

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

Product Manual

Constellation® ES.2 SAS

Standard Model

ST33000650SS

Self-Encrypting Drive Model

ST33000651SS

**SED FIPS 140-2 Model
Review Pending**

ST33000652SS

100628615
Rev. D
April 2011

Revision history

Revision	Date	Sheets affected or comments
Rev. A	11/29/10	Initial release
Rev. B	02/23/11	6-8, 10, 22, 27, 32 42 & 46-48.
Rev. C	03/23/11	8-9, 13-14, 35-36 & 40.
Rev. D	04/28/11	fc, 2, 23 & 47.

© 2011, Seagate Technology LLC All rights reserved.
Publication number: 100628615, Rev. D April 2011

Seagate, Seagate Technology and the Wave logo are registered trademarks of Seagate Technology LLC in the United States and/or other countries. Constellation ES.2 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. The FIPS logo is a certification mark of NIST, which does not imply product endorsement by NIST, the U.S., or Canadian governments. 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.

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. 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
3.1.4	China Restriction of Hazardous Substances (RoHS) Directive	5
3.2	Reference documents	6
4.0	General description	7
4.1	Standard features	8
4.2	Media description	8
4.3	Performance	8
4.4	Reliability	9
4.5	Formatted capacities	9
4.6	Programmable drive capacity	9
4.7	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	13
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	PowerChoice™ power management	21
7.1.1	PowerChoice reporting methods	22
7.2	AC power requirements	22
7.3	DC power requirements	22
7.3.1	Conducted noise immunity	24
7.3.2	Power sequencing	24
7.3.3	Current profiles	25

7.4	Power dissipation	26
7.5	Environmental limits	28
7.5.1	Temperature	28
7.5.2	Relative humidity	28
7.5.3	Effective altitude (sea level)	29
7.5.4	Shock and vibration	29
7.5.5	Acoustics	31
7.5.6	Air cleanliness	31
7.5.7	Corrosive environment	31
7.5.8	Electromagnetic susceptibility	32
7.6	Mechanical specifications	33
8.0	About FIPS	34
9.0	About self-encrypting drives	35
9.1	Data encryption	35
9.2	Controlled access	35
9.2.1	Admin SP	35
9.2.2	Locking SP	35
9.2.3	Default password	36
9.3	Random number generator (RNG)	36
9.4	Drive locking	36
9.5	Data bands	36
9.6	Cryptographic erase	36
9.7	Authenticated firmware download	37
9.8	Power requirements	37
9.9	Supported commands	37
9.10	RevertSP	37
10.0	Defect and error management	38
10.1	Drive internal defects/errors	38
10.2	Drive error recovery procedures	38
10.3	SAS system errors	39
10.4	Background Media Scan	39
10.5	Media Pre-Scan	40
10.6	Deferred Auto-Reallocation	40
10.7	Idle Read After Write	40
10.8	Protection Information (PI)	41
10.8.1	Levels of PI	41
10.8.2	Setting and determining the current Type Level	41
10.8.3	Identifying a Protection Information drive	41
11.0	Installation	42
11.1	Drive orientation	42
11.2	Cooling	43
11.3	Drive mounting	44
11.4	Grounding	44
12.0	Interface requirements	45
12.1	SAS features	45
12.1.1	task management functions	45
12.1.2	task management responses	45
12.2	Dual port support	46
12.3	SCSI commands supported	47
12.3.1	Inquiry data	51
12.3.2	Mode Sense data	51
12.4	Miscellaneous operating features and conditions	54

12.4.1	SAS physical interface	55
12.4.2	Physical characteristics	58
12.4.3	Connector requirements.	58
12.4.4	Electrical description	58
12.4.5	Pin descriptions	58
12.4.6	SAS transmitters and receivers	59
12.4.7	Power.	59
12.5	Signal characteristics.	59
12.5.1	Ready LED Out	59
12.5.2	Differential signals	60
12.6	SAS-2 Specification Compliance.	60
12.7	Additional information	60

List of Figures

Figure 1.	3TB model current profiles	25
Figure 2.	3TB models (3Gb) DC current and power vs. input/output operations per second.	26
Figure 3.	3TB models (6Gb) DC current and power vs. input/output operations per second.	27
Figure 4.	Location of the HDA temperature check point	28
Figure 5.	Recommended mounting	30
Figure 6.	Mounting configuration dimensions	33
Figure 7.	Example of FIPS tamper evidence labels.	34
Figure 8.	Physical interface	42
Figure 9.	Air flow	43
Figure 10.	Physical interface	55
Figure 11.	SAS device plug dimensions	56
Figure 12.	SAS device plug dimensions (detail)	57
Figure 13.	SAS transmitters and receivers	59

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, Constellation® ES.2 SAS (Serial Attached SCSI) disk drives.

Constellation ES.2 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. The Self-Encrypting Drive Reference Manual, part number 100515636, describes the interface, general operation, and security features available on Self-Encrypting Drive models.

Note. Previous generations of Seagate Self-Encrypting Drive models were called Full Disk Encryption (FDE) models before a differentiation between drive-based encryption and other forms of encryption was necessary.

Product data communicated in this manual is specific only to the model numbers listed in this manual. The data listed in this manual may not be predictive of future generation specifications or requirements. If you are designing a system which will use one of the models listed or future generation products and need further assistance, please contact your Field Applications Engineer (FAE) or our global support services group as shown in Section 1.0.

Unless otherwise stated, the information in this manual applies to standard and Self-Encrypting Drive models.

Model Number	Self-Encrypting Drive (SED)	FIPS 140-2 Level 2 (Review Pending)
ST33000650SS	No	No
ST33000651SS	Yes	No
ST33000652SS	Yes	Yes

For more information on FIPS 140-2 Level 2 certification see Section 8.0 on page 34.

For product certification status visit - <http://csrc.nist.gov/groups/STM/cmvp/documents/140-1/1401vend.htm>.

3.0 Applicable standards and reference documentation

The drives documented in this manual have been developed as system peripherals to the highest standards of design and construction. The drives depends on 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 11.3.

3.1 Standards

The Constellation ES.2 family complies with Seagate standards as noted in the appropriate sections of this manual and the Seagate *SAS Interface Manual*, part number 100293071.

The drives are recognized in accordance with UL 60950-1 as tested by UL, CSA 60950-1 as tested by CSA, and EN60950-1 as tested by TUV.

The security features of Self-Encrypting Drive models are based on the “TCG Storage Architecture Core Specification” and the “TCG Storage Workgroup Security Subsystem Class: Enterprise_A” specification with additional vendor-unique features as noted in this product manual.

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 Table 2, 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
- Keyboard
- Monitor/display
- Printer
- 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 2004/108/EC as put into place on 20 July 2007.

Australian C-Tick

If this model has the C-Tick Marking it complies with the Australia/New Zealand Standard AS/NZ CISPR22 and meets the Electromagnetic Compatibility (EMC) Framework requirements of Australia's Spectrum Management Agency (SMA).

Korean KCC

If these drives have the Korean Communications Commission (KCC) logo, they comply with paragraph 1 of Article 11 of the Electromagnetic Compatibility control Regulation and meet the Electromagnetic Compatibility (EMC) Framework requirements of the Radio Research Laboratory (RRL) Communications Commission, Republic of Korea.

These drives have been tested and comply with the Electromagnetic Interference/Electromagnetic Susceptibility (EMI/EMS) for Class B products. Drives are tested in a representative, end-user system by a Korean-recognized lab.

- Family name: Constellation ES SAS
- Certificate number: STX-ST33000650SS
- Manufacturing date: July 2, 2010 (Date of Certification)
- Manufacturer/nationality: USA, Singapore and China

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, Cadmium, Mercury, Hexavalent Chromium, PBB and PBDE, in electronic products, effective July 2006. This drive is manufactured with components and materials that comply with the RoHS Directive.

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.

3.1.4 China Restriction of Hazardous Substances (RoHS) Directive 中国限制危险物品的指令

This product has an Environmental Protection Use Period (EPUP) of 20 years. The following table contains information mandated by China's "Marking Requirements for Control of Pollution Caused by Electronic Information Products" Standard.



