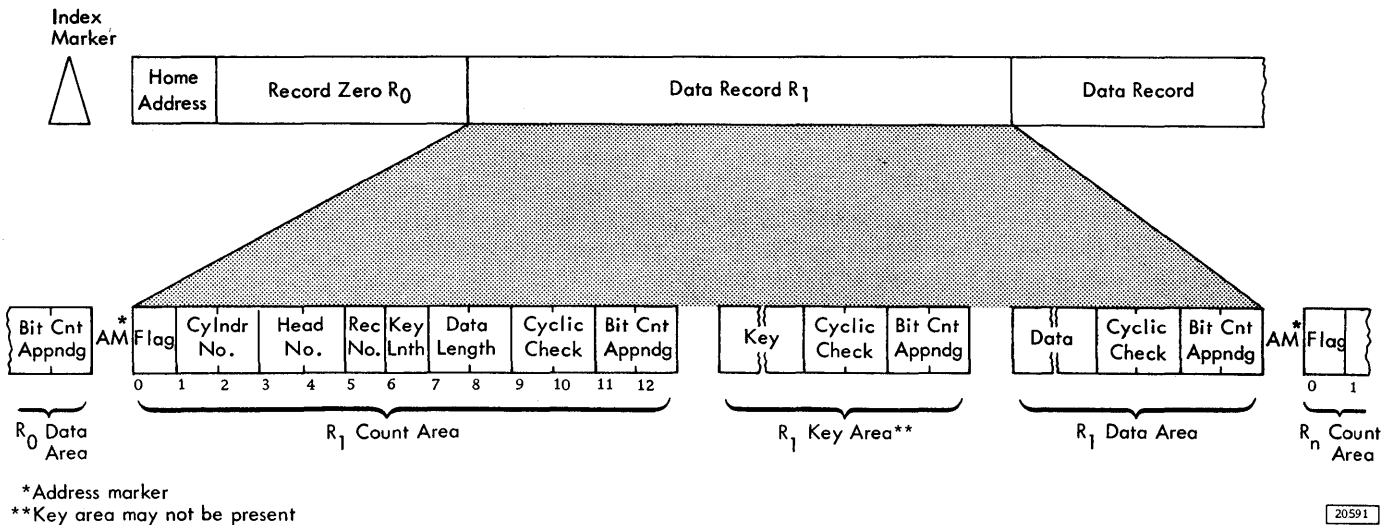


Figure 11. Data Records

Count Area (R1-Rn Thirteen Bytes)

This 13 byte area (Figure 11) is composed of the flag, cylinder number, read/write head number, record number, key length, data length, and four data check bytes. Bytes 1 through 8 (i.e., not the flag or data check bytes) are developed in CPU main storage by the program used to write the record. As in R0, the ID (identifier) of a record is the cylinder number, read/write head number, and record number bytes of the count area.

Flag (One Byte): Byte 0 of the count area is generated by the 2314 control unit as each record is written on a track. This byte is not sent from the CPU. Bit significance in the flag byte is:

Flag Byte Bit	Function or Setting
0	0 for even records (i.e., R0, R2, R4, etc.). 1 for odd records (i.e., R1, R3, R5, etc.). This bit is used by the 2314 control unit when reading adjacent records (command chaining) to ensure that no single record is missed. The 2314 signals a missing address marker (sense information) when two consecutive, identical bits are encountered (unless an index point intervenes).

Flag Byte Bit	Function or Setting
1	Used for record-overflow operations: 0 for all non-overflow records and for the last segment of an overflow record. 1 for each segment (except the last) of an overflow record.
2	Zero
3	Zero
4	Zero
5	Zero
6	Track Condition: 0 indicates operative track. 1 indicates defective track.
7	Track Use: 0 indicates primary track. 1 indicates alternate track.

Bits 6 and 7 are set to the values of the corresponding bits in the home address flag byte as each record is written on the track. The 2314 control unit performs this operation.

Cylinder Number (Two Bytes): Bytes 1 and 2 contain the cylinder number in which the record is stored. Byte 1 is set to zero, and byte 2 can have a value in the range of 0 to 202.

Read/Write Head Number (Two Bytes): Bytes 3 and 4 contain the read/write head number for the disk

surface on which the record is stored. Byte 3 is set to zero, and byte 4 can have a value in the range of 0 to 19.

Record Number (One Byte): Byte 5 designates the sequential number of the record on the track. The values represented by this byte can be from 1 to 255. This is used as a record identifier and its application is up to the programmer. (R0 has a record number of zero.)

Key Length (One Byte): Byte 6 specifies the number of bytes in the key area of the record (excluding data check bytes). If the record has no key, the key length is zero. The key area length values represented by this byte can be from 0 to 255.

Data Length (Two Bytes): Bytes 7 and 8 specify the number of bytes (excluding data check bytes) in the data area of the record. The data length represented by these two bytes can be from one to track capacity. (See Record Capacities. For formatting records that exceed the capacity of a track, refer to Record Overflow.)

Zero data length indicates the end of a logical file. Unit exception status is sent to the CPU when an end-of-file record is read or written.

Data Check (Four Bytes): Bytes 9, 10, 11, and 12 are used for data checking information in the count area. (See Data Checking.)

Gap (G1)

This is a fixed length gap of 41 bytes, generated by the 2314 control unit to separate recorded areas. It contains one byte which indicates to the control unit, that the next recorded area is a key area. This is the same G1 gap that precedes all R0-Rn key areas (Figure 8).

Key Area (R1-Rn 1 to 255 Bytes)

The key area can contain record identifying information such as serial number, social security number, or policy number.

Commands are provided to search key areas for this identifying information. When the desired key is found, a read or write command can be issued and the data area read or written during the same revolution that read the key. The data in the key area may be repeated in the data area.

For a search key operation, key information is sent from the CPU to the 2314 control unit. A comparison is then made between the key read from the

track and the key in the control unit. If the comparison is satisfied, the operation is completed. (See Search Commands.)

Key area length ranges from 1 to 255 bytes. Four data check bytes are added to the key area by the 2314 control unit. If the key length (in the count area of the record) is zero, a key area is not written for that record.

Gap (G2)

This is a fixed length gap of 41 bytes, generated by the 2314 control unit to separate recorded areas. It contains one byte which indicates, to the control unit, that the next recorded area is a data area. This is the same G2 gap that precedes all R0-Rn data areas (Figure 8).

Data Area R1-Rn

This area contains the information identified by the count and key areas. Data information is organized and arranged by the programmer.

Four error check bytes are appended to the data area by the 2314 control unit. If the data length in the count field is zero, the 2314 writes four data check bytes and indicates end-of-file by unit-exception status.

Data Checking (Four Bytes)

Cyclic Checking (Two Bytes)

In the 2314, data is stored into and retrieved from areas that contain one or more fields. Capacity is more effectively used by associating check bits (in cyclic check bytes), with each area, rather than with each byte (Figure 11).

When data is transferred from the CPU to a disk pack, the 2314 control unit removes the parity bit from each byte. The 2314 control unit then computes two cyclic check bytes that are placed at the end of each area. The cyclic check bytes are arithmetically coded from the information to be placed in the associated area. Subsequently, when the area is read from the disk pack, the 2314 control unit recomputes the check information and compares it to the check bytes read from the associated area. Because the cyclic check bytes are at the end of each area, the 2314 cannot complete the check until the entire area is read.

Bit Count Appendage (BCA) (Two Bytes)

Bit count appendage (BCA) consists of one indicator byte and one bit count byte (Figure 12). The indicator byte is used to indicate the presence of BCA. The

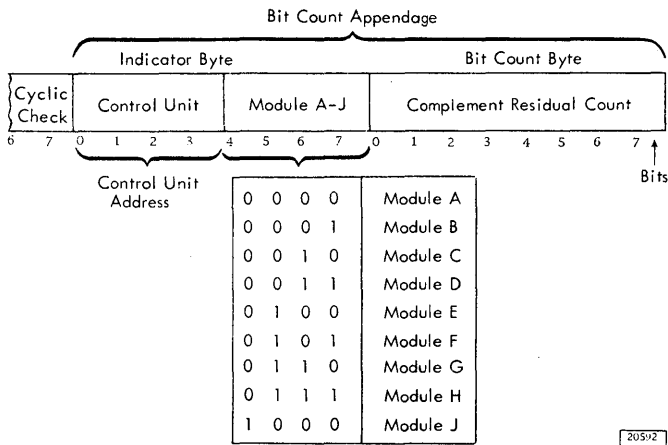


Figure 12. Bit Count Appendage (BCA)

bit count byte consists of a residual count of the total number of data bits in the field. During a write operation, the bit count is accumulated and appended with the indicator byte to the cyclic code bytes. During a read operation, the bit count is accumulated and compared with the bit count byte recorded. Data check is indicated if either the bit count does not compare or the cyclic code check registers are non-zero.

Indicator Byte: The indicator byte, which is located between the bit count byte and the cyclic code bytes, is provided for the purpose of maintaining compatibility between disk packs without BCA and disk packs with BCA. (Records were formerly formatted without BCA.) No comparison of the bit count is accomplished if the indicator byte is not present. The indicator byte consists of the following format. Bits 0, 1, 2, 3 define the control unit address. Bits 4, 5, 6, 7 define the physical drive address as specified in bits 4 through 7 of sense byte 4 (refer to Figure 12). The control unit and physical drive identification specify the physical drive on which

the field was written. If BCA is not present (old format), the byte following the cyclic codes bytes contains a hexadecimal "CC." In no case will the indicator byte of BCA contain a hexadecimal "CC."

Bit Count Byte: The bit count byte contains the complement of the residual count of the total number of data bits in the field to which it is appended. Included in the bit count are the bits in the byte preceding the data (part of gap) and the bits in the first cyclic code check byte.

RECORD CAPACITIES

Because each record has certain non-data areas (such as count fields and gaps) the net data storage capacity of tracks varies with the number of records (Figure 13). With one record per track, each track has a data capacity of 7294 bytes when the key length is zero and R0 is used as specified by IBM programming systems.

Because of the high-density recording method used by the 2314, minute contamination particles can affect data reading and writing and may cause loss of bits. Therefore, 203 tracks per disk surface are provided to ensure that the stated capacity of 29.17 million bytes (based on 200 tracks) is maintained for the life of the disk pack (Figure 2).


Record capacities of the 2314 are listed for records with keys and without keys in Figure 14. Appendix A lists record capacities and transmission time for 1 to 20 records per track.

TIMING CONSIDERATIONS

The data transfer rate to or from the 2314 is 312,000 bytes per second through either the basic 2314 control or the 2844.

Access Times

Seek operations for individual 2314 disk modules can be overlapped, thereby allowing any combination of



* G3 gap contains an address marker

** Key area may not be present

20593

Figure 13. Track Capacity

Track Capacity in Bytes, When R ₀ is Used as Specified By IBM Programming Systems.	Track Capacity in Bytes When R ₀ is Used for Data	Bytes Per Data Record			
		Data Records (except for last record)		Last Record	
		Without Key	With Key	Without Key	With Key
7294	7403	$101 + \frac{534}{512} (DL)^*$	$146 + \frac{534}{512} (K_L + D_L)^*$	D _L	$45 + K_L + D_L$

Record R ₀ used as specified by IBM Programming Systems. No application data; K _L = 0; D _L = 8	Number of Equal Length Records that can be stored on a Track																			
	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
Bytes per record without key	7294	3521	2298	1693	1332	1092	921	793	694	615	550	496	450	411	377	347	321	298	276	258
Bytes per record with key	7249	3476	2254	1649	1288	1049	878	750	650	571	506	452	407	368	333	304	277	254	233	215

* Carry to four places past decimal when calculating.

30293A

Figure 14. Record Capacities

the on-line access mechanisms to be in motion simultaneously. Figure 15 shows the approximate access times in either the forward or reverse direction for the number of cylinders that the access mechanism travels.

Note that these access times take into account the movement of the access mechanism only. During an application run, the effective time for a seek may be greater than the access times listed here. The reason for this is that the selector channel with which the 2314 is operating may be masked against interruptions at the time the seek is completed. The channel cannot recognize the interruption associated with completion of the seek (i.e., the 2314 presents device end to the channel) until the masking condition is removed.

Also, the selector channel may be engaged in an operation with another I/O unit at the time the seek is finished. Ending of the seek operation is then not recognized by the channel until the operation that the channel is performing with the other I/O unit is completed. For descriptions of the masking of channels and for other details of channel operations, refer to IBM System/360 Principles of Operation, Order No. GA22-6821

Rotational Delay

After the read/write heads are positioned at a cylinder, a search operation to find a specific record can be performed. The search is started on a track read by a particular read/write head. One revolution of the disk surface takes 25 ms. (The disk surfaces make 2,400 revolutions per minute.) Therefore, when the search is started, the read/

write head may be at the record to be found, or it may be 25 ms away from the record. Hence, the rotational delay can vary from about 0 ms to 25 ms. Considering a large number of different search operations, the average rotational delay to find a record on a particular track is 12.5 ms (i.e., one-half of the revolution time).

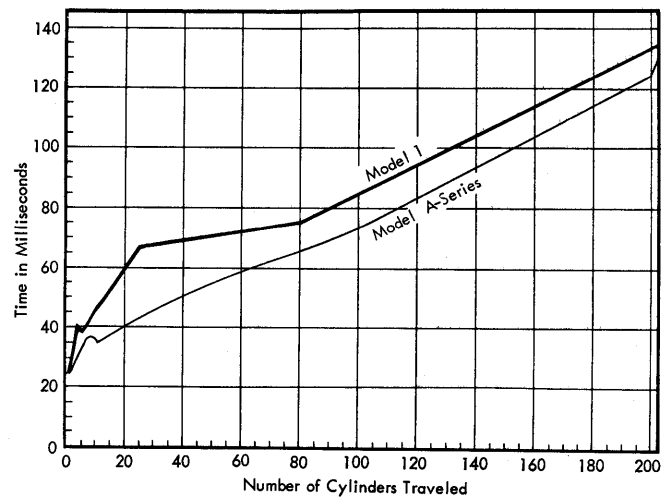


Figure 15. IBM 2314 Access Time

25425C

A multi-track search through a cylinder takes place at the rate of 25 milliseconds per track. No rotational time is lost while switching heads between tracks.

2. A system reset or malfunction reset causes any file mask set in the 2314 control unit to reset to all 0's.
3. If the 2314 is reserved for a channel and a system reset occurs, the reservation condition is removed. (See Two-Channel Switch Special Feature.)
4. 2314 status information is reset when a system reset occurs.

IN-LINE MODE OF OPERATION

To reduce facility down time, the 2314 can run diagnostic tests on a failing disk module concurrent with normal operation of the on-line modules. This mode of control unit operation, termed "in-line", is effective only on the service module. A failing on-line module must have its module plug replaced by the service module plug (thereby making the failing module the service module) before the customer engineer can service it with the in-line diagnostics. During in-line mode, the 2314 control unit is time-shared between diagnostic and customer operations.

The diagnostic state is entered at the completion of a CCW chain. During this state, the 2314 responds with control unit busy to any start I/O or test I/O. At the completion of a single diagnostic pass, the 2314 determines if control unit busy status was presented. If so, the control unit will enter a 200 millisecond wait loop and attempt to present control unit end status. If control unit busy did not occur, the 2314 will determine if any seeks have completed at the on-line modules. If so, the 200 millisecond loop is entered in an attempt to present device end status. If neither control unit end nor device end status is outstanding, the 2314 will re-enter the diagnostic state.

Upon receiving control unit end status, the control program should re-execute the start I/O at which the control unit busy sequence occurred before executing any other start I/O against the same module. This will ensure against the possibility of the 2314 in-line mode synchronizing with the I/O control program thus conflicting with customer throughput.

RESETS

System reset and malfunction reset apply to the 2314 as generally described in IBM System/360 Principles of Operation Order No. GA22-6821.

However, the following notes apply specifically to the 2314:

1. The 2314 control unit appears busy for up to 500 microseconds following a system reset.

2314 CHANNEL ATTACHMENT

Because of the high data rate of the 2314 and the chaining timing requirements of any 2314 program, the 2314 is attached to a selector channel, not to a multiplexer channel.

The 2314 should be attached only to selector channel 1 on a System/360 Model 30 and either selector channel 1 or selector channel 2, but not both on a System/360 Model 40. If the 2314, used with a model 30 or 40, has 2844 Auxiliary Storage Control, the other channel of the model 30 or 40 must not be attached to the 2844 channel interface. The 2314/2844 may attach to both selector channels on a system/360 Model 40 by use of the two-channel switch special feature. (See 2844 Auxiliary Storage Control.)

Enable/Disable Switch

The enable/disable switch (located on the single drive module on the 2314 model 1, and on the control unit of the 2314 A-series) enables or disables communications between the 2314 and the channel. It also enables or disables the usage meter.

An additional switch is provided if the 2844 Auxiliary Storage Control is attached. If the two-channel switch special feature is installed in either (or both) the 2314 or the 2844, each channel has its own enable/disable switch.

Regardless of the number of switches, all processing time is recorded on a single usage meter.

To enable communications between a channel and the 2314/2844 the following conditions must exist:

1. The enable/disable switch associated with the channel being enabled must be in the enable position.
2. The CPU must be in a stopped or wait state.

To disable communications between a channel and the 2314/2844 the following conditions must exist:

1. The enable/disable switch associated with the channel to be disabled must be in the disable position.
2. The channel selection switch in the 2314/2844 must not be selected to the channel being disabled.
3. The CPU must be in a stopped or wait state.

If the CPU enters the stopped or wait state when the 2314 contains outstanding status, the status may be made unavailable if the enable/disable switch

associated with the channel has previously been set to disable.

Usage Meter

A single usage meter is provided for the 2314/2844 (located on the single drive module of the 2314 model 1 or on the control unit of the 2314 A-series). The usage meter operates whenever the 2314/2844 is enabled and operating with any of the attached channels. See Enable/Disable switch.

INPUT/OUTPUT OPERATIONS

A general description of input/output (I/O) operations related to the 2314 is presented here. For detailed information regarding CPU and channel program control of I/O operations, refer to IBM System/360 Principles of Operation, Order No. GA22-6821.


Input/output (I/O) operations are initiated by I/O instructions in the CPU program and subsequently controlled by commands fetched from main storage by a channel. Problem program arithmetical and logical decision operations are performed while the processing unit is in the problem state. For an input/output operation to be started, the processing unit must be placed in the supervisor state. The transition from problem state to supervisor state is normally effected when a supervisor call instruction is encountered in the program, or when an I/O interruption causes a new PSW (Figure 16) to be loaded and causes the old one to be stored.

Some way must be provided to return to the right place in the program to resume operations in the problem state as soon as the functions performed in the supervisor state are completed. The problem state is restored by loading a program status word (PSW) specifying the problem state. If the supervisor state was entered because of an I/O interruption, the program normally reloads the PSW stored at the time of the I/O interruption.

The operating program initiates a supervisor call (SVC) and the CPU supervisor program initiates I/O operations with a start I/O instruction. Bit positions 16-31 of this instruction identify the channel and I/O device to be used in the operation. Start I/O causes the channel to fetch the channel address word (CAW) from main storage location 72 (Figure 17). The command address portion of the CAW designates the location in main storage from which the channel subsequently fetches the first channel command word (CCW) (Figure 18). The CCW specifies the command to be executed and the storage area to be used.

If the channel is available, the channel attempts to select the device by sending the address of the device to all attached control units. The control unit specified in the address responds to its selection and awaits further instructions. The command code is sent to the selected control unit; the control unit then responds with an initial status byte to the channel.


The start I/O instruction is now finished, releasing the CPU to perform the next instruction. The results of the attempt to initiate the execution of the



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

30068

Figure 16. Program Status Word



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 17. Channel Address Word

command are indicated by the condition code in the program status word (PSW). If the I/O operation was not started, new status information to show why is normally set in the channel status word (CSW).


All data transfers from the channel to the control unit are checked for correct parity. If a parity error is detected, a unit check signal is sent to the channel by the control unit. If a parity error is detected in the command bytes, the command is not executed.

When command chaining is specified, a series of commands is executed by the device as the result of a single start I/O instruction.

The end of an I/O device operation normally is indicated to the program by an I/O interrupt. If the

channel is masked, the I/O interrupt remains pending. When the interrupt occurs, the CSW is stored with the status bits indicating channel end and device end. The channel end condition indicates that the I/O device has received or provided all information associated with the operation and no longer needs channel facilities. The device end condition indicates that the I/O device has finished the operation. The device end condition can occur concurrently with the channel end, or later.

If command chaining has been specified in the CCW, the next CCW is fetched by the channel and the operation designated is commenced. Unusual conditions and errors terminate the chain, cause an I/O interrupt, and cause the CSW to be stored.



Fields in the CCW are allocated for the following purposes:

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 control unit.
8-31	Data Address	Specify the location of an 8-bit byte in main storage. This is the address of the area designated by the CCW.
32	Chain Data (CD) Flag	When set to one, specifies chaining of data. Make sure the data rate of the I/O device permits chaining by the particular System/360 model before using.
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.

CCW Bit Position	Field Designation	Function
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 No-op and Recalibrate commands.
35	Skip Flag	When set to one, specifies suppression of a transfer of information to storage during a Read or Sense operation. Cyclic checking takes place as though the information had been placed in storage. When bit 35 is zero, normal transfer of data takes place.
36	Program-Control-Interruption (PCI) Flag	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.
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. For additional information, see Control Command-Transfer-in-Channel.
40-47		Not used.
48-63	Count	Specify the number of 8-bit byte locations in the storage area designated by the CCW. Must be nonzero except for TIC command.

25524

Figure 18. Channel Command Word

CHAINING

Obtaining a new CCW upon completion of the operation for the current CCW without the issuance of a new start I/O by the CPU program is called chaining. The CCW's belonging to such a sequence are said to be chained. There are two types of chaining, data and command.

In command chaining, each CCW used specifies a new I/O operation. It is practically mandatory to perform 2314 read or write operations by chaining the commands. One example is the search operation followed by a read or write.

A search command searches only the next record that happens to be passing the R/W head. In order to search all records on a track, the search command must be repeated as many times as there are records. This is done by chaining the search to a Transfer In Channel (TIC) Command. This TIC causes a branch back to the CCW location containing the search command. The search is thereby repeated until it is satisfied. A satisfied search causes the channel program to skip the next sequential CCW (in this case, the TIC) and execute the read or write CCW just beyond.

Time is made available by the 2314 to execute command chaining functions in the gaps between record areas. Certain restrictions exist, however, with regard to sequences of commands. (For further information, see the descriptions, in this manual, of the commands used with the 2314.)

Disk I/O Operations (General Information)

A variety of commands and command sequences (chains) can be used with the 2314 (Figure 19). This section is a brief, general description of operations basic to most 2314 channel programs. Detailed information regarding use of specific commands and handling of unusual conditions is presented in other sections of this manual.

Disk initialization, consisting of writing home address and record R0, must be performed before a disk pack can be used for application data. If any defective tracks are found, they must be flagged, and alternate tracks must be assigned. Also, if the disk pack is to contain an IPL (initial program load) record, then that record must be placed on the pack at track zero in cylinder zero. Utility programs, provided by IBM, can be used to perform the initializing operations. (Refer to IBM System/360 Bibliography Order No. GA22-6822, for a listing of publications that describe IBM programming systems available for System/360.)

Generally, most 2314 operations (whether for initialization or for application runs) include the following two steps:


1. First, a set file mask command is issued. This command transfers a file mask, previously placed in main storage by the program, to the 2314 control unit. The file mask (stored in the 2314 until the end of the chain) is used to prevent execution of commands that should not operate on the file area accessed. For example, the file mask can be set up to prevent execution of all write commands. This procedure could be used to prevent altering a master file that contains constant data. If any write command were inadvertently issued in the chain, the 2314 would be notified of that fact.
2. Next, a seek operation is performed to access the desired disk pack area. A specific track and cylinder can be accessed through use of the seek or seek cylinder command (in the 2314 these two commands perform the same function). A specific read/write head (within the currently selected cylinder) can be accessed by the seek-head command.

An operation that does not require a search (such as a write home address command) can now be executed as long as it is permitted by the file mask. If, however, a specific record must be found, some type of search must be performed first.

If an equal comparison is not obtained for any ID on the track, the search can be programmed to continue to the next track. (This operation can be done through a new seek to the next track or through use of the multiple track mode. See Multiple Track Operation for further details.) A no-record-found sense indication is available in the 2314 control unit if the ID cannot be found on the track (non-multiple track mode operation). The unit check status bit is set on in the CSW (channel status word) when a no-record-found condition occurs. Subsequent inspection of the sense information (placed in main storage via a sense command) would indicate that the no-record-found condition exists.

If the record is known to exist, the no-record-found condition indicates a malfunction; however, in many applications, the record may or may not be present. In this case, the no-record-found condition is normal.

For multiple-track and record-overflow operations, an end-of-cylinder indication is provided if the desired record is not located within a cylinder.



Command	0 M/T	1 HI	2 =	3 C	4 K	5 D	6	7
Write Data	0	0	0	0	0	1	0	1
Write Key, Data	0	0	0	0	1	1	0	1
Write Count, Key, Data	0	0	0	1	1	1	0	1
Write Home Address	0	0	0	1	1	0	0	1
Write Record Zero	0	0	0	1	0	1	0	1
Write Special CKD	0	0	0	0	0	0	0	1
Erase	0	0	0	1	0	0	0	1
Search Equal ID	X	0	1	1	0	0	0	1
Search Equal Key	X	0	1	0	1	0	0	1
Search Equal Home Address	X	0	1	1	1	0	0	1
Search Hi ID	X	1	0	1	0	0	0	1
Search Hi Key	X	1	0	0	1	0	0	1
Search Hi Equal ID	X	1	1	1	0	0	0	1
Search Hi Equal Key	X	1	1	0	1	0	0	1
Read Data	X	0	0	0	0	1	1	0
Read Key, Data	X	0	0	0	1	1	1	0
Read Count, Key, Data	X	0	0	1	1	1	1	0
Read Home Address	X	0	0	1	1	0	1	0
Read Record Zero	X	0	0	1	0	1	1	0
Read Count	X	0	0	1	0	0	1	0
Read Initial Program Load (IPL)	0	0	0	0	0	0	1	0
Control Seek (BBCCHH)	0	0	0	0	0	1	1	1
Control Recalibrate	0	0	0	1	0	0	1	1
Control Restore *	0	0	0	1	0	1	1	1
Control No Op	0	0	0	0	0	0	1	1
Control Cyl. Seek (CCHH)	0	0	0	0	1	0	1	1
Control Head Seek (HH)	0	0	0	1	1	0	1	1
Control Set File Mask	0	0	0	1	1	1	1	1
Test I/O (Not written by programmer)	0	0	0	0	0	0	0	0
Sense I/O	0	0	0	0	0	1	0	0

X: On Search and Read commands B0 can be either "0" or "1". If "0", head switching will not take place when Index Point is detected. If "1", head switching will take place when Index Point is detected (Multiple Track Mode).

