

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 blue and grey.

Product Manual

Pulsar[®] XT.2 SAS

Standard Models

ST400FX0002
ST200FX0002
ST100FX0002

Self-Encrypting Drive Models

ST400FX0012

100647497
Rev. B
June 2011

Revision history

Revision	Date	Sheets affected or comments
Rev. A	03/16/11	Initial release.
Rev. B	06/01/11	34. (7mm weight correction)

© 2011 Seagate Technology LLC. All rights reserved.

Publication number: 100647497, Rev. B June 2011

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

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

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)	4
3.1.4	China Restriction of Hazardous Substances (RoHS) Directive	5
3.2	Reference documents	5
4.0	General description	6
4.1	Standard features	6
4.2	Media description	7
4.3	Performance	7
4.4	Reliability	7
4.5	Formatted capacities	8
4.6	Programmable drive capacity	8
4.7	Factory-installed options	8
4.8	Thin Provisioning	9
4.8.1	Logical Block Provisioning	9
4.8.2	Thin Provisioning capabilities	9
4.8.3	UNMAP	9
4.8.4	FORMAT UNIT command	9
4.8.5	Protection Information (PI) and Security (SED)	9
5.0	Performance characteristics	11
5.1	Internal drive characteristics	11
5.2	Performance characteristics	11
5.2.1	Access time	11
5.2.2	FORMAT UNIT command execution time for 512-byte LBA's (minutes)	12
5.2.3	Performance	12
5.3	Start/stop time	13
5.4	Cache control	13
5.4.1	Caching write data	13
6.0	Reliability specifications	14
6.1	Error rates	14
6.1.1	Unrecoverable Errors	15
6.1.2	Interface errors	15
6.2	Endurance Management	15
6.2.1	Wear Leveling	15
6.2.2	Garbage Collection	15
6.2.3	Write Amplification	15
6.2.4	UNMAP	15
6.2.5	Data Retention	16
6.2.6	Lifetime Endurance Management	16
6.3	Reliability and service	16
6.3.1	Annualized Failure Rate (AFR) and Mean Time Between Failure (MTBF)	16
6.3.2	Preventive maintenance	16
6.3.3	Hot plugging the drive	16
6.3.4	S.M.A.R.T.	17
6.3.5	Thermal monitor	18
6.3.6	Drive Self Test (DST)	19

6.3.7	Product warranty	21
7.0	Physical/electrical specifications	22
7.1	Power specifications	22
7.1.1	Power consumption	22
7.2	AC power requirements	22
7.3	DC power requirements	23
7.3.1	Conducted noise immunity	26
7.3.2	Power sequencing	26
7.3.3	Current profiles	26
7.4	Power dissipation	28
7.5	Environmental limits	29
7.5.1	Temperature	29
7.5.2	Relative humidity	30
7.5.3	Effective altitude (sea level)	30
7.5.4	Shock and vibration	30
7.5.5	Air cleanliness	32
7.5.6	Corrosive environment	32
7.5.7	Electromagnetic susceptibility	32
7.6	Mechanical specifications	33
8.0	About self-encrypting drives	35
8.1	Data encryption	35
8.2	Controlled access	35
8.2.1	Admin SP	35
8.2.2	Locking SP	36
8.2.3	Default password	36
8.3	Random number generator (RNG)	36
8.4	Drive locking	36
8.5	Data bands	36
8.6	Cryptographic erase	37
8.7	Authenticated firmware download	37
8.8	Power requirements	37
8.9	Supported commands	37
8.10	RevertSP	37
9.0	Defect and error management	38
9.1	Drive internal defects/errors	38
9.2	Drive error recovery procedures	39
9.3	SAS system errors	39
9.4	Background Media Scan	39
9.5	Auto-Reallocation	39
9.6	Protection Information (PI)	40
9.6.1	Levels of PI	40
9.6.2	Setting and determining the current Type Level	40
9.6.3	Identifying a Protection Information drive	40
10.0	Installation	41
10.1	Drive orientation	41
10.2	Cooling	42
10.3	Drive mounting	43
10.4	Grounding	43
11.0	Interface requirements	44
11.1	SAS features	44
11.1.1	Task management functions	44
11.1.2	Task management responses	44

11.2	Dual port support	45
11.3	SCSI commands supported.	46
	11.3.1 INQUIRY data	50
	11.3.2 MODE SENSE data	50
11.4	Miscellaneous operating features and conditions	55
	11.4.1 SAS physical interface	55
	11.4.2 Physical characteristics	58
	11.4.3 Connector requirements.	58
	11.4.4 Electrical description	58
	11.4.5 Pin descriptions	58
	11.4.6 SAS transmitters and receivers	59
	11.4.7 Power	59
11.5	Signal characteristics.	59
	11.5.1 Ready LED Out	59
	11.5.2 Differential signals	60
11.6	SAS-2 Specification compliance	60
11.7	Additional information	60

List of Figures

Figure 1.	Current profiles for 400GB models	26
Figure 2.	Current profiles for 200GB models	27
Figure 3.	Current profiles for 100GB models	27
Figure 4.	400GB (at 6Gb) DC current and power vs. input/output operations per second	28
Figure 5.	200GB (at 6Gb) DC current and power vs. input/output operations per second	28
Figure 6.	100GB (at 6Gb) DC current and power vs. input/output operations per second	29
Figure 7.	Temperature check point location - 15mm drives	30
Figure 8.	Temperature check point location - 7mm drives	30
Figure 9.	Recommended mounting	31
Figure 10.	Mounting configuration dimensions (400GB models)	33
Figure 11.	Mounting configuration dimensions (200 & 100GB models)	34
Figure 12.	Physical interface	41
Figure 13.	Air flow	42
Figure 14.	Physical interface	55
Figure 15.	SAS device plug dimensions	56
Figure 16.	SAS device plug dimensions (detail)	57
Figure 17.	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, Pulsar® XT.2 SAS (Serial Attached SCSI) drives.

Pulsar XT.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.

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.

Standard models	Standard SED models
ST400FX0002	ST400FX0012
ST200FX0002	
ST100FX0002	

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.

Note. 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).

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 depend 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 drives must be securely mounted to guarantee the specified performance characteristics. Mounting by bottom holes must meet the requirements of Section 10.3.

3.1 Standards

The Pulsar XT.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 and CSA 60950 as tested by UL(CSA) and EN60950 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 Tables 8 through 10, 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 KN22 and KN61000.

Taiwanese BSMI

If this model has the Taiwanese certification mark then it complies with Chinese National Standard, CNS13438.

3.1.3 European Union Restriction of Hazardous Substances (RoHS)

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

A number of parts and materials in Seagate products are procured from external suppliers. We rely on the representations of our suppliers regarding the presence of RoHS substances in these parts and materials. Our supplier contracts require compliance with our chemical substance restrictions, and our suppliers document their compliance with our requirements by providing material content declarations for all parts and materials for the 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.

4.0 General description

Pulsar XT.2 drives provide high performance, high capacity data storage for a variety of systems with a Serial Attached SCSI (SAS) interface. The Serial Attached SCSI interface is designed to meet next-generation computing demands for performance, scalability, flexibility and high-density storage requirements.

Pulsar XT.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. Pulsar XT.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 drives. This allows users to continue to leverage 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.trustedcomputing-group.org).

Note. Never disassemble and do not attempt to service items in the enclosure. The drive does not contain user-replaceable parts. Opening for any reason voids the drive warranty.

4.1 Standard features

Pulsar XT.2 SAS drives have the following standard features:

- 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, 524, 528, 4096, 4160, 4192, or 4224 bytes per logical block)
- Industry standard SFF 2.5-inch dimensions
- ECC maximum burst correction length of 90 bits
- No preventive maintenance or adjustments required
- Self diagnostics performed when power is applied to the drive
- Vertical, horizontal, or top down mounting
- Drive Self Test (DST)
- Background Media Scan (BMS)
- Parallel flash access channels
- Power loss data protection
- Thin Provisioning with Block Unmap Support
- Silent operation
- Lifetime Endurance Management

Pulsar XT.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 consists of Single Layer Cell (SLC) NAND Flash for improved reliability and performance.

4.3 Performance

- Programmable multi-segmentable cache buffer
- 600MB/s maximum instantaneous data transfers.
- Background processing of queue
- Non-Volatile Write Cache

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 0.44%
- Mean time between failures (MTBF) of 2,000,000 hours
- 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 UNIT 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:

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 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 shown below. The Capacity stated is identical when the drive is formatted with or without PI enabled.