该产品具有20年的环境保护使用周期（EPUP）。下表包含了中国“电子产品所导致的污染的控制的记号要求”所指定的信息。

Name of Parts 部件名称	Toxic or Hazardous Substances or Elements 有毒有害物质或元素					
	Lead 铅 (Pb)	Mercury 汞 (Hg)	Cadmium 镉 (Cd)	Hexavalent Chromium 六价铬 (Cr6+)	Polybrominated Diphenyl 多溴联苯 (PBB)	Polybrominated Diphenyl Ether 多溴二苯醚 (PBDE)
PCBA	X	O	O	O	O	O
HDA	X	O	O	O	O	O

"O" indicates the hazardous and toxic substance content of the part (at the homogenous material level) is lower than the threshold defined by the China RoHS MCV Standard.

“O”表示该部件（于同类物品程度上）所含的危险和有毒物质低于中国RoHS MCV标准所定义的门槛值。

"X" indicates the hazardous and toxic substance content of the part (at the homogenous material level) is over the threshold defined by the China RoHS MCV Standard.

“X”表示该部件（于同类物品程度上）所含的危险和有毒物质超出中国RoHS MCV标准所定义的门槛值。

3.2 Reference documents

SAS Interface Manual

Seagate part number: 100293071

SCSI Commands Reference Manual

Seagate part number: 100293068

Self-Encrypting Drives Reference Manual

Seagate part number: 100515636

ANSI SAS Documents

SFF-8323	3.5" Drive Form Factor with Serial Connector
SFF-8460	HSS Backplane Design Guidelines
SFF-8470	Multi Lane Copper Connector
SFF-8482	SAS Plug Connector
ANSI INCITS.xxx	Serial Attached SCSI (SAS-2) Standard (T10/1562-D)
ISO/IEC 14776-xxx	SCSI Architecture Model-3 (SAM-4) Standard (T10/1561-D)
ISO/IEC 14776-xxx	SCSI Primary Commands-3 (SPC-4) Standard (T10/1416-D)
ISO/IEC 14776-xxx	SCSI Block Commands-2 (SBC-3) Standard (T10/1417-D)

ANSI Small Computer System Interface (SCSI) Documents

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

Trusted Computing Group (TCG) Documents (apply to Self-Encrypting Drive models only)

TCG Storage Architecture Core Specification, Rev. 1.0

TCG Storage Security Subsystem Class Enterprise Specification, Rev. 1.0

Specification for Acoustic Test Requirement and Procedures

Seagate part number: 30553-001

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

4.0 General description

Constellation ES.2 drives 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.

Constellation ES.2 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. Constellation ES.2 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 6Gb/s serial data transfer rate.

The Self-Encrypting Drive models indicated on the cover of this product manual have provisions for “Security of Data at Rest” based on the standards defined by the Trusted Computing Group (see www.trustedcomputinggroup.org).

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.

Constellation ES.2 drives use a dedicated load/unload zone at the outermost 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 ramp load/unload when power is removed from the drive.

An automatic shipping lock prevents potential damage to the heads and discs 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.

Constellation ES.2 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.

Note. Seagate recommends validating your configuration with the selected HBA/RAID controller manufacturer to ensure full 3TB capacity capabilities.

4.1 Standard features

Constellation ES.2 drives have the following standard features:

- Perpendicular recording technology
- 1.5 / 3.0 / 6.0 Gb 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 3.5-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 400 bits
- No preventive maintenance or adjustments required
- Embedded servo design
- Dedicated head load/unload zone
- Self diagnostics performed when power is applied to the drive
- Vertical, horizontal, or top down mounting
- 64 MB data buffer (see Section 5.5).
- Drive Self Test (DST)
- Background Media Scan (BMS)
- Idle Read After Write (IRAW)
- Power Save

Constellation® ES.2 SAS Self-Encrypting Drive models have the following additional features:

- Automatic data encryption/decryption
- Controlled access
- Random number generator
- Drive locking
- 16 independent data bands
- Cryptographic erase of user data for a drive that will be repurposed or scrapped
- Authenticated firmware download

4.2 Media description

The media used on the drive has a 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
- 600MB/s maximum instantaneous data transfers.
- 7200 RPM spindle. Average latency = 4.16ms
- Background processing of queue
- Supports start and stop commands (spindle stops spinning)
- Adaptive seek velocity; improved seek performance

Note. There is no significant performance difference between Self-Encrypting Drive and standard (non-Self-Encrypting Drive) models.

4.4 Reliability

- Annualized Failure Rate (AFR) of 0.73%
- Mean time between failures (MTBF) of 1,200,000 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	ST33000650SS ST33000651SS ST33000652SS	
	Decimal	Hex
512	5,860,533,168	15D50A3B0h
520	5,736,538,480	155ECA170
528	5,578,747,784	14C84EF88

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 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 *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 Constellation ES.2 drives.

5.1 Internal drive characteristics

	ST33000650SS	ST33000651SS	ST33000652SS	
Drive capacity	3			TB (formatted, rounded off value)
Read/write data heads	10			
Bytes per track	1,419,776			Bytes (average, rounded off values)
Bytes per surface	300,000			MB (unformatted, rounded off value)
Tracks per surface (total)	284,399			Tracks (user accessible)
Tracks per inch	270,000			TPI (average)
Peak bits per inch	1,638,000			BPI
Areal density	444			Gb/in ²
Internal data rate	68.7 - 155			MB/s (variable with zone)
disk rotation speed	7200			rpm
Avg rotational latency	4.16			ms

5.2 Seek performance characteristics

See Section 12.4.1, "SAS physical interface" on page 55 and the *SAS Interface Manual* (part number 100293071) for additional timing details.

5.2.1 Access time

		Not including controller overhead ^{1, 2} (ms)		Including controller overhead ^{1, 2} (ms)	
		Read	Write	Read	Write
Average	Typical ^{3,4}	8.3	9.3	8.5	9.5
Single track	Typical ^{3,4}	0.5	0.5	0.7	0.7
Full stroke	Typical ^{3,4}	15.5	16.2	15.7	16.4

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.

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

	ST33000650SS
	ST33000651SS
	ST33000652SS
Maximum (with verify)	827
Maximum (without verify)	402

Note. There is approximately a 1.5 increase in time to format a SED drive versus a non-SED drive of the same capacity.

5.2.3 General performance characteristics

Minimum sector interleave	1 to 1
Data buffer to/from disk media (one 512-byte logical block)*	113 to 239 MB/s
Sustained transfer rate	68.7 to 155 MB/s
SAS Interface maximum instantaneous transfer rate	600MB/s* per port (dual port = 1200MB/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	4.16ms

*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 30 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 30 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 8.

5.5 Cache operation

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

Of the 64MB physical buffer space in the drive, approximately 30,000 kbytes 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 8 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 ¹⁵ 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):	1,200,000 hours
Annualized Failure Rate (AFR):	0.73%
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 7.3, "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 Constellation ES.2 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 11.2 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 0.73% (MTBF of 1,200,000 hours) when operated in an environment that ensures the HDA case temperatures do not exceed the values specified in Section 7.5. Operation at case temperatures outside the specifications in Section 7.5 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 7.5 are not exceeded. Temperatures outside the specifications in Section 7.5 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 initiates 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 one hour.

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 one hour subject to an idle period on the drive interfaces. The process of measuring off-line attribute data and saving data to the disk is interruptible. 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	210 ms	75 ms

Reporting control

Reporting is controlled by the MRIE bits in the Informational Exceptions Control mode page (1Ch). An example, if the MRIE is set to one, the firmware will issue to the host an 01-5D00 sense code. The FRU field contains the type of predictive failure that occurred. 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

Constellation ES.2 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 1) 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 1: 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 whenever 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 PowerChoice™ power management

Drives using the load/unload architecture provide programmable power management to tailor systems for performance and greater energy efficiency.

The table below lists the supported PowerChoice modes. The further you go down in the table, the more power savings you get. For example, Idle_B mode results in greater power savings than Idle_A mode. Standby_Z mode results in the greatest power savings.

PowerChoice modes

Mode	Description
Idle_A	Reduced electronics
Idle_B	Heads unloaded. Disks spinning at full RPM
Idle_C	Heads unloaded. Disks spinning at reduced RPM
Standby_Y	Heads unloaded. Disks spinning at reduced RPM. Recovery requires the NOTIFY (Enable Spinup) command.
Standby_Z	Heads unloaded. Motor stopped (disks not spinning) Recovery requires the NOTIFY (Enable Spinup) command.

