FUNCTIONAL SPECIFICATION
ULTRASTAR XP (DFHS) SCSI MODELS
1.12/2.25 GB - 1.0" HIGH
4.51 GB - 1.6" HIGH
3.5 INCH DISK DRIVES
Version 5.0

November 9, 1994

1.0 Preface

This document details the Hardware Specifications for the ***DFHS*** High Performance Family of 3.5-inch Direct Access Storage Devices. The capacity model offerings are 1.12 or 2.25 GByte in a 1-inch high form factor and a 4.51 GByte model in a 1.6-inch high form factor. Please refer to section 4.1.1, “Capacity Equations” on page 13 for exact capacities based on model and block size.

The product description and other data found in this document represent IBM's design objectives and is provided for information and comparative purposes. Actual results may vary based on a variety of factors and the information herein is subject to change. THIS PRODUCT DATA DOES NOT CONSTITUTE A WARRANTY, EXPRESS OR IMPLIED. Questions regarding IBM's warranty terms or the methodology used to derive the data should be referred to your IBM representative.

1.1.1 Document Distribution

Distribution of this document should come directly from your IBM Customer Representative to ensure you are receiving the current specification.

Replacement and disposal of down level versions is the responsibility of the holder.

2.0 Contents

3. Description

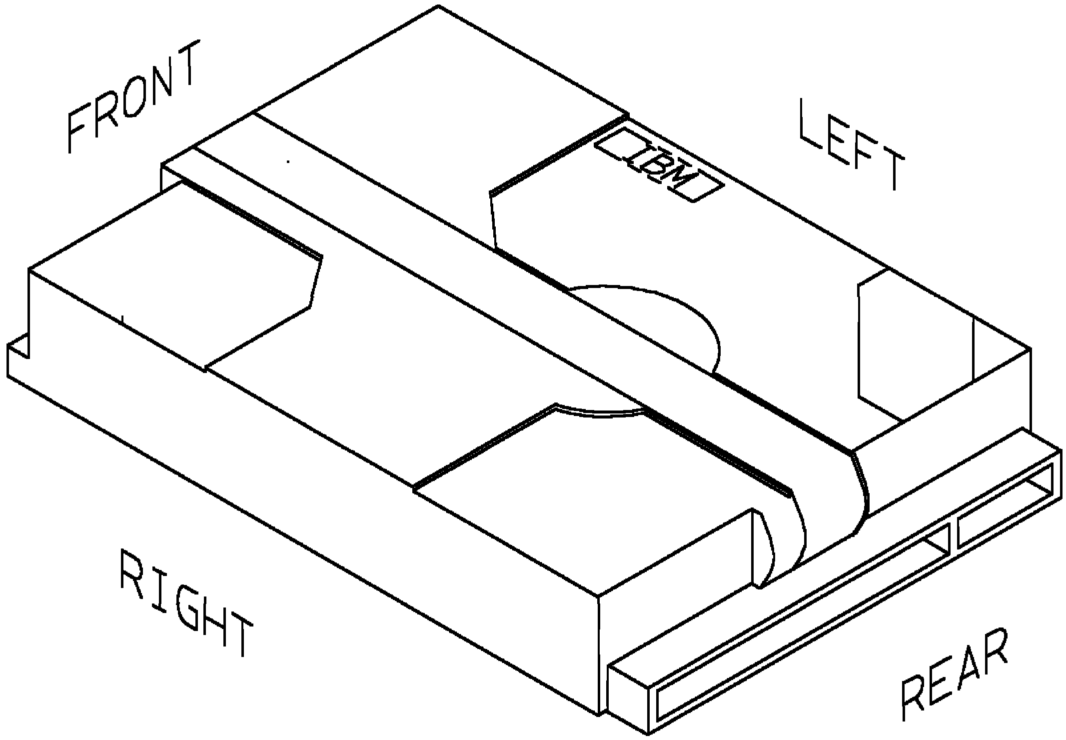

Figure 1. ***DFHS*** Drive Assembly

3.1 Features

General Features

- 1.12/2.25/4.51 GigaBytes formatted capacity (512 bytes/sector)
- Industry-standard interface:
  - 50 pin ANSI SCSI-2
  - 68 pin ANSI SCSI-3
  - Single Ended (50/68 pin) or Differential (68 pin)
  - Single Connector Attachment (SCA 80 pin)
- Rotary voice coil motor actuator
- Closed-loop digital actuator servo (embedded sector servo)
- Magnetoresistive (MR) heads
- (0,8,6,infinity) 8/9 rate encoding
- Partial Response Maximum Likelihood (PRML) data channel with digital filter
- All mounting orientations supported
- Jumperable auto spindle motor start
- Jumperable drive supplied terminator power
- Jumperable on board active SCSI terminators (Optional)
- Jumperable write protection
- Spindle synchronization
- LED Driver

- Bezel (optional)

Performance Summary

- Average read seek time (1.12 GB) : 6.9 milliseconds
- Average read seek time (2.25 GB) : 7.5 milliseconds
- Average read seek time (4.51 GB) : 8.0 milliseconds
- Average Latency: 4.17 milliseconds
- Media data transfer rate: 9.59 to 12.58 MegaBytes/second (10 bands)
- SCSI data transfer rate: 10 MegaTransfers/second (sustained synchronous)
- SCSI Bus Overhead: < 40 microseconds

Interface Controller Features

- SCSI bus parity
- SCSI disconnect and reconnect capability
- Multiple initiator support
- Fast SCSI supported
- Wide SCSI models
- Supports blocksizes from 256 to 5952 bytes
- 512K byte, multi-segmented, dual port data buffer
- Read-ahead caching
- Adaptive caching algorithms
- Write Cache supported (write-back & write-through)
- Tagged and untagged command queuing
- Command reordering
- Back-to-back writes (merged writes)
- Automatic retry and data correction on read errors
- Automatic sector reallocation
- In-line alternate sector assignment for high-performance
- Improved technique for down-loadable SCSI firmware
- SCSI behavior customizing jumpers

  For example -
  - Disable Target Initiated Synchronous Negotiation
  - Disable Unit Attentions
  - Disable SCSI Parity
  - Auto Start Delay

Reliability Features

- Self-diagnostics on power up
- Dedicated head landing zone
- Automatic actuator latch
- Embedded Sector Servo for improving on-track positioning capability
- Buffer memory parity
- Longitudinal Redundancy Check (LRC) on Customer Data
- ECC on the fly
- Error logging and analysis
- Data Recovery Procedures (DRP)
- Predictive Failure Analysis ™ (PFA ™)
- No preventative maintenance required
- Two Field Replaceable Units (FRU's): Electronics Card and Head Disk Assembly (HDA)
- Probability of not recovering data: **10 in 10^{15}** bits read

3.2 Models

The ***DFHS*** disk drive is available in various models as shown below:

The ***DFHS*** data storage capacities vary as a function of model and user block size. An emerging industry trend is capacity points in multiples of 1.08 GB (ie. 1.08/2.16/4.32) at a block size of 512 bytes. Future IBM products will plan to provide capacities that are consistent with this trend. Users who choose to make full use of the ***DFHS*** drive capacity above the standard capacity points may not find equivalent capacity breakpoints in future products.

DFHS-31080 Models	HDA Type	Capacity GB	SCSI Pins/Connector Type	SCSI Electrical Signal Type
-S1F	F2	1.12	50/'A' Connector	Single Ended Fast
-S1W	F2	1.12	68/'Unitized Connector	Single Ended Fast/Wide
-S1E	F2	1.12	68/Unitized Connector	Differential Fast/Wide
-S1S	F2	1.12	80 SCA	Single Ended Fast/Wide

DFHS-32160 Models	HDA Type	Capacity GB	SCSI Pins/Connector Type	SCSI Electrical Signal Type
-S2F	F4	2.25	50/'A' Connector	Single Ended Fast
-S2W	F4	2.25	68/Unitized Connector	Single Ended Fast/Wide
-S2E	F4	2.25	68/Unitized Connector	Differential Fast/Wide
-S2S	F4	2.25	80 SCA	Single Ended Fast/Wide

DFHS-34320 Models	HDA Type	Capacity GB	SCSI Pins/Connector Type	SCSI Electrical Signal Type
-S4F	F8	4.51	50/'A' Connector	Single Ended Fast
-S4W	F8	4.51	68/Unitized Connector	Single Ended Fast/Wide
-S4E	F8	4.51	68/Unitized Connector	Differential Fast/Wide
-S4S	F8	4.51	80 SCA	Single Ended Fast/Wide

Note: 50 pin SCSI connector models offer an 8 bit SCSI bus using the SCSI 'A' connector. 68 pin SCSI connector models offer an 8/16 bit SCSI bus using the SCSI 'P' connector which supports Wide data transfers.

Note: 80 pin SCSI connector models offer an 8/16 bit SCSI bus using the 'SCA' connector.

Note: All models support Fast SCSI data transfers

Note:
Please refer to section 4.1.1, "Capacity Equations" on page 13 for exact capacities based on user block size.

4. Specifications

All specifications are nominal values unless otherwise noted.

The ***DFHS*** data storage capacities vary as a function of model and user block size. An emerging industry trend is capacity points in multiples of 1.08 GB (ie. 1.08/2.16/4.32) at a block size of 512 bytes. Future IBM products will plan to provide capacities that are consistent with this trend. Users who choose to make full use of the ***DFHS*** drive capacity above the standard capacity points may not find equivalent capacity breakpoints in future products.

4.1 General

Note: The recording band located nearest the disk outer diameter (OD) is referred to as 'Notch #1'. While the recording band located nearest the inner diameter (ID) is called 'Notch #10'. 'Average' values are weighted with respect to the number of LBAs per notch when the drive is formatted with 512 byte blocks.

Data transfer rates

	Notch #1	Notch #10	Average	
Buffer to/from media	12.58	9.59	12.07	MB/s (instantaneous)

Host to/from buffer up to 20.0 MB/s (synchronous) (sustained)

Data Buffer Size (bytes) 512 K (See 5, “Performance” on page 39 for user data capacity.)

Rotational speed (RPM) 7202.7

Average latency (milliseconds) 4.17

Track Density (TPI) 4352

	Minimum	Maximum
Recording density (BPI)	96,567	124,970
Areal density (Megabits/square inch)	420.3	543.9

(model numbers - >)	S4x	S2x	S1x
Disks	8	4	2
User Data Heads (trk/cyl)	16	8	4
Seek times (in milliseconds)			
Single cylinder (Read)	0.5	0.5	0.5
(Write)	2.0	2.0	2.0
Average (weighted) (Read)	8.0	7.5	6.9
(Write)	9.5	9.0	9.0
Full stroke (Read)	16.5	15.0	14.0
(Write)	18.0	16.5	15.5

Note: Times are typical for a drive population under nominal voltages and casting temperature of 25 C. Weighted seeks are seeks to the cylinders of random logical block addresses (LBAs).

Total Cylinders (tcyl) & User Cylinders (ucyl)

	All models tcyl	S4x Models ucyl	S2x Models ucyl	S1x Models ucyl
. **Notch #1**	1893	1879	1877	1872
. **Notch #2**	956	955	955	955
. **Notch #3**	49	48	48	48
. **Notch #4**	310	309	309	309
. **Notch #5**	349	348	348	348
. **Notch #6**	116	115	115	115
. **Notch #7**	214	213	213	213
. **Notch #8**	190	189	189	189
. **Notch #9**	131	130	130	130
. **Notch #10**	208	206	206	206
. **Sum of all Notches**	4416	4392	4390	4385

Spares Sectors/cylinder (spr/cyl)

	S4x Models	S2x Models	S1x Models
Notch #1	40	20	10
Notch #2	40	20	10
Notch #3	38	19	10
Notch #4	37	19	9
Notch #5	36	18	9
Notch #6	34	17	9
Notch #7	33	17	8
Notch #8	32	16	8
Notch #9	31	16	8
Notch #10	30	15	7
Last cylinder extra spares (lcspr)	60	30	14

User bytes/sector (ub/sct)
256 - 744 (even number of bytes only)

. **Sectors/logical block (sct/lba)**
1-8

The lowest sct/lba that satisfies the following rules is used...

1. Block Length is evenly divisible by a number 2-8.
2. Quotient of previous equation is evenly divisible by 2.
3. Quotient must be $\geq$ 256 and $\leq$ 744.

User bytes/logical block (ub/lba)
256 - 5952 (See rules for determining sct/lba above for determining supported ub/lba values.)

Sectors/track (sct/trk)
(See Table 1 on page 13 or contact an IBM Customer Representative for other block lengths.)

	Notch #									
User bytes / logical block	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**
256	216	216	216	202	195	180	180	180	180	162
512	135	135	130	126	120	115	112	108	105	100
520	128	128	128	123	115	112	108	105	102	99
522	128	128	128	122	115	112	108	105	102	90
524	128	128	128	120	115	112	108	105	102	90
528	128	128	126	120	112	112	108	105	101	90
688	102	102	102	98	90	90	90	90	81	78
744	96	96	96	90	90	90	81	78	77	73

Table 1. Gross sectors per track for several block lengths

	S4x Models		S2x Models		S1x Models	
User bytes / logical block	**formatted capacity (bytes)**	**logical blocks / drive**	**formatted capacity (bytes)**	**logical blocks / drive**	**formatted capacity (bytes)**	**logical blocks / drive**
256	3,654,540,800	14,275,550	1,826,312,448	7,134,033	912,135,680	3,563,030
512	4,512,701,440	8,813,870	2,255,098,368	4,404,489	1,126,337,536	2,199,878
520	4,375,536,880	8,414,494	2,186,554,760	4,204,913	1,092,119,600	2,100,230
522	4,374,300,492	8,379,886	2,185,931,898	4,187,609	1,091,803,716	2,091,578
524	4,385,878,952	8,369,998	2,191,716,460	4,182,665	1,094,691,544	2,089,106
528	4,408,629,984	8,349,678	2,203,082,640	4,172,505	1,100,365,728	2,084,026
688	4,604,578,976	6,692,702	2,300,969,904	3,344,433	1,149,310,880	1,670,510
744	4,675,830,192	6,284,718	2,336,559,528	3,140,537	1,167,099,408	1,568,682

Table 2. User capacity for several block lengths

4.1.1 Capacity Equations

4.1.1.1 For Each Notch

The next group of equations must be calculated separately for each notch.

$$\text{user bytes/sector (ub/sct)} = \frac{\text{ub/lba}}{\text{sct/lba}}$$

user sectors/cyl (us/cyl) = (sct/trk)(trk/cyl) - spr/cyl

spares/notch (spr/nch) = (spr/cyl)(ucyl)

Note: Add lcspr to the equation above for the notch closest to the inner diameter (#10).

user sectors/notch (us/nch) = (us/cyl)(ucyl)

Note: Subtract lcspr from the equation above for the notch closest to the inner diameter (#10).

4.1.1.2 For Entire Drive

$$\text{spares/drive (spr/drv)} = \sum_{\text{notch}=1}^{10} \text{spr/nch}$$

$$\text{user sectors/drive (us/drv)} = \sum_{\text{notch}=1}^{10} \text{us/nch}$$

$$\text{logical blocks/drive (lba/drv)} = \text{INT}\left[\frac{\text{us/drv}}{\text{sct/lba}}\right]$$

user capacity (fcap) = (lba/drv)(ub/lba)

4.2 Power Requirements by Model

4.2.1 S1x Model

The following voltage specifications apply at the drive power connector. There is no special power on/off sequencing required. The extra power needed for drives with differential SCSI is shown in the section called 4.2.7, “Additional 5V Power Requirements for Differential” on page 35.

Input Voltage

+5 Volts Supply	5V (± 5% during run and spin-up)
+12 Volts Supply	12V (± 5% during run) (+5% / -7% during spin-up)

The following current values are measured values. Safety factors have not been applied.

Power Supply Current	Notes	Population Mean	Population Stand. Dev.
+5VDC (power-up)	Minimum voltage slew rate = 4.5 V/sec		
+5VDC (idle avg)		0.73 Amps[1]	0.02 Amps
+5VDC (R/W baseline)		0.96 Amps[2]	0.05 Amps
+5VDC (R/W pulse)	Base-to-peak	.36 Amps	0.06 Amps
+12VDC (power-up)	Minimum voltage slew rate = 7.4 V/sec		
+12VDC (idle avg)		0.28 Amps	0.02 Amps
+12VDC (seek avg)	1 op/sec	0.0027 Amps	0.0002 Amps
+12VDC (seek peak)		1.2 Amps[3]	0.02 Amps
+12VDC (spin-up)	3.0 sec max	1.5 Amps[4]	0.1 Amps

Drive power	Notes	Population Mean	Population Stand. Dev.
Avg idle power		7.0 Watts	.35 Watts
Avg R/W power	30 ops/sec	9.1 Watts	.35 Watts

[1] 5 Volt Current is given with termination power provided by the using system.

[2] See Figure 2 on page 18 for a plot of how the read/write baseline and read/write pulse sum together.

[3] The idle average and seek peek should be added together to determine the total 12 volt peak current. See Figure 3 on page 19
for a typical buildup of these currents. Refer to examples on the following page to see how to combine these values.

[4] The current at start is the total 12 volt current required (ie. the motor start current, module current and voice coil retract current). See Figure 4 on page 20 for typical 12 volt current during spindle motor start.

4.2.1.1 Power Calculation Examples

Note: The above formulas assume all system ops as a 1 block read or write transfer from a random cylinder while at nominal voltage conditions.

Example 1. Calculate the mean 12 volt average current.

If we assume a case of 30 operations/second then to compute the sum of the 12 volt mean currents the following is done.

		mean	
+12VDC (idle average)		0.28	amps
+12VDC (seek average)	0.0027 * 30 =	0.081	amps
TOTAL		0.361	amps

Example 2. Calculate the mean plus 3 sigma 12 volt average current.

To compute the sum of the 12 volt mean current's 1 sigma value assume all the distributions are normal. Therefore the square root of the sum of the squares calculation applies. Assume a case of 30 operations/second.

		sigma	
+12VDC (idle average)		0.02	amps
+12VDC (seek average)	sqrt(30*((0.0002)**2))=	0.001	amps
TOTAL	sqrt((0.02)**2+(.001)**2))=	0.02	amps

So the mean plus 3 sigma mean current is 0.361 + 3*0.02 = 0.42 amps

Example 3. Power Calculation.

Nominal idle drive power = (.73 Amps * 5 Volts) + (0.28 Amps * 12 Volts) = 7.01 Watts

Nominal R/W drive power at 30 ops/sec = (0.96 Amps * 5 Volts) + (0.361 Amps * 12 Volts) = 9.13 Watts

Mean plus 3 sigma drive power for 30 random R/W operations/second. Assume that the 5 volt and 12 volt distributions are independent therefore the square root of the sum of the squares applies.

+5VDC (1 sigma power)	0.05 * 5	= 0.25	watts
+12VDC (1 sigma power)	0.02 * 12	= 0.240	watts
Total (1 sigma power)	sqrt((0.25)**2+(0.24)**2)	= 0.347	watts
Total power	9.13 + 3 * 0.347	= 10.2	watts

Example 4. Calculate the 12 volt peak current. To compute the sum of the 12 volt peak currents the following is done.

	mean
+12VDC (idle avg)	0.28 amps
+12VDC (seek peak)	1.20 amps
TOTAL	1.48 amps

Example 5. Calculate the mean plus 3 sigma 12 volt peak current.

To compute the sum of the 12 volt peak current's 1 sigma value assume all distributions are normal. Therefore the square root of the sum of the squares calculation applies.

	sigma
+12VDC (idle avg)	0.02 amps
+12VDC (seek peak)	0.02 amps

TOTAL sqrt((0.02)**2+(0.02)**2)=0.028 amps

So the mean plus 3 sigma peak current is 1.48 + 3*0.028= 1.56 amps

The above formulas assume all system ops as a 1 block read or write transfer from a random cylinder while at nominal voltage conditions.

Things to check when measuring 12 V supply current:

- Null the current probe frequently. Be sure to let it warm up.
- Adjust the power supply to 12.00 V *at the drive terminals.*
- Use a proper window width, covering an integral number of spindle revolutions.
- Measure values at 25 degree C casting temperature.
- Get a reliable trigger for Seek Peak readings.

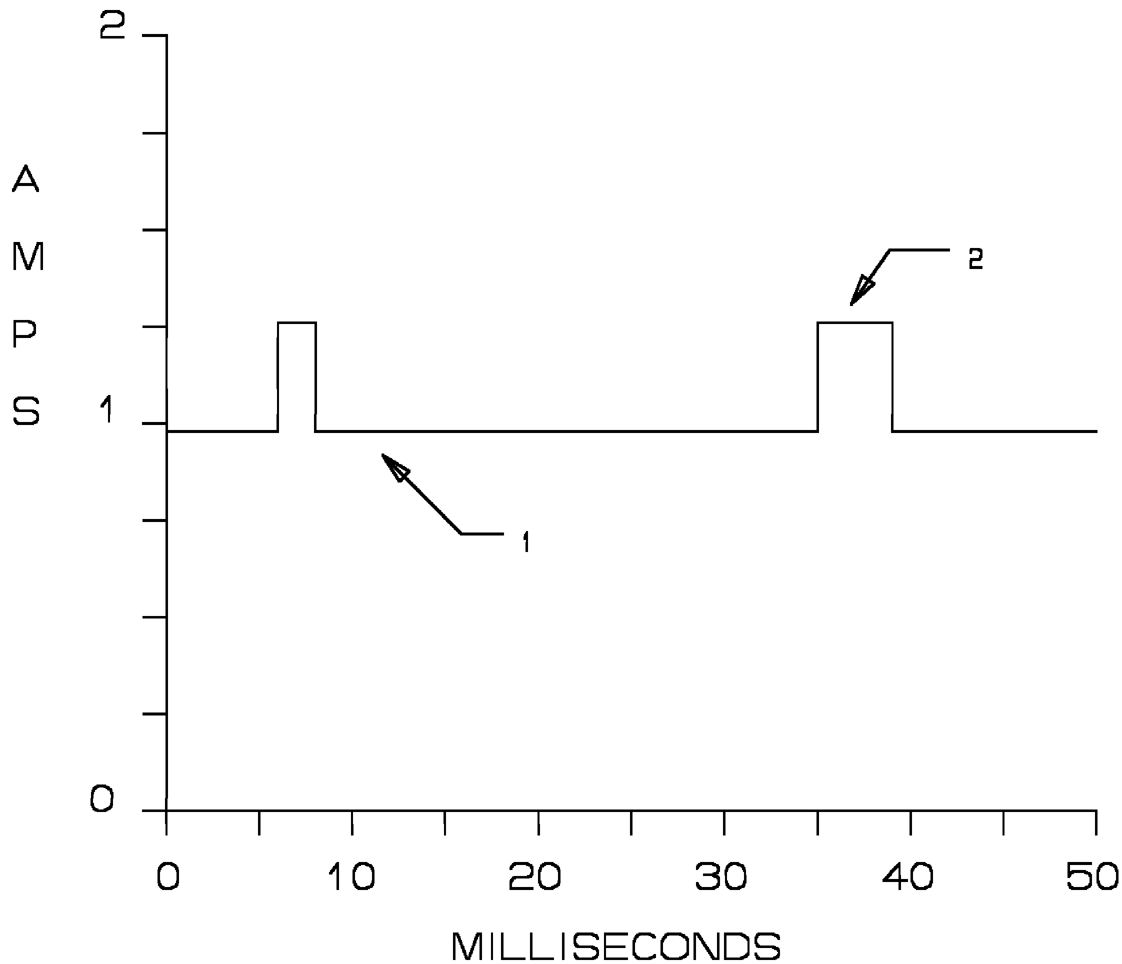

5 volt current during read/write operations.

Figure 2.

1. Read/write baseline voltage.

2. Read/write pulse. The width of the pulse is proportional to the number of consecutive blocks read or written. The 5 volt supply must be able to provide the required current during this event.