Table 1: Formatted Capacity LBA Count

Capacity (LBAs)						
LBA Size	400GB		200GB		100GB	
	Decimal	Hex	Decimal	Hex	Decimal	Hex
512	781,422,768	2E9390B0h	390,721,968	1749F1B0h	195,371,568	BA52230h
520	764,871,800	2D970478h	382,435,904	16CB8240h	191,217,952	B65C120h
524	754,677,072	2CFB7550h	377,338,536	167DBAA8h	188,669,272	B3EDD58h
528	743,833,040	2C55FDD0h	371,916,520	162AFEE8h	185,958,264	B157F78h
4096	97,677,846	5D27216h	48,840,246	2E93E36h	24,421,446	174A446h
4160	96,153,847	5BB30F7h	48,076,924	2DD987Ch	24,038,462	16ECC3Eh
4192	95,419,848	5AFFDC8h	47,709,924	2D7FEE4h	23,854,962	16BFF72h
4224	94,696,970	5A4F60Ah	47,348,485	2D27B05h	23,674,243	1693D83h

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

OEMs 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 LBA 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.

4.8 Thin Provisioning

4.8.1 Logical Block Provisioning

The drive is designed with a feature called Thin Provisioning. Thin Provisioning is a technique which does not require Logical Blocks to be associated to Physical Blocks on the storage medium until such a time as needed. The use of Thin Provisioning is a major factor in SSD products because it reduces the amount of wear leveling and garbage collection that must be performed. The result is an increase in the products endurance. For more details on Logical Block Provisioning and Thin Provisioning, Reference the SBC-3 document provided by the T-10 committee.

4.8.2 Thin Provisioning capabilities

The level of Thin Provisioning support may vary by product model. Devices that support Thin Provisioning are allowed to return a default data pattern for read requests made to Logical Blocks that have not been mapped to Physical Blocks by a previous WRITE command.

In order to determine if Thin Provisioning is supported and what features of it are implemented requires the system to send a READ CAPACITY 16 (9Eh) command to the drive. Thin Provisioning and the READ CAPACITY 16 (9Eh) command is defined in the Seagate SCSI Command Reference 100293068.

Table 2: Thin Provisioning Product Configuration

Product Configuration	LBPME	LBPRZ
Non-SED	Supported	Supported
SED	Supported	Not Supported

A logical block provisioning management enabled (LBPME) bit set to one indicates that the logical unit implements logical block provisioning management. An LBPME bit set to zero indicates that the logical unit is fully provisioned and does not implement logical block provisioning management.

A logical block provisioning read zeros (LBPRZ) bit set to one indicates that, for an unmapped LBA specified by a read operation, the device server sends user data with all bits set to zero to the data-in buffer. An LBPRZ bit set to zero indicates that, for an unmapped LBA specified by a read operation, the device server may send user data with all bits set to any value to the data-in buffer.

4.8.3 UNMAP

The UNMAP command requests that the device server break the association of a specific Logical Block address from a Physical Block, thereby freeing up the Physical Block from use and no longer requiring it to contain user data. An unmapped block will respond to a READ command with data that is determined by the setting of the LBPRZ bit in the READ CAPACITY parameter data.

4.8.4 FORMAT UNIT command

A device which supports Thin Provisioning will be capable of performing a SCSI FORMAT UNIT command which allocates Logical Blocks Addresses that are not linked to Physical Block Locations. A FORMAT command will cause all LBAs to become unmapped.

4.8.5 Protection Information (PI) and Security (SED)

The requirements in this section apply to any device which supports LBA unmapping.

In SCSI devices, unmapped LBAs are defined as part of the Thin Provisioning model. Support of the Thin Provisioning model is indicated by the LBPME bit having a value of '1' in the READ CAPACITY (16) parameter data.

When a region of LBA's are erased via cryptographic erase, as part of the erase, the drive shall unmap those LBAs.

If the host attempts to access an unmapped or trimmed LBA, the drive shall return scrambled data. For a given LBA, the data shall be identical from access to access, until that LBA is either updated with actual data from the host or that LBA is cryptographically erased. The drive shall report a value of '0' in the LBPRZ field returned in the READ CAPACITY (16) parameter data.

If the host attempts to access an unmapped LBA on a drive that has been formatted with Protection Information (PI), the drive shall return scrambled PI data for that LBA. Depending on the value of the RDPROTECT field in the data-access command CDB, this may result in the drive returning a standard PI error to the host.

If the host reduces the addressable capacity of the drive via a MODE SELECT command, the drive shall unmap or trim any LBA within the inaccessible region of the device.

Additionally, an UNMAP command is not permitted on a locked band.

Table 3: PI and SED Drive Configuration

	Drive Configuration					
	Standard			SED		
	Disabled	Enabled		Disabled	Enabled	
PI Setting	Disabled	Enabled		Disabled	Enabled	
PROT_EN bit	0	1		0	1	
LBPME bit	1	1		1	1	
LBPRZ bit	1	1		0	0	
PI Check Requested	N/A	Yes	No	N/A	Yes	No
DATA Returned for Thin Provisioned LBA	0x00	0x00	0x00	Random	None	Random
PI Returned for Thin Provisioned LBA	None	0xFF	0xFF	None	None	Scrambled PI data
PI Check Performed	N/A	No	No	N/A	Yes	No
Error reported to Host	No	No	No	No	Yes	No

5.0 Performance characteristics

This section provides detailed information concerning performance-related characteristics and features of Pulsar XT.2 drives.

Note. Data provided is based on format at 512-bytes.

5.1 Internal drive characteristics

	ST400FX0002 ST400FX0012	ST200FX0002	ST100FX0002	
Drive capacity	400	200	100	GB (formatted, rounded off value)
Flash Memory Type	NAND SLC			
Emulated LBA Size	512, 520, 524, 528, 4096, 4160, 4192, or 4224			
Native Programmable Page Size	4096 User Bytes			
Default Transfer Alignment Offset	0			

5.2 Performance characteristics

See Section 11.4.1, "SAS physical interface" and the *SAS Interface Manual* (part number 100293071) for additional timing details.

5.2.1 Access time

Access measurements are taken with nominal power at 25°C ambient temperature. All times are measured using drive diagnostics. The specifications in the table below are defined as follows:

- Page-to-page access time is an average of all possible page-to-page accesses in both directions for a sequentially preconditioned drive.
- Average access time is a true statistical random average of at least 5000 measurements of accesses between programmable pages on a randomly preconditioned drive.

Table 4: Typical Access Time (µsec)

		400GB ^{1,2}		100/200 GB ^{1,2}	
		Read	Write	Read	Write
Average	Typical ³	268	133	208	121
Page to Page		268	133	207	121
Average Latency		247		188	

1. Execution time measured from receipt of the Command to the Response.
2. Assumes no errors.
3. Typical access times are measured under nominal conditions of temperature, voltage, and horizontal orientation as measured on a representative sample of drives.

Note. These drives are designed to provide the highest possible performance under typical conditions. However, due to the nature of Flash memory technologies there are many factors that can result in values different than those stated in this specification

5.2.2 FORMAT UNIT command execution time for 512-byte LBA's (minutes)

The device may be formatted as either a Thin Provisioned device or a Fully Provisioned device. The default format is Thin Provisioned and is recommended for most applications. Thin Provisioning provides the most flexibility for the device to manage the flash medium to maximize endurance.

Table 5: Maximum FORMAT UNIT Times (minutes)

Configuration	Format Mode	DCRT Bit	IP Bit	400GB	200GB	100GB
Non-SED	(Default) Thin Provisioned	DCRT = 0	IP = 0	5	5	5
Non-SED	(Default) Thin Provisioned	DCRT = 1	IP = 0	5	5	5
Non-SED	Fully Provisioned	DCRT = 0	IP = 1	140	60	30
Non-SED	Fully Provisioned	DCRT = 1	IP = 1	100	40	20
SED	(Default) Thin Provisioned	DCRT = 0	IP = 0	5	N/A	N/A
SED	(Default) Thin Provisioned	DCRT = 1	IP = 0	5	N/A	N/A
SED	Fully Provisioned	DCRT = 0	IP = 1	140	N/A	N/A
SED	Fully Provisioned	DCRT = 1	IP = 1	100	N/A	N/A

5.2.3 Performance

Table 6: Performance

	Notes	400GB	200GB	100GB
Maximum Burst Transfer Rate		600MB/s		
Peak sequential 128KB read/write data transfer rate (MB/s max)	[1]	360/300		
Sustained sequential 128KB read/write data transfer rate (MB/s)	[1]	300/300	300/200	300/100
Peak 4KB random read/write command rate (IOPs)	[2]	48,000/22,000		
Sustained 4KB random read/write command rate (IOPs)	[2]	48,000/ 22,000	48,000/ 16,000	48,000/ 8000
Sustainable 4KB Random combined IOPS for 5 year Endurance (65%/35% R/W, 70% Duty Cycle)	[3]	31,000	31,000	31,000

[1] Testing performed at Queue Depth = 32, Sequentially Preconditioned drive, using IOMeter 2006.7.27.