* Executed as a No Op in 2314.

Figure 19. 2314 Command Codes

TERMINATION OF I/O OPERATIONS

Termination of the input/output data transfer operation is signalled by channel end, which indicates that the channel is free to transfer information between other attached devices and the processing unit. Device end signals that the device is no longer busy with an operation. Also at the end of the operation, a CSW is stored in core storage to indicate the status of the input/output device.

In summary, to initialize an I/O operation requires an instruction, such as start I/O, which in turn causes the channel to fetch a CCW. The channel then takes over and controls the input/output operation as defined by the CCW. The CCW contains the necessary information for the completion of the operation. Upon completion, the input/output device initiates an interruption to inform the processing unit of the completion. The processing unit, in accepting the interruption, stores the CSW in a fixed storage location. This CSW contains the status of the input/output device that caused the interruption. The CSW is available in its fixed storage location until another CSW is later stored by another interruption or another I/O instruction.

I/O INSTRUCTIONS

The System/360 uses only four I/O instructions: start I/O, halt I/O, test I/O, and test channel. The format and bit designations concerning 2314 operations are shown in Figure 20.

Start I/O

All I/O operations are initiated by a start I/O instruction. If the channel facilities are free, start I/O is accepted and the CPU continues its program. The channel independently selects the I/O device specified by the instruction.

The CAW at main storage location 72 contains the protection key for the subchannel and the address of the first CCW. The CCW so designated specifies

the operation to be performed, the main-storage area to be used, and the action to be taken when the operation is completed.

If any unusual conditions exist, start I/O causes the status portion (bit positions 32-47) of the CSW at main storage location 64 to be replaced by a new set of status bits. The status bits pertain to the device addressed by the instruction. The contents of the other fields of the CSW are not changed.

Test I/O

The test I/O instruction sets the condition code in the PSW to indicate the state of the addressed channel, sub-channel, and device. When the status byte from the 2314 is not an all 0 byte, the condition code in the PSW is set to 1, and the CSW is stored at main storage location 64.

Test I/O clears any pending status but does not reset previous error conditions.


Test Channel

Execution of the test channel instruction sets the condition code in the PSW to indicate the state of the channel addressed by the instruction. It does not affect the 2314 Direct Access Storage Facility.

Halt I/O

When a halt I/O instruction is issued, the channel operation is terminated and the 2314 is disconnected from the channel. A halt I/O does not cause a command byte to be sent to the 2314.


The operation at the file continues to its normal stopping point and the control unit appears busy until the operation is complete. The ending status for the operation causes an interrupt when the operation is complete, unless the halt I/O is given during initial selection. If the 2314 operation in progress (when the halt I/O is executed) is the result of a write or erase command, the 2314 completes the command by inserting zeros to the end of the field or track.



I/O Instruction Format

Fields in the instruction are allocated as follows:

Bit Position	Field Designation	Function
0-7	Operation (OP) Code	Designates the operation to be performed.
8-15	Not Used	
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)	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 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.
8-15	Not Used	
16-20	Must be Zero	
21-23	Channel Address	000 - designates multiplexer channel. 001 - 110 - designates selector channel 1-6. 111 - invalid combination.
24	Shared Channel Indicator	For the 2314, a 1 indicates multiplex channel or sub-channel. On a selector channel, this bit is included in the control unit address.
25-27	Control Unit	0-7 control units per channel.
28	Must be Zero	
29-31	Module 0-7	

20620

Figure 20. I/O Instruction Format

Addressing

The address of the channel and I/O unit are contained in the start I/O instruction. The address is obtained from the sum of the contents of the general (base) register specified in the start I/O instruction, and the value of the displacement (in the start I/O instruction). Only the 11 low-order bits of the sum are used. The three high-order bits of this 11-bit address designate the channel to which the I/O unit is attached. The eight low-order bits identify the I/O unit. When the 2314 is addressed, the eight low-order bits specify the 2314 and the access mechanism (i. e., disk module location) in the 2314 as follows:

Bits of Eight-Bit Unit Address	Specify
0	(Always set to zero)
1 through 3	The 2314 control unit (0-7)
4	(Always set to zero)
5 through 7	The module (0-7)

|

Bits 1 through 3 can designate a 2314 control unit address in the range 000 through 111 (0 through 7 decimal). A maximum of eight adapter, also called control unit, connections can be made to any System/360 channel.

Bits 5 through 7 can be in the range 000 to 111 (0 to 7 decimal). This provides a unique address for each of the on-line disk modules in a 2314.

An attempt to address a disk module that is not attached to the 2314 (i. e., module 6 on a 2314 A-series facility with five modules) will result in a unit check (intervention required) condition. See "Sense Bytes."

The unit address of the 2314 is determined at installation, at which time it is internally wired into the 2314 control unit by an IBM Customer Engineer.

The address of any disk module can be changed manually. See Module Identifier Plugs in this manual.

Channel Status Word (CSW)

The channel status word informs the program of the status of an I/O device or the conditions under which an I/O operation has been terminated. The CSW is formed, or parts of it are replaced, during I/O inter-

ruptions and during execution of I/O instructions. The CSW is placed in main storage at location 64. It is available to the program at this location until the next I/O interruption occurs or until another I/O instruction generates a new CSW, whichever occurs first.

When the CSW is stored as a result of an I/O interruption, the I/O device is identified by the I/O address in the old PSW. The information placed in the CSW by an I/O instruction pertains to the device addressed by the instruction.

The CSW format is shown in Figure 21.




Figure 21. Channel Status Word

Fields in the CSW are allocated for the following purposes:

CSW Bit Position	Field Designation	Function
0-3	Protection Key	Form the storage protection key used in the chain of operations.
4-7	Not Used	Always zero.
8-31	Command Address	Form an address eight positions higher than the address of the last CCW used.
32	Attention	Not Used.
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 2314 is busy. This bit, in conjunction with the busy bit, signifies control unit busy.

<u>CSW Bit Position</u>	<u>Field Designation</u>	<u>Function</u>	<u>CSW Bit Position</u>	<u>Field Designation</u>	<u>Function</u>																		
34	Control Unit End	Set if a control unit busy status has been generated previously and the busy condition has been terminated.			A Sense I/O command may then be used to identify the condition.																		
35	Busy	Indicates that the selected device is busy. It is set when a new command chain is initiated while the selected access mechanism is still in motion due to a previous Seek command. In conjunction with the status modifier bit, indicates the control unit is busy. It is set when a new command chain is initiated while the 2314 is causing a track to be erased following a Format Write command. See Two Channel Switch Special Feature.	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, search ID, 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. If key length is not zero, the key area is transferred.																		
36	Channel End	Set at the end of each channel command.	40-47	Channel Status	Indicate channel conditions as follows:																		
37	Device End	Indicates that an access mechanism is free to be used.			<table border="1"> <thead> <tr> <th><u>Bit</u></th> <th><u>Designation</u></th> </tr> </thead> <tbody> <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> </tbody> </table>	<u>Bit</u>	<u>Designation</u>	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
<u>Bit</u>	<u>Designation</u>																						
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																						
38	Unit Check	Set whenever an unusual or error condition is detected in the 2314.	48-63	Count	The residual count from the last CCW used.																		

The information presented here relates specifically to 2314 operations. For further details about over-all I/O programming and the fields in channel command words, refer to IBM System/360 Principles of Operation, Order No. GA22-6821.

TEST I/O

The status byte from the 2314 control unit is stored in the CSW in CPU main storage when a test I/O instruction is accepted and processed by the 2314.

A command (operation code 00000000) is generated automatically by the channel when the channel requires status information from the 2314. It is generated as a result of processing a test I/O instruction, or it is initiated as an internal channel function when the channel requires status. (A command code of 00000000 should never be written into a channel program; it is not a valid command.)

Status Byte

The status byte contains information that reflects the operational condition of the 2314. The name of each bit and the corresponding bit positions of the CSW (channel status word) where the status appears are:

Status Byte Bit (from 2314)	CSW Bit Location (in CPU Main Storage)	Name
0	32	Not used
1	33	Status Modifier
2	34	Control Unit End
3	35	Busy
4	36	Channel End
5	37	Device End
6	38	Unit Check
7	39	Unit Exception

Status stored in the CSW remains unchanged until a subsequent I/O interruption occurs or an I/O instruction is processed. See IBM System/360 Principles of Operation, Order No. GA22-6821.

SENSE COMMAND

The sense command (Figure 22) should be sent to the 2314 whenever unit check status is detected, even if the sense information will not be used. The 2314 sends one to six sense bytes as specified by the count (bits 48 through 63) in the sense command.

The sense bytes (Figure 23) are placed in main storage starting at the data address specified in the sense command. That is, sense byte 0 is placed in the first address byte specified, sense byte 1 in the next higher sequential byte, etc.

Sense I/O

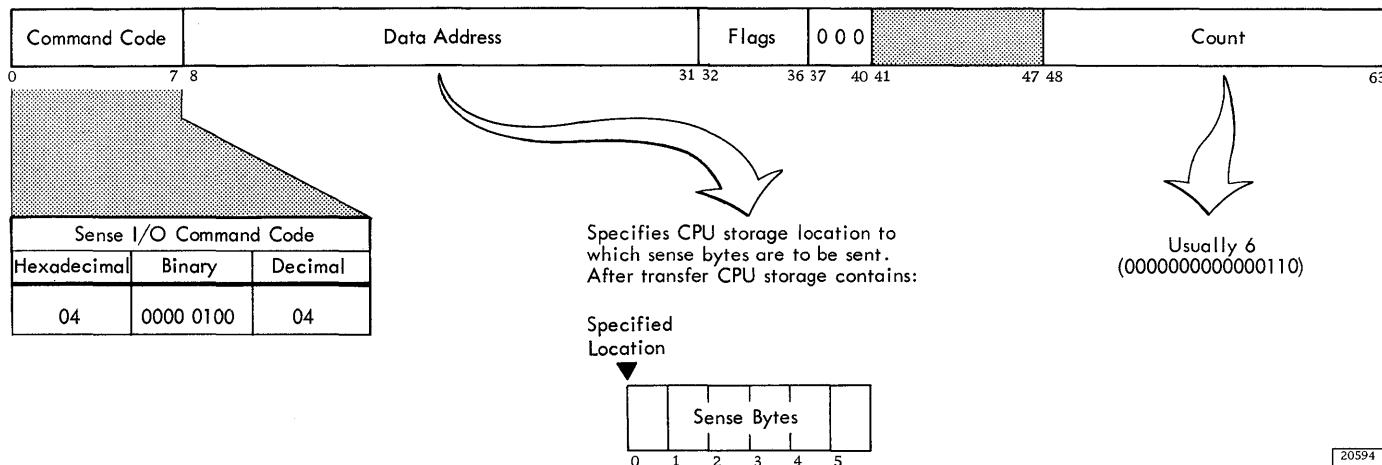


Figure 22. Sense I/O Command

Sense Byte Position	Sense Bit Condition (Status Byte = Unit Check)
Byte 0	
Bit 0	Command Reject
Bit 1	Intervention Required
Bit 2	Bus Out Parity
Bit 3	Equipment Check
Bit 4	Data Check
Bit 5	Overrun
Bit 6	Track Condition Check
Bit 7	Seek Check
Byte 1	
Bit 0	Data Check in Count Field. (Also causes Byte 0, Bit 4 to be turned on.)
Bit 1	Track Overrun. Indicated on Write
Bit 2	End of Cylinder
Bit 3	Invalid Sequence. (Also causes Byte 0, Bit 0 to be turned on.)
Bit 4	No Record Found
Bit 5	File Protected. (Also causes Byte 0, Bit 0 to be turned on.)
Bit 6	Missing Address Marker
Bit 7	Overflow Incomplete
Byte 2	
Bit 0	Unsafe
Bit 1	(Not Used)
Bit 2	SERDES Check
Bit 3	(Not Used)
Bit 4	2314 ALU Check
Bit 5	Unselected Status
Bit 6	(Reserved)
Bit 7	(Reserved)
	} Also causes Byte 0, Bit 3 to be turned on.
Byte 3	
Bit 0	Busy
Bit 1	On Line
Bit 2	Unsafe
Bit 3	Wr Current Sense
Bit 4	Pack Change
Bit 5	End of Cylinder
Bit 6	Multi-Module Select
Bit 7	Seek Incomplete
Byte 4	
Bit 0	Wrong Length Record - Multiplex Buffer Feature
Bit 1	Pending Status - Multiplex Buffer Feature
Bit 2	Not Used
Bit 3	Not Used
Bit 4	
Bit 5	} Used for Module Identification
Bit 6	} Refer to following chart
Bit 7	

Module Identification - Sense Byte 4

Four bits in sense byte 4 are used for module identification. These bits are now used on all 2314's. Since each of the nine modules can be assigned any of the logical channel addresses, the following sense bit indications are provided to allow the system program to identify which particular module is assigned a given address. Sense byte 4 will contain the identity of the disk drive corresponding to the address to which the sense command was issued. The following shows the code used for module identification of disk drives. Bits 0 and 1 (shown as "X") were described earlier.

Sense Byte 4	Disk Drive
XX000000	A
XX000001	B
XX000010	C
XX000011	D
XX000100	E
XX000101	F
XX000110	G
XX000111	H
XX001000	J
XX001111	Module not Defined

Sense Byte 5

This byte is used only when overflow incomplete occurs (sense byte 1, bit 7). The codes in byte 5 indicate the type of command being executed when the overflow incomplete occurs. The codes and their meanings when the unit check (overflow incomplete sense) interruption occurs are:

Sense Byte 5 Code	Meaning
00000110	A read command was in progress
00000101	A write command was in progress
00100101	A search key-data-equal command was in progress, and the comparison is equal to this point
01000101	A search key-data-high command was in progress, and the comparison is equal to this point
01100101	A search key-data-high-or-equal command was in progress, and the comparison is equal to this point.
01010101	Any search key-data operation was in progress, and the comparison is low; or a search key-data equal was in progress, and the record is high
01110101	A search key-data-high or high-equal command was in progress, and the comparison is high.
00000000	Overflow incomplete did not occur

24275C

Figure 23. Sense Byte Summary

Sense Bytes

Unit check status (bit 6 of the status byte) indicates that the I/O device or control unit has detected programming errors (such as invalid command sequences) equipment malfunctions, or conditions requiring operator or program intervention. The conditions causing unit check are detailed by sense information. For the 2314 there are six sense bytes, numbered 0 through 5. All are sent from the 2314 to the channel when a sense command is executed by the 2314. The data address portion of the sense command specifies the address in CPU main storage where the sense bytes are to be placed. When a sense command is sent to the 2314, the sense bytes returned pertain to the module for which the unit checked occurred.

Any information in sense bytes 0, 1, 2, 5, and bits 0 and 1 of byte 4, is reset to all 0's whenever an initial status byte of zero is given in response to a test I/O instruction or a command other than sense, no op, reserve, or release. Also, whenever another unit check is generated, any outstanding sense information is reset and replaced by the new sense information. It is important to issue a sense command after every unit check indication even if the sense information is of no interest. If a sense command is not executed, expected future interruptions may not occur, and some access paths may be unavailable.

In the following sense bit descriptions, the condition(s) applies only if the corresponding bit is on (set to a value of 1). See also Figure 23.

The significance of a "1" condition for each bit is:

<u>Byte</u>	<u>Bit</u>	<u>Designation</u>	<u>Significance of "1"</u>
0	0	Command Reject	Indicates that the 2314 has received an invalid operation code, an invalid sequence of commands, an invalid seek address,* a command associated with a special feature that is not installed; or that a file mask has been violated on a write command. (See Set File Mask.)
0	1	Intervention Required	Indicates that the specified module is not physically attached to the system or, if physically attached to the system, is not available for use because the file motor is not on, a cover interlock is open, the module identifier plug is not installed, etc.
0	2	Bus Out Parity Check	Indicates that the 2314 has detected a parity error during the transfer of a command or data from the channel to the 2314.
0	3	Equipment Check	Indicates that an unusual condition is detected in the control unit or module. Conditions covered by this bit are defined by sense byte 2.
0	4	Data Check	Indicates that a data error has been detected in the information received by the 2314 from the module. Also indicates if index is detected after address mark and before detecting the key, or data field when reading or searching.
0	5	Overrun	Indicates that a chained CCW was issued but it was received too late to be properly executed; or that a byte was received too late (during a read or write operation) to be transferred properly. When writing, the remaining portion of the record area is filled with zeros and the overrun check is generated. When reading, the remaining portion of the record is ignored and the overrun check is generated immediately.
0	6	Track Condition Check	Indicates defective track (bit 6 of the flag byte is 1) to a search command other than search HA, a read command other than read HA or read R0, or any write command, Also indicates that a switch to a defective track was attempted as a result of executing a multiple-track command or a command on a segment of an overflow record (bit 1 of the flag byte is 1) other than search HA, read HA, or read R0. In the case of overflow, overflow-incomplete is also set. Also indicates that an attempt was made to switch from an alternate track (bit 7 of the flag byte) or a defective track (bit 6 of the flag byte) by means of overflow or by a multiple-track command.
0	7	Seek Check	Indicates that the file has been unable to complete a seek because: 1. Transferred seek address is outside the valid address boundaries of the storage device. Unused seek address bytes must be a valid address for the device selected. Command reject is also set. 2. Less than six seek address bytes were sent. Command reject is set.

*A track condition interrupt on an overflow record occurs during a write command.

<u>Byte</u>	<u>Bit</u>	<u>Description</u>	<u>Significance of "1"</u>	<u>Byte</u>	<u>Bit</u>	<u>Description</u>	<u>Significance of "1"</u>														
			3. The equipment failed, which resulted in the access mechanism going to either the inner or outer stop. In this case, command reject is not set.	1	7	Overflow Incomplete	This bit is used with the record overflow operation. It is set with other indicators to signal conditions as follows:														
1	0	Count Area Check	Indicates that a data error has been detected in a count area read from the module. Data check (bit 4) in byte 0 is also turned on. Error detection is the same as described for data check.				<table border="0"> <tr> <td></td> <td>Sets Overflow Incomplete and Other Indicator:</td> </tr> <tr> <td><u>Condition</u></td> <td><u>Track condition (byte 0, bit 6)</u></td> </tr> <tr> <td>Overflow to a defective track</td> <td>Track condition (byte 0, bit 6)</td> </tr> <tr> <td>Overflow from an alternate track</td> <td>Track condition (byte 0, bit 6)</td> </tr> <tr> <td>Data check in data area of overflow record other than last record.</td> <td>Data check (byte 0, bit 4)</td> </tr> <tr> <td>Overflow to file protected boundary</td> <td>File protected (byte 1, bit 5) Command reject (byte 0, bit 0)</td> </tr> <tr> <td>Overflow to wrong track (Head number unequal)</td> <td>Seek check (byte 0, bit 7).</td> </tr> </table>		Sets Overflow Incomplete and Other Indicator:	<u>Condition</u>	<u>Track condition (byte 0, bit 6)</u>	Overflow to a defective track	Track condition (byte 0, bit 6)	Overflow from an alternate track	Track condition (byte 0, bit 6)	Data check in data area of overflow record other than last record.	Data check (byte 0, bit 4)	Overflow to file protected boundary	File protected (byte 1, bit 5) Command reject (byte 0, bit 0)	Overflow to wrong track (Head number unequal)	Seek check (byte 0, bit 7).
	Sets Overflow Incomplete and Other Indicator:																				
<u>Condition</u>	<u>Track condition (byte 0, bit 6)</u>																				
Overflow to a defective track	Track condition (byte 0, bit 6)																				
Overflow from an alternate track	Track condition (byte 0, bit 6)																				
Data check in data area of overflow record other than last record.	Data check (byte 0, bit 4)																				
Overflow to file protected boundary	File protected (byte 1, bit 5) Command reject (byte 0, bit 0)																				
Overflow to wrong track (Head number unequal)	Seek check (byte 0, bit 7).																				
1	1	Track Overrun	Indicates that writing has not been completed by the time the index point is detected. This type of error is created during a write R0; write count, key, and data; or an erase operation.																		
1	2	Cylinder End	Indicates that the CCW command chain has not been completed, and cylinder end has been detected during multiple-track or record overflow operation.																		
1	3	Invalid Sequence	Indicates that an attempt has been made to execute an invalid sequence of CCWs or that two set file mask commands appear in the same command chain. Valid command sequences are defined in the individual command descriptions. Command reject (byte 0 bit 0) is also set when an invalid sequence is detected.	2	0	Unsafe	Indicates that a disk drive malfunction was detected.														
				2	1	Reserved															
1	4	No Record Found	Indicates that while executing a chain of CCWs, the 2314 has detected two index points without completing an intervening command to read or write the data area, read home address, or read R0. It is also set in conjunction with missing address marker if there is no data on the track. No record found is never set if the multi-track bit in the command (bit 0) is on.	2	2	Serializer/ Deserializer Check	Indicates that a bit was either lost or gained when the parallel channel data was converted to serial data during a write operation.														
				2	3	Not Used															
				2	4	ALU Check	Indicates the microprogram has detected an invalid condition, indicating an equipment malfunction.														
1	5	File Protected	Indicates that a seek or write CCW was issued contrary to the file mask. The command reject bit is also set by this condition if the operation violates the write portion of the file mask.	2	5	Unselected Status	Indicates that some bit on the file status lines is on without any module being selected. This condition indicates a device malfunction because no bit should be on prior to selection.														
1	6	Missing Address Marker	A missing address marker, which may indicate a missing record, is detected during the execution of a command or chain of commands which operate on successive count areas on a track. The condition detected is two successive records on a track with equal bit conditions in bit 0 of the flag bytes, with no intervening index point. Missing address marker is set in conjunction with no record found if there is no data on the track.	2	6	Reserved															
				2	7	Reserved															
				3	0	Busy	Indicates that the access mechanism is in the seeking process.														
				3	1	On Line	Indicates that the heads are extended and ready to read or write.														
				3	2	Unsafe	Indicates that the drive is unsafe. No read, write, erase, seek, or head select operations can be performed.														

Byte	Bit	Description	Significance of "
3	3	Write Current Sense	Indicates that the selected drive is writing.
3	4	Pack Change	Indicates that a pack change has taken place.
3	5	End of Cylinder	Indicates that during a cylinder operation the head select register has advanced from 19 to 20 (19 equals head 20, and 20 equals a nonexistent head).
3	6	Multi-Module Select	Indicates that more than one module is selected.
3	7	Seek Incomplete	Indicates that the access mechanism has failed to detent after a determined period of time.
4	0	Wrong Length Record	Multiplex Storage Control feature.
4	1	Pending Status	Multiplex Storage Control feature.
4	2	Not used	
4	3	Not used	
4	4	Module Identification	Each of the nine physical disk modules may be assigned any of the eight logical drive (module) addresses. Sense byte 4 allows the system program to identify which particular module is assigned a given address. This byte contains the identity of the physical drive to which the sense command was issued, according to the following:
4	5	Module Identification	
4	6	Module Identification	
4	7	Module Identification	

Sense Byte 5

This byte is used only when overflow incomplete occurs (sense byte 1, bit 7). The codes in byte 5 indicate the type of command being executed when the overflow incomplete occurs. The codes and their meanings when the unit check (overflow incomplete sense) interruption occurs are:

Sense Byte 5 Code	Meaning
00000110	A read command was in progress.
00000101	A write command was in progress.
00100101	A search key-data-equal command was in progress, and the comparison is equal to this point.
01000101	A search key-data-high command was in progress, and the comparison is equal to this point.
01100101	A search key-data-high-or-equal command was in progress, and the comparison is equal to this point.
01010101	A search key-data operation was in progress, and the comparison is low; or a search key-data equal was in progress, and the record is high.
01110101	A search key-data-high or high-equal command was in progress, and the comparison is high.
00000000	Overflow incomplete did not occur.

For further information, refer to the Record Overflow section in this manual.

BRANCHING IN 2314 CHANNEL PROGRAMS

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

1. If command chaining is specified in the current CCW and execution of this CCW results in a status modifier indication (without detection of other unusual conditions), the channel fetches the next CCW from a main-storage location 16 positions higher than the current CCW (i.e., one CCW is skipped).

Because all search commands cause a status-modifier condition (when the search is satisfied), branching from the chain is accomplished when the search is satisfied.

Sense Byte Four

Physical Drive

Bit: 01234567

XX000000	A
XX000001	B
XX000010	C
XX000011	D
XX000100	E
XX000101	F
XX000110	G
XX000111	H
XX001000	J
XX001111	Module Not Defined

Note: Bits 0 and 1 (XX) are used in the Multiplex Storage Control Feature.

20622

2. Also, the programmer can modify the CCW chain sequence by using the TIC (transfer in channel) command. This command directs the channel to fetch the next CCW from an address specified in the data address field of the TIC. (See Transfer in Channel for additional information.)

these errors causes generation of a unit check interruption.

Note that a TIC is the only CCW that may have a count field of zero. An incorrect length indication cannot occur during execution of a TIC, and the SLI flag is ignored.

The following notation is used in this manual.

These methods of modifying the sequence of a chain of CCWs provide branching capability in the channel program.

TIC*-8

Transfer In Channel

The TIC (transfer in channel) command provides for chaining between CCWs not located in adjacent CPU storage locations. The next CCW is fetched from the location specified by the data address field in the TIC command.