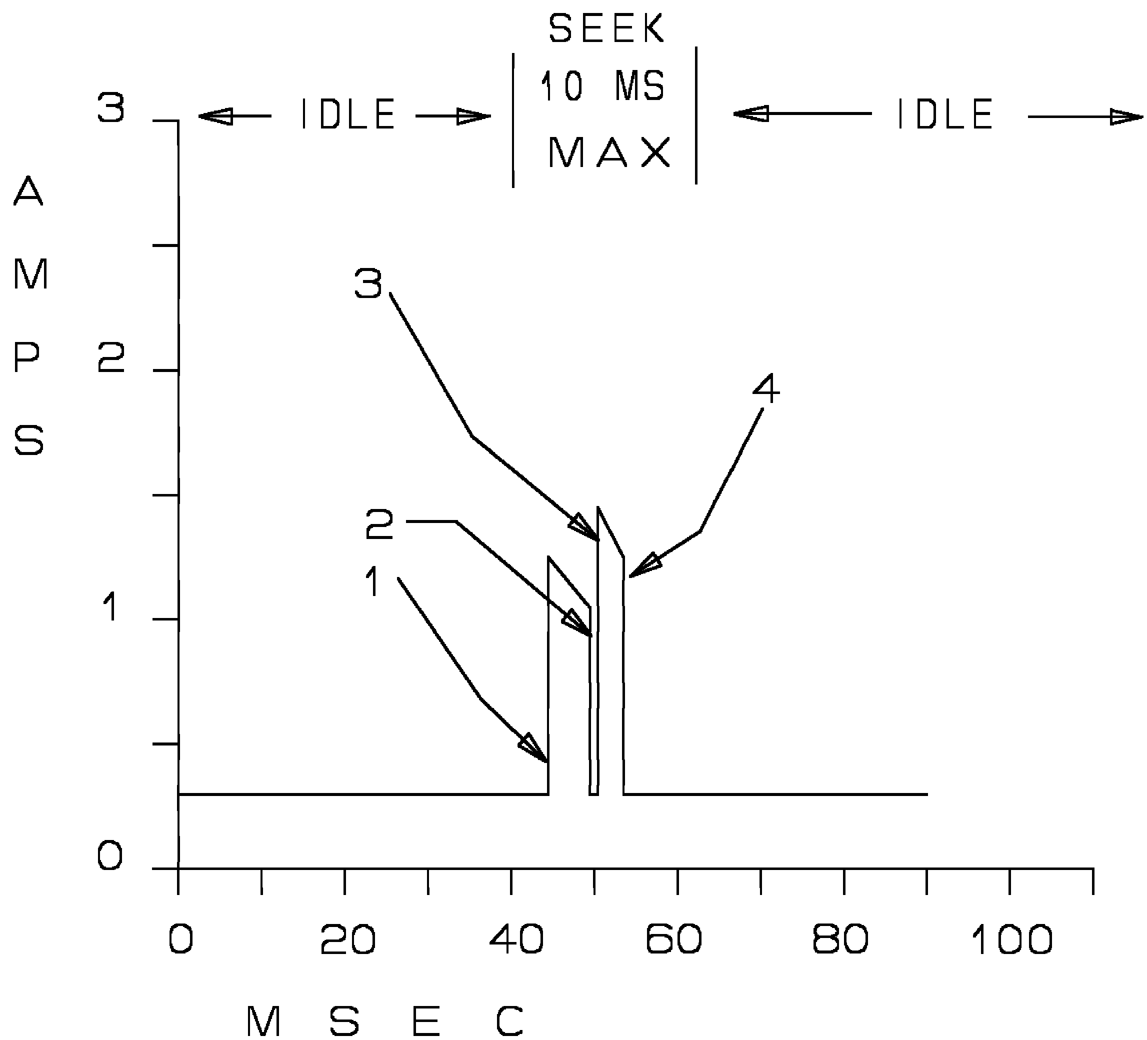

Typical 12 volt current.

Figure 3.

1. Maximum slew rate is 7 amps/millisecond.

2. Maximum slew rate is 100 amps/millisecond.

3. Maximum slew rate is 7 amps/millisecond.

4. Maximum slew rate is 3 amps/millisecond.

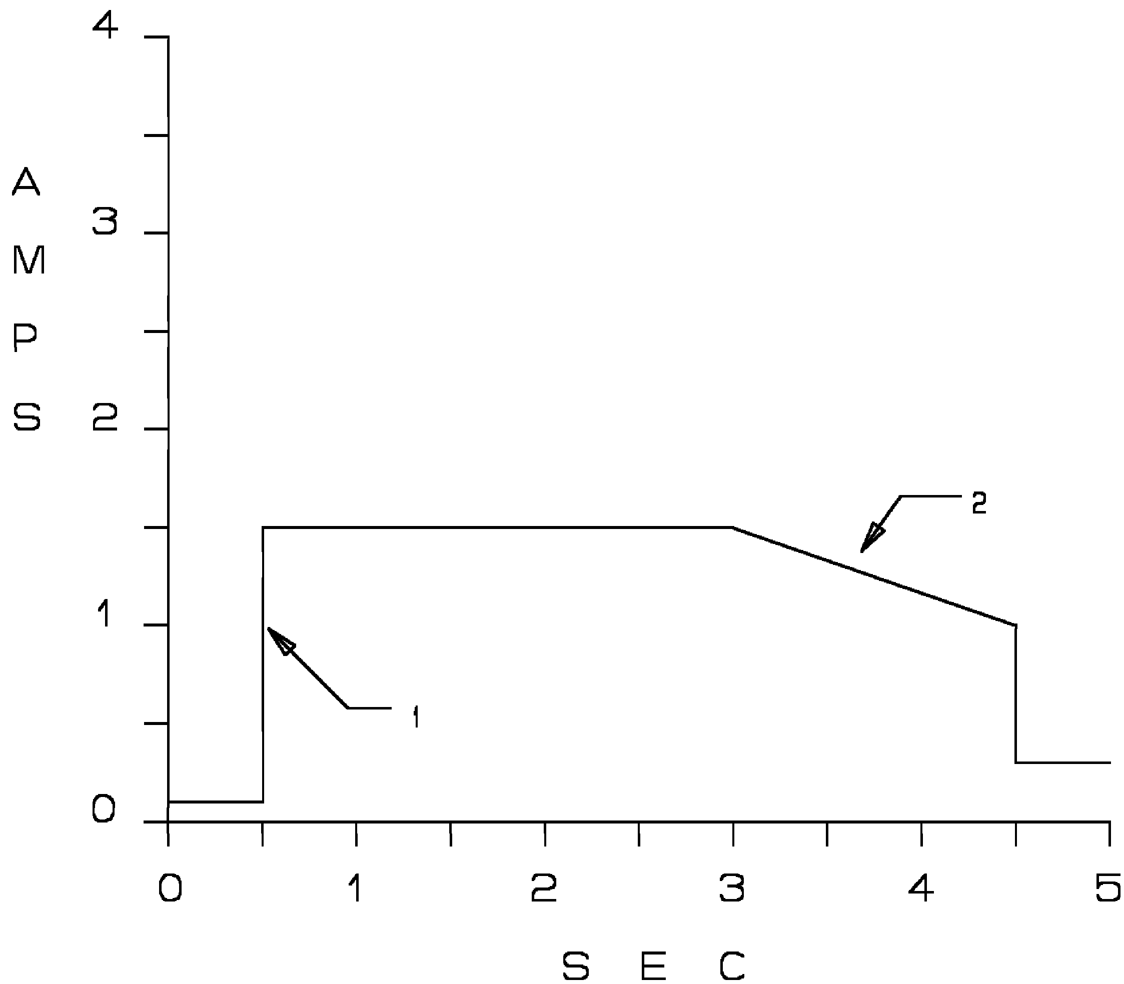

Typical 12 volt spin-up current.

Figure 4.

1. Maximum slew rate is 20 amps/millisecond.

2. Current drops off as motor comes up to speed.

4.2.2 S2x Model

The following voltage specifications apply at the drive power connector. There is no special power on/off sequencing required. The extra power needed for drives differential SCSI is shown in the section called 4.2.7, “Additional 5V Power Requirements for Differential” on page 35.

Input Voltage

+5 Volts Supply	5V (± 5% during run and spin-up)
+12 Volts Supply	12V (± 5% during run) (+5% / -7% during spin-up)

The following current values are measured values. Safety factors have not been applied.

Power Supply Current	Notes	Population Mean	Population Stand. Dev.
+5VDC (power-up)	Minimum voltage slew rate = 4.5 V/sec		
+5VDC (idle avg)		0.73 Amps[1]	0.02 Amps
+5VDC (R/W baseline)		0.96 Amps[5]	0.05 Amps
+5VDC (R/W pulse)	Base-to-peak	.36 Amps	0.06 Amps
+12VDC (power-up)	Minimum voltage slew rate = 7.4 V/sec		
+12VDC (idle avg)		0.41 Amps	0.02 Amps
+12VDC (seek avg)	1 op/sec	0.0031 Amps	0.0002 Amps
+12VDC (seek peak)		1.20 Amps[6]	0.02 Amps
+12VDC (spin-up)	4.2 sec max	1.5 Amps[7]	0.1 Amps

Drive power

Avg idle power		8.6 Watts	.35 Watts
Avg R/W power	30 ops/sec	10.8 Watts	.35 Watts

[5] See Figure 5 on page 24 for a plot of how the read/write baseline and read/write pulse sum together.

[6] The idle average and seek peek should be added together to determine the total 12 volt peak current. See Figure 6 on page 25
for a typical buildup of these currents. Refer to examples on the following page to see how to combine these values.

[7] The current at start is the total 12 volt current required (ie. the motor start current, module current and voice coil retract current). See Figure 7 on page 26 for typical 12 volt current during spindle motor start.

4.2.2.1 Power Calculation Examples

Note: The above formulas assume all system ops as a 1 block read or write transfer from a random cylinder while at nominal voltage conditions.

Example 1. Calculate the mean 12 volt average current.

If we assume a case of 30 operations/second then to compute the sum of the 12 volt mean currents the following is done.

	mean
+12VDC (idle average)	0.41 amps
+12VDC (seek average) 0.0031 * 30 =	0.09 amps
TOTAL	0.50 amps

Example 2. Calculate the mean plus 3 sigma 12 volt average current.

To compute the sum of the 12 volt mean current's 1 sigma value assume all the distributions are normal. Therefore the square root of the sum of the squares calculation applies. Assume a case of 30 operations/second.

	sigma
+12VDC (idle average)	0.02 amps
+12VDC (seek average) sqrt(30*((0.0002)**2))=	0.001 amps

TOTAL sqrt((0.02)**2+(.001)**2))=0.02 amps

So the mean plus 3 sigma mean current is 0.50 + 3*0.02 = 0.56 amps

Example 3. Power Calculation.

Nominal idle drive power = (.73 Amps * 5 Volts) + (0.41 Amps * 12 Volts) = 8.6 Watts

Nominal R/W drive power at 30 ops/sec = (0.96 Amps * 5 Volts) + (0.50 Amps * 12 Volts) = 10.8 Watts

Mean plus 3 sigma drive power for 30 random R/W operations/second. Assume that the 5 volt and 12 volt distributions are independent therefore the square root of the sum of the squares applies.

+5VDC (1 sigma power) 0.05 * 5	=	0.25 watts
+12VDC (1 sigma power) 0.02 * 12	=	0.24 watts
Total (1 sigma power) sqrt((0.25)**2+(0.24)**2)	=	0.35 watts
Total power 10.8 + 3 * 0.35	=	11.9 watts

Example 4. Calculate the 12 volt peak current.

To compute the sum of the 12 volt peak currents the following is done.

	mean	
+12VDC (idle avg)	0.41	amps
+12VDC (seek peak)	1.20	amps
TOTAL	1.61	amps

Example 5. Calculate the mean plus 3 sigma 12 volt peak current.

To compute the sum of the 12 volt peak current's 1 sigma value assume all distributions are normal. Therefore the square root of the sum of the squares calculation applies.

	sigma	
+12VDC (idle avg)	0.03	amps
+12VDC (seek peak)	0.02	amps

TOTAL sqrt((0.03)**2+(0.02)**2)=0.036 amps

So the mean plus 3 sigma peak current is 1.61 + 3*0.036= 1.72 amps

Things to check when measuring 12 V supply current:

- Null the current probe frequently. Be sure to let it warm up.
- Adjust the power supply to 12.00 V *at the drive terminals.*
- Use a proper window width, covering an integral number of spindle revolutions.
- Measure values at 25 degree C casting temperature.
- Get a reliable trigger for Seek Peak readings.

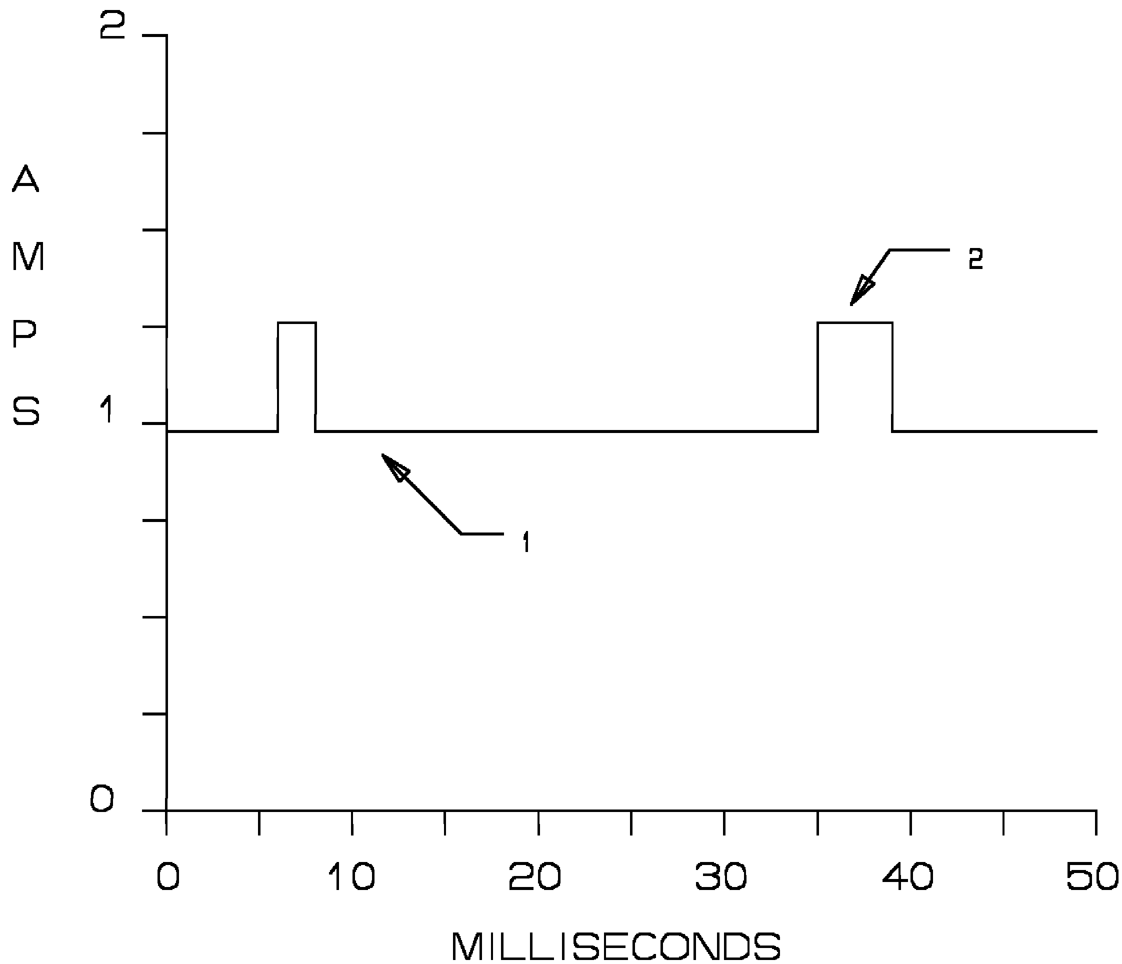

5 volt current during read/write operations.

Figure 5.

1. Read/write baseline voltage.

2. Read/write pulse. The width of the pulse is proportional to the number of consecutive blocks read or written. The 5 volt supply must be able to provide the required current during this event.

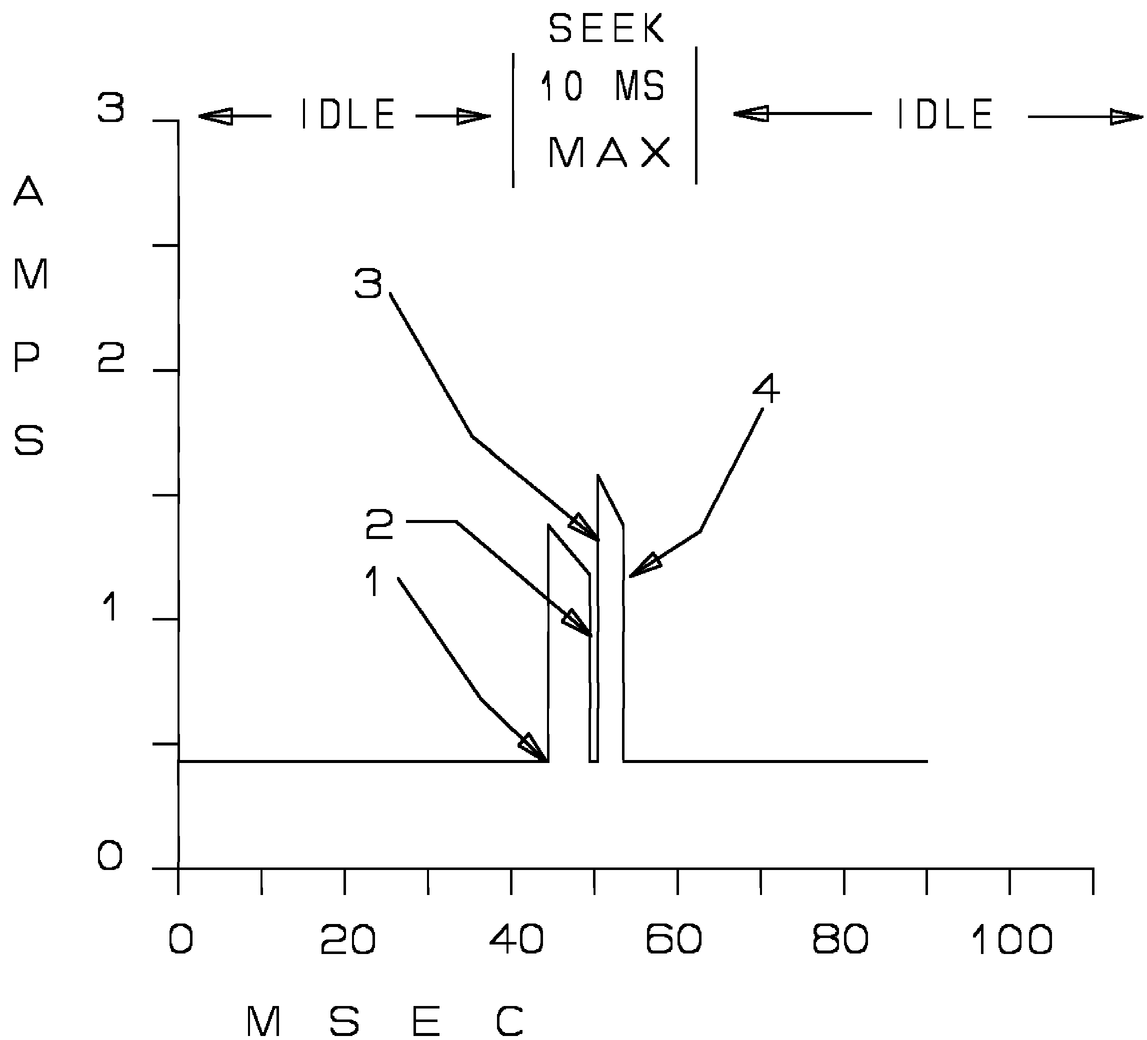

Typical 12 volt current.

Figure 6.

1. Maximum slew rate is 7 amps/millisecond.

2. Maximum slew rate is 100 amps/millisecond.

3. Maximum slew rate is 7 amps/millisecond.

4. Maximum slew rate is 3 amps/millisecond.

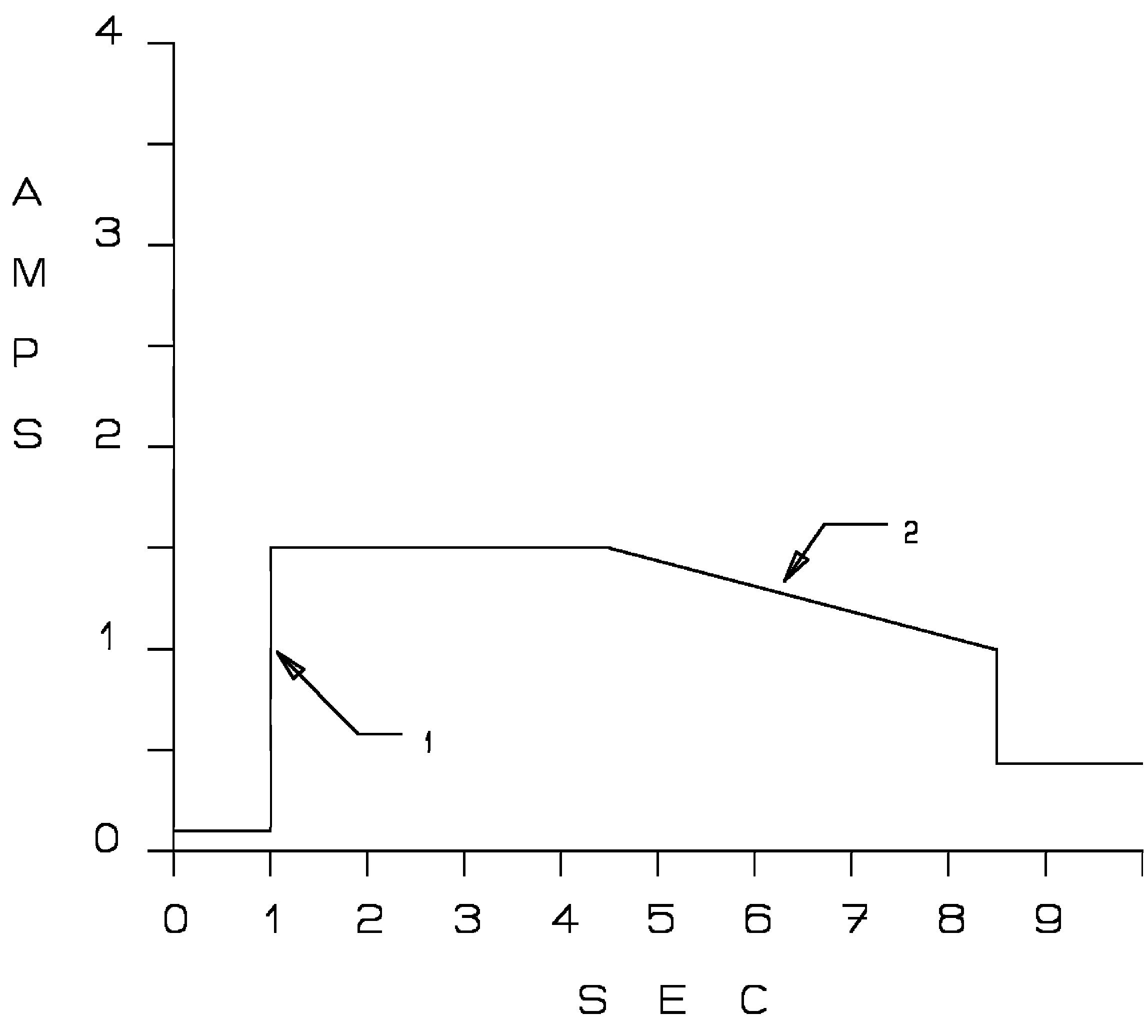

Typical 12 volt spin-up current.

Figure 7.

1. Maximum slew rate is 20 amps/millisecond.

2. Current drops off as motor comes up to speed.

4.2.3 S4x Model

The following voltage specifications apply at the drive power connector. There is no special power on/off sequencing required. The extra power needed for drives with differential SCSI is shown in the section called 4.2.7, “Additional 5V Power Requirements for Differential” on page 35.

Input Voltage

+5 Volts Supply	5V (± 5% during run and spin-up)
+12 Volts Supply	12V (± 5% during run) (+5% / -7% during spin-up)

The following current values are measured values. Safety factors have not been applied.

Power Supply Current	Notes	Population Mean	Population Stand. Dev.
+5VDC (power-up)	Minimum voltage slew rate = 4.5 V/sec		
+5VDC (idle avg)		0.76 Amps[1]	0.02 Amps
+5VDC (R/W baseline)		0.98 Amps[8]	0.05 Amps
+5VDC (R/W pulse)	Base-to-peak	.36 Amps	0.06 Amps
+12VDC (power-up)	Minimum voltage slew rate = 7.4 V/sec		
+12VDC (idle avg)		0.77 Amps	0.03 Amps
+12VDC (seek avg)	1 op/sec	0.0036 Amps	0.0002 Amps
+12VDC (seek peak)		1.3 Amps[9]	0.02 Amps
+12VDC (spin-up)	8.5 sec max	2.2 Amps[10]	0.1 Amps

Drive power

Avg idle power		13.0 Watts	.44 Watts
Avg R/W power	30 ops/sec	15.5 Watts	.44 Watts

[8] See Figure 8 on page 30 for a plot of how the read/write baseline and read/write pulse sum together.

[9] The idle average and seek peek should be added together to determine the total 12 volt peak current. See Figure 9 on page 31 for a typical buildup of these currents. Refer to examples on the following page to see how to combine these values.

[10] The current at start is the total 12 volt current required (ie. the motor start current, module current and voice coil retract current). See Figure 10 on page 32 for typical 12 volt current during spindle motor start.

4.2.3.1 Power Calculation Examples

Note: The above formulas assume all system ops as a 1 block read or write transfer from a random cylinder while at nominal voltage conditions.

Example 1. Calculate the mean 12 volt average current.

If we assume a case of 30 operations/second then to compute the sum of the 12 volt mean currents the following is done.

	mean
+12VDC (idle average)	0.77 amps
+12VDC (seek average) 0.0036 * 30 =	0.11 amps
TOTAL	0.88 amps

Example 2. Calculate the mean plus 3 sigma 12 volt average current.

To compute the sum of the 12 volt mean current's 1 sigma value assume all the distributions are normal. Therefore the square root of the sum of the squares calculation applies. Assume a case of 30 operations/second.

	sigma
+12VDC (idle average)	0.02 amps
+12VDC (seek average) sqrt(30*((0.0002)**2))=	0.001 amps

TOTAL sqrt((0.02)**2+(.001)**2))=0.02 amps

So the mean plus 3 sigma mean current is 0.88 + 3*0.02 = 0.94 amps

Example 3. Power Calculation.

