A decorative graphic on the left side of the page, featuring a dashed circle, a solid blue circle, and several overlapping curved lines in blue and grey.

Product Manual

Savvio® 10K.2 SAS

ST9146802SS
ST973402SS

100384760
Rev. E
April 2012

© 2012, Seagate Technology LLC All rights reserved.
Publication number: 100384760, Rev. E April 2012

Seagate, Seagate Technology and the Wave logo are registered trademarks of Seagate Technology LLC in the United States and/or other countries. Savvio and SeaTools are either trademarks or registered trademarks of Seagate Technology LLC or one of its affiliated companies in the United States and/or other countries. All other trademarks or registered trademarks are the property of their respective owners.

No part of this publication may be reproduced in any form without written permission of Seagate Technology LLC. Call 877-PUB-TEK1(877-782-8351) to request permission.

When referring to drive capacity, one gigabyte, or GB, equals one billion bytes and one terabyte, or TB, equals one trillion bytes. Your computer's operating system may use a different standard of measurement and report a lower capacity. In addition, some of the listed capacity is used for formatting and other functions, and thus will not be available for data storage. Actual quantities will vary based on various factors, including file size, file format, features and application software. Actual data rates may vary depending on operating environment and other factors. The export or re-export of hardware or software containing encryption may be regulated by the U.S. Department of Commerce, Bureau of Industry and Security (for more information, visit www.bis.doc.gov), and controlled for import and use outside of the U.S. Seagate reserves the right to change, without notice, product offerings or specifications.

Contents

1.0	Seagate Technology support services	1
2.0	Scope	2
3.0	Applicable standards and reference documentation	3
3.1	Standards	3
3.1.1	Electromagnetic compatibility	3
3.1.2	Electromagnetic compliance	4
3.1.3	European Union Restriction of Hazardous Substances (RoHS)	5
3.2	Reference documents	5
4.0	General description	6
4.1	Standard features	7
4.2	Media description	7
4.3	Performance	7
4.4	Reliability	8
4.5	Formatted capacities	8
4.6	Programmable drive capacity	8
4.7	Factory-installed accessories	8
4.8	Factory-installed options	9
5.0	Performance characteristics	10
5.1	Internal drive characteristics	10
5.2	Seek performance characteristics	10
5.2.1	Access time	10
5.2.2	Format command execution time for 512-byte sectors (minutes)	11
5.2.3	General performance characteristics	11
5.3	Start/stop time	11
5.4	Prefetch/multi-segmented cache control	12
5.5	Cache operation	12
5.5.1	Caching write data	12
5.5.2	Prefetch operation	13
6.0	Reliability specifications	14
6.1	Error rates	14
6.1.1	Recoverable Errors	14
6.1.2	Unrecoverable Errors	14
6.1.3	Seek errors	15
6.1.4	Interface errors	15
6.2	Reliability and service	15
6.2.1	Annualized Failure Rate (AFR) and Mean Time Between Failure (MTBF)	15
6.2.2	Preventive maintenance	15
6.2.3	Hot plugging the drive	15
6.2.4	S.M.A.R.T.	16
6.2.5	Thermal monitor	17
6.2.6	Drive Self Test (DST)	18
6.2.7	Product warranty	20
7.0	Physical/electrical specifications	21
7.1	AC power requirements	21
7.2	DC power requirements	21
7.2.1	Conducted noise immunity	23
7.2.2	Power sequencing	23
7.2.3	Current profiles	23
7.3	Power dissipation	25
7.4	Environmental limits	26

7.4.1	Temperature	26
7.4.2	Relative humidity	27
7.4.3	Effective altitude (sea level)	27
7.4.4	Shock and vibration	27
7.4.5	Air cleanliness	29
7.4.6	Corrosive environment	29
7.4.7	Acoustics	30
7.4.8	Electromagnetic susceptibility	30
7.5	Mechanical specifications	31
8.0	Defect and error management	32
8.1	Drive internal defects/errors	32
8.2	Drive error recovery procedures	32
8.3	SAS system errors	33
8.4	Background Media Scan	34
8.5	Media Pre-Scan	34
8.6	Deferred Auto-Reallocation	34
8.7	Idle Read After Write	35
9.0	Installation	36
9.1	Drive orientation	36
9.2	Cooling	37
9.3	Drive mounting	38
9.4	Grounding	38
10.0	Interface requirements	39
10.1	SAS features	39
10.1.1	task management functions	39
10.1.2	task management responses	39
10.2	Dual port support	40
10.3	SCSI commands supported	41
10.3.1	Inquiry data	45
10.3.2	Mode Sense data	45
10.4	Miscellaneous operating features and conditions	49
10.4.1	SAS physical interface	50
10.4.2	Physical characteristics	53
10.4.3	Connector requirements	53
10.4.4	Electrical description	53
10.4.5	Pin descriptions	53
10.4.6	SAS transmitters and receivers	54
10.4.7	Power	54
10.5	Signal characteristics	54
10.5.1	Ready LED Out	54
10.5.2	Differential signals	55

List of Figures

Figure 1.	Savvio 10K.2 SAS disk drive	2
Figure 2.	ST9146802SS current profile	23
Figure 3.	ST973402SS current profile	24
Figure 4.	ST9146802SS (3Gb) DC current and power vs. input/output operations per second	25
Figure 5.	ST973402SS (3Gb) DC current and power vs. input/output operations per second	25
Figure 6.	Location of the HDA temperature check point	26
Figure 7.	Recommended mounting	28
Figure 8.	Mounting configuration dimensions	31
Figure 9.	Physical interface	36
Figure 10.	Air flow	37
Figure 11.	Physical interface.	50
Figure 12.	SAS device plug dimensions	51
Figure 13.	SAS device plug dimensions (detail)	52
Figure 14.	SAS transmitters and receivers	54
Figure 15.	Receive eye mask	56
Figure 16.	Reveive tolerance eye mask.	56
Figure 17.	Sinusoidal jitter mask	57
Figure 18.	Compliance interconnect test load	63
Figure 19.	Zero-length test load.	63
Figure 20.	ISI loss example at 3.0 Gbps	63
Figure 21.	ISI loss example at 1.5 Gbps	64

1.0 Seagate Technology support services

SEAGATE ONLINE SUPPORT and SERVICES

For information regarding products and services, visit http://www.seagate.com/www/en-us/about/contact_us/

Available services include:

- Presales & Technical support
- Global Support Services telephone numbers & business hours
- Authorized Service Centers

For information regarding Warranty Support, visit

http://www.seagate.com/www/en-us/support/warranty_&_returns_assistance

For information regarding Data Recovery Services, visit <http://www.i365.com>

For Seagate OEM & Distribution partner portal, visit <https://direct.seagate.com/portal/system>

For Seagate reseller portal, visit <http://spp.seagate.com>

2.0 Scope

This manual describes Seagate Technology® LLC, Savvio® 10K.2 SAS (Serial Attached SCSI) disk drives.

Savvio drives support the SAS Protocol specifications to the extent described in this manual. The *SAS Interface Manual* (part number 100293071) describes the general SAS characteristics of this and other Seagate SAS drives.




Figure 1. Savvio 10K.2 SAS disk drive

3.0 Applicable standards and reference documentation

The drive has been developed as a system peripheral to the highest standards of design and construction. The drive depends on its host equipment to provide adequate power and environment for optimum performance and compliance with applicable industry and governmental regulations. Special attention must be given in the areas of safety, power distribution, shielding, audible noise control, and temperature regulation. In particular, the drive must be securely mounted to guarantee the specified performance characteristics. Mounting by bottom holes must meet the requirements of Section 9.3.

3.1 Standards

The Savvio family complies with Seagate standards as noted in the appropriate sections of this manual and the Seagate *SAS Interface Manual*, part number 100293071.

The Savvio disk drive is a UL recognized component per UL1950, CSA certified to CAN/CSA C22.2 No. 950-95, and VDE or TUV certified to EN60950.

3.1.1 Electromagnetic compatibility

The drive, as delivered, is designed for system integration and installation into a suitable enclosure prior to use. The drive is supplied as a subassembly and is not subject to Subpart B of Part 15 of the FCC Rules and Regulations nor the Radio Interference Regulations of the Canadian Department of Communications.

The design characteristics of the drive serve to minimize radiation when installed in an enclosure that provides reasonable shielding. The drive is capable of meeting the Class B limits of the FCC Rules and Regulations of the Canadian Department of Communications when properly packaged; however, it is the user's responsibility to assure that the drive meets the appropriate EMI requirements in their system. Shielded I/O cables may be required if the enclosure does not provide adequate shielding. If the I/O cables are external to the enclosure, shielded cables should be used, with the shields grounded to the enclosure and to the host controller.

3.1.1.1 Electromagnetic susceptibility

As a component assembly, the drive is not required to meet any susceptibility performance requirements. It is the responsibility of those integrating the drive within their systems to perform those tests required and design their system to ensure that equipment operating in the same system as the drive or external to the system does not adversely affect the performance of the drive. See Tables 12 and 13, DC power requirements.

3.1.2 Electromagnetic compliance

Seagate uses an independent laboratory to confirm compliance with the directives/standards for CE Marking and C-Tick Marking. The drive was tested in a representative system for typical applications. The selected system represents the most popular characteristics for test platforms. The system configurations include:

- Typical current use microprocessor
- 3.5-inch floppy disk drive
- Keyboard
- Monitor/display
- Printer
- External modem
- Mouse

Although the test system with this Seagate model complies with the directives/standards, we cannot guarantee that all systems will comply. The computer manufacturer or system integrator shall confirm EMC compliance and provide the appropriate marking for their product.

Electromagnetic compliance for the European Union

If this model has the CE Marking it complies with the European Union requirements of the Electromagnetic Compatibility Directive 89/336/EEC of 03 May 1989 as amended by Directive 92/31/EEC of 28 April 1992 and Directive 93/68/EEC of 22 July 1993.

Australian C-Tick

If this model has the C-Tick Marking it complies with the Australia/New Zealand Standard AS/NZS3548 1995 and meets the Electromagnetic Compatibility (EMC) Framework requirements of Australia's Spectrum Management Agency (SMA).

Korean MIC

If this model has the Korean Ministry of Information and Communication (MIC) logo, it complies with paragraph 1 of Article 11 of the Electromagnetic Compatibility (EMC) Control Regulation and meets the Electromagnetic Compatibility Framework requirements of the Radio Research Laboratory (RRL) Ministry of Information and Communication Republic of Korea.

Taiwanese BSMI

If this model has two Chinese words meaning "EMC certification" followed by an eight digit identification number, as a Marking, it complies with Chinese National Standard (CNS) 13438 and meets the Electromagnetic Compatibility (EMC) Framework requirements of the Taiwanese Bureau of Standards, Metrology, and Inspection (BSMI).

3.1.3 European Union Restriction of Hazardous Substances (RoHS)

The European Union Restriction of Hazardous Substances (RoHS) Directive restricts the presence of chemical substances, including Lead (Pb), in electronic products effective July 2006.

A number of parts and materials in Seagate products are procured from external suppliers. We rely on the representations of our suppliers regarding the presence of RoHS substances in these parts and materials. Our supplier contracts require compliance with our chemical substance restrictions, and our suppliers document their compliance with our requirements by providing material content declarations for all parts and materials for the disk drives documented in this publication. Current supplier declarations include disclosure of the inclusion of any RoHS-regulated substance in such parts or materials.

Seagate also has internal systems in place to ensure ongoing compliance with the RoHS Directive and all laws and regulations which restrict chemical content in electronic products. These systems include standard operating procedures that ensure that restricted substances are not utilized in our manufacturing operations, laboratory analytical validation testing, and an internal auditing process to ensure that all standard operating procedures are complied with.

3.2 Reference documents

Savvio 10K.2 SAS Installation Guide

Seagate part number: 100384761

SAS Interface Manual

Seagate part number: 100293071

SCSI Commands Reference Manual

Seagate part number: 100293068

ANSI SAS Documents

SFF-82232.5" Drive Form Factor with Serial Connector

SFF-8460HSS Backplane Design Guidelines

SFF-8470Multi Lane Copper Connector

SFF-8482SAS Plug Connector

ANSI INCITS.xxx Serial Attached SCSI (SAS) Standard (T10/1562-D)

ISO/IEC 14776-xxxSCSI Architecture Model-3 (SAM-3) Standard (T10/1561-D)

ISO/IEC 14776-xxxSCSI Primary Commands-3 (SPC-3) Standard (T10/1416-D)

ISO/IEC 14776-xxxSCSI Block Commands-2 (SBC-2) Standard (T10/1417-D)

ANSI Small Computer System Interface (SCSI) Documents

X3.270-1996(SCSI-3) Architecture Model

Specification for Acoustic Test Requirement and Procedures

Seagate part number: 30553-001

Package Test Specification Seagate P/N 30190-001 (under 100 lb.)

Package Test Specification Seagate P/N 30191-001 (over 100 lb.)

In case of conflict between this document and any referenced document, this document takes precedence.

4.0 General description

Savvio drives combine Tunneling Magnetoresistive (TMR) heads and a Serial Attached SCSI (SAS) interface to provide high performance, high capacity data storage for a variety of systems including engineering workstations, network servers, mainframes, and supercomputers. The Serial Attached SCSI interface is designed to meet next-generation computing demands for performance, scalability, flexibility and high-density storage requirements.

Savvio drives are random access storage devices designed to support the Serial Attached SCSI Protocol as described in the ANSI specifications, this document, and the *SAS Interface Manual* (part number 100293071) which describes the general interface characteristics of this drive. Savvio drives are classified as intelligent peripherals and provide level 2 conformance (highest level) with the ANSI SCSI-1 standard. The SAS connectors, cables and electrical interface are compatible with Serial ATA (SATA), giving future users the choice of populating their systems with either SAS or SATA hard disk drives. This allows you to continue to leverage your existing investment in SCSI while gaining a 3Gb/s serial data transfer rate.

The head and disk assembly (HDA) is sealed at the factory. Air recirculates within the HDA through a non-replaceable filter to maintain a contamination-free HDA environment.

Note. Never disassemble the HDA and do not attempt to service items in the sealed enclosure (heads, media, actuator, etc.) as this requires special facilities. The drive does not contain user-replaceable parts. Opening the HDA for any reason voids your warranty.

Savvio drives use a dedicated landing zone at the innermost radius of the media to eliminate the possibility of destroying or degrading data by landing in the data zone. The heads automatically go to the landing zone when power is removed from the drive.

An automatic shipping lock prevents potential damage to the heads and disks that results from movement during shipping and handling. The shipping lock disengages and the head load process begins when power is applied to the drive.

Savvio drives decode track 0 location data from the servo data embedded on each surface to eliminate mechanical transducer adjustments and related reliability concerns.

The drives also use a high-performance actuator assembly with a low-inertia, balanced, patented, straight arm design that provides excellent performance with minimal power dissipation.