TIC (Figure 24) does not initiate any channel I/O operation, and the I/O device is not signalled that this command is being executed.

TIC cannot be the first CCW designated by a CAW. One TIC command cannot transfer directly to a second TIC command. When either of these programming errors is detected or when an invalid address is specified, a program-check condition is generated. During command chaining, detection of

This is IBM assembler programming notation that indicates that if the TIC is performed, an unconditional branch occurs to the TIC storage address (*) minus a count of eight. In other words, the address in the TIC CCW specifies that location. An example of TIC is shown in the following sequence, in which the TIC branches back to repeat the search ID command.

- Search ID
- TIC*-8
- Read data

If, however, the search is successful, the status modifier bit indicates that the read data CCW is

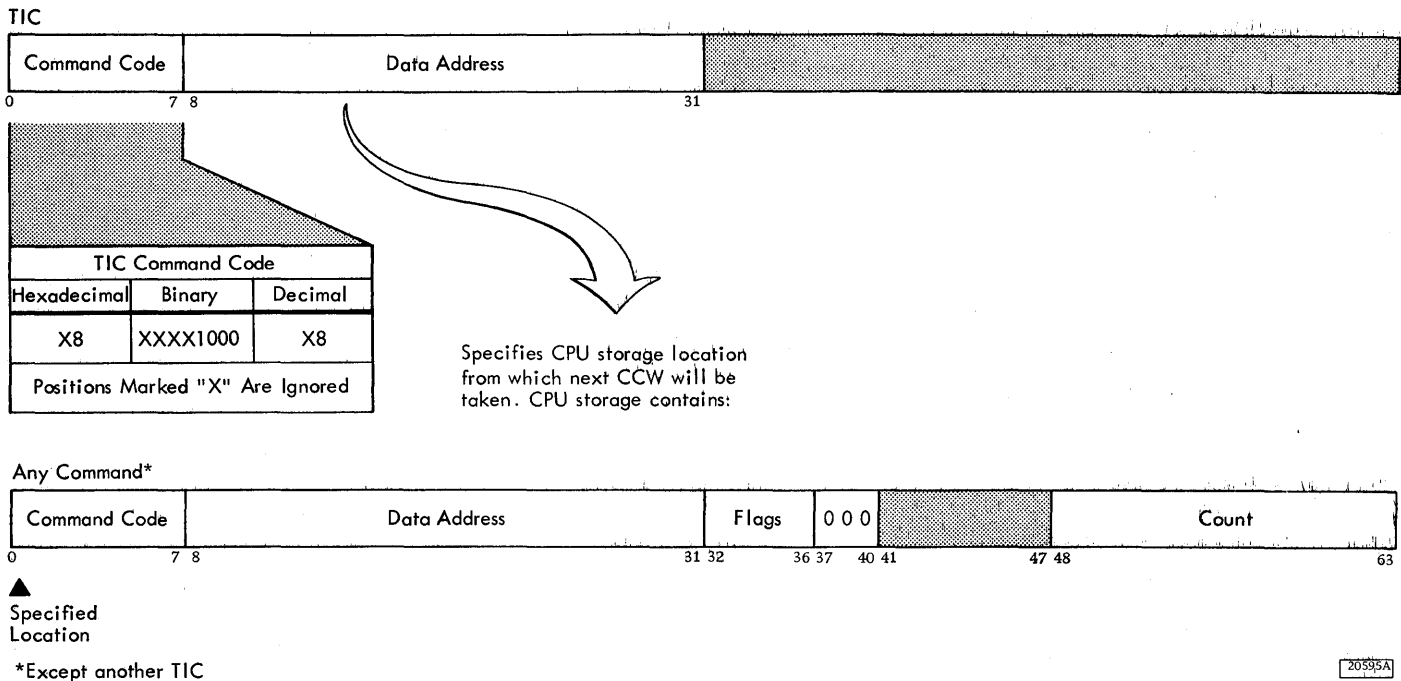


Figure 24. Transfer in Channel Command

the next to be used (i.e., the TIC is not performed because one CCW, the TIC, is skipped).

CONTROL COMMANDS

Control operations do not involve a transfer of data records between the 2314 and the CPU. However, in certain operations, control bytes are transferred between the CPU and the 2314 to enable the operation to take place. These bytes are parity checked during transfer.

No Operation (No Op)

No op (Figure 25) causes no action at the 2314. This command is the only command that may be processed as an immediate command by the 2314. Normally, channel end and device end status are presented by the 2314 at initial selection in response to a no op, except when the 2314 is completing a format write operation and the no op is chained to it. In this latter case, zero initial status is given. Channel end and device end are then presented to the channel after the control unit has completed erasing the remainder of the track.

No op should not be used indiscriminately because it resets orientation information in the 2314. For example, a no op inserted between a read count command and a read data command changes the operation so that the data field read is in the next record (i.e., the one after the record in which the count field is read). Also, any no-record-found (sense byte 1, bit 4) and missing-address-marker (sense byte 1, bit 6) indications are suppressed by insertion of no ops after search commands in command chains.

The SLI flag must be on to avoid an incorrect length check.

Recalibrate

This command (Figure 26) causes the designated module access mechanism to seek to head zero and cylinder zero. Channel end is generated about 15 milliseconds after the 2314 has accepted the command. Device end is presented to the channel when the operation is completed.

The file mask setting affects the recalibrate command in the same way as it does seek commands.

Note the recalibrate is not an immediate command (i.e., at initial selection the 2314 returns a zero status byte if the recalibrate CCW is accepted). The SLI flag must be on in any recalibrate CCW.

Set File Mask

Execution of the set file mask command (Figure 27) causes one byte of data to be transferred from main storage to the 2314 control unit. This byte specifies the write and seek commands that the 2314 is allowed to execute (see Figure 27). If a seek or write command that violates the mask is subsequently issued (i.e., following the set file mask CCW in the same chain), that command is not executed. Unit check is then presented to the channel by the 2314. A subsequent sense command obtains sense information with the file-protect and command-reject bits set on if the command was a write. If the command not executed was a seek, only the file-protect sense bit is set on.

The file protect indication takes precedence over the end-of-cylinder indication. Hence, when file protect is detected, the end-of-cylinder sense bit is not set.

The set file mask command can be issued once at any point in a command chain. At the completion of the chain, the file mask is reset to all zeros.

No-Op




Figure 25. No-Operation Command

Recalibrate




Figure 26. Recalibrate Command

Set File Mask

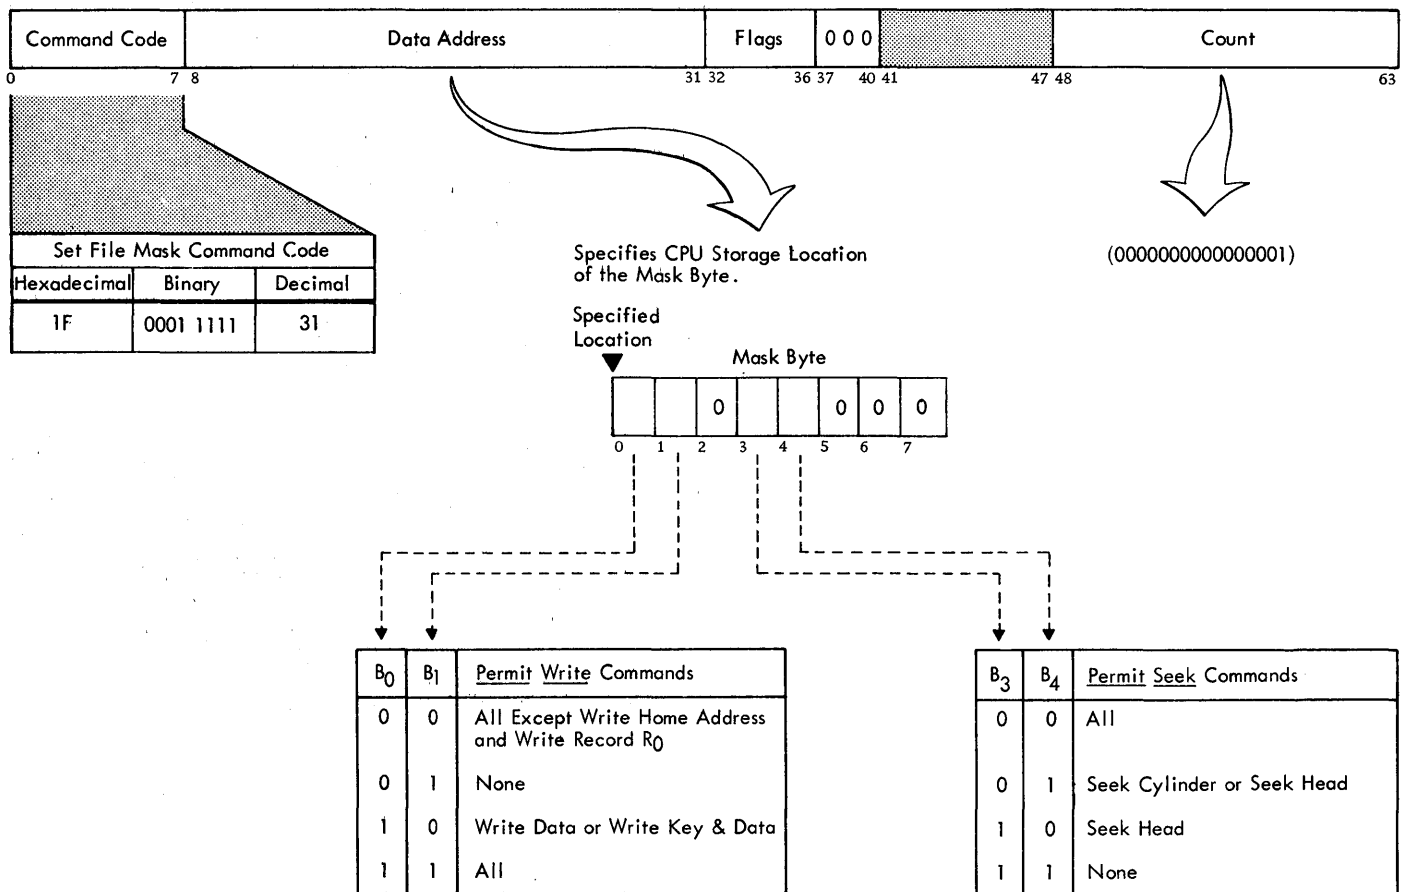


Figure 27. Set File Mask Command

A system reset (results from pressing the CPU system reset key) or a selective reset (i.e., malfunction reset initiated internally by the channel to reset a unit or units that have a malfunction) causes any file mask in the 2314 to be reset to all zeros. If a start I/O is subsequently issued, and a set file mask CCW is not in the chain then used:

- All seek commands are permitted.
- Write count-key-data, write key-data, and write data commands are permitted.
- Write home address and write R0 are not permitted (i.e., bits 0 and 1 of the mask are set to zeros).
- Any command not controlled by the file mask is permitted.

The set file mask command resets the defective and alternate track bit indications in the 2314. Therefore, when included in a CCW chain, the set file mask command should be given prior to any read or search command.

If an attempt is made to issue a set file mask command more than once in any CCW chain, unit check status is signaled by the 2314. A subsequent sense operation indicates command reject and invalid sequence.

Seek Commands

Three seek commands (Figure 28) are used with the 2314:

1. Seek
2. Seek cylinder
3. Seek head

After a start I/O instruction has selected the proper disk module in the 2314, a seek CCW is used to locate the proper area on the disk pack. When necessary, and depending on the seek command used, the access mechanism in the module is moved to the specified cylinder and the designated head is selected.

Execution of a seek CCW results in transfer of a 6-byte seek address from main storage to the 2314.

Any seek CCW count field (bit positions 48 through 63) should specify a 6-byte field. If the count

Seek
Seek Cylinder
Seek Head

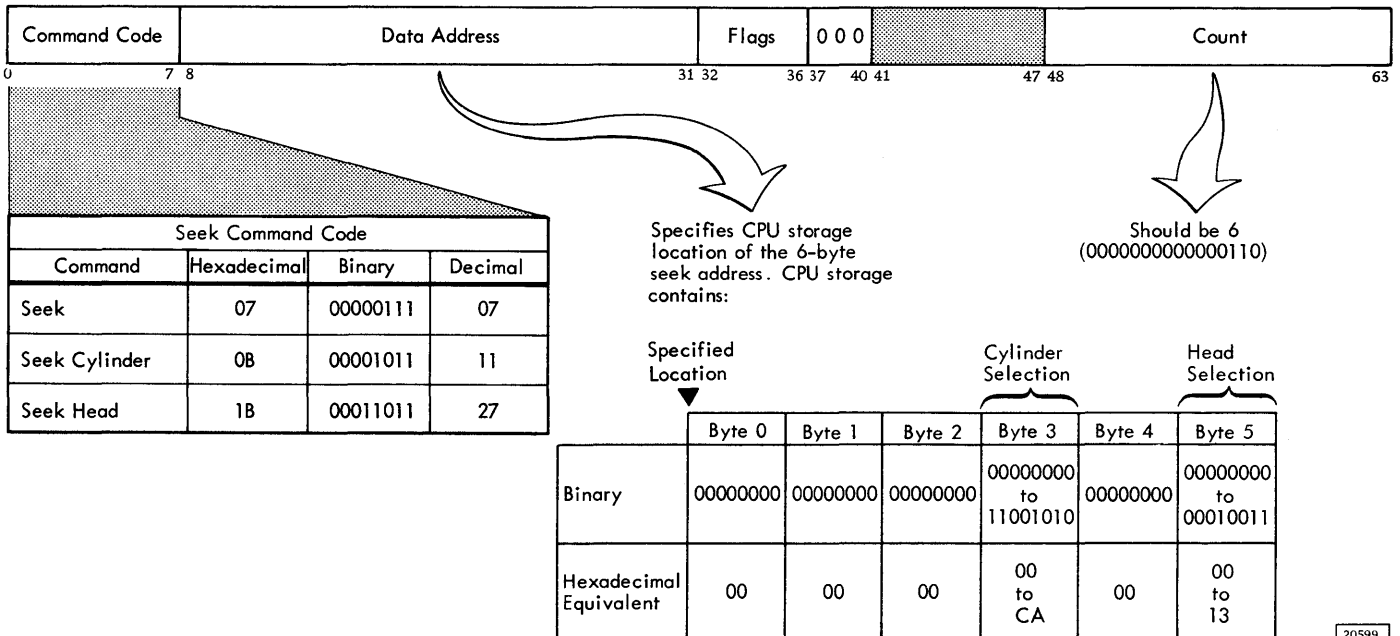


Figure 28. Seek Commands

is greater than six, the 2314 operates on the first six bytes transferred and, if the seek CCW SLI flag bit is zero, a wrong length record is signalled to the CPU (i.e., incorrect length, bit 41, in the CSW is set on). If the count is less than six, the seek command is not executed and the seek check and command reject sense bits are set on. In this case, unit check, channel end, and device end conditions are set in the CSW.

The six bytes sent from main storage must indicate a valid cylinder and head address. If the address is not valid, the seek check and command reject sense bits are set on. Unit check, channel end, and device end are presented to the channel.

If the seek address is valid, channel end is presented to the channel after the address is transferred. If no mechanical motion is required, device end is presented with channel end. If mechanical motion is necessary, device end is made available to the channel after the mechanical movement is completed.

If a parity error occurs during transfer of the seek address, the command is not executed. The bus-out parity sense bit is set and unit check, channel end, and device end are presented to the channel.

If the 2314 is unable to complete the seek for any other reason, the 2314 presents unit check, channel end, and device end when polled by the channel. The seek check sense bit is also set.

A seek command does not have to be preceded by any other CCW in order to be executed. However, see Set File Mask for the effect the file mask has on seek commands.

The three different seek commands enable restriction by the file mask and provide compatibility with other direct access devices (see Set File Mask). For the 2314, seek and seek cylinder have the same effect. Seek head is used when the file mask permits seeks only on the cylinder where the access mechanism is already positioned.

The valid addresses are shown in Figure 28.

Seek

All six seek address bytes referenced by the CCW are used to determine the cylinder and track to be accessed. For the 2314, bytes 0, 1, and 2 must be zero.

Seek Cylinder

Only the four low-order bytes (bytes 2 through 5) referenced by the CCW are used to determine the seek address. (All six bytes must be present to constitute a valid seek address. Bytes 0, 1, and 2 must be zero.)

Seek Head

Only the two low-order bytes (bytes 4 and 5) referenced by the CCW are used to determine the seek address. (All six bytes must constitute a valid seek address.)

SEARCH COMMANDS

On all search operations, the channel operates in the write mode while the 2314 operates in the read mode. The 2314 compares the information coming from CPU storage with the information read from the record on a track. Each search command operates on only one record. The search can then be repeated for the next record by branching back to the search CCW with a transfer-in-channel command.

If the search condition is satisfied, a status modifier indication occurs and the channel fetches the next CCW in the command chain from a location 16 positions higher than the current search CCW (i.e., one CCW, usually a TIC, is skipped). This allows modification of a command chain as a function of the data read from the track.

For all search commands, command code bit 0 determines whether or not a multiple-track operation is to occur (see Multiple-Track Operation); that is, whether switching to the next read/write head in the cylinder is to occur when the index point is detected. If bit 0 is not set (0), head switching does not take place; if bit 0 is set (1), head switching does take place. If head switching does take place and if the search command is repeated, the next track is used. This procedure allows sequential searching of an entire cylinder by repeating the search command once

for each record to be searched. For a multiple-track search, it is important to start the search at the beginning of the track by starting the chain with a read HA or read R0 command.

The following command chain illustrates the procedure for reading record N, identified by a key stored at location X in the CPU, on a single track.

<u>Command Chain</u>	<u>Function</u>
Search Key X	Compare key with search argument.
TIC*-8	Transfer back to search, if search is unsuccessful.
Read Data N	Read data area if status modifier was returned from search.

The channel is busy during a search operation.

Search Home Address Equal (Search HA)

This command (Figure 29) causes the 2314 to search for the index point and then to compare the cylinder and head address coming from main storage with the cylinder and head in the home address area of the track. The flag byte is not transferred or compared during this command.

If the logical comparison is equal, then channel end, device end, and status modifier signals are generated and sent to the CSW. If the logical comparison is unequal, then channel end, device end, and unit check are generated.

Search home address generates a no-record-found sense indication if the specified home address search is not satisfied. If no field is found, the no-record-found and missing-address-marker sense bits are set on.

If the CCW count is greater than four bytes, the search operation is completed when the 2314 count equals zero (four bytes transferred). The 2314 terminates the command with a channel end and device end. The status modifier is generated if the logical comparison is satisfied.

If the CCW count is less than four bytes, the logical comparison between the data coming from CPU storage and the data coming from the track continues until the CCW count reaches zero. At the time the 2314 count reaches zero, channel end and device end are generated. A status modifier is generated if the search condition is satisfied on the short field.

If a parity check, or data check is detected, unit check, channel end, and device end signals are sent to the channel at completion of the command. If overrun is detected, the operation is terminated and channel end, device end, and unit check are also set on.

Search HA Equal

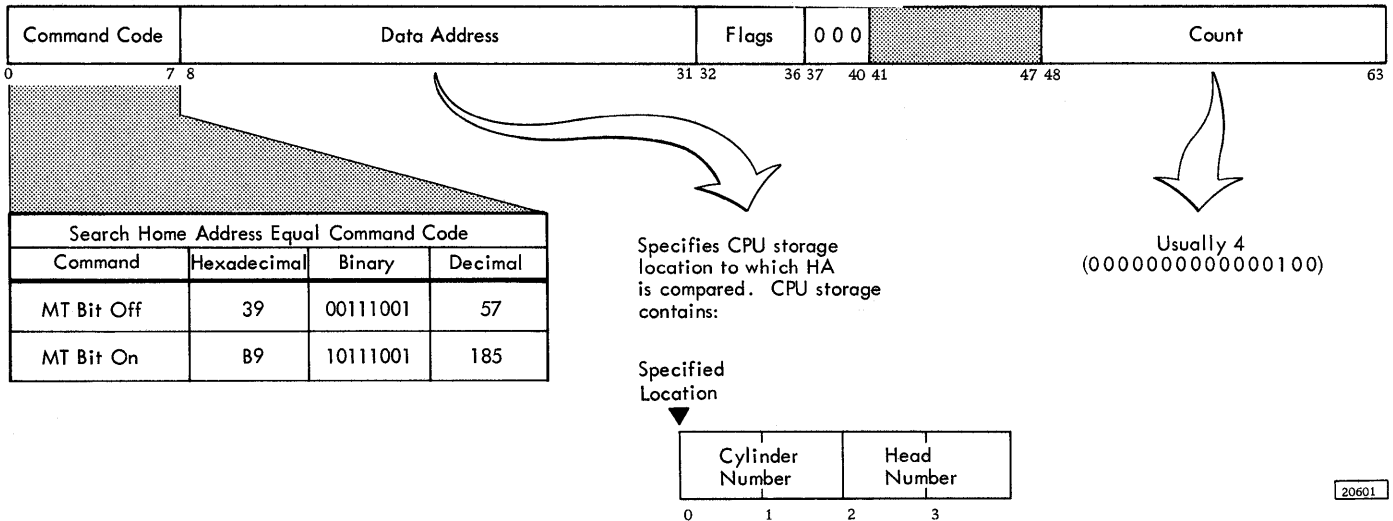


Figure 29. Search Home Address Command

Search Identifier (Search ID)

Search ID commands (Figure 30) cause a comparison to be made between five bytes of data from CPU storage and the five-byte record identifier portion of a count area from the storage unit.

The ID to be searched is the next ID that is encountered on the track, including the ID of R0.

If the CCW count is greater than five bytes, the search operation is completed when the 2314 count of five is decreased to zero. The 2314 terminates the command with a channel end and device end. The status modifier is generated if the logical comparison is satisfied.

If the CCW count is less than five bytes, the logical comparison between the data coming from main storage and the data coming from the track continues until the CCW count reaches zero. Channel end and device end are generated at the end of the

count area. A status modifier is generated if the search condition is satisfied on the short field.

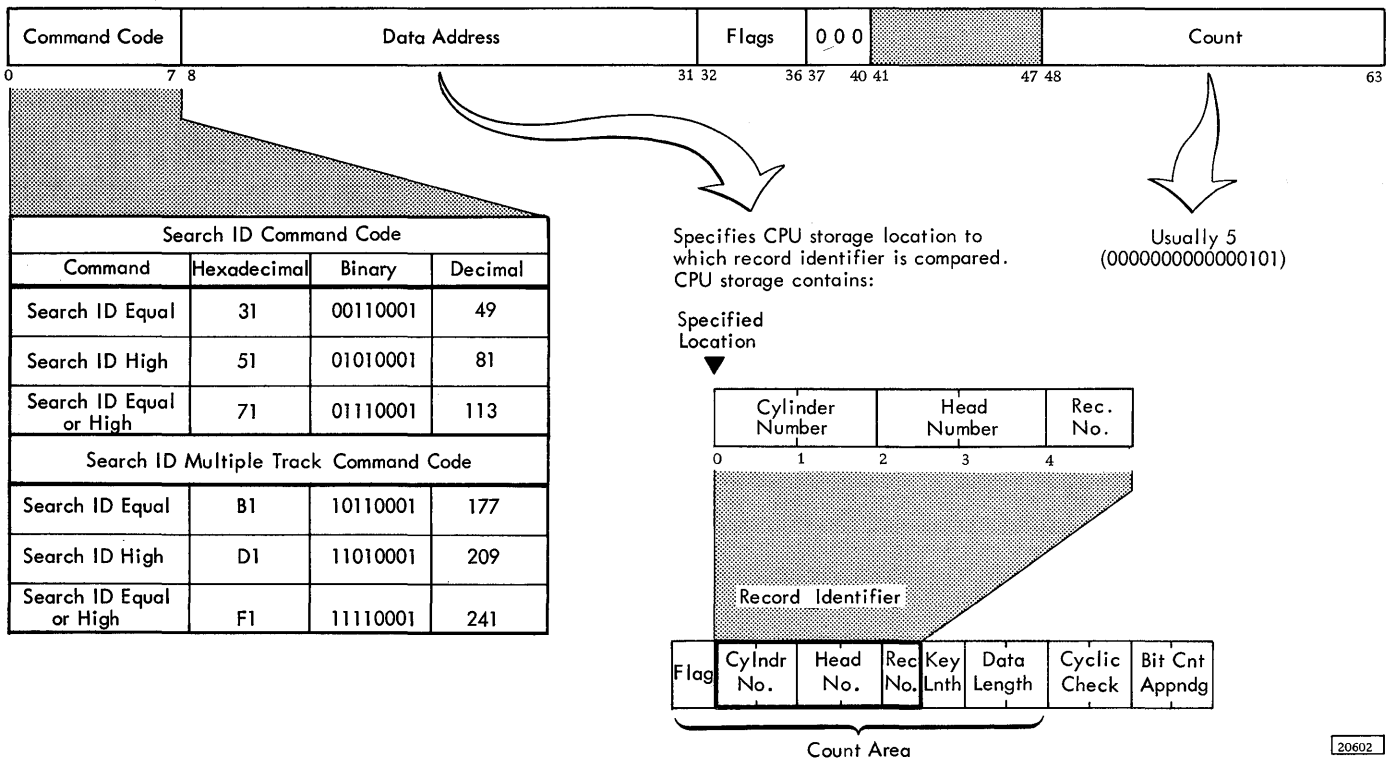
If a parity check or data check is detected during a search-ID operation, unit check, channel end, and device end signals are generated at the completion of the command. If overrun is detected, the operation is terminated and unit check, channel end, and device end are set on.

A search ID command does not have any chaining requirements.

If command code bit 0 (multiple track) is 0, the search is confined to one track and can be repeated until either the search condition is satisfied or until two index points are sensed; at which time unit check, channel end, and device end signals are generated. The no-record-found sense bit is also set on.