PowerChoice™ can be invoked using one of these two methods:

- Power Condition mode page method—Enable and initialize the idle condition timers and/or the standby condition timers. The timer values are based on the values set in the Power Condition mode page.
- START STOP UNIT command method—Use the START STOP UNIT command (OPERATION CODE 1Bh). This allows the host to directly transition the drive to any supported PowerChoice mode.

If both the Power Condition mode page and START STOP UNIT command methods are used, the START STOP UNIT command request takes precedence over the Power Condition mode page power control and may disable the idle condition and standby condition timers. The REQUEST SENSE command reports the current PowerChoice state if active and also the method by which the drive entered the PowerChoice state.

When the drive receives a command, all power condition timers are suspended if they were enabled via the Power Condition mode page. Once all outstanding commands are processed, the power condition timers are reinitialized to the values defined in the Power Condition mode page.

7.1.1 PowerChoice reporting methods

PowerChoice™ provides these reporting methods for tracking purposes:

Request Sense command reports

- Current power condition
- Method of entry

Note. Processing the Request Sense command does not impact the drive's power save state.

Mode Sense command reports (mode page 0x1A)

- Idle conditions enabled / disabled
- Idle condition timer values (100ms increments) (default, saved, current, changeable)

Power Condition Vital Product Data (VPD) Page (VPD page 0x8A)

- Supported power conditions
- Typical recovery time from power conditions (1ms increments)

Start/Stop Cycle Counter Log Page reports (log page 0x0E)

- Specified and accumulated Start/Stops and Load/Unload cycles

Power Condition Transitions Log Page reports (log page 0x1A, subpage 0x00)

- Accumulated transitions to Active, Idle_A, Idle_B, Idle_C, Standby_Y, Standby_Z

7.2 AC power requirements

None.

7.3 DC power requirements

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

The standard drive models and the SED drive models have identical hardware, however the security and encryption portion of the drive controller ASIC is enabled and functional in the SED models. This represents a small additional drain on the 5V supply of about 30mA and a commensurate increase of about 150mW in power consumption. There is no additional drain on the 12V supply.

Table 2: 3000GB drive (Standard & SED model) DC power requirements

	Notes	3.0Gb mode		6.0Gb mode	
		(Amps)	(Amps)	(Amps)	(Amps)
Voltage		+5V	+12V [4]	+5V	+12V [4]
Regulation	[5]	± 5% [2]		± 5% [2]	
Avg idle current $DC\bar{X}$	[1] [7]	0.28	0.50	0.28	0.50
Advanced idle current					
Idle_A		0.28	0.50	0.28	0.50
Idle_B		0.26	0.41	0.26	0.41
Idle_C/ Standby_Y		0.26	0.24	0.26	0.24
Standby_Z		0.25	0.01	0.25	0.01
Maximum starting current					
(peak DC) DC	3 σ [5]	0.60	1.71	0.60	1.70
(peak AC) AC	3 σ [5]	0.82	2.29	0.88	2.23
Delayed motor start (max) DC	3 σ [1] [6]	0.39	0.01	0.39	0.01
Peak operating current (random read):					
Typical $DC\bar{X}$	[1] [6]	0.48	0.74	0.49	0.74
Maximum DC	3 σ [1]	0.49	0.75	0.50	0.76
Maximum (peak) DC	3 σ	1.14	1.88	1.16	1.86
Peak operating current (random write)					
Typical $DC\bar{X}$		0.58	0.60	0.59	0.61
Maximum DC	3 σ	0.60	0.62	0.60	0.62
Maximum (peak) DC	3 σ	1.20	1.86	1.18	1.88
Peak operating current (sequential read)					
Typical $DC\bar{X}$		0.79	0.48	0.79	0.49
Maximum DC	3 σ	0.83	0.49	0.84	0.50
Maximum (peak) DC	3 σ	1.16	0.88	1.12	0.90
Peak operating current (sequential write)					
Typical $DC\bar{X}$		0.94	0.48	0.94	0.48
Maximum DC	3 σ	0.98	0.49	0.98	0.50
Maximum (peak) DC	3 σ	1.20	0.88	1.20	0.88

[1] Measured with average reading DC ammeter.

[2] Instantaneous +12V current peaks will exceed these values.

[3] Power supply at nominal voltage. N (number of drives tested) = 6, 35 Degrees C ambient.

[4] For +12 V, a –10% tolerance is allowed during initial spindle start but must return to ± 5% before reaching 7200 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.

- [5] See +12V current profile in Figure 1.
- [6] This condition occurs after OOB and Speed Negotiation completes but before the drive has received the Notify Spinup primitive.
- [7] See paragraph 7.3.1, "Conducted noise immunity." Specified voltage tolerance includes ripple, noise, and transient response.
- [8] Operating condition is defined as random 8 block reads.
- [9] 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.

7.3.1 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 100 Hz 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.

7.3.2 Power sequencing

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

7.3.3 Current profiles

The +12V (top) and +5V (bottom) current profiles for the Constellation ES drives are shown below.




Figure 1. 3TB model current profiles

Note: All times and currents are typical. See Table 2 for maximum current requirements.

7.4 Power dissipation

3TB models in 3Gb operation

Please refer to Table 2 for power dissipation numbers.

To obtain operating power for typical random read operations, refer to the following I/O rate curve (see Figure 2). 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 2. 3TB models (3Gb) DC current and power vs. input/output operations per second

Note. For power details about SED vs. non-SED drive, please refer to section 7.8.

3TB models in 6Gb operation

Please refer to Table 2 for power dissipation numbers.

To obtain operating power for typical random read operations, refer to the following I/O rate curve (see Figure 3.). 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 3. 3TB models (6Gb) DC current and power vs. input/output operations per second

For power details about SED vs. non-SED drive, please refer to section 7.8.

7.5 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).

7.5.1 Temperature

a. Operating

The drive meets the operating specifications over a 41°F to 140°F (5°C to 60°C) drive case temperature range with a maximum temperature gradient of 36°F (20°C) per hour.

The maximum allowable drive case temperature is 60°C. See Figure 4 for HDA case temperature measurement location.

The MTBF specification for the drive assumes the operating environment is designed to maintain nominal case temperature. The rated MTBF is based upon a sustained case temperature of 104°F (40°C). Occasional excursions in operating temperature between the rated MTBF temperature and the maximum drive operating case temperature may occur without impact to the rated MTBF temperature. However, continual or sustained operation at case temperatures beyond the rated MTBF temperature will degrade the drive MTBF and reduce product reliability.

Air flow may be required to achieve consistent nominal case temperature values (see Section 11.2). To confirm that the required cooling is provided for the electronics and HDA, place the drive in its final mechanical configuration, and perform random write/read operations. After the temperatures stabilize, measure the case temperature of the drive.

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 4. Location of the HDA temperature check point

Note. Image is for reference only, may not represent actual drive.

7.5.2 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.

7.5.3 Effective altitude (sea level)

- a. Operating
–1000 to +10,000 feet (–304.8 to +3,048 meters)
- b. Non-operating
–1000 to +40,000 feet (–304.8 to +12,192 meters)

7.5.4 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 5, and in accordance with the restrictions of Section 11.3.

7.5.4.1 Shock

- a. Operating—normal

The drive, as installed for normal operation, shall operate error free while subjected to intermittent shock not exceeding 70 Gs (read) and 40 Gs (write) at a maximum duration of 2ms (half sinewave). Shock may be applied in the X, Y, or Z axis. Shock is not to be repeated more than once every 2 seconds.

- 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 once every 2 seconds.

- 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 80 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 5. Recommended mounting

Note. Image is for reference only, may not represent actual drive.

7.5.4.2 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

5 - 22 Hz	0.25 Gs, limited displacement
22 - 350 Hz	0.5 Gs
350 - 500 Hz	0.25 Gs

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

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

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

5-500 Hz @ 0.75 G (X, Y, or Z axis)

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	2 Gs (0 to peak, linear, swept sine, 0.5 octave/min)
22 - 350 Hz	5 Gs (0 to peak, linear, swept sine, 0.5 octave/min)
350 - 500 Hz	2 Gs (0 to peak, linear, swept sine, 0.5 octave/min)

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

7.5.5 Acoustics

Sound power during idle mode shall be 2.8 bels typical when measured to ISO 7779 specification. Sound power while operating shall be 3.0 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.

7.5.6 Air cleanliness

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

7.5.7 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.

7.5.8 Electromagnetic susceptibility

See Section 3.1.1.1.