[2] Testing performed at Queue Depth = 32, Randomly Preconditioned drive, using IOMeter 2006.7.27.

[3] Testing performed at Queue Depth = 32, Non-Preconditioned drive, using IOMeter 2006.7.27.

Note. IOMeter is available at <http://www.iometer.org/> or <http://sourceforge.net/projects/iometer/>. IOMeter is licensed under the Intel Open Source License and the GNU General Public License. Intel does not endorse any IOMeter results.

Peak performance is defined as the typical best case performance that the product will be able to achieve when the product is preconditioned as mentioned and host commands are aligned on 4KB boundaries.

Sustained performance is defined as the typical worst case performance that the product will be able to achieve when the product is preconditioned as mentioned and host commands are aligned on 4KB boundaries. Write values also take into account the worst case performance throttling that may occur to ensure the product meets specified reliability specifications.

Due to the nature of Flash memory technologies there are many factors that can result in values different than those stated in this specification. Some discrepancies can be caused by bandwidth limitations in the host adapter, operating system, or driver limitations. It is not the intent of this manual to cover all possible causes of performance discrepancies.

When evaluating performance of SSD devices, it is recommended to measure performance of the device in a method that resembles the targeted application using real world data and workloads. Test time should also be adequately large to ensure that sustainable metrics and measures are obtained.

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 13 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 13 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. Stop time is 3 seconds (maximum) from removal of DC power. SCSI stop time is 3 seconds. There is no power control switch on the drive.

5.4 Cache control

All default cache mode parameter values (Mode Page 08h) for standard OEM versions of this drive family are given in Table 17 and 18.

5.4.1 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 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.

Data that has not been written to the medium is protected by a back up power source which provides the ability of the data to be written to non-volatile medium in the event of an unexpected power loss.

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. Tables 16, 17 and 18 show the mode default settings for the drive.

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.

Read Error Rates ¹

Unrecovered Data	Less than 1 LBA in 10^{16} bits transferred
Miscorrected Data	Less than 1 LBA in 10^{21} bits transferred
Interface error rate:	Less than 1 error in 10^{12} bits transferred
Mean Time Between Failure (MTBF):	2,000,000 hours
Annualized Failure Rate (AFR):	0.44%
Preventive maintenance:	None required
Typical Data Retention with Power removed (at 40C) ²	3 months
Endurance Rating: ³	Method 1: Full drive writes per day 35
	Method 2: TBW (per JEDEC JESD218)
	400GB = 24,800 TB
	200GB = 12,400 TB
	100GB = 6,200 TB

1. Error rate specified with automatic retries and data correction with ECC enabled and all flaws reallocated.
2. As NAND Flash devices age with use, the capability of the media to retain a programmed value begins to deteriorate. This deterioration is affected by the number of times a particular memory cell is programmed and subsequently erased. When a device is new, it has a powered off data retention capability of up to several years. With use the retention capability of the device is reduced. Temperature also has an effect on how long a Flash component can retain its programmed value with power removed. At high temperature the retention capabilities of the device are reduced. Data retention is not an issue with power applied to the SSD. The SSD drive contains firmware and hardware features that can monitor and refresh memory cells when power is applied.
3. Endurance rating is the expected amount of host data that can be written by product when subjected to a specified workload at a specified operating and storage temperature. For the specific workload to achieve this level of endurance, please reference JEDEC Specification JESD218. TBW is defined as 1×10^{12} Bytes.

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 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 read 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.

6.1.2 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 Endurance Management

Customer satisfaction with Solid State Drives can be directly related to the internal algorithms which an SSD uses to manage the limited number of Program-Erase (PE) cycles that NAND Flash can withstand. These algorithms consist of Wearleveling, Garbage Collection, Write Amplification, Unmap, Data Retention, Lifetime Endurance Management.

6.2.1 Wear Leveling

Wear Leveling is a technique used by the drive to ensure that all Flash cells are written to or exercised as evenly as possible to avoid any hot spots where some cells are used up faster than other locations. Wear Leveling is automatically managed by the drive and requires no user interaction. The Seagate algorithm is tuned to operate only when needed to ensure reliable product operation.

6.2.2 Garbage Collection

Garbage Collection is a technique used by the drive to consolidate valid user data into a common cell range freeing up unused or obsolete locations to be erased and used for future storage needs. Garbage Collection is automatically managed by the drive and requires no user interaction. The Seagate algorithm is tuned to operate only when needed to ensure reliable product operation.

6.2.3 Write Amplification

While Write Amplification is not an algorithm, it is a major characteristic of SSD's that must be accounted for by all the algorithms that the SSD implements. The Write Amplification Factor of an SSD is defined as the ratio of Host/User data requested to be written to the actual amount of data written by the SSD internal to account for the user data and the housekeeping activities such as Wear Leveling and Garbage Collection. The Write Amplification Factor of an SSD can also be directly affected by the characteristics of the host data being sent to the SSD to write. The best Write Amplification Factor is achieved for data that is written in sequential LBA's that are aligned on 4KB boundaries. The worst case Write Amplification Factor typically occurs for randomly written LBA's of transfer sizes that are less than 4KB and that originate on LBA's that are not on 4KB boundaries.

6.2.4 UNMAP

A new SCSI command has been added to the SSD as part of the Thin Provisioning feature set. Use of the UNMAP command reduces the Write Amplification Factor of the drive during housekeeping tasks such as Wear Leveling and Garbage Collection. This is accomplished because the drive does not need to retain data which has been classified by the host as obsolete.

6.2.5 Data Retention

Data Retention is another major characteristic of SSD's that must be accounted for by all the algorithms that the SSD implements. While powered up, the Data Retention of SSD cells are monitored and rewritten if the cell levels decay to an unexpected level. Data Retention when the drive is powered off is affected by Program and Erase (PE) cycles and the temperature of the drive when stored.

6.2.6 Lifetime Endurance Management

As stated in Section 6.2, an SSD has a limited number of Program and Erase (PE) cycles that are capable. In worse case applications, the write workload could be such that the drive experiences a high Write Amplification Factor that could lead to potential wear out prior to the drive achieving its expected field life. Additionally, the Data Retention spec of the SSD needs to be considered to ensure the spec is met once the drive is worn out. Seagate has implemented a Lifetime Endurance Management technique which helps OEMS and user to avoid early wear out. By monitoring the write workload being sent to the drive, the drive can add additional response time to WRITE commands to provide a sustainable level of performance that is capable of being sustained for the life of the drive. Most users may never see this added response time in their applications.

6.3 Reliability and service

Integrators can enhance the reliability of Pulsar XT.2 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 10.2 provides recommended air-flow information.

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

The production drive shall achieve an AFR of 0.44% (MTBF of 2,000,000 hours) when operated in an environment that ensures the 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.

6.3.2 Preventive maintenance

No routine scheduled preventive maintenance is required.

6.3.3 Hot plugging the drive

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

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

Note. It is the responsibility of the systems integrator to assure that no temperature, energy, voltage hazard, or ESD potential hazard is presented during the hot connect/disconnect operation. Discharge

the static electricity from the drive carrier prior to inserting it into the system.

6.3.4 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 administrators 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.

An application 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.

An application 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 applications to control when S.M.A.R.T. interruptions occur. Forcing S.M.A.R.T. with the REZERO UNIT command resets the timer.

Performance impact

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

Maximum processing delay

**Fully-enabled delay
DEXCPT = 0**

S.M.A.R.T. delay times 75 ms

Reporting control

Reporting is controlled by the MRIE bits in the Informational Exceptions Control mode page (1Ch). Subject to the reporting method. For 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 appropriate 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.

6.3.5 Thermal monitor

Pulsar XT.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. (i.e., the reference temperature 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. Administrators 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 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 7) can be used to set this trip point. The default value for this drive is listed in the table, however, applications can set it to any value in the range defined. If a temperature is specified that is greater than the maximum allowed in this field, the temperature is rounded down to the maximum allowed. A sense code is sent to the host to indicate the rounding of the parameter field.

Table 7: Temperature Log Page (0Dh)

Parameter Code	Description		400GB	200/100GB
0000h	Primary Temperature		Drive Temperature	
0001h	Reference Temperature	Default Setting	70°C	65°C
		Changeable Range	0 to 70°C	0 to 65°C

6.3.6 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 contents, but does some fundamental tests and scans portions of the media.

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

6.3.6.1 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 drive parameters are not modified to test the drive more stringently, and the recovery capabilities 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 recovery processes required to recover the data.

The following conditions are considered DST failure conditions:

- Read error after recovery attempts are exhausted
- Write error after recovery attempts are exhausted

Recovered errors will not be reported as diagnostic failures.

6.3.6.2 Implementation

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

6.3.6.2.1 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 UNIT, or another DST. It is the responsibility of the host application to determine the "not ready" cause.

6.3.6.2.2 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.