If the multiple track bit is a 1, it is important to start the chain with a read HA or read R0. The search can then be repeated until the search

Search ID Equal
Search ID High
Search ID Equal or High



20602

Figure 30. Search ID Commands

condition is satisfied, or until the end-of-cylinder is detected. At this time a unit check signal is generated and the end-of-cylinder sense bit is set on.

Search ID Equal

If a logical equal comparison is encountered, channel end, device end, and status modifier signals are generated.

If the logical comparison is unequal, channel end and device end signals are generated.

Search ID High

This command operates in a manner similar to that of the search ID equal command, except that the comparison is made for a high condition. The high condition indicates that the ID on the track is higher than the ID in main storage. The comparison is made byte by byte until the first unequal comparison is detected.

Search ID Equal or High

This command operates in a manner similar to that of the search ID equal command except that the comparison is made for either an equal or high condition. The equal or high condition indicates that the ID on the track is equal to or higher than the ID in main storage.

Search Key

Execution of a search key command (Figure 31) causes a comparison to be made between the key in main storage and the key read from the track. The key to be searched is the key of the record following the next address marker detected by the read/write head. However, when the search key CCW is chained from a search ID or read count CCW, the key searched is in the record in which the ID is read or searched. Search key passes over R0 unless chained from a search ID that has searched the ID of R0.

Search Key Equal
Search Key High
Search Key Equal or High

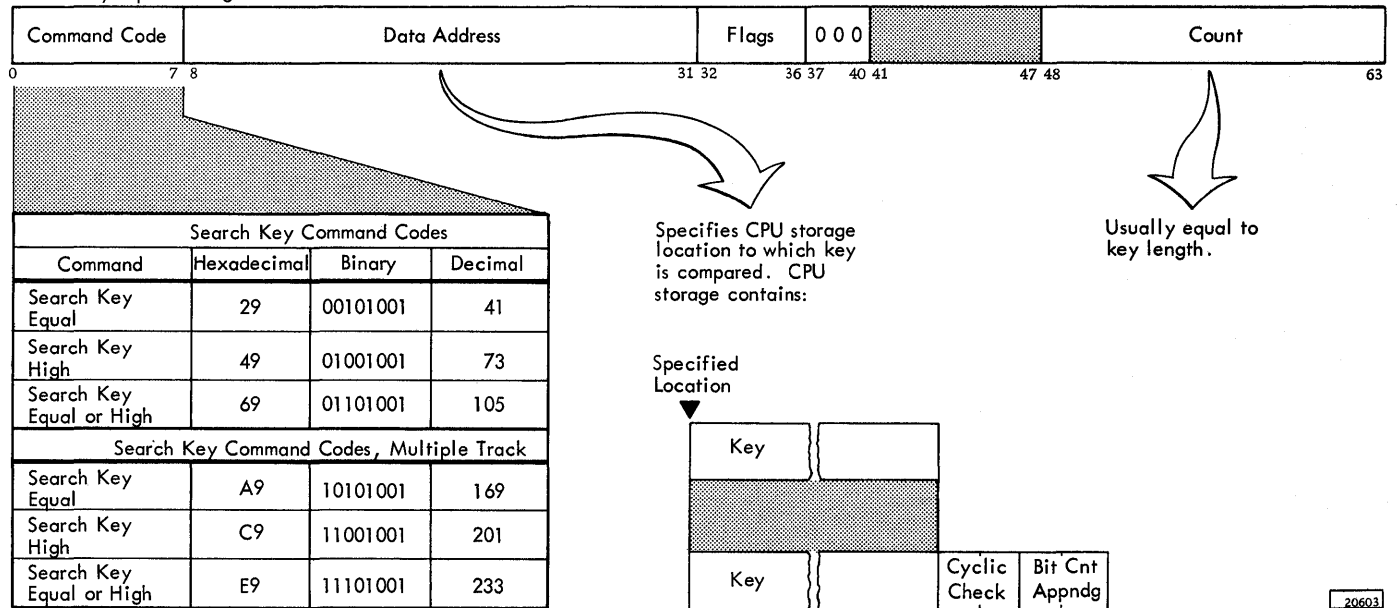


Figure 31. Search Key Commands

If the CCW count is greater than the key length, the search operation is completed when the 2314 count equals zero. The 2314 terminates the command with a channel end and device end. The status modifier is generated if the logical comparison is satisfied.

If the CCW count is less than the key length, the logical comparison between data from CPU storage and the data from the track continues until the CCW count reaches zero. When the 2314 count reaches zero, channel end and device end are generated. A status modifier is generated if the search condition is satisfied on the short field.

If a parity check or data check is detected during the search-key operation, unit check, channel end, and device end signals are generated at the completion of the command. If overrun is detected, the operation is terminated and unit check, channel end, and device end are set on.

A search-key command does not have any chain-requirements.

If the multiple track bit is 0, the search is confined to one track and can be repeated until either the search condition is satisfied or until two index points are sensed, at which time unit check (and the no-record-found sense bit), channel end, and device end are set on. If the multiple track bit is 1, it is important to start the chain with a read HA or read R0. The search can be repeated until either the search condition is satisfied or until end-of-cylinder is detected. If end-of-cylinder is detected, the end-of-cylinder sense bit is set on.

The search key command never returns a status modifier if the key length of the record searched is zero.

Search Key Equal

If a logical comparison on equal is encountered, channel end, device end, and status modifier signals are generated. If the logical comparison is unequal or the record has no key area, then channel end and device end signals are generated.

Search Key High

This command operates in a manner similar to that of the search-key equal command except that the comparison is made for a high condition. The high condition indicates that the key in the storage unit

is higher than the key in CPU storage. The comparison is made byte by byte until the first high comparison is detected.

Search Key Equal or High

This command operates in a manner similar to that of the search key equal command except that the comparison is made for either an equal or high condition. The equal or high condition indicates that the key in the storage unit is equal to or higher than the key in CPU storage.

Search Key and Data (File Scan Commands)

The search key and data commands (Figure 32) provide a rapid means of locating and selecting records on the basis of a comparison of selected parts of the key and data fields of a record. Execution of these commands differs from that of the search key command, in that:

1. The comparison continues through the data field, and
2. For any bytes in the key or data fields on which a comparison is not desired, an FF (hexadecimal) byte should be inserted in the scan argument in main storage.

The search key and data commands are normally used in the same way as the search key commands. A TIC command is used to cause re-execution of the search key and data command until a record is found that satisfies the comparison criteria. When a record is found that satisfies the search comparison, the 2314 presents status modifier, channel end, and device end status to the channel. The status modifier indication causes the channel to skip over one CCW (normally, the TIC) and continue with the channel program. The search argument must be set up in main storage and the address of the argument inserted in the search key and data CCW.

If the multiple track bit is off (0), the search continues until the condition is satisfied (indicated by status modifier) or until two index points are sensed. At that time, unit check (with the no-record-found sense bit set on), channel end, and device end signals are generated. If the multiple track bit is

Search Key and Data Equal
 Search Key and Data High
 Search Key and Data Equal or High




Figure 32. Search Key and Data Commands

1, the command chain must be started with a read HA or a read R0. The search can be repeated until the specified condition is met or until end-of-cylinder is encountered, at which time the end-of-cylinder sense bit is set on.

No more positions than the number specified by the CCW count are compared. If the CCW count is greater than key length plus data length, the search operation is limited by the record key and data lengths. A zero key length causes comparison of data only. If the channel count is less than the data length plus the key length, a truncated search is performed. Status modifier is generated (if the comparison is satisfied) on the basis of the comparison up to where the CCW count reaches zero.

If a parity check or data check is detected during a search key and data operation, unit check, channel end, and device end signals are generated at the

completion of the command. If overrun is detected, the operation is terminated and unit check, channel end, and device end are set on.

Search Key and Data Equal

This command causes the 2314 to make a logical equal comparison between the key and data information from the storage device with the information coming from main storage.

Search Key and Data High

This command operates in a manner similar to that of the search key and data equal command except that the comparison is made for a high condition. The high condition indicates that the storage device information is higher than the information in main

storage. The comparison is made byte by byte until the first unequal comparison is detected.

Search Key and Data Equal or High

This command differs from the search key and data equal command in that the comparison is made for an equal or high condition. The equal or high condition indicates that the storage device information is equal to or higher than the information in main storage.

A file scan function over an entire cylinder can be executed by the following sequence of CCWs:

<u>Command</u>	<u>Remark</u>
1. Seek	Position access.
*2. Search ID equal (Previous)	Find record prior to beginning of scan area.
3. TIC	On unequal, transfer back 8 bytes to repeat search.
*4. Read count <i>a</i>	Read count of record into <i>a</i> .
**5. Search key and data equal	Scan key and data.
*6. TIC	On unequal, transfer back 16 bytes to repeat steps 4 and 5.
7. Seek head <i>a</i>	Reposition access. (Necessary for overflow record environment only.)
8. Read Home Address	HA of last track read on step 4, to prevent spurious no-record-found if index point is passed twice (necessary only if step 7 is omitted.)
***9. Search ID <i>a</i> equal	Search for ID read in step 4.
10. TIC	On unequal, transfer back 8 bytes to repeat step 9.
**11. Read key and data	Read key and data.

- *Multiple track bit on.
- **Multiple track bit is irrelevant.
- ***Multiple track bit off.

Step 4 of the sequence causes the count information coming from the module to be read into CPU storage. Key and data information from the module are compared with the information from main storage on step 5. The comparison is based on the condition specified by the search key and data CCW. When the specified condition is encountered, the program must reorient to the beginning of the record with steps 9 and 10. Then the desired key and data is read in step 11.

A search for a low, low or equal, or unequal record can be done with the following sequence:

- Read count.
- Search key data (high equal, high, or equal respectively).
- TIC*+16.
- TIC*-24.

Space Count

This command (Figure 33) provides a means of bypassing a defective area on a track in order to recover data immediately following on that track. This command can be used for two different situations:

1. A defective home address or record zero area.
2. A defective count area, other than record zero.

Defective Home Address or Record R0: This situation sets unit-check and data check, or data check, in the count field to interrupt the program. In this case, record R1 is read by the following command chain:

Space Count (SLI bit on)
Read Count Key Data

If two successive space count commands are followed by read CKD, record R2 is read. (If space count is the first command in a chain, it searches for index point, then counts over the home address and record-zero areas. The function is the same if a no op is the first, and space count the second, command in a chain.)

Defective Count Area (Record N): This situation sets the data-check-in-count-field sense bit. In this case, key and data for record N are read by the following command chain:

Read Home Address
Search ID (record N-1)
Transfer-in-Channel (to repeat search if unsuccessful)
Read Data (record N-1; Skip and SLI flag bits on)
Space Count (over record N count area)
Read Key Data (record N)

The space count CCW must contain a count of three bytes and the main storage address where the record N key length and data length are located.

If index point occurs before an address mark is found, no record found is set. If index point occurs while spacing over data, track overrun is set. The three bytes of data sent by the channel are used by the 2314 for key length (the first byte) and data length (remaining two bytes) if the space count

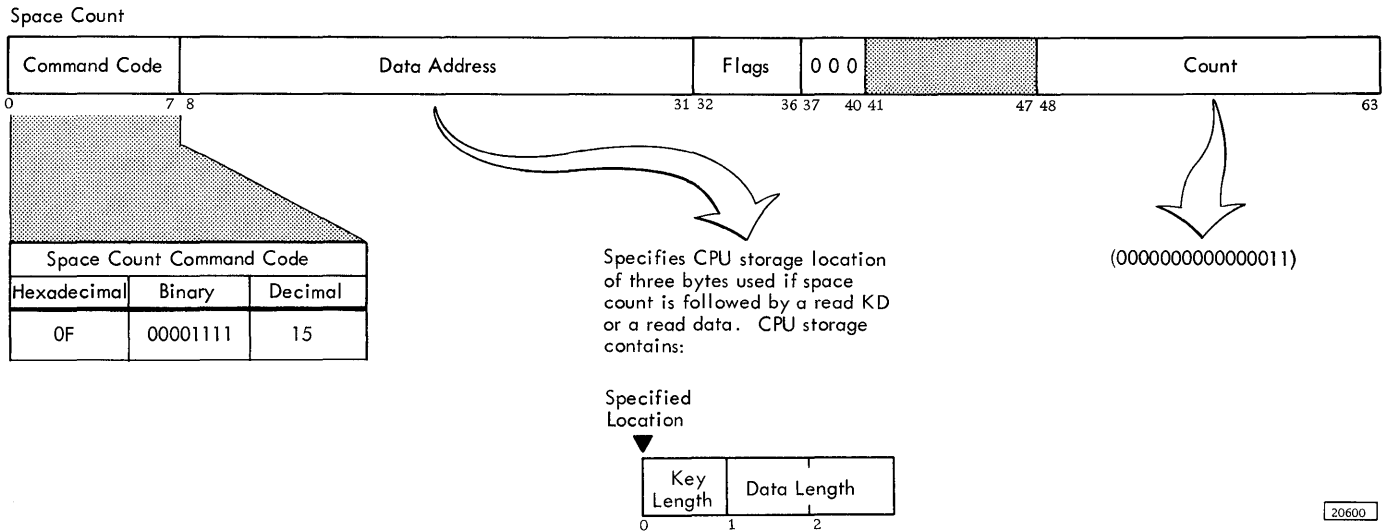


Figure 33. Space Count Command

command is followed by a read key and data or a read data.

Command Sequence: When space count is chained from a read, search, or space count command, it searches for the next address marker. It then spaces over the count area of that record, reading the values of the key length and data length from main storage.

If a space count is used to space over record N count area and the 2314 fails to detect the address marker prior to the count area of record N, missing address marker will not be posted. The 2314 will skip over record N and orient on record N+1. The next key and/or data command will operate on record N+1 instead of record N.

If the last command in the previous chain is a read-count-key-data (instead of read-key-data), record N+1 is read normally and record N is ignored. The space count CCW key length and data length are ignored when the next command in the chain reads the count area.

The space count command is rejected, with invalid-sequence and command-reject bits set, if it is chained from a write command. Also, a space count command sets the file mask so that a following write command in the same chain will be rejected.

READ COMMANDS

Execution of a read command results in transfer of information from the 2314 to main storage. On all read commands, the 2314 checks (by means of the data check bytes) the validity of each area of a

record as the record is read from a track. A parity bit is added to each byte as it is sent to the channel.

If a data check is detected, the command is terminated at the end of the field in which the error occurs. The appropriate sense bits are set (i.e., data-check-in-count-field and/or data check) and unit check, channel end, and device end are set in the CSW. Note that a data check will not be set for the key field in a read data command. Also, data check will not be set for an HA field on a read R0 command.

If an overrun is detected, the operation is terminated and unit check, channel end, and device end are set in the CSW.

No read command needs to be preceded by any other CCW in order to be executed although read commands are normally chained from search commands. Read IPL, however, cannot be preceded by a set file mask command in the same chain.

Read Home Address (Read HA)

This command (Figure 34) causes the 2314 to search for the index point. Detection of the index point causes the flag, cylinder, and head bytes of home address information to be transferred from the 2314 to CPU storage. Exactly five bytes should be transferred, including the flag byte. If the channel count is less than five, only that number of bytes is transferred.

Read Count

This command (Figure 35) causes bytes 1 through 8 of the count area following the next address marker

Read HA

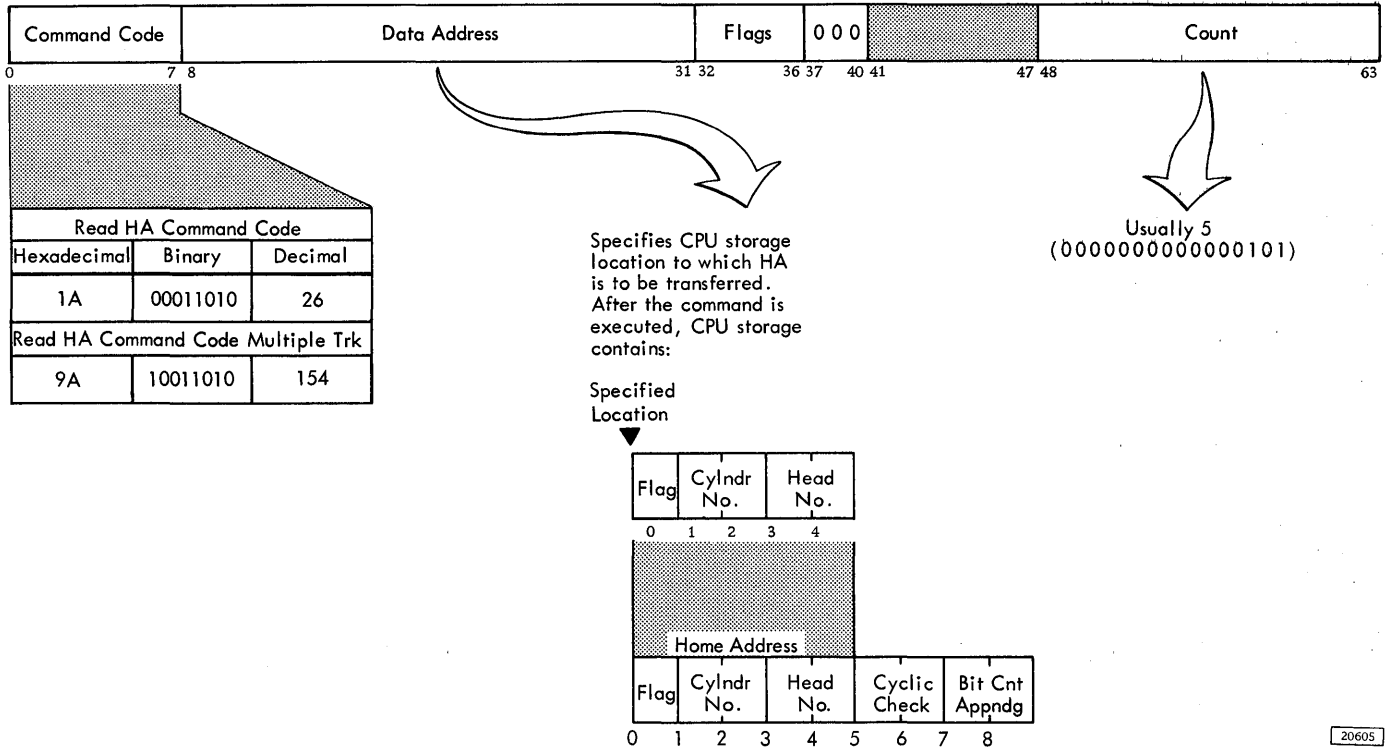


Figure 34. Read Home Address Command

Read Count

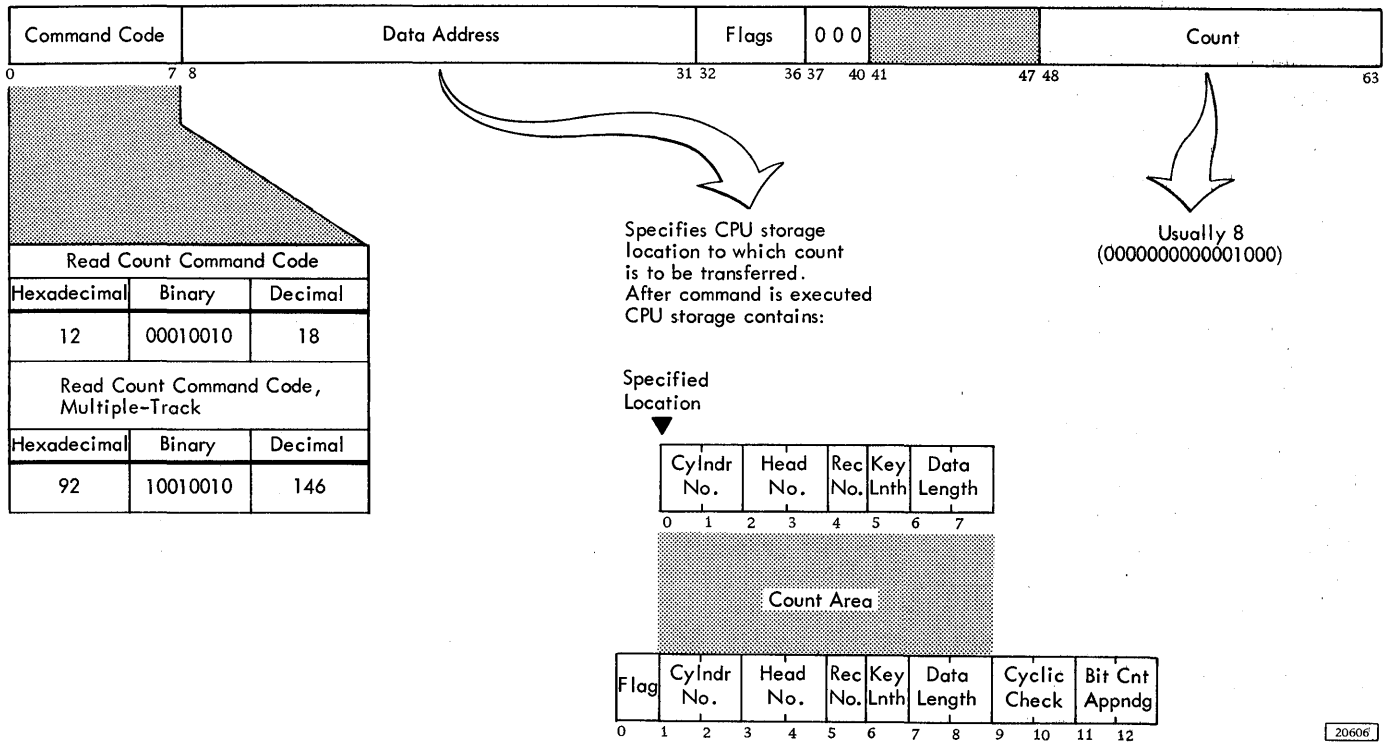


Figure 35. Read Count Command

to be transferred from the 2314 to main storage. The number of bytes of information to be read should always be eight. If the channel count is less than eight, only that number of bytes is transferred.

The count area of R0 cannot be read by a read count command because R0 has no address marker.

Read Track Descriptor Record (Read R0)

This command (Figure 36) causes the 2314 to search for index point. Detection of index point causes the 2314 to read (internally) the home address and the following gap. When these areas have been passed, record R0 (count, key, and data) is transferred from the 2314 to main storage.

If the read R0 command is chained from a search HA or read HA, it is executed immediately and does not cause a search for index marker.


Read Data

Execution of a read data command (Figure 37) results in transfer of the data area of a record from the 2314 to main storage. The data to be read is: (1) the data area of the same record read by a search ID or search key command from which the read data command is chained (this is the normal method), or (2) the data area of the record read by a read count command from which the read data command is chained, or (3) the data area of the record following the next address marker encountered on the track.

Read Key and Data

This command (Figure 37) causes the key and data areas of a record to be transferred from the 2314 to main storage. If key length is zero, this command

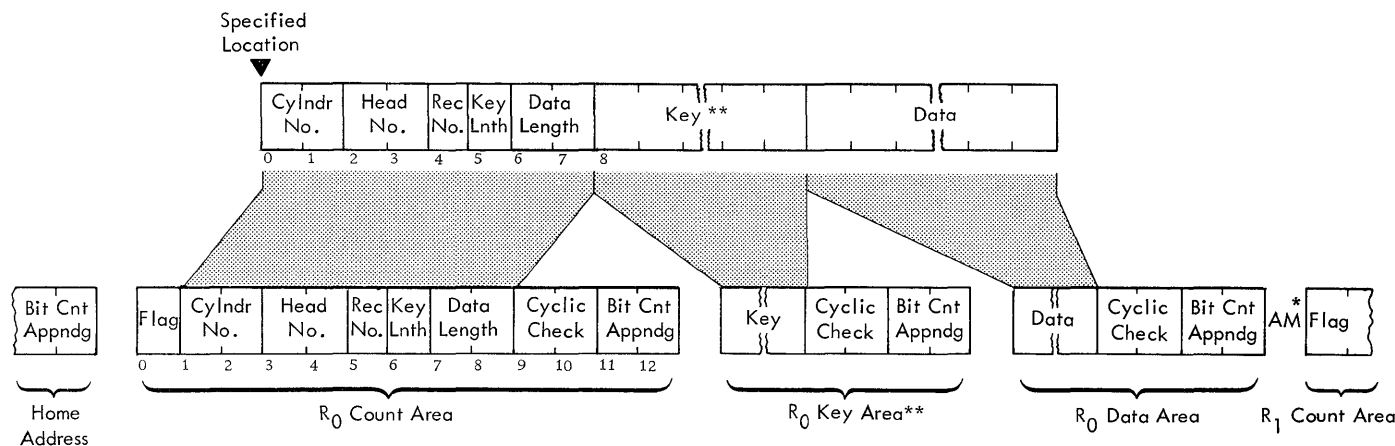
Read R0



Read R0 Command Code		
Hexadecimal	Binary	Decimal
16	00010110	22
Read R0 Command Code, Multiple-Trk		
96	10010110	150

Specifies CPU storage location to which R0 count, key, and data are to be transferred. After command is executed CPU storage contains:

Specifies the number of bytes to be transferred to CPU storage from R0.