7.6 Mechanical specifications

Refer to Figure 6 for detailed mounting configuration dimensions. See Section 11.3, "Drive mounting."

Weight: 1.543 lb 700 g

Note. These dimensions conform to the Small Form Factor Standard documented in SFF-8301 and SFF-8323, found at www.sffcommittee.org.




Figure 6. Mounting configuration dimensions

8.0 About FIPS

The Federal Information Processing Standard (FIPS) Publication 140-2, FIPS PUB 140-2, is a U.S. government computer security standard used to accredit cryptographic modules. It is titled “Security Requirements for Cryptographic Modules”. The initial publication was on May 25, 2001 and was last updated December 3, 2002.

Purpose

The National Institute of Standards and Technology (NIST) issued the FIPS 140 Publication Series to coordinate the requirements and standards for cryptography modules that include both hardware and software components.

Federal Information Processing Standard (FIPS) 140-2 Level 2 Certification requires drives to go through government agencies certifications to add requirements for physical tamper-evidence and role-based authentication.

Level 2 security

Level 2 improves upon the physical security mechanisms of a Level 1 (lowest level of security) cryptographic module by requiring features that show evidence of tampering, including tamper-evident coatings or seals that must be broken to attain physical access to the plaintext cryptographic keys and critical security parameters (CSPs) within the module, or pick-resistant locks on covers or doors to protect against unauthorized physical access.



Figure 7. Example of FIPS tamper evidence labels.

Note. Does not represent actual drive.

9.0 About self-encrypting drives

Self-encrypting drives (SEDs) offer encryption and security services for the protection of stored data, commonly known as “protection of data at rest.” These drives are compliant with the Trusted Computing Group (TCG) Enterprise Storage Specifications as detailed in Section 3.2.

The Trusted Computing Group (TCG) is an organization sponsored and operated by companies in the computer, storage and digital communications industry. Seagate’s SED models comply with the standards published by the TCG.

To use the security features in the drive, the host must be capable of constructing and issuing the following two SCSI commands:

- Security Protocol Out
- Security Protocol In

These commands are used to convey the TCG protocol to and from the drive in their command payloads.

9.1 Data encryption

Encrypting drives use one inline encryption engine for each port, employing AES-128 data encryption in Cipher Block Chaining (CBC) mode to encrypt all data prior to being written on the media and to decrypt all data as it is read from the media. The encryption engines are always in operation, cannot be disabled, and do not detract in any way from the performance of the drive.

The 32-byte Data Encryption Key (DEK) is a random number which is generated by the drive, never leaves the drive, and is inaccessible to the host system. The DEK is itself encrypted when it is stored on the media and when it is in volatile temporary storage (DRAM) external to the encryption engine. A unique data encryption key is used for each of the drive's possible 16 data bands (see Section 9.5).

9.2 Controlled access

The drive has two security partitions (SPs) called the "Admin SP" and the "Locking SP." These act as gatekeepers to the drive security services. Security-related commands will not be accepted unless they also supply the correct credentials to prove the requester is authorized to perform the command.

9.2.1 Admin SP

The Admin SP allows the drive's owner to enable or disable firmware download operations (see Section 9.4). Access to the Admin SP is available using the SID (Secure ID) password or the MSID (Makers Secure ID) password.

9.2.2 Locking SP

The Locking SP controls read/write access to the media and the cryptographic erase feature. Access to the Locking SP is available using the BandMasterX or EraseMaster passwords. Since the drive owner can define up to 16 data bands on the drive, each data band has its own password called BandMasterX where X is the number of the data band (0 through 15).

9.2.3 Default password

When the drive is shipped from the factory, all passwords are set to the value of MSID. This 32-byte random value is printed on the drive label and it can be read by the host electronically over the I/O. After receipt of the drive, it is the responsibility of the owner to use the default MSID password as the authority to change all other passwords to unique owner-specified values.

9.3 Random number generator (RNG)

The drive has a 32-byte hardware RNG that it uses to derive encryption keys or, if requested to do so, to provide random numbers to the host for system use, including using these numbers as Authentication Keys (passwords) for the drive's Admin and Locking SPs.

9.4 Drive locking

In addition to changing the passwords, as described in Section 9.2.3, the owner should also set the data access controls for the individual bands.

The variable "LockOnReset" should be set to "PowerCycle" to ensure that the data bands will be locked if power is lost. This scenario occurs if the drive is removed from its cabinet. The drive will not honor any data read or write requests until the bands have been unlocked. This prevents the user data from being accessed without the appropriate credentials when the drive has been removed from its cabinet and installed in another system.

When the drive is shipped from the factory, the firmware download port is unlocked.

9.5 Data bands

When shipped from the factory, the drive is configured with a single data band called Band 0 (also known as the Global Data Band) which comprises LBA 0 through LBA max. The host may allocate Band1 by specifying a start LBA and an LBA range. The real estate for this band is taken from the Global Band. An additional 14 Data Bands may be defined in a similar way (Band2 through Band15) but before these bands can be allocated LBA space, they must first be individually enabled using the EraseMaster password.

Data bands cannot overlap but they can be sequential with one band ending at LBA (x) and the next beginning at LBA (x+1).

Each data band has its own drive-generated encryption key and its own user-supplied password. The host may change the Encryption Key (see Section 9.6) or the password when required. The bands should be aligned to 4K LBA boundaries.

9.6 Cryptographic erase

A significant feature of SEDs is the ability to perform a cryptographic erase. This involves the host telling the drive to change the data encryption key for a particular band. Once changed, the data is no longer recoverable since it was written with one key and will be read using a different key. Since the drive overwrites the old key with the new one, and keeps no history of key changes, the user data can never be recovered. This is tantamount to an instantaneous data erase and is very useful if the drive is to be scrapped or repositioned.

9.7 Authenticated firmware download

In addition to providing a locking mechanism to prevent unwanted firmware download attempts, the drive also only accepts download files which have been cryptographically signed by the appropriate Seagate Design Center.

Three conditions must be met before the drive will allow the download operation:

1. The download must be an SED file. A standard (base) drive (non-SED) file will be rejected.
2. The download file must be signed and authenticated.
3. As with a non-SED drive, the download file must pass the acceptance criteria for the drive. For example it must be applicable to the correct drive model, and have compatible revision and customer status.

9.8 Power requirements

The standard drive models and the SED drive models have identical hardware, however the security and encryption portion of the drive controller ASIC is enabled and functional in the SED models. This represents a small additional drain on the 5V supply of about 30mA and a commensurate increase of about 150mW in power consumption. There is no additional drain on the 12V supply. See the tables in Section 7.3 for power requirements on the standard (non-SED) drive models.

9.9 Supported commands

The SED models support the following two commands in addition to the commands supported by the standard (non-SED) models as listed in Table 6:

- Security Protocol Out (B5h)
- Security Protocol In (A2h)

9.10 RevertSP

The SED models will support RevertSP feature where it erases all data in all bands on the device and returns the contents of all SPs (Security Providers) on the device to their Original Factory State.

10.0 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 10.4)
- Media Pre-Scan (see Section 10.5)
- Deferred Auto-Reallocation (see Section 10.6)
- Idle Read After Write (see Section 10.7)

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

10.1 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*.

10.2 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 consist of 12 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 3 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 3: Read and write retry count maximum recovery times

Read retry count*	Maximum recovery time per LBA (cumulative, msec)	Write retry count	Maximum recovery time per LBA (cumulative, msec)
		0	23.94
1	124.32	1	35.91
5	621.62	2	55.86
10	1243.23	3	67.83
15	1864.85	4	119.79
20 (default)	2486.47	5 (default)	147.72

* For read retry count, every tick ~ 5% of total error recovery. Valid range setting is 1-20.

e.g. 1 ~ 5%

5 ~ 25%

20 ~ 100%

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 and if the read retry count is set to 5, this means ~ 25% of error recovery will be executed which consumes 621.62 ms (please refer to the table above). If the limit is reached and a LBA has not yet been recovered (i.e. requires retries beyond 621.62 ms), the command will end with Check Condition status report and unrecoverable read error will be reported.

10.3 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.

10.4 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.

10.5 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.

10.6 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.

10.7 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.

