



X18 SAS Product Manual

512E* models	512E* models	512E* models
Standard	Self-Encrypting	SED-FIPS 140-2
ST18000NM004J	ST18000NM005J	ST18000NM007J
ST16000NM004J	ST16000NM005J	ST16000NM007J

* Default configuration is 512E for 512E / 4KN drives.
See [Section 4.1.2](#) to Fast Format to 4KN in seconds

100865853, Rev. A
August 2020

Document Revision History

Revision	Date	Pages affected and Description of changes
Rev. A	08/06/2020	Initial release.

© 2020 Seagate Technology LLC. All rights reserved.

Publication number: 100865853, Rev. A August 2020

Seagate, Seagate Technology and the Spiral logo are registered trademarks of Seagate Technology LLC in the United States and/or other countries. Raid Rebuild and SeaTools are either trademarks or registered trademarks of Seagate Technology LLC or one of its affiliated companies in the United States and/or other countries. The FIPS logo is a certification mark of NIST, which does not imply product endorsement by NIST, the U.S., or Canadian governments. All other trademarks or registered trademarks are the property of their respective owners.

No part of this publication may be reproduced in any form without written permission of Seagate Technology LLC.

Call 877-PUB-TEK1 (877-782-8351) to request permission.

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

Contents

Seagate® Technology Support Services	5
<hr/>	
1.0 Scope	6
<hr/>	
2.0 HDD and SSD Regulatory Compliance and Safety	7
2.0.1 Regulatory Models	7
2.1 Reference documents	7
<hr/>	
3.0 General description	8
3.1 Standard features	9
3.2 Performance	9
3.3 Reliability	9
3.4 Media description	10
3.5 Formatted capacities	10
3.6 Programmable drive capacity	10
3.7 Factory-installed options	10
<hr/>	
4.0 Performance characteristics	11
4.1 Internal drive characteristics	11
4.1.1 Format command execution time	11
4.1.2 Fast Format	11
4.1.3 General performance characteristics	12
4.2 Start/stop time	12
4.3 Prefetch/multi-segmented cache control	13
4.4 Cache operation	13
4.4.1 Caching write data	14
4.4.2 Prefetch operation	14
<hr/>	
5.0 Reliability specifications	15
5.1 Error rates	15
5.1.1 Recoverable Errors	15
5.1.2 Unrecoverable Errors	15
5.1.3 Seek errors	15
5.1.4 Interface errors	15
5.2 Reliability and service	16
5.2.1 Annualized Failure Rate (AFR) and Mean Time Between Failure (MTBF)	16
5.2.2 Hot plugging the drive	16
5.2.3 S.M.A.R.T.	17
5.2.4 Thermal monitor	18
5.2.5 Drive Self Test (DST)	18
5.2.6 Product warranty	20

6.0	Physical/electrical specifications	21
6.1	PowerChoice™ power management	21
6.1.1	PowerChoice reporting methods	22
6.2	Power Balance	22
6.3	AC power requirements	22
6.4	DC power requirements	23
6.4.1	Conducted noise immunity	24
6.4.2	Power sequencing	24
6.4.3	Current profiles	25
6.5	Power dissipation	26
6.6	Environmental limits	27
6.6.1	Temperature	27
6.6.2	Humidity	27
6.6.3	Effective altitude (sea level)	27
6.6.4	Shock and Vibration	28
6.6.5	Acoustics	29
6.6.6	Air cleanliness	29
6.6.7	Corrosive environment	29
6.7	Mechanical specifications	30

7.0	About self-encrypting drives	31
7.1	Data encryption	31
7.2	Controlled access	31
7.2.1	Admin SP	31
7.2.2	Locking SP	31
7.2.3	Default password	31
7.3	Random number generator (RNG)	31
7.4	Drive locking	31
7.5	Data bands	32
7.6	Cryptographic erase	32
7.7	Authenticated firmware download	32
7.8	Power requirements	32
7.9	Supported commands	32
7.10	Sanitize - CRYPTOGRAPHIC ERASE	32
7.11	RevertSP	32