* Address Marker
** Key may not be present

20607

Figure 36. Read Track Descriptor Record (R0) Command

Read Data
 Read Key and Data
 Read Count, Key and Data

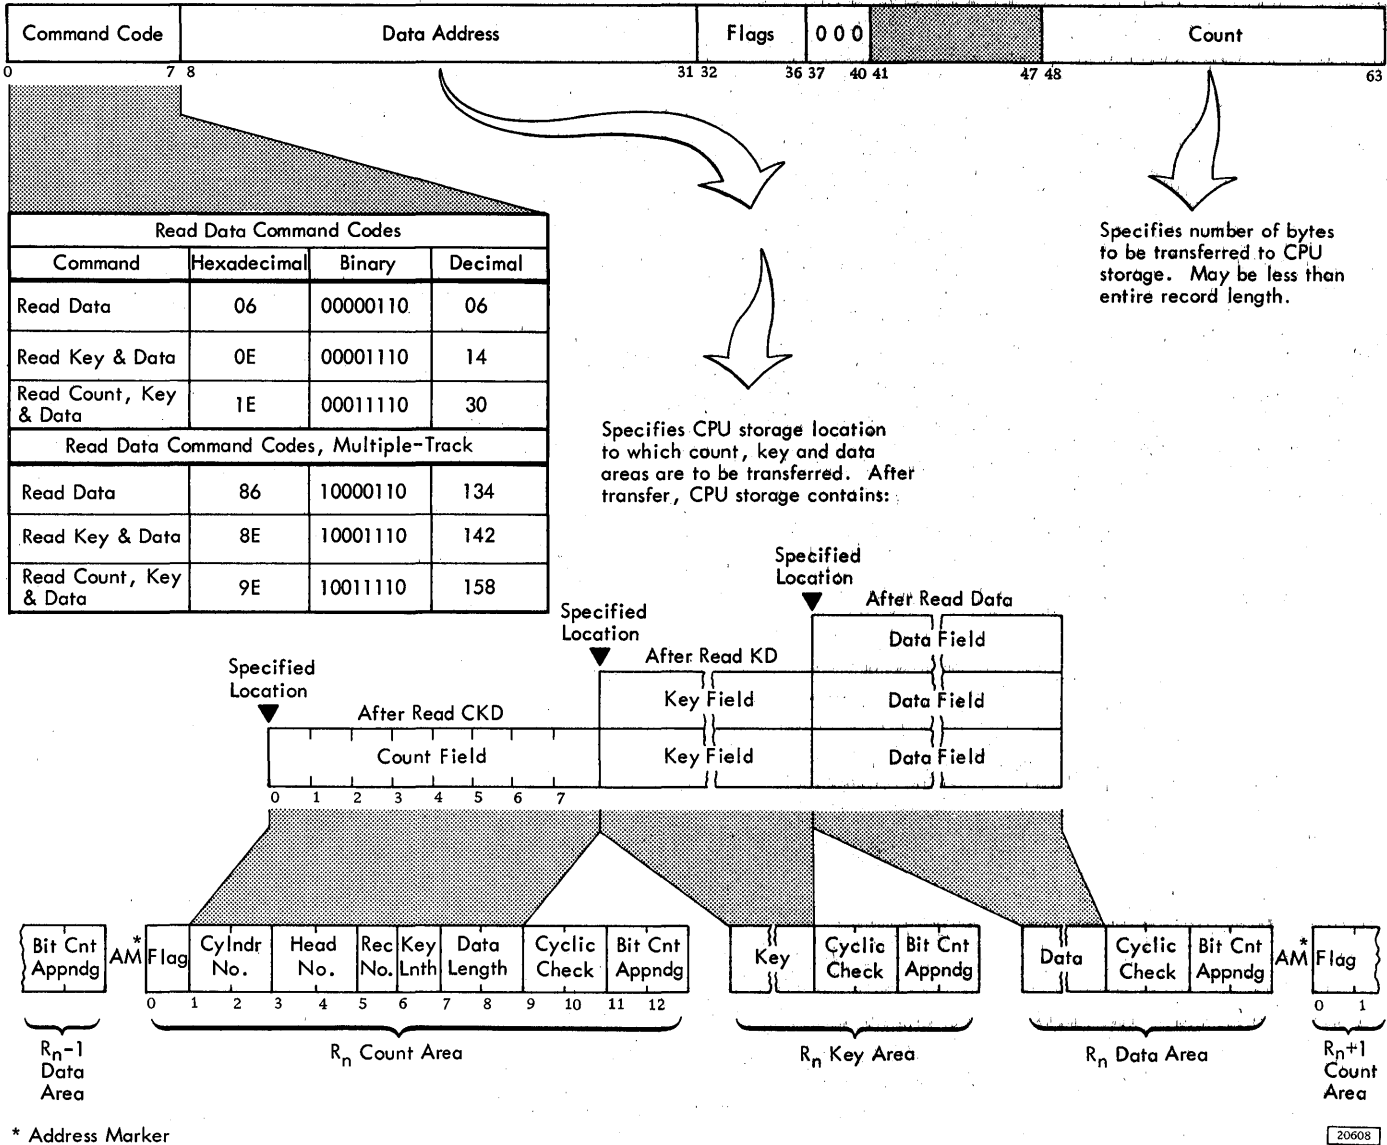


Figure 37. Read Data, Read Key and Data, and Read Count, Key, and Data Commands

operates like a read data command. The key and data areas to be read are (1) the key and data areas of the same record read by the search ID command from which the read key and data command is chained (this is the normal method), or (2) the key and data areas of the same record read by the read count command from which the read key and data command is chained, or (3) the key and data areas of the

record following the next address marker encountered on the track.

Read Count, Key, and Data

This command (Figure 37) causes the entire record (count, key, and data) following the next address

marker (AM) to be transferred from the storage device to main storage. Record R0 is bypassed because it is not preceded by an address marker. Therefore, this command is normally chained from a search for the ID of the preceding record.

Read IPL (Initial Program Load)

The read IPL command (Figure 38) causes the 2314 to seek to cylinder 0, head 0 in the addressed module. The 2314 then searches for index point. After index point is detected, the read IPL command operates as if it were a read data command. The data area to be read is in the first record after R0.

The IPL operation is initiated when the load key on the system console is pressed.

Read IPL cannot be preceded by a set file mask command in the same chain; otherwise, unit check and command reject are set.

WRITE COMMANDS

Write commands cause data to be transferred from main storage to the module. During the transfer, the 2314 adds appropriate data check bits to each area written. At the completion of the command, channel end and device end signals are sent to the channel. Write commands can be grouped into (1) formatting write commands, used to initialize tracks, and (2) data write commands.

Formatting Write Commands

The following are formatting write commands:

1. Write home address
2. Write R0

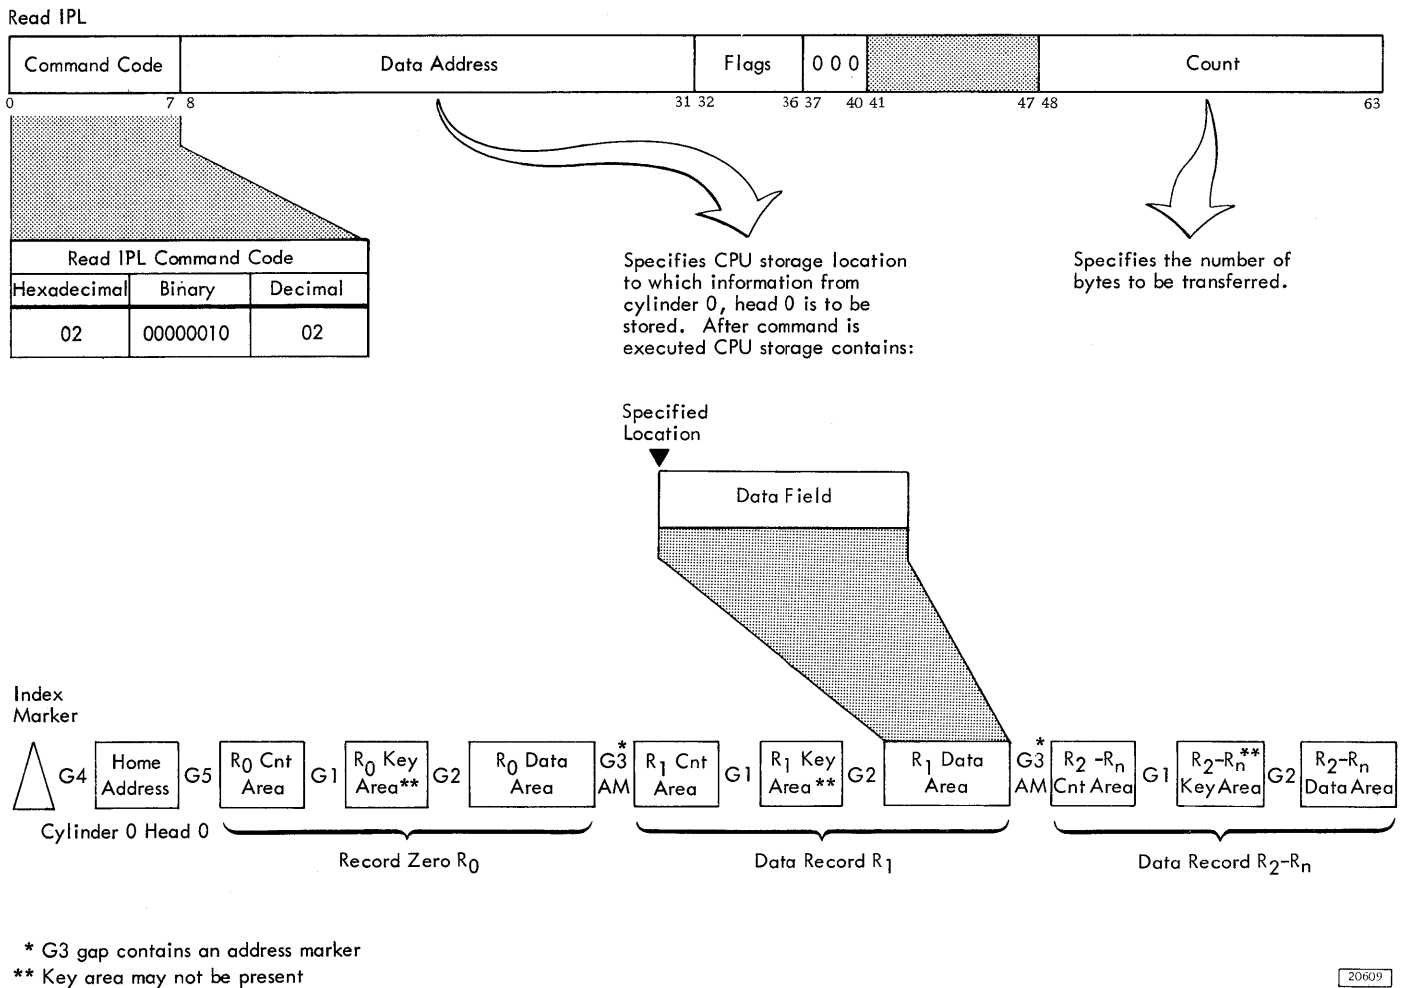


Figure 38. Read Initial Program Load Command

3. Write count, key, and data
4. Erase
5. Write special count, key, and data

A format write command is used to initialize tracks and records and to establish the lengths of the areas within each record. If a command other than a format write is command chained from a format write command, the 2314 retains the command and executes it after the track has been written and the index marker is detected.

Detection of a parity check or an overrun condition during a format write operation causes a unit check to be signalled to the channel at the end of the operation. The 2314 writes valid zeros from the time the overrun condition is detected to the end of the record.

Write Checking

To achieve the maximum level of performance of which the 2314 is capable, the program should provide for repeating commands which show a unit-check error. These errors are often due to temporary conditions which are not present in subsequent tries.

A write command which does not write correctly because of temporary or intermittent conditions can be detected by immediately issuing a read command (with skip and SLI bits on). In this way, any such "soft" write error can be corrected while the data is still available in main storage. If this write-check procedure is not followed, the "soft" write

error becomes a "hard" error, which can be corrected only by a reconstruction or adjustment.

In IBM programming systems, use of write check is optional ("verify" option in IBM programming systems). However, the programmer should weigh the system down time and the time consumed by the procedure to recover from an error against the time consumed in write checking, before deciding not to write check.

Write Home Address

This command (Figure 39) causes the 2314 to search for the index marker (index point). When the index marker is detected, the specified data is transferred from main storage to the 2314. The 2314 receives five bytes of data (FCCHH: flag cylinder and head bytes) from the CPU, and adds four data check bytes.

At this point, channel end and device end are signalled to the channel.

If the CCW count is less than five, the 2314 records 0's until five bytes have been written. If the CCW count is greater than five, the 2314 transfers only the first five bytes from CPU storage. A write home address command is normally used to establish track identity within a storage device. Each track must be initialized with a home address before a data operation that involves that track can take place.

If the home address area of a track is defective, the home address/R0 areas can be moved down the track. To do this, use the following procedures:

Write HA

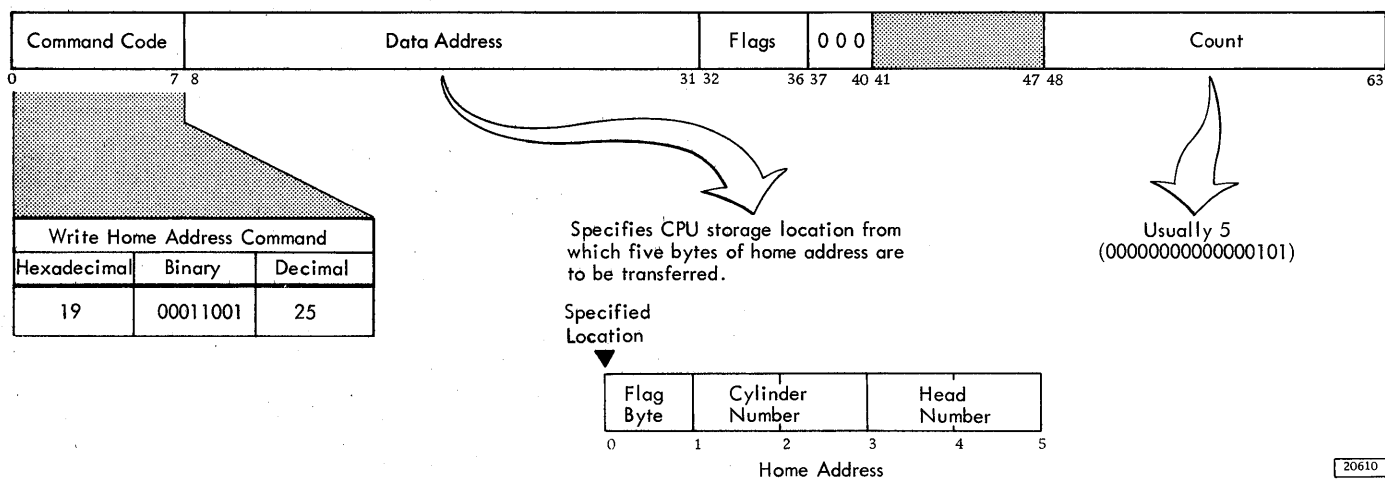


Figure 39. Write Home Address Command

1. Assign an alternate track.
2. Write the home address on the defective track, using a flag byte with bits 1 and 6 set to 1 and all other bits set to 0. Because bit 1 in the flag byte is set to 1, the 2314 adds 705 bytes to the gap area that is between index marker detection and the home address.
3. Write record R0. No other records can be written on the defective track. R0 should contain the address of the alternate track.
4. Seek to the alternate track and write the home address and set the flag bytes bits 6 and 7 to 0 and 1, respectively.
5. Write R0 on the alternate track. The contents and length of the data field are immaterial on the alternate track.

Note that sending the flag byte with bit 1 set to 1 (step 2) does not cause the 1 to be written on the track. The only function this bit performs is to signal the 2314 control unit to increase the home address gap by 705 bytes.

Chaining Requirements: Execution of this command is dependent upon execution of a correct set file mask command (preceding the home address CCW) in the same command chain.

Write Track Descriptor Record (Write R0)

This command (Figure 40) causes data specified in the CCW to be transferred from main storage to the storage device.

The first eight bytes transferred from main storage make up the count area. The flag byte is generated by the 2314, not transferred from main storage. The remaining data is written in the key and data areas as specified by key length and data length in the count area. The 2314 writes the data check bytes at the end of each area. Channel end and device end signals are generated after the data check bytes of the data area have been written.

The CCW count field specifies the number of bytes to be transferred from main storage by the channel. This should be eight bytes plus key length plus data length. If the CCW count is less than this sum, the 2314 writes 0's so that the key and data areas conform to the lengths specified in the record count area.

Chaining Requirements: This command causes invalid sequence and command reject signals to be generated, unless it is chained from a write home address or search home address CCW that compared equal on all four bytes of the home address.

Write R0

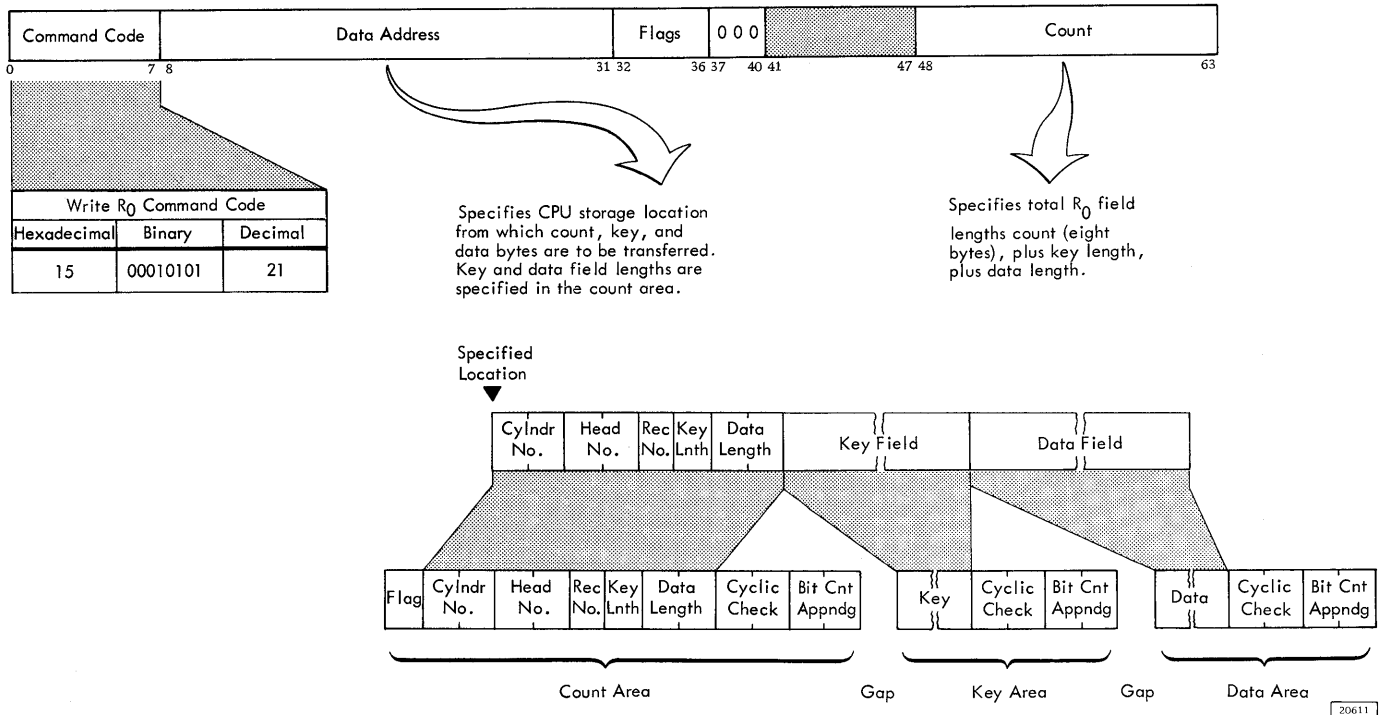


Figure 40. Write Track Descriptor Record (R0) Command

Write Count, Key, and Data

This command (Figure 41) differs from write R0 in that the 2314 causes an address marker to be written on the track immediately before the count area.

This command causes an invalid sequence signal to be generated unless it is chained from a write R0, another write count, key, and data, or from a successful search ID equal or search key equal command. This search must not be a truncated search (one in which the CCW count is less than the length of the area). A read-data or read-key-and-data CCW may be inserted between a search equal CCW and a write count, key, and data CCW.

Erase

This command (Figure 42) is used to erase the end of a track after a track overrun has occurred. It is the same as a write CKD except that no address mark is written. Data is written on the selected storage device, but cannot be recovered.

Both the channel and the 2314 are busy during the execution of this command.

Chaining requirement: This command must be chained from a successful search equal identifier or search equal key command, or from a write R0 or a write count, key, and data command. The search must not be a truncated search, or one in which the CCW count is less than the length of the area to be searched. A read data or a read key and data CCW may be inserted between a search CCW and an erase CCW.

Write Special Count, Key, and Data

This command (Figure 43) functions like the write count, key, and data command except that it causes a 1 bit to be placed in bit position 1 of the flag byte when the 2314 generates and writes the flag byte. It is used with record overflow operations to write overflow record segments prior to the last segment.

Write Data
Write Key and Data
Write Count Key and Data

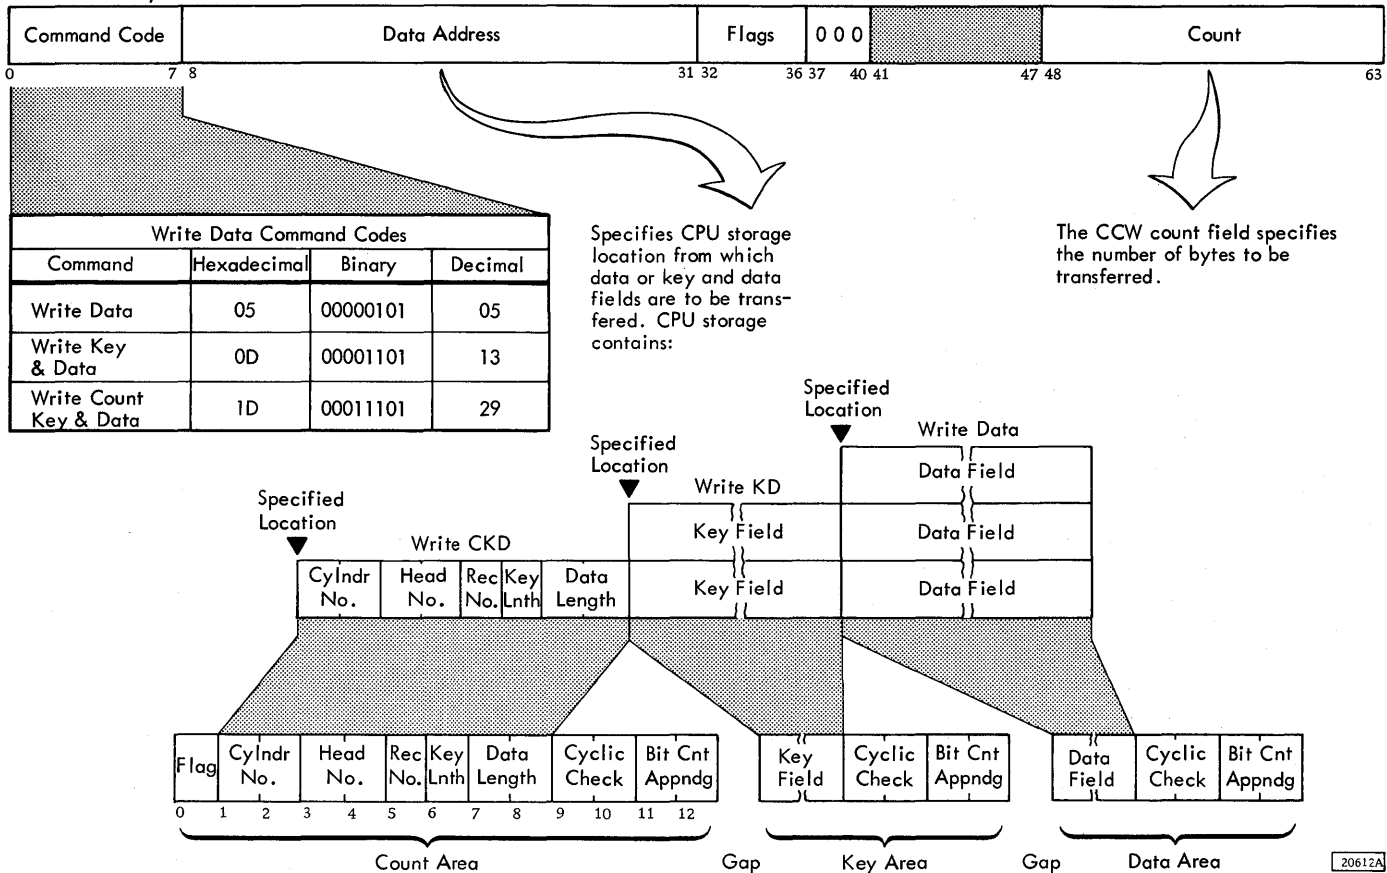


Figure 41. Write Data, Write Key and Data, and Write Count, Key, and Data Commands

Erase

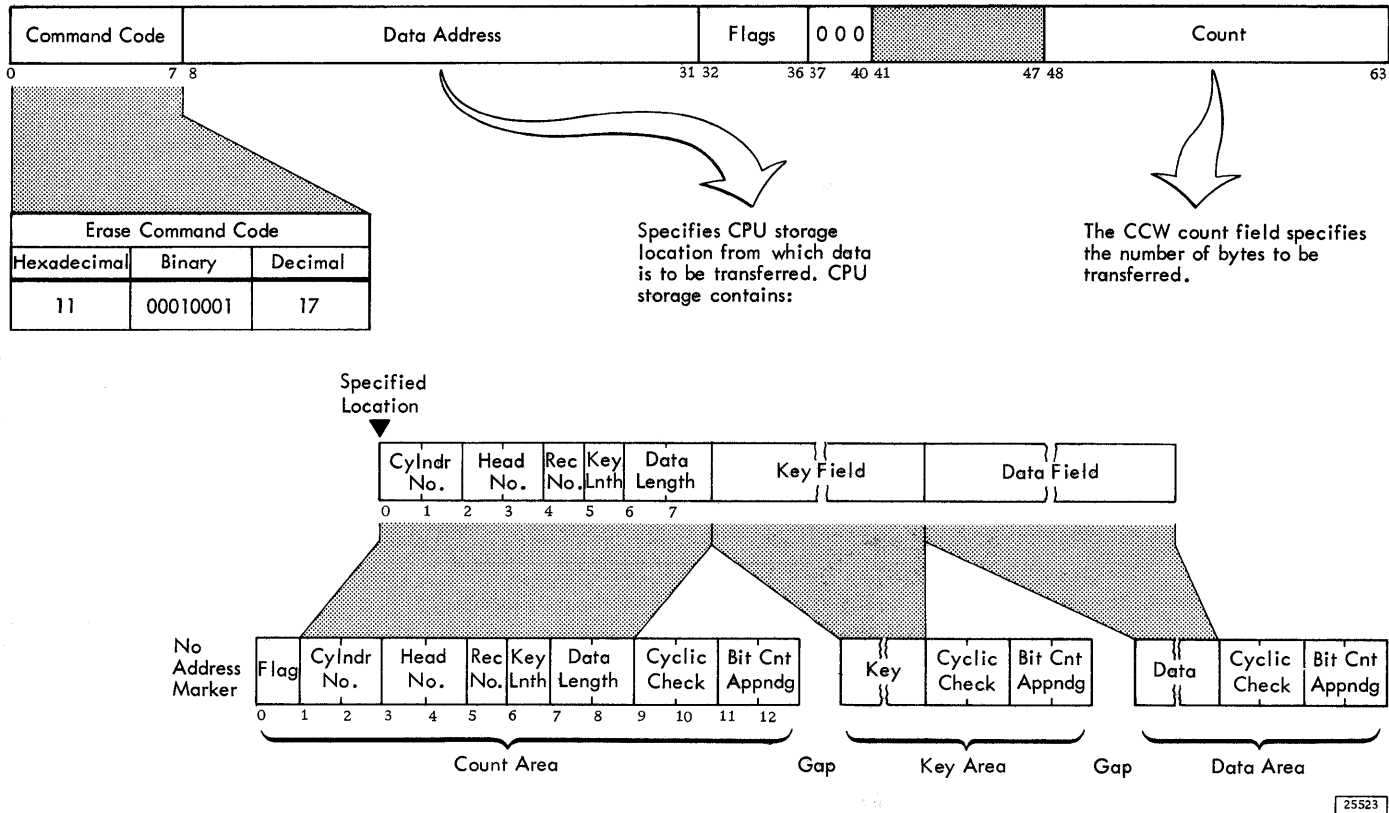


Figure 42. Erase Command

Data Write Commands

A data write operation is used for normal record updating after the track has been formatted. Detection of a parity error or an overrun condition causes a unit check to be signaled to the CSW at the completion of the command. A data check occurring in record areas that must be passed over but not written terminates any write command before data is written.