10.8 Protection Information (PI)

Protection Information is intended as a standardized approach to system level LRC traditionally provided by systems using 520 byte formatted LBAs. Drives formatted with PI information provide the same, common LBA count (i.e. same capacity point) as non-PI formatted drives. Sequential performance of a PI drive will be reduced by approximately 1.56% due to the extra overhead of PI being transferred from the media that is not calculated as part of the data transferred to the host. To determine the full transfer rate of a PI drive, transfers should be calculated by adding the 8 extra bytes of PI to the transferred LBA length, i.e. $512 + 8 = 520$. PI formatted drives are physically formatted to 520 byte sectors that store 512 bytes of customer data with 8 bytes of Protection Information appended to it. The advantage of PI is that the Protection Information bits can be managed at the HBA and HBA driver level. Allowing a system that typically does not support 520 LBA formats to integrate this level of protection.

Protection Information is valid with any supported LBA size. 512 LBA size is used here as common example.

10.8.1 Levels of PI

There are 4 types of Protection Information.

Type 0 - Describes a drive that is not formatted with PI information bytes. This allows for legacy support in non-PI systems.

Type 1 - Provides support of PI protection using 10 and 16 byte commands. The RDPROTECT and WRTPROTECT bits allow for checking control through the CDB. Eight bytes of Protection Information are transmitted at LBA boundaries across the interface if RDPROTECT and WRTPROTECT bits are nonzero values. Type 1 does not allow the use of 32 byte commands.

Type 2 - Provides checking control and additional expected fields within the 32 byte CDBs. Eight bytes of Protection Information are transmitted at LBA boundaries across the interface if RDPROTECT and WRTPROTECT bits are nonzero values. Type 2 does allow the use of 10 and 16 byte commands with zero values in the RDPROTECT and WRTPROTECT fields. The drive will generate 8 bytes (e.g.0xFFFF) 8 bytes of Protection Information to be stored on the media, but the 8 bytes will not be transferred to the host during a read command.

Type 3 - Seagate products do not support Type 3.

10.8.2 Setting and determining the current Type Level

A drive is initialized to a type of PI by using the format command on a PI capable drive. Once a drive is formatted to a PI Type, it may be queried by a Read Capacity (16) command to report the PI type which it is currently formatted to. PI Types cannot coexist on a single drive. A drive can only be formatted to a single PI Type. It can be changed at anytime to a new Type but requires a low level format which destroys all existing data on the drive. No other vehicle for changing the PI type is provided by the T10 SBC3 specification.

Type 1 PI format CDB command: 04 90 00 00 00 00, Write Buffer: 00 A0 00 00

Type 2 PI format CDB command: 04 D0 00 00 00 00, Write Buffer: 00 A0 00 00

10.8.3 Identifying a Protection Information drive

The Standard Inquiry provides a bit to indicate if PI is support by the drive. Vital Product Descriptor (VPD) page 0x86 provides bits to indicate the PI Types supported and which PI fields the drive supports checking.

Note. For further details with respect to PI, please refer to SCSI Block Commands - 3 (SBC-3) Draft Standard documentation.

11.0 Installation

Constellation ES.2 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 the 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 12.4.1 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 8. Physical interface

Note. Image is for reference only, may not represent actual drive.

11.1 Drive orientation

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

11.2 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 7.5.1, "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 7.5.1.

Forced air flow may be required to keep temperatures at or below the temperatures specified in Section 7.5.1 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 9. 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 7.5.1 are met.




Figure 9. Air flow

Note. Image is for reference only, may not represent actual drive.

11.3 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.

Note. Do not cover breather hole on top cover.



Breather hole location - top cover

11.4 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.

12.0 Interface requirements

This section partially describes the interface requirements as implemented on Constellation ES.2 drives. Additional information is provided in the *SAS Interface Manual* (part number 100293071).

12.1 SAS features

This section lists the SAS-specific features supported by Constellation ES.2 drives.

12.1.1 task management functions

Table 4 lists the SAS task management functions supported.

Table 4: 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

12.1.2 task management responses

Table 5 lists the SAS response codes returned for task management functions supported.

Table 5: 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

12.2 Dual port support

Constellation ES.2 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 have the capability of independent port clocking (e.g. both ports can run at 6Gb/s or the first port can run at 6Gb/s while the second port runs at 3Gb/s. The supported link rates are 1.5, 3.0, or 6.0 Gb/s.

Subject to buffer availability, the Constellation ES.2 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.

12.3 SCSI commands supported

Table 6 lists the SCSI commands supported by Constellation ES.2 drives.

Table 6: Supported commands

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
DCRT bit supported		Y
DPRY bit supported		N
DSP bit supported		Y
IMMED bit supported		Y
IP bit supported		Y
SI (Security Initialize) bit supported		N
STPF 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

Table 6: Supported commands

Command name	Command code	Supported
Information Exceptions Log page (2Fh)		N
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	Y
Read (32)	7Fh/0009h	N

Table 6: Supported commands

Command name	Command code	Supported
Read Buffer (modes 0, 2, 3, Ah and Bh supported)	3Ch	Y (non-SED drives only)
Read Capacity (10)	25h	Y
Read Capacity (16)	9Eh/10h	Y
Read Defect Data (10)	37h	Y
Read Defect Data (12)	B7h	Y
Read Long	3Eh	Y (non-SED drives only)
Read Long (16)	9Eh/11h	Y
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
Security Protocol In	A2h	Y (SED models only)
Security Protocol Out	B5h	Y (SED models only)
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	Y

Table 6: Supported commands

Command name	Command code	Supported
Test Unit Ready	00h	Y
Verify (10)	2Fh	Y
BYTCHK bit		Y
Verify (12)	AFh	N
Verify (16)	AFh	Y
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	Y
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	Y
Write and Verify (32)	7Fh/000Ch	N
Write Buffer (modes 0, 2, supported)	3Bh	Y (non-SED drives only)
Write Buffer	3Bh	
Firmware Download option (modes 5, 7, Ah and Bh) [3]		Y (non-SED drives only)
Firmware Download option (modes 4, 5, 7)		Y (SED drives only)
Write Long (10)	3Fh	Y
Write Long (16)	9Fh/11h	Y
Write Same (10)	41h	Y
PBdata		N
LBdata		N
Write Same (16)	93h	Y
Write Same (32)	7Fh/000Dh	N
XDRead	52h	N
XDWrite	50h	N
XPWrite	51h	N

- [1] Constellation ES.2 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.

12.3.1 Inquiry data

Table 7 lists the Inquiry command data that the drive should return to the initiator per the format given in the *SAS Interface Manual*.

Table 7: Constellation ES.2 inquiry data

Bytes	Data (hex)																
0-15	00	00	xx**	12	8B	00	30	02	53	45	41	47	41	54	45	20	Vendor ID
16-31	[53	54	33	33	30	30	30	36	35	30	53	53]	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*	39*	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 ST33000650SS.

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

ST33000651SS 53 54 33 33 30 30 30 36 35 31 53 53

ST33000652SS 53 54 33 33 30 30 30 36 35 32 53 53

12.3.2 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 8, 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 8: Mode Sense data changeable and default values for 3TB drives

MODE DATA HEADER:

01 9a 00 10 01 00 00 10

BLOCK DESCRIPTOR:

00 00 00 01 5d 50 a3 b0 00 00 00 00 00 00 02 00

MODE PAGES:

DEF 81 0a c0 14 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 00 01 00 00 00 00 00 02 00 03 02 00 00 01 00 00 00 00 40 00 00 00
CHG 83 16 00

DEF 84 16 00 00 04 0a 00 00 00 00 00 00 00 00 00 00 00 00 00 00 1c 20 00 00
CHG 84 16 00

DEF 87 0a 00 14 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 06 00 00 80 00 00 00 00 6b d0
CHG 8a 0a 07 f6 00 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 0e 46 00 07 d0 00 00 00 00 00 00 00 00 00 00
CHG 99 0e 50 00 ff ff ff ff ff ff 00 00 00 00 00 00

DEF 9a 26 00 06 00 00 00 0a 00 00 8c a0 00 00 17 70 00 00 46 50 00 00 46 50 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
CHG 9a 26 01 0f ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff ff 00 00
00 00 00 00 00 00 00 00 00 00 00 00 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 ff

DEF dc 01 00 0c 01 01 00 48 00 18 01 f4 00 00 00 00
CHG dc 01 00 0c 00 01 ff ff ff ff ff ff ff ff 00 00

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

12.4 Miscellaneous operating features and conditions

Table 9 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 9: 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 10: 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