Nominal idle drive power = (.76 Amps * 5 Volts) + (0.77 Amps * 12 Volts) = 13.04 Watts

Nominal R/W drive power at 30 ops/sec = (0.98 Amps * 5 Volts) + (0.88 Amps * 12 Volts) = 15.46 Watts

Mean plus 3 sigma drive power for 30 random R/W operations/second. Assume that the 5 volt and 12 volt distributions are independent therefore the square root of the sum of the squares applies.

+5VDC (1 sigma power) 0.05 * 5	= 0.25 watts
+12VDC (1 sigma power) 0.03 * 12	= 0.36 watts
Total (1 sigma power) sqrt((0.25)**2+(0.36)**2)	= 0.44 watts
Total power 15.46 + 3 * 0.44	= 16.8 watts

Example 4. Calculate the 12 volt peak current.

To compute the sum of the 12 volt peak currents the following is done.

	mean
+12VDC (idle avg)	0.77 amps
+12VDC (seek peak)	1.30 amps
TOTAL	2.07 amps

Example 5. Calculate the mean plus 3 sigma 12 volt peak current.

To compute the sum of the 12 volt peak current's 1 sigma value assume all distributions are normal. Therefore the square root of the sum of the squares calculation applies.

	sigma
+12VDC (idle avg)	0.02 amps
+12VDC (seek peak)	0.02 amps

TOTAL sqrt((0.02)**2+(0.02)**2)=0.028 amps

So the mean plus 3 sigma peak current is 2.07 + 3*0.028= 2.1 amps

Things to check when measuring 12 V supply current:

- Null the current probe frequently. Be sure to let it warm up.
- Adjust the power supply to 12.00 V *at the drive terminals.*
- Use a proper window width, covering an integral number of spindle revolutions.
- Measure values at 25 degree C casting temperature.
- Get a reliable trigger for Seek Peak readings.

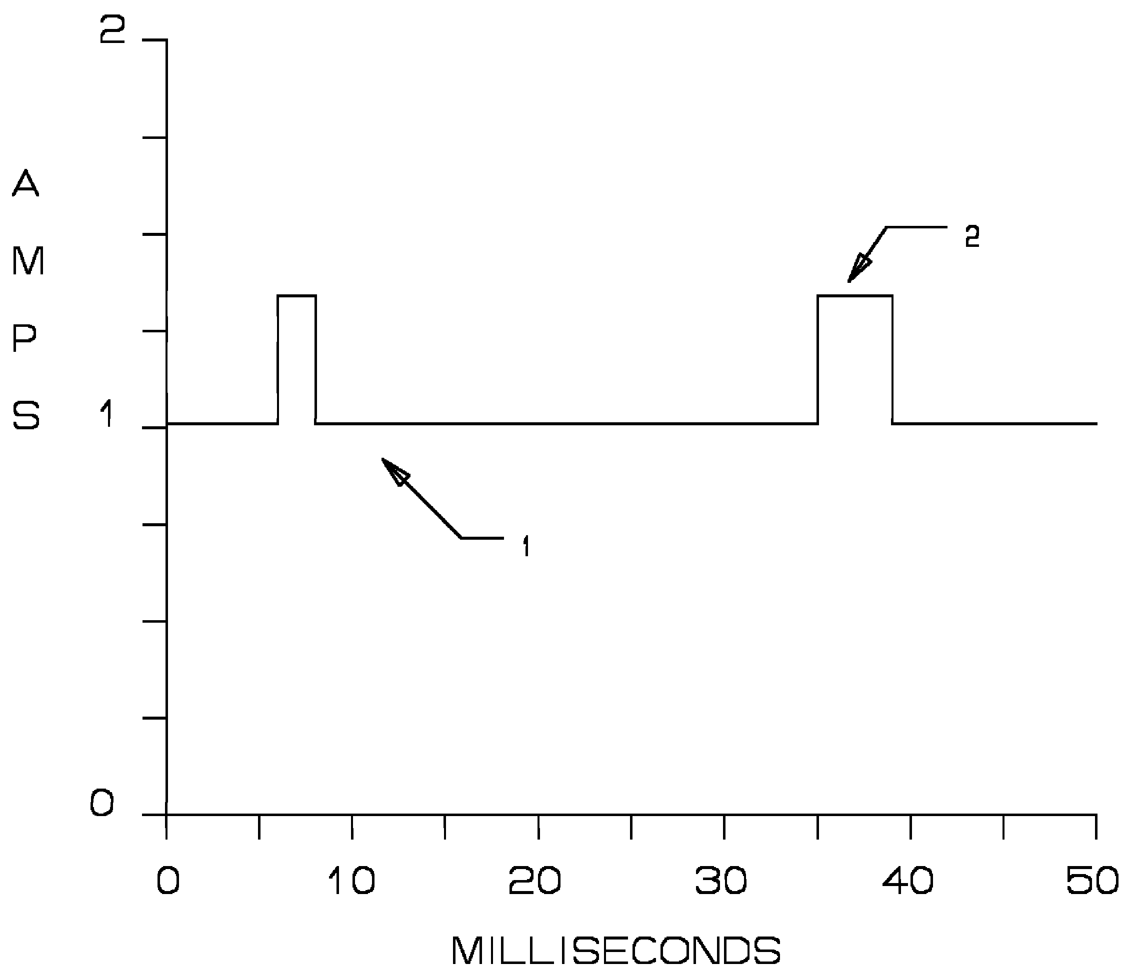

5 volt current during read/write operations.

Figure 8.

1. Read/write baseline voltage.

2. Read/write pulse. The width of the pulse is proportional to the number of consecutive blocks read or written. The 5 volt supply must be able to provide the required current during this event.

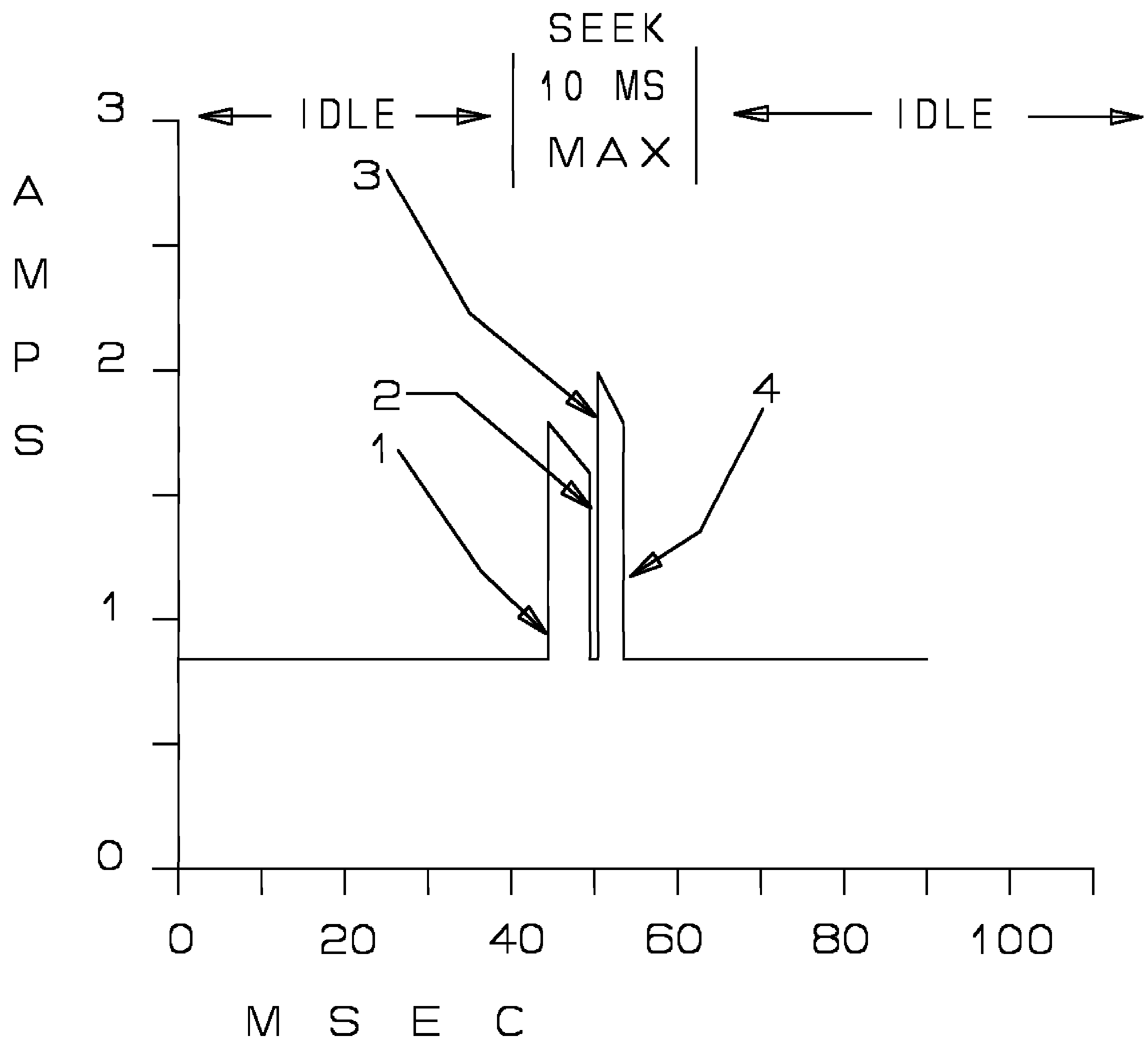

Typical 12 volt current.

Figure 9.

1. Maximum slew rate is 7 amps/millisecond.

2. Maximum slew rate is 100 amps/millisecond.

3. Maximum slew rate is 7 amps/millisecond.

4. Maximum slew rate is 3 amps/millisecond.

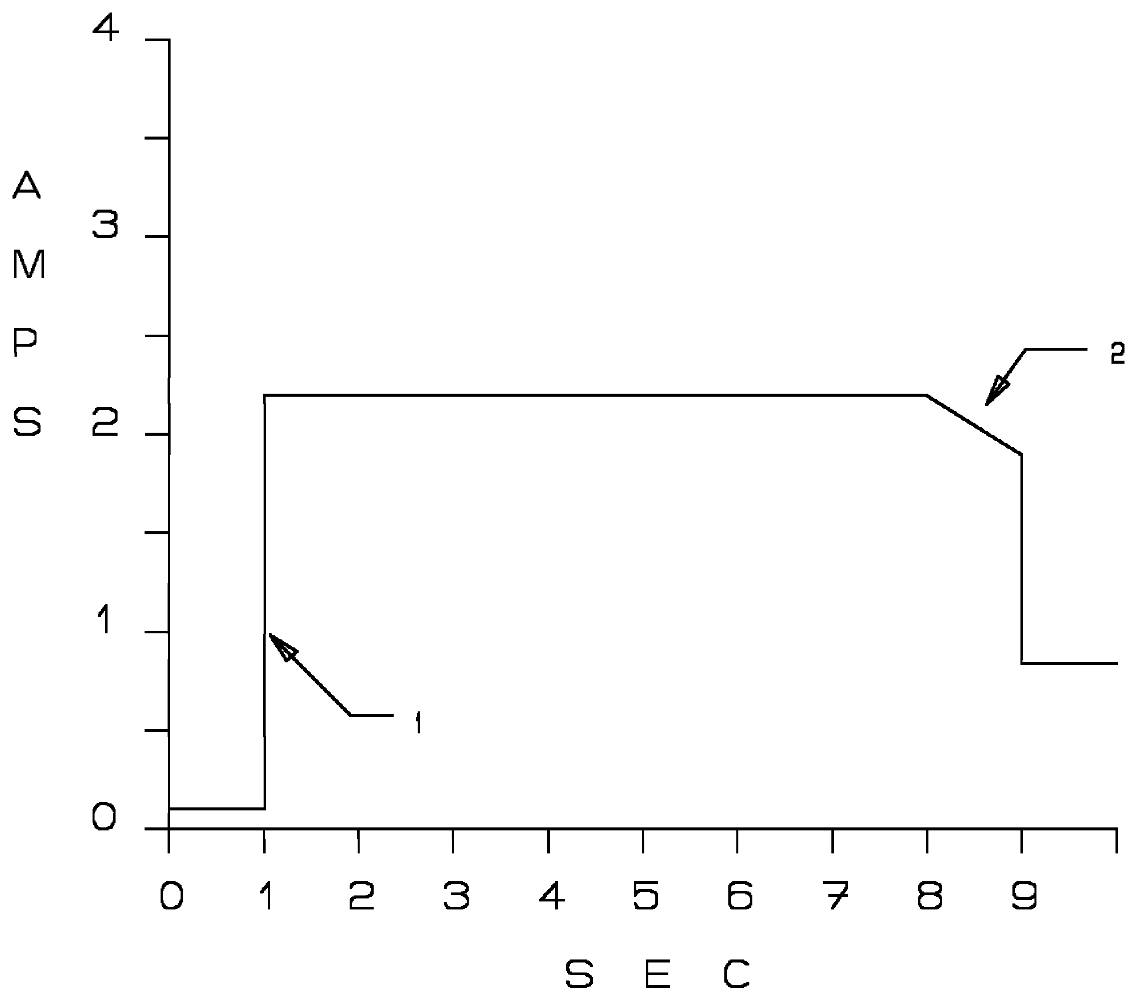

Typical 12 volt spin-up current.

Figure 10.

1. Maximum slew rate is 20 amps/millisecond.

2. Current drops off as motor comes up to speed.

4.2.4 Power Supply Ripple

Externally Generated Ripple[11] **as seen at drive power connector**	Maximum	Notes
+5VDC	200 mV peak-to-peak	0-20 MHz
+12VDC	200 mV peak-to-peak	0-20 MHz

During drive start up and seeking, 12 volt ripple is generated by the drive (referred to as dynamic loading). If several drives have their power daisy chained together then the power supply ripple plus other drive's dynamic loading must remain within the regulation tolerance window of +/- 5%. A common supply with separate power leads to each drive is a more desirable method of power distribution.

4.2.5 Grounding Requirements of the Disk Enclosure

The disk enclosure is at Power Supply ground potential. It is allowable for the user mounting scheme to common the Disk Enclosure to Frame Ground potential or to leave it isolated from Frame Ground.

From a ElectroMagnetic Compatibility (EMC) standpoint it will, in most cases be preferable to common the Disk Enclosure to the system's mounting frame. With this in mind, it is important that the Disk Enclosure not become an excessive return current path from the system frame to power supply. The drive's mounting frame must be within +/- 150 millivolts of the drive's power supply ground. At no time should more than 35 milliamps of current (0 to 100Mhz) be injected into the disk encloure.

Please contact your IBM Customer Representative if you have questions on how to integrate this drive in your system.

4.2.6 Hot plug/unplug support

Power supply and SCSI bus hot plug and un-plug is allowed for the 50 and 68 pin connector drives only. Hot plugging the 80 pin single connector model is not recommended since provisions are not available to determine which electrical contact is made first. There is no special sequence required for connecting 5 or 12 volts. During a hot plug-in event the drive being plugged will draw a large amount of current at the instant of plug-in. This current spike is due to charging the bypass capacitors on the drive. This current pulse may cause the power supply to go out of regulation. If this supply is shared by other drives then a low voltage power on reset may be initiated on those drives. Therefore the recommendation for hot plugging is to have one supply for each drive. Never daisy chain the power leads if hot plugging is planned. Hot plugging should be minimized to prevent wear on the power connector and Tantalum capacitor (surge current) stress.

Hot plugging the SCSI bus may cause glitches on the bus. To minimize the chance of glitching, it is recommended to plug in the SCSI bus before the power is applied. In addition proper ESD procedures should be followed prior to plugging the SCSI bus to insure that the Drive ground is at the same voltage potential as the SCSI bus Ground.

During hot plugging, the supplies must not go over the upper voltage limit.

During hot un-plugging if the operating shock limit specification can be exceeded then the drive should be issued a SCSI Stop Unit command that is allowed to complete before un-plugging.

[11] This ripple must not cause the power supply to the drive to go outside of the +/-5 % regulation tolerance.

4.2.7 Additional 5V Power Requirements for Differential

The following current values are measured values. Safety factors have not been applied.

4.2.7.1 Additional Current Required for 68 pin models:

Additional Power Supply Current	Notes	Population Mean
+5VDC (idle avg)		0.10 Amps
+5VDC (R/W baseline)	Typical	0.06 Amps
+5VDC (R/W pulse)	Base-to-peak	1.0 Amps

As a reminder, when determining the current necessary to supply the differential cards, only one drive per SCSI bus may be selected therefore only one drive per SCSI bus may be sinking the maximum 5V current.

4.2.8 Bring-up Sequence (and Stop) Times

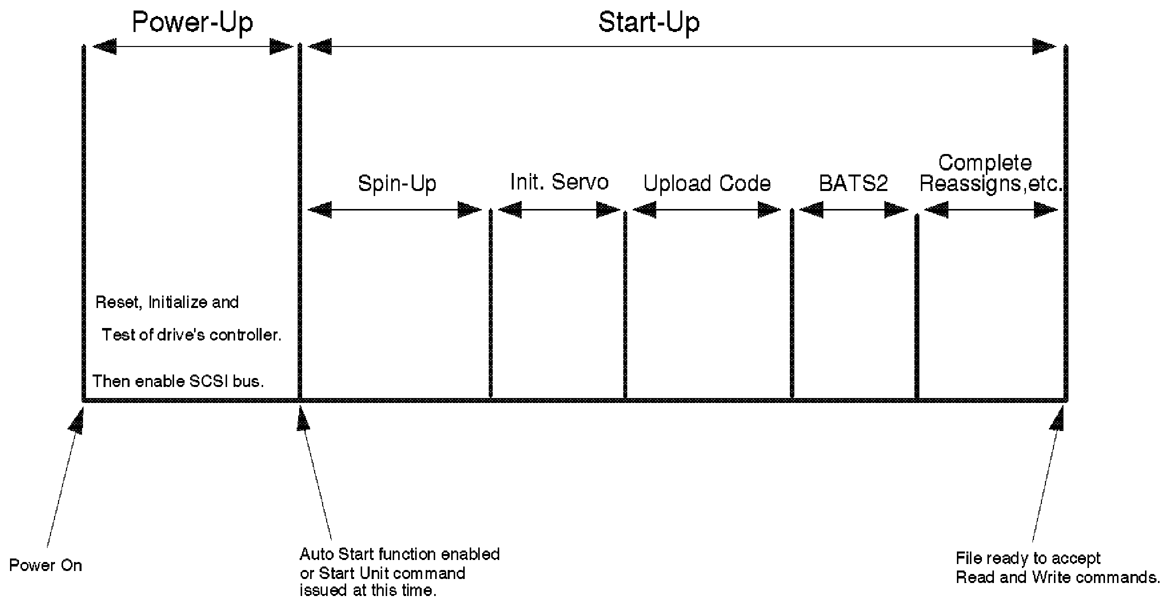

Figure 11. Start Time Diagram

Note: BATS is the abbreviation for Basic Assurance Tests. Start-up sequence spins up the spindle motor, initializes the servo subsystem, uploads code, performs BATS2 (verifies read/write hardware), resumes "Reassign in Progress" operations, and more. For more information on the startup sequence, refer to the **DFHS Interface Specification**.

Note: If a RESET is issued before the drive comes ready, the power on sequence will start again. In all other cases when a RESET is issued the present state of the motor is not altered.

Note: Reference 5.10.1.1.3, “Start/Stop Unit Time” on page 55 for additional details.

Note: See 7.4, “Spindle Synchronization” on page 84 for details about Start-up time increases when the device is requested via Mode Parameters to synchronize the spindle motor to another device.

Note: A 12 second timer is started when a motor stop command is issued. A motor start command must wait until that time has expired before attempting to start the motor. Therefore if you issue a stop motor command and then immediately issue a start motor command it will take an additional 12 seconds.

Event	Nominal	Maximum	Notes
Power-up	1.5 sec	2.0 sec	*See Figure 11
Start-up	12.4 sec	45 sec	*See Figure 11
Spin-up	8.2 sec	29.2 sec	*See Figure 11
Spindle Stop	6.0 sec	12.0 sec	

Table 3. Bring-up Sequence Times and Stop Time for S1x Models

Event	Nominal	Maximum	Notes
Power-up	1.5 sec	2.0 sec	*See Figure 11
Start-up	17.6 sec	45 sec	*See Figure 11
Spin-up	13.2 sec	29.2 sec	*See Figure 11
Spindle Stop	9.0 sec	12.0 sec	

Table 4. Bring-up Sequence Times and Stop Time for S2x Models

Event	Nominal	Maximum	Notes
Power-up	1.5 sec	2.0 sec	*see Figure 11 on page 36
Start-up	16.5 sec	45 sec	*see Figure 11 on page 36
Spin-up	11.7 sec	30.9 sec	*see Figure 11 on page 36
Spindle Stop	8.0 sec	12.0 sec	

Table 5. Bring-up Sequence Times and Stop Time for S4x Models

5. Performance

Drive performance characteristics are dependant upon the workloads run and the environments in which they are run.

All times listed in this chapter are typical values provided for information only, so that the performance for environments and workloads other than those shown as examples can be approximated. Actual minimum and maximum values will vary depending upon factors such as workload, logical and physical operating environments.

5.1 Environment Definition

Drive performance criteria is based on the following operating environments. Deviations from these environments may cause deviations from values listed in this specification.

- Block lengths are formatted at 512 bytes per block.
- The number of data buffer cache segments is 8. The total data buffer length is 512k bytes. Each segment is of equal length. Therefore each cache segment is 64k bytes.

  The number of blocks of customer data that can fit into one segment is reduced because 2 bytes of LRC information is also stored in the segment for each block of customer data stored in the segment. Therefore, use the following equation to determine how many blocks can fit into one segment.

$$\frac{\left(\dfrac{512\text{KB}}{\textbf{\# of segments}}\right)}{\textbf{ub/lba} + 2}$$

- Ten byte SCSI Read and Write commands are used.
- SCSI environment consists of a single initiator and single target with no SCSI Bus contention.
- Buffer full/empty ratios are set to their optimum values such that a minimum number of intermediate disconnects occur during the SCSI data transfer and the overlap of the SCSI and disk data transfer is maximized. This minimizes Command Execution Times with no bus contention.
- Read Caching and Read Ahead functions are enabled except where noted.
- The initiator delay while transferring SCSI command, status, message, and data bytes is assumed to be zero.
- The media is formatted with the skew definition that optimizes the disk data transfer rate for un-synchronized spindle operation.
- Tagged Command Queuing is not used, unless otherwise specified.
- All Current Mode Parameters are set to their Default values except where noted.
- SCSI data transfers are successfully negotiated to be 20 MB/sec.
- Averages are based on a sample size of 10,000 operations.

5.2 Workload Definition

The drive's performance criteria is based on the following command workloads. Deviations from these workloads may cause deviations from this specification.

- Operations are either all Reads or all Writes. The specifications for Command Execution Time with Read Ahead describe exceptions to this restriction. For that scenario all commands are preceded by a Read command, except for sequential write commands.
- The time between the end of an operation, and when the next operation is issued is 50 msec, +/- a random value of 0 to 50 msec, unless otherwise noted.

5.2.1 Sequential

- No Seeks. The target LBA for all operations is the previous LBA + Transfer Length.

5.2.2 Random

- All operations are to random LBAs. The average seek is an average weighted seek.

5.3 Command Execution Time

Command Execution, or Service, Times are the sum of several Basic Components. Those Components are -

1. Seek
2. Latency
3. Command Execution Overhead
4. Data Transfer to/from Disk
5. Data Transfer to/from SCSI Bus

The impact or contribution of those Basic Components to Command Execution Time is a function of the workload being sent to the drive and the environment in which the drive is being operated.

The following graphs show Command Execution Times for four generic workloads

- Sequential Reads
- Sequential Writes
- Random Reads
- Random Writes

with several different requested Transfer Lengths while running in various environments whose key factors are identified within each graph.

Note: Times are calculated with Typical Data Sector Transfer Rates for S4x models and are averaged over the entire drive.

. **Note:** Below, "TCQing" means Tagged Command Queueing and "Qd" is the average number of commands
. queued by the drive at one time.