4.1 Standard features

Savvio drives have the following standard features:

- Perpendicular recording technology
- 1.5 / 3Gb Serial Attached SCSI (SAS) interface
- Integrated dual port SAS controller supporting the SCSI protocol
- Support for SAS expanders and fanout adapters
- Firmware downloadable using the SAS interface
- 128 - deep task set (queue)
- Supports up to 32 initiators
- Jumperless configuration.
- User-selectable logical block size (512, 520, or 528 bytes per logical block).
- Industry standard SFF 2.5-inch dimensions
- Programmable logical block reallocation scheme
- Flawed logical block reallocation at format time
- Programmable auto write and read reallocation
- Reallocation of defects on command (Post Format)
- ECC maximum burst correction length of 320 bits
- No preventive maintenance or adjustments required
- Dedicated head landing zone
- Embedded servo design
- Automatic shipping lock
- Self diagnostics performed when power is applied to the drive
- Zone bit recording (ZBR)
- Vertical, horizontal, or top down mounting
- Dynamic spindle brake
- 16MB data buffer (see Section 5.5).
- Drive Self Test (DST)
- Background Media Scan (BGMS)
- Idle Read After Write (IRAW)
- Power Save

4.2 Media description

The media used on the drive has an aluminum substrate coated with a thin film magnetic material, overcoated with a proprietary protective layer for improved durability and environmental protection.

4.3 Performance

- Programmable multi-segmentable cache buffer
- 300MB/s maximum instantaneous data transfers.
- 10k RPM spindle. Average latency = 3.0ms
- Background processing of queue
- Supports start and stop commands (spindle stops spinning)
- Adaptive seek velocity; improved seek performance

4.4 Reliability

- Annualized Failure Rate (AFR) or 0.55%
- Mean time between failures (MTBF) of 1,600,000 hours
- Balanced low mass rotary voice coil actuator
- Incorporates industry-standard Self-Monitoring Analysis and Reporting Technology (S.M.A.R.T.)
- 5-year warranty

4.5 Formatted capacities

Standard OEM models are formatted to 512 bytes per block. The block size is selectable at format time and must be a multiple of 4 bytes. Users having the necessary equipment may modify the data block size before issuing a format command and obtain different formatted capacities than those listed.

To provide a stable target capacity environment and at the same time provide users with flexibility if they choose, Seagate recommends product planning in one of two modes:

1. Seagate designs specify capacity points at certain block sizes that Seagate guarantees current and future products will meet. We recommend customers use this capacity in their project planning, as it ensures a stable operating point with backward and forward compatibility from generation to generation. The current guaranteed operating points for this product are:

Capacity (Blocks)				
Sector Size	ST9146802SS		ST973402SS	
	Decimal	Hex	Decimal	Hex
512	286749488	11177330h	143374744	88BB998h
520	280790192	10BC84B0h	140395096	85E4258h
528	272662936	10408198h	136331468	82040CCh

2. Seagate drives also may be used at the maximum available capacity at a given block size, but the excess capacity above the guaranteed level will vary between other drive families and from generation to generation, depending on how each block size actually formats out for zone frequencies and splits over servo bursts. This added capacity potential may range from 0.1 to 1.3 percent above the guaranteed capacities listed above. Using the drives in this manner gives the absolute maximum capacity potential, but the user must determine if the extra capacity potential is useful, or whether their assurance of backward and forward compatibility takes precedence.

4.6 Programmable drive capacity

Using the Mode Select command, the drive can change its capacity to something less than maximum. See the Mode Select (6) parameter list table in the *SAS Interface Manual*, part number 100293071. A value of zero in the Number of Blocks field indicates that the drive will not change the capacity it is currently formatted to have. A number other than zero and less than the maximum number of LBAs in the Number of Blocks field changes the total drive capacity to the value in the Number of Blocks field. A value greater than the maximum number of LBAs is rounded down to the maximum capacity.

4.7 Factory-installed accessories

OEM standard drives are shipped with the *Savvio 10K.2 SAS Installation Guide* (part number 100384761).

4.8 Factory-installed options

You may order the following items which are incorporated at the manufacturing facility during production or packaged before shipping. Some of the options available are (not an exhaustive list of possible options):

- Other capacities can be ordered depending on sparing scheme and sector size requested.
- Single-unit shipping pack. The drive is normally shipped in bulk packaging to provide maximum protection against transit damage. Units shipped individually require additional protection as provided by the single unit shipping pack. Users planning single unit distribution should specify this option.
- The *Savvio 10K.2 SAS Installation Guide*, part number 100384761, is usually included with each standard OEM drive shipped, but extra copies may be ordered.
- The *Safety and Regulatory Agency Specifications*, part number 75789512, is usually included with each standard OEM drive shipped, but extra copies may be ordered.

5.0 Performance characteristics

This section provides detailed information concerning performance-related characteristics and features of Savvio drives.

5.1 Internal drive characteristics

	ST9146802SS	ST973402SS	
Drive capacity	146.8	73.4	Gbytes (formatted, rounded off value)
Read/write data heads	4	2	
Bytes per track	489,984	489,984	Bytes (average, rounded off values)
Bytes per surface	36,704	36,704	Mbytes (unformatted, rounded off value)
Tracks per surface (total)	75,373	75,373	Tracks (user accessible)
Tracks per inch	150,000	150,000	TPI (average)
Peak bits per inch	925,000	925,000	BPI
Areal density	135	135	Gbit/in ²
Internal data rate	660-1054	660-1054	Mbits/sec (variable with zone)
Disk rotation speed	10k	10k	rpm
Avg rotational latency	3.0	3.0	ms

5.2 Seek performance characteristics

See Section 10.4.1, "SAS physical interface" on page 50 and the *SAS Interface Manual* (part number 100293071) for additional timing details.

5.2.1 Access time

		Including controller overhead ^{1,2} (ms)		Not including controller overhead ^{1,2} (ms)	
		Read	Write	Read	Write
Average	Typical ^{3,4}	4.0	4.6	3.8	4.4
Single track	Typical ^{3,4}	0.4	0.6	0.2	0.4
Full stroke	Typical ^{3,4}	8.17	8.76	7.97	8.56

1. Execution time measured from receipt of the Command to the Response.
2. Assumes no errors and no sector has been relocated.
3. Typical access times are measured under nominal conditions of temperature, voltage, and horizontal orientation as measured on a representative sample of drives.
4. Access time = controller overhead + average seek time and applies to all data transfer commands.
Access to data = access time + latency time.

Data sheet data except full stroke which was Gen 2.0 data.

5.2.2 Format command execution time for 512-byte sectors (minutes)

	ST9146802SS	ST973402SS
Maximum (with verify)	104	52
Maximum (without verify)	70	35

5.2.3 General performance characteristics

Minimum sector interleave	1 to 1
Data buffer to/from disk media (one 512-byte logical block)*	82 to 132MB/s
Sustained transfer rate	55 to 89MB/s
SAS Interface maximum instantaneous transfer rate	300MB/s* per port (dual port = 600MB/s*)
Logical block sizes	
512 (default), 520 or 528.	
Read/write consecutive sectors on a track	Yes
Flaw reallocation performance impact (for flaws reallocated at format time using the spare sectors per sparing zone reallocation scheme.)	Negligible
Average rotational latency	3.0ms

*Assumes no errors and no relocated logical blocks. Rate measured from the start of the first logical block transfer to or from the host.

5.3 Start/stop time

The drive accepts the commands listed in the SAS Interface Manual less than 3 seconds after DC power has been applied.

If the drive receives a NOTIFY (ENABLE SPINUP) primitive through either port and has not received a START STOP UNIT command with the START bit equal to 0, the drive becomes ready for normal operations within 20 seconds (excluding the error recovery procedure).

If the drive receives a START STOP UNIT command with the START bit equal to 0 before receiving a NOTIFY (ENABLE SPINUP) primitive, the drive waits for a START STOP UNIT command with the START bit equal to 1. After receiving a START STOP UNIT command with the START bit equal to 1, the drive waits for a NOTIFY (ENABLE SPINUP) primitive. After receiving a NOTIFY (ENABLE SPINUP) primitive through either port, the drive becomes ready for normal operations within 20 seconds (excluding the error recovery procedure).

If the drive receives a START STOP UNIT command with the START bit and IMMED bit equal to 1 and does not receive a NOTIFY (ENABLE SPINUP) primitive within 5 seconds, the drive fails the START STOP UNIT command.

The START STOP UNIT command may be used to command the drive to stop the spindle. Stop time is 20 seconds (maximum) from removal of DC power. SCSI stop time is 20 seconds. There is no power control switch on the drive.

5.4 Prefetch/multi-segmented cache control

The drive provides a prefetch (read look-ahead) and multi-segmented cache control algorithms that in many cases can enhance system performance. Cache refers to the drive buffer storage space when it is used in cache operations. To select this feature, the host sends the Mode Select command with the proper values in the applicable bytes in page 08h. Prefetch and cache operations are independent features from the standpoint that each is enabled and disabled independently using the Mode Select command; however, in actual operation, the prefetch feature overlaps cache operation somewhat as described in sections 5.5.1 and 5.5.2.

All default cache and prefetch mode parameter values (Mode Page 08h) for standard OEM versions of this drive family are given in Table 19.

5.5 Cache operation

Note. Refer to the *SAS Interface Manual* for more detail concerning the cache bits.

Of the 16MB physical buffer space in the drive, approximately 13,000KB can be used as a cache. The buffer is divided into logical segments from which data is read and to which data is written.

The drive keeps track of the logical block addresses of the data stored in each segment of the buffer. If the cache is enabled (see RCD bit in the *SAS Interface Manual*), data requested by the host with a read command is retrieved from the buffer, if possible, before any disk access is initiated. If cache operation is not enabled, the buffer is still used, but only as circular buffer segments during disk medium read operations (disregarding Prefetch operation for the moment). That is, the drive does not check in the buffer segments for the requested read data, but goes directly to the medium to retrieve it. The retrieved data merely passes through some buffer segment on the way to the host. All data transfers to the host are in accordance with buffer-full ratio rules. See the explanation provided with the information about Mode Page 02h (disconnect/reconnect control) in the *SAS Interface Manual*.

The following is a simplified description of the prefetch/cache operation:

Case A—read command is received and all of the requested logical blocks are already in the cache:

1. Drive transfers the requested logical blocks to the initiator.

Case B—A Read command requests data, and at least one requested logical block is not in any segment of the cache:

1. The drive fetches the requested logical blocks from the disk and transfers them into a segment, and then from there to the host in accordance with the Mode Select Disconnect/Reconnect parameters, page 02h.
2. If the prefetch feature is enabled, refer to section 5.5.2 for operation from this point.

Each cache segment is actually a self-contained circular buffer whose length is an integer number of logical blocks. The drive dynamically creates and removes segments based on the workload. The wrap-around capability of the individual segments greatly enhances the cache's overall performance.

Note. The size of each segment is not reported by Mode Sense command page 08h, bytes 14 and 15. The value 0XFFFF is always reported regardless of the actual size of the segment. Sending a size specification using the Mode Select command (bytes 14 and 15) does not set up a new segment size. If the STRICT bit in Mode page 00h (byte 2, bit 1) is set to one, the drive responds as it does for any attempt to change an unchangeable parameter.

5.5.1 Caching write data

Write caching is a write operation by the drive that makes use of a drive buffer storage area where the data to be written to the medium is stored while the drive performs the Write command.

If read caching is enabled (RCD=0), then data written to the medium is retained in the cache to be made available for future read cache hits. The same buffer space and segmentation is used as set up for read functions. The buffer segmentation scheme is set up or changed independently, having nothing to do with the state of RCD. When a write command is issued, if RCD=0, the cache is first checked to see if any logical blocks that are to be written are already stored in the cache from a previous read or write command. If there are, the respective cache segments are cleared. The new data is cached for subsequent Read commands.

If the number of write data logical blocks exceed the size of the segment being written into, when the end of the segment is reached, the data is written into the beginning of the same cache segment, overwriting the data that was written there at the beginning of the operation; however, the drive does not overwrite data that has not yet been written to the medium.

If write caching is enabled (WCE=1), then the drive may return Good status on a write command after the data has been transferred into the cache, but before the data has been written to the medium. If an error occurs while writing the data to the medium, and Good status has already been returned, a deferred error will be generated.

The Synchronize Cache command may be used to force the drive to write all cached write data to the medium. Upon completion of a Synchronize Cache command, all data received from previous write commands will have been written to the medium.

Table 19 shows the mode default settings for the drive.

5.5.2 Prefetch operation

If the Prefetch feature is enabled, data in contiguous logical blocks on the disk immediately beyond that which was requested by a Read command are retrieved and stored in the buffer for immediate transfer from the buffer to the host on subsequent Read commands that request those logical blocks (this is true even if cache operation is disabled). Though the prefetch operation uses the buffer as a cache, finding the requested data in the buffer is a prefetch hit, not a cache operation hit.

To enable Prefetch, use Mode Select page 08h, byte 12, bit 5 (Disable Read Ahead - DRA bit). DRA bit = 0 enables prefetch.

The drive does not use the Max Prefetch field (bytes 8 and 9) or the Prefetch Ceiling field (bytes 10 and 11).

When prefetch (read look-ahead) is enabled (enabled by DRA = 0), the drive enables prefetch of contiguous blocks from the disk when it senses that a prefetch hit will likely occur. The drive disables prefetch when it decides that a prefetch hit is not likely to occur.

6.0 Reliability specifications

The following reliability specifications assume correct host and drive operational interface, including all interface timings, power supply voltages, environmental requirements and drive mounting constraints.

Seek error rate:	Less than 10 errors in 10^8 seeks
Read Error Rates ¹	
Recovered Data	Less than 10 errors in 10^{12} bits transferred (OEM default settings)
Unrecovered Data	Less than 1 sector in 10^{16} bits transferred
Miscorrected Data	Less than 1 sector in 10^{21} bits transferred
Interface error rate:	Less than 1 error in 10^{12} bits transferred
Mean Time Between Failure (MTBF):	1,600,000 hours
Annualized Failure Rate (AFR):	0.55%
Preventive maintenance:	None required

1. Error rate specified with automatic retries and data correction with ECC enabled and all flaws reallocated.

6.1 Error rates

The error rates stated in this manual assume the following:

- The drive is operated in accordance with this manual using DC power as defined in paragraph 7.2, "DC power requirements."
- Errors caused by host system failures are excluded from error rate computations.
- Assume random data.
- Default OEM error recovery settings are applied. This includes AWRE, ARRE, full read retries, full write retries and full retry time.

6.1.1 Recoverable Errors

Recoverable errors are those detected and corrected by the drive, and do not require user intervention.

Recoverable Data errors will use correction, although ECC on-the-fly is not considered for purposes of recovered error specifications.

Recovered Data error rate is determined using read bits transferred for recoverable errors occurring during a read, and using write bits transferred for recoverable errors occurring during a write.

6.1.2 Unrecoverable Errors

An unrecoverable data error is defined as a failure of the drive to recover data from the media. These errors occur due to head/media or write problems. Unrecoverable data errors are only detected during read operations, but not caused by the read. If an unrecoverable data error is detected, a MEDIUM ERROR (03h) in the Sense Key will be reported. Multiple unrecoverable data errors resulting from the same cause are treated as 1 error.

6.1.3 Seek errors

A seek error is defined as a failure of the drive to position the heads to the addressed track. After detecting an initial seek error, the drive automatically performs an error recovery process. If the error recovery process fails, a seek positioning error (Error code = 15h or 02h) will be reported with a Hardware error (04h) in the Sense Key. Recoverable seek errors are specified at Less than 10 errors in 10^8 seeks. Unrecoverable seek errors (Sense Key = 04h) are classified as drive failures.