12.4.1 SAS physical interface

Figure 10 shows the location of the SAS device connector J1. Figures 11 and 12 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 10. Physical interface

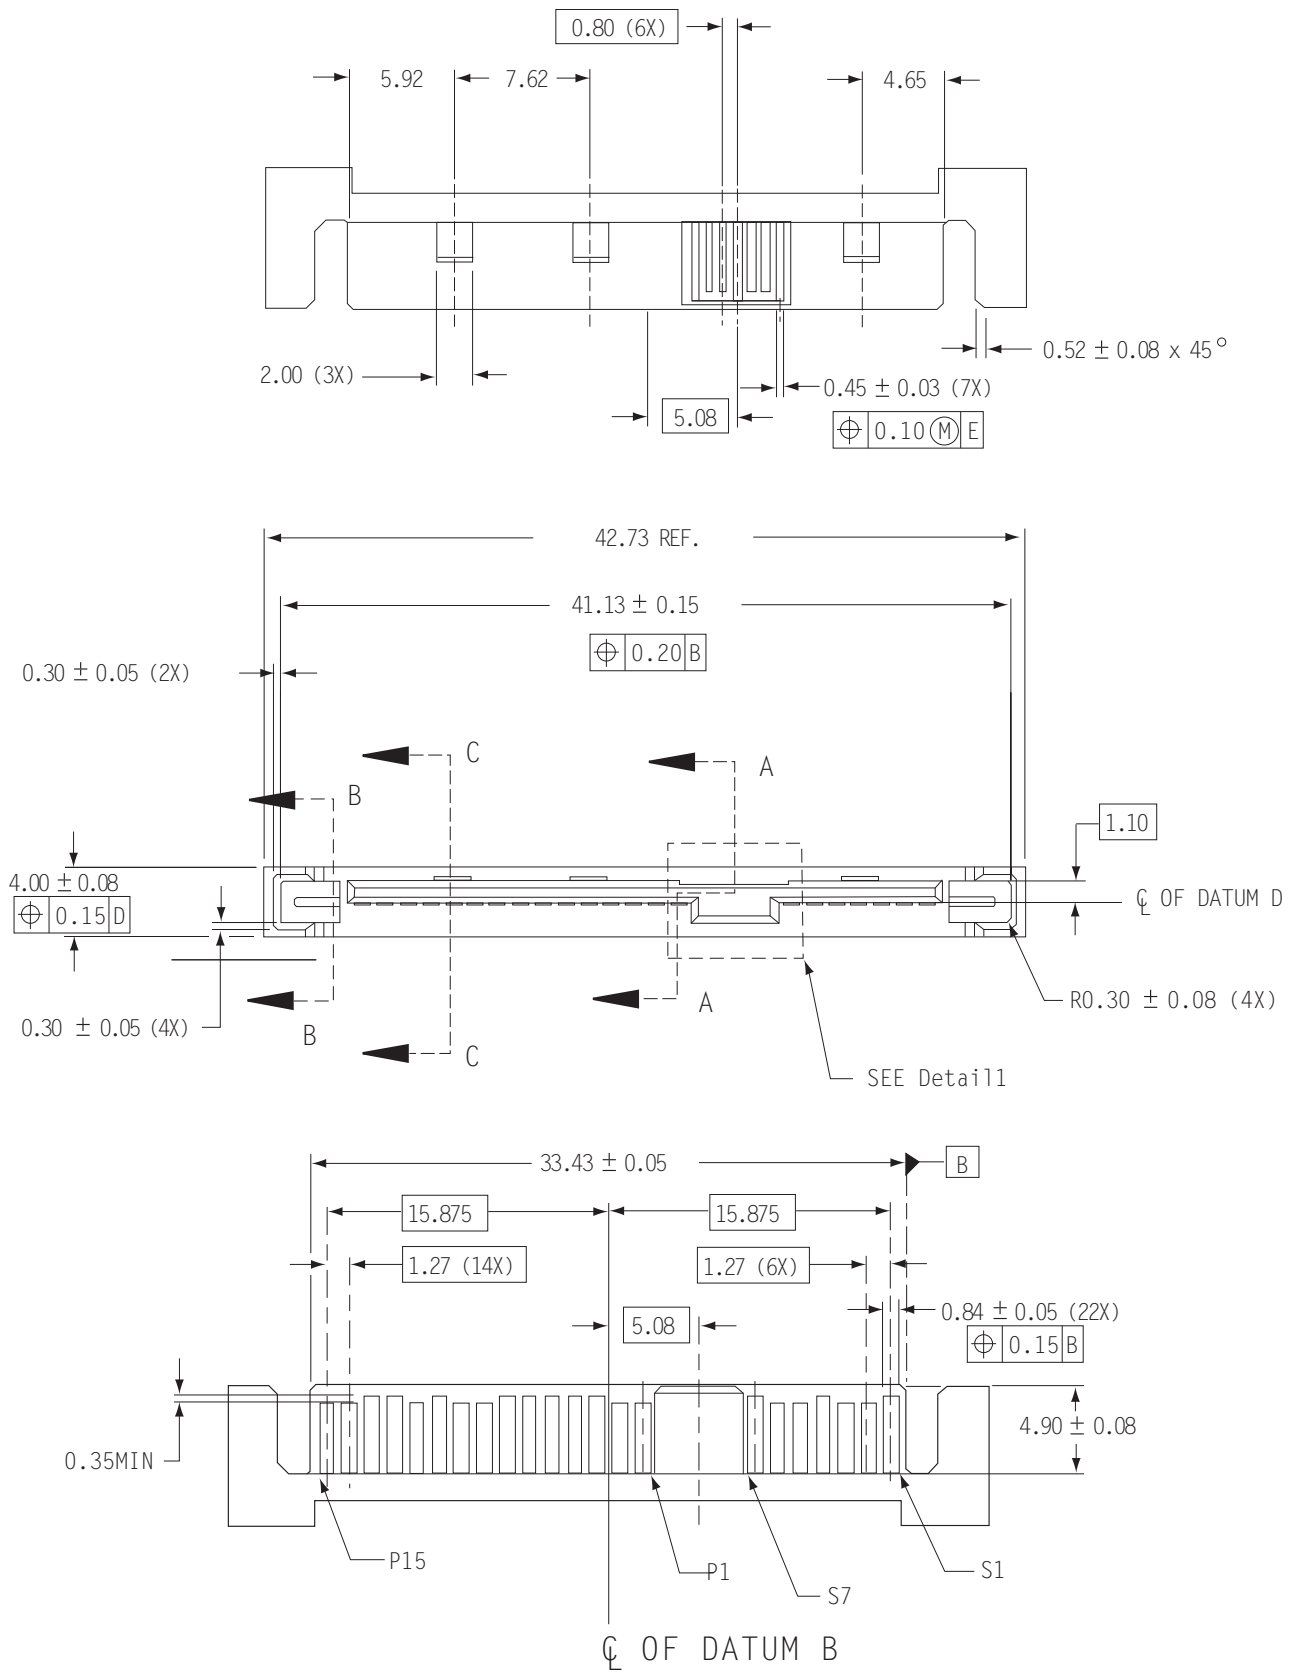


Figure 11. SAS device plug dimensions

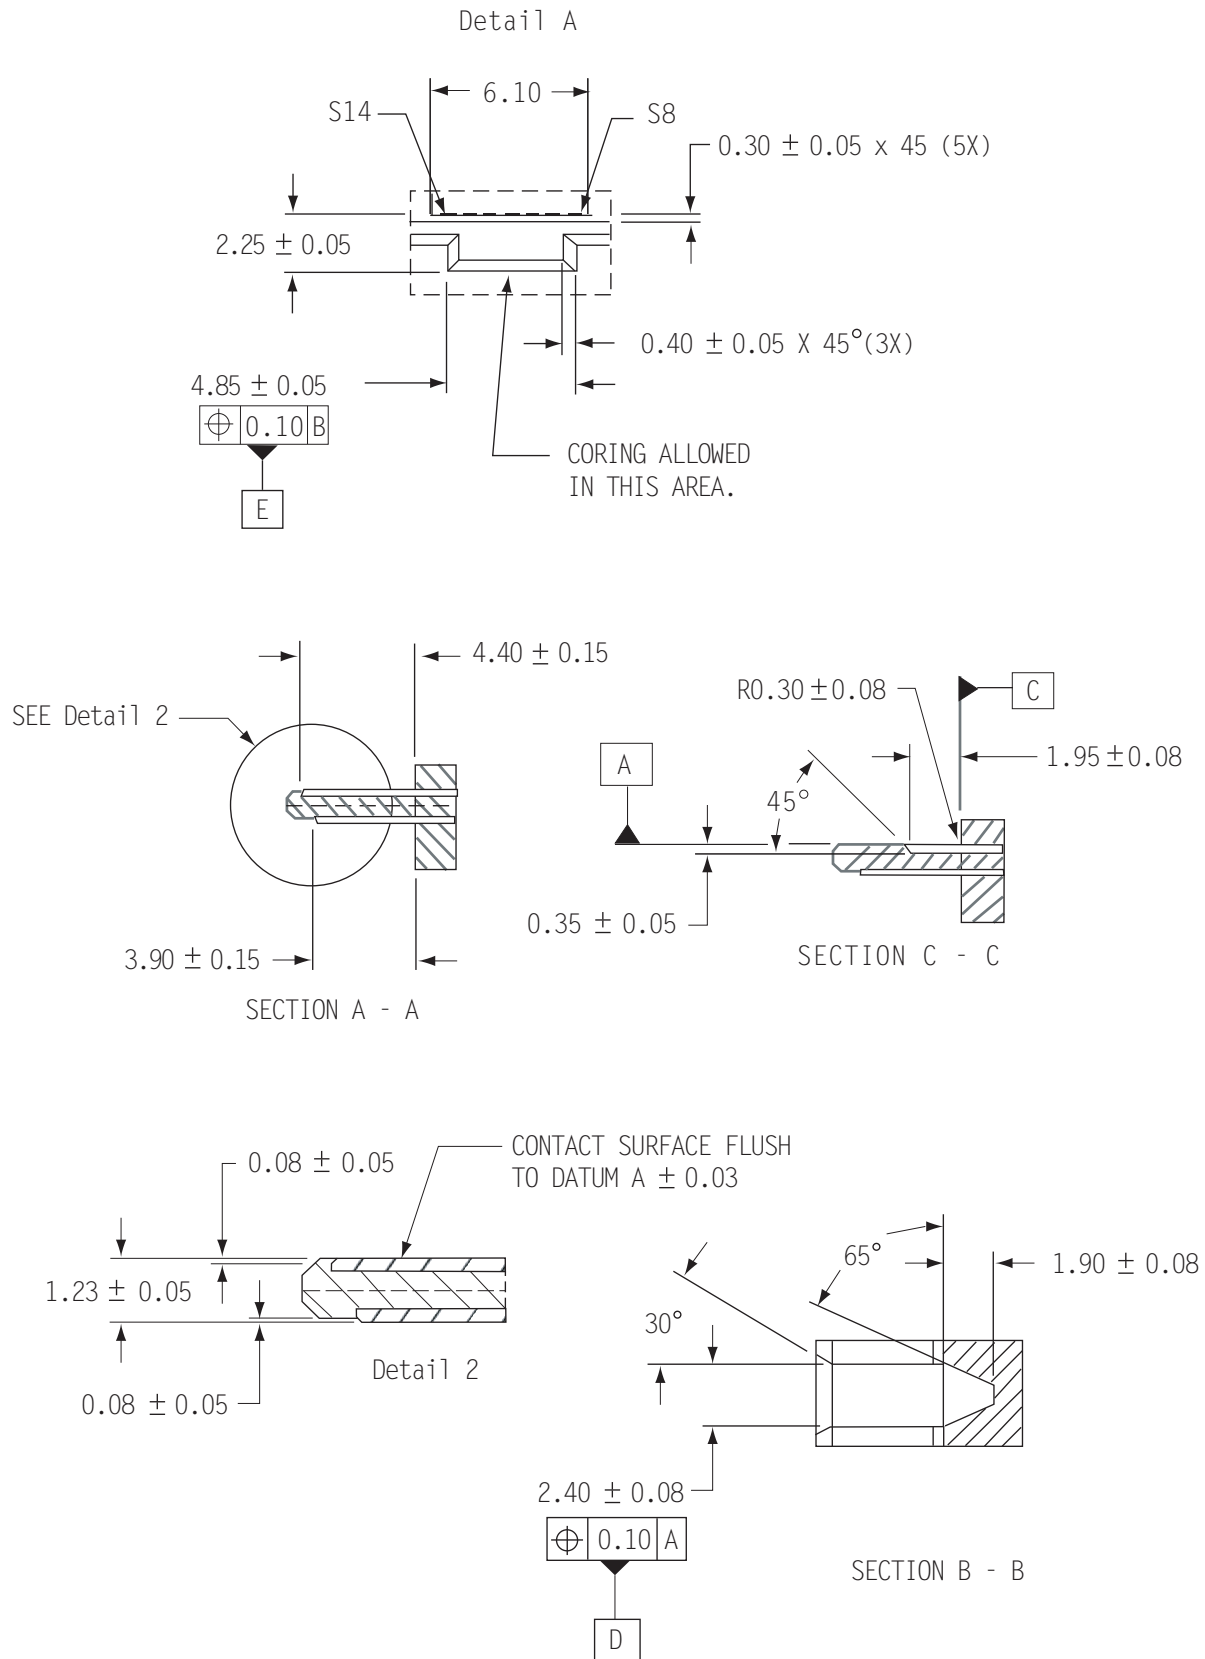


Figure 12. SAS device plug dimensions (detail)

12.4.2 Physical characteristics

This section defines physical interface connector.

12.4.3 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 11 and 12.

12.4.4 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.

12.4.5 Pin descriptions

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

Table 11: 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.

12.4.6 SAS transmitters and receivers

A typical SAS differential copper transmitter and receiver pair is shown in Figure 13. The receiver is AC coupling to eliminate ground shift noise.




Figure 13. SAS transmitters and receivers

12.4.7 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.

12.5 Signal characteristics

This section describes the electrical signal characteristics of the drive's input and output signals. See Table 11 for signal type and signal name information.

12.5.1 Ready LED Out

The Ready LED Out signal is driven by the drive as indicated in Table 12.

Table 12: 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)	
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 13 for the output characteristics of the LED drive signals.

Table 13: 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	$I_{OL} = 15\text{ mA}$	$0 \leq V_{OL} \leq 0.225\text{ V}$

12.5.2 Differential signals