8.0	Defect and error management	33
8.1	Drive internal defects/errors	33
8.2	Drive error recovery procedures	33
8.3	SAS system errors	34
8.4	Deferred Auto-Reallocation	34
8.5	Idle Read After Write	34
8.6	Protection Information (PI)	35
8.6.1	Levels of PI	35
8.6.2	Setting and determining the current Type Level	35
8.6.3	Identifying a Protection Information drive	35

9.0	Installation	36
9.1	Drive orientation	36
9.2	Cooling	37
9.3	Drive mounting	38
9.4	Grounding	38

10.0	Interface requirements	39
10.1	SAS features	39
10.1.1	task management functions	39
10.1.2	task management responses	39
10.2	Dual port support	39
10.3	SCSI commands supported	40
10.3.1	Inquiry data	45
10.3.2	Mode Sense data	46
10.4	Miscellaneous operating features and conditions	48
10.4.1	SAS physical interface	48
10.4.2	Physical characteristics	51
10.4.3	Connector requirements	51
10.4.4	Electrical description	51
10.4.5	Pin descriptions	51
10.4.6	SAS transmitters and receivers	52
10.4.7	Power	52
10.5	Signal characteristics	52
10.5.1	Ready LED Out	52
10.5.2	Differential signals	53
10.6	SAS-3 Specification Compliance	53
10.7	Additional information	53

Seagate® Technology Support Services

For Seagate Product Support, visit: <https://www.seagate.com/support>

For Seagate Compliance, Safety, and Disposal, visit: <https://www.seagate.com/support>

For Firmware Download and Tools Download for Secure Erase, visit: <https://www.seagate.com/support/downloads/>

For information regarding online support and services, visit: <http://www.seagate.com/contacts/>

For information regarding Warranty Support, visit: <http://www.seagate.com/support/warranty-and-replacements/>

For information regarding data recovery services, visit: <http://www.seagate.com/services-software/recover/>

For Seagate OEM and Distribution partner and Seagate reseller portal, visit: <http://www.seagate.com/partners>

1.0 Scope

This manual describes Seagate® Exos® X18 SAS (Serial Attached SCSI) disk drives.

Exos X18 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.

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 designing a system which will use one of the models listed or future generation products and need further assistance, please contact the Field Applications Engineer (FAE) or our global support services group as shown in **“Seagate® Technology Support Services” on page 5.**

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

Model type	5xxE
Standard	ST18000NM004J, ST16000NM004J
Self-Encrypting Drive (SED)	ST18000NM005J, ST16000NM005J
SED-FIPS 140-2	ST18000NM007J, ST16000NM007J

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).

2.0 HDD and SSD Regulatory Compliance and Safety

For the latest regulatory and compliance information see: <https://www.seagate.com/support/> scroll to bottom of page and click the Compliance, Safety and Disposal Guide link.

2.0.1 Regulatory Models

The following regulatory model number represent all features and configurations within the series:

Regulatory Model Numbers: STL015

2.1 Reference documents

SAS Interface Manual

Seagate part number: 100293071

SCSI Commands Reference Manual

Seagate part number: 100293068

Self-Encrypting Drives Reference Manual

Seagate part number: 100515636

ANSI SAS Documents

SFF-8323	3.5" Drive Form Factor with Serial Connector
SFF-8460	HSS Backplane Design Guidelines
SFF-8470	Multi Lane Copper Connector
SFF-8482	SAS Plug Connector
INCITS 538	SCSI Protocol Layer-4 (SPL-4) Rev. 08
INCITS 534	Serial Attached SCSI (SAS-4)
INCITS 506	SCSI Block Commands-4 (SBC-4) Rev. 10
INCITS 502	SCSI Primary Commands-5 (SPC-5) Rev. 10

ANSI Small Computer System Interface (SCSI) Documents

INCITS 515 SCSI Architecture Model (SAM-5) Rev. 11

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

TCG Storage Architecture Core Specification, Rev. 1.0

TCG Storage Security Subsystem Class Enterprise Specification, Rev. 1.0

Specification for Acoustic Test Requirement and Procedures

Seagate part number: 30553-001

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

3.0 General description

Exos X18 drives provide high performance, high capacity data storage for a variety of systems including engineering workstations, network servers, mainframes, and supercomputers. The Serial Attached SCSI interface is designed to meet next-generation computing demands for performance, scalability, flexibility and high-density storage requirements.