6.3.6.2.3 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 two segments: an electrical 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 contents, 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 "diagnostic failed" 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. The read operation tests the media contents. The integrity of the media is checked through a read/verify scan of the media.

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

6.3.6.2.4 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.

6.3.6.2.5 Abort

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

Applications 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).

6.3.7 Product warranty

Beginning on the date of shipment to the customer and continuing for the period specified in the purchase contract, Seagate warrants that each product (including components and subassemblies) that fails to function properly under normal use due to defect in materials or workmanship or due to nonconformance to the applicable specifications will be repaired or replaced, at Seagate's option and at no charge to the customer, if returned by customer at customer's expense to Seagate's designated facility in accordance with Seagate's warranty procedure. Seagate will pay for transporting the repair or replacement item to the customer. For more detailed warranty information, refer to the standard terms and conditions of purchase for Seagate products stated in purchase documentation.

The remaining warranty for a particular drive can be determined by calling Seagate Customer Service at 1-800-468-3472. Customers can also determine remaining warranty using the Seagate web site (www.seagate.com). The drive serial number is required to determine remaining warranty information.

Shipping

When transporting or shipping a drive, use only a Seagate-approved container. Keep the 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.

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.

Storage

The maximum recommended storage period for the drive in a non-operational environment is 90 days. Drives should be stored in the original unopened Seagate shipping packaging when ever possible. Once the drive is removed from the Seagate original packaging the recommended maximum period between drive operation cycles is 30 days. During any storage period the drive non-operational temperature, humidity, wet bulb, atmospheric conditions, shock, vibration, magnetic and electrical field specifications should be followed. (see Section 7.0)

7.0 Physical/electrical specifications

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

7.1 Power specifications

The drive receives DC power (+5V and +12V) through the standard SAS interface.

7.1.1 Power consumption

Power requirements for the drives are listed in the tables beginning on page 23. Typical power measurements are based on an average of drives tested, under nominal conditions, using +5V and +12V input voltage at 60°C ambient temperature.

- **Startup power**

Startup power is measured from the time of power-on to the time that the drive reaches operating condition and can process media access commands.

- **Peak operating mode**

During peak operating mode, the drive is tested in various read and write access patterns to simulate the worst-case power consumption.

- **Idle mode power**

Idle mode power is measured with the drive powered up and ready for media access commands, with no media access commands having been received from the host.

7.2 AC power requirements

None.

7.3 DC power requirements

Table 8: 400GB standard model DC power requirements

Parameter		400GB (6.0Gb)			
Regulation			±5%	±5%	
Voltage			+5V	+12V	
			Current (A)	Current (A)	Power (W)
Average idle current		DC \bar{x}	0.44	0.31	5.92
Maximum starting current					
	(peak DC) DC	3 σ	1.15	0.45	
	(peak AC) AC	3 σ	1.21	0.63	
Delayed start (max) DC		3 σ	0.44	0.31	5.92
Peak operating current (random read):					
	Typical DC	DC \bar{x}	0.47	0.36	6.67
	Maximum DC	3 σ	0.48	0.37	6.84
	Maximum (peak) DC	3 σ	0.47	0.46	
Peak operating current (random write)					
	Typical DC	DC \bar{x}	0.45	0.42	7.29
	Maximum DC	3 σ	0.46	0.43	7.46
	Maximum (peak) DC	3 σ	0.60	0.53	
Peak operating current (sequential read)					
	Typical DC	DC \bar{x}	0.58	0.44	8.18
	Maximum DC	3 σ	0.59	0.45	8.35
	Maximum (peak) DC	3 σ	0.64	0.52	
Peak operating current (sequential write)					
	Typical DC	DC \bar{x}	0.55	0.53	9.11
	Maximum DC	3 σ	0.56	0.54	9.28
	Maximum (peak) DC	3 σ	0.57	0.63	

Table 9: 200GB standard model DC power requirements

Parameter		200GB (6.0Gb)			
Regulation			±5%	±5%	
Voltage			+5V	+12V	
			Current (A)	Current (A)	Power (W)
Average idle current		DC \bar{x}	0.40	0.17	4.04
Maximum starting current					
	(peak DC) DC	3 σ	1.13	0.30	
	(peak AC) AC	3 σ	1.20	0.64	
Delayed start (max) DC		3 σ	0.43	0.19	4.43
Peak operating current (random read):					
	Typical DC	DC \bar{x}	0.47	0.22	4.99
	Maximum DC	3 σ	0.50	0.22	5.14
	Maximum (peak) DC	3 σ	0.77	0.55	
Peak operating current (random write)					
	Typical DC	DC \bar{x}	0.45	0.27	5.49
	Maximum DC	3 σ	0.49	0.28	5.81
	Maximum (peak) DC	3 σ	0.71	0.74	
Peak operating current (sequential read)					
	Typical DC	DC \bar{x}	0.54	0.27	5.94
	Maximum DC	3 σ	0.58	0.28	6.26
	Maximum (peak) DC	3 σ	0.79	0.55	
Peak operating current (sequential write)					
	Typical DC	DC \bar{x}	0.52	0.35	6.80
	Maximum DC	3 σ	0.56	0.36	7.12
	Maximum (peak) DC	3 σ	0.79	0.85	

Table 10: 100GB standard model DC power requirements

Parameter		100GB (6.0Gb)			
Regulation			±5%	±5%	
Voltage			+5V	+12V	
			Current (A)	Current (A)	Power (W)
Average idle current		DC \bar{x}	0.40	0.17	4.04
Maximum starting current					
	(peak DC) DC	3 σ	1.09	0.31	
	(peak AC) AC	3 σ	1.16	0.73	
Delayed start (max) DC		3 σ	0.43	0.18	4.31
Peak operating current (random read):					
	Typical DC	DC \bar{x}	0.46	0.21	4.82
	Maximum DC	3 σ	0.49	0.23	5.21
	Maximum (peak) DC	3 σ	0.80	0.72	
Peak operating current (random write)					
	Typical DC	DC \bar{x}	0.45	0.25	5.25
	Maximum DC	3 σ	0.49	0.26	5.57
	Maximum (peak) DC	3 σ	0.71	0.54	
Peak operating current (sequential read)					
	Typical DC	DC \bar{x}	0.54	0.28	6.06
	Maximum DC	3 σ	0.58	0.29	6.38
	Maximum (peak) DC	3 σ	0.80	0.58	
Peak operating current (sequential write)					
	Typical DC	DC \bar{x}	0.53	0.36	6.97
	Maximum DC	3 σ	0.57	0.37	7.29
	Maximum (peak) DC	3 σ	0.82	0.82	

- [1] Measured with average reading DC ammeter. Instantaneous +12V current peaks will exceed these values. Power supply at nominal voltage. N (number of drives tested) = 6, 60 Degrees C ambient.
- [2] For +12 V, a -10% tolerance is allowed during initial start but must return to ±5% before reaching ready state. The ±5% must be maintained after the drive signifies that its power-up sequence has been completed and that the drive is able to accept selection by the host initiator.
- [3] See +12V current profile in Figure 4 (for 400GB models) and Figure 5 (for 200GB models).
- [4] See +12V current profile in Figure 6 (for 100GB models).
- [5] This condition occurs after OOB and Speed Negotiation completes but before the drive has received the Notify Spinup primitive.
- [6] See paragraph 7.3.1, "Conducted noise immunity." Specified voltage tolerance includes ripple, noise, and transient response.

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 start, are measured after a 10-minute warm up.

7.3.1 Conducted noise immunity

Noise is specified as a periodic and random distribution of frequencies covering a defined frequency. 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 = 450 mV pp from 100 Hz to 100 KHz.
250 mV pp from 100 KHz to 20 MHz.
150 mV pp from 20 MHz to 80 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 and +5V current profiles for the Pulsar 10K.4 drives are shown below.




Figure 1. Current profiles for 400GB models




Figure 2. Current profiles for 200GB models




Figure 3. Current profiles for 100GB models

7.4 Power dissipation

400GB models in 6Gb operation

Typical power dissipation under idle conditions in 6Gb operation is 5.92 watts 20.20 BTUs per hour).

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




Figure 4. 400GB (at 6Gb) DC current and power vs. input/output operations per second

200GB models in 6Gb operation

Typical power dissipation under idle conditions in 6Gb operation is 4.04 watts (13.79 BTUs per hour).

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




Figure 5. 200GB (at 6Gb) DC current and power vs. input/output operations per second

100GB models in 6Gb operation

Typical power dissipation under idle conditions in 6Gb operation is 4.04 watts (13.79 BTUs per hour).

To obtain operating power for typical random write operations, refer to the following I/O rate curve (see Figure 6). 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.