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

Table 14 defines the general interface characteristics.

Table 14: General interface characteristics

Characteristic	Units	1.5Gb/s	3.0Gb/s	6.0Gb/s
Bit rate (nominal)	Mbaud	1,500	3,000	6,000
Unit interval (UI)(nominal)	ps	666.6	333.3	166.6
Impedance (nominal, differential)	ohm	100	100	100
Transmitter transients, maximum	V	± 1.2	± 1.2	± 1.2
Receiver transients, maximum	V	± 1.2	± 1.2	± 1.2

12.6 SAS-2 Specification Compliance

Seagate SAS-2 drives are entirely compatible with the latest SAS-2 Specification (T10/1760-D) Revision 16.

The most important characteristic of the SAS-2 drive at 6Gb/s is that the receiver is capable of adapting the equalizer to optimize the receive margins. The SAS-2 drive has two types of equalizers:

1. A Decision Feedback Equalizer (DFE) which utilizes the standard SAS-2 training pattern transmitted during the SNW-3 training gap. The DFE circuit can derive an optimal equalization characteristic to compensate for many of the receive losses in the system.
2. A Feed Forward Equalizer (FFE) optimized to provide balanced receive margins over a range of channels bounded by the best and worst case channels as defined by the relevant ANSI standard.

12.7 Additional information

Please contact your Seagate representative for SAS electrical details, if required.

For more information about the Phy, Link, Transport, and Applications layers of the SAS interface, refer to the Seagate SAS Interface Manual, part number 100293071.

For more information about the SCSI commands used by Seagate SAS drives, refer to the Seagate SCSI Commands Reference Manual, part number 100293068.