Write Data

This command (Figure 41) causes the specified data to be transferred from CPU storage to the storage device. Writing occurs only in the data area of the record. The number of bytes to be written is specified by the write data CCW and may be less than the data length specified in the count area when the track was formatted. If the CCW count is less, the control unit fills the remainder with 0's. When the CCW count is greater, the control unit transfers data only until the data length specified by the data length portion of the count area has been reached. The control unit appends the data check code bytes, then signals the channel with device end. The remaining data in

the CCW residual count will not be written. If the SLI (suppress length indicator) bit 34 in the CCW is off the device end will be accompanied by unit check and incorrect length (bit 41) in the CSW.

Chaining requirements: A write data command must be chained from a successful search equal identifier command or from a search equal key command. The search must not be truncated. Unit check (invalid sequence and command reject sense bits) is signaled to the channel if the write data command is not properly chained.

Write Key and Data

This command (Figure 41) causes the specified data to be transferred from CPU storage to the storage device. Writing occurs in the key area and the data area of the record. The number of bytes to be written is specified by the write key and data CCW. The length of the key and data areas may be less than the lengths specified in the count area when the record was formatted; if so, the control unit fills in the remainder with 0's.

Write Special Count, Key, and Data

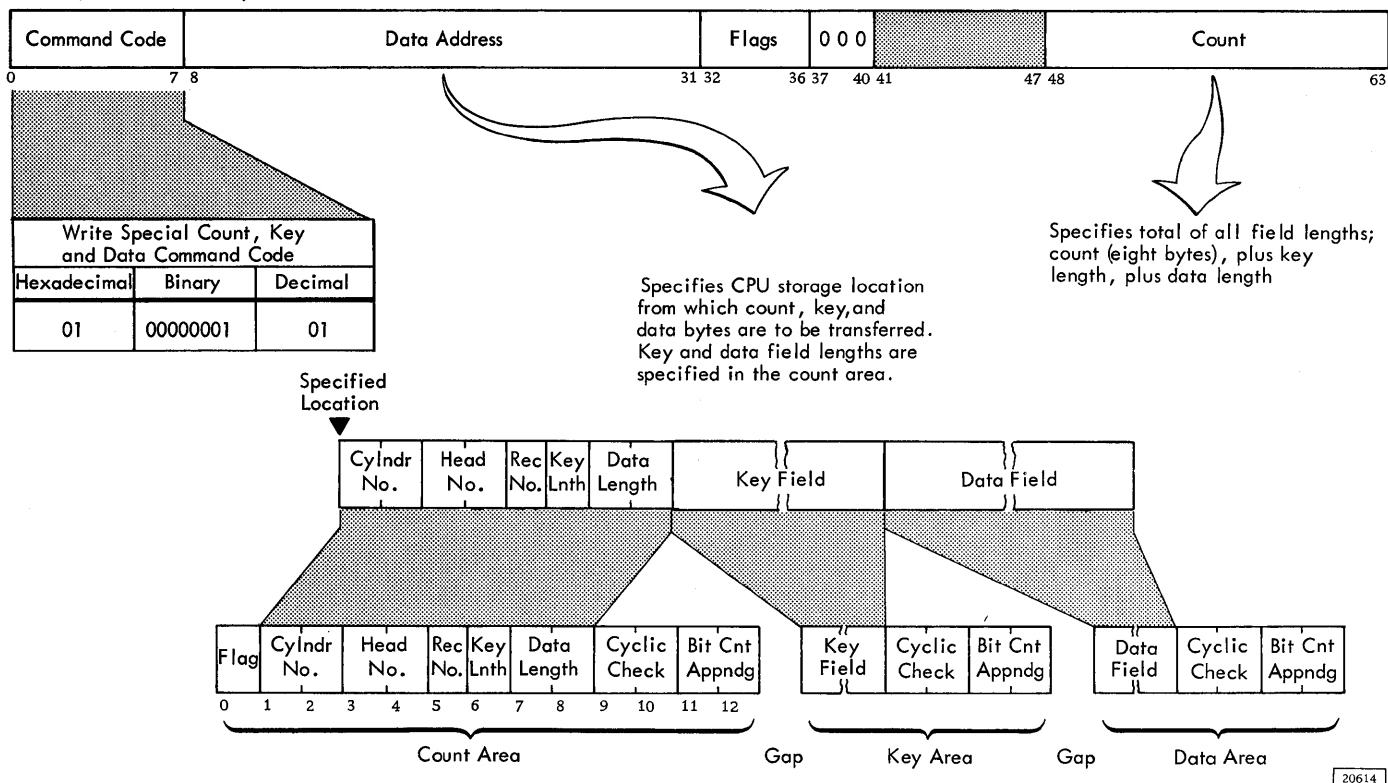


Figure 43. Write Special Count, Key, and Data Command

If a CCW count exceeds the previously formatted area, the control unit transfers data to the device only until the key and data length specified in the count area have been reached. The control unit then signals the channel with device end. The CCW count will not be zero; therefore, the channel sets a bit 41 (incorrect length) in the CSW unless the SLI (bit 34) is present in the CCW.

The number of bytes specified in the CCW count should be key length plus data length. If the record was formatted with a key length of zero, the number of bytes specified are all written in the data area, as in a write data command. If a key area is needed in a record which was formatted without one, the record must be reformatted.

When the number of bytes specified has been written, the control unit appends the data check bytes and signals channel end and device end to the channel.

Chaining requirements: A write key and data command must be chained from a successful search equal identifier command. The search must not be truncated. Unit check is signaled to the channel if the write KD command is not properly chained.

END OF FILE

The end of a logical file is indicated by a record whose data area length is zero. This may be any record on a track.

This record is normally written by a write count, key, and data CCW. The indicated zero data length causes the 2314 to record a data area of four data check bytes.

An end-of-file area may also have a key area.

The 2314 does not require any data transfer to support the writing of an end of file data field. Operations that cause unit exception when a zero data length is detected are: read IPL, read R0, read CKD, read KD, read D, write KD, write D, or search KD. Unit exception is not set for read count, write CKD, search key, or search ID.

MULTIPLE TRACK (M-T) OPERATION

The 2314 has the ability to automatically select the next sequentially numbered head on an access:

mechanism under control of bit 0 of the operation code in the CCW. This is accomplished without loss of a disk revolution. Switching to the next head number takes place at index point if this bit is set to 1. The operation continues until the search is satisfied or the end-of-cylinder is detected. The M-T bit is recognized on all read and search commands.

A recalibrate command must be followed by another seek command before a command calling for a multiple track read or search is attempted.

The M-T bit is effective only if it is in the command under execution when the index point is detected.

Operation of home address commands that have the M-T bit on results in switching to the next head

number before the home address on the first track is read.

RECORD OVERFLOW

Record overflow is provided to allow a logical record to overflow from one track to another. It is useful in processing records which exceed the capacity of a track. The cylinder boundary then becomes the limiting factor to the size of a record.

Formatting Overflow Records

A portion of an overflow record that is written on one track is called a record segment. The write special

(Text continued on page 49.)

(This page intentionally left blank.)

count, key, and data CCW is used for formatting all segments of an overflow record except the last segment. The last segment is written by the normal write count, key, and data CCW.

The write special CKD CCW causes a bit to be written in bit position 1 of the flag byte in the record segment being written. This bit indicates that there is another part of this record located on the next track of a cylinder. Otherwise, the write special CKD CCW functions just like the normal write command.

All overflow segments must be recorded as the first record following R0 on the overflow track.

Overflow segments are normally records without a key field (though key fields may be used), since only the key field of the first segment has significance. All overflow record segments, except the first and the last one, are full-track records. Each segment has a count field and a data field. The key length and data length in a count field always pertain to the segment and not to the entire overflow record.

Processing Overflow Records

If bit 1 of the flag byte is set on, the following commands do not terminate at the end of a record segment, but head switching to the next track in the cylinder occurs when index point is detected during the following operations:

1. Read count, key, and data
2. Read key and data
3. Read data
4. Write key and data
5. Write data
6. Search key and data
7. Continue scan commands (see Continue Scan Commands).

Head switching may not occur, however, if an error or some type of head advance interruption (such as end-of-cylinder) occurs.

If the next record segment is marked with bit 1 on in the flag byte, the operation continues to the next track. When a segment is found that is not marked, the operation terminates at the end of the data field.

If an overrun is detected during a read, search, or continue scan operation, it is signaled immediately. But if it is detected during a write operation, it is signaled at the end of the present record segment. Zeros are then used to fill out the segment operated on by the write command.

If a data check or bus-out parity error occurs, it is indicated at the end of the associated field, unless writing is being done. In that case, the operation continues to the end of the record segment.

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

The record overflow feature is not effective when a space count command is used, because the flag byte is spaced over and the segment is not detected as an overflow segment. Spacing over an overflow segment must be considered carefully by the programmer. Spacing does not occur automatically. For example, consider the following sequence.

```
Search ID (for the first segment of an overflow
record)
TIC *-8 (transfer back eight bytes to repeat
search command if unsuccessful)
Read CKD (Multi-track)
```

In this sequence, the read count, key, and data command does not read the next logical record in the cylinder, but starts reading in the count area of the second segment of the overflow record. Contrast this sequence with the following.

```
Search ID
TIC *-8 (transfer back eight bytes to repeat
search command if unsuccessful)
Read KD (skip and SLI bits on)
Read CKD
```

This sequence does read the count, key, and data areas of the next logical record.

Unusual Conditions: In addition to the checks provided in normal processing of any record, certain conditions can occur which are unique to overflow records. The commands stop immediately on detecting the following conditions:

1. Overflow to a Defective Track: Overflow incomplete and track condition check sense bits are set if an overflow occurs to a track which has been flagged as defective.
2. Overflow From an Alternate Track: Overflow incomplete and track condition check sense bits are set if an attempt is made to overflow from a track flagged as an alternate.
3. Overflow Violating a File Mask: Attempting to overflow by issuing a command in violation of a file mask sets overflow incomplete, file protected.
4. Overflow to a Track with Incorrect Head Number: Overflow incomplete and seek check sense bits are set if the head number compare is unequal during an overflow.

5. **Overflow Across an End-of-Cylinder Boundary:** Overflow incomplete and end-of-cylinder sense bits are set if an overflow is attempted across an end-of-cylinder boundary.

When one of these conditions occurs while processing an overflow record, unit check, channel end, or device end is signalled to the CSW at the completion of the operation. Additional information on the cause of the unit check in the CSW is contained in byte 5 of the sense bytes, and sense byte 1, bit seven is turned on to indicate overflow incomplete.

Sense Byte 5	Meaning of Setting
0000 0110	A read command was in progress.
0000 0101	A write command was in progress.
0010 0101	A search key and data equal was in progress, and the record is equal to the point of error.
0100 0101	A search key and data high was in progress, and the record is equal to the point of error.
0110 0101	A search key and data equal or high was in progress and the record is equal to the point of error.
0101 0101	A search key and data was in progress and the record is low or a search key and data equal was in progress and the record is high to the point of error.
0111 0101	A search key and data high or a search key and data equal or high was in progress and the record is high to the point of error.

24328

Continue Scan Commands (Record Overflow)

These commands (Figure 44) are used to assist in continuing a search key and data command that is interrupted during an overflow record. The continue scan commands are essentially search data commands that operate like a read data command, except that data is searched rather than read. The file mask affects these commands in the same way as it does the search key and data commands. These commands are developed by the use of sense byte 5.

If a data check is detected, the command is terminated at the end of the field in which the error occurs. Unit check, channel end, and device end are presented to the channel, and the appropriate sense bits (i.e., data-check-in-count-field and/or data check are set).

If an overrun is detected, it is signalled immediately (i.e., the operation is terminated), and unit check, channel end, and device end are presented to the channel.

No continue scan commands have chaining requirements.

CHANNEL PROGRAMS

The following channel programs are typical examples of how CCW's are arranged to format, read, and write records on the 2314 DASD. The examples given do not include the CPU program which would be used to initiate the channel program.

The format for each of the CCW's is as follows:

Com- mand Code	Data Address	Flags	Count
0	78	31 32	47 48
			63

25498

All numbers used are hexadecimal numbers.

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 0 and a data length of 8 bytes and 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
Write Home Address
Write Record Zero
Write CKD
Write CKD
Write CKD

Summary of Example 1: The first CCW in the channel program is a seek command. All seek commands transfer six bytes of data from main storage to the 2314. (Thus the byte count of six.) The first three bytes of the seek address are always 0's, the cylinder number (6A) is specified in the fourth byte, byte five is always 0, and byte six specifies the desired head (08 at 03ED).


Seek

07	0003E8	4000	0006
		B B C C H H	
	03E8	00 00 00 6A 00 08	


25499

The next command in the channel program is a set file mask command. The mask byte in this case (1100 0000 at address 03EE) permits all write and seek commands. The mask is reset to all 0's at the end of each chain of commands.

Set File Mask




Write Record 0



The write home address command creates the home address area on the track. The home address area is five bytes long (FCCHH). When formatting tracks, the flag byte is normally zero, the first cylinder byte is always zero, the second cylinder byte contains the cylinder number (6A), the first head byte is always zero, and the last head byte contains the read/write head number.

Note: 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 2314 for other write commands.

Write Home Address



Following the home address area is record 0. The write record 0 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 in the count field. In this example, the data address is at 07D0 and a byte count of 16 is specified.

Since the key length specified is 0, address 07D5 is coded 00 and no key area is written. The data length specified is eight bytes so addresses 07D6 and 07D7 are coded 0008 and the data in the following eight main storage locations is written in the data area.

Note that the byte count in the write record 0 command is sixteen and the 2314 requested sixteen bytes (eight for the count area and eight for data). Therefore no incorrect length error is generated.

The next command in the channel program is a write CKD command. Execution of this command causes an address marker, 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 field, to be written on the disk. In this example, the data address is at 0BB8 and a byte count of eight is specified.

As in the write record 0 command, the first byte written is the first cylinder byte. Main storage locations 0BB8 to 0BBF are coded with track number, record number, key length, and data length of record R1. Since the key length specified is six, a key area six bytes long will be created. The data length specified is 03E8 (1000 bytes). Although a byte count of eight is specified in the CCW, and the channel byte count will go to zero after eight bytes have been written, the 2314 is committed to writing a key area six bytes long and a data area 1000 bytes long. Therefore the 2314 inserts 0's into the applicable byte positions on the track until the 2314 byte count equals zero.


The difference in the channel byte count and the 2314 byte count will cause an incorrect length indication. Therefore, the SLI bit (bit 34) is on in the CCW.

In this example, six bytes of 0's will be recorded in the key area followed by three check bytes, one indicator byte, a gap, 1000 bytes of 0's, three more check bytes and one indicator byte. At a later time, valid data could be recorded in the key and data areas with the following CCW sequence:


- Search ID equal (R1)
- TIC*-8
- Write Key and Data
- Search ID equal (R2)
- TIC*-8
- Write Key and Data
- etc.

The next two commands in the channel program perform the same function as the write CKD (R1). The only differences are that the record numbers are 2 and 3, and the chain flag is not on in the last CCW.


Write CKD (R1)



Write CKD (R2)



Write CKD (R3)



To write valid data in records R1, R2, and R3, increase the byte count in the write CKD CCW's. The specified data from main storage is then re-recorded in the key and data areas of the records.

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


Summary of Example 2: The first command in this channel program is a seek command. The seek command transfers the track address to the 2314, moves the access mechanism, and selects the specified head.

Seek




After locating the proper track, it is necessary to find which record on the track is Frank Smith's. Since the disk is organized by key areas, a search key equal command is executed. Execution of this command causes the 2314 to search the key field after the first address marker it finds on cylinder 0C head 04. If the key is not equal to Frank Smith's man number, (main storage locations 07D0 to 07D5) the 2314 signals channel end and device end to the channel and the TIC command (back to the search key equal) is executed. Subsequent key areas are searched until Frank Smith's record is found. The 2314 then signals channel end, device end, and status modifier to the channel. The status modifier condition in the ending status byte causes the channel to skip the next command (TIC) in the chain, and execute the write data command.

Search Key Equal



TIC




X = Positions Ignored

25505

The write data command then writes 64 (100₁₀) bytes of data into Frank Smith's payroll record from main storage locations 0BB8 to 0C1C.

Write Data



If Frank Smith's payroll record had not been on cylinder 0C, head 04, or some other programming error had been made which would not have resulted in satisfying the search key equal, the program would loop between the search key equal and the TIC until every key on the track had been searched (index passed twice). The 2314 would then signal unit check to the channel. A subsequent sense command would indicate no record found.

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

1. The disk is organized by ID - no key areas.
2. Joe Brown's man number is 12345.
3. The data length of each record is 015E (350 bytes).
4. His policy number is in the data area.
5. The data set begins on cylinder 0A track 00.

Using the record capacity charts in Appendix A, it is known that fifteen 350 byte records can be written on one 2314 track. Since the disk is organized by IDs, (Joe Brown's = 12345) the track and record location can be determined by dividing the ID by the number of records per track. In this case.

$$\frac{12345}{15} = 823 \quad \text{Note: Add 1 to the remainder to establish the address of the specific record.}$$

Thus Joe Brown's ID is 823 tracks from the beginning of the data set. There is no remainder so the first record on track 823 is Joe Brown's.

The CCHHR for the seek command is then determined by converting the 823 tracks to cylinders and adding the results to the beginning of the data set.

	Cylinder	Track	Record	C	C	H	H	R
Starting Address:	10	00	0	00	0A	00	00	00
Displacement:*	41	03	1	00	29	00	03	01
Result:	51	03	1	00	33	00	03	01


*Determined by dividing 823 by 20.

The channel program used is:



Seek
 Search ID Equal
 TIC*-8
 Read Data

Summary of Example 3: The track and record ID of Joe Brown's record is computed, the channel program seeks to the track, searches for the record, and reads the data.

The first command in the channel program is a seek to cylinder 33 (decimal 51) to position the access mechanism for the search.




The search ID equal command causes the first ID encountered on the track to be compared to Joe Brown's ID. All unequal comparisons of IDs cause the 2314 to signal channel end and 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 2314 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.



X = Positions Ignored

25508


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 0D16.



If the search ID equal is not satisfied, and the index point is passed twice, unit check is set in the status byte. A subsequent sense command would

indicate no record found. The course of action would then be determined by the error recovery procedures (Figures 47, 48, and 49).

Continue Scan



Continue Scan Command Codes			
Command	Hexadecimal	Binary	Decimal
Continue Scan Equal	25	00100101	37
Continue Scan High	45	01000101	69
Continue Scan Equal or High	65	01100101	101
Continue Scan, No Compare	55	01010101	85
Continue Scan, Set Compare	75	01110101	117
Continue Scan, Set Compare	35	00110101	53


Continue Scan Command Codes, Multiple-Track			
Command	Hexadecimal	Binary	Decimal
Continue Scan Equal	A5	10100101	165
Continue Scan High	C5	11000101	197
Continue Scan Equal or High	E5	11100101	229
Continue Scan, No Compare	D5	11010101	213
Continue Scan, Set Compare	F5	11110101	245
Continue Scan, Set Compare	85	10110101	181

Specifies number of bytes in CPU storage (including "Mask Bytes" to be used as Search Argument). If CCW count equals key length, command is effectively a "Search Key"

Specifies CPU storage location which contains search argument. CPU storage contains:

Specified Location



Search Argument in CPU Storage



Key Area Byte 4 is not compared because CPU Storage is set to all 1's

Data Area Bytes 2 and 3 and 8-11 are not compared because of "Mask Bytes" in CPU storage. CCW count of 20 limits search to 20 position span.

This example assumes CCW count 20



Key and Data Areas From Storage Unit

20615

Figure 44. Continue Scan Commands

TWO-CHANNEL SWITCH SPECIAL FEATURE

The Two-Channel Switch feature provides the ability for the 2314 to be shared by two channels and also allows individual 2314 module locations (i.e., any of the on-line modules) to be reserved for the exclusive use of either of the channels. The two channels can be attached to the same CPU (for the model 40 and up) or to different CPU's. Channel switching and device reservation in the 2314 are controlled by the channel program.

Channel Selection Switch

This program-controlled switch has three positions: channel A, channel B, and neutral. With the switch in the neutral position, the 2314 is selected by the first channel to complete the selection sequence. Once the 2314 is selected by a channel, it remains connected to that channel until channel end is presented for that operation or chain of operations. The channel selection switch then returns to neutral unless there is a unit check condition and the sense information has not been read.

If channel A, for example*, attempts to select the 2314 while the 2314 is selected to channel B, the 2314 responds to channel A with control unit busy (status modifier plus busy status bits). This signal, in turn, causes the 2314 to attempt to present to channel A a status byte containing control unit end after the channel selection switch returns to the neutral position. The address byte associated with this status condition is the base address of the 2314 on that channel. This control unit busy condition may occur on any attempt to select the 2314, including initial program load, test I/O, and all initial commands of a CCW chain.

Device Reservation

When a module is busy with channel A (for any reason) or reserved to channel A, any command from channel B addressed to that module is rejected with a busy indication in the initial status byte. After the reservation has been terminated, the busy indication causes the 2314 to attempt to present to channel B a status byte containing device end. The device address byte presented with this ending status byte is the same as that associated with the busy status byte.

If device end status results from the completion of any command that causes mechanical motion of an access mechanism, that status is presented to the channel that issued the command.

A module that generates device end status when it changes from the not-ready to the ready state causes the 2314 to present device-end status to both channels. Each channel must accept this device end before it can use the device.

Device Reserve Command

With the two-channel switch feature, a device reserve command (Figure 45) causes the addressed module to be reserved to the channel issuing the command. The module then remains reserved to the same channel until that channel executes a device release command addressed to the specific device, or until a system reset occurs on that channel.

This command also transfers sense information (same as that which is transferred by the sense command**) to the channel.

A device reserve command is rejected with a busy indication in the CSW if the 2314 is busy. However, a device reserve command is executed regardless of any abnormal file status condition, such as off-line or unsafe.

A device reserve command is rejected when a set file mask command precedes it in the same command chain. The unit check bit in the CSW is set when the command is rejected, and the command reject and invalid sequence bits are set to indicate the conditions which caused the unit check.

Device Release Command

With the two-channel switch feature, a device release command (Figure 46) terminates the reservation of the addressed device to the channel. This command is rejected with a busy indication in the CSW if the device is busy. However, a device release command is executed regardless of any abnormal file status condition, such as off-line or unsafe.

A device is normally reserved to a particular channel whenever that channel executes a device reserve command. The device remains reserved

*Channels A and B could be reversed in this explanation.

**CCW should have a data address and count the same as a sense command.

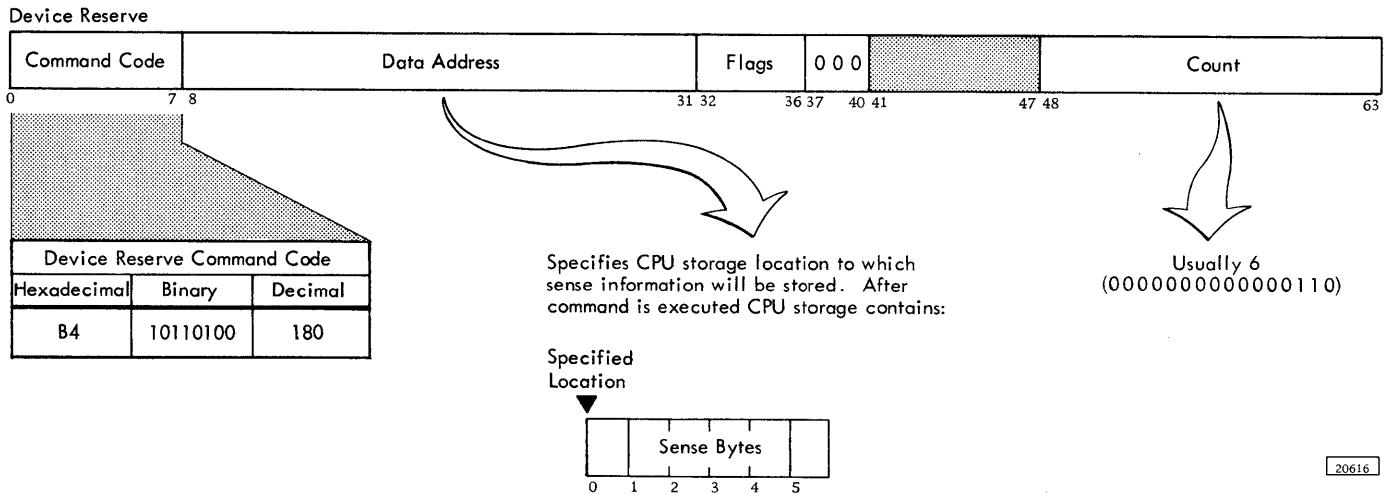


Figure 45. Device Reserve Command




Figure 46. Device Release Command

to the same channel until that channel causes the 2314 to execute a device release command, or until a system reset occurs on that channel.