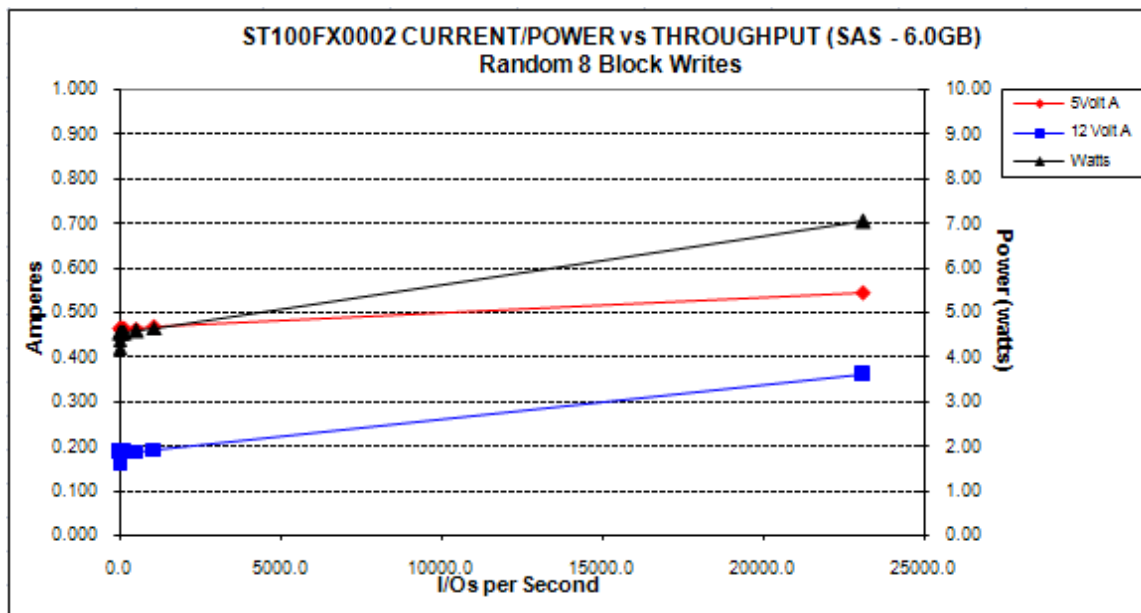


Figure 6. 100GB (at 6Gb) DC current and power vs. input/output operations per second

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).

Note. To maintain optimal performance drives should be run at nominal case temperatures.

7.5.1 Temperature

a. Operating

The drive meets the operating specifications over a 32°F to 140°F (0°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.

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 122°F (50°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 7.5). To confirm that the required cooling is provided, 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. See Figure 7 and 8 for temperature checkpoint.

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 7. Temperature check point location - 15mm drives




Figure 8. Temperature check point location - 7mm drives

Note. Images may not represent actual product, for reference only.

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

–200 to +10,000 feet (–60.96 to +3048 meters)

b. Non-operating

–200 to +40,000 feet (–60.96 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 9, and in accordance with the restrictions of Section 10.3.

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:

- 1000 Gs at a maximum duration of 0.5ms (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.

Note. This specification does not cover connection issues that may result from testing at this level.

b. 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 the three values below, shall not exhibit device damage or performance degradation.

- 1000 Gs at a maximum duration of 0.5ms (half sinewave)

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

c. Packaged

Seagate finished drive bulk packs are designed and tested to meet or exceed applicable ISTA and ASTM standards. Volume finished drives will be shipped from Seagate factories on pallets to minimize freight costs and ease material handling. Seagate finished drive bulk packs may be shipped individually. For less than full shipments, instructions are printed on the bulk pack carton for minimum drive quantities and proper drive placement.




Figure 9. Recommended mounting

Note. Image may not represent actual product, for reference only.

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 vibration:

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

Operating normal translational random flat profile

20 - 2000 Hz (translational random flat profile) 16.3 GRMS

Note. This specification does not cover connection issues that may result from testing at this level.

b. Operating—abnormal

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

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

20 - 2000 Hz (translational random flat profile) 16.3 GRMS

Note. This specification does not cover connection issues that may result from testing at this level.

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 vibration.

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

Non-operating translational random flat profile

20 - 2000 Hz (translational random flat profile) 16.3 GRMS

7.5.5 Air cleanliness

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

7.5.6 Corrosive environment

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

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

7.5.7 Electromagnetic susceptibility

See Section 3.1.1.1.

7.6 Mechanical specifications

Refer to Figure 10 or 11 for detailed mounting configuration dimensions. See Section 10.3, "Drive mounting."

Weight: 0.441 pounds 200 grams

Note. These dimensions conform to the Small Form Factor Standard documented in SFF-8201 and SFF-8223 found at www.sffcommittee.org.




Figure 10. Mounting configuration dimensions (400GB models)

Weight: 0.220 pounds 100 grams

Note. These dimensions conform to the Small Form Factor Standard documented in SFF-8201 and SFF-8223 found at www.sffcommittee.org




Figure 11. Mounting configuration dimensions (200 & 100GB models)

8.0 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 the appropriate command payloads.

8.1 Data encryption

Encrypting drives use one in-line encryption engine for each port, employing AES-256 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 8.5).

8.2 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.

8.2.1 Admin SP

The Admin SP allows the drive's owner to enable or disable firmware download operations (see Section 8.4). Access to the Admin SP is available using the SID (Secure ID) password or the MSID (Makers Secure ID) password.

8.2.2 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).

8.2.3 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.

8.3 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.

8.4 Drive locking

In addition to changing the passwords, as described in Section 8.2.3, 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 allowing the drive to accept any attempt to download new firmware. The drive owner must use the SID credential to lock the firmware download port before firmware updates will be rejected.

8.5 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 8.6) or the password when required. The bands shall be aligned to 4KB LBA boundaries.

8.6 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 redispositioned.

8.7 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.

8.8 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.

8.9 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 14:

- SECURITY PROTOCOL OUT (B5h)
- SECURITY PROTOCOL IN (A2h)

8.10 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.

9.0 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 9.4)
- Auto-Reallocation (see Section 9.5)

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

9.1 Drive internal defects/errors

During the initial drive manufacturing test 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). At factory format time, these known defects are also deallocated, that is, marked as retired 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*.

The drive uses a vendor unique format to report defects via the READ DEFECT DATA command pending T10 standardization of a format for Solid State Devices. This format defect type is defined as 110b in the SCSI FORMAT UNIT command. The definition of the 110b format is defined in the following table.

Table 11: SSD Physical format address descriptor

Bit Byte	7	6	5	4	3	2	1	0
0	(MSB) MEDIA ID (LSB)							
1								
2	CHANNEL							
3	DIE							
4	(MSB) BLOCK (LSB)							
5								
6	RESERVED							
7	VENDOR UNIQUE							

The MEDIA ID field contains an identifier for the flash controller for devices that utilize more than one flash controller.

The CHANNEL field contains the channel number within the corresponding Flash Controller.

The DIE field contains the die number within channel.

The BLOCK field contains the block number within the die.

The VENDOR UNIQUE field may contain vendor unique information.

9.2 Drive error recovery procedures

When an error occurs during drive operation, the drive performs error recovery procedures to attempt to recover the data. The error recovery procedures used are not user changeable.

9.3 SAS system errors

Information on the reporting of operational errors across the interface is given in the *SAS Interface Manual*. The SSP Response returns information to the host about numerous kinds of errors. 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.

9.4 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 reads across the entire addressable space 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 and reallocated.

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 BMS scan is performed after 500ms of idle time. Other features that normally use idle time to function will function normally because BMS functions for bursts of 500ms 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 1ms.

9.5 Auto-Reallocation

Auto-Reallocation allows the drive to reallocate unreadable locations on a subsequent write command if the recovery process deems the location to be defective. The drive performs auto-reallocation on every WRITE command. With each write to a Logical LBA, the drive writes the data to a different physical media location. Physical locations that return unrecoverable errors are retired during future WRITE attempts and associated recovery process.

This is in contrast to the system having to use the REASSIGN BLOCKS command to reassign a location that was unreadable and then generate a WRITE command to rewrite the data. This operation requires that AWRE and ARRE are enabled—this is the default setting from the Seagate factory.

9.6 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 LBA's 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.

9.6.1 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 of Protection Information (e.g. 0xFFFFFFFF) 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.

9.6.2 Setting and determining the current Type Level

A drive is initialized to a type of PI by using the FORMAT UNIT 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. A drive can only be formatted to a single PI Type. It can be changed at anytime to a new Type but requires a FORMAT UNIT command 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 UNIT CDB command: 04 90 00 00 00 00, parameter data: 00 A0 00 00

Type 2 PI FORMAT UNIT CDB command: 04 D0 00 00 00 00, parameter data: 00 A0 00 00

9.6.3 Identifying a Protection Information drive

The Standard INQUIRY data 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.

10.0 Installation

Pulsar XT.2 drive installation is a plug-and-play process. There are no jumpers 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 the drive must be mounted. 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. The drive can be mounted in any orientation.