Index

Numerics

12 volt
 pins 59
5 volt pins 59
6 Gbps 60

A

abort task set function 45
AC coupling 59
AC power requirements 22
ACA active status 54
ACA active, faulted initiator status 54
acoustics 31
active LED Out signal 59
actuator 9
 assembly design 7
adaptive caching 54
Admin SP 35
AES-128 data encryption 35
air cleanliness 31
air flow 43
 illustrated 43
air inlet 43
altitude 29
ambient 28
ambient temperature 43
ANSI documents
 SCSI 6
 Serial Attached SCSI 6
asynchronous event notification 54
audible noise 3
Australian C-Tick 4
auto write and read reallocation
 programmable 8
automatic contingent allegiance 54
average idle current 23
average rotational latency 10

B

Background Media Scan 39
backpanel 58
Band 0 36
BandMasterX 35
BMS 39
BSMI 4
buffer
 data 8
 space 12
busy status 54
bytes per surface 10
bytes per track 10

C

cache operation 12
cache segments 12
caching write data 13
Canadian Department of Communications 3
capacity
 unformatted 10
CBC 35
CE Marking 4
check condition status 54
China RoHS directive 5
Cipher Block Chaining 35
class B limit 3
clear ACA function 45
clear task set function 45
commands supported 47
condensation 28
condition met/good status 54
connector
 illustrated 58
 requirements 58
continuous vibration 31
cooling 43
CRC
 error 15
Cryptographic erase 36
C-Tick 4
Current profiles 25
customer service 20

D

DAR 40
Data Bands 36
data bands 35
data block size
 modifying the 9
data buffer to/from disc media 11
Data encryption 35
Data Encryption Key 35
data heads
 read/write 10
data rate
 internal 10
data transfer rate 11
DC power 58
 requirements 22
Decision Feedback Equalizer 60
decrypt 35
default MSID password 36
defect and error management 38
defects 38
Deferred Auto-Reallocation 40
deferred error handling 54
DEK 35

- description 7
- DFE 60
- dimensions 33
- disc rotation speed 10
- drive 31
- drive characteristics 10
- Drive Locking 36
- drive mounting 33, 44
- drive select 58
- dual port support 46

E

- electrical
 - description of connector 58
 - signal characteristics 59
 - specifications 21
- electromagnetic compatibility 3
- Electromagnetic Compatibility (EMC) 4
- Electromagnetic Compatibility control Regulation 4
- Electromagnetic compliance for the European Union 4
- electromagnetic susceptibility 32
- EMI requirements 3
- encryption engine 35
- encryption key 36
- environment 43
- environmental
 - limits 28
 - requirements 14
- environmental control 31
- EraseMaster 35
- error
 - management 38
 - rates 14
- errors 38
- European Union 4

F

- FCC rules and regulations 3
- features 8
 - interface 45
- feed forward equalizer 60
- FFE 60
- FIPS 34
- firmware 8
 - corruption 50
- firmware download port 36
- flawed sector reallocation 8
- Format command execution time 11
- function
 - complete, code 00 45
 - not supported, code 05 45
 - reject, code 04 45

G

- Global Data Band 36
- Good status 54
- gradient 28
- ground shift noise 59
- grounding 44

H

- HDA 43, 44
- head and disc assembly (HDA) 7
- head and disc assembly. See HDA
- heads
 - read/write data 10
- heat removal 43
- heat source 43
- host equipment 44
- hot plugging the drive 15
- humidity 28
- humidity limits 28

I

- Identifying a PI drive 41
- Idle Read After Write 40
- Idle1 21
- Idle2 21
- Idle3 21
- inquiry data 51
- installation 42
 - guide 6
- interface
 - commands supported 47
 - error rate 14
 - errors 15
 - illustrated 55
 - physical 55
 - requirements 45
- interleave
 - minimum 11
- intermediate/condition met/good status 54
- intermediate/good status 54
- internal data rate 10
- internal defects/errors 38
- internal drive characteristics 10
- IRAW 40

J

- jumpers 42

K

- KCC 4
- Korean Communications Commission 4
- Korean KCC 4

L

- latency
 - average rotational 10, 11
- Locking SP 35
- LockOnReset 36
- logical block address 12
- logical block reallocation scheme 8
- logical block size 8, 11
- logical segments 12

M

- maintenance 14
- Makers Secure ID 35
- maximum delayed motor start 23
- maximum start current 23
- mean time between failure. See MTBF
- media description 8
- Media Pre-Scan 40
- minimum sector interleave 11
- miscellaneous feature support
 - Adaptive caching 54
 - Asynchronous event notification 54
 - Automatic contingent allegiance 54
 - Deferred error handling 54
 - Parameter rounding 54
 - Queue tagging 54
 - Reporting actual retry count 54
 - Segmented caching 54
 - SMP = 1 in Mode Select command 54
 - Synchronized (locked) spindle operation 54
 - Zero latency read 54
- miscellaneous status support
 - ACA active 54
 - ACA active, faulted initiator 54
 - Busy 54
 - Check condition 54
 - Condition met/good 54
 - Good 54
 - Intermediate/condition met/good 54
 - Intermediate/good 54
 - Reservation conflict 54
 - Task set full 54
- miscorrected media data 14
- Mode sense
 - data, table 51, 53
- mounting 44
 - holes 44
 - orientations 42
- mounting configuration 33
- mounting configuration dimensions 33
- MSID 35, 36
- MTBF 14, 15

N

- noise
 - audible 3
 - noise immunity 24
 - non-operating 28, 29, 31
 - temperature 28
 - non-operating vibration 31

O

- office environment 31
- operating 28, 29, 31
- option selection 58
- options 9
- out-of-plane distortion 44

P

- package size 30
- packaged 30
- parameter rounding 54
- password 35, 36
- passwords 35
- PCBA 44
- peak bits per inch 10
- peak operating current 23
- peak-to-peak measurements 24
- performance characteristics
 - detailed 10
 - general 11
- performance degradation 29
- performance highlights 8
- physical damage 31
- physical interface 55
- physical specifications 21
- PI level - Type 0 41
- PI level - Type I 41
- PI level - Type II 41
- PI level - Type III 41
- PI Levels 41
- pin descriptions 58
- power 59
 - dissipation 26
 - requirements, AC 22
 - requirements, DC 22
 - sequencing 24
- Power Condition mode page 21
- power distribution 3
- power management 21
- PowerChoice 21
- PowerChoice reports 22
- PowerCycle 36
- prefetch/multi-segmented cache control 12
- preventive maintenance 14
- protection information 41
- protection of data at rest 35

Q

queue tagging 54

R

radio interference regulations 3

Random number generator 36

RCD bit 12

read error rates 14, 38

read/write data heads 10

receivers 59

recommended mounting 30

Recoverable Errors 14

recovered media data 14

reference

documents 6

relative humidity 28

reliability 9

specifications 14

reliability and service 15

repair and return information 20

reporting actual retry count 54

reservation conflict status 54

resonance 29

return information 20

RNG 36

RoHS 5

rotation speed 10

S

safety 3

SAS

interface 58

physical interface 55

task management functions 45

SAS documents 6

SAS Interface Manual 3, 6

SAS-2 specification compliance 60

SCSI interface

commands supported 47

Secure ID 35

security partitions 35

Security Protocol In 35

Security Protocol Out 35

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 54

self-encrypting drives 35

Self-Monitoring Analysis and Reporting Technology

9, 16

Serial Attached SCSI (SAS) Interface Manual 2

shielding 3

shipping 20

shipping container 28

shock 29

and vibration 29

shock mount 44

SID 35

signal

characteristics 59

single-unit shipping pack kit 9

SMART 9, 16

SMP = 1 in Mode Select command 54

SNW-3 training gap 60

standards 3

Standby1 21

Standby2 21

START STOP UNIT command 21

start/stop time 11

support services 1

surface stiffness

allowable for non-flat surface 44

switches 42

synchronized spindle

operation 54

system chassis 44

T

Taiwanese BSMI 4

task management functions 45

Abort task set 45

Clear ACA 45

Clear task set 45

terminate task 45

task management response codes 45

Function complete 00 45

Function not supported 05 45

Function reject 04 45

task set full status 54

TCG 35

technical support services 1

temperature 28, 43

limits 28

non-operating 28

regulation 3

See also cooling

terminate task function 45

terminators 42

tracks per inch 10

tracks per surface 10

transmitters 59

transporting the drive 20

Trusted Computing Group 35

Type 1 PI format 41

Type 2 PI format 41

U

unformatted 9

Unrecoverable Errors 14

unrecovered media data 14

V

vibration 29, 31

W

warranty 20

Z

zero latency read 54



Seagate Technology LLC
920 Disc Drive, Scotts Valley, California 95066-4544, USA
Publication Number: 100628615, Rev. D