6.1.4 Interface errors

An interface error is defined as a failure of the receiver on a port to recover the data as transmitted by the device port connected to the receiver. The error may be detected as a running disparity error, illegal code, loss of word sync, or CRC error.

6.2 Reliability and service

You can enhance the reliability of Savvio disk drives by ensuring that the drive receives adequate cooling. Section 7.0 provides temperature measurements and other information that may be used to enhance the service life of the drive. Section 9.2 provides recommended air-flow information.

6.2.1 Annualized Failure Rate (AFR) and Mean Time Between Failure (MTBF)

The production disk drive shall achieve an AFR of 0.55% (MTBF of 1,600,000 hours) when operated in an environment that ensures the HDA case temperatures do not exceed the values specified in Section 7.4. Operation at case temperatures outside the specifications in Section 7.4 may increase the product AFR (decrease the MTBF). The AFR (MTBF) is a population statistic not relevant to individual units.

The AFR (MTBF) specification is based on the following assumptions for Enterprise Storage System environments:

- 8760 power-on hours per year.
- 250 average on/off cycles per year.
- Operations at nominal voltages.
- Systems will provide adequate cooling to ensure the case temperatures specified in Section 7.4 are not exceeded. Temperatures outside the specifications in Section 7.4 will increase the product AFR and decrease the MTBF.

6.2.2 Preventive maintenance

No routine scheduled preventive maintenance is required.

6.2.3 Hot plugging the drive

When a disk is powered on by switching the power or hot plugged, the drive runs a self test before attempting to communicate on its' interfaces. When the self test completes successfully, the drive initiates a Link Reset starting with OOB. An attached device should respond to the link reset. If the link reset attempt fails, or any time the drive loses sync, the drive initiated link reset. The drive will initiate link reset once per second but alternates between port A and B. Therefore each port will attempt a link reset once per 2 seconds assuming both ports are out of sync..

If the self-test fails, the drive does not respond to link reset on the failing port.

Note. It is the responsibility of the systems integrator to assure that no temperature, energy, voltage hazard, or ESD potential hazard is presented during the hot connect/disconnect operation. Discharge the static electricity from the drive carrier prior to inserting it into the system.

Caution. The drive motor must come to a complete stop prior to changing the plane of operation. This time is required to insure data integrity.

6.2.4 S.M.A.R.T.

S.M.A.R.T. is an acronym for Self-Monitoring Analysis and Reporting Technology. This technology is intended to recognize conditions that indicate imminent drive failure and is designed to provide sufficient warning of a failure to allow you to back up the data before an actual failure occurs.

Note. The drive's firmware monitors specific attributes for degradation over time but can't predict instantaneous drive failures.

Each monitored attribute has been selected to monitor a specific set of failure conditions in the operating performance of the drive and the thresholds are optimized to minimize "false" and "failed" predictions.

Controlling S.M.A.R.T.

The operating mode of S.M.A.R.T. is controlled by the DEXCPT and PERF bits on the Informational Exceptions Control mode page (1Ch). Use the DEXCPT bit to enable or disable the S.M.A.R.T. feature. Setting the DEXCPT bit disables all S.M.A.R.T. functions. When enabled, S.M.A.R.T. collects on-line data as the drive performs normal read and write operations. When the PERF bit is set, the drive is considered to be in "On-line Mode Only" and will not perform off-line functions.

You can measure off-line attributes and force the drive to save the data by using the Rezero Unit command. Forcing S.M.A.R.T. resets the timer so that the next scheduled interrupt is in two hours.

You can interrogate the drive through the host to determine the time remaining before the next scheduled measurement and data logging process occurs. To accomplish this, issue a Log Sense command to log page 0x3E. This allows you to control when S.M.A.R.T. interruptions occur. Forcing S.M.A.R.T. with the RTZ command resets the timer.

Performance impact

S.M.A.R.T. attribute data is saved to the disk so that the events that caused a predictive failure can be recreated. The drive measures and saves parameters once every two hours subject to an idle period on the drive interfaces. The process of measuring off-line attribute data and saving data to the disk is uninterruptable. The maximum on-line only processing delay is summarized below:

Maximum processing delay

	On-line only delay DEXCPT = 0, PERF = 1	Fully-enabled delay DEXCPT = 0, PERF = 0
S.M.A.R.T. delay times	220ms	455ms

Reporting control

Reporting is controlled by the MRIE bits in the Informational Exceptions Control mode page (1Ch). Subject to the reporting method, the firmware will issue to the host an 01-5Dxx sense code. The error code is preserved through bus resets and power cycles.

Determining rate

S.M.A.R.T. monitors the rate at which errors occur and signals a predictive failure if the rate of degraded errors increases to an unacceptable level. To determine rate, error events are logged and compared to the number of total operations for a given attribute. The interval defines the number of operations over which to measure the rate. The counter that keeps track of the current number of operations is referred to as the Interval Counter.

S.M.A.R.T. measures error rates. All errors for each monitored attribute are recorded. A counter keeps track of the number of errors for the current interval. This counter is referred to as the Failure Counter.

Error rate is the number of errors per operation. The algorithm that S.M.A.R.T. uses to record rates of error is to set thresholds for the number of errors and their interval. If the number of errors exceeds the threshold before the interval expires, the error rate is considered to be unacceptable. If the number of errors does not exceed the threshold before the interval expires, the error rate is considered to be acceptable. In either case, the interval and failure counters are reset and the process starts over.

Predictive failures

S.M.A.R.T. signals predictive failures when the drive is performing unacceptably for a period of time. The firmware keeps a running count of the number of times the error rate for each attribute is unacceptable. To accomplish this, a counter is incremented each time the error rate is unacceptable and decremented (not to exceed zero) whenever the error rate is acceptable. If the counter continually increments such that it reaches the predictive threshold, a predictive failure is signaled. This counter is referred to as the Failure History Counter. There is a separate Failure History Counter for each attribute.

6.2.5 Thermal monitor

Savvio drives implement a temperature warning system which:

1. Signals the host if the temperature exceeds a value which would threaten the drive.
2. Signals the host if the temperature exceeds a user-specified value.
3. Saves a S.M.A.R.T. data frame on the drive which exceeds the threatening temperature value.

A temperature sensor monitors the drive temperature and issues a warning over the interface when the temperature exceeds a set threshold. The temperature is measured at power-up and then at ten-minute intervals after power-up.

The thermal monitor system generates a warning code of 01-0B01 when the temperature exceeds the specified limit in compliance with the SCSI standard. The drive temperature is reported in the FRU code field of mode sense data. You can use this information to determine if the warning is due to the temperature exceeding the drive threatening temperature or the user-specified temperature.

This feature is controlled by the Enable Warning (EWasc) bit, and the reporting mechanism is controlled by the Method of Reporting Informational Exceptions field (MRIE) on the Informational Exceptions Control (IEC) mode page (1Ch).

The current algorithm implements two temperature trip points. The first trip point is set at 65°C which is the maximum temperature limit according to the drive specification. The second trip point is user-selectable using the Log Select command. The reference temperature parameter in the temperature log page (see Table 11)

can be used to set this trip point. The default value for this drive is 65°C, however, you can set it to any value in the range of 0 to 65°C. If you specify a temperature greater than 65°C in this field, the temperature is rounded down to 65°C. A sense code is sent to the host to indicate the rounding of the parameter field.

Table 11: Temperature Log Page (0Dh)

Parameter Code	Description
0000h	Primary Temperature
0001h	Reference Temperature

6.2.6 Drive Self Test (DST)

Drive Self Test (DST) is a technology designed to recognize drive fault conditions that qualify the drive as a failed unit. DST validates the functionality of the drive at a system level.

There are two test coverage options implemented in DST:

1. Extended test
2. Short test

The most thorough option is the extended test that performs various tests on the drive and scans every logical block address (LBA) of the drive. The short test is time-restricted and limited in length—it does not scan the entire media surface, but does some fundamental tests and scans portions of the media.

If DST encounters an error during either of these tests, it reports a fault condition. If the drive fails the test, remove it from service and return it to Seagate for service.

6.2.6.1 DST failure definition

The drive will present a “diagnostic failed” condition through the self-tests results value of the diagnostic log page if a functional failure is encountered during DST. The channel and servo parameters are not modified to test the drive more stringently, and the number of retries are not reduced. All retries and recovery processes are enabled during the test. If data is recoverable, no failure condition will be reported regardless of the number of retries required to recover the data.

The following conditions are considered DST failure conditions:

- Seek error after retries are exhausted
- Track-follow error after retries are exhausted
- Read error after retries are exhausted
- Write error after retries are exhausted

Recovered errors will not be reported as diagnostic failures.

6.2.6.2 Implementation

This section provides all of the information necessary to implement the DST function on this drive.

6.2.6.2.1 State of the drive prior to testing

The drive must be in a ready state before issuing the Send Diagnostic command. There are multiple reasons why a drive may not be ready, some of which are valid conditions, and not errors. For example, a drive may be in process of doing a format, or another DST. It is the responsibility of the host application to determine the “not ready” cause.

While not technically part of DST, a Not Ready condition also qualifies the drive to be returned to Seagate as a failed drive.

A Drive Not Ready condition is reported by the drive under the following conditions:

- Motor will not spin
- Motor will not lock to speed
- Servo will not lock on track
- Drive cannot read configuration tables from the disk

In these conditions, the drive responds to a Test Unit Ready command with an 02/04/00 or 02/04/03 code.

6.2.6.2.2 Invoking DST

To invoke DST, submit the Send Diagnostic command with the appropriate Function Code (001b for the short test or 010b for the extended test) in bytes 1, bits 5, 6, and 7.

6.2.6.2.3 Short and extended tests

DST has two testing options:

1. short
2. extended

These testing options are described in the following two subsections.

Each test consists of three segments: an electrical test segment, a servo test segment, and a read/verify scan segment.

Short test (Function Code: 001b)

The purpose of the short test is to provide a time-limited test that tests as much of the drive as possible within 120 seconds. The short test does not scan the entire media surface, but does some fundamental tests and scans portions of the media. A complete read/verify scan is not performed and only factual failures will report a fault condition. This option provides a quick confidence test of the drive.

Extended test (Function Code: 010b)

The objective of the extended test option is to empirically test critical drive components. For example, the seek tests and on-track operations test the positioning mechanism. The read operation tests the read head element and the media surface. The write element is tested through read/write/read operations. The integrity of the media is checked through a read/verify scan of the media. Motor functionality is tested by default as a part of these tests.

The anticipated length of the Extended test is reported through the Control Mode page.

6.2.6.2.4 Log page entries

When the drive begins DST, it creates a new entry in the Self-test Results Log page. The new entry is created by inserting a new self-test parameter block at the beginning of the self-test results log parameter section of the log page. Existing data will be moved to make room for the new parameter block. The drive reports 20 parameter blocks in the log page. If there are more than 20 parameter blocks, the least recent parameter block will be deleted. The new parameter block will be initialized as follows:

1. The Function Code field is set to the same value as sent in the DST command
2. The Self-Test Results Value field is set to Fh
3. The drive will store the log page to non-volatile memory

After a self-test is complete or has been aborted, the drive updates the Self-Test Results Value field in its Self-Test Results Log page in non-volatile memory. The host may use Log Sense to read the results from up to the last 20 self-tests performed by the drive. The self-test results value is a 4-bit field that reports the results of the test. If the field is set to zero, the drive passed with no errors detected by the DST. If the field is not set to zero, the test failed for the reason reported in the field.

The drive will report the failure condition and LBA (if applicable) in the Self-test Results Log parameter. The Sense key, ASC, ASCQ, and FRU are used to report the failure condition.

6.2.6.2.5 Abort

There are several ways to abort a diagnostic. You can use a SCSI Bus Reset or a Bus Device Reset message to abort the diagnostic.

You can abort a DST executing in background mode by using the abort code in the DST Function Code field. This will cause a 01 (self-test aborted by the application client) code to appear in the self-test results values log. All other abort mechanisms will be reported as a 02 (self-test routine was interrupted by a reset condition).

6.2.7 Product warranty

See Section 1.0 for warranty contact information.

Shipping

When transporting or shipping a drive, use only a Seagate-approved container. Keep your original box. Seagate approved containers are easily identified by the Seagate Approved Package label. Shipping a drive in a non-approved container voids the drive warranty.

Seagate repair centers may refuse receipt of components improperly packaged or obviously damaged in transit. Contact your authorized Seagate distributor to purchase additional boxes. Seagate recommends shipping by an air-ride carrier experienced in handling computer equipment.

Storage

The maximum recommended storage period for the drive in a non-operational environment is 90 days. Drives should be stored in the original unopened Seagate shipping packaging when ever possible. Once the drive is removed from the Seagate original packaging the recommended maximum period between drive operation cycles is 30 days. During any storage period the drive non-operational temperature, humidity, wet bulb, atmospheric conditions, shock, vibration, magnetic and electrical field specifications should be followed.

Product repair and return information

Seagate customer service centers are the only facilities authorized to service Seagate drives. Seagate does not sanction any third-party repair facilities. Any unauthorized repair or tampering with the factory seal voids the warranty.

7.0 Physical/electrical specifications

This section provides information relating to the physical and electrical characteristics of the drive.

7.1 AC power requirements

None.

7.2 DC power requirements

The voltage and current requirements for a single drive are shown below. Values indicated apply at the drive connector.

Table 12: ST9146802SS DC power requirements

	Notes	ST9146802SS 1.5Gb mode		ST9146802SS 3.0Gb mode	
		(Amps)	(Amps)	(Amps)	(Amps)
Voltage		+5V	+12V [2]	+5V	+12V [2]
Regulation	[5]	±5%	±5% [2]	±5%	±5% [2]
Avg idle current $DC\bar{X}$	[1] [7]	0.79	0.15	0.77	0.15
Maximum starting current					
(peak DC) DC	3 σ [3]	0.87	1.59	0.85	1.59
(peak AC) AC	3 σ [3]	1.22	2.01	1.43	2.02
Delayed motor start (max) DC	3 σ [1] [4]	0.65	0.02	0.68	0.03
Operating current:					
Typical $DC\bar{X}$	[1] [6]	0.76	0.37	0.79	0.39
Maximum DC	3 σ [1]	0.78	0.40	0.83	0.43
Maximum (peak) DC	3 σ	1.34	1.19	1.41	1.21

Gen 3.5 LPSS measured data (latest available) plus 12V max start current margins of .15A for DC and .20A for AC. All others have 10% margin added.

Table 13: ST973402SS DC power requirements

	Notes	ST973402SS 1.5 Gb mode		ST973402SS 3.0 Gb mode	
		(Amps)	(Amps)	(Amps)	(Amps)
Voltage		+5V	+12V [2]	+5V	+12V [2]
Regulation	[5]	±5%	±5% [2]	±5%	±5% [2]
Avg idle current $\overline{DC\bar{X}}$	[1] [7]	0.81	0.13	0.79	0.12
Maximum starting current					
(peak DC) DC	3 σ [3]	0.88	1.38	0.87	1.38
(peak AC) AC	3 σ [3]	1.25	1.91	1.36	1.91
Delayed motor start (max) DC	3 σ [1] [4]	0.66	0.03	0.69	0.03
Operating current:					
Typical $\overline{DC\bar{X}}$	[1] [6]	0.65	0.31	0.81	0.34
Maximum DC	3 σ [1]	0.79	0.37	0.84	0.39
Maximum (peak) DC	3 σ	1.41	1.17	1.43	1.17

Gen 3.5 LPSS measured data (latest available) plus 12V max start current margins of .15A for DC and .20A for AC. All others have 10% margin added.