Note. SAS drives are designed to be attached to the host system without I/O or power cables. If the intent is to use the drive in a non-backplane host system, connecting the drive using high-quality cables is acceptable as long as the I/O cable length does not exceed 10 meters (32.8 feet).

Slide the carrier or tray into the appropriate bay in the host system using the instructions provided by the host system. This connects the drive directly to the system's SAS connector. The SAS connector is normally located on a SAS backpanel. See Section 11.4.1 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. Reformatting the drive is only required if the application requires a different logical block size.




Figure 12. Physical interface

10.1 Drive orientation

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


10.2 Cooling

Cabinet cooling must be designed by the customer so that the temperature of 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 assembly. The system designer 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. 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 13. 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.



Note. Air flows in the direction shown (back to front) or in reverse direction (front to back)



Note. Air flows in the direction shown or in reverse direction (side to side)

Figure 13. Air flow

Note. Image may not represent actual product, for reference only.

10.3 Drive mounting

Mount the drive using the bottom or side mounting holes. If mounting 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:

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

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.

10.4 Grounding

Signal ground (PCBA) and case 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 drive with no electrically isolating shock mounts. If it is desired for the system chassis to not be connected to the drive 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 designers 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.

11.0 Interface requirements

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

11.1 SAS features

This section lists the SAS-specific features supported by Pulsar XT.2 drives.

11.1.1 Task management functions

Table 12 lists the SAS task management functions supported.

Table 12: SAS task management functions supported

Task name	Supported
Abort Task	Yes
Abort task set	Yes
Clear ACA	Yes
Clear task set	Yes
I_T Nexus Reset	Yes
Logical Unit Reset	Yes
Query Task	Yes
Query Task Set	Yes
Query Asynchronous Event	Yes

11.1.2 Task management responses

Table 13 lists the SAS response codes returned for task management functions supported.

Table 13: 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

11.2 Dual port support

Pulsar XT.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 Pulsar XT.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.

11.3 SCSI commands supported

Table 14 lists the SCSI commands supported by Pulsar XT.2 drives.

Table 14: Supported commands

Command name	Command code	Supported
CHANGE DEFINITION	40h	N
FORMAT UNIT [1]	04h	Y
DPRY bit supported		N
DCRT bit supported		Y
STPF bit supported		Y
IP bit supported		Y
DSP bit supported		Y
IMMED bit supported		Y
VS (vendor specific)		N
INQUIRY	12h	Y
Block Limits page (B0h)		Y
Block Device Characteristics page (B1h)		Y
Date Code page (C1h)		Y
Device Behavior page (C3h)		Y
Device Identification page (83h)		Y
Extended Inquiry Data page (86h)		Y
Firmware Numbers page (C0h)		Y
Jumper Settings page (C2h)		N
Power Conditions page (8Ah)		Y
Supported Vital Product Data page (00h)		Y
Thin Provisioning page (B2h)		Y
Unit Serial Number page (80h)		Y
Vendor Unique page (D1h)		Y
Vendor Unique page (D2h)		Y
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
LOG SENSE	4Dh	Y
Application Client Log page (0Fh)		N
Background Scan Results log page (15h)		Y
Buffer Over-run/Under-run page (01h)		N
Cache Statistics page (37h)		Y
Factory Log page (3Eh)		Y
Information Exceptions Log page (2Fh)		Y

Table 14: Supported commands

Command name	Command code	Supported
Last n Deferred Errors or Asynchronous Events page (0Bh)		N
Last n Error Events page (07h)		N
Non-medium Error page (06h)		Y
Pages Supported list (00h)		Y
Protocol-Specific Port log pages (18h)		Y
Read Error Counter page (03h)		Y
Read Reverse Error Counter page (04h)		N
Self-test Results page (10h)		Y
Solid State Media log page (11h)		Y
Start-stop Cycle Counter page (0Eh)		Y
Temperature page (0Dh)		Y
Vendor Unique page (3Ch)		Y
Verify Error Counter page (05h)		Y
Write error counter page (02h)		Y
MODE SELECT (6) (same pages as MODE SENSE (6))	15h	Y [3]
MODE SELECT (10) (same pages as MODE SENSE (6))	55h	Y
MODE SENSE (6)	1Ah	Y [3]
Caching Parameters page (08h)		Y
Control Mode page (0Ah)		Y
Disconnect/Reconnect (02h)		Y
Error Recovery page (01h)		Y
Format page (03h)		N
Information Exceptions Control page (1Ch)		Y
Background Scan mode subpage (1Ch/01h)		Y
Notch and Partition Page (0Ch)		N
Protocol-Specific LUN mode page (18h)		Y
Protocol-Specific Port page (19h)		Y
Power Condition page (1Ah)		Y
Rigid Disc Drive Geometry page (04h)		N
Unit Attention page (00h)		Y
Verify Error Recovery page (07h)		Y
Xor Control page (10h)		N
MODE SENSE (10) (same pages as MODE SENSE (6))	5Ah	Y
PERSISTENT RESERVE IN	5Eh	Y
PERSISTENT RESERVE OUT	5Fh	Y
PRE-FETCH (10)	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

Table 14: Supported commands

Command name	Command code	Supported
READ (32)	7Fh/0009h	Y
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 (10)	3Eh	Y (non-SED drives only)
READ LONG (16)	9Eh/11h	Y (non-SED drives only)
REASSIGN BLOCKS	07h	Y
RECEIVE DIAGNOSTIC RESULTS	1Ch	Y
Supported Diagnostics pages (00h)		Y
Translate page (40h)		Y
RELEASE (6)	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 (6)	16h	Y
3rd Party Reserve		Y
Extent Reservation		N
RESERVE (10)	56h	Y
3rd Party Reserve		Y
Extent Reservation		N
REZERO UNIT	01h	Y
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)		N
START UNIT/STOP UNIT	1Bh	Y
SYNCHRONIZE CACHE	35h	Y
SYNCHRONIZE CACHE (16)	91h	Y
TEST UNIT READY	00h	Y
UNMAP	42H	Y
VERIFY (10)	2Fh	Y
BYTCHK bit		Y
VERIFY (12)	AFh	N
VERIFY (16)	AFh	Y

Table 14: Supported commands

Command name	Command code	Supported
VERIFY (32)	7Fh/000Ah	Y
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	Y
WRITE AND VERIFY (10)	2Eh	Y
DPO bit		Y
WRITE AND VERIFY (12)	AEh	N
WRITE AND VERIFY (16)	8Eh	Y
WRITE AND VERIFY (32)	7Fh/000Ch	Y
WRITE BUFFER (modes 0, 2, supported)	3Bh	Y (non-SED drives only)
WRITE BUFFER	3Bh	
Firmware Download option (modes 5, 7, Ah and Bh) [2]		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	Y
XDREAD	52h	N
XDWRITE	50h	N
XPWRITE	51h	N

- [1] Pulsar XT.2 drives can format to 512, 520, 524, 528, 4096, 4160, 4192 and 4224 bytes per logical block.
- [2] **Warning.** Power loss during a firmware upgrade 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.

11.3.1 INQUIRY data

Table 15 lists the INQUIRY command data that the drive should return to the initiator per the format given in the *SAS Interface Manual*.

Table 15: Pulsar XT.2 INQUIRY data

Bytes	Data (hex)																
0-15	00	00	xx**	12	8B	01	PP	02	53	45	41	47	41	54	45	20	Vendor ID
16-31	[53	54	34	30	30	46	58	30	30	31	32]	20	20	20	20	20	Product ID
32-47	R#	R#	R#	R#	S#	S#	S#	S#	S#	S#	S#	S#	00	00	00	00	
48-63	00	00	00	00	00	00	00	00	00	00	00	00	00	00	00	00	
64-79	00	00	00	00	00	00	00	00	00	00	00	00	00	00	00	00	
80-95	00	00	00	00	00	00	00	00	00	00	00	00	00	00	00	00	
96-111	00	43	6F	70	79	72	69	67	68	74	20	28	63	29	20	32*	*Copyright
112-127	30*	31*	31*	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 ST400FX0012.
Refer to the values below for the values of bytes 16 through 26 for a particular model:

ST400FX0002	53	54	34	30	30	46	58	30	30	30	32
ST200FX0002	53	54	32	30	30	46	58	30	30	30	32
ST100FX0002	53	54	31	30	30	46	58	30	30	30	32

11.3.2 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 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 16, 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.