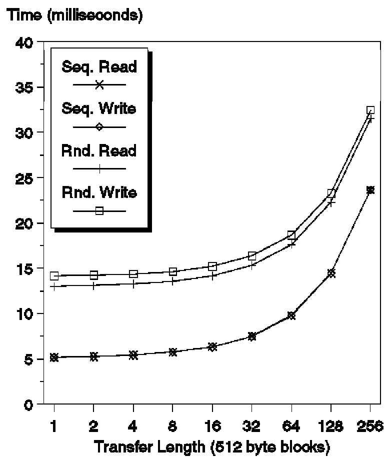

. Figure 12. All caching disabled, Re-instruction Times
. = 0-100ms, no TCQing

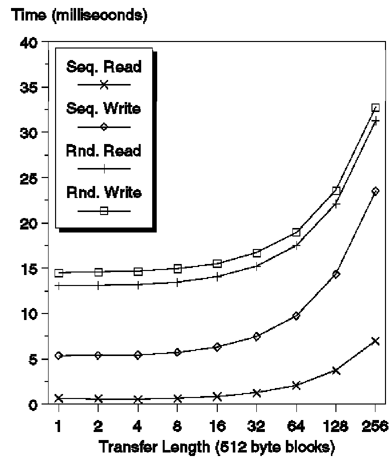

Figure 13. Read Caching enabled, Re-instruction Times = 50-150ms, no TCQing

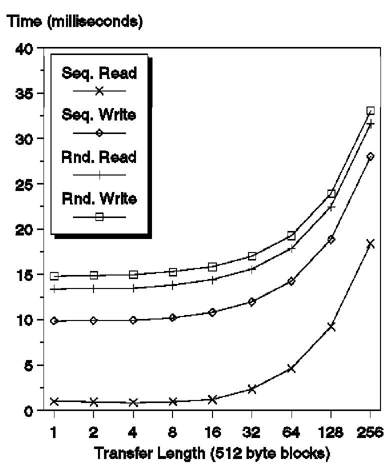

. Figure 14. Read Caching enabled, Re-instruction Times
. = 0ms, no TCQing

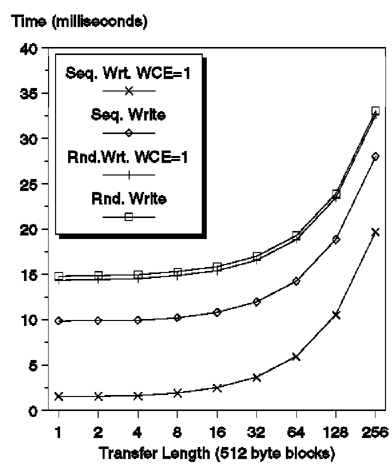

Figure 15. Read Caching enabled, Re-instruction Times = 0ms, no TCQing

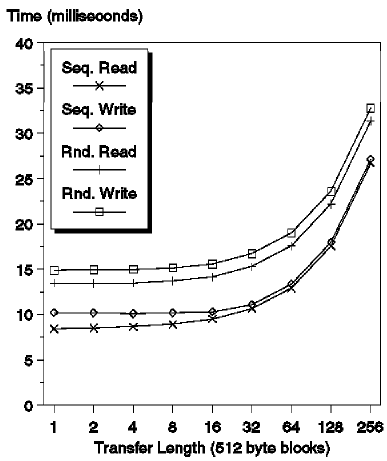

. Figure 16. All caching disabled, TCQing used

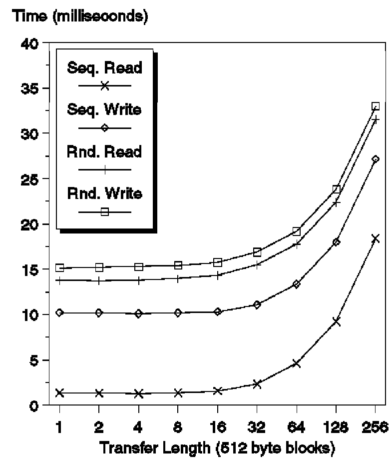

Figure 17. Read Caching enabled, TCQing used

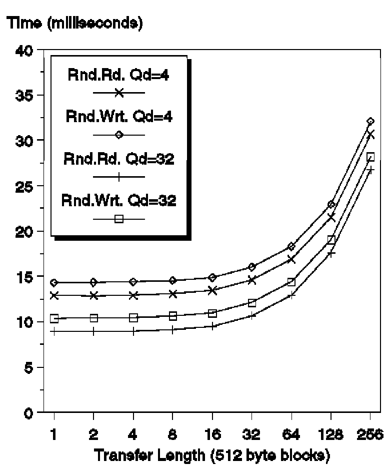

. Figure 18. Read Caching & Unrestricted Command
. Reordering enabled, TCQing

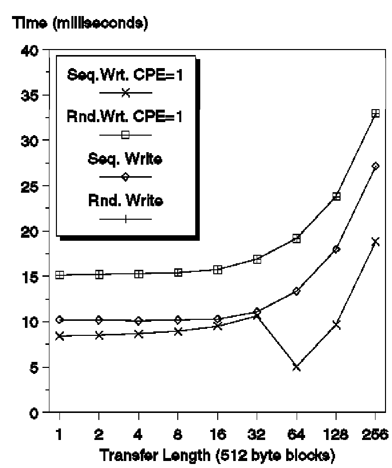

Figure 19. Read Caching & Concurrent Command
Processing enabled, TCQing

5.3.1 Basic Component Descriptions

Seek

The average time from the initiation of the seek, to the acknowledgement that the R/W head is on the track that contains the first requested LBA. Values are population averages, and vary as a function of operating conditions. The values used in the graphs showing Command Execution Times for sequential commands is 0 ms and the values for random commands are shown in section 4, "Specifications" on page 11.

Latency

The average time required from the activation of the read/write hardware until the target sector has rotated to the head and the read/write begins. This time is 1/2 of a revolution of the disk, or 4.17 ms.

Command Execution Overhead

The average time added to the Command Execution Time due to the processing of the SCSI command. It includes all time the drive spends not doing a disk operation or SCSI data transfer, whether or not it is connected to the bus. (See 5.6, "Read Command Performance" on page 49 and 5.7, "Write Command Performance" on page 51 for examples of detailed descriptions of the components of Command Execution Overhead.) The value of this parameter varies greatly depending upon workloads and environments.

The following values are used when calculating the Command Execution Times.

Workload	Command Execution	SCSI Bus
Sequential Read (wo/RA - w/RA)	.85 - .41	.03
Sequential Write	.88	.04
Random Read (wo/RA - w/RA)	.28 - .58	.03
Random Write	.38	.04
Note: "w/RA" means a Read Ahead operation is in progress when command is received. "wo/RA" means a Read Ahead is not in progress.		

Table 6. Overhead Values. (All times are in milliseconds.)

Other Initiator controlled factors such as use of disconnects, Tagged Command Queueing and the setting of Mode Parameters like DWD, DRD, DPSDP and ASDPE also affect Command Execution Overhead. They also affect SCSI Bus Overhead which is partially a subset of Command Execution Overhead.

SCSI Bus Overhead is defined as the time the device is connected to the bus transferring all SCSI Command, Status and Message phase information bytes. This includes any processing delays between SCSI Bus phases while remaining connected to the SCSI Bus. Initiator delays while transferring information bytes are assumed to be zero. This time does not include the SCSI Data phase transfer. (See 5.6, "Read Command Performance" on page 49 and 5.7, "Write Command Performance" on page 51 for more detailed descriptions of the components of SCSI Bus Overhead.)

Post Command Processing time of .33 ms is defined as the average time required for process cleanup after the command has completed. If a re-instruct period faster than this time is used, the difference is added to the Command Execution Overhead of the next operation.

Data Transfer to/from Disk

The average time used to transfer the data between the media and the drive's internal data buffer. This is calculated from:

(Data Transferred)/(Media Transfer Rate).

There are four interpretations of Media Transfer Rate. How it is to be used helps decide which interpretation is appropriate to use.

1. Instantaneous Data Transfer Rate

   The same for a given notch formatted at any of the supported logical block lengths. It varies by notch only and does not include any overhead.

2. Track Data Sector Transfer Rate

   Varies depending upon the formatted logical block length and varies from notch to notch. It includes the overhead associated with each individual sector. This is calculated from:

   (user bytes/sector)/(individual sector time)

   (Contact an IBM Customer Representative for individual sector times of the various formatted block lengths.)

   Note: These rates are used to help estimate optimum SCSI Buffer Full/Empty Ratios.

3. Theoretical Data Sector Transfer Rate

   Also includes time required for track and cylinder skew and overhead associated with each track.

   Each group of cylinders with a different number of gross sectors per track is called a notch. The following shows how the value for notch #1 for S4x models is calculated. For the other notches and block lengths use values that correspond to those notches and block lengths.

   Data Sector Transfer Rate =

   $$\frac{\text{Bytes/cylinder}}{\text{time for 1 cyl} + \text{track skews} + \text{1 cyl skew}}$$

   Bytes/cylinder = {(tracks/cyl)(gross sectors/track) - spares/cyl}(user bytes/sector)
   = {(16)(135) - 40}(512)
   = 1,085,440 Bytes/cyl

   time for 1 cyl of data = {(tracks/cyl)(gross sectors/track) - spares/cyl}(avg. sector time)
   = {(16)(135) - 40}(.061705)
   = 130.815 msec/cyl

   time for track skews = (tracks/cyl - 1)(track skew)(avg. sector time)
   = (16-1)(13)(.061705)
   = 12.032 msec/cyl

   time for 1 cyl skew = (cylinder skew)(avg. sector time)
   = (25)(.061705)
   = 1.543 msec/cyl

   Data Sector Transfer Rate =

   $$\frac{\text{1,085,440 Bytes}}{\text{130.815 msec} + \text{12.032 msec} + \text{1.543 msec}}$$

   = 7.517 MB/sec (Notch #1)

   Note: See 4, "Specifications" on page 11 for the descriptions of

   - tracks/cyl (trk/cyl)
   - gross sectors/track (gs/trk)
   - spares/cyl (b1spr/cyl and b2spr/cyl)
   - user bytes/sector (ub/sct)
   - gross bytes/sector (gb/sct)

   See 5.8, "Skew" on page 52 for the descriptions of

   - track skew (tss)
   - cylinder skew (css)

Average sector times per notch can be calculated as follows:

- average sector time (ast) =

$$\frac{1 \text{ sec}}{120.045 \times \text{gs/trk}}$$

4. Typical Data Sector Transfer Rates

   Also includes the effects of defective sectors and skipped revolutions due to error recovery. (See Appendix B. of the **DFHS Interface Specification** for a description of error recovery procedures.)

Rates for drives formatted at 512 bytes/block are located in Table 7.

Model Type	**All**		**S4x**		**S2x**		**S1x**	
Notch #	**Instantaneous**	**Track**	**Theoretical**	**Typical**	**Theoretical**	**Typical**	**Theoretical**	**Typical**
Average	12.07	7.91	7.17	7.13	7.13	7.10	7.06	7.03
1	12.58	8.30	7.52	7.48	7.48	7.44	7.40	7.37
2	12.58	8.30	7.52	7.48	7.48	7.44	7.40	7.37
3	12.51	7.99	7.22	7.18	7.18	7.15	7.11	7.08
4	11.96	7.74	7.02	6.99	6.99	6.95	6.92	6.89
5	11.26	7.38	6.66	6.63	6.63	6.60	6.57	6.54
6	11.05	7.07	6.41	6.38	6.38	6.35	6.31	6.28
7	10.64	6.88	6.23	6.20	6.19	6.16	6.13	6.10
8	10.29	6.64	6.03	6.00	6.00	5.97	5.94	5.91
9	10.01	6.45	5.85	5.83	5.83	5.80	5.77	5.74
10	9.59	6.15	5.55	5.53	5.53	5.50	5.48	5.45

Note: The values for Typical Data Sector Transfer Rates assume a typically worst case value of 3.16 errors in 10^9 bits read at nominal conditions for soft error rate.

Note: Contact an IBM Customer Representative for values when formatted at other block lengths.

Note: "Average" values are sums of the individual notch values weighted by the number of LBAs in the associated notches.

Table 7. Data Sector Transfer Rates. (All rates are in MB/sec)

Data Transfer to/from SCSI Bus

The time required to transfer data between the SCSI bus and the drive's internal data buffer, that is not overlapped with the time for the Seek, Latency or Data Transfer to/from Disk. This time is based on a SCSI synchronous data transfer rate of 20.0 MB/sec.

The SCSI data transfer rate is dependent on the mode, either synchronous or asynchronous. It also depends upon the width of the data path used. 8 and 16 bit transfers are supported.

When the drive is configured for an 8 bit wide transfer a synchronous data transfer rate of 10 MB/sec can be realized. The 16 bit wide maximum synchronous data transfer rate is 20 MB/sec.

The asynchronous data transfer rate is dependent on both the initiator and target delays to the assertion and negation of the SCSI REQ and ACK signals. It is also dependent on SCSI cable delays. The drive is capable of supporting up to 5 MegaTransfer/sec asynchronous data transfer rates.

The SCSI data transfer rate specification only applies to the Data phase for logical block data for Read, Write, Write and Verify, etc... commands. The data rate for parameter/sense data for Request Sense, Mode Select, etc... commands is not specified.

5.3.2 Comments

Overlap has been removed from the Command Execution Time calculations. The components of the Command Execution Times are truly additive times to the entire operation. For example,

- The SCSI Bus Overhead data is not included in the calculation since some of it's components are also components of Command Execution Overhead and the remaining components overlap the Data Transfer to/from Disk. (See 5.6, “Read Command Performance” on page 49 and 5.7, “Write Command Performance” on page 51 for details.)
- The Post Command Processing times are not components of the Command Execution time therefore they are not included in the calculation of environments were the re-instruct period exceeds the Post Command Processing time.

With Read Ahead enabled, this specification measures a Read or Write command when the immediately preceding command is a Read command (which starts up the Read Ahead function). If the preceding command is a Write command, then the time difference due to Read Ahead is zero.

Longer inter-op delay, or low re-instruction rate, environments are such that the Read Ahead function has filled the drive's internal data cache segment before the next Read or Write command is received.

Environments with inter-op delays less than 1 revolution period, or high re-instruction rates, are such that the Read Ahead function is still in the process of filling the drive's internal data cache segment when the next Read or Write command is received. For sequential reads, Command Execution Time is 1 revolution less than similar operations with equal inter-op delays and Read Ahead disabled.

The effects of idle time functions are not included in the above examples. The 5.2.1, “Sequential” on page 40 and 5.2.2, “Random” on page 40 both define environments where the effects due to increased command overhead of Idle Time Functions upon Command Execution time are less than 0.15 %.

5.4 Disconnection During Read/Write Data Phase

If a nonzero Maximum Burst Size parameter is specified, the drive disconnects after transferring the number of blocks specified by the Maximum Burst Size parameter. This disconnection requires approximately 33 µsec and the subsequent reconnection requires approximately 20 µsec.

The drive also disconnects prior to completion of the Data phase if the drive's internal data buffer cache segment becomes empty during a Read command or full during a Write command. This disconnection occurs regardless of the Maximum Burst Size parameter. This intermediate disconnect causes a pause of approximately 0.24 msec during the Data phase. This disconnection requires approximately 33 µsec and the subsequent reconnection requires approximately 20 µsec.

5.5 Effects of Different Environments

5.5.1 When Read Caching is Enabled

For read commands with Read Caching enabled Command Execution time can be approximated by deleting Seek, Latency and Data Transfer to/from Disk times from those shown on the graphs if all of the requested data is available in a cache segment (cache hit). When some, but not all, of the requested data is available in a cache segment (partial cache hit) Data Transfer to/from Disk will be reduced but not eliminated. Seek and Latency may or may not be reduced depending upon the location of requested data not in the cache and location of the read/write heads at the time the command was received. The contribution of the Data

Transfer to/from SCSI Bus to the Command Execution time may increase since a larger, or entire, portion of the transfer may no longer be overlapped with the components that were reduced.

5.5.2 When Write Caching is Enabled

For write commands with the Write Caching Enabled (WCE) Mode parameter bit set, Command Execution time can be approximated by deleting Seek, Latency and Data Transfer to/from Disk times from those shown in the graphs. The contribution of the Data Transfer to/from SCSI Bus to the Command Execution time may increase since a larger, or entire, portion of the transfer may no longer be overlapped with the components that were reduced.

The reduced times effectively are added to the Post Command Processing Time.

Like Tagged Command Queuing, the potential to reduce Command Execution Overhead exists due to concurrent command processing.

Like Tagged Command Queuing, when the WCE bit is set Back-To-Back write commands are supported. See 5.5.4.2, “Back-To-Back Write Commands” on page 48 for more information.

5.5.3 When Adaptive Caching is Enabled

The Adaptive Caching feature attempts to increase Read Cache hit ratios by monitoring workload and adjusting cache control parameters, normally determined by the using system via the SCSI Mode Parameters, with algorithms using the collected workload information.

5.5.4 For Queued Commands

The effects of Command Execution Overhead can be reduced significantly if Tagged Command Queuing is enabled since more than 1 command can be operated on concurrently. For instance, while a disk operation is being performed for one command another command can be received via the SCSI bus and placed in the device command queue. Certain environments may cause Command Execution Overhead to increase if the added function to process the queue and the messages associated with queueing is not permitted to overlap with a disk operation.

5.5.4.1 Reordered Commands

If the Queue Algorithm Modifier Mode Parameter field is set to allow it, commands in the device command queue may be executed in a different order than they were received. Commands are reordered so that the Seek portion of Command Execution time is minimized. The amount of reduction is a function of the location of the 1st requested block per command and the rate at which the commands are sent to the drive.

A Queue Algorithm Modifier Mode Parameter value of 9 enables an algorithm that gives the using system the ability to place new commands into the drive command queue execution order relative to the outstanding commands in the queue. For example, if a request is sent to the drive that the using system prioritizes such that it's completion time is more important than one or more of the outstanding commands, the using system can increase the likelihood that command is executed before those others by using a tag value greater than those outstanding commands. See the **DFHS Interface Specification** for more details about this reorder algorithm.

5.5.4.2 Back-To-Back Write Commands

If all of the requirements are met as stated in the **DFHS Interface Specification** section describing Back-To-Back write commands, contiguous data from 2 or more consecutive write commands can be written to the disk without requiring any disk Latency.

Note: There is a minimum write command transfer length for a given environment where continuous writing to the disk can not be maintained without missing a motor revolution. When Write Caching is enabled the likelihood is increased that shorter transfer write commands can fulfill the requirements needed to maintain continuous writing to the disk.

5.6 Read Command Performance

Note: This case is for Random SCSI Read commands, with Read Ahead disabled.

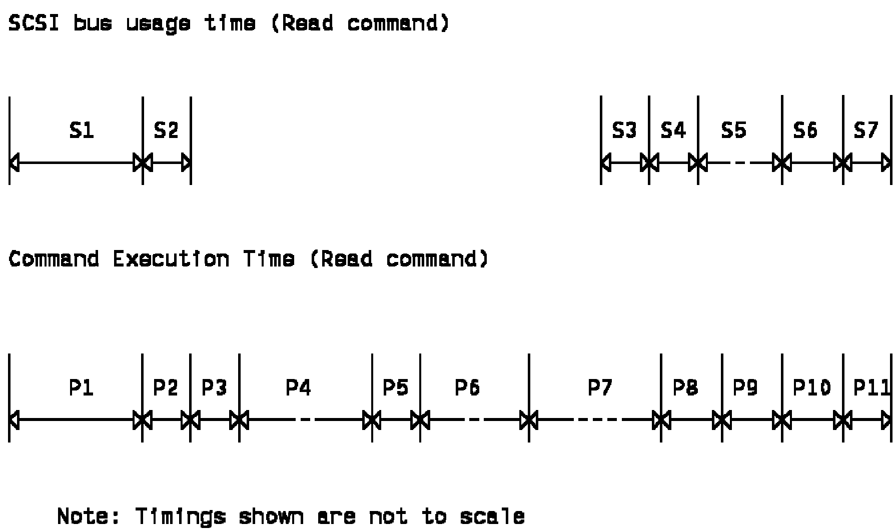

Figure 20. SCSI Read command performance measurements

5.6.1.1 SCSI Bus Overhead

Note: All times listed in this section are provided for information only so that the performance for other environments/workloads can be approximated. These component times should not be measured against the specification.

S1	**Selection, Identify Msg., Command Descriptor Block (CDB)**	12 µsec
S2a	**Save Data Pointers (SDP) Msg.**	1 µsec
S2b	**Disconnect Msg., Bus Free**	1 µsec
S3	**Arbitrate, Reselect, Identify Msg.**	7 µsec
S4	**Start SCSI transfer in**	3 µsec
S5	**SCSI bus data transfer in**	(Transfer size)/(SCSI Data Transfer Rate)
S6	**SCSI read ending processing**	2 µsec
S7	**Status, Command Complete Msg., Bus Free**	2 µsec

Note: The SCSI bus overhead for a Read Command is composed of S1,S2(a&b),S3,S4,S6,and S7. (0.03 msec total).

5.6.1.2 Command Execution Overhead

: **P1**	**Selection, Identify Msg., CDB**	12 µsec
P2a	**SDP Msg.**	1 µsec
P2b	**Disconnect Msg., Bus Free**	1 µsec
P3	**Start seek or head switch**	258 µsec
: **P4**	**Seek or head switch (for example, average seek)**	(Read Seek = 6.9, 7.5 or 8.0 msec)
P5	**Set up read disk transfer**	0 µsec
P6	**Latency (for example, half revolution)**	(latency = 4.17 msec)
P7	**Disk data transfer**	(Data transferred)/(Typical Data Sector Transfer Rate)
P8	**End read disk transfer**	(Sector size)/(SCSI Data Transfer Rate)
P9	**Transfer last few SCSI blocks in**	(5)(Sector size)/(SCSI Data Transfer Rate)
P10	**SCSI read ending processing**	2 µsec
: **P11**	**Status, Command Complete Msg., Bus Free**	2 µsec

Note: The Command execution overhead for a read command is composed of P1,P2(a&b),P3,P5,P10,and P11. (0.28 msec total).

Time to Read data = P1+P2+P3+P4+P5+P6+P7+P8+P9+P10+P11

5.7 Write Command Performance

Note: This case is for Random SCSI Write commands, with Read Ahead disabled.

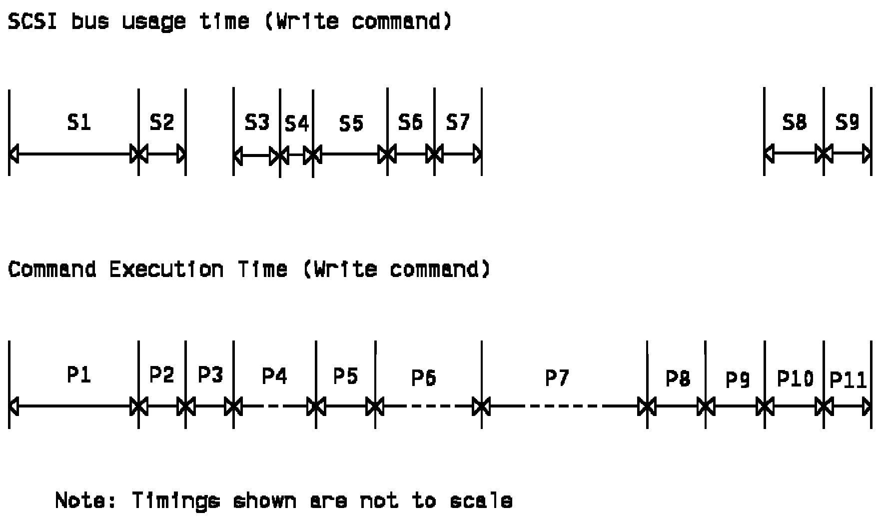

Figure 21. SCSI Write command performance measurements

5.7.1.1 SCSI Bus Overhead

Note: All times listed in this section are provided for information only so that the performance for other environments can be approximated. These component times should not be measured against the specification.

:	**S1 Selection, Identify Msg., CDB**	12 µsec
	S2a SDP Msg.	1 µsec
	S2b Disconnect Msg., Bus Free	1 µsec
:	**S3 Arbitrate, Reselect, Identify Msg.**	7 µsec
:	**S4 start SCSI transfer out**	2 µsec
	S5 SCSI bus data transfer out	(Transfer size)/(SCSI Data Transfer Rate)
:	**S6 End SCSI transfer out**	3 µsec
	S7A SDP Msg.	1 µsec
	S7B Disconnect Msg., Bus Free	1 µsec
:	**S8 Arbitrate, Reselect, Identify Msg.**	9 µsec
:	**S9 Status, Command Complete Msg., Bus Free**	2 µsec

Note: The SCSI bus overhead for a write command is composed of S1,S2(a&b),S3,S4,S6,S7,S8 and S9. (0.04 msec total).

5.7.1.2 Command Execution Overhead

P1 Selection, Identify Msg., CDB	12 µsec
P2a SDP Msg.	1 µsec
P2b Disconnect Msg., Bus Free	1 µsec
P3 Start seek	258 µsec
P4 Seek (for example, average seek)	(Write Seek = 9.0, 9.0 or 9.5 msec)
P5 Set up write disk transfer	0 µsec
P6 Latency (for example, half revolution)	(Latency = 4.17 msec)
P7 Disk data transfer	(Data transferred)/(Typical Data Sector Transfer Rate)
P8 End write disk transfer	75 µsec
P9 SCSI write ending processing	25 µsec
P10 Arbitrate, Reselect, Identify Msg.	9 µsec
P11 Status, Command Complete Msg., Bus Free	2 µsec

Note: The Command execution overhead for a write command is composed of P1, P2(a&b), P3, P5, P8, P9, P10 and P11. (0.38 msec total).

Time to Write data = P1 + P2 + P3 + P4 + P5 + P6 + P7 + P8 + P9 + P10 + P11

5.8 Skew

5.8.1 Cylinder to Cylinder Skew

Cylinder skew is the sum of the sectors required for physically moving the heads (csms), which is a function of the formatted block length and recording density (notch #), and reassign allowance sectors (ras = 3) used to maintain optimum performance over the normal life of the drive.