- [1] Measured with average reading DC ammeter. Instantaneous +12V current peaks will exceed these values. Power supply at nominal voltage. N (number of drives tested) = 6, 35 Degrees C ambient.
- [2] For +12 V, a –10% tolerance is allowed during initial spindle start but must return to ±5% before reaching 9936 RPM. The ±5% must be maintained after the drive signifies that its power-up sequence has been completed and that the drive is able to accept selection by the host initiator.
- [3] See +12V current profile in Figure 2.
- [4] This condition occurs after OOB and Speed Negotiation completes but before the drive has received the Notify Spinup primitive.
- [5] See paragraph 7.2.1, "Conducted noise immunity." Specified voltage tolerance includes ripple, noise, and transient response.
- [6] Operating condition is defined as random 8 block reads.
- [7] During idle, the drive heads are relocated every 60 seconds to a random location within the band from three-quarters to maximum track.

General DC power requirement notes.

1. Minimum current loading for each supply voltage is not less than 1.7% of the maximum operating current shown.
2. The +5V and +12V supplies should employ separate ground returns.
3. Where power is provided to multiple drives from a common supply, careful consideration for individual drive power requirements should be noted. Where multiple units are powered on simultaneously, the peak starting current must be available to each device.
4. Parameters, other than spindle start, are measured after a 10-minute warm up.
5. No terminator power.

7.2.1 Conducted noise immunity

Noise is specified as a periodic and random distribution of frequencies covering a band from DC to 10 MHz. Maximum allowed noise values given below are peak-to-peak measurements and apply at the drive power connector.

- +5v = 250 mV pp from 0 to 100 kHz to 20 MHz.
- +12v = 800 mV pp from 100 Hz to 8 KHz.
450 mV pp from 8 KHz to 20 KHz.
250 mV pp from 20 KHz to 5 MHz.

7.2.2 Power sequencing

The drive does not require power sequencing. The drive protects against inadvertent writing during power-up and down.

7.2.3 Current profiles

The +12V and +5V current profiles for the ST9146802SS and ST973402SS models are shown below.




Figure 2. ST9146802SS current profile

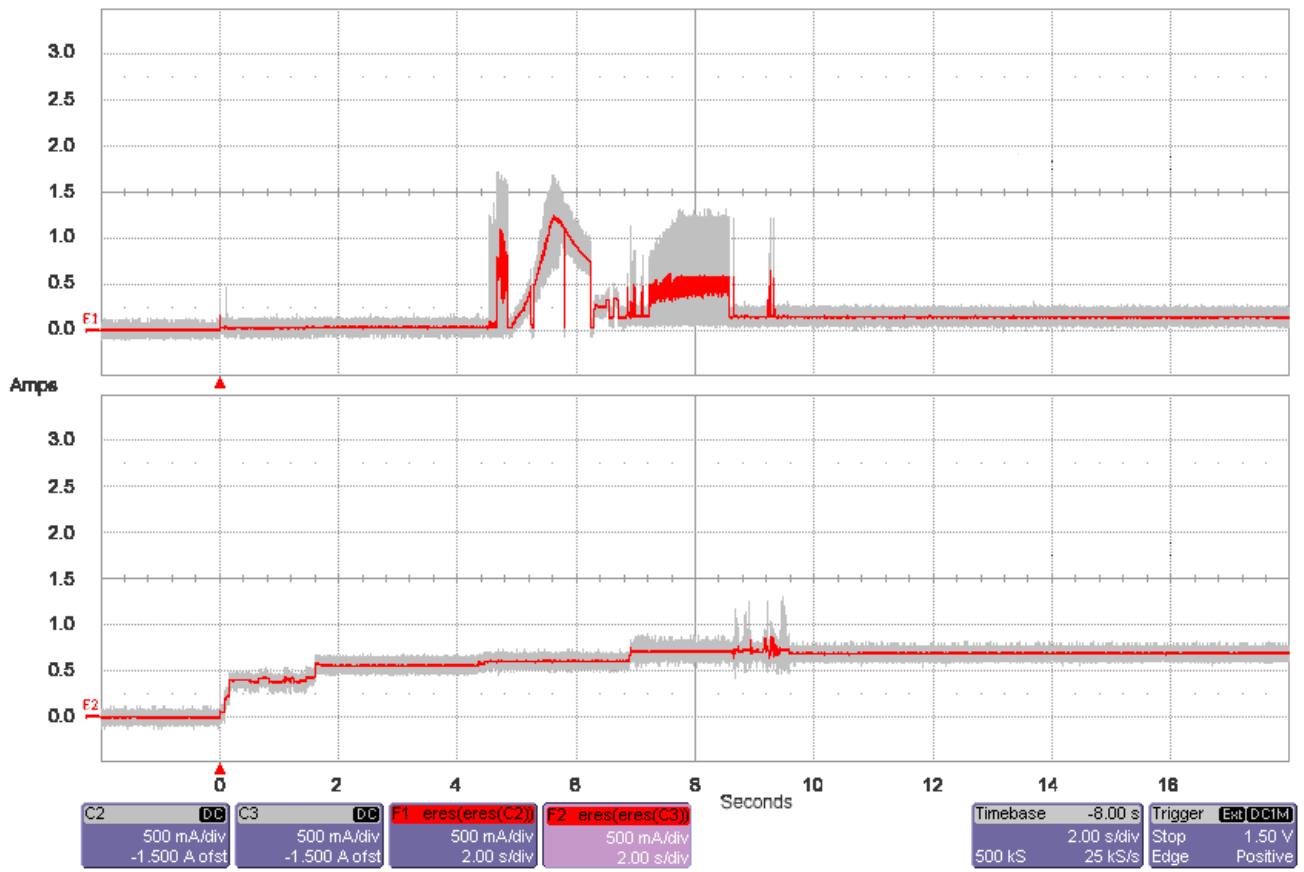


Figure 3. ST973402SS current profile

Note: All times and currents are typical. See Tables 12 and 13 for maximum current requirements.

7.3 Power dissipation

ST9146802SS in 3Gb operation

Typical power dissipation under idle conditions in 3Gb operation is 5.18W (17.68 BTUs per hour).

To obtain operating power for typical random read operations, refer to the following I/O rate curve (see Figure 4). Locate the typical I/O rate for a drive in your system on the horizontal axis and read the corresponding +5 volt current, +12 volt current, and total watts on the vertical axis. To calculate BTUs per hour, multiply watts by 3.4123.




Figure 4. ST9146802SS (3Gb) DC current and power vs. input/output operations per second

ST973402SS in 3Gb operation

Typical power dissipation under idle conditions in 3Gb operation is 5.14W (17.54 BTUs per hour).

To obtain operating power for typical random read operations, refer to the following I/O rate curve (see Figure 4). Locate the typical I/O rate for a drive in your system on the horizontal axis and read the corresponding +5 volt current, +12 volt current, and total watts on the vertical axis. To calculate BTUs per hour, multiply watts by 3.4123.

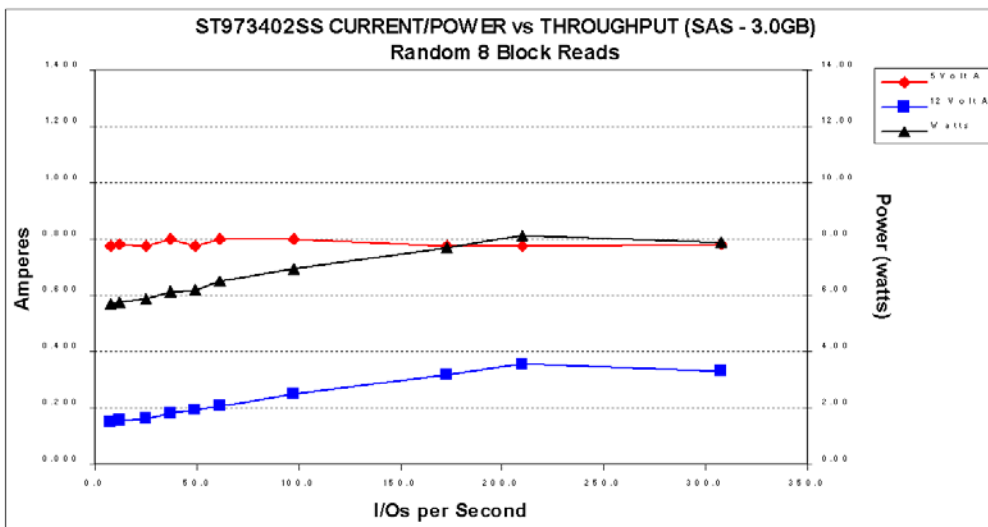


Figure 5. ST973402SS (3Gb) DC current and power vs. input/output operations per second

7.4 Environmental limits

Temperature and humidity values experienced by the drive must be such that condensation does not occur on any drive part. Altitude and atmospheric pressure specifications are referenced to a standard day at 58.7°F (14.8°C). Maximum wet bulb temperature is 82°F (28°C).

7.4.1 Temperature

a. Operating

With cooling designed to maintain the case temperatures, the drive meets all specifications over a 41°F to 131°F (5°C to 55°C) drive ambient temperature range with a maximum temperature gradient of 36°F (20°C) per hour. The enclosure for the drive should be designed such that these temperatures not exceeded. Air flow may be needed to achieve these temperature values (see Section 9.2). Operation at case temperatures above these values may adversely affect the drives ability to meet specifications.

The MTBF specification for the drive is based on operating in an environment that ensures that the case temperatures are not exceeded. Occasional excursions to drive ambient temperatures of 131°F (55°C) or 41°F (5°C) may occur without impact to specified MTBF. Air flow may be needed to achieve these temperatures. Continual or sustained operation at case temperatures above these values may degrade MTBF. The maximum allowable continuous or sustained HDA case temperature for the rated MTBF is 122°F (50°C).

To confirm that the required cooling for the electronics and HDA is provided, place the drive in its final mechanical configuration, perform random write/read operations. After the temperatures stabilize, measure the case temperature of the drive.

The maximum allowable HDA case temperature is 60°C. Operation of the drive at the maximum case temperature is intended for short time periods only. Continuous operation at the elevated temperatures will reduce product reliability.

b. Non-operating

–40° to 158°F (–40° to 70°C) package ambient with a maximum gradient of 36°F (20°C) per hour. This specification assumes that the drive is packaged in the shipping container designed by Seagate for use with drive.




Figure 6. Location of the HDA temperature check point

7.4.2 Relative humidity

The values below assume that no condensation on the drive occurs.

- a. Operating
5% to 95% non-condensing relative humidity with a maximum gradient of 20% per hour.
- b. Non-operating
5% to 95% non-condensing relative humidity.

7.4.3 Effective altitude (sea level)

- a. Operating
–200 to +10,000 feet (–60.96 to +3048 meters)
- b. Non-operating
–200 to +40,000 feet (–60.96 to +12,210 meters)

7.4.4 Shock and vibration

Shock and vibration limits specified in this document are measured directly on the drive chassis. If the drive is installed in an enclosure to which the stated shock and/or vibration criteria is applied, resonances may occur internally to the enclosure resulting in drive movement in excess of the stated limits. If this situation is apparent, it may be necessary to modify the enclosure to minimize drive movement.

The limits of shock and vibration defined within this document are specified with the drive mounted by any of the four methods shown in Figure 7, and in accordance with the restrictions of Section 9.3. Orientation of the side nearest the LED may be up or down.

7.4.4.1 Shock

- a. Operating—normal

The drive, as installed for normal operation, shall operate error free while subjected to intermittent shock not exceeding 15 Gs at a maximum duration of 11ms (half sinewave). The drive, as installed for normal operation, shall operate error free while subjected to intermittent shock not exceeding 20 Gs at a maximum duration of 2ms (half sinewave). Shock may be applied in the X, Y, or Z axis.

- b. Operating—abnormal

Equipment, as installed for normal operation, does not incur physical damage while subjected to intermittent shock not exceeding 40 Gs at a maximum duration of 11ms (half sinewave). Shock occurring at abnormal levels may promote degraded operational performance during the abnormal shock period. Specified operational performance will continue when normal operating shock levels resume. Shock may be applied in the X, Y, or Z axis. Shock is not to be repeated more than two times per second.

- c. Non-operating

The limits of non-operating shock shall apply to all conditions of handling and transportation. This includes both isolated drives and integrated drives.

The drive subjected to nonrepetitive shock not exceeding 75 Gs at a maximum duration of 11ms (half sinewave) shall not exhibit device damage or performance degradation. Shock may be applied in the X, Y, or Z axis.

The drive subjected to nonrepetitive shock not exceeding 300 Gs at a maximum duration of 2ms (half sinewave) does not exhibit device damage or performance degradation. Shock may be applied in the X, Y, or Z axis.

The drive subjected to nonrepetitive shock not exceeding 200 Gs at a maximum duration of 0.5ms (half sinewave) does not exhibit device damage or performance degradation. Shock may be applied in the X, Y, or Z axis.

d. Packaged

Disk drives shipped as loose load (not palletized) general freight will be packaged to withstand drops from heights as defined in the table below. For additional details refer to Seagate specifications 30190-001 (under 100 lbs/45 kg) or 30191-001 (over 100 lbs/45 Kg).

Package size	Packaged/product weight	Drop height
<600 cu in (<9,800 cu cm)	Any	60 in (1524 mm)
600-1800 cu in (9,800-19,700 cu cm)	0-20 lb (0 to 9.1 kg)	48 in (1219 mm)
>1800 cu in (>19,700 cu cm)	0-20 lb (0 to 9.1 kg)	42 in (1067 mm)
>600 cu in (>9,800 cu cm)	20-40 lb (9.1 to 18.1 kg)	36 in (914 mm)

Drives packaged in single or multipacks with a gross weight of 20 pounds (8.95 kg) or less by Seagate for general freight shipment shall withstand a drop test from 48 in (1070 mm) against a concrete floor or equivalent.




Figure 7. Recommended mounting

7.4.4.2 Vibration

a. Operating—normal

The drive as installed for normal operation, shall comply with the complete specified performance while subjected to continuous vibration not exceeding

1-500 Hz @ 0.5 G (zero to peak)

Vibration may be applied in the X, Y, or Z axis.

Operating normal translational random flat profile

10 - 400 Hz 0.4 GRMS

b. Operating—abnormal

Equipment as installed for normal operation shall not incur physical damage while subjected to periodic vibration not exceeding:

15 minutes of duration at major resonant frequency

10-500 Hz @ 1.5 G (X, Y, or Z axis)

Vibration occurring at these levels may degrade operational performance during the abnormal vibration period. Specified operational performance will continue when normal operating vibration levels are resumed. This assumes system recovery routines are available.

Operating abnormal translational random flat profile

10 - 500 Hz 1.2 GRMS

c. Non-operating

The limits of non-operating vibration shall apply to all conditions of handling and transportation. This includes both isolated drives and integrated drives.

The drive shall not incur physical damage or degraded performance as a result of continuous vibration not exceeding

5-22 Hz @ 0.040 in (1.02 mm) displacement (zero to peak)

20-500 Hz @ 3.0 G (zero to peak)

Vibration may be applied in the X, Y, or Z axis.