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 16: MODE SENSE data for 400GB drives

MODE DATA HEADER:

01 3e 00 10 01 00 00 10

BLOCK DESCRIPTOR:

00 00 00 00 2e 93 90 b0 00 00 00 00 00 00 02 00

MODE PAGES:

DEF 81 0a c0 01 5a 00 00 00 0b 00 ff ff
 CHG 81 0a 38 00 00 00 00 00 ff 00 00 00

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 87 0a c0 01 5a 00 00 00 00 00 ff ff
 CHG 87 0a 38 00 00 00 00 00 00 ff ff

DEF 88 12 14 00 ff ff 00 00 ff ff ff ff a0 20 00 00 00 00 00
 CHG 88 12 a5 00 00 00 ff ff ff ff 00 00 01 00 00 00 00 00

DEF 8a 0a 00 00 00 80 00 00 00 00 7f ff
 CHG 8a 0a 07 f0 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
 CHG 99 0e 50 00 ff ff ff ff ff ff 00 00 00 00 00

DEF 59 01 00 64 00 06 00 02 00 00 00 00 14 1a 0e 00 50 00 c5 00 00 1a b2 81 50 06 05
 b0 01 49 c2 60 02 00 00 00 00 00 00 88 aa 00 00 00 00 00 00 00 00 00 00 00 00
 00 00 00 01 00 00 10 29 0e 00 50 00 c5 00 00 1a b2 82 50 06 05 b0 00 00 fe e4 06
 00 00 00 00 00 00 00 88 aa 00 00 00 00 00 00 00 00 00 00 00 00 00 00
 CHG 59 01 00 64 00
 00
 00
 00

DEF 59 03 00 2c 00 06 00 02 00 00 00 10 80 ac 00 01 80 ac 00 01 80 bc 00 00 00 00 1a
 00 00 01 00 10 80 ac 00 01 80 ac 00 01 00 00 00 00 00 00 09 00
 CHG 59 03 00 2c 00
 00

DEF 9a 0a 00 02 00 00 00 05 00 00 8c a0
 CHG 9a 0a 00 02 ff ff ff ff 00 00 00 00

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

DEF dc 01 00 0c 01 00 01 50 00 18 00 00 00 00 00
 CHG dc 01 00 0c 00 00 ff ff ff ff 00 00 00 00 00

DEF 80 06 00 80 00 00 00 00
 CHG 80 06 b7 80 00 00 00 00

Table 17: MODE SENSE data for 200GB drives

MODE DATA HEADER:

01 3e 00 10 01 00 00 10

BLOCK DESCRIPTOR:

00 00 00 00 17 49 f1 b0 00 00 00 00 00 00 02 00

MODE PAGES:

DEF 81 0a c0 01 5a 00 00 00 0b 00 ff ff

CHG 81 0a 38 00 00 00 00 00 ff 00 00 00

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 87 0a c0 01 5a 00 00 00 00 00 ff ff

CHG 87 0a 38 00 00 00 00 00 00 ff ff

DEF 88 12 14 00 ff ff 00 00 ff ff ff ff a0 20 00 00 00 00 00

CHG 88 12 a5 00 00 00 ff ff ff ff 00 00 01 00 00 00 00 00

DEF 8a 0a 00 00 00 80 00 00 00 00 7f ff

CHG 8a 0a 07 f0 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 59 01 00 64 00 06 00 02 00 00 00 00 14 1a 0e 00 50 00 c5 00 00 1a 4b 59 50 06 05
 b0 01 49 c2 60 02 00 00 00 00 00 00 88 aa 00 00 00 00 00 00 00 00 00 00 00 00
 00 00 00 01 00 00 10 29 0e 00 50 00 c5 00 00 1a 4b 5a 50 06 05 b0 00 00 fe e4 06
 00 00 00 00 00 00 00 88 aa 00 00 00 00 00 00 00 00 00 00 00 00 00 00

CHG 59 01 00 64 00
 00
 00
 00

DEF 59 03 00 2c 00 06 00 02 00 00 00 10 80 ac 00 01 80 ac 00 01 80 bc 00 00 00 00 1a
 00 00 01 00 10 80 ac 00 01 80 ac 00 01 00 00 00 00 00 00 09 00

CHG 59 03 00 2c 00
 00

DEF 9a 0a 00 02 00 00 00 05 00 00 8c a0

CHG 9a 0a 00 02 ff ff ff ff 00 00 00 00

DEF 9c 0a 10 00 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 00 01 50 00 18 00 00 00 00 00 00

CHG dc 01 00 0c 00 00 ff ff ff ff 00 00 00 00 00 00

DEF 80 06 00 80 00 00 00 00

CHG 80 06 b7 80 00 00 00 00

Table 18: MODE SENSE values for 100GB drives

MODE DATA HEADER:

01 3e 00 10 01 00 00 10

BLOCK DESCRIPTOR:

00 00 00 00 0b a5 22 30 00 00 00 00 00 00 02 00

MODE PAGES:

DEF 81 0a c0 01 5a 00 00 00 0b 00 ff ff
 CHG 81 0a 38 00 00 00 00 00 ff 00 00 00

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 87 0a c0 01 5a 00 00 00 00 00 ff ff
 CHG 87 0a 38 00 00 00 00 00 00 ff ff

DEF 88 12 14 00 ff ff 00 00 ff ff ff ff a0 20 00 00 00 00 00
 CHG 88 12 a5 00 00 00 ff ff ff ff 00 00 01 00 00 00 00 00

DEF 8a 0a 00 00 00 80 00 00 00 00 7f ff
 CHG 8a 0a 07 f0 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
 CHG 99 0e 50 00 ff ff ff ff ff ff 00 00 00 00 00

DEF 59 01 00 64 00 06 00 02 00 00 00 00 14 1a 0e 00 50 00 c5 00 00 1a 3a ed 50 06 05
 b0 01 49 c2 60 02 00 00 00 00 00 00 88 aa 00 00 00 00 00 00 00 00 00 00 00 00
 00 00 00 01 00 00 10 29 0e 00 50 00 c5 00 00 1a 3a ee 50 06 05 b0 00 00 fe e4 06
 00 00 00 00 00 00 00 88 aa 00 00 00 00 00 00 00 00 00 00 00 00 00 00
 CHG 59 01 00 64 00
 00
 00
 00

DEF 59 03 00 2c 00 06 00 02 00 00 00 10 80 ac 00 01 80 ac 00 01 80 bc 00 00 00 00 1a
 00 00 01 00 10 80 ac 00 01 80 ac 00 01 00 00 00 00 00 00 09 00
 CHG 59 03 00 2c 00
 00

DEF 9a 0a 00 02 00 00 00 05 00 00 8c a0
 CHG 9a 0a 00 02 ff ff ff ff ff ff ff ff

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

DEF dc 01 00 0c 01 00 01 50 00 18 00 00 00 00 00
 CHG dc 01 00 0c 00 00 ff ff ff ff 00 00 00 00 00

DEF 80 06 00 80 00 00 00 00
 CHG 80 06 b7 80 00 00 00 00

11.4 Miscellaneous operating features and conditions

Table 19 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 19: Miscellaneous features

Supported	Feature or condition
N	Automatic contingent allegiance
N	Asynchronous event notification
Y	Segmented caching
N	Zero latency read
Y	Queue tagging (up to 128 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

Table 20: Miscellaneous status

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

11.4.1 SAS physical interface

Figure 14 shows the location of the SAS device connector J1. Figures 15 and 16 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 14. Physical interface




Figure 15. SAS device plug dimensions




Figure 16. SAS device plug dimensions (detail)

11.4.2 Physical characteristics

This section defines physical interface connector.

11.4.3 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 15 and 16.

11.4.4 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.

11.4.5 Pin descriptions

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

Table 21: SAS pin descriptions

Pin	Signal name	Signal type
S1	Port A Ground	
S2*	+Port A_in	Diff. input pair
S3*	-Port A_in	
S4	Port A Ground	
S5*	-Port A_out	Diff output pair
S6*	+Port A_out	
S7	Port A Ground	
S8	Port B Ground	
S9*	+Port B_in	Diff. input pair
S10*	-Port B_in	
S11	Port A Ground	
S12*	-Port B_out	Diff output pair
S13*	+Port B_out	
S14	Port B Ground	

Pin	Signal name	Signal type
P1*	NC (reserved 3.3Volts)	
P2*	NC (reserved 3.3Volts)	
P3	NC (reserved 3.3Volts)	
P4	Ground	
P5	Ground	
P6	Ground	
P7	5 Volts charge	
P8*	5 Volts	
P9*	5 Volts	
P10	Ground	
P11*	Ready LED	Open collector out
P12	Ground	
P13	12 Volts charge	
P14*	12 Volts	
P15*	12 Volts	

* - Short pin to support hot plugging

NC - No connection in the drive.

11.4.6 SAS transmitters and receivers

A typical SAS differential copper transmitter and receiver pair is shown in Figure 17. The receiver is AC coupling to eliminate ground shift noise.




Figure 17. SAS transmitters and receivers

11.4.7 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.

11.5 Signal characteristics

This section describes the electrical signal characteristics of the drive's input and output signals. See Table 21 for signal type and signal name information.