Note: The values in the SCSI Mode Page 3 'Cylinder Skew Factor' are notch specific non-synchronized spindle mode values. The value for notch 1 is returned when the Active Notch is set to 0.

	Notch #									
User bytes / logical block	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**
256	39	39	39	37	35	33	33	33	33	29
512	25	25	24	23	22	21	21	20	19	18
520	23	23	23	23	21	21	20	19	19	18
522	23	23	23	22	21	21	20	19	19	17
524	23	23	23	22	21	21	20	19	19	17
528	23	23	23	22	21	21	20	19	19	17
688	19	19	19	18	17	17	17	17	15	15
744	18	18	18	17	17	17	15	15	14	14
Note: Contact an IBM Customer Representative for values at other formatted block lengths.										

Table 8. Optimal Cylinder Skew for several block lengths

In order to increase the likelihood that equivalent LBA's on two or more devices are located at the same relative physical position when the devices are used in a synchronized spindle mode, cylinder skew is calculated differently. The cylinder skew calculations do not take into account known defective sites. To prohibit revolutions from being missed on cylinder crossings by drives formatted while in a synchronized spindle mode, an extra allowance for 6 defects is added that is not added when optimally formatted in a non-synchronized mode.

5.8.2 Track to Track Skew

Note: The values in the SCSI Mode Page 3 'Track Skew Factor' are notch specific values. The value for notch 1 is returned when the Active Notch is set to 0.

	Notch #									
User bytes / logical block	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**
256	20	20	20	19	19	17	17	17	17	15
512	13	13	13	12	12	11	11	10	10	10
520	12	12	12	12	11	11	10	10	10	10
522	12	12	12	12	11	11	10	10	10	9
524	12	12	12	12	11	11	10	10	10	9
528	12	12	12	12	11	11	10	10	10	9
688	10	10	10	10	9	9	9	9	8	8
744	9	9	9	9	9	9	8	8	8	7
Note: Contact an IBM Customer Representative for values at other formatted block lengths.										

Table 9. Track (or Head) Skew for several block lengths

5.9 Idle Time Functions

The execution of various functions by the drive during idle times may result in delays of commands requested by SCSI initiators. 'Idle time' is defined as time spent by the drive not executing a command requested by a SCSI initiator. The functions performed during idle time are:

1. Servo Run Out Measurements
2. Servo Bias Measurements
3. Predictive Failure Analysis (PFA)
4. Channel Calibration
5. Save Logs and Pointers
6. Disk Sweep

The command execution time for SCSI commands received while performing idle time activities may be increased by the amount of time it takes to complete the idle time activity. Arbitration, Selection, Message and Command phases, and disconnects controlled by the drive are not affected by idle time activities.

Note: Command Timeout Limits do not change due to idle time functions.

All Idle Time Functions have mechanisms to lessen performance impacts for critical response time periods of operation. And in some cases virtually eliminate those impacts from an Initiator's point of view. All Idle Time Functions will only be started if the drive has not received a SCSI command for at least 5 seconds (40 seconds for Sweep). This means that multiple SCSI commands are accepted and executed without delay if the commands are received by the drive within 5 seconds after the completion of a previous SCSI command. This mechanism has the benefit of not requiring special system software (such as issuing SCSI Rezero Unit commands at known & fixed time intervals) in order to control if and when this function executes.

Note: Applications which can only accommodate Idle Time Function delays at certain times, but can not guarantee a 5 second reinstruction period, may consider synchronizing idle activities to the system needs through use of the LITF bit in Mode Select Page 0, and the Rezero Unit command. See the **DFHS Interface Specification** for more details.

Following are descriptions of the various types of idle functions, how often they execute and their duration. Duration is defined to be the maximum amount of time the activity can add to a command when no errors occur. No more than one idle function will be interleaved with each SCSI command.

Following the descriptions is a summary of the possible impacts to performance.

5.9.1 Servo Run Out Measurements

The drive periodically measures servo run out, the amount of wobble on each disk, to track follow more precisely.

Servo run out for all heads is measured every 60 minutes, therefore the frequency of run out measurements is dependent on the number of heads a particular model has. The drive attempts to spread the measurements evenly in time and each measurement takes 100 milliseconds.

```
For example, a model with 8 heads performs one run out
measurement every 7 1/2 minutes (60 / 8).
```

5.9.2 Servo Bias Measurements

The drive periodically measures servo bias, the amount of resistance to head movement as a function of disk radius. It also helps prevent disk lubrication migration by moving the heads over the entire disk surface.

Servo bias is measured every 12 minutes during the first hour after a power cycle, and every 60 minutes after that. The measurement takes 200 milliseconds.

5.9.3 Predictive Failure Analysis

Predictive ***F***ailure ***A***nalysis measures drive parameters and can predict if a drive failure is imminent.

Eight different PFA measurements are taken for each head. All measurements for all heads are taken over a period of 4 hours, therefore the frequency of PFA is dependent on the number of heads a particular model has. The drive attempts to spread the measurements evenly in time and each measurement takes about 80 milliseconds.

```
For example, a model with 8 heads will perform one PFA measurement
every 3.7 minutes (240 / 8*8).
```

For the last head tested for a particular measurement type (once every 1/2 hour), the data is analyzed and stored. The extra execution time for those occurrences is approximately 40 milliseconds.

This measurement/analysis feature can be disabled for critical response time periods of operation by setting the Page 0h Mode Parameter LITF = 1. The using system also has the option of forcing execution at known times by issuing the SCSI Rezero Unit command if the Page 0h Mode Parameter TCC = 1. All tests for all heads occur at those times. See the **DFHS Interface Specification** for more details about PFA, LITF and TCC.

5.9.4 Channel Calibration

The drive periodically calibrates the channel to insure that the read and write circuits function optimally, thus reducing the likelihood of soft errors.

Channel calibration is done once every 4 hours and typically completes in 20 milliseconds, but may take up to 64 milliseconds per measurement.

5.9.5 Save Logs and Pointers

The drive periodically saves data in logs in the reserved area of the disks. The information is used by the drive to support various SCSI commands and for the purpose of failure analysis.

Logs are saved every 35 minutes. The amount of time it takes to update the logs varies depending on the number of errors since the last update. In most cases, updating those logs and the pointers to those logs will occur in less than 30 milliseconds.

5.9.6 Disk Sweep

The heads are moved to another area of the disk if the drive has not received a SCSI command for at least 40 seconds. After flying in the same spot for 9 minutes, the heads are moved to another position. Execution time is less less than 1 full stroke seek.

5.9.7 Summary

Idle Time Function Type	Period of Occurance (minutes)	Duration (ms)	Mechanism to Delay/Disable
Servo Run Out	60/(trk/cyl)	100	Re-instruction Period
Servo Bias (< 1st hour)	12	200	Re-instruction Period
Servo Bias (> 1st hour)	60	200	Re-instruction Period
PFA	30/(trk/cyl)	80	Re-instruction Period / LITF
Channel Calibration	240	64	Re-instruction Period
Save Logs & Pointers	35	30	Re-instruction Period
Note: "Re-instruction Period" is the time between consecutive SCSI command requests.			

Table 10. Summary of Idle Time Function Performance Impacts

5.10 Command Timeout Limits

The 'Command Timeout Limit' is defined as the time period from the SCSI Arbitration phase through the SCSI Command Complete message, associated with a particular command.

The following times are for environments where Automatic Reallocation is disabled and there are no queued commands.

5.10.1.1.1 Reassignment Time: The drive should be allowed a minimum of 30 sec to complete a "Reassign Blocks" command.

5.10.1.1.2 Format Time: S4x models should be allowed a minimum of 45 minutes to complete a "Format Unit" command.

S2x models should be allowed a minimum of 25 minutes to complete a "Format Unit" command.

S1x models should be allowed a minimum of 15 minutes to complete a "Format Unit" command.

5.10.1.1.3 Start/Stop Unit Time: The drive should be allowed a minimum of 45 sec to complete a "Start/Stop Unit" command (with Immed bit = 0).

Initiators should also use this time to allow start-up sequences initiated by auto start ups and "Start/Stop Unit" commands (with Immed bit = 1) to complete and place the drive in a "ready for use" state.

Note: A timeout of one minute or more is recommended but NOT required. The larger system timeout limit allows the system to take advantage of the extensive ERP/DRP that the drive may attempt in order to successfully complete the start-up sequence.

Note: A 60 second minimum is required if electronics card replacement is required as a service practice. Please contact an IBM Customer Representative for more details if required.

5.10.1.1.4 Medium Access Command Time: The timeout limit for medium access commands that transfer user data and/or non-user data should be a minimum of 30 sec. These commands are:

- Log Sense
- Mode Select (6)
- Mode Sense (6)
- Pre-Fetch
- Read (6)
- Read (10)
- Read Capacity
- Read Defect Data
- Read Long
- Release
- Reserve
- Rezero Unit
- Seek (6)
- Seek (10)
- Send Diagnostic
- Skip Mask (Read)
- Skip Mask (Write)
- Write (6)
- Write (10)
- Write and Verify
- Write Buffer
- Write Long
- Write Same
- Verify

Note: The 30 sec limit assumes the absence of bus contention and user data transfers of 64 blocks or less. This time should be adjusted for anticipated bus contention and if longer user data transfers are requested.

5.10.1.1.5 Timeout limits for other commands: The drive should be allowed a minimum of 5 sec to complete these commands:

- Inquiry
- Request Sense
- Read Buffer
- Start/Stop Unit (with Immed bit = 1)
- Synchronize Cache
- Test Unit Ready

When Automatic Reallocation is enabled add 45 sec to the timeout of the following commands; Read (6), Read (10), Write (6), Write (10), Write and Verify, and Write Same.

The command timeout for a command that is not located at the head of the command queue should be increased by the sum of command timeouts for all of the commands that are performed before it is.

6. Mechanical

6.1 Weight and Dimensions

	S1x & S2x Models		S4x Models	
	U.S.	S.I. Metric	U.S.	S.I.Metric
Weight	1.00 pounds	0.46 kilograms	1.80 pounds	0.82 kilograms
Height	1.00 inches	25.4 millimeters	1.63 inches	41.3 millimeters
Width	4.00 inches	101.6 millimeters	4.00 inches	101.6 millimeters
Depth	5.75 inches	146.0 millimeters	5.75 inches	146.0 millimeters

6.2 Clearances

A minimum of 2 mm clearance should be given to the bottom surface except for a 10 mm maximum diameter area around the bottom mounting holes. Figure 22 and Figure 23 show the clearance requirements (Note 1). For proper cooling it is suggested that a clearance of 6 mm be provided under the drive and on top of the drive.

. There should be 7 mm of clearance between the drives that are mounted with their top sides (see Figure 39 on page 92 for top view of drive) facing each other.

6.3 Mounting

The drive can be mounted with any surface facing down.

The drive is available with both side and bottom mounting holes. Refer to Figure 22 to Figure 24 for the location of these mounting holes for each configuration.

The maximum allowable penetration of the mounting screws is 3.8 mm.

: The recommended torque to be applied to the mounting screws is 0.8 Newton-meters +/- 0.4
: Newton-meters. IBM will provide technical support to users that wish to investigate different mounting
: torques in their application.

WARNING: The drive may be sensitive to user mounting implementation due to frame distortion effects. IBM will provide technical support to assist users to overcome mounting sensitivity.

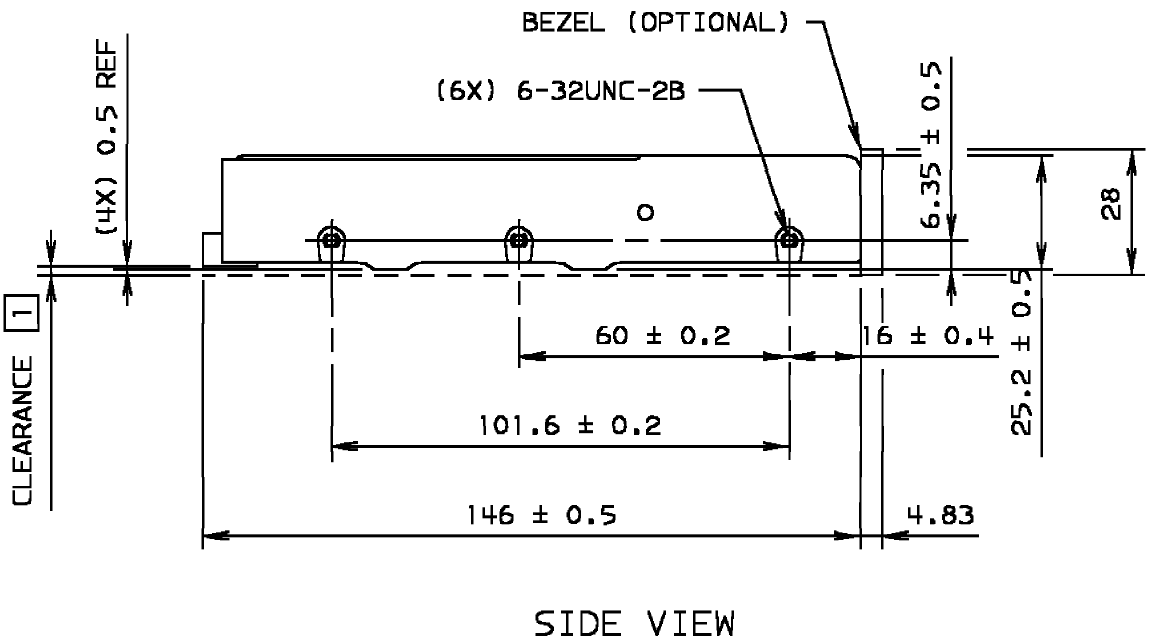

. **Notes:** 1) Bottom clearance required by 6.2, “Clearances” on page 57.

. 2) Dimensions are in millimeters.

Figure 22. Location of Side Mounting Holes of S1x & S2x Models

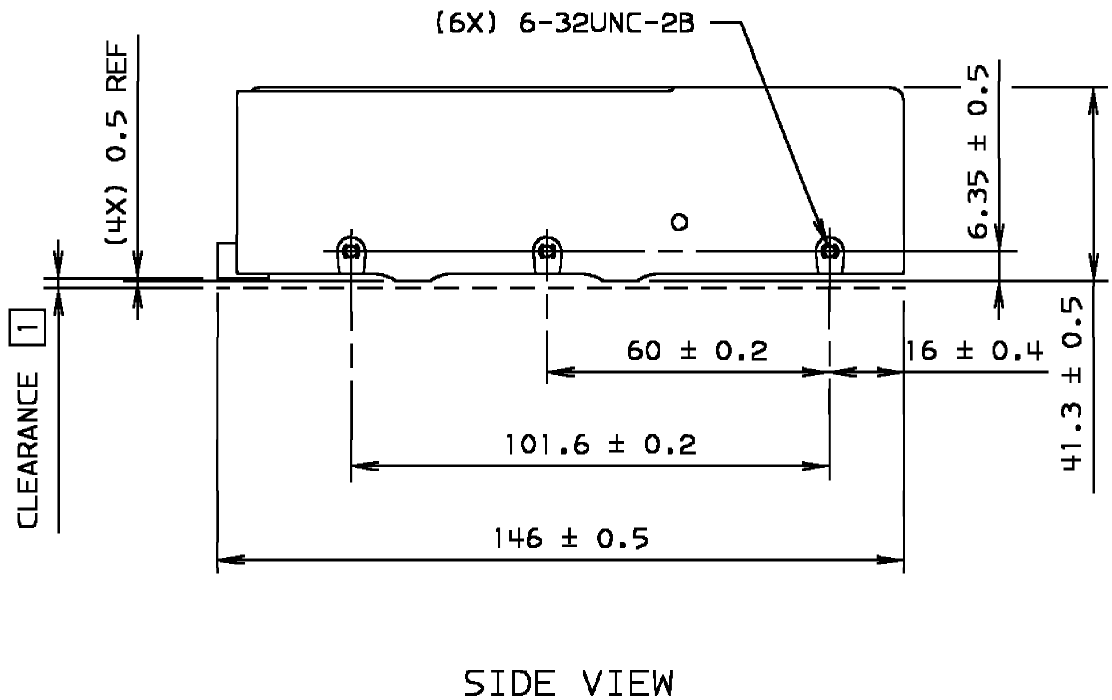

Notes: 1) Bottom clearance required by 6.2, “Clearances” on page 57.

2) Dimensions are in millimeters.

Figure 23. Location of Side Mounting Holes of S4x Models

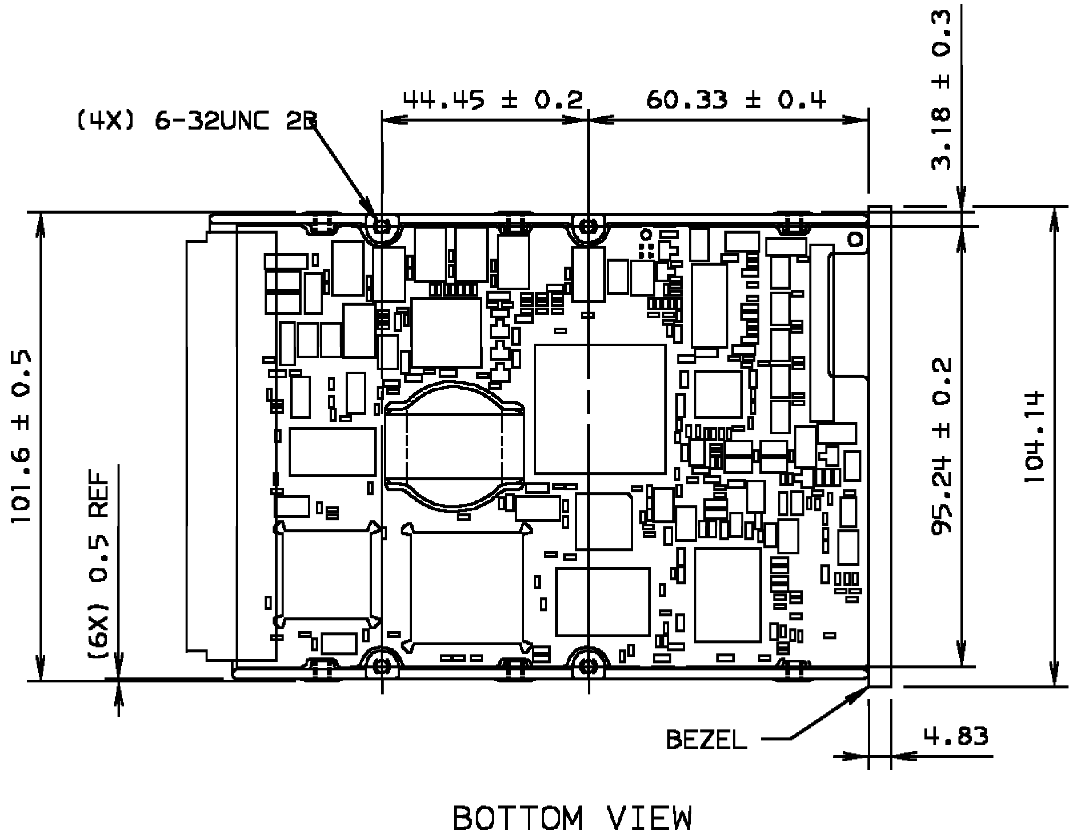

Notes: 1) The purpose of this drawing is to show the bottom hole pattern. The 50 pin connector is shown, other connectors will look different as shown on Figure 25 to Figure 28.

2) Dimensions are in millimeters.

Figure 24. Location of Bottom Mounting Holes of all Models

6.4 Electrical Connector Locations

The electrical connectors are located as shown in Figure 25 to Figure 28. The front jumper pin locations are shown in Figure 29

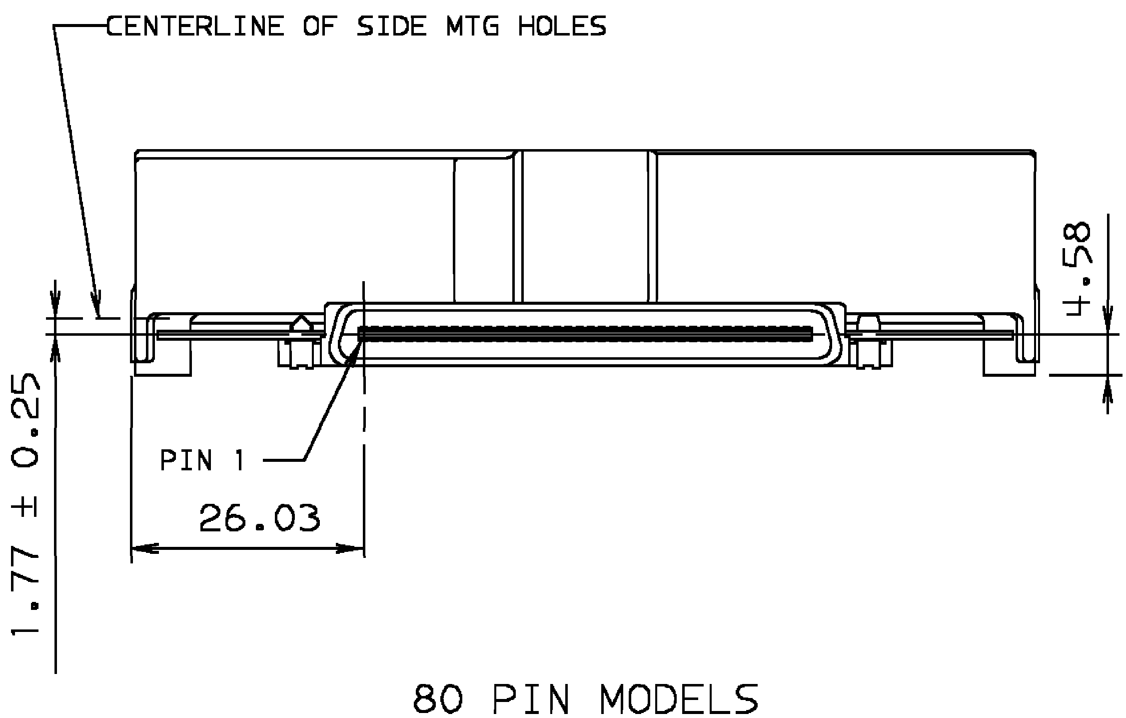

. **Notes:** 1) Dimensions are in millimeters.

. Figure 25. Electrical connectors (rear view) -- 80 pin SCA models.

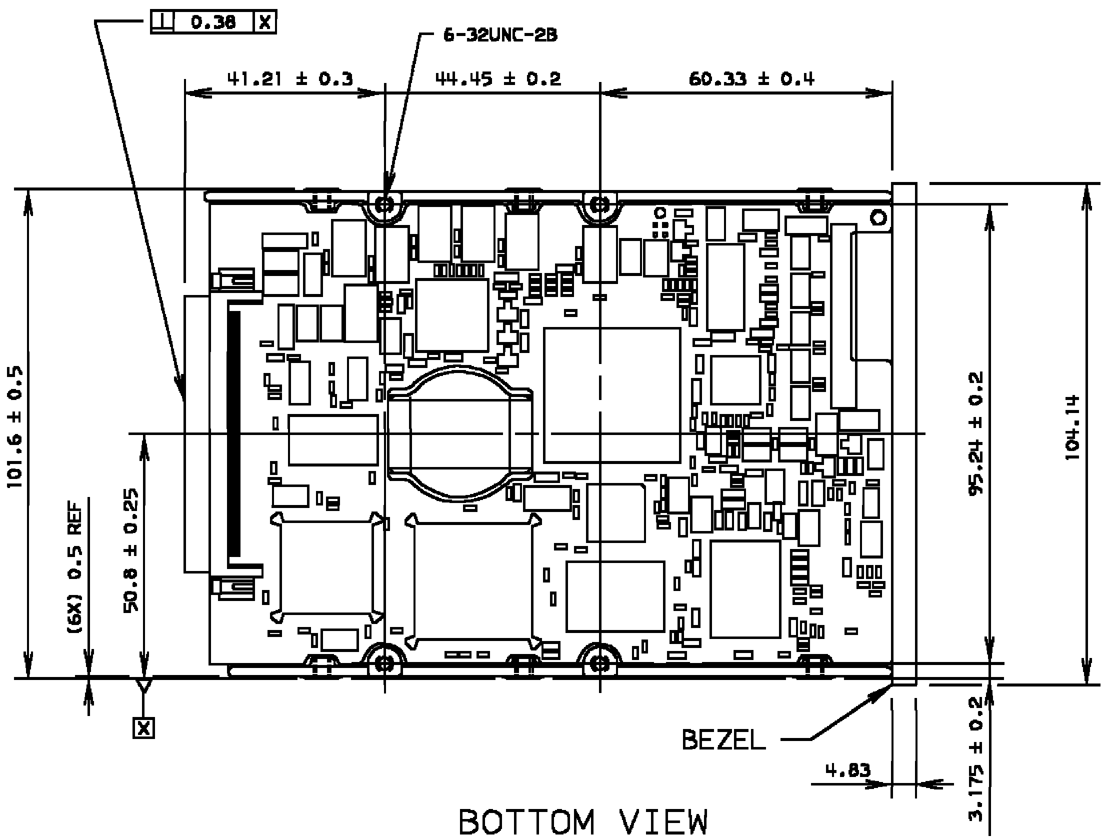

Notes: 1) Dimensions are in millimeters.

Figure 26. Electrical connectors (bottom view) -- 80 pin SCA models.