Non-operating translational random flat profile

10 - 500 Hz 1.2 GRMS

7.4.5 Air cleanliness

The drive is designed to operate in a typical office environment with minimal environmental control.

7.4.6 Corrosive environment

Seagate electronic drive components pass accelerated corrosion testing equivalent to 10 years exposure to light industrial environments containing sulfurous gases, chlorine and nitric oxide, classes G and H per ASTM B845. However, this accelerated testing cannot duplicate every potential application environment.

Users should use caution exposing any electronic components to uncontrolled chemical pollutants and corrosive chemicals as electronic drive component reliability can be affected by the installation environment. The silver, copper, nickel and gold films used in Seagate products are especially sensitive to the presence of sulfide, chloride, and nitrate contaminants. Sulfur is found to be the most damaging. In addition, electronic components should never be exposed to condensing water on the surface of the printed circuit board assembly (PCBA) or exposed to an ambient relative humidity greater than 95%. Materials used in cabinet fabrication, such as vulcanized rubber, that can outgas corrosive compounds should be minimized or eliminated. The useful life of any electronic equipment may be extended by replacing materials near circuitry with sulfide-free alternatives.

7.4.7 Acoustics

Sound power during idle mode shall be 3.1 bels typical when measured to ISO 7779 specification.

There will not be any discrete tones more than 10 dB above the masking noise on typical drives when measured according to Seagate specification 30553-001. There will not be any tones more than 24 dB above the masking noise on any drive.

7.4.8 Electromagnetic susceptibility

See Section 3.1.1.1.

7.5 Mechanical specifications

The following nominal dimensions are exclusive of the decorative front panel accessory. However, dimensions of the front panel are shown in figure below. Refer to Figure 8 for detailed mounting configuration dimensions. See Section 9.3, "Drive mounting."

Height:	0.583 in	14.8 mm
Width:	2.76 in	70 mm
Depth:	3.957 in	100.5 mm
Weight:	0.50 pounds	0.227 kilograms




Figure 8. Mounting configuration dimensions

8.0 Defect and error management

Seagate continues to use innovative technologies to manage defects and errors. These technologies are designed to increase data integrity, perform drive self-maintenance, and validate proper drive operation.

SCSI defect and error management involves drive internal defect/error management and SAS system error considerations (errors in communications between the initiator and the drive). In addition, Seagate provides the following technologies used to increase data integrity and drive reliability:

- Background Media Scan (see Section 8.4)
- Media Pre-Scan (see Section 8.5)
- Deferred Auto-Reallocation (see Section 8.6)
- Idle Read After Write (see Section 8.7)

The read error rates and specified storage capacities are not dependent on host (initiator) defect management routines.

8.1 Drive internal defects/errors

During the initial drive format operation at the factory, media defects are identified, tagged as being unusable, and their locations recorded on the drive primary defects list (referred to as the “P” list and also as the ETF defect list). At factory format time, these known defects are also reallocated, that is, reassigned to a new place on the medium and the location listed in the defects reallocation table. The “P” list is not altered after factory formatting. Locations of defects found and reallocated during error recovery procedures after drive shipment are listed in the “G” list (defects growth list). The “P” and “G” lists may be referenced by the initiator using the Read Defect Data command.

Details of the SCSI commands supported by the drive are described in the *SAS Interface Manual*. Also, more information on the drive Error Recovery philosophy is presented in the *SAS Interface Manual*.

8.2 Drive error recovery procedures

When an error occurs during drive operation, the drive, if programmed to do so, performs error recovery procedures to attempt to recover the data. The error recovery procedures used depend on the options previously set in the Error Recovery Parameters mode page. Error recovery and defect management may involve using several SCSI commands described in the *SAS Interface Manual*. The drive implements selectable error recovery time limits required in video applications.

The error recovery scheme supported by the drive provides a way to control the total error recovery time for the entire command in addition to controlling the recovery level for a single LBA. The total amount of time spent in error recovery for a command can be limited using the Recovery Time Limit bytes in the Error Recovery mode page. The total amount of time spent in error recovery for a single LBA can be limited using the Read Retry Count or Write Retry Count bytes in the Error Recovery mode page.

The drive firmware error recovery algorithms consists of 11 levels for read recoveries and five levels for write. Each level may consist of multiple steps, where a step is defined as a recovery function involving a single re-read or re-write attempt. The maximum level used by the drive in LBA recovery is determined by the read and write retry counts.

Table 14 equates the read and write retry count with the maximum possible recovery time for read and write recovery of individual LBAs. The times given do not include time taken to perform reallocations. Reallocations are performed when the ARRE bit (for reads) or AWRE bit (for writes) is one, the RC bit is zero, and the recovery time limit for the command has not yet been met. Time needed to perform reallocation is not counted against the recovery time limit.

When the RC bit is one, reallocations are disabled even if the ARRE or AWRE bits are one. The drive will still perform data recovery actions within the limits defined by the Read Retry Count, Write Retry Count, and Recovery Time Limit parameters. However, the drive does not report any unrecovered errors.

Table 14: Read and write retry count maximum recovery times

Read retry count	Maximum recovery time per LBA (cumulative, ms)	Write retry count	Maximum recovery time per LBA (cumulative, ms)
0	67.43	0	31.12
1	77.81	1	46.68
2	264.54	2	72.62
3	300.85	3	88.18
4	344.60	4	155.73
5	378.72	5 (default)	192.04
6	420.21		
7	529.22		
8	617.62		
9	659.11		
10	705.80		
11 (default)	1693.67		

Setting these retry counts to a value below the default setting could result in degradation of the unrecovered error rate. For example, suppose the read/write recovery page has the RC bit = 0, the read retry count set to 4, and the recovery time limit set to 450. A 4-block read command can take up to 371ms recovery time for each block and a maximum of 450ms recovery for all four blocks. If either of these limits is reached and a block has not yet been recovered, the command will end with Check Condition status and an unrecoverable read error will be reported.

8.3 SAS system errors

Information on the reporting of operational errors or faults across the interface is given in the *SAS Interface Manual*. The SSP Response returns information to the host about numerous kinds of errors or faults. The Receive Diagnostic Results reports the results of diagnostic operations performed by the drive.

Status returned by the drive to the initiator is described in the *SAS Interface Manual*. Status reporting plays a role in systems error management and its use in that respect is described in sections where the various commands are discussed.

8.4 Background Media Scan

Background Media Scan (BMS) is a self-initiated media scan. BMS is defined in the T10 document SPC-4 available from the T10 committee. BMS performs sequential reads across the entire pack of the media while the drive is idle. In RAID arrays, BMS allows hot spare drives to be scanned for defects prior to being put into service by the host system. On regular duty drives, if the host system makes use of the BMS Log Page, it can avoid placing data in suspect locations on the media. Unreadable and recovered error sites will be logged or reallocated per ARRE/AWRE settings.

With BMS, the host system can consume less power and system overhead by only checking BMS status and results rather than tying up the bus and consuming power in the process of host-initiated media scanning activity.

Since the background scan functions are only done during idle periods, BMS causes a negligible impact to system performance. The first BMS scan for a newly manufactured drive is performed as quickly as possible to verify the media and protect data by setting the “Start time after idle” to 5ms, all subsequent scans begin after 500ms of idle time. Other features that normally use idle time to function will function normally because BMS functions for bursts of 800ms and then suspends activity for 100ms to allow other background functions to operate.

BMS interrupts immediately to service host commands from the interface bus while performing reads. BMS will complete any BMS-initiated error recovery prior to returning to service host-initiated commands. Overhead associated with a return to host-servicing activity from BMS only impacts the first command that interrupted BMS, this results in a typical delay of about 1 ms.

8.5 Media Pre-Scan

Media Pre-Scan is a feature that allows the drive to repair media errors that would otherwise have been found by the host system during critical data accesses early in the drive’s life. The default setting for Media Pre-Scan is enabled on standard products. Media Pre-Scan checks each write command to determine if the destination LBAs have been scanned by BMS. If the LBAs have been verified, the drive proceeds with the normal write command. If the LBAs have not been verified by BMS, Pre-Scan will convert the write to a write verify to certify that the data was properly written to the disk.

Note. During Pre-Scan write verify commands, write performance may decrease by 50% until Pre-Scan completes. Write performance testing should be performed after Pre-Scan is complete. This may be checked by reading the BMS status.

To expedite the scan of the full pack and subsequently exit from the Pre-Scan period, BMS will begin scanning immediately when the drive goes to idle during the Pre-Scan period. In the event that the drive is in a high transaction traffic environment and is unable to complete a BMS scan within 24 power on hours BMS will disable Pre-Scan to restore full performance to the system.

8.6 Deferred Auto-Reallocation

Deferred Auto-Reallocation (DAR) simplifies reallocation algorithms at the system level by allowing the drive to reallocate unreadable locations on a subsequent write command. Sites are marked for DAR during read operations performed by the drive. When a write command is received for an LBA marked for DAR, the auto-reallocation process is invoked and attempts to rewrite the data to the original location. If a verification of this rewrite fails, the sector is re-mapped to a spare location.

This is in contrast to the system having to use the Reassign Command to reassign a location that was unreadable and then generate a write command to rewrite the data. DAR is most effective when AWRE and ARRE are enabled—this is the default setting from the Seagate factory. With AWRE and ARRE disabled DAR is unable to reallocate the failing location and will report an error sense code indicating that a write command is being attempted to a previously failing location.

8.7 Idle Read After Write

Idle Read After Write (IRAW) utilizes idle time to verify the integrity of recently written data. During idle periods, no active system requests, the drive reads recently written data from the media and compares it to valid write command data resident in the drives data buffer. Any sectors that fail the comparison result in the invocation of a rewrite and auto-reallocation process. The process attempts to rewrite the data to the original location. If a verification of this rewrite fails, the sector is re-mapped to a spare location.

9.0 Installation

Savvio disk drive installation is a plug-and-play process. There are no jumpers, switches, or terminators on the drive.

SAS drives are designed to be used in a host system that provides a SAS-compatible backplane with bays designed to accommodate the drive. In such systems, the host system typically provides a carrier or tray into which you need to mount the drive. Mount the drive to the carrier or tray provided by the host system using four M3 x 0.5 metric screws. When tightening the screws, use a maximum torque of 4.5 in-lb +/- 0.45 in-lb. Do not over-tighten or force the screws. You can mount the drive in any orientation.

Note. SAS drives are designed to be attached to the host system without I/O or power cables. If you intend to use the drive in a non-backplane host system, connecting the drive using high-quality cables is acceptable as long as the I/O cable length does not exceed 4 meters (13.1 feet).

Slide the carrier or tray into the appropriate bay in your host system using the instructions provided by the host system. This connects the drive directly to your system's SAS connector. The SAS connector is normally located on a SAS backpanel. See Section 10.4.1 for additional information about these connectors.

Power is supplied through the SAS connector.

The drive is shipped from the factory low-level formatted in 512-byte logical blocks. You need to reformat the drive only if you want to select a different logical block size.




Figure 9. Physical interface

9.1 Drive orientation

The drive may be mounted in any orientation. All drive performance characterizations, however, have been done with the drive in horizontal (disks level) and vertical (drive on its side) orientations, which are the two preferred mounting orientations.

9.2 Cooling

Cabinet cooling must be designed by the customer so that the ambient temperature immediately surrounding the drive will not exceed temperature conditions specified in Section 7.4.1, "Temperature."

The rack, cabinet, or drawer environment for the drive must provide heat removal from the electronics and head and disk assembly (HDA). You should confirm that adequate heat removal is provided using the temperature measurement guidelines described in Section 7.4.1.

Forced air flow may be required to keep temperatures at or below the temperatures specified in Section 7.4.1 in which case the drive should be oriented, or air flow directed, so that the least amount of air flow resistance is created while providing air flow to the electronics and HDA. Also, the shortest possible path between the air inlet and exit should be chosen to minimize the travel length of air heated by the drive and other heat sources within the rack, cabinet, or drawer environment.

If forced air is determined to be necessary, possible air-flow patterns are shown in Figure 10. The air-flow patterns are created by one or more fans, either forcing or drawing air as shown in the illustrations. Conduction, convection, or other forced air-flow patterns are acceptable as long as the temperature measurement guidelines of Section 7.4.1 are met.




Figure 10. Air flow

9.3 Drive mounting

Mount the drive using the bottom or side mounting holes. If you mount the drive using the bottom holes, ensure that you do not physically distort the drive by attempting to mount it on a stiff, non-flat surface.

The allowable mounting surface stiffness is 80 lb/in (14.0 N/mm). The following equation and paragraph define the allowable mounting surface stiffness:

$$K \times X = F \leq 15\text{lb} = 67\text{N}$$

where K is the mounting surface stiffness (units in lb/in or N/mm) and X is the out-of-plane surface distortion (units in inches or millimeters). The out-of-plane distortion (X) is determined by defining a plane with three of the four mounting points fixed and evaluating the out-of-plane deflection of the fourth mounting point when a known force (F) is applied to the fourth point.

9.4 Grounding

Signal ground (PCBA) and HDA ground are connected together in the drive and cannot be separated by the user. The equipment in which the drive is mounted is connected directly to the HDA and PCBA with no electrically isolating shock mounts. If it is desired for the system chassis to not be connected to the HDA/PCBA ground, the systems integrator or user must provide a nonconductive (electrically isolating) method of mounting the drive in the host equipment.

Increased radiated emissions may result if you do not provide the maximum surface area ground connection between system ground and drive ground. This is the system designer's and integrator's responsibility.

10.0 Interface requirements

This section partially describes the interface requirements as implemented on Savvio drives. Additional information is provided in the *SAS Interface Manual* (part number 100293071).

10.1 SAS features

This section lists the SAS-specific features supported by Savvio drives.

10.1.1 task management functions

Table 15 lists the SAS task management functions supported.

Table 15: SAS task management functions supported

Task name	Supported
Abort Task	Yes
Clear ACA	Yes
Clear task set	Yes
Abort task set	Yes
Logical Unit Reset	Yes
Query Task	Yes

10.1.2 task management responses

Table 16 lists the SAS response codes returned for task management functions supported.

Table 16: Task management response codes

Function name	Response code
Function complete	00
Invalid frame	02
Function not supported	04
Function failed	05
Function succeeded	08
Invalid logical unit	09

10.2 Dual port support

Savvio SAS drives have two independent ports. These ports may be connected in the same or different SCSI domains. Each drive port has a unique SAS address.

The two ports run at the same link rate. The first port to successfully complete speed negotiation sets the link rate support by both ports. When the second port participates in speed negotiation, it indicates the only supported speed is the speed selected by the first port. If the first port to complete speed negotiation loses sync before the second port completes speed negotiation, both ports revert back to the power on condition of allowing either link rate (1.5 or 3.0Gb/s).

Subject to buffer availability, the Savvio drives support:

- Concurrent port transfers—The drive supports receiving COMMAND, TASK management transfers on both ports at the same time.
- Full duplex—The drive supports sending XFER_RDY, DATA and RESPONSE transfers while receiving frames on both ports.

10.3 SCSI commands supported

Table 17 lists the SCSI commands supported by Savvio drives.