11.5.1 Ready LED Out

The Ready LED Out signal is driven by the drive as indicated in Table 22.

Table 22: Ready LED Out conditions

Normal command activity	LED status	
	0	1
Ready LED Meaning bit mode page 19h	0	1
Drive stopped, not ready, and no activity	Off	Off
Drive stopped, not ready, and activity (command executing)	On	On
Drive started, ready, and no activity	On	Off
Drive started, ready, and activity (command executing)	Off	On
Drive transitioning from not-ready state to ready state or the reverse.	Blinks steadily (50% on and 50% off, 0.5 seconds on and off for 0.5 seconds)	
FORMAT UNIT in progress,	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 23 for the output characteristics of the LED drive signals.

Table 23: 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	$I_{OL} = 15\text{ mA}$	$0 \leq V_{OL} \leq 0.225\text{ V}$

11.5.2 Differential signals

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

Table 24 defines the general interface characteristics

Table 24: 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

11.6 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.

11.7 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 44
AC coupling 59
AC power requirements 22
ACA active status 55
ACA active, faulted initiator status 55
access time
 average latency 11
 average typical 11
 page to page typical 11
active LED Out signal 59
Admin SP 35
AES-128 data encryption 35
air cleanliness 32
air flow 42
 illustrated 42
air inlet 42
altitude 30
ambient 30
ambient temperature 22
ANSI documents
 SCSI 5
 Serial Attached SCSI 5
asynchronous event notification 55
audible noise 3
automatic contingent allegiance 55
Auto-Reallocation 39
average idle current 23, 24, 25

B

Background Media Scan 39
backpanel 58
Band 0 36
BandMasterX 36
BMS 39
busy status 55

C

cache control 13
caching write data 13
Canadian Department of Communications 3
capacity
 unformatted 11
case 43

CBC 35
CE Marking 4
check condition status 55
China RoHS directive 5
Cipher Block Chaining 35
class B limit 3
clear ACA function 44
clear task set function 44
commands supported 46
condensation 30
condition met/good status 55
connector
 illustrated 58
 requirements 58
cooling 42
CRC
 error 15
Cryptographic erase 37
Current profiles 26
customer service 21

D

Data Bands 36
data bands 35
data block size
 modifying the 8
Data encryption 35
Data Encryption Key 35
Data Retention 14, 16
DATA Returned for Thin Provisioned LBA 10
DC power 22, 58
 requirements 23
Decision Feedback Equalizer 60
decrypt 35
default MSID password 36
defects 38
deferred error handling 55
DEK 35
description 6
DFE 60
dimensions 33
drive 32
drive characteristics 11
drive failure 16
Drive Locking 36
drive mounting 33, 43
drive select 58
dual port support 45

E

electrical
 description of connector 58
 signal characteristics 59
 specifications 22

- electromagnetic compatibility 3
- electromagnetic susceptibility 32
- EMI requirements 3
- encryption engine 35
- encryption key 36
- Endurance 14
- Endurance Management 15
- environment 42
- environmental
 - limits 29
 - requirements 14
- environmental control 32
- EraseMaster 36
- error
 - management 38
 - rates 14
- Error reported to Host 10
- errors 38
- European Union Restriction of Hazardous Substances 4

F

- FCC rules and regulations 3
- features 6
 - interface 44
- feed forward equalizer 60
- FFE 60
- firmware 6
 - corruption 49
- firmware download port 36
- Format command execution time 12
- function
 - complete, code 00 44
 - not supported, code 05 44
 - reject, code 04 44

G

- Garbage Collection 15
- Global Data Band 36
- Good status 55
- gradient 30
- ground shift noise 59
- grounding 43

H

- heat removal 42
- heat source 42
- host equipment 43
- hot plugging the drive 16
- humidity 30
- humidity limits 29

I

- Identifying a PI drive 40
- Idle mode 22
- Idle mode power 22
- input voltage 22
- inquiry data 50
- installation 41
 - guide 5
- interface
 - commands supported 46
 - error rate 14
 - errors 15
 - illustrated 55
 - physical 55
 - requirements 44
- internal defects/errors 38
- internal drive characteristics 11

J

- jumpers 41

L

- LBPME bit 10
- LBPRZ bit 10
- Lifetime Endurance Management 16
- Locking SP 35, 36
- LockOnReset 36
- Logical Block Provisioning 9
- logical block size 6

M

- maintenance 14
- Makers Secure ID 35
- maximum delayed motor start 23, 24, 25
- maximum start current 23, 24, 25
- mean time between failure. See MTBF
- media description 7
- miscellaneous feature support
 - Asynchronous event notification 55
 - Automatic contingent allegiance 55
 - Deferred error handling 55
 - Parameter rounding 55
 - Queue tagging 55
 - Reporting actual retry count 55
 - Segmented caching 55
 - Zero latency read 55
- miscellaneous status support
 - ACA active 55
 - ACA active, faulted initiator 55
 - Busy 55
 - Check condition 55
 - Condition met/good 55
 - Good 55

- Reservation conflict 55
- Task set full 55
- miscorrected media data 14
- Mode sense
 - data, table 50, 52, 53
- mounting 43
 - holes 43
 - orientations 41
- mounting configuration 33
- mounting configuration dimensions 33, 34
- MSID 35, 36
- MTBF 14, 16

N

- noise
 - audible 3
- noise immunity 26
- non-operating 30, 31, 32
 - temperature 30
- non-operating vibration 32

O

- office environment 32
- operating 30, 31, 32
- option selection 58
- options 8
- out-of-plane distortion 43

P

- packaged 31
- parameter rounding 55
- password 35, 36
- passwords 36
- PCBA 43
- peak operating current 23, 24, 25
- Peak operating mode 22
- peak-to-peak measurements 26
- Performance 12
- performance characteristics 11
 - detailed 11
- performance highlights 7
- physical damage 32
- physical interface 55
- physical specifications 22
- PI Check Performed 10
- PI Check Requested 10
- PI level - Type 0 40
- PI level - Type I 40
- PI level - Type II 40
- PI level - Type III 40
- PI Levels 40
- PI Returned for Thin Provisioned LBA 10
- PI Setting 10
- pin descriptions 58

- power 59
 - dissipation 28
 - requirements, AC 22
 - requirements, DC 23
 - sequencing 26
- Power consumption 22
- power distribution 3
- Power specifications 22
- PowerChoice 22
- PowerCycle 36
- preventive maintenance 14
- PROT_EN bit 10
- protection information 40
- protection of data at rest 35

Q

- queue tagging 55

R

- radio interference regulations 3
- Random number generator 36
- read error rates 14, 38
- receivers 59
- recommended mounting 31
- reference
 - documents 5
- relative humidity 30
- reliability 7
 - specifications 14
- reliability and service 16
- repair and return information 21
- reporting actual retry count 55
- reservation conflict status 55
- resonance 30
- return information 21
- RNG 36
- RoHS 4, 5

S

- safety 3
- SAS
 - interface 58
 - physical interface 55
 - task management functions 44
- SAS documents 5
- SAS Interface Manual 3, 5
- SAS-2 Specification 60
- SCSI interface
 - commands supported 46
- Secure ID 35
- security partitions 35
- Security Protocol In 35
- Security Protocol Out 35
- segmented caching 55

- self-encrypting drives 35
- Self-Monitoring Analysis and Reporting Technology
 - 7, 17
- Serial Attached SCSI (SAS) Interface Manual 2
- shielding 3
- shipping 21
- shipping container 30
- shock 31
 - and vibration 30
- shock mount 43
- SID 35
- signal
 - characteristics 59
- single-unit shipping pack kit 8
- SMART 7, 17
- SNW-3 training gap 60
- Specification 60
- SSD Physical format address descriptor 38
- standards 3
- start/stop time 13
- Startup power 22
- support services 1
- surface stiffness
 - allowable for non-flat surface 43
- system chassis 43

T

- Task Aborted 55
- task management functions 44
 - Abort task set 44
 - Clear ACA 44
 - Clear task set 44
 - terminate task 44
- task management response codes 44
 - Function complete 00 44
 - Function not supported 05 44
 - Function reject 04 44
- task set full status 55
- TCG 35
- technical support services 1
- temperature 29, 42
 - limits 29
 - non-operating 30
 - regulation 3
 - See *also* cooling
- terminate task function 44
- Thin Provisioning 9
- transmitters 59
- transporting the drive 21
- Trusted Computing Group 35
- Type 1 PI format 40
- Type 2 PI format 40

U

- unformatted 8
- Unmap 9, 15
- Unrecoverable Errors 15
- unrecovered media data 14

V

- vibration 30, 32
- voltage 22

W

- warranty 21
- Wear Leveling 15
- Write Amplification 15

Z

- zero latency read 55



Seagate Technology LLC

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

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

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

Publication Number: 100647497, Rev. B

June 2011