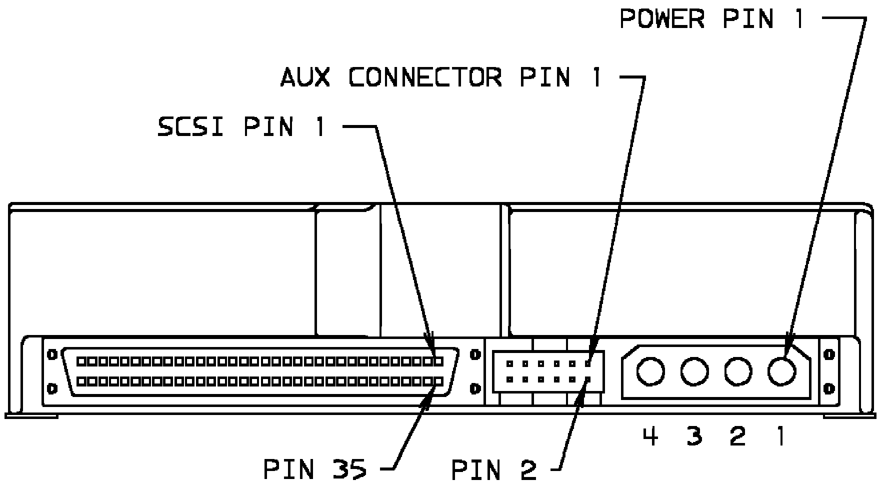

POWER PIN ASSIGNMENTS	
PIN	VOLTAGE
1	+ 12 V
2	GND
3	GND
4	+ 5 V

Figure 27. Electrical connectors (rear view) -- 68 pin models.

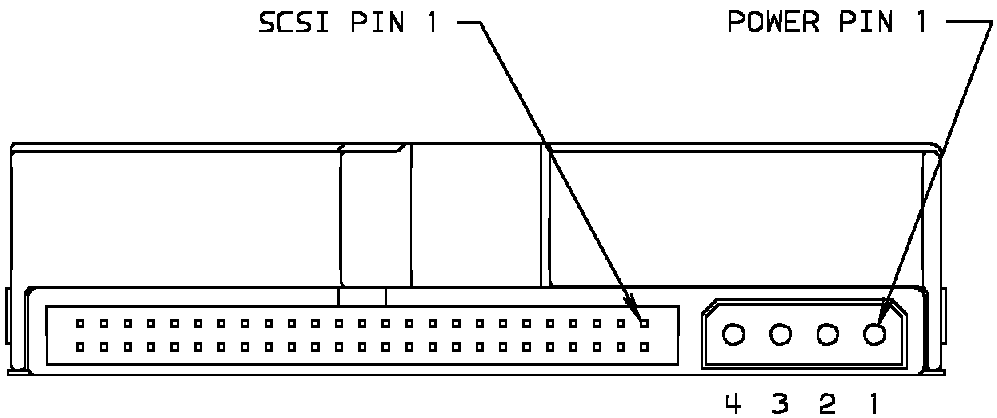

POWER PIN ASSIGNMENTS	
PIN	VOLTAGE
1	+ 12 V
2	GND
3	GND
4	+ 5 V

Figure 28. Electrical connectors (rear view) -- 50 pin models.

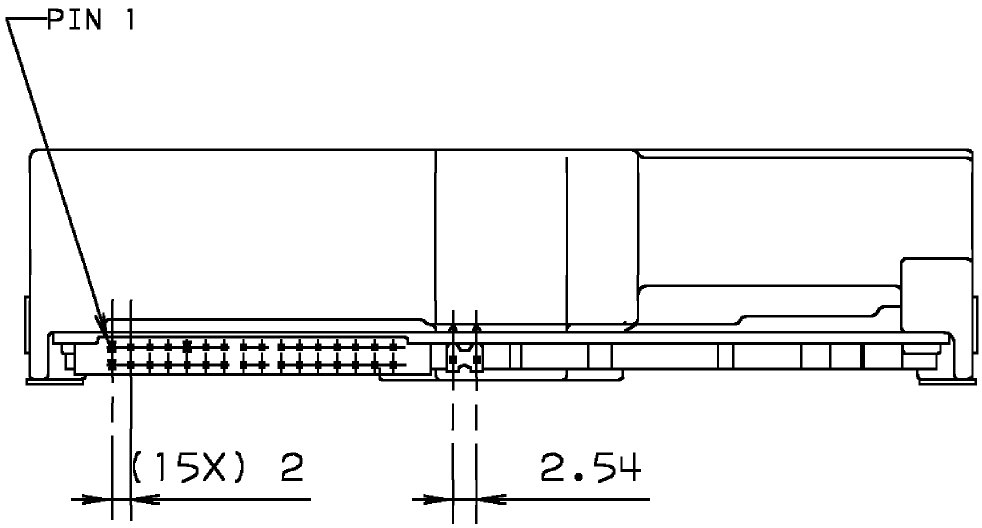

Notes: 1) Dimensions are in millimeters.

2) 80 pin models do not have the terminator power (2 pin) jumper.

Figure 29. Jumper pin locations (front view)

7. Electrical Interface

7.1 Power Connector

. The DC power connectors used on all models (50,68 and 80 pin SCA) are an integral portion of the 50, 68
. pin SCSI 'Unitized' Connectors or the 80 pin 'Single Connector Attachment' (SCA) Connector.

. 50 pin models use an AMP connector (PN 84160-1) that is compatible with the ANSI SCSI "A" connector
. specifications.

. 68 pin models use an AMP connector (PN 786963-1) that is compatible with the ANSI SCSI "P" connector
. specifications.

. 80 pin SCA models use a Molex connector (PN 87091-0001)
. that is compatible with the Specification of: 'Single Connector Attachment for Small SCSI Disk Drives'
: SFF-8015 draft document Revision 3.3.

. The user must insure Electrical and Physical compatibility of the mating connector.

. Pin assignments for the 50 and 68 pin models are shown in Table 11.

Pin #	Voltage Level
1	+12V
2	Ground
3	Ground
4	+5V

Table 11. Power connector pin assignments

Refer to Figure 28 on page 64, for 50 pin power connector pin locations.

Refer to Figure 27 on page 63, for 68 pin power connector pin locations.

Refer to Figure 25 on page 61 and Table 15 on page 71 for power pin locations on the 80pin SCA connector.

Refer to the section entitled 4.2, “ Power Requirements by Model” on page 15 and 4.2.7, “Additional 5V Power Requirements for Differential” on page 35, for details on drive power requirements.

7.2 SCSI Bus Connector

DFHS has different model types that support 50 or 68 pin SCSI connectors in single-ended or 68 pin differential driver/receiver configurations. Also supported is the 80 pin SCA in a single-ended driver/receiver configuration.

This section describes the signal assignments of the ***DFHS*** SCSI connectors.

7.2.1 50 Pin Signal Connector

50 pin models use an AMP PN 84160-1 connector. The connector is compatible with the ANSI SCSI "A" connector specifications. It is limited to 8 bit data transfers only. Refer to Figure 28 on page 64 for a rear view of the 50 pin model connector.

The SCSI connector contact assignments for the 50 pin single-ended model is shown in Table 12 on page 68.

Signal Name	Connector Contact Number		Signal Name
GROUND	1	2	-DB(0)
GROUND	3	4	-DB(1)
GROUND	5	6	-DB(2)
GROUND	7	8	-DB(3)
GROUND	9	10	-DB(4)
GROUND	11	12	-DB(5)
GROUND	13	14	-DB(6)
GROUND	15	16	-DB(7)
GROUND	17	18	-DB(P0)
GROUND	19	20	GROUND
GROUND	21	22	GROUND
GROUND	23	24	GROUND
OPEN	25	26	TERMPWR
GROUND	27	28	GROUND
GROUND	29	30	GROUND
GROUND	31	32	-ATN
GROUND	33	34	GROUND
GROUND	35	36	-BSY
GROUND	37	38	-ACK
GROUND	39	40	-RST
GROUND	41	42	-MSG
GROUND	43	44	-SEL
GROUND	45	46	-C/D
GROUND	47	48	-REQ
GROUND	49	50	-I/O

Table 12. 50 Pin Single-Ended SCSI Connector Contact Assignments

7.2.2 68 Pin Signal Connector

68 pin models use an AMP connector (PN 786963-1) that is compatible with the ANSI SCSI "P" connector specifications. It can transfer data in both 8 bit (narrow) and 16 bit (wide) modes. Refer to Figure 27 on page 63 for a rear view of a 68 pin model.

The SCSI connector contact assignments for the 68 pin single-ended models are shown in Table 13 on page 69.

The SCSI connector contact assignments for the 68 pin differential models are shown in Table 14 on page 70.

Signal Name	Connector Contact Number		Signal Name
GROUND	1	35	-DB(12)
GROUND	2	36	-DB(13)
GROUND	3	37	-DB(14)
GROUND	4	38	-DB(15)
GROUND	5	39	-DB(P1)
GROUND	6	40	-DB(0)
GROUND	7	41	-DB(1)
GROUND	8	42	-DB(2)
GROUND	9	43	-DB(3)
GROUND	10	44	-DB(4)
GROUND	11	45	-DB(5)
GROUND	12	46	-DB(6)
GROUND	13	47	-DB(7)
GROUND	14	48	-DB(P0)
GROUND	15	49	GROUND
GROUND	16	50	GROUND
TERMPWR	17	51	TERMPWR
TERMPWR	18	52	TERMPWR
OPEN	19	53	OPEN
GROUND	20	54	GROUND
GROUND	21	55	-ATN
GROUND	22	56	GROUND
GROUND	23	57	-BSY
GROUND	24	58	-ACK
GROUND	25	59	-RST
GROUND	26	60	-MSG
GROUND	27	61	-SEL
GROUND	28	62	-C/D
GROUND	29	63	-REQ
GROUND	30	64	-I/O
GROUND	31	65	-DB(8)
GROUND	32	66	-DB(9)
GROUND	33	67	-DB(10)
GROUND	34	68	-DB(11)

Note: 8 bit devices which connect to the P-cable should tie the following signals inactive (high): -DB(8), -DB(9), -DB(10), -DB(11), -DB(12), -DB(13), -DB(14), -DB(15), -DB(P1). All other signals shall be connected as defined.

Table 13. 68 Pin Single-Ended SCSI Connector Contact Assignments

Signal Name	Connector Contact Number		Signal Name
+DB(12)	1	35	-DB(12)
+DB(13)	2	36	-DB(13)
+DB(14)	3	37	-DB(14)
+DB(15)	4	38	-DB(15)
+DB(P1)	5	39	-DB(P1)
GROUND	6	40	GROUND
+DB(0)	7	41	-DB(0)
+DB(1)	8	42	-DB(1)
+DB(2)	9	43	-DB(2)
+DB(3)	10	44	-DB(3)
+DB(4)	11	45	-DB(4)
+DB(5)	12	46	-DB(5)
+DB(6)	13	47	-DB(6)
+DB(7)	14	48	-DB(7)
+DB(P0)	15	49	-DB(P0)
DIFFSENS	16	50	GROUND
TERMPWR	17	51	TERMPWR
TERMPWR	18	52	TERMPWR
OPEN	19	53	OPEN
+ATN	20	54	-ATN
GROUND	21	55	GROUND
+BSY	22	56	-BSY
+ACK	23	57	-ACK
+RST	24	58	-RST
+MSG	25	59	-MSG
+SEL	26	60	-SEL
+C/D	27	61	-C/D
+REQ	28	62	-REQ
+I/O	29	63	-I/O
GROUND	30	64	GROUND
+DB(8)	31	65	-DB(8)
+DB(9)	32	66	-DB(9)
+DB(10)	33	67	-DB(10)
+DB(11)	34	68	-DB(11)

Note: 8 bit devices which connect to the P-cable should tie the following signals inactive: +/-DB(8), +/-DB(9), +/-DB(10), +/-DB(11), +/-DB(12), +/-DB(13), +/-DB(14), +/-DB(15), +/-DB(P1), or select "ENABLE NARROW MODE" on the Front Option Jumper Block and 'float' the same signals. All other signals shall be connected as defined.

Table 14. 68 Pin Differential SCSI Connector Contact Assignments

7.2.3 80 Pin (Single Connector Attachment) Connector

80 pin SCA models use a Molex connector (PN 87091-0001) that is compatible with the Specification of: 'Single Connector Attachment for Small SCSI Disk Drives' SFF-8015 draft document Revision 3.2 dated June 10,1994.

It can transfer data in both 8 bit (narrow) and 16 bit (wide) modes. Refer to Figure 25 on page 61 for a rear view of an 80 pin model.

Note that the "SCA" connector is not mechanically compatible with the 50-pin "A" connector or the 68-pin "P" connector as defined in the ANSI SCSI standard. This connector is intended for direct backplane attachment and is not intended to be cable attached to the bus.

The SCSI connector contact assignments for the 80 pin single-ended models are shown in Table 15.

Signal Name	**Connector Contact Number**		**Signal Name**
12 Volt	1	41	12V Ground
12 Volt	2	42	12V Ground
12 Volt	3	43	12V Ground
12 Volt	4	44	12V Ground
Reserved /NC	5	45	Reserved /NC
Reserved /NC	6	46	Ground
-DB(11)	7	47	Ground
-DB(10)	8	48	Ground
-DB(9)	9	49	Ground
-DB(8)	10	50	Ground
-I/O	11	51	Ground
-REQ	12	52	Ground
-C/D	13	53	Ground
-SEL	14	54	Ground
-MSG	15	55	Ground
-RST	16	56	Ground
-ACK	17	57	Ground
-BSY	18	58	Ground
-ATN	19	59	Ground
-DB(P0)	20	60	Ground
-DB(7)	21	61	Ground
-DB(6)	22	62	Ground
-DB(5)	23	63	Ground
-DB(4)	24	64	Ground
-DB(3)	25	65	Ground
-DB(2)	26	66	Ground
-DB(1)	27	67	Ground
-DB(0)	28	68	Ground
-DB(P1)	29	69	Ground
-DB(15)	30	70	Ground
-DB(14)	31	71	Ground
-DB(13)	32	72	Ground
-DB(12)	33	73	Ground
5 Volt	34	74	5 V Ground
5 Volt	35	75	5 V Ground
5 Volt	36	76	5 V Ground
SLAVE SYNC	37	77	Active LED Out
AUTO START	38	78	AUTO START DELAY
SCSI ID(0)	39	79	SCSI ID(1)
SCSI ID(2)	40	80	SCSI ID(3)

Note: 8bit devices which connect to the SCA connector should tie the following signals inactive high: DB(8), DB(9), DB(10), DB(11), DB(12), DB(13), DB(14), DB(15), DB(P1) or select "ENABLE NARROW MODE" on the Front Option Jumper Block and 'float' the same signals. All other signals shall be connected as defined.

Table 15. 80 Pin SCA Connector Contact Assignments

7.2.4 SCSI Bus Cable

Single-ended models permit cable lengths of up to 6 meters (19.68 feet). It should be noted however that users who plan to use "Fast" data transfers with single-ended models should follow all of the ANSI SCSI guidelines for single-ended "Fast" operations. This may include a cable length of less than 6 meters.

SCA connector models are not designed for direct cable attachment due to the combination of power and SCSI bus signals. "Fast" data transfers with SCA models should follow all of the ANSI SCSI guidelines for single-ended "Fast" operations.

Differential models permit cable lengths of up to 25 meters (82.02 feet). Cables must meet the requirements for differential cables as set forth in the ANSI SCSI standard under "Cable Requirements - Differential Cable".

The ANSI SCSI standard states that any stub from main cable must not exceed 0.1 meters for single-ended cables and 0.2 meters for differential cables. ***DFHS*** has a maximum internal stub length of 0.05 meters on all 'single ended' SCSI signals, and 0.1 meters on all 'differential' SCSI signals. To remain compliant with the standard the SCSI bus cable must not add more than 0.05 meters additional stub length to any of the single-ended SCSI signals or 0.1 meters to any differential SCSI signals.

7.2.5 SCSI Bus Terminators (Optional)

Upon request, Single Ended 50 and 68 pin models are available with on card SCSI bus Active terminators. (Please contact your IBM Customer Representative for the appropriate card PN).

For those cards having the Active Termination feature, this function can be enabled by installing a jumper between pins 13 and 14 of the Front Option Jumper Block or connecting pins 9 and 10 of the Auxiliary Connector on 68 SCSI pin models. (Refer to Figure 30 on page 75, Figure 31 on page 76, and Figure 34 on page 79.) The using system is responsible for making sure that all required signals are terminated at both ends of the cable. See 7.2.7, "SCSI Bus Electrical Characteristics" on page 73 for input capacitance values when terminators are disabled and when terminators are not populated on the card.

80 pin SCA models do not have internal SCSI bus terminators.

Some external terminator possibilities for single ended cabled systems are listed below: .

50 Pin Model Terminators	68 Pin Model Terminators
Methode Data Mate DM550-06-0	Methode Data Mate DM2050-02-68S

Table 16. Single Ended SCSI Terminators

Differential models do not have internal SCSI bus terminators. Some external terminator possibilities are listed below: .

50 Pin Model Terminators	68 Pin Model Terminators
Methode Data Mate DM550-05-0	Methode Data Mate DM2050-01-68D
Methode Data Mate DM1050-02-0	

Table 17. Differential SCSI Terminators

7.2.6 SCSI Bus Termination Power

Termination power is optionally provided for systems that desire to use it. In order to use the termination power, the user needs to install a jumper between pins 1 and 2 of the TermPower Block. (Refer to Figure 30 on page 75, Figure 31 on page 76.) The jumper should only be installed on one device, which should be the last device on the SCSI bus (i.e. the drive that is physically closest to a terminator). 68 pin models can source up to 2.0 Amps of current at 5.0 Volts (+- 5%) for termination power. 50 pin models can source up to 1.5 Amps of current at 5.0 Volts (+- 5%) for termination power.

7.2.6.1 SCSI Bus Termination Power Short Circuit Protection

The ANSI SCSI specification recommends for devices that optionally supply TERMPWR, to include current limiting protection for accidental short circuits. It also recommends that the maximum current available for TERMPWR should be 2 Amps. UL has a different requirement that they call the 8 Amp rule. This rule states that when a power source leaves an enclosure (like SCSI TERMPWR in the SCSI cable), it must trip 8 Amps of current within 1 minute.

The drive limits current to 5.0 Amps thru the use of a permanent fuse mounted on the electronics card.

Systems may also provide short circuit protection for drive supplied TERMPWR by limiting the current of the 5 Volt power it supplies to the drive.

7.2.7 SCSI Bus Electrical Characteristics

The following DC operating characteristics pertain to the single-ended SCSI bus transceivers. All of these parameters meet the ANSI SCSI-2 requirements.

- Ta = 0 to 70 deg. C

Symbol	SCSI I/O Parameters	min	max	Units	Notes
Vol	low level output voltage		0.4	V	Iout = 48 mA (REQ & ACK), Iout = 48 mA (all others)
Voh	high level output voltage	2.5		V	Iout = -400 uA
Vil	low level input voltage	-0.2	0.8	V	
Vih	high level input voltage	2.0	5.5	V	
Iil	low level input current		10	uA	
Iih	high level input current		50	uA	
Vihys	input hysteresis	0.3		V	
Ci	input capacitance		25	pF	w/terminators disabled
Ci	input capacitance		19	pF	w/o terminators

Table 18. Single-Ended Bus Electrical Characteristics

Differential models meet all electrical requirements as defined in the ANSI SCSI-2 standard under "Electrical Description - Differential Alternative".

7.2.8 Recommendations For SCSI Bus Noise Reduction

The SCSI committee has spent a large amount of resource looking into what needs to be done to assure SCSI devices will work as specified in the SCSI-2 standard. As a result of this, the committee is recommending the following approach:

- Use regulated 110 ohm terminator
- Use AWG 28 polyolefin shielded cables

- Make sure data and parity are on the outer ring of the cable and that REQ and ACK are in the core of the cable.

7.3 Option Block Connector (Jumper Blocks)

DFHS models contain a jumper block that can be used to enable certain features and select the SCSI ID of the drive. This jumper block is refered to as the 'Front' Option Jumper Block due to it's location on the drive (opposite the SCSI connector). This jumper block varies in pin definition based on interface type (50, 68, Differential, SCA).

The Option Block connector (2x16) used on 50 pin models is an AMP connector (PN 84156-4) having a pin spacing of 2mm. The pins on this connector have 0.030" Gold plating.

The Option Block connector (2x16) used on 68 and 80 pin models is an AMP connector (PN 84156-5) having a pin spacing of 2mm. The pins on this connector have 0.030" Gold plating.

The front block for these connector types are shown in Figure 30, Figure 31 on page 76, Figure 33 on page 78, and Figure 32 on page 77, note the differences in pins 1,2 and 14 based on interface type.

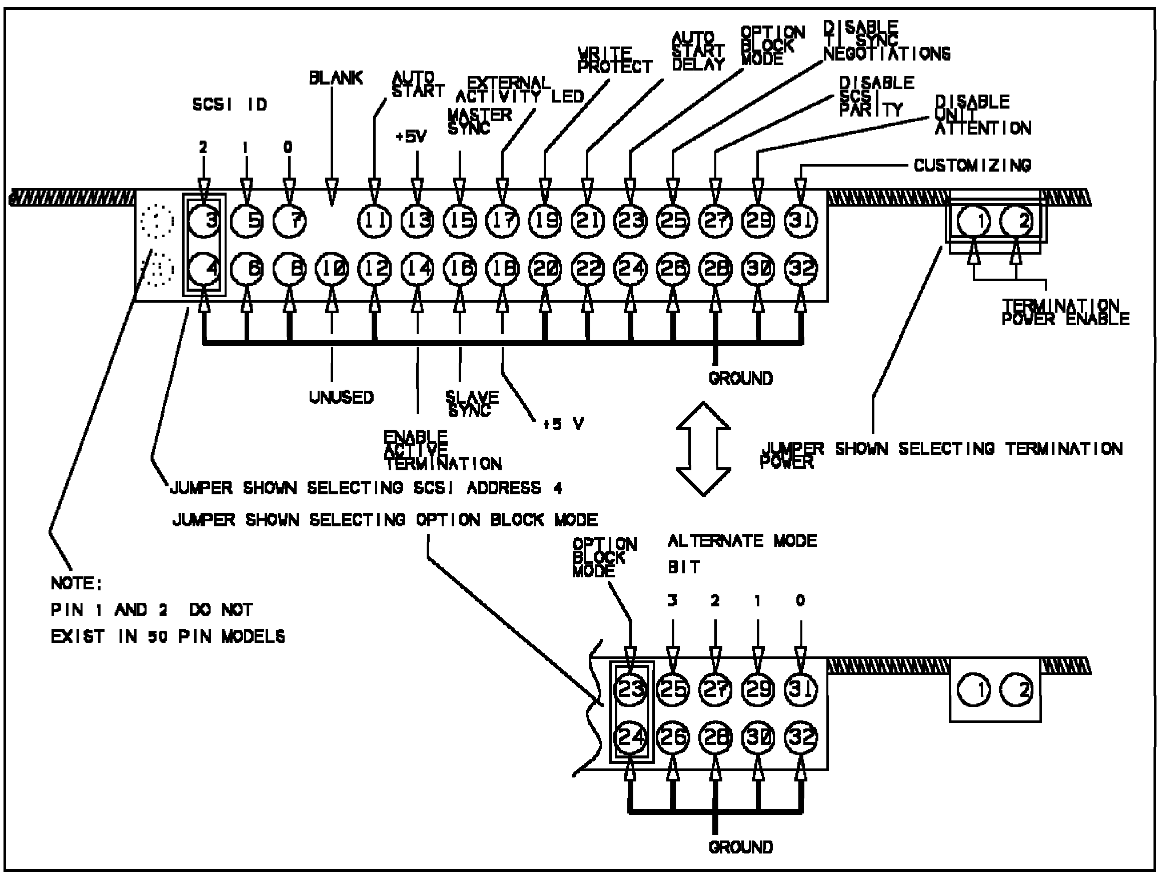

Figure 30. 50 pin Single Ended Front Option Jumper Block (& TermPower Block)