Table 17: Commands supported by Savvio SAS family drive

Command name	Command code	Supported
Change Definition	40h	N
Compare	39h	N
Copy	18h	N
Copy and Verify	3Ah	N
Format Unit [1]	04h	Y
DPRY bit supported		N
DCRT bit supported		Y
STPF bit supported		Y
IP bit supported		Y
DSP bit supported		Y
IMMED bit supported		Y
VS (vendor specific)		N
Inquiry	12h	Y
Date Code page (C1h)		Y
Device Behavior page (C3h)		Y
Firmware Numbers page (C0h)		Y
Implemented Operating Def page (81h)		Y
Jumper Settings page (C2h)		Y
Supported Vital Product Data page (00h)		Y
Unit Serial Number page (80h)		Y
Lock-unlock cache	36h	N
Log Select	4Ch	Y
PCR bit		Y
DU bit		N
DS bit		Y
TSD bit		Y
ETC bit		N
TMC bit		N
LP bit		N
Protocol-specific Log Page for SAS (18h)		Y
Log Sense	4Dh	Y
Application Client Log page (0Fh)		Y
Buffer Over-run/Under-run page (01h)		N
Cache Statistics page (37h)		Y
Factory Log page (3Eh)		Y
Information Exceptions Log page (2Fh)		N

Table 17: Commands supported by Savvio SAS family drive (continued)

Command name	Command code	Supported
Last n Deferred Errors or Asynchronous Events page (0Bh)		N
Last n Error Events page (07h)		N
Non-medium Error page (06h)		Y
Pages Supported list (00h)		Y
Read Error Counter page (03h)		Y
Read Reverse Error Counter page (04h)		N
Self-test Results page (10h)		Y
Start-stop Cycle Counter page (0Eh)		Y
Temperature page (0Dh)		Y
Verify Error Counter page (05h)		Y
Write error counter page (02h)		Y
Mode Select (same pages as Mode Sense 1Ah)	15h	Y [2]
Mode Select (10) (same pages as Mode Sense 1Ah)	55h	Y
Mode Sense	1Ah	Y [2]
Caching Parameters page (08h)		Y
Control Mode page (0Ah)		Y
Disconnect/Reconnect (02h)		Y
Error Recovery page (01h)		Y
Format page (03h)		Y
Information Exceptions Control page (1Ch)		Y
Notch and Partition Page (0Ch)		N
Protocol-Specific Port page (19h)		Y
Power Condition page (1Ah)		Y
Rigid Disk Drive Geometry page (04h)		Y
Unit Attention page (00h)		Y
Verify Error Recovery page (07h)		Y
Xor Control page (10h)		N
Mode Sense (10) (same pages as Mode Sense 1Ah)	5Ah	Y
Persistent Reserve In	5Eh	Y
Persistent Reserve Out	5Fh	Y
Prefetch	34h	N
Read (6)	08h	Y
Read (10)	28h	Y
DPO bit supported		Y
FUA bit supported		Y
Read (12)	A8h	N
Read (16)	88h	N
Read (32)	7Fh/0009h	N
Read Buffer (modes 0, 2, 3, Ah and Bh supported)	3Ch	Y

Table 17: Commands supported by Savvio SAS family drive (continued)

Command name	Command code	Supported
Read Capacity (10)	25h	Y
Read Capacity (16)	9Eh/10h	N
Read Defect Data (10)	37h	Y
Read Defect Data (12)	B7h	Y
Read Long	3Eh	Y
Read Long (16)	9Eh/11h	N
Reassign Blocks	07h	Y
Receive Diagnostic Results	1Ch	Y
Supported Diagnostics pages (00h)		Y
Translate page (40h)		Y
Release	17h	Y
Release (10)	57h	Y
Report LUNs	A0h	Y
Request Sense	03h	Y
Actual Retry Count bytes		Y
Extended Sense		Y
Field Pointer bytes		Y
Reserve	16h	Y
3rd Party Reserve		Y
Extent Reservation		N
Reserve (10)	56h	Y
3rd Party Reserve		Y
Extent Reservation		N
Rezero Unit	01h	Y
Search Data Equal	31h	N
Search Data High	30h	N
Search Data Low	32h	N
Seek (6)	0Bh	Y
Seek (10)	2Bh	Y
Send Diagnostics	1Dh	Y
Supported Diagnostics pages (00h)		Y
Translate page (40h)		Y
Set Limits	33h	N
Start Unit/Stop Unit (spindle ceases rotating)	1Bh	Y
Synchronize Cache	35h	Y
Synchronize Cache (16)	91h	N
Test Unit Ready	00h	Y
Verify (10)	2Fh	Y
BYTCHK bit		Y

Table 17: Commands supported by Savvio SAS family drive (continued)

Command name	Command code	Supported
Verify (12)	AFh	N
Verify (16)	AFh	N
Verify (32)	7Fh/000Ah	N
Write (6)	0Ah	Y
Write (10)	2Ah	Y
DPO bit		Y
FUA bit		Y
Write (12)	AAh	N
Write (16)	8Ah	N
Write (32)	7Fh/000Bh	N
Write and Verify (10)	2Eh	Y
DPO bit		Y
Write and Verify (12)	A Eh	N
Write and Verify (16)	8Eh	N
Write and Verify (32)	7Fh/000Ch	N
Write Buffer (modes 0, 2, supported)	3Bh	Y
Firmware Download option (modes 5, 7, Ah and Bh supported) [3]		Y
Write Long (10)	3Fh	Y
Write Long (16)	9Fh/11h	N
Write Same (10)	41h	Y
PBdata		N
LBdata		N
Write Same (16)	93h	N
Write Same (32)	7Fh/000Dh	N
XDRead	52h	N
XDWrite	50h	N
XPWrite	51h	N

[1] Savvio drives can format to 512, 520, or 528 bytes per logical block.

[2] **Warning.** Power loss during flash programming can result in firmware corruption. This usually makes the drive inoperable.

[3] Reference Mode Sense command 1Ah for mode pages supported.

[4] Y = Yes. Command is supported.

N = No. Command is not supported.

A = Support is available on special request.

10.3.1 Inquiry data

Table 18 lists the Inquiry command data that the drive should return to the initiator per the format given in the *SAS Interface Manual*.

Table 18: Savvio inquiry data

Bytes	Data (hex)																
0-15	00	00	xx**	12	8B	00	PP	0A	53	45	41	47	41	54	45	20	Vendor ID
16-31	[53	54	39	31	34	36	38	30	32	53	53]	20	20	20	20	20	Product ID
32-47	R#	R#	R#	R#	S#	S#	S#	S#	S#	S#	S#	S#	00	00	00	00	
48-63	00	00	00	00	00	00	00	00	00	00	00	00	00	00	00	00	
64-79	00	00	00	00	00	00	00	00	00	00	00	00	00	00	00	00	
80-95	00	00	00	00	00	00	00	00	00	00	00	00	00	00	00	00	
96-111	00	43	6F	70	79	72	69	67	68	74	20	28	63	29	20	32*	*Copyright
112-127	30*	30*	36*	20	53	65	61	67	61	74	65	20	41	6C	6C	20	notice
128-143	72	69	67	68	74	73	20	72	65	73	65	72	76	65	64	20	

* Copyright year (changes with actual year).

** SCSI Revision support. See the appropriate SPC release documentation for definitions.

PP 10 = Inquiry data for an Inquiry command received on Port A.

30 = Inquiry data for an Inquiry command received on Port B.

R# Four ASCII digits representing the last four digits of the product firmware release number.

S# Eight ASCII digits representing the eight digits of the product serial number.

[] Bytes 16 through 26 reflect model of drive. The table above shows the hex values for Model ST9146802SS.

Refer to the values below for the values of bytes 16 through 26 of you particular model:

ST9146802SS 53 54 39 31 34 36 38 30 32 53 53

ST973402SS 53 54 39 37 33 34 30 32 53 53 20

10.3.2 Mode Sense data

The Mode Sense command provides a way for the drive to report its operating parameters to the initiator. The drive maintains four sets of mode parameters:

1. Default values

Default values are hard-coded in the drive firmware stored in flash E-PROM (nonvolatile memory) on the drive's PCB. These default values can be changed only by downloading a complete set of new firmware into the flash E-PROM. An initiator can request and receive from the drive a list of default values and use those in a Mode Select command to set up new current and saved values, where the values are changeable.

2. Saved values

Saved values are stored on the drive's media using a Mode Select command. Only parameter values that are allowed to be changed can be changed by this method. Parameters in the saved values list that are not changeable by the Mode Select command get their values from default values storage.

When power is applied to the drive, it takes saved values from the media and stores them as current values in volatile memory. It is not possible to change the current values (or the saved values) with a Mode Select command before the drive achieves operating speed and is "ready." An attempt to do so results in a "Check Condition" status.

On drives requiring unique saved values, the required unique saved values are stored into the saved values storage location on the media prior to shipping the drive. Some drives may have unique firmware with unique default values also.

On standard OEM drives, the saved values are taken from the default values list and stored into the saved values storage location on the media prior to shipping.

3. Current values

Current values are volatile values being used by the drive to control its operation. A Mode Select command can be used to change the values identified as changeable values. Originally, current values are installed from saved or default values after a power on reset, hard reset, or Bus Device Reset message.

4. Changeable values

Changeable values form a bit mask, stored in nonvolatile memory, that dictates which of the current values and saved values can be changed by a Mode Select command. A one (1) indicates the value can be changed. A zero (0) indicates the value is not changeable. For example, in Table 19, refer to Mode page 81, in the row entitled "CHG." These are hex numbers representing the changeable values for Mode page 81. Note in columns 5 and 6 (bytes 04 and 05), there is 00h which indicates that in bytes 04 and 05 none of the bits are changeable. Note also that bytes 06, 07, 09, 10, and 11 are not changeable, because those fields are all zeros. In byte 02, hex value FF equates to the binary pattern 11111111. If there is a zero in any bit position in the field, it means that bit is not changeable. Since all of the bits in byte 02 are ones, all of these bits are changeable.

The changeable values list can only be changed by downloading new firmware into the flash E-PROM.

Note. Because there are often several different versions of drive control firmware in the total population of drives in the field, the Mode Sense values given in the following tables may not exactly match those of some drives.

The following tables list the values of the data bytes returned by the drive in response to the Mode Sense command pages for SCSI implementation (see the *SAS Interface Manual*).

DEF = Default value. Standard OEM drives are shipped configured this way.

CHG = Changeable bits; indicates if default value is changeable.

Table 19: Mode Sense data changeable and default values for ST9146802SS drives

MODE DATA HEADER:

03 fa 00 10 00 00 00 08

BLOCK DESCRIPTOR:

11 17 73 30 00 00 02 00

MODE PAGES:

DEF 81 0a c0 0b ff 00 00 00 05 00 ff ff

CHG 81 0a ff ff 00 00 00 00 ff 00 ff ff

DEF 82 0e 00 00 00 00 00 00 00 00 01 3a 00 00 00 00

CHG 82 0e 00 00 00 00 00 00 00 00 ff ff 00 00 00 00

DEF 83 16 1e 3c 00 00 00 1c 00 00 03 03 02 00 00 01 00 6c 00 4e 40 00 00 00

CHG 83 16 00 00 00 00 ff ff 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00

DEF 84 16 01 08 fc 04 00 00 00 00 00 00 00 00 00 00 00 00 00 00 27 31 00 00

CHG 84 16 00

DEF 87 0a 00 0b ff 00 00 00 00 00 ff ff

CHG 87 0a 0f ff 00 00 00 00 00 ff ff

DEF 88 12 14 00 ff ff 00 00 ff ff ff ff 80 20 00 00 00 00 00 00

CHG 88 12 a5 00 00 00 ff ff ff ff 00 00 20 00 00 00 00 00 00 00

DEF 8a 0a 02 00 00 00 00 00 00 08 16

CHG 8a 0a 03 f0 00 00 00 00 00 00 00

DEF 18 06 06 00 00 00 00 00

CHG 18 06 00 00 00 00 00 00

DEF 99 06 06 00 07 d0 00 00

CHG 99 06 10 00 ff ff ff ff

DEF 9a 0a 00 02 00 00 00 05 00 00 00 04

CHG 9a 0a 00 03 ff ff ff ff 00 00 00 00

DEF 9c 0a 10 00 00 00 00 00 00 00 01

CHG 9c 0a 9d 0f ff ff ff ff ff ff ff

DEF 80 06 00 80 0f 00 00 00

CHG 80 06 b7 c0 0f 00 00 00

Table 20: Mode Sense data changeable and default values for ST973402SS drives

MODE DATA HEADER:

03 fa 00 10 00 00 00 08

BLOCK DESCRIPTOR:

08 8b b9 98 00 00 02 00

MODE PAGES:

DEF 81 0a c0 0b ff 00 00 00 05 00 ff ff
 CHG 81 0a ff ff 00 00 00 00 ff 00 ff ff

DEF 82 0e 00 00 00 00 00 00 00 00 01 3a 00 00 00 00
 CHG 82 0e 00 00 00 00 00 00 00 00 ff ff 00 00 00 00

DEF 83 16 0f 1e 00 00 00 0e 00 00 03 03 02 00 00 01 00 6c 00 4e 40 00 00 00
 CHG 83 16 00 00 00 00 ff ff 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00

DEF 84 16 01 08 fc 02 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 27 31 00 00
 CHG 84 16 00

DEF 87 0a 00 0b ff 00 00 00 00 00 ff ff
 CHG 87 0a 0f ff 00 00 00 00 00 ff ff

DEF 88 12 14 00 ff ff 00 00 ff ff ff ff 80 20 00 00 00 00 00 00
 CHG 88 12 a5 00 00 00 ff ff ff ff 00 00 20 00 00 00 00 00 00 00

DEF 8a 0a 02 00 00 00 00 00 00 04 2e
 CHG 8a 0a 03 f0 00 00 00 00 00 00 00

DEF 18 06 06 00 00 00 00 00
 CHG 18 06 00 00 00 00 00 00

DEF 99 06 06 00 07 d0 00 00
 CHG 99 06 10 00 ff ff ff ff

DEF 9a 0a 00 02 00 00 00 05 00 00 00 04
 CHG 9a 0a 00 03 ff ff ff ff 00 00 00 00

DEF 9c 0a 10 00 00 00 00 00 00 00 01
 CHG 9c 0a 9d 0f ff ff ff ff ff ff ff

DEF 80 06 00 80 0f 00 00 00
 CHG 80 06 b7 c0 0f 00 00 00

10.4 Miscellaneous operating features and conditions

Table 21 lists various features and conditions. A “Y” in the support column indicates the feature or condition is supported. An “N” in the support column indicates the feature or condition is not supported.

Table 21: Miscellaneous features

Supported	Feature or condition
N	Automatic contingent allegiance
N	Asynchronous event notification
N	Synchronized (locked) spindle operation
Y	Segmented caching
N	Zero latency read
Y	Queue tagging (up to 64 queue tags supported)
Y	Deferred error handling
Y	Parameter rounding (controlled by Round bit in Mode Select page 0)
Y	Reporting actual retry count in Extended Sense bytes 15, 16, and 17
N	Adaptive caching
Y	SMP = 1 in Mode Select command needed to save RPL and rotational offset bytes

Table 22: Miscellaneous status

Supported	Status
Y	Good
Y	Check condition
Y	Condition met/good
Y	Busy
Y	Intermediate/good
Y	Intermediate/condition met/good
Y	Reservation conflict
Y	Task set full
N	ACA active
N	ACA active, faulted initiator

10.4.1 SAS physical interface

Figure 11 shows the location of the SAS device connector J1. Figures 12 and 13 provide the dimensions of the SAS connector.