This command also transfers sense information (same as that which is transferred by the sense command*) to the channel.

A device release command is rejected when a set file mask command precedes it in the same command chain. The unit check bit in the CSW is set when the command is rejected, and the command reject and invalid sequence bits in the sense bytes are set to indicate the conditions which caused the unit check.

Addressing with Two-Channel Switch

The base address (four high-order bits) of the 2314 on one channel is independent of the base address of the 2314 on the other channel. However, the four low-order address bits for any attached module must be the same on both channels. Bit position 4 of all device addresses must be zero.

Resets with Two-Channel Switch

A system reset occurs when the CPU system reset key is pressed. A system reset causes all 2314 reservations and status conditions related to the resetting system to be reset. Reservations related to another system are not affected.

A system reset occurring when the channel selection switch is neutral causes a machine reset. A selective reset (i.e., malfunction reset) has no effect on device reservations or status.

Untagged-Tagged Switch

The 2314 provides an operator-accessible switch to control the tagging of unsolicited interrupts. This switch is located on a panel behind the front cover of the single-drive module. The operation of this switch is described in the 2844 Auxiliary Storage Control portion of this manual.

Usage Meter with Two-Channel Switch

A single usage meter records processing time in the 2314. However, a separate enable/disable switch is provided on the 2314 for each channel. These switches provide separate disabling of 2314 operations with each channel. Disabling is normally done by the Customer Engineer for servicing. In order to disable 2314 operations with a channel, the following conditions must exist:

1. The enable switch for that channel must be in the disable position.
2. The channel selection switch (on the 2314) must not select that channel.
3. The CPU must be in the stop or wait state.

To enable 2314 operations with a channel, the following conditions must exist.

1. The enable switch for that channel must be in the enable position.
2. The CPU must be in the stop or wait state.

If the CPU enters the stop or wait state when the 2314 contains outstanding status, the status may be unavailable if the enable/disable switch associated with the channel has previously been set to disable.

Power Control with Two-Channel Switch

If either channel provides power, the 2314 turns on. The 2314 turns off only if both channels indicate power off.

When a 2314 with the two-channel switch feature is connected in a multi-system configuration, a special emergency power off control box must be installed in the multi-system, so that when the emergency power off switch in either system is pulled, both systems will execute the emergency power off sequence.

When a 2314 with the two-channel switch feature is shared by two CPUs and one CPU is to be powered down, the corresponding enable/disable switch on the single drive module must be placed in the disable position before turning off the CPU power.

REMOTE SWITCH ATTACHMENT SPECIAL FEATURE

This feature consists of locating the enable switches (of the 2314 with the two-channel switch feature) on an IBM 2167 Configuration Unit instead of on the 2314. The 2167 is used in System/360 Model 67 when that system has multiple processing units. The enable switches cause a single usage meter to record processing time.

When the 2167 is used, the enable switches for the two-channel switch feature are not on the 2314. Manual enabling or disabling of the 2314 with either of the two channels with which it can operate is under control of the switches on the 2167.

*CCW should have a data address and count the same as a sense command.

Program control of 2314 disk modules is accomplished in the same manner through either the basic 2314 control or the IBM 2844 Auxiliary Storage Control. Therefore, the descriptions of control also apply to the 2844. Additional information pertaining to control of the 2314/2844 complex is presented here.

Attachment and Addressing (2314/2844)

One 2844 can be attached to the system in conjunction with a 2314 model 1 or an A-series facility with eight on-line and one service modules. The 2844 and the basic 2314 control can be attached to the same selector channel or, on the model 50 and up, to two different selector channels. If two channels are used, the 2844 is attached only to one, and the basic 2314 control is attached only to the other (except when the two-channel switch feature is used).

The 2314/2844 complex appears to the controlling programs as two separate control units, each of which can control an operation for any of the eight 2314 on-line disk modules. Operation with any one of the modules is initiated by a start I/O instruction that specifies the address of the channel, control unit, and a particular module. The command specifying the operation is then sent by the channel, to the addressed control. Simultaneously, another start I/O may be used to initiate an operation with one of the other on-line modules through the other control.

The two channels, with which the 2314/2844 complex communicates, can be attached to the same or to two different systems. (Note, however, that the 2314 can be operated -- through either the basic control or the 2844 -- only on selector channel 1 in IBM System/360 Model 30, and only on one of the two selector channels in System/360 Model 40.) The basic 2314 and the 2844 control can have the same or different unit addresses, which are assigned at installation time and internally wired into the control unit by an IBM Customer Engineer.

The two-channel switch special feature can be installed in either control or in both the basic 2314 control and in the 2844 control. (Note, however, that the 2314/2844 complex can communicate with only one selector channel at a given time on System/360 Model 40. Therefore, the two-channel switch for the 2314 cannot be set to one model 40 selector channel while the 2844 is set to the other selector channel.) If the two-channel switch is installed in both controls, operations with any 2314 module can be initiated through any of the four (maximum) separ-

ate channels to which the 2314/2844 complex is attached. Any combination of these four channels can be attached to the same or to different systems, provided the channel restrictions of the system permit.

A two-channel switch for either control can be used to switch communications for that control between two selector channels of the System/360 Models 40 and above, but not between two selector channels on the same System/360 Model 30.

If the two-channel switch feature is installed in either or in both controls and the 2314/2844 complex is attached to a System/360 Model 67, the remote switch attachment special feature can be installed. When this feature is used, the enable switches for the storage control(s) with the two-channel switch are placed on an IBM 2167 Configuration Unit, and the switches on the unit(s) that has a two-channel switch feature become inactive. Manual enabling or disabling of the control unit(s) is then under control of the switches on the 2167. The 2167 is used in System/360 Model 67 when that system has multiple processing units.

Program Control of 2314/2844

The 2844 is operated by means of the same channel commands specified for the basic 2314. Operation of the 2844 requires no new channel commands except that the reserve and release commands, for reserving and releasing modules to either control (2314 or 2844), can be used in the same manner as they are used with a two-channel switch feature even though no two-channel switch is used in the complex. (See Two-Channel Switch Special Feature in this manual for a description of the reserve and release commands.)

In the 2844, the file scan and record overflow standard features are controlled in the same way as in the basic 2314. Also, a 2844 with the two-channel switch special feature is controlled in the same manner as is the basic 2314 with a two-channel switch.

Program control of the eight 2314 on-line disk modules is achieved in the following manner. An I/O operation is initiated in the normal manner.

1. The start I/O instruction contains the channel and unit address of the desired control (basic 2314 or 2844) and disk module.

2. The channel fetches the CAW (channel address word) that was previously loaded into main storage locations 72 through 75.
3. The CAW designates the main storage address of the first CCW to be executed by the channel program.

If the controls are on separate channels, the control on the addressed channel (specified in the start I/O instruction) is selected.

In any case, the eight-bit unit address specifies:

<u>Bit Positions of Unit Address</u>	<u>Specifies</u>	<u>Setting</u>
0	Always 0.
1-3	A control.	000 to 111. (0 to 7 decimal).
4	Always 0
5-7	One of the eight 2314 on-line modules.	(0 to 7 decimal).

If no status conditions are outstanding for the addressed module and if that module is not already busy with the other control or performing a mechanical access, the specified operation is initiated.

Status Presentation (2314/2844)

When a channel addresses a module (through either of the controls) for command execution, or when a module is momentarily connected to a channel for establishing interrupt priority or clearing interrupts, the channel is presented with busy status if:

1. The addressed module is currently selected by another channel through the other control (i. e., the control not addressed by the command), or
2. The addressed module is currently reserved to another channel, or
3. The access mechanism of the addressed module is performing mechanical motion.

After the busy condition is terminated, the addressed control unit attempts to send device-end status, which indicates that the module is no longer busy, to the channel that received the busy status. This device-end status must be accepted by that channel before it can use the module that was busy with the other channel.

When a channel issues a command that results in access motion for the addressed module, then the device end for the operation is presented only to the channel that issued the command.

Other status combinations are presented to a channel in the normal manner (i. e., in the

same manner as in a basic 2314 without the 2844).

Untagged-Tagged Toggle Switch

When the 2844 is attached to the 2314, an untagged-tagged switch is placed on the front of the ninth disk module in the 2314. (Also, a separate untagged-tagged switch is provided on the ninth module for each two-channel switch feature used in the 2314/2844 complex.) The manner of presentation of a device-end resulting from making the 2314/2844 ready is determined by the setting of this switch. Such a not-ready-to-ready-sequence device end occurs after a disk pack change is completed and the affected module is returned to an on-line condition.

When the switch is in the tagged position, device end for the not-ready-to-ready-sequence is presented to all channels connected to the 2314/2844. After the device end is accepted, the module can be addressed for command execution by the channel that accepted the device end. Before any other channel can use the module, it must also receive the not-ready-to-ready-sequence device end.

If the toggle switch is in the untagged position, the 2314/2844 attempts to present the not-ready-to-ready-sequence device end to all channels connected to the 2314/2844 complex. This device end is cleared as soon as it is accepted by any channel, and no further attempt is made to present the device end to any other channel.

The switch is set to the tagged position for most applications. However, if a single supervisor program controls all 2314/2844 operations, and if the not-ready-to-ready-sequence device end interruption can be cleared by any channel, the switch can be set to the untagged position.

If the untagged setting is used, the first channel that accepts the device end clears the affected module for operation with any channel to which the 2314/2844 complex is attached. Suppose, for example, that a disk pack is changed and the module made ready. The not-ready-to-ready-sequence device end must be accepted before the module can be addressed for a new operation. (Assume that the basic 2314 control is attached to channel A and the 2844 is attached to channel B.) If, at the time of the interruption, channel A is masked against interruptions and channel B is not so masked, channel B may immediately accept the interruption (assuming channel B is not engaged in another operation). This action clears the device-end interruption for both channel A and channel B. If channel A subsequently becomes unmasked, the interruption is not processed a second time. Channel A can then immediately initiate an I/O operation with the same module (assuming that the module is operational and not already busy with channel B).

Caution

If the switch is set to the untagged position, the programming environment must be such that the device-end interruption can logically be cleared by any channel to which the 2314/2844 complex is attached. If separate supervisor programs are used and the device end interruption must be indicated separately to each supervisor, then the switch must be set to the tagged position.

Timing and Control (2314/2844)

Use of the 2844 does not change the time required for:

1. Command chaining functions for the 2314.
2. The data rate (312 kilobytes per second) of the 2314.
3. Mechanical access for modules in the 2314.
4. Signal responses between the 2314 and a channel.

A tie-breaking circuit is provided for the situation in which the basic 2314 control and the 2844 control request service at the same time for the same module.

Power Control (2314/2844)

The IBM 2844 Auxiliary Storage Control contains its own power supplies. The disk storage modules receive their required DC voltages from either the basic 2314 control or the 2844 control, depending upon the 2844 power switch setting.

Power turns on for the 2314/2844 whenever any one of the channels, to which the 2314/2844 is attached, indicates power on. All channels to which the 2314/2844 is attached must then indicate power-off before power is removed from the 2314/2844.

When a 2314/2844 with the two-channel switch feature is connected in a multi-system configuration, a special emergency power off control box must be installed in the multi-system, so that when the emergency power off switch in either system is

pulled, both systems will execute the emergency power off sequence.

A manually operated power switch is provided in the 2844 for the Customer Engineer to control disk module power in the completely duplexed power system in the 2314/2844. Either control unit may be powered down and the system processing continued on the active control unit. This continuation is accomplished by ensuring that the disk module power is switched to the active control unit. Each system attached to the 2314/2844 must be in the halt or wait state before power is switched at the 2314/2844.

The 2314/2844 will respond to any attempted selection by an attached channel, with control unit busy status when the Customer Engineer is performing a power switch. At completion of the power switch, either control which presented control unit busy will attempt to interrupt and present a status byte containing control unit end.

When a 2844 with the two-channel switch feature is shared by two CPUs and one CPU is to be powered down, the corresponding enable/disable switch on the single drive module must be placed in the disable position before turning off the CPU power.

Configuration Control

The 2314/2844 provides an enable/disable switch for each channel to which it is attached. (See Enable/Disable Switch.) Whenever a switch is placed in the disable position and the CPU associated with that switch goes into the wait state, the 2314/2844 thereafter appears off-line to that CPU. When the enable/disable switch is returned to the enable position and the CPU associated with that switch goes into the wait state, the 2314/2844 returns to an on-line state for the CPU involved.

Usage Meter

A single usage meter is provided for the 2314/2844 complex. The meter is located with the enable/disable switches on the single file module for 2314 model 1 or on the control unit for an A-series facility. The meter operates whenever the 2314/2844 is enabled and operating with any of the attached channels.

SUGGESTED ERROR RECOVERY PROCEDURES

The suggested error recovery procedures are organized in the following manner:

1. The error condition table (Figure 47) defines the errors that can occur and references the action table.
2. The action table (Figure 48) indicates the appropriate procedure to be followed when an error occurs.

3. Error messages (Figure 49) that relate to the action table are listed separately.

If you are using an IBM programming system, refer to the appropriate programming systems reference library (SRL) publications for information about error recovery procedures used by the particular programming system. Programming SRL's are listed by Order No. and are briefly described in IBM System/360 Bibliography, Order No. GA22-6822.

Sense Bit(s) and Name(s)	Explanation	General Class	Action
Byte 0, Bit 0 Command Reject	The 2314 has received an invalid command code.	Program Error	1
Byte 0, Bit 0 Command Reject and and Byte 0, Bit 7 Seek Check	An invalid seek address was received by the 2314.	Program Error	1
Byte 0, Bit 0 Command Reject and Byte 1, Bit 3 Invalid Sequence	The 2314 has received an invalid sequence of commands.	Program Error	1
Byte 0, Bit 0 Command Reject and Byte 1, Bit 5 File Protected	The 2314 has received a command that violates the inhibit write portion of the file mask.	Program Error	1
Byte 0, Bit 1 Intervention Required	The specified file is: 1. Not on line 2. Not available for use due to a cover interlock open, file motor off, etc.	Equipment Error	2
Byte 0, Bit 2 Bus-Out Parity	A parity error has been detected during the transfer of information from the channel to the 2314.	Equipment Error	2
Byte 0, Bit 3 Equipment Check and Byte 2, Bit 0 Unsafe	An unusual condition has been detected in the 2314. The condition is indicated in sense byte 2. A file malfunction has been detected.	Equipment Error	2
Byte 0, Bit 4 Data Check	A data error has been detected by the 2314 in the data field received from the file.	Equipment Error	5
Byte 1, Bit 6 Missing address Marker	The 2314 has detected: 1. Two index points without an intervening address marker, or 2. Two successive count fields with equal bit settings in bit 0 of the flag bytes.	Equipment Error	5
Byte 0, Bit 5 Overrun	A byte from the 2314 was not accepted (during reading) by the channel before the 2314 was ready with the next byte; or a byte was not available from the channel to the 2314 during writing; or a chained CCW was received too late to be properly executed.	Equipment Error	3
Byte 0, Bit 6 Track Condition Check	1. A search, read, or write was attempted on a flagged defective track. (Track condition-bit 6=1.) <u>Exception: HA and RQ</u> 2. Command chaining and multiple track mode signals indicate that operations from an alternate track are to continue on the next higher ordered track.	Program Error Program Error	6 6
Byte 0, Bit 6 Track Condition Check, and Byte 1, Bit 7 Overflow Incomplete	Overflow has been attempted, either to a flagged defective track or from an alternate track.	Program Error	6
Byte 0, Bit 7 Seek Check	File has been unable to complete a seek due to: 1. Access mechanism has failed to reposition properly. 2. Home address compare fails after automatic head switching on a multiple-track operation.	Equipment Error	4
Byte 0, Bit 7 Seek Check and and Byte 1, Bit 7 Overflow Incomplete	Overflow to a wrong track has occurred. Head number compare failed.	Equipment Error	4
Byte 1, Bit 0 Data Check in Count Field, and Byte 0, Bit 4 Data Check	A data error has been detected by the 2314 in the count field received from the file.	Equipment Error	5
Byte 1, Bit 1 Track Overrun	Writing has not been completed by the time index point is detected.	Program Error	1
Byte 1, Bit 2 End of Cylinder	An end-of-cylinder was detected before the CCW command chain was completed in a multiple track operation.	Program Error	1
Byte 1, Bit 4 No Record Found, and Byte 1, Bit 6 is <u>not</u> Present	Two index points were detected while executing a chain of CCW's with no intervening read or write operation (on the data field of any record) or read HA or read R0 CCW. (Note: This could be an expected condition on search command chains.)	Program Error	7
Byte 1, Bit 5 File Protect	The 2314 has received a seek multi-track read, or multi-track search that violates the file mask.	Program Error	1
Byte 1, Bit 5 File Protect and Byte 1, Bit 7 Overflow Incomplete.	Overflow was attempted to a file-protected boundary.	Program Error	1

Figure 47. Error Condition Table

11534

Action	Procedure
1	Exit with program error indication
2	<ul style="list-style-type: none"> a. Repeat the original sequence once. b. If the error condition still occurs, print message 1 for operator and/or Customer Engineer notification. This is considered an uncorrectable error and the procedure should be dependent upon the application.
3	<ul style="list-style-type: none"> a. Repeat the original sequence 10 times if the error persists. b. After 10 unsuccessful retries, print message 1 for operator and/or Customer Engineer notification.
4	<ul style="list-style-type: none"> a. Issue a recalibrate command. b. Seek for the original address. c. Repeat the original sequence 10 times if the error persists. d. Print message 1 and exit with data error indication.
5	<ul style="list-style-type: none"> a. Repeat the original sequence 10 times if the error persists. b. After 10 unsuccessful retries, the data error is said to be "hard". At this point the error recovery operations are dependent upon the user's application procedures. c. Print message 1 and exit with data error indication and other information that depends upon the user's procedures.
6	<ul style="list-style-type: none"> a. If this is an alternate track, use the address of the defective track plus one in a seek command. (This will be found in the ID field of the R0 count area.) Resume the operation after searching to the desired track position. b. If this is a defective track, use the address of the alternate track in a seek command. (This will be found in the ID field of the R0 record.) Resume the operation after searching to the desired track position.
7	<ul style="list-style-type: none"> a. Issue a read home address to verify that the correct track has been reached. (The correct cylinder will suffice on multiple-track operations.) b. If the correct cylinder (and/or track) is found, perform action 1. c. If the incorrect cylinder (and/or track) is found, perform action 4.

11541

Figure 48. Suggested-Action Table

<p>The following two error messages should be included in the operating environment of all 2314 users:</p>
<p>Message 1 (should be printed on all uncorrectable errors.)</p> <ul style="list-style-type: none"> A. Message Code b. Error type -- read, write, or control. c. Module designation, cylinder number, and head number (i.e., device addressed and seek address). d. Channel designation. e. Status and sense bytes sent to CPU.
<p>Message 2 (Should be printed periodically, upon completion of an application run or in response to operator request.)</p> <ul style="list-style-type: none"> a. Unit designation b. Number of entries into error routine. c. Number of uncorrectable errors.

11542

Figure 49. Suggested Messages

DISK PACK HANDLING

For complete details on disk pack handling, refer to IBM Disk Pack and Cartridge Handling Procedures, Order No. GA26-5756.

Each IBM 2316 Disk Pack is protected in transit by a special container.

When the disk pack is received, examine the container closely. If its condition is acceptable, remove the disk pack and store it. Keep the container; you may need it later.

If the container or its contents show any unusual shipping damage, do not use the disk pack. Retain the damaged container and disk pack in its "as received" condition and notify an IBM Customer Engineer immediately.

Disk packs have been designed to be easily transported from location to location. For best shipping results:

1. Be sure the pack is secure in its two-piece cover.
2. Use only the specially designed IBM shipping container. If the original container is worn or damaged, a new carton may be ordered from the local IBM office.

Handle a disk pack only with its cover on. If the disk pack is accidentally dropped, or receives a sharp impact of any kind, have it inspected by an IBM Customer Engineer before you use it.

Disk Pack Labeling

For identification purposes, IBM recommended labels can be placed on the disk pack center hub. Labels in this location can be read through the transparent disk cover.

The following procedures should be used when labeling disk packs:

1. Use adhesive-backed labels that can be applied and removed easily.
2. Use a writing implement, such as a pen or felt-tip marker, that does not produce a loose residue. Do not use a lead pencil.
3. Write on the label before it is applied to the disk pack.
4. Place the label only on the center hub, not on the disk pack cover or top disk surface.

5. Use a new label if changes to the label information are necessary. Never erase a label that is on a disk pack. Microscopic eraser particles can damage disk surfaces and read/write heads.
6. Removal or placement of labels can be done only when a disk pack is in place in a 2314 module location. Do not attempt to remove the top cover when the disk pack is out of the 2314.

Disk Pack Cover

A two-piece plastic cover protects the disk pack from dust and moisture when the disk pack is not in the 2314. If the cover is damaged, it may lose its protective capability and should be replaced.

On the top section of the plastic cover is a handle that is used for carrying the pack. A self-locking device in the handle permits removal of the disk pack from the 2314 only if the top part of the cover is first locked on the disk pack. The plastic cover should never be left on a disk pack that is mounted in a 2314 disk module location.

The lower half of the plastic cover should be secured to any 2316 that is not mounted in the 2314.

Disk Pack Storage

To assure maximum disk pack life and reliability:

1. Store the disk packs flat, not on edge.
2. Each pack should rest on a shelf, not on another disk pack.
3. Store in a clean, enclosed metal cabinet or a similar fire-resistant container, never in direct sunlight.
4. Store disk packs in a machine-room atmosphere (60° F to 90° F, 10% to 80% humidity).
5. If a disk pack must be stored in a different environment, allow two hours for adjustment to machine room atmosphere before use.
6. Always be sure that both the top and bottom plastic covers are on a disk pack whenever it is out of a 2314 module location.

Error Free Disk Packs

All 2316 disk packs manufactured after January 1, 1969 are error free using the following utilities as a test:

OS-"DASDI"-360S-UT-507

The error free disk packs can be recognized by a new flagged track label.

OPERATOR PROCEDURES

Disk Pack Loading

1. Open the module drawer and remove the mounted disk pack, if any, using the procedure described in Disk Pack Unloading.
2. Remove the bottom cover of the disk pack to be mounted by turning the bottom locking knob.
3. It may be necessary to move (gently) the disk brush assembly out of the way if it extends into the disk pack area (Figure 50).
4. Place the 2316 disk pack with top cover, on the spindle.

NOTE:

The disks are not shown in this illustration so that the Disk Brush assembly can be shown. During operation, the brushes extend into the Disk Pack and sweep the disk surface.

5. Turn the top cover in the clockwise direction until firm resistance is met.
6. Lift the top cover from the disk pack.

Caution

Do not close a module drawer if the top cover is still on the disk pack.

7. Slide the module drawer back into the 2314. To assure that the 2314 drawer has been properly closed after the mounting of a pack, an attempt to open the drawer while maintaining a downward pressure on the handle should be made. If properly closed, it will not open without an upward lift on the handle.
8. Check to make sure that the correct module identifier plug is in the module's receptacle.
9. Place the start/stop switch for the module in the start position to return the module to operation. The module returns to operational status




Figure 50. Disk Brush Assembly

11426

about 60 seconds after the start/stop switch is placed in the start position.

Disk Pack Unloading

1. Place the start/stop switch of the affected module in the stop position. (Note that power is removed from the disk drive motor when the module drawer is opened, even if the start/stop switch is in the start position.)
2. Wait about 15 seconds and then open the module drawer.
3. Be sure that disk pack motion has stopped before you perform any of the following steps.
4. Separate the top and bottom disk pack covers.
5. If the brush assembly extends into the disk pack, gently move it out of the way (Figure 50) and place the disk pack top cover on the disk pack.
6. Turn the top cover counterclockwise at least two full turns.
7. Lift the top cover, now containing the disk pack, from the spindle.
8. Firmly fasten the bottom cover to the disk pack.
9. Be sure that the module drawer is not left open if you do not intend to load a new pack immediately.
10. Store the disk pack in a clean cabinet or on a clean shelf. (See Disk Pack Storage in this manual.)

Changing a Module Address

The address of a file module is normally changed only when a module requires maintenance. To change the address of any module:

1. Make sure that the program(s) controlling the 2314 is in a wait state or that the existing condition allows removal of module identifier.
2. Remove the module identifier plug from the affected modules and perform any necessary pack changes.
3. Place the module identifier plugs in the appropriate receptacles.
4. Program control is now in effect.

Always ensure that the module plug is completely in or completely out. If the module plug is left part way out for an extended period of time it may cause damage to the 2314.

OPERATOR CONTROLS AND INDICATORS

Module Identifier Plugs

Each of the file modules has a module identifier plug receptacle. If a module does not have a module

identifier plug in this receptacle, then that module is not logically attached to the control unit and cannot be controlled by the program. If a module identifier plug is placed in the module's receptacle, the access mechanism of that module is positioned at cylinder 000 and a not-ready-to-ready sequence device end is then presented to the channel from the control unit. The affected module is then in the same physical and logical state as if its disk pack had been changed and then the module made ready, or in the same state as when the system is initially powered on.