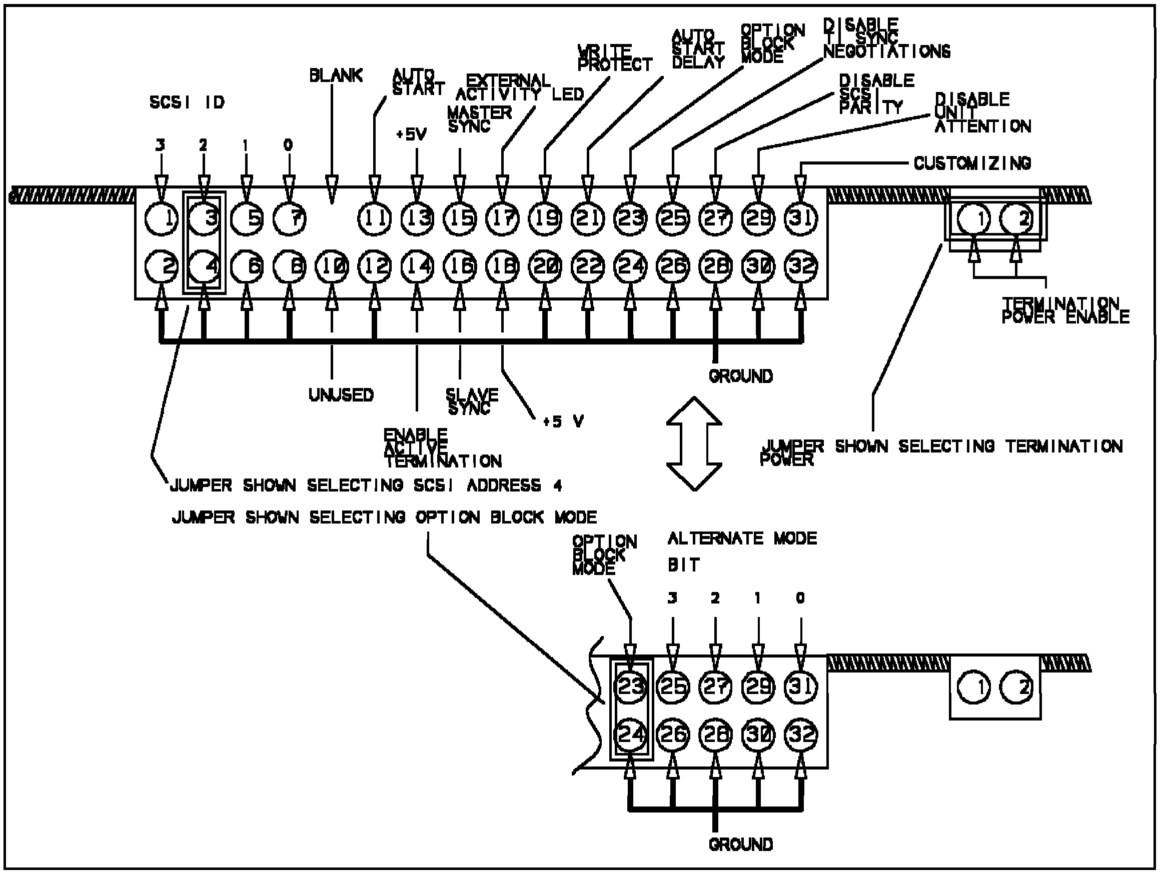

. Figure 31. 68 pin Single Ended Front Option Jumper Block (& TermPower Block)

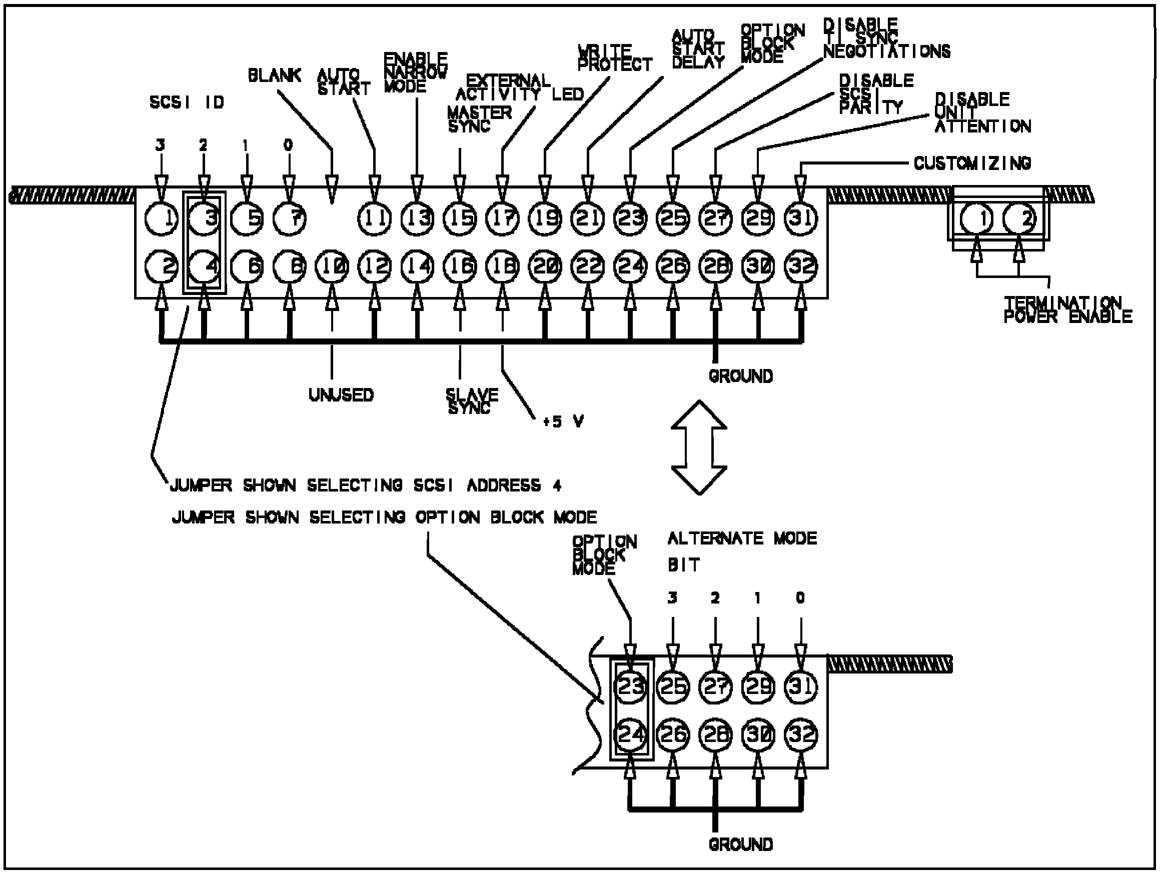

. Figure 32. 68 pin Differential Front Option Jumper Block

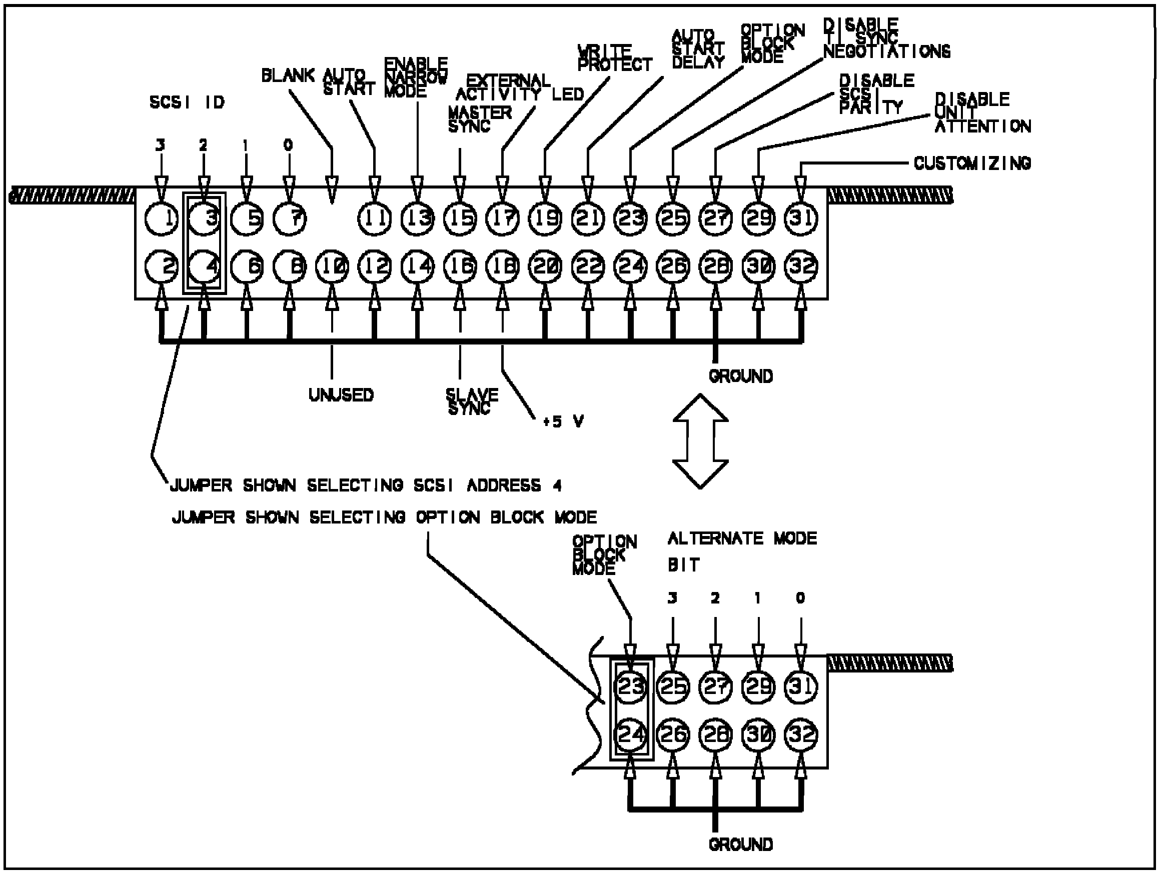

. Figure 33. 80 pin Single Ended Front Option Jumper Block

7.3.1 68 Pin Auxiliary Connector

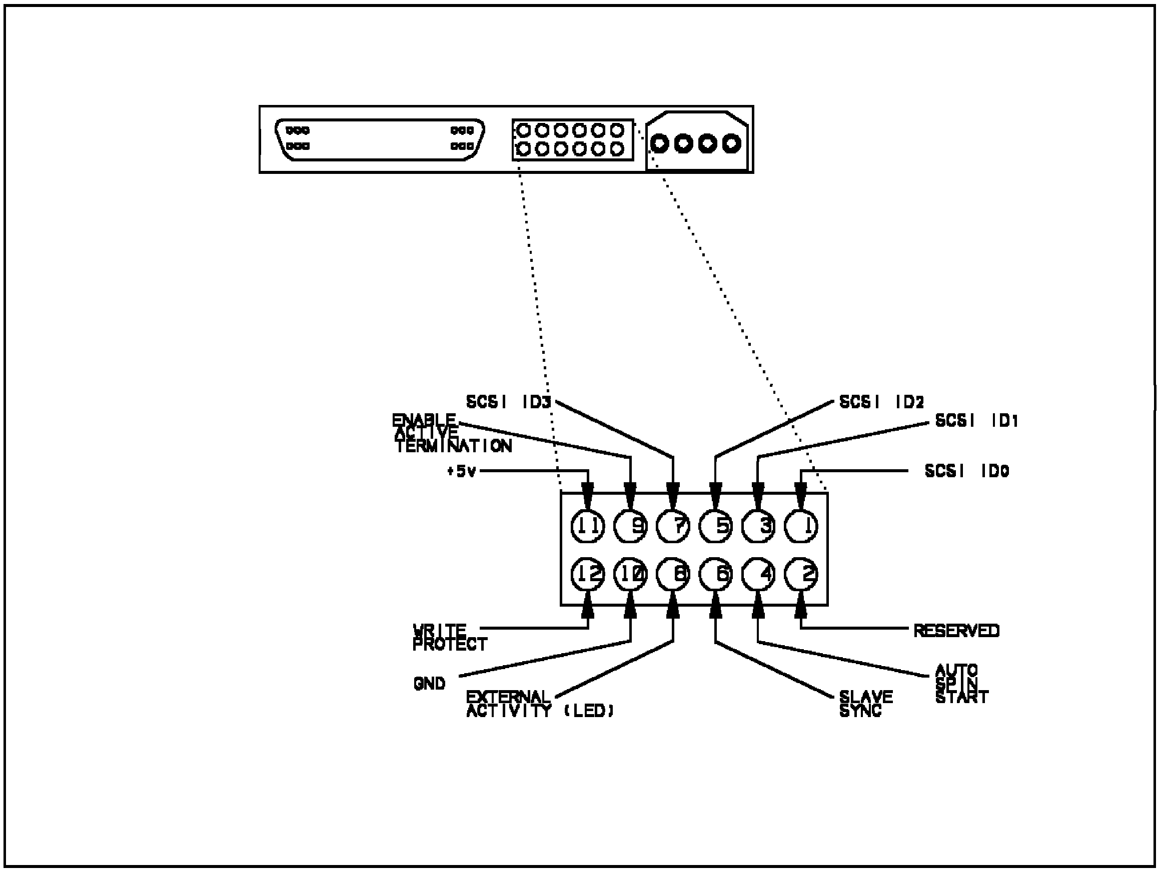

Figure 34. Auxiliary Connector on the 68 pin Connector

The 68 pin models contain an 'Auxiliary' connector that replicates some of the functions contained in the Front Option Jumper Block. This Connector is shown in figure Figure 34. The Auxiliary connector signal definition conforms to the SCSI document: SFF-8009 Rev 3.0 definition with the following exceptions:

1. EXTERNAL FAULT (XTFALT-) is not supported on pin 2
2. AUTO SPIN START was chosen as the 'vendor unique' signal assignment (on pin 4.) (This signal is an input to the drive. The SCSI spec (SCSI SFF-8009) specifies this pin as an output.) This signal should be useful for those applications that want to "auto-start" the drive based on location dependent SCSI ID.

   This pin should be handled in one of the following ways:

   a. tied to ground (auto spin start enabled)
   b. allowed to 'float' (no connection)
   c. driven with an open collector driver (>1mA sink capability)
3. This connector does not support the direct jumpering of SCSI ID at this connector due to Spindle Sync contention with this mode of operation. Please contact your IBM Customer Representative if your application requires the direct jumpering provision as defined in the SFF-8009 documentation.

7.3.2 SCSI ID (Address) Pins

Information on how to select a particular address for the SCSI device ID is given in Table 19 and Table 20.

Note: In the address determination tables, "off" means jumper is not in place and "on" means jumper is in place.

: **Note:** The jumpered state of these pins immediately after power up will determine the SCSI ID recognized
: by the drive.

BIT 3	BIT 2	BIT 1	BIT 0	ADDRESS
off	off	off	off	0
off	off	off	on	1
off	off	on	off	2
off	off	on	on	3
off	on	off	off	4
off	on	off	on	5
off	on	on	off	6
off	on	on	on	7
on	off	off	off	8
on	off	off	on	9
on	off	on	off	10
on	off	on	on	11
on	on	off	off	12
on	on	off	on	13
on	on	on	off	14
on	on	on	on	15

Table 19. Address Determination of 68 and 80 Pin Models

BIT 2	BIT 1	BIT 0	ADDRESS
off	off	off	0
off	off	on	1
off	on	off	2
off	on	on	3
on	off	off	4
on	off	on	5
on	on	off	6
on	on	on	7

Table 20. Address Determination of 50 Pin Models

7.3.3 Auto Start (& Delay) Pins

The Auto Start and Auto Start Delay pins control when and how the drive can spin up and come ready. When configured for Auto-Startup, the motor spins up after power is applied without the need of a SCSI Start Unit command. For no Auto-Startup, a SCSI Start Unit command is required to make the drive spin and be ready for media access operations. When in Auto-Startup mode, the drive will delay it's start time by a period of time multiplied by it's own SCSI address. Table 21 on page 81 and Table 22 on page 81 show whether or not Auto-Startup mode is active and the delay periods, where applicable, for all combinations of the pins.

Pins (50 & 68 interface pin models)		**Drive Behavior**	
AUTO START DELAY	**AUTO START**	**Auto-Startup Mode ?**	**Delay (sec) Multiplier**
off	off	NO	na
off	on	YES	0
on	off	YES	10
on	on	YES	4

. Table 21. Auto-Startup Modes selectable by Auto-Start/Delay Pin Combinations

Pins (80 interface pin models)		**Drive Behavior**	
AUTO START DELAY	**AUTO START**	**Auto-Startup Mode ?**	**Delay (sec) Multiplier**
off	off	YES	0
off	on	NO	na
on	off	YES	10
on	on	NO	na

. Table 22. Auto-Startup Modes selectable by Auto-Start/Delay Pin Combinations

7.3.4 External Activity (LED) Pins

The LED pins can be used to drive an external Light Emitting Diode. Please refer to the LED pin section of the **DFHS Interface Specification** for a detailed functional description of this pin.

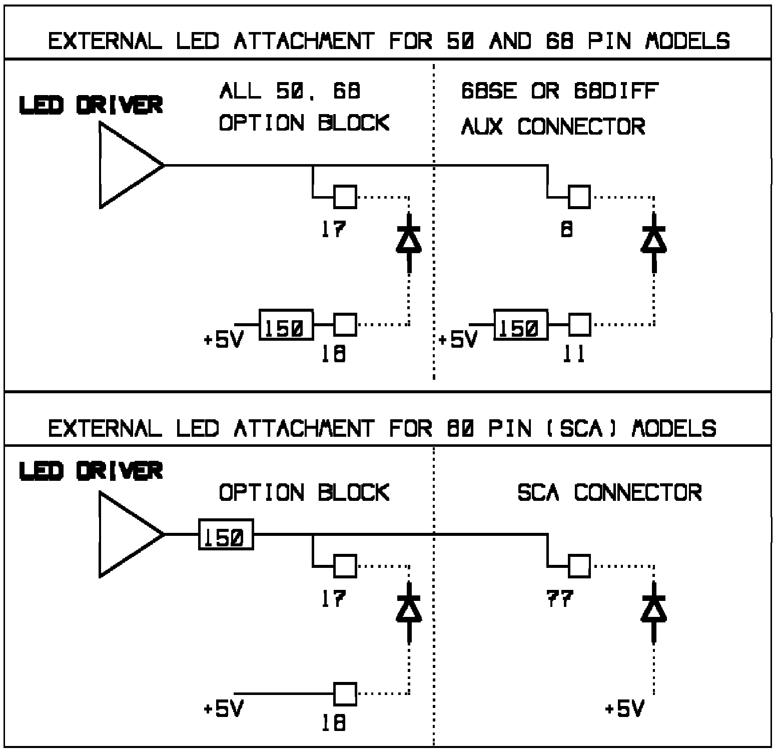

| Figure 35. External LED Connections

| Up to 33 mA of sink current capability is provided. Current limiting for the LED is provided on the
| electronics card. The LED Anode must be tied to the +5V source provided on pin 18 of the Front Option
| Jumper Block, pin 11 of the Auxiliary connector on the 68pin Unitized connector or the 5V power source
| on the 80 pin SCA model to obtain the specified current limit. The LED Cathode is then connected to the
| EXTERNAL ACTIVITY or ACTIVE LED OUT Pin to complete the circuit.

Note: This set of pins can be used to drive an LED located in a bezel connected to the front of the drive or to an external LED in systems where the front of the drive can not be easily seen.

Note: 68 pin and 80 pin SCA SCSI models have two sets of pins, a set on the front and a set on the back, that are connected to the same LED driver circuit. The combined drive capability is stated above.

7.3.5 Write Protect Pin

If the Write Protect pin is jumpered to ground the drive will prohibit SCSI commands that alter the customer data area portion of the media from being performed. The state of this pin is monitored on a per command basis.
See the **DFHS Interface Specification** for functional details.

7.3.6 Option Block Mode Pin

The Option Block Mode pin is used to modify the function of the Front Option Jumper Block. When the Option Block Mode Pin is not grounded (refer to Figure 30 on page 75, Figure 31 on page 76, Figure 33 on page 78, and Figure 32 on page 77)

the pin function of the Front Option Jumper Block will be as defined by the upper portion of the referenced figures. When the Option Block Mode Pin is grounded the pins (25,27,29,31) will be redefined to control what is referred to as 'ALTERNATE MODE'. ALTERNATE MODE allows tailoring of SCSI options such as 'boot up' characteristics and other operational features on a per customer basis.

Please contact your IBM Customer Representative for functional details on the customization of function allowed by this mode.

The jumpered state of these pins immediately after power is applied to the drive will determine their function.

7.3.7 Disable T.I.Sync. Negotiation Pin

If a Disable Target Initiated Synchronous Negotiation pin is grounded then an Initiator is required to start a negotiation handshake if Synchronous and/or 'Wide' (Double Byte) SCSI transfers are desired. Please refer to the **DFHS Interface Specification** for more details on this feature.

7.3.8 Disable SCSI Parity Pin

Grounding this pin will disable SCSI Parity checking.

7.3.9 Disable Unit Attention Pin

Grounding this pin will disable the drive from building Unit Attention Sense information for commands immediately following a Power On Reset (POR) or SCSI Bus Reset. Any pending Unit Attention conditions will also be cleared at POR or SCSI Reset times.

7.3.10 Customizing Pin

The customizing pin is currently reserved for future use.

7.3.11 Enable Narrow Mode

| Jumpering pin 14 to pin 13 on Figure 32 on page 77, or Figure 33 on page 78 will cause the 68 pin
| Differential or the 80 pin SCA to operate in a Single Byte mode. The drive will not negotiate for 'Wide'(Double Byte) operation. The drive will terminate the unused upper byte and upper byte parity on the SCSI bus.

7.3.12 Enable Active Termination

: Single Ended 50 and 68 pin models are available with or without on-card SCSI bus Active terminators.
: (Please contact your IBM Customer Representative for the appropriate card PN).

For those cards having the Active Termination feature, this function can be enabled by installing a jumper
: between pins 13 and 14 of the Front Option Jumper Block or connecting pins 9 and 10 of the Auxiliary
: Connector on 68 SCSI pin models. (Refer to Figure 30 on page 75, Figure 31 on page 76, and Figure 34 on page 79.)

7.4 Spindle Synchronization

7.4.1 Spindle Synchronization Overview

There are four modes of spindle synchronization. Reference Figure 36 for a list of how the -MASTER SYNC and -SLAVE SYNC pins on the Option Jumper Block are used for the different modes. The following paragraphs give a short description of each spindle synchronization mode:

- The Slave drive (Slave Sync mode) receives the index from the Master drive on the -SLAVE SYNC line and synchronizes its INDEX (Slave index) to it.
- Should the drive be the Master drive, (Master Sync mode), it outputs its INDEX on the -MASTER SYNC and the -SLAVE SYNC lines. The Master drive does not synchronize its index to any other device. It simply outputs its INDEX.
- In the Master Sync Control mode, a drive will synchronize its spindle to the signal it receives on the -SLAVE SYNC input. It outputs to -MASTER SYNC a pulse that has the same period as the drive INDEX, but is not synchronized to the drive INDEX generated from the disk.
- In the non-sync mode, the drive will receive the -SLAVE SYNC signal, but it is not used by the drive.

Reference the DFHS Interface Specification for further information on the different synchronization modes.

Spindle Synchronization Control Lines		Functional Mode
-MASTER SYNC	-SLAVE SYNC	
released	receive	Slave sync
drive	drive	Master Sync
drive	receive	Master Sync Control
released	receive	non sync

Figure 36. Spindle Synchronization Functional Modes

- Drive Synchronization with Offset.

  The drive electronics receives the Master Index and creates the delayed Slave Index from the drive INDEX. The delay is determined by using the SCSI Mode Select command, Rigid Disk Drive Geometry Parameters (Page 4). A rotational offset of 0/256 of a revolution up to 255/256 of a revolution may be selected in increments of 1/256 of a revolution. Reference the DFHS Interface Specification for further information on the rotational offset of synchronized spindles.

- Synchronization Time

  The SCSI Mode Select command is used to select the Spindle Sync mode. It will take 6 seconds to synchronize the Slave drive to the Master drive. While the Slave drive is synchronizing to the Master, it will not be able to read and write data. Once synchronized, the drive will maintain +/- 20 usec synchronization tolerance.

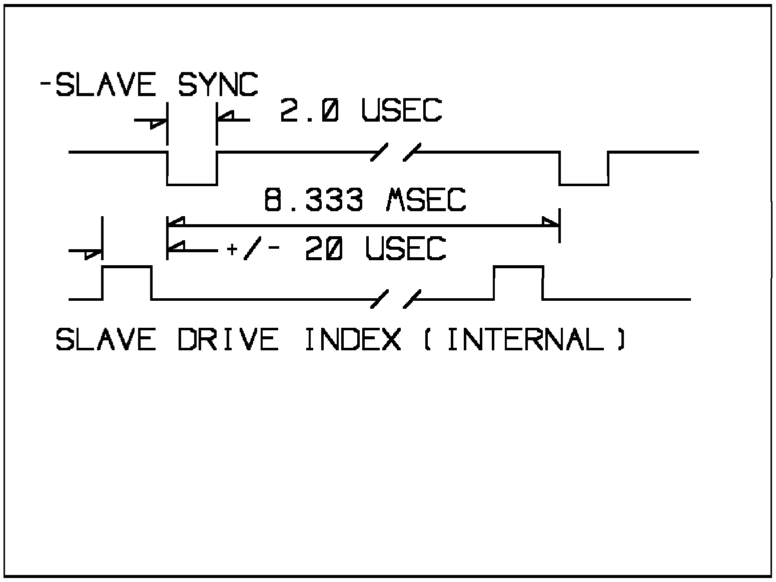

Figure 37. Slave-Sync to Slave-Index Timing

7.4.2 Spindle Synchronization Bus

The spindle synchronization Bus consists of the two signal lines, -MASTER SYNC and -SLAVE SYNC. Reference 7.3, "Option Block Connector (Jumper Blocks)" on page 75 which show the location of these signal lines on the Option Jumper Block connector. One potential configuration of this bus for drives that are to be used in a synchronized mode is shown in the following figure. This example requires the -SLAVE SYNC lines to be daisy chained together.