Details of the physical, electrical, and logical characteristics are provided within this section. The operational aspects of Seagate's SAS drives are provided in the *SAS Interface Manual*.




Figure 11. Physical interface




Figure 12. SAS device plug dimensions




Figure 13. SAS device plug dimensions (detail)

10.4.2 Physical characteristics

This section defines physical interface connector.

10.4.3 Connector requirements

Contact your preferred connector manufacturer for mating part information. Part numbers for SAS connectors will be provided in a future revision of this publication when production parts are available from major connector manufacturers.

The SAS device connector is illustrated in Figures 12 and 13.

10.4.4 Electrical description

SAS drives use the device connector for:

- DC power
- SAS interface
- Activity LED

This connector is designed to either plug directly into a backpanel or accept cables.

10.4.5 Pin descriptions

This section provides a pin-out of the SAS device and a description of the functions provided by the pins.

Table 23: SAS pin descriptions

Pin	Signal name	Signal type	Pin	Signal name	Signal type
S1	Port A Ground		P1*	NC (reserved 3.3Volts)	
S2*	+Port A_in	Diff. input pair	P2*	NC (reserved 3.3Volts)	
S3*	-Port A_in		P3	NC (reserved 3.3Volts)	
S4	Port A Ground		P4	Ground	
S5*	-Port A_out	Diff output pair	P5	Ground	
S6*	+Port A_out		P6	Ground	
S7	Port A Ground		P7	5 Volts charge	
S8	Port B Ground		P8*	5 Volts	
S9*	+Port B_in	Diff. input pair	P9*	5 Volts	
S10*	-Port B_in		P10	Ground	
S11	Port A Ground		P11*	Ready LED	Open collector out
S12*	-Port B_out	Diff output pair	P12	Ground	
S13*	+Port B_out		P13	12 Volts charge	
S14	Port B Ground		P14*	12 Volts	
			P15*	12 Volts	

* - Short pin to support hot plugging

NC - No connection in the drive.

10.4.6 SAS transmitters and receivers

A typical SAS differential copper transmitter and receiver pair is shown in Figure 14. The receiver is AC coupling to eliminate ground shift noise.




Figure 14. SAS transmitters and receivers

10.4.7 Power

The drive receives power (+5 volts and +12 volts) through the SAS device connector.

Three +12 volt pins provide power to the drive, 2 short and 1 long. The current return for the +12 volt power supply is through the common ground pins. The supply current and return current must be distributed as evenly as possible among the pins.

Three +5 volt pins provide power to the drive, 2 short and 1 long. The current return for the +5 volt power supply is through the common ground pins. The supply current and return current must be distributed as evenly as possible among the pins.

Current to the drive through the long power pins may be limited by the system to reduce inrush current to the drive during hot plugging.

10.5 Signal characteristics

This section describes the electrical signal characteristics of the drive's input and output signals. See Table 23 for signal type and signal name information.

10.5.1 Ready LED Out

The Ready LED Out signal is driven by the drive as indicated in Table 24.

Table 24: Ready LED Out conditions

Normal command activity	LED status	
	0	1
Ready LED Meaning bit mode page 19h	0	1
Spun down and no activity	Off	Off
Spun down and activity (command executing)	On	On
Spun up and no activity	On	Off
Spun up and activity (command executing)	Off	On
Spinning up or down	Blinks steadily (50% on and 50% off, 0.5 seconds on and off for 0.5 seconds)	

Table 24: Ready LED Out conditions

Normal command activity	LED status
Format in progress, each cylinder change	Toggles on/off

The Ready LED Out signal is designed to pull down the cathode of an LED. The anode is attached to the proper +3.3 volt supply through an appropriate current limiting resistor. The LED and the current limiting resistor are external to the drive. See Table 25 for the output characteristics of the LED drive signals.

Table 25: LED drive signal

State	Test condition	Output voltage
LED off, high	$0\text{ V} \leq V_{OH} \leq 3.6\text{ V}$	$-100\ \mu\text{A} < I_{OH} < 100\ \mu\text{A}$
LED on, low	$I_{OL} = 15\text{ mA}$	$0 \leq V_{OL} \leq 0.225\text{ V}$

10.5.2 Differential signals

The drive SAS differential signals comply with the intra-enclosure (internal connector) requirements of the SAS standard.

Table 26 defines the general interface characteristics.

Table 26: General interface characteristics

Characteristic	Units	1.5 Gbps	3.0 Gbps
Bit rate (nominal)	Mbaud	1,500	3,000
Unit interval (UI)(nominal)	ps	666.6	333.3
Impedance (nominal, differential)	ohm	100	100
Transmitter transients, maximum	V	± 1.2	± 1.2
Receiver transients, maximum	V	± 1.2	± 1.2

10.5.2.1 Eye masks

10.5.2.1.1 Eye masks overview

The eye masks are graphical representations of the voltage and time limits on the signal at the compliance point. The time values between X1 and (1 - X1) cover all but 10^{-12} of the jitter population. The random content of the total jitter population has a range of ± 7 standard deviations.

10.5.2.1.2 Receive eye mask

Figure 15 describes the receive eye mask. This eye mask applies to jitter after the application of a single pole high-pass frequency-weighting function that progressively attenuates jitter at 20 dB/decade below a frequency of $((\text{bit rate}) / 1.667)$.




Figure 15. Receive eye mask

Verifying compliance with the limits represented by the receive eye mask should be done with reverse channel traffic present in order that the effects of crosstalk are taken into account.

10.5.2.1.3 Jitter tolerance masks

Figure 16 describes the receive tolerance eye masks and is constructed using the $X2$ and $Z2$ values given in table 29. $X1_{OP}$ is half the value for total jitter in table 29 and $X1_{TOL}$ is half the value for total jitter in table 30, for jitter frequencies above $((\text{bit rate}) / 1.667)$.




Figure 16. Receive tolerance eye mask

The leading and trailing edge slopes of figure 15 shall be preserved. As a result the amplitude value of Z1 is less than that given in table 28 and $Z1_{TOL}$ and $Z1_{OP}$ shall be defined from those slopes by the following equation:

$$Z1_{TOL} = Z1_{OP} \times \frac{X2_{OP} - ((0, 5) \times \text{additional sinusoidal jitter}) - X1_{OP}}{X2_{OP} - X1_{OP}}$$

where:

$Z1_{TOL}$ is the value for Z1 to be used for the tolerance masks; and

$Z1_{OP}$, $X1_{OP}$, and $X2_{OP}$ are the values in table 28 for Z1, X1, and X2.

The X1 points in the receive tolerance masks are greater than the X1 points in the receive masks, due to the addition of sinusoidal jitter.

Figure 17 defines the sinusoidal jitter mask.




Figure 17. Sinusoidal jitter mask

10.5.2.2 Transmitter signal characteristics

Table 27 specifies the signal requirements at the transmitter end of a TxRx connection as measured into the zero-length test load. All specifications are based on differential measurements.

The OOB sequence is performed at signal voltage levels corresponding to the lowest supported transfer rate. Table 27 specifies the signal characteristics.

Table 27: Transmitter signal characteristics

Signal characteristic ^a	Units	1.5 Gbps	3.0 Gbps
Skew ^b	ps	20	15
Tx Off Voltage ^c	mV(P-P)	< 50	< 50
Maximum rise/fall time ^d	ps	273	137
Minimum rise/fall time ^d	ps	67	67
Maximum transmitter output imbalance ^e	%	10	10
OOB offset delta ^f	mV	± 25	± 25
OOB common mode delta ^g	mV	± 50	± 50

a All tests in this table shall be performed with zero-length test load shown in figure 19.

b The skew measurement shall be made at the midpoint of the transition with a repeating 0101b pattern on the physical link. The same stable trigger, coherent to the data stream, shall be used for both the Tx+ and Tx- signals. Skew is defined as the time difference between the means of the midpoint crossing times of the Tx+ signal and the Tx- signal.

c The transmitter off voltage is the maximum A.C. voltage measured at compliance points when the transmitter is unpowered or transmitting D.C. idle (e.g., during idle time of an OOB signal).

d Rise/fall times are measured from 20 % to 80 % of the transition with a repeating 0101b pattern on the physical link.

e The maximum difference between the V+ and V- A.C. RMS transmitter amplitudes measured on a CJTPAT test pattern (see 10.5.2.3.3) into the test load shown in figure 19, as a percentage of the average of the V+ and V- A.C. RMS amplitudes.

f The maximum difference in the average differential voltage (D.C. offset) component between the burst times and the idle times of an OOB signal.

g The maximum difference in the average of the common mode voltage between the burst times and the idle times of an OOB signal.

10.5.2.3 Receiver signal characteristics

Table 28 defines the compliance point requirements of the signal at the receiver end of a TxRx connection as measured into the test loads specified in figure 18 and figure 19.

Table 28: Receiver signal characteristics

Signal characteristic	Units	1.5 Gbps	3.0 Gbps
Jitter (see figure 15) ^b	N/A	See table 29	See table 29
2 x Z2	mV(P-P)	1,200	1,600
2 x Z1	mV(P-P)	325	275
X1 ^a	UI	0.275	0.275
X2	UI	0.50	0.50
Skew ^d	ps	80	75
Max voltage (non-op)	mV(P-P)	2.000	2.000
Minimum OOB ALIGN burst amplitude ^c	mV(P-P)	240	240
Maximum noise during OOB idle time ^c	mV(P-P)	120	120
Max near-end crosstalk ^e	mV(P-P)	100	100

- a The value for X1 shall be half the value given for total jitter in table 29. The test or analysis shall include the effects of a single pole high-pass frequency-weighting function that progressively attenuates jitter at 20 dB/decade below a frequency of ((bit rate) / 1,667).
- b The value for X1 applies at a total jitter probability of 10^{-12} . At this level of probability direct visual comparison between the mask and actual signals is not a valid method for determining compliance with the jitter output requirements.
- c With a measurement bandwidth of 1.5 times the baud rate (i.e. 4.5 GHz for 3.0 Gbps).
- d The skew measurement shall be made at the midpoint of the transition with a repeating 0101b pattern on the physical link. The same stable trigger, coherent to the data stream, shall be used for both the Rx+ and Rx- signals. Skew is defined as the time difference between the means of the midpoint crossing times of the Rx+ signal and the Rx- signal.
- e Near-end crosstalk is the unwanted signal amplitude at receiver terminals DR, CR, and XR coupled from signals and noise sources other than the desired signal. Refer to SFF-8410.

10.5.2.3.1 Jitter

Table 29 defines the maximum allowable jitter.

Table 29: Maximum allowable jitter

1.5 Gbps ^{m, n}		3.0 Gbps ^{m, n}	
Deterministic jitter ^d	Total jitter ^{c,d,e,f}	Deterministic jitter ^e	Total jitter ^{c,d,e,f}
0.35	0.55	0.35	0.55

- a Units are in UI.
- b The values for jitter in this section are measured at the average amplitude point.
- c Total jitter is the sum of deterministic jitter and random jitter. If the actual deterministic jitter is less than the maximum specified, then the random jitter may increase as long as the total jitter does not exceed the specified maximum total jitter.
- d Total jitter is specified at a probability of 10^{-12} .
- e The deterministic and total values in this table apply to jitter after application of a single pole high-pass frequency-weighting function that progressively attenuates jitter at 20 dB/decade below a frequency of ((bit rate) / 1 667).
- f If total jitter received at any point is less than the maximum allowed, then the jitter distribution of the signals is allowed to be asymmetric. The total jitter plus the magnitude of the asymmetry shall not exceed the allowed maximum total jitter. The numerical difference between the average of the peaks with a BER < 10^{-12} and the average of the individual events is the measure of the asymmetry. Jitter peak-to-peak measured < (maximum total jitter - |Asymmetry|).

10.5.2.3.2 Receiver jitter tolerance

Table 30 defines the amount of jitter the receiver shall tolerate .

Table 30: Receiver jitter tolerance

1.5 Gbps ^a			3.0 Gbps ^a		
Sinusoidal jitter ^{b,c}	Deterministic jitter ^{e,f,h}	Total jitter ^h	Sinusoidal jitter ^{b,d}	Deterministic jitter ^{e,g,h}	Total jitter ^h
0.10	0.35	0.65	0.10	0.35	0.65

a Units are in UI.

b The jitter values given are normative for a combination of deterministic jitter, random jitter, and sinusoidal jitter that receivers shall be able to tolerate without exceeding a BER of 10^{-12} . Receivers shall tolerate sinusoidal jitter of progressively greater amplitude at lower frequencies, according to the mask in figure 17 with the same deterministic jitter and random jitter levels as were used in the high frequency sweep.

c Sinusoidal swept frequency: 900 kHz to > 5 MHz.

d Sinusoidal swept frequency: 1.800 kHz to > 5 MHz.

e No value is given for random jitter. For compliance with this standard, the actual random jitter amplitude shall be the value that brings total jitter to the stated value at a probability of 10^{-12} . The additional 0.1 UI of sinusoidal jitter is added to ensure the receiver has sufficient operating margin in the presence of external interference.

f Deterministic jitter: 900 kHz to 750 MHz.

g Deterministic jitter: 1.800 kHz to 1.500 MHz.

h The deterministic and total values in this table apply to jitter after application of a single pole high-pass frequency-weighting function that progressively attenuates jitter at 20 dB/decade below a frequency of ((bit rate) / 1.667).

10.5.2.3.3 Compliant jitter test pattern (CJTPAT)

The CJTPAT within a compliant protocol frame is used for all jitter testing unless otherwise specified. See the SAS Interface Manual for definition of the required pattern on the physical link and information regarding special considerations for scrambling and running disparity.

10.5.2.3.4 Impedance specifications

Table 31 defines impedance requirements.