Each removable module identifier plug has a unique address which is one of the values in the range 0 through 7. The printed address on each identifier plug is in decimal notation and represents the four low-order bits of the device address. The module identifier plug for the service drive is labeled "S". Although the A-series facility may not have a service module, a service module identifier plug (labeled "S") is provided. To change the address of a module, remove its module identifier plug and insert the plug that has the desired address. Always ensure that the module plug is completely in or completely out; not left part way out. (See Changing a Module Address.) Local installation practices must determine which 2316 disk pack is to be mounted at each specific module address.

Start/Stop Toggle Switches

Each module has a start/stop toggle switch that controls application of voltage to the associated module. If a switch is set to the stop position, the associated module cannot be started from either control (2844 or 2314) or from the system console. Also, placing a start/stop switch in the stop position causes the access mechanism for the affected module to retract from the disk pack and causes power to be removed from the disk drive motor. Disk pack motion then stops in about 15 seconds. If the switch is then placed in the start position and the module drawer is closed with a disk pack mounted, the disk drive motor turns on and an automatic seek to cylinder 000 is performed. The time for the individual module to become operational is then about 60 seconds.

Select Lock Indicators

This red indicator (separate indicator on each module) lights when an equipment malfunction occurs for the associated module. If this light is on, call the IBM Customer Engineer.

Ready Indicator (Labeled A-J)

This is a green lens with the physical device location (A-J) imprinted on it. This indicator will turn on following completion of the head loading sequence. This indicator turns off during the following times:

1. When a seek operation is being performed - this is a normal condition.
2. When the access mechanism has not detented properly - this is an abnormal condition. In this case unit check status and seek check (sense byte 0, bit 7) indications are made available to the channel.

Notify the Customer Engineer if the ready indicator for any otherwise operational module turns off.

Enable/Disable Switch

When the CPU is in the wait state (and not servicing requests for the 2314) or stop state, this switch enables or disables the communication of the 2314 with the CPU. It also enables or disables the equipment usage meter.

If the CPU is running when the switch setting is changed, the 2314 operating state is not changed until the CPU is placed in the stop or wait state.

An additional enable/disable switch is provided if the 2844 Auxiliary Storage Control is attached. If the two-channel switch special feature is installed in either (or both) the 2314 or the 2844, each channel has its own enable/disable switch.

For the 2314 model 1, the switch(es) are located on the single drive module. For the 2314 A-series facilities, the switches are located on the 2314 control unit.

APPENDIX A. RECORD CAPACITIES

Record Length (Bytes)		Track Capacity in Records	Cylinder Capacity in Records	Module Capacity in Records	Facility Capacity in Records	Trans. Time in MS/Record	
Lower Length	Upper Length					Minimum	Maximum
3477	7249	1	20	4000	32000	11.13	23.20
2255	3476	2	40	8000	64000	7.22	11.13
1650	2254	3	60	12000	96000	5.28	7.22
1289	1649	4	80	16000	128000	4.13	5.28
1050	1288	5	100	20000	160000	3.36	4.13
879	1049	6	120	24000	192000	2.82	3.36
751	878	7	140	28000	224000	2.41	2.81
651	750	8	160	32000	256000	2.09	2.40
572	650	9	180	36000	288000	1.83	2.08
507	571	10	200	40000	320000	1.63	1.83
453	506	11	220	44000	352000	1.45	1.62
408	452	12	240	48000	384000	1.31	1.45
369	407	13	260	52000	416000	1.18	1.31
334	368	14	280	56000	448000	1.07	1.18
305	333	15	300	60000	480000	0.98	1.07
278	304	16	320	64000	512000	0.89	0.98
255	277	17	340	68000	544000	0.82	0.89
234	254	18	360	72000	576000	0.75	0.82
216	233	19	380	76000	608000	0.70	0.75
199	215	20	400	80000	640000	0.64	0.69

11535 A

Figure A-1. Record Capacities when Keys are Used

(This page intentionally left blank.)

Record Length (Bytes)		Track Capacity in Records	Cylinder Capacity in Records	Module Capacity in Records	Facility Capacity in Records	Trans. Time in MS/Record	
Lower Length	to Upper Length					Minimum	Maximum
3522	7294	1	20	4000	32000	11.27	23.34
2299	3521	2	40	8000	64000	7.36	11.27
1694	2298	3	60	12000	96000	5.42	7.36
1333	1693	4	80	16000	128000	4.27	5.42
1093	1332	5	100	20000	160000	3.50	4.27
922	1092	6	120	24000	192000	2.95	3.50
794	921	7	140	28000	224000	2.54	2.95
695	793	8	160	32000	256000	2.23	2.54
616	694	9	180	36000	288000	1.98	2.22
551	615	10	200	40000	320000	1.77	1.97
497	550	11	220	44000	352000	1.59	1.76
451	496	12	240	48000	384000	1.45	1.59
412	450	13	260	52000	416000	1.32	1.44
378	411	14	280	56000	448000	1.21	1.32
348	377	15	300	60000	480000	1.12	1.21
322	347	16	320	64000	512000	1.03	1.11
299	321	17	340	68000	544000	0.96	1.03
277	298	18	360	72000	576000	0.89	0.96
259	276	19	380	76000	608000	0.83	0.89
242	258	20	400	80000	640000	0.78	0.83

11536A

Figure A-2. Record Capacities when Keys are Not Used

APPENDIX B. COMMAND SUMMARY

COMMAND		COMMAND CODE						DATA ADDRESS	COUNT
		Multiple Track Off			Multiple Track On (if Applicable)				
		Decimal	Hexadecimal	Binary	Decimal	Hexadecimal	Binary		
Control	No Op	0 3	0 3	0000 0011				X	Not Zero
	Release*	1 4 8	9 4	1001 0100				X	Not Zero
	Reserve*	1 8 0	B 4	1011 0100				X	Not Zero
	Recalibrate	1 9	1 3	0001 0011				X	Not Zero
	Seek	0 7	0 7	0000 0111				} CPU storage location of seek address	6
	Seek Cylinder	1 1	0 B	0000 1011					6
	Seek Head	2 7	1 B	0001 1011				} CPU storage location to which 6 sense bytes are sent.	6
	Sense	0 4	0 4	0000 0100					6
	Set File Mask	3 1	1 F	0001 1111				} CPU storage location of mask byte.	1
	Space Count	1 5	0 F	0000 1111				X	Not Zero
Transfer in Channel	X 8	X 8	XXXX 1000				} CPU storage location of next CCW (Must be divisible by 8.)	X	
Search	Home Address Equal	5 7	3 9	0011 1001	1 8 5	B 9	1011 1001	} CPU storage location of search argument	4 (usually)
	Identifier Equal	4 9	3 1	0011 0001	1 7 7	B 1	1011 0001		5 (usually)
	Identifier High	8 1	5 1	0101 0001	2 0 9	D 1	1101 0001		5 (usually)
	Identifier Equal or High	1 1 3	7 1	0111 0001	2 4 1	F 1	1111 0001		5 (usually)
	Key Equal	4 1	2 9	0010 1001	1 6 9	A 9	1010 1001		From 1 to 255
	Key High	7 3	4 9	0100 1001	2 0 1	C 9	1100 1001		From 1 to 255
	Key Equal or High	1 0 5	6 9	0110 1001	2 3 3	E 9	1110 1001		From 1 to 255
	Key and Data Equal	4 5	2 D	0010 1101	1 7 3	A D	1010 1101		} Number of bytes (including mask bytes) in search argument
	Key and Data High	7 7	4 D	0100 1101	2 0 5	C D	1100 1101		
Key and Data Equal or High	1 0 9	6 D	0110 1101	2 3 7	E D	1110 1101			
Read	Home Address	2 6	1 A	0001 1010	1 5 4	9 A	1001 1010	} CPU storage location to which areas read will be transferred	5
	Count	1 8	1 2	0001 0010	1 4 6	9 2	1001 0010		8
	Record R0	2 2	1 6	0001 0110	1 5 0	9 6	1001 0110		Number of bytes to be transferred
	Data	0 6	0 6	0000 0110	1 3 4	8 6	1000 0110		Number of bytes to be transferred
	Key and Data	1 4	0 E	0000 1110	1 4 2	8 E	1000 1110		Number of bytes to be transferred
	Count, Key and Data	3 0	1 E	0001 1110	1 5 8	9 E	1001 1110		Number of bytes to be transferred
	IPL	0 2	0 2	0000 0010					Number of bytes to be transferred
Continue Scan	Search Equal	3 7	2 5	0010 0101	1 6 5	A 5	1010 0101	} CPU storage location of search argument	} Number of bytes in search argument
	Search High	6 9	4 5	0100 0101	1 9 7	C 5	1100 0101		
	Search High or Equal	1 0 1	6 5	0110 0101	2 2 9	E 5	1110 0101		
	Set Status Modifier	5 3	3 5	0011 0101	1 8 1	B 5	1011 0101		
	Do not Set Status Modifier	8 5	5 5	0101 0101	2 1 3	D 5	1101 0101		
	Set Status Modifier	1 1 7	7 5	0111 0101	2 4 5	F 5	1111 0101		
Write	Home Address	2 5	1 9	0001 1001				} CPU storage location from which areas to be written will be transferred	5 (usually)
	Record R0	2 1	1 5	0001 0101					8+Key Length + Data Length of Record R0
	Erase	1 7	1 1	0001 0001					8 + Key Length + Data Length
	Count, Key and Data	2 9	1 D	0001 1101					8+Key Length + Data Length
	Special Count Key and Data	0 1	0 1	0000 0001					8+Key Length + Data Length
	Data	0 5	0 5	0000 0101					Data Length
Key and Data	1 3	0 D	0000 1101				Key Length + Data Length		

* Two channel switch special feature or 2314/2844 configuration only.
X Not significant.

APPENDIX C. TRACK ORIENTATION

As there is no unique physical address associated with each record on a given track, the 2314 must have some means of locating these records. There are seven states of orientation with respect to a track in the 2314 that are used to accomplish this.

1. Index Point State – Gap between Index Point and Home Address.
2. Home Address State – Gap between Home Address and R0.
3. Count State – Gap between Count and Key Fields.
4. Key State – Gap between Key and Data Field.
5. Data State – Gap between Data Field and succeeding Address Mark or Index Point if this is the last record on the track.
6. Address Marker State – Gap between Address Mark and Count Field.
7. Reset Orientation State – None of the above.

The reset Orientation State, while not associated with any specific area of a track, does not necessarily imply that the 2314 has lost orientation. Any time a CCW chain is broken or a Control Command is performed, the 2314 is set to this state. The next data command (i.e., read, write or search) further

defines this state to one of the three sub-states below.

1. Reset Orientation to Index Point State (ROIP). Orientation state is set to Index Point State upon detection of the Index Point on the track.
2. Reset Orientation to Address Marker State (ROAM). Orientation State is set to Address Marker State upon detection of any Address Marker.
3. Reset Orientation to Address Marker or Index Point State (ROAM or IP). Orientation State is set to either Index Point State upon detection of the Index Point, or to Address Marker State upon detection of any Address Markers, whichever occurs first.

Read, Write, Search and some Control commands in the 2314 have two types of prerequisites that must be satisfied to insure proper operation of the command. By considering command sequence restrictions and orientation requirements, all valid command sequences can be constructed and the result of their execution predetermined. The following table illustrates these two prerequisites and the resulting orientation state for all data commands:

Command	Command Prerequisite	Valid Orientation State of Beginning of Command	Orientation State at Completion of Command
Read CKD	None	ROAM	Data
Read KD	None	Count ROAM	Data
Read D	None	Count Key ROAM	Data
Write CKD (also Write Special CKD)	Search Equal Count or Key Write CKD Write R0	Count Key Data	Data
Write KD	Search Equal Count	Count ROAM	Data
Write D	Search Equal Count or Key	Count Key	Data
Search ID	None	ROAM or IP	Count
Search Key	None	ROAM Count	Key
Search Home Address	None	ROIP	Home Address
Read R0	None	Home Address ROIP	Data
Write R0	Search Equal Home Address Write HA	Home Address	Data
Read HA	None	ROIP	Home Address
Write HA	None	ROIP	Home Address
Read IPL	None	ROAM	Data
Read Count	None	ROAM	Count
Control Space Count	Search (any) Read (any)	Count Key	Data if Not Chained. Count if Chained.
Control Erase	Write CKD Write R0	Count Key Data	Reset Orientation
Control NOP	None	None	Reset Orientation

11540B

INDEX

- Access Mechanism 3
- Access Times 12
- Action Table 60
- Address Marker, Gap (G3) 5,9
- Addressing 21
- Addressing with Two-Channel Switch 54
- Alternate Track 7
- ALU Check 26
- AM (Address Marker) 9
- Appendage (BCA), Bit Count 11
- Appendix A. Record Capacities 65
- Appendix B. Command Summary 66
- Appendix C. Track Orientation 67
- Attachment and Addressing (2314/2844) 55
- Auxiliary Storage Control, 2844 1,55

- BCA (Bit Count Appendage) 11
- Bit 3
- Bit Count Appendage (BCA) 11
- Bit Count Byte 12
- Bit Positions 4
- Branching in Channel Programs 27
- Bus Out Parity Check 25
- Busy 22,26
- Byte 3,4

- CAW (Channel Address Word) 16
- Capacities, Record 12,65
- Capacity, Module 1
- Capacity, Storage 4
- CCW (Channel Command Word) 17
- Chain Command (CC) Flag 17
- Chain Data (CD) Flag 17
- Chaining 18
- Chaining Check 22
- Chaining Requirements (see Specific Command)
- Changing a Module Address 63
- Channel Address Word (CAW) 16
- Channel Attachment,
 - Enable/Disable Switch 14
 - Usage Meter 15
- Channel Command Word (CCW) 17
- Channel Commands, 23
 - Branching in Channel Programs 27
 - Control Commands 29
 - End of File 47
 - Multiple Track (M-T) Operation 48
 - Read Commands 39
 - Record Overflow 48.1
 - Search Commands 32
 - Write Commands 43
- Channel Control Check 22
- Channel Data Check 22
- Channel End 22
- Channel Programs, examples 50
- Channel Requirements 14
- Channel Selection Switch 52
- Channel Status 22
- Channel Status Word (CSW) 21
- Checks,
 - Chaining 22
 - Channel Control 22
 - Channel Data 22
 - Interface Control 22
 - Program 22
 - Protection 22
- Checking, Cyclic 11
- Checking, Data 11
- CKD (Count, Key, and Data) 8
- Command Code 17
- Command Codes, Table of 19
- Command Reject 25
- Command Sequence, Space Count 39
- Command Summary 66
- Command, Sense 23
- Command, Set File Mask 18
- Commands (see Specific Command)
- Commands, Channel (also see Channel Commands) 23
- Comparison with 2311 2
- Configuration Control (2314/2844) 57
- Configuration Unit (2167) 54
- Continue Scan Commands 50
- Control Commands,
 - No Operation (No Op) 29
 - Recalibrate 29
 - Seek Commands 31
 - Set File Mask 29
- Control Unit End 22
- Count Area (R0) 8
- Count Area (R1-Rn) 10
- Count Area Check 26
- Count, CCW 17
- CPU (Central Processing Unit) 14
- CSW (Channel Status Word) 21
- Cyclic Checking 11
- Cylinder 3
- Cylinder Concept 4
- Cylinder End 26
- Cylinder Number 7,8,10

- Data Address 17
- Data Area (R0) 9
- Data Area (R1-Rn) 11
- Data Check 7,9,11,25
- Data Checking,
 - Bit Count Appendage (BCA), 11
 - Bit Count Byte 12
 - Indicator Byte 12
 - Cyclic Checking 11
- Data Format, 3
 - Bit Positions 4
 - Byte 4

Data Format (Continued)
 Cylinder Concept 4
 Field 4
 Logical File 4
 Record 4
 Data Length 9,11
 Data Records (R1-Rn) 9
 Count Area (R1-Rn),
 Cylinder Number 10
 Data Check 11
 Data Length 11
 Flag Byte 10
 Key Length 11
 Read/Write Head Number 10
 Record Number 11
 Data Area (R1-Rn) 11
 Key Area (R1-Rn) 11
 Data Transfer 4
 Data Transfer Rate 12
 Data Write Commands 47
 Defective Count Area 38
 Defective Home Address or Record R0 38
 Defective Track 7
 Delay, Rotational 13
 Description, Functional 2
 Detection, Index Marker 5
 Device End 22
 Device Release Command 52
 Device Reservation 52
 Device Reserve Command 52
 Diagnostic Operations 14
 Diagnostic State 14
 Direct Access Storage Facility, 2314 1
 Disabling 14
 Disk I/O Operations 18
 Disk Organization 3
 Disk Pack Cover 61
 Disk Pack Handling,
 Labeling 61
 Loading 62
 Storage 61
 Unloading 63
 Disk Pack, 2316 3
 Disk Packs, Error Free 61

 Emergency Power Off Switch 54,57
 Enable/Disable Switch 14,64
 Enabling 14
 End of Cylinder 18,27
 End of File 47
 EPO (Emergency Power Off) 57
 Equipment Check 25
 Erase 46
 Error Checking 11
 Error Condition Table 59
 Error Free Disk Packs 61
 Error Recovery Procedures 58

 Field 4
 File Mask 29
 File Protected 26,29

 File Scan 38
 File Scan Commands 36
 Flag Byte 6,8,10
 Format, I/O Instruction 20
 Format, Track 5
 Formatting Overflow Records 48.1
 Formatting Write Commands 43
 Functional Description,
 Access Mechanism 3
 Comparison with 2311 2
 Disk Organization 3
 2316 Disk Pack 3

 G (Gap) 5
 Gaps (G), 5
 Gap (G1) 9,11
 Gap (G2) 9,11
 Gap (G3) Address Marker 9
 Gap (G4) 5
 Gap (G5) 7

 HA (Home Address) 5
 Halt I/O 20
 Handling, Disk Pack 61
 Head Number, Read/Write 7,8,10
 Head Switching 32
 Home Address Area, 5
 Cylinder Number 7
 Data Check 7
 Home Address Flag Byte 6
 Read/Write Head Number 7
 Home Address Gap (also see Gaps) 5

 ID (Identifier) 10
 Identifier Plugs, Module 63
 Incorrect Length 22
 Index Marker 5
 Index Point 5
 Indicator Byte 12
 Indicators and Operator Controls 63
 Indicators, Select Lock 63
 Indicators, Status 63
 Information Transfer 4
 Information, Operator 61
 Initialization 18
 Initialize (also see Format Write Commands 43
 In-Line Mode 14
 Input/Output Operations 16
 Instructions, I/O 19
 Interface Control Check 22
 Interrupt, I/O 17
 Intervention Required 25
 Introduction,
 2314 Direct Access Storage Facility 1
 2844 Auxiliary Storage Control 1
 Invalid Sequence 26
 I/O Instruction Format 20
 I/O Instructions,
 Addressing 21
 Halt I/O 20

I/O Instructions (Continued)
 Start I/O 19
 Test I/O 20
 I/O Interrupt 17
 I/O Operations, Disk 18
 I/O Operations, Termination of 19
 IPL Record 18

Key Area (R0) 9
 Key Area (R1-Rn) 11
 Key Length 9,11

Labeling, Disk Pack 61
 Loading, Disk Pack 62
 Logical File 4

Malfunction Reset 14
 Mask, File 29
 Messages 60
 Meter, Usage 15
 Missing Address Marker 26
 Module Address, Changing a 63
 Module Identification 27
 Module Identifier Plugs 63
 Multi-Module Select 27
 Multiple Track (M-T) Operation 32, 48

No Operation (No-Op) 29
 No Record Found 18, 26

On Line 26
 OP Codes (Operation Codes) 19
 Operations, Input/Output 16
 Operator Controls and Indicators 63
 Module Identifier Plugs 63
 Select Lock Indicators 63
 Start/Stop Switches 63
 Status Indicators 63
 Operator Information 61
 Operator Procedures,
 Changing a Module Address 63
 Disk Pack Loading 62
 Disk Pack Unloading 63
 Organization, Disk 3
 Orientation and Synchronization, Disk Pack 5
 Orientation, Track 67
 Overflow Incomplete 26
 Overflow Records,
 Formatting 48.1
 Processing 49
 Overrun 25

Pack Change 27
 Packed-Decimal Digits 4
 Parity 4
 Parity Check 29
 Pending Status 27
 Pending, Interrupt 17
 Power Control (2314/2844) 57
 Power Control with Two-Channel Switch 54
 Primary Track 7
 Problem State 16
 Procedures, Error Recovery 58
 Procedures, Operator 62
 Processing Overflow Records 49
 Program Check 22
 Program Control of 2314/2844 55
 Program-Control-Interruption (PCI) Flag 17
 Program-Controlled Interruption 22
 Program Status Word (PSW) 16
 Protection Check 22
 PSW (Program Status Word) 16

R0 (Record Zero), Track Descriptor Record 7
 Read Commands,
 Read Count 39
 Read Count, Key and Data 42
 Read Data 41
 Read Home Address (Read HA) 39
 Read IPL (Initial Program Load) 43
 Read Key and Data 41
 Read Track Descriptor Record (R0) 41
 Read/Write Head Number 7, 8, 10
 Recalibrate 29
 Record 4
 Record Capacities 12, 65
 Record Identifier (also see Count Area) 8, 10, 34
 Record Number 9, 11
 Record Overflow,
 Continue Scan Commands 50
 Formatting Overflow Records 48.1
 Processing Overflow Records 49
 Remote Switch Attachment 54
 Requirements, Channel 14
 Reset 31
 Resets,
 Malfunction Reset 14
 System Reset 14
 Resets with Two-Channel Switch 54
 Residual Count 22
 Rotational Delay 13

Scan, File 38
 Search Commands, 32
 Search Home Address Equal (Search HA) 33
 Search Identifier (Search ID) 34
 Search Key 35
 Search Key and Data 36
 Seek Check 25
 Seek Commands,

Seek Commands (Continued),

- Seek 32
 - Seek Cylinder 32
 - Seek Head 32
- Seek Incomplete 27
- Segments, Record Overflow 48.1
- Select Lock Indicators 63
- Sense Byte Summary 24
- Sense Bytes 25
- Sense Command 23
- Serializer/Deserializer Check 26
- Service Module 1,14
- Set File Mask 29
- Set File Mask Command 18
- Skip Flag 17
- SLI (Suppress Length Indicator) 17
- Space Count 38
- Spare Module (also see Service Module) 1,14
- Specifications, 2314 2
- Specifications, 2316 Disk Pack 3
- Start I/O 19
- Start/Stop Switches 63
- Status,
 - Byte 23
 - Indicators 63
 - Modifier 21
 - Presentation (2314/2844) 56
- Storage Capacity 4
- Storage, Disk Pack 61
- Summary, Command 66
- Summary, Sense Byte 24
- Supervisor Call (SVC) 16
- Supervisor State 16
- Suppress Length Indicator (SLI) 17
- SVC (Supervisor Call) 16
- Switches,
 - Channel Selection 52
 - Emergency Power Off 57
 - Enable/Disable 14,64
 - Start/Stop 63
 - Two-Channel 52
 - Untagged-Tagged 54
 - Untagged-Tagged (2314/2844) 56
- Synchronization, Disk Pack, Orientation and 5
- System Reset 14

- Table, Action 60
- Table, Error Condition 59
- Termination of I/O Operations 19
- Test I/O 20,23
- TIC (Transfer in Channel) 28
- Timing and Control (2314/2844) 57
- Timing Considerations,
 - Access Times 12
 - Rotational Delay 13
- Track Condition 7
- Track Condition Check 25
- Track Descriptor Record (R0), 7
 - Count Area (R0), 8
 - Cylinder Number 8
 - Data Check 9
 - Data Length 9

Track Descriptor Record (R0) (Continued)

- Count Area (R0), (Continued)
 - Flag Byte 8
 - Key Length 9
 - Read/Write Head Number 8
 - Record Identifier 8
 - Record Number 9
 - Data Area (R0) 9
 - Key Area (R0) 9
- Track Format 5
 - Data Records (R1-Rn) 9
 - Gaps (G) 5
 - Home Address Area 5
 - Index Marker 5
 - Track Descriptor Record (R0) 7
- Track Orientation 67
- Track Overrun 26
- Transfer in Channel 28
- Transfer, Data 4
- Two-Channel Switch,
 - Addressing 54
 - Channel Selection Switch 52
 - Device Reservation,
 - Device Release Command 52
 - Device Reserve Command 52
 - Power Control 54
 - Resets 54
 - Untagged-Tagged Switch 54
 - Usage Meter 54
- Unit Check 22
- Unit Exception 22
- Unloading, Disk Pack 63
- Unsafe 26
- Unselected Status 26
- Untagged-Tagged Switch 54,55
- Usage Meter 15,57
- Usage Meter with Two-Channel Switch 54
- Write Checking 44
- Write Commands,
 - Data Write Commands,
 - Write Data 47
 - Write Key and Data 47
 - Formatting Write Commands,
 - Erase 46
 - Write Count, Key, and Data 46
 - Write Home Address 44
 - Write Special Count, Key, and Data 46
 - Write Track Description Record (Write R0) 45
 - Write Checking 44
- Write Current Sense 27
- Write Record Zero 45
- Wrong Length Record 27
- Zoned-Format 3
- 2167 Configuration Unit 54
- 2314 Direct Access Storage Facility 1
- 2316 Disk Pack 3

2844 Auxiliary Storage Control 1, 55
Attachment and Addressing (2314/2844) 55
Configuration Control (2314/2844) 57
Power Control (2314/2844) 57
Program Control of 2314/2844 55

2844 Auxiliary Storage Control (Continued)
Status Presentation (2314/2844) 56
Timing and Control (2314/2844) 57
Untagged-Tagged Switch (2314/2844) 56
Usage Meter (2314/2844) 57

READER'S COMMENT FORM

IBM System/360 Component
 Descriptions -- 2314 and 2844

Form GA26-3599-6

- 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	
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
95114

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)

IBM System/360 Component Publications - 2314 Direct Access
Storage Facility and 2844 Auxiliary Storage Control

Printed in U.S.A. GA26-3599-6



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)