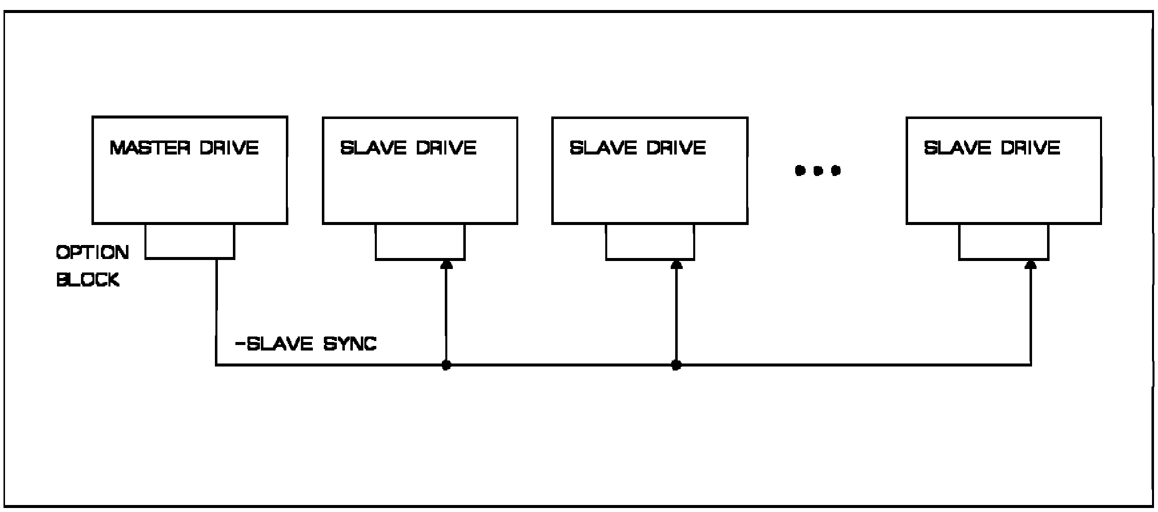

Figure 38. Daisy-Chain Connection of Synchronization Bus

- Termination

Bus termination of the -MASTER SYNC and -SLAVE SYNC signals is internal to the drive. These two signals each have a 5.1K ohm pullup to the +5 volt supply. A maximum of 30 drives can have their -MASTER SYNC or -SLAVE SYNC lines daisy chained together. Violating this could damage the Master drive line driver on the -MASTER SYNC and/or -SLAVE SYNC line.

It is the using system's responsibility to provide the cable to connect to the -SLAVE SYNC lines, and -MASTER SYNC lines where needed, of the synchronized drives.

- Bus Characteristics

  – maximum Bus length = 6 meters
  – 2 micro-second negative active pulse (when sourced by drive)
  – minimum of 1 micro-second negative active pulse when externally sourced
  – 0.8 volts = valid low input
  – 2.2 volts = valid high input
  – 0.4 volts = low output
  – Vcc volts = High output
  – 30 milli-amps = maximum output low level sink current

  The driver used for these two signal lines is a Open Drain buffer.

- -SLAVE SYNC Input Timing Requirements

If the input to -SLAVE SYNC is supplied by an external source other than a ***DFHS*** drive, the period of the input signal must be 8.333 msec with a tolerance of +/-.025%.

8. Reliability and Serviceability

Note: The reliability projections are based on the conditions stated below. All of the SCSI models will meet the projections as long as reliability operating conditions are not exceeded.

8.1 Error Detection

Error reporting ≥ 99%

All detected errors excluding interface and BATs #1 (Basic Assurance Test) errors

Error detection ≥ 99%

FRU isolation = 100%

To the device when the "Recommended Initiator Error Recovery Procedures" in the DFHS Interface Specification are followed.

No isolation to sub-assemblies within the device are specified.

8.2 Data Reliability

Probability of not recovering data	10 in 10^{15} bits read
Recoverable read errors	10 in 10^{13} bits read (Measured at nominal DC conditions and room environment with default error recovery - QPE* enabled.)[12]

Probability of miscorrecting unrecoverable data

Note: Eighteen bytes of ECC and two bytes of LRC are provided for each data block.

8.3 Seek Error Rate

Maximum allowed seek errors	10 in 10^{8} Seeks

8.4 Power On Hours Examples:

Maximum power on hours (with minimum power on/off cycles)

43,800 hours for life based on:

- 5 Power on/off cycles per month

- 730 power on hours per month

Nominal power on hours (with nominal power on/off cycles)

30,000 hours for life based on:

[12] * Please reference QPE (qualify post error) definition in the DFHS Interface Specification.

- 25 Power on/off cycles per month
- 500 power on hours per month

8.5 Power on/off cycles

| **Maximum on/off cycles** 1080/year (see 8.7, " *Mean Time Between Failure (*MTBF)" for *MTBF assumptions)

8.6 Useful Life

| **Product Life** 5 Years (see 8.7, " *Mean Time Between Failure (*MTBF)" for *MTBF assumptions)

8.7 *Mean Time Between Failure (*MTBF)

The mean time to failure target is 1,000,000 device hours per fail (3.0% CDF) based on the following assumptions:

- 6000 power on hours per year (500 power on hours per month times 12 months)
- 300 average on/off cycles per year (25 power cycles per month times 12 months)
- Seeking/Reading/Writing is assumed to be 20% of power on hours (Approximately 10 read/write operations per second)
- Operating at or below the Reliability temperature specifications (See Table 25 on page 92) and nominal voltages (See 4.2, " Power Requirements by Model" on page 15)
- 5 Year product life

Note: *MTBF - is a measure of the failure characteristics over **total product life**. *MTBF includes normal integration, installation, early life (latent), and intrinsic failures. *MTBF is predicated on supplier qualification, product design verification test, and field performance data.

8.8 Sample Failure Rate Projections

The following tables are for reference only. The tables contain failure rate projections for a given set of user conditions. Similar projections will be provided, upon request, for each user specific power on hour and power cycles per month condition. Contact your IBM customer representative for a customized projection.

Application	**Electronics only - (RA/MM)**				
	1st 30 days	**1st 60 days**	**1st 90 days**	**30 day average over life**	**CDF**
500POH/MM	0.000800	0.001362	0.001780	0.000253	1.5%
730POH/MM	0.001080	0.001764	0.002388	0.000427	2.6%

Table 23. Projected failure rates for the electronics only.

Application	Electronics and HDA - (RA/MM)				
	1st 30 days	1st 60 days	1st 90 days	30 day average over life	CDF
500POH/MM	0.001518	0.002510	0.003203	0.000500	3.0%
730POH/MM	0.001682	0.002724	0.003715	0.000730	4.38%

Table 24. Projected failure rates for the entire drive. (Electronics and HDA).

8.9 SPQL (Shipped product quality level)

	LA vintage	Ultimate (13th month)
Targets	.25%	.10%

8.10 Install Defect Free

Install Defect Free percentage 99.99 percent

8.11 Periodic Maintenance

None required

8.12 ESD Protection

The ***DFHS*** disk drives contain electrical components sensitive to damage due to electrostatic discharge (ESD). Proper ESD procedures must be followed during handling, installation, and removal. This includes the use of ESD wrist straps and ESD protective shipping containers.

8.13 Service

Replacement of the drive is the prefered service method. For those cases where customer data is critical, the system service strategy may be to replace the ***DFHS*** logic card. See 5.10, “Command Timeout Limits” on page 55, for system design requirements to support logic card replacement.

9. Operating Limits

The IBM Corporate specifications and bulletins, such as C-S 1-9700-000 in the contaminants section, that are referenced in this document are available for review (Please contact your IBM Customer Representative).

9.1 Environmental

Temperature	
Operating Ambient	41 to 131˚F (5 to 55˚C)
Operating Casting Temperature	41 to 158˚F (5 to 70˚C)
Storage	34 to 149˚F (1 to 65˚C) See Note
Shipping	-40 to 149˚F (-40 to 65˚C)
Temperature Gradient	
Operating	36˚F (20˚C) per hour
Shipping and storage	below condensation

Humidity		
Operating	5% to 90% noncondensing	
Storage	5% to 95% noncondensing	
Shipping	5% to 100%	(Applies at the packaged level)

Wet Bulb Temperature	
Operating	80˚F (26.7˚C) maximum
Shipping and Storage	85˚F (29.4˚C) maximum

Elevation	
Operating and Storage	-1000 to 10,000 feet (-304.8 to 3048 meters)
Shipping	-1000 to 40,000 feet (-304.8 to 12,192 meters)

Note: Guidelines for storage below 1˚C are given in IBM Technical Report TR 07.2112.

9.1.1 Temperature Measurement Points

The following is a list of measurement points and their temperatures (maximum and reliability). Maximum temperatures must not be exceeded at the worst case drive and system operating conditions with the drive randomly seeking, reading and writing. Reliability temperatures must not be exceeded at the nominal drive and system operating conditions with the drive randomly seeking, reading, and writing.

There must be sufficient air flow through the drive so that the casting and module temperature limits defined in Table 25 are not exceeded.

Table 25. Maximum and Reliability Operating Temperature Limits		
	Maximum	**Reliability**
Disk Enclosure Top	158°F (70°C)	131°F (55°C)
Disk Enclosure Bottom	158°F (70°C)	131°F (55°C)
PRDF Prime Module	203°F (95°C)	176°F (80°C)
WD Modules	185°F (85°C)	167°F (75°C)
Microprocessor Module	194°F (90°C)	167°F (75°C)
VCM FET	194°F (90°C)	167°F (75°C)
SMP FET	194°F (90°C)	167°F (75°C)

Figure 39 defines where measurements should be made to determine the top casting temperature during drive operation. Figure 40 defines the modules that are located on the bottom side of the card and the measurement location on the bottom of the casting.

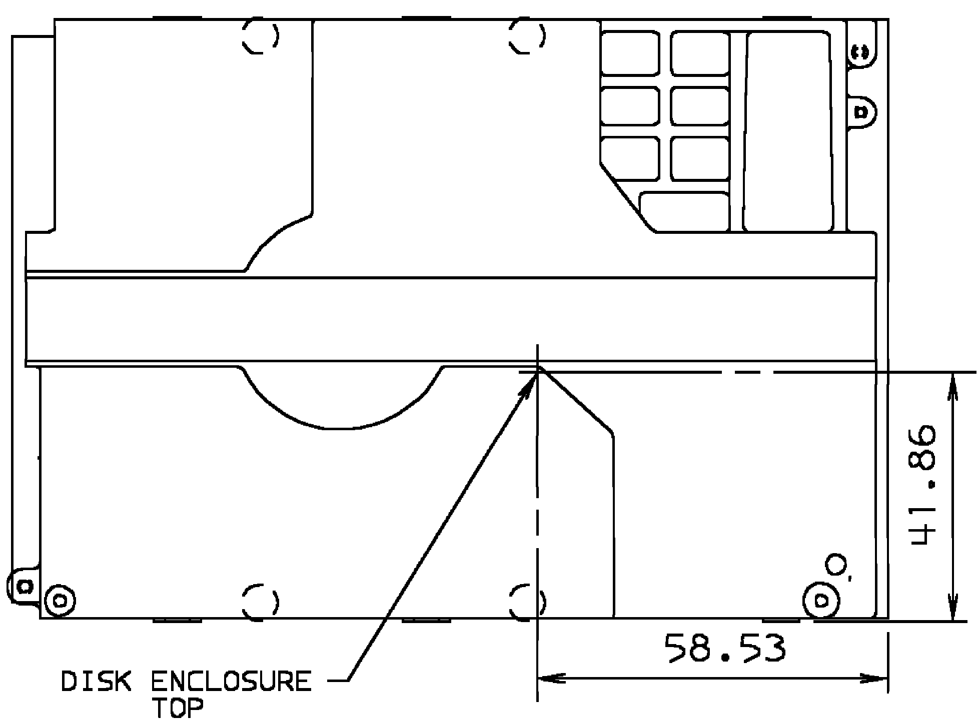

Notes: 1) Dimensions are in millimeters.

Figure 39. Temperature Measurement Points (top view)

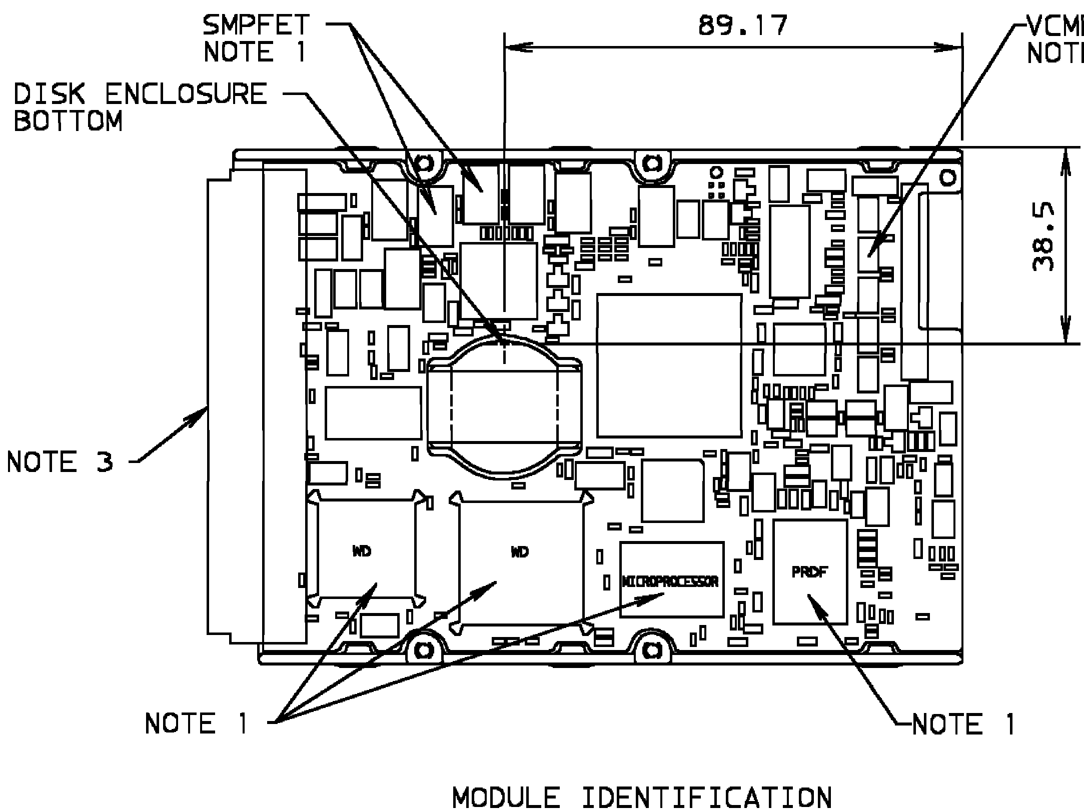

Notes:

1) Center thermocouple on the top surface of the module.

2) If copper tape is used to attach temperature sensors, it should be no larger than 6 mm square.

3) This drawing shows the drive with the 50 pin connector. Other connectors will look different as shown on Figure 25 to Figure 28.

4) Dimensions are in millimeters.

Figure 40. Temperature Measurement Points (bottom view)

9.2 Vibration and Shock

The operating vibration and shock limits in this specification are verified in two mount configurations:

1. By mounting with the 6-32 bottom holes on the drive with 2 mm clearance as required by 6.2, “Clearances” on page 57.
2. By mounting on any two opposing pairs of the 6-32 side mount holes.

Other mount configurations may result in different operating vibration and shock performance.

9.2.1 Drive Mounting Guidelines

The following guidelines may be helpful as drive mounting systems are being designed.

1. Mount the drive to its carrier/rack using the four extreme side holes to ensure that the drive's center of gravity is as close as possible to the center of stiffness of the mounting.
2. Do not permit any metal-to-metal impacts or chattering between the carrier/rack and the drive or between the carrier/rack and anything else. Metal-to-metal impacts create complex shock waveforms with short periods; such waveforms can excite high frequency modes of the components inside the drive.
3. The carrier/rack should not allow the drive to rotate in the plane of the disk and the carrier/rack itself should be mounted so that it does not rotate in the plane of the disk when the drive is running. Even though the drive uses a balanced rotatory actuator, its position can still be influenced by rotational acceleration.
4. Keep the rigid body resonances of the drive away from harmonics of the spindle speed. Consider not only the drive as mounted on its carrier but also when the drive is mounted to a carrier and then the carrier is mounted in a rack, the resonances of the drive in the entire system must be considered.

   7200 RPM Harmonics: 120 hz, 240 hz, 360 hz, 480 hz, ...
5. When the entire system/rack is vibration tested, the vibration amplitude of the drive as measured in all axis should decrease significantly for frequencies above 300 hz.
6. Consider the use of plastics or rubber in the rack/carrier design. Unlike metal, these materials can dampen vibration energy from other drives or fans located elsewhere in the rack.
7. Rather that creating a weak carrier/rack that flexes to fit the drive/carrier, hold the mounting gap to tighter tolerances. A flexible carrier/rack may contain resonances that cause operational vibration and/or shock problems.

9.2.2 Output Vibration Limits

Spindle Imbalance	1.0 gram-millimeters maximum for S1x, S2x Models
	1.5 gram-millimeters maximum for S4x Model

9.2.3 Operating Vibration

The vibration is applied in each of the three mutually perpendicular axis, one axis at a time. Referring to Figure 1 on page 7, the x-axis is defined as a line normal to the front/rear faces, the y-axis is defined as a line normal to the left side/right side faces, and the z-axis is normal to the x-y plane.

WARNING: The drives are sensitive to rotary vibration. Mounting within using systems should minimize the rotational input to the drive mounting points due to external vibration. IBM will provide technical support to assist users to overcome problems due to vibration.

Random Vibration

For excitation in the x-direction and the y-direction, the drive meets the required throughput specifications when subjected to vibration levels not exceeding the V4 vibration level defined below.

For excitation in the z-direction, the drive meets the required throughput specifications when subjected to vibration levels not exceeding the V4S vibration level defined below.

Note: The RMS value in the table below is obtained by taking the square root of the area defined by the g^2/hz spectrum from 5 to 500 hz.

Table 26. Random Vibration Levels										
Class	5 hz	17 hz	45 hz	48 hz	62 hz	65 hz	150 hz	200 hz	500 hz	RMS
V4	2.0E-5	1.1E-3	1.1E-3	8.0E-3	8.0E-3	1.0E-3	1.0E-3	8.0E-5	8.0E-5	0.56
V4S	2.0E-5	1.1E-3	1.1E-3	8.0E-3	8.0E-3	1.0E-3	1.0E-3	4.0E-5	4.0E-5	0.55
units	g^2/hz									g

<u>Swept Sine Vibration</u>

The drive will operate without hard errors when subjected to the swept sine vibration of 1.0 G peak from 5 to 300 hz in the x- and y direction. For input in the z-direction, an input of 1.0 G peak amplitude can be applied from 5 hz to 250 hz, the amplitude at 300 hz is 0.5 G peak. Linear interpolation is used to determine the acceleration levels between 250 hz and 300 hz.

The test will consist of a sweep from 5 to 300 hz and back to 5 hz. The sweep rate will be one hz per second.

Note: 1.0 G acceleration at 5 hz requires 0.78 inch double amplitude displacement.

9.2.3.1 Nonoperating Vibration

No damage will occur as long as vibration at the unpackaged drive in all three directions defined above does not exceed the levels defined in the table below. The test will consist of a sweep from 5 hz to 200 hz and back to 5 hz at a sweep rate of eight decades per hour.

Table 27. Non-operating Vibration Levels		
Frequency	5 hz to 7 hz	7 hz to 200 hz
Amplitude	0.8 inch DA	2.0 G peak

9.2.4 Operating Shock

No permanent damage will occur to the drive when subjected to a 10 G half sine wave shock pulse of 11 milliseconds duration.

: No permanent damage will occur to the drive when subjected to a 10 G half sine wave shock pulse of
: 2 millisecond duration.

The shock pulses are applied in either direction in each of three mutually perpendicular axis, one axis at a time.

9.2.5 Nonoperating Shock

<u>Translational Shock</u>

No damage will occur if the unpackaged drive is not subjected to a square wave shock greater than a "faired" value of 35 Gs applied to all three axis for a period of 20 milliseconds, one direction at a time.

No damage will occur if the unpackaged drive is not subjected to an 11 millisecond half sine wave shock greater than 70 Gs applied to all three axis, one direction at a time.

No damage will occur if the unpackaged drive is not subjected to a 2 millisecond half sine wave shock greater than 125 Gs applied to all three axis, one direction at a time.

<u>Rotational Shock</u>

No damage will occur if the unpackaged drive is not subjected to an 11 millisecond half sine wave shock greater than 7,000 radians per second squared applied to all three axis, one direction at a time.

No damage will occur if the unpackaged drive is not subjected to a 2 millisecond half sine wave shock greater than 15,000 radians per second squared applied to all three axis, one direction at a time.

9.3 Contaminants

The corrosive gas concentration expected to be typically encountered is Subclass G1; the particulate environment is expected to be P1 of C-S 1-9700-000 (1/89).

9.4 Acoustic Levels

Upper Limit Sound Power Requirements (Bels) for S1x & S2x Models									
	Octave Band Center Frequency (Hz)							A-weighted (see notes)	
	125	250	500	1K	2K	4K	8K	Maximum	Mean
Idle	4.5	3.5	3.3	3.5	4.5	4.5	4.5	5.00	4.7
Operating	4.5	4.0	3.6	4.1	4.8	4.8	4.5	5.25	5.0

Additionally, the population average of the sound pressure measured one meter above the center of the drive in idle mode will not exceed 36 dBA.

Upper Limit Sound Power Requirements (Bels) for S4x Models									
	Octave Band Center Frequency (Hz)							A-weighted (see notes)	
	125	250	500	1K	2K	4K	8K	Maximum	Mean
Idle	4.6	3.5	3.3	3.5	4.5	4.8	4.8	5.0	4.7
Operating	4.6	4.0	3.6	4.1	5.1	4.8	4.8	5.3	5.0

Additionally, the population average of the sound pressure measured one meter above the center of the drive in idle mode will not exceed 41 dBA.

Notes:

1. The above octave band and maximum sound power levels are statistical upper limits of the sound power levels. See C-B 1-1710-027 and C-S 1-1710-006 for further explanation.
2. The drive's are tested after a minimum of 20 minutes warm-up in idle mode.
3. The operating mode is simulated by seeking at a rate between 28 and 32 seeks per second.
4. The mean of a sample size of 10 or greater will be less than or equal to the stated mean with 95% confidence.

10. Standards

10.1 Safety

- UNDERWRITERS LABORATORY (UL) APPROVAL:

  The product is approved as a Recognized Component for use in Information Technology Equipment according to UL 1950 (without any D3 deviations). The UL Recognized Component marking is located on the product.

- CANADIAN STANDARDS ASSOCIATION (CSA) APPROVAL:

  The product is certified to CAN/CSA-C22.2 No. 950-M89 (without any Code 3 deviations). The CSA certification mark is located on the product.

- INTERNATIONAL ELECTROTECHNICAL COMMISSION (IEC) STANDARDS

  The product is certified to comply to EN60950 (IEC 950 with European additions) by TUV Rheinland. The TUV Rheinland Bauart mark is located on the product.

- SAFE HANDLING:

  The product is conditioned for safe handling in regards to sharp edges and corners.

- ENVIRONMENT:

  IBM will not knowingly or intentionally ship any units which during normal intended use or foreseeable misuse, would expose the user to toxic, carcinogenic, or otherwise hazardous substances at levels above the limitations identified in the current publications of the organizations listed below.

  International Agency for Research on Cancer (IARC)

  National Toxicology Program (NTP)

  Occupational Safety and Health Administration (OSHA)

  American Conference of Governmental Industrial Hygienists (ACGIH)

  California Governor's List of Chemical Restricted under California Safe Drinking Water and Toxic Enforcement Act 1986 (Also known as California Proposition 65)

- SECONDARY CIRCUIT PROTECTION REQUIRED IN USING SYSTEMS

  IBM has exercised care not to use any unprotected components or constructions that are particularly likely to cause fire. However, adequate secondary overcurrent protection is the responsibility of the user of the product. Additional protection against the possibility of sustained combustion due to circuit or component failure may need to be implemented by the user with circuitry external to the product. Overcurrent limits into the drive of 10 Amps or less should provide sufficient protection.

10.2 Electromagnetic Compatibility (EMC)

- FCC Requirements

  Pertaining to the STARFIRE disk drive, IBM will provide technical support to assist users in complying with the **United States Federal Communications Commission (FCC) Rules and Regulations, Part 15, Subpart J Computing Devices "Class A and B Limits"**. Tests for conformance to this requirement are performed with the disk drive mounted in the using system.

- VDE Requirements

Pertaining to the STARFIRE disk drive, IBM will provide technical support to assist users in complying with the requirements of the **German Vereingung Deutcher Elektriker (VDE) 0871/6.78 both the Individual Operation Permit(IOP) and the General Operation Permit (GOP) Limits**.

- CISPR 22 Requirements

  Pertaining to the STARFIRE disk drive, IBM will provide technical support to assist users in complying with the **Comite International Special des Perturbations Radio Electriques(International Special Committee on Radio Interference) CISPR 22 "Class A and B Limits"** .

- European Declaration of Conformity.

  Pertaining to the STARFIRE disk drive, IBM will provide technical support to assist users in complying with the **European Council Directive 89/336/EEC** so the final product can thereby bear the "CE" Mark of Conformity.