Table 31: Impedance requirements (Sheet 1 of 2)

Requirement	Units	1.5 Gbps	3.0 Gbps
Time domain reflectometer rise time 20 % to 80 % ^{a,b}	ps	100	50
Media (PCB or cable)			
Differential impedance ^{b,c,d}	ohm	100 ± 10	100 ± 10
Differential impedance imbalance ^{b,c,d,g}	ohm	5	5
Common mode impedance ^{b,c,d}	ohm	32.5 ± 7.5	32.5 ± 7.5
Mated connectors			
Differential impedance ^{b,c,d}	ohm	100 ± 15	100 ± 15
Differential impedance imbalance ^{b,c,d,g}	ohm	5	5
Common mode impedance ^{b,c,d}	ohm	32.5 ± 7.5	32.5 ± 7.5
Receiver termination			
Differential impedance ^{b,e,f}	ohm	100 ± 15	100 ± 15
Differential impedance imbalance ^{b,e,f,g}	ohm	5	5
Receiver termination time constant ^{b,e,f}	ps	150 max	100 max

Table 31: Impedance requirements (Sheet 2 of 2)

Requirement	Units	1.5 Gbps	3.0 Gbps
Common mode impedance ^{b,e}	ohm	20 min/40 max	20 min/40 max
Transmitter source termination			
Differential impedance ^b	ohm	60 min/115 max	60 min/115 max
Differential impedance imbalance ^{b,g}	ohm	5	5
Common mode impedance ^b	ohm	15 min/40 max	15 min/40 max

- a All times indicated for time domain reflectometer measurements are recorded times. Recorded times are twice the transit time of the time domain reflectometer signal.
- b All measurements are made through mated connector pairs.
- c The media impedance measurement identifies the impedance mismatches present in the media when terminated in its characteristic impedance. This measurement excludes mated connectors at both ends of the media, when present, but includes any intermediate connectors or splices. The mated connectors measurement applies only to the mated connector pair at each end, as applicable.
- d Where the media has an electrical length of > 4 ns the procedure detailed in SFF-8410, or an equivalent procedure, shall be used to determine the impedance.
- e The receiver termination impedance specification applies to all receivers in a TxRx connection and covers all time points between the connector nearest the receiver, the receiver, and the transmission line terminator. This measurement shall be made from that connector.
- f At the time point corresponding to the connection of the receiver to the transmission line the input capacitance of the receiver and its connection to the transmission line may cause the measured impedance to fall below the minimum impedances specified in this table. The area of the impedance dip (amplitude as ρ , the reflection coefficient, and duration in time) caused by this capacitance is the receiver termination time constant. The receiver time constant shall not be greater than the values shown in this table. An approximate value for the receiver termination time constant is given by the product of the amplitude of the dip (as ρ) and its width (in ps) measured at the half amplitude point. The amplitude is defined as being the difference in the reflection coefficient between the reflection coefficient at the nominal impedance and the reflection coefficient at the minimum impedance point. The value of the receiver excess input capacitance is given by the following equation:

$$C = \frac{\text{receiver termination time constant}}{(R0 \parallel RR)}$$

where $(R0 \parallel RR)$ is the parallel combination of the transmission line characteristic impedance and termination resistance at the receiver.

- g The difference in measured impedance to ground on the plus and minus terminals on the interconnect, transmitter or receiver, with a differential test signal applied to those terminals.

10.5.2.4 Electrical TxRx connections

TxRx connections may be divided into TxRx connection segments. In a single TxRx connection individual TxRx connection segments may be formed from differing media and materials, including traces on printed wiring boards and optical fibers. This subclause applies only to TxRx connection segments that are formed from electrically conductive media.

Each electrical TxRx connection segment shall comply with the impedance requirements of table 31 for the media from which they are formed. An equalizer network, if present, shall be part of the TxRx connection.

TxRx connections that are composed entirely of electrically conducting media shall be applied only to homogeneous ground applications (e.g., between devices within an enclosure or rack, or between enclosures interconnected by a common ground return or ground plane).

10.5.2.4.1 Transmitter characteristics

The drive are D.C. coupled.

A combination of a zero-length test load and the transmitter compliance transfer function (TCTF) test load methodology is used for the specification of transmitter characteristics. This methodology specifies the transmitter signal at the test points on the required test loads. The transmitter uses the same settings (e.g., pre-emphasis, voltage swing) with both the zero-length test load and the TCTF test load. The signal specifications at IR are met under each of these loading conditions.

The TCTF is the mathematical statement of the transfer function through which the transmitter shall be capable of producing acceptable signals as defined by a receive mask. The transmission magnitude response of the TCTF in dB is given by the following equation for 1.5 Gbps:

$$|S_{21}| = -20 \times \log_{10}(e) \times ((6,5 \times 10^{-6} \times f^{0,5}) + (2,0 \times 10^{-10} \times f) + (3,3 \times 10^{-20} \times f^2)) \text{ dB}$$

for 50 MHz < f < 1.5 GHz, and:

$$|S_{21}| = -5,437 \text{ dB}$$

for 1.5 GHz < f < 5.0 GHz,

where:

- a) f is the signal frequency in hertz.

The transmission magnitude response of the TCTF in dB is given by the following equation for 3.0 Gbps:

$$|S_{21}| = -20 \times \log_{10}(e) \times ((6,5 \times 10^{-6} \times f^{0,5}) + (2,0 \times 10^{-10} \times f) + (3,3 \times 10^{-20} \times f^2)) \text{ dB}$$

for 50 MHz < f < 3.0 GHz, and:

$$|S_{21}| = -10,884 \text{ dB}$$

for 3.0 GHz < f < 5.0 GHz,

where:

- a) f is the signal frequency in hertz.

The TCTF is used to specify the requirements on transmitters that may or may not incorporate pre-emphasis or other forms of compensation. A compliance interconnect is any physical interconnect with loss equal to or greater than that of the TCTF at the above frequencies that also meets the ISI loss requirements shown in figure 20 and figure 21.

Compliance with the TCTF test load requirement is verified by measuring the signal produced by the transmitter through a physical compliance interconnect attached to the transmitter.

Compliance with the zero-length test load requirement verified by measurement made across a load equivalent to the zero-length load shown in figure 19.

For both test load cases, the transmitter delivers the output voltages and timing listed in table 28 at the designated compliance points. The default mask is IR for intra-cabinet TxRx connections. The eye masks are shown in 10.5.2.1.

Figure 18 shows the compliance interconnect test load.




Figure 18. Compliance interconnect test load

Figure 19 shows the zero-length test load.




Figure 19. Zero-length test load

Figure 20 shows an ISI loss example at 3.0 Gbps.

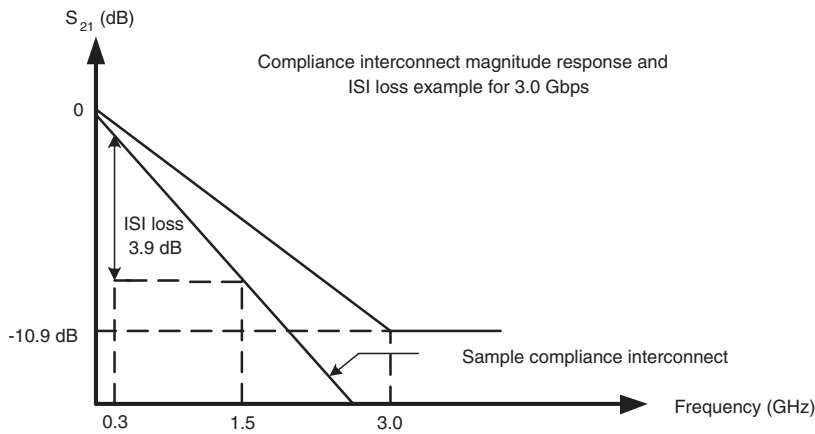


Figure 20. ISI loss example at 3.0 Gbps

Figure 21 shows an ISI loss example at 1.5 Gbps.




Figure 21. ISI loss example at 1.5 Gbps

10.5.2.5 Receiver characteristics

The drive receiver is A.C. coupled. The receive network terminates the TxRx connection by a 100 ohm equivalent impedance as specified in table 31.

The receiver operates within a BER of 10^{-12} when a SAS signal with valid voltage and timing characteristics is delivered to the compliance point from a 100 ohm source. The received SAS signal are considered valid if it meets the voltage and timing limits specified in table 28.

Additionally the receiver operates within the BER objective when the signal at a receiving phy has the additional sinusoidal jitter present that is specified in table 30 and the common mode signal V_{CM} over frequency range F_{CM} as specified in table 26. The jitter tolerance is shown in Figure 17 for all Rx compliance points in a TxRx connection. The figure given assumes that any external interference occurs prior to the point at which the test is applied. When testing the jitter tolerance capability of a receiver, the additional 0.1 UI of sinusoidal jitter may be reduced by an amount proportional to the actual externally induced interference between the application point of the test and the input to the receiving phy. The additional jitter reduces the eye opening in both voltage and time.

Index

Numerics

12 volt
 pins 54
5 volt pins 54

A

abort task set function 39
AC coupling 54
AC power requirements 21
ACA active status 49
ACA active, faulted initiator status 49
acoustics 30
active LED Out signal 54
actuator 8
 assembly design 6
adaptive caching 49
air cleanliness 29
air flow 26, 37
 illustrated 37
air inlet 37
altitude 27
ambient 26
ambient temperature 26, 37
ANSI documents
 SCSI 5
 Serial Attached SCSI 5
asynchronous event notification 49
audible noise 3
auto write and read reallocation
 programmable 7
automatic contingent allegiance 49
average idle current 21, 22
average rotational latency 10

B

Background Media Scan 34
backpanel 53
BMS 34
buffer
 data 7
 space 12
busy status 49
bytes per surface 10
bytes per track 10

C

cache operation 12
cache segments 12
caching write data 12
Canadian Department of Communications 3

capacity
 unformatted 10
case temperature 26
check condition status 49
class B limit 3
clear ACA function 39
clear task set function 39
commands supported 41
condensation 27
condition met/good status 49
connector
 illustrated 53
 requirements 53
continuous vibration 29
cooling 37
CRC
 error 15
Current profiles 23
customer service 20

D

DAR 34
data block size
 modifying the 8
data buffer to/from disc media 11
data heads
 read/write 10
data rate
 internal 10
data transfer rate 11
DC power 53
 requirements 21
defect and error management 32
defects 32
Deferred Auto-Reallocation 34
deferred error handling 49
description 6
dimensions 31
disc rotation speed 10
drive 29
drive characteristics 10
drive failure 15
drive mounting 31, 38
drive select 53
dual port support 40

E

electrical
 description of connector 53
 signal characteristics 54
 specifications 21
electromagnetic compatibility 3
electromagnetic susceptibility 30
EMI requirements 3

- environment 37
- environmental
 - limits 26
 - requirements 14
- environmental control 29
- error
 - management 32
 - rates 14
- errors 32
- European Union Restriction of Hazardous Substances 5

F

- FCC rules and regulations 3
- features 7
 - interface 39
- firmware 7
 - corruption 44
- flawed sector reallocation 7
- Format command execution time 11
- front panel 31
- function
 - complete, code 00 39
 - not supported, code 05 39
 - reject, code 04 39

G

- Good status 49
- gradient 26, 27
- ground shift noise 54
- grounding 38

H

- HDA 37, 38
- head and disc assembly (HDA) 6
- head and disc assembly. *See* HDA
- heads
 - read/write data 10
- heat removal 37
- heat source 37
- host equipment 38
- hot plugging the drive 15
- humidity 27
- humidity limits 26

I

- Idle Read After Write 35
- inquiry data 45
- installation 36
 - guide 5
- interface
 - commands supported 41
 - error rate 14

- errors 15
 - illustrated 50
 - physical 50
 - requirements 39
- interleave
 - minimum 11
- intermediate/condition met/good status 49
- intermediate/good status 49
- internal data rate 10
- internal defects/errors 32
- internal drive characteristics 10
- IRAW 35

J

- jumpers 36

L

- latency
 - average rotational 10, 11
- logical block address 12
- logical block reallocation scheme 7
- logical block size 7, 11
- logical segments 12

M

- maintenance 14
- maximum delayed motor start 21, 22
- maximum start current 21, 22
- mean time between failure. *See* MTBF
- media description 7
- Media Pre-Scan 34
- minimum sector interleave 11
- miscellaneous feature support
 - Adaptive caching 49
 - Asynchronous event notification 49
 - Automatic contingent allegiance 49
 - Deferred error handling 49
 - Parameter rounding 49
 - Queue tagging 49
 - Reporting actual retry count 49
 - Segmented caching 49
 - SMP = 1 in Mode Select command 49
 - Synchronized (locked) spindle operation 49
 - Zero latency read 49
- miscellaneous status support
 - ACA active 49
 - ACA active, faulted initiator 49
 - Busy 49
 - Check condition 49
 - Condition met/good 49
 - Good 49
 - Intermediate/condition met/good 49
 - Intermediate/good 49
 - Reservation conflict 49

- Task set full 49
- miscorrected media data 14
- Mode sense
 - data, table 45, 47
- mounting 38
 - holes 38
 - orientations 36
- mounting configuration 31
- mounting configuration dimensions 31
- MTBF 14, 15, 26

N

- noise
 - audible 3
- noise immunity 23
- non-operating 27, 29
 - temperature 26
- non-operating vibration 29

O

- office environment 29
- operating 27, 29
- option selection 53
- options 9
- orientation 27
- out-of-plane distortion 38

P

- package size 28
- package test specification 5
- packaged 28
- parameter rounding 49
- PCBA 38
- peak bits per inch 10
- peak operating current 21, 22
- peak-to-peak measurements 23
- performance characteristics
 - detailed 10
 - general 11
- performance degradation 27
- performance highlights 7
- physical damage 29
- physical interface 50
- physical specifications 21
- pin descriptions 53
- power 54
 - dissipation 25
 - requirements, AC 21
 - requirements, DC 21
 - sequencing 23
- power distribution 3
- prefetch/multi-segmented cache control 12
- preventive maintenance 14

Q

- queue tagging 49

R

- radio interference regulations 3
- RCD bit 12
- read error rates 14, 32
- read/write data heads 10
- receivers 54
- recommended mounting 28
- Recoverable Errors 14
- recovered media data 14
- reference
 - documents 5
- relative humidity 27
- reliability 8
 - specifications 14
- reliability and service 15
- repair and return information 20
- reporting actual retry count 49
- reservation conflict status 49
- resonance 27
- return information 20
- RoHS 5
- rotation speed 10

S

- safety 3
- SAS
 - interface 53
 - physical interface 50
 - task management functions 39
- SAS documents 5
- SAS Interface Manual 3, 5
- SCSI interface
 - commands supported 41
- seek error
 - defined 15
 - rate 14
- seek performance characteristics 10
- seek time
 - average typical 10
 - full stroke typical 10
 - single track typical 10
- segmented caching 49
- Self-Monitoring Analysis and Reporting Technology 8, 16
- Serial Attached SCSI (SAS) Interface Manual 2
- shielding 3
- shipping 20
- shipping container 26
- shock 27
 - and vibration 27
- shock mount 38

- signal
 - characteristics 54
- single-unit shipping pack kit 9
- SMART 8, 16
- SMP = 1 in Mode Select command 49
- spindle brake 7
- standards 3
- start/stop time 11
- support services 1
- surface stiffness
 - allowable for non-flat surface 38
- switches 36
- synchronized spindle
 - operation 49
- system chassis 38

T

- task management functions 39
 - Abort task set 39
 - Clear ACA 39
 - Clear task set 39
 - terminate task 39
- task management response codes 39
 - Function complete 00 39
 - Function not supported 05 39
 - Function reject 04 39
- task set full status 49
- technical support services 1
- temperature 26, 37
 - ambient 26
 - case 26
 - gradient 26
 - limits 26
 - non-operating 26
 - regulation 3
 - See also* cooling
- terminate task function 39
- terminators 36
- tracks per inch 10
- tracks per surface 10
- transmitters 54
- transporting the drive 20

U

- unformatted 8
- Unrecoverable Errors 14
- unrecovered media data 14

V

- vibration 27, 29

W

- warranty 20

Z

- zero latency read 49
- zone bit recording (ZBR) 7



Seagate Technology LLC

AMERICAS Seagate Technology LLC 10200 South De Anza Boulevard, Cupertino, California 95014, United States, 408-658-1000

ASIA/PACIFIC Seagate Singapore International Headquarters Pte. Ltd. 7000 Ang Mo Kio Avenue 5, Singapore 569877, 65-6485-3888

EUROPE, MIDDLE EAST AND AFRICA Seagate Technology SAS 16-18 rue du Dôme, 92100 Boulogne-Billancourt, France, 33 1-4186 10 00

Publication Number: 100384760, E

April 2012