Exos X18 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. Exos X18 drives are classified as intelligent peripherals and provide level 2 conformance (highest level) with the ANSI SCSI-1 standard. The SAS connectors, cables and electrical interface are compatible with Serial ATA (SATA), giving future users the choice of populating their systems with either SAS or SATA hard disk drives. This allows users to continue to leverage existing investment in SCSI while gaining a 12Gb/s serial data transfer rate.

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

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

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

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

Exos X18 drives decode track 0 location data from the servo data embedded on each surface to eliminate mechanical transducer adjustments and related reliability concerns.

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

Note	Seagate recommends validating the configuration with the selected HBA/RAID controller manufacturer to ensure use of full capacity is supported.
-------------	---

3.1 Standard features

Exos X18 drives have the following standard features:

- 128 - deep task set (queue)
- 256MB data buffer (See Section 4.4).
- 3.0 / 6.0/12.0 Gb Serial Attached SCSI (SAS) interface
- Drive Self Test (DST)
- Embedded servo design
- Firmware downloadable using the SAS interface
- Flawed logical block reallocation at format time
- Idle Read After Write (IRAW)
- Industry standard SFF 3.5-in dimensions
- Integrated dual port SAS controller supporting the SCSI protocol
- Jumperless configuration.
- No preventive maintenance or adjustments required
- Perpendicular recording technology
- Power Balance supported (see Section 6.2 on page 22)
- Power Save
- Programmable auto write and read reallocation
- Programmable logical block reallocation scheme
- Reallocation of defects on command (Post Format)
- SAS Power Disable
- Self diagnostics performed when power is applied to the drive
- Support for SAS expanders and fanout adapters
- Supports up to 32 initiators
- T10 Fast Format supported (see Section 4.1.2)
- User-selectable logical block sizes for 4096 native models (4096, 4160 or 4224 bytes per logical block)
- User-selectable logical block sizes for 512E (512, 520 or 528 bytes per logical block).
- Vertical, horizontal, or top down mounting

Seagate Exos X18 SAS Self-Encrypting Drive models have the following additional features:

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

3.2 Performance

- 1200MB/s maximum instantaneous data transfers.
- 7200 RPM spindle. Average latency = 4.16ms
- Adaptive seek velocity; improved seek performance
- Background processing of queue
- Programmable multi-segmentable cache buffer
- Supports start and stop commands (spindle stops spinning)

Note

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

3.3 Reliability

- 5-year warranty
- Annualized Failure Rate (AFR) of 0.35%
- Balanced low mass rotary voice coil actuator
- Incorporates industry-standard Self-Monitoring Analysis and Reporting Technology (S.M.A.R.T.)
- Mean time between failures (MTBF) of 2,500,000 hours

3.4 Media description

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

3.5 Formatted capacities

Standard OEM models are formatted to 512 bytes per block for 512 emulation drives and 4096 bytes per block for 4096 native drives. The block size is selectable at format time. Supported block sizes are 512, 520 and 528 for 512 emulation drives and 4096, 4160 and 4224 for 4096 native drives. Users having the necessary equipment may modify the data block size before issuing a format command and obtain different formatted capacities than those listed.

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

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

Sector Size	18TB with PI bytes		18TB w/o PI bytes		16TB with PI bytes		16TB w/o PI bytes	
	Decimal	Hex	Decimal	Hex	Decimal	Hex	Decimal	Hex
512	35,156,656,128	82F800000	34,441,527,128	804E00000	30,616,322,048	720E00000	31,251,759,104	746C00000
520	34,441,527,296	804E00000	33,919,336,448	7E5C00000	30,152,851,456	705400000	30,616,322,048	720E00000
528	33,919,336,448	7E5C00000	33,413,922,816	7C7A00000	29,701,963,776	6EA600000	30,152,851,456	705400000
4096	4,394,582,016	105F00000	4,363,911,168	1041C0000	3,879,206,912	E73800000	3,906,469,888	E8D800000
4160	4,305,190,912	1009C0000	4,297,064,448	100200000	3,819,700,224	E3AC00000	3,827,040,256	E41C00000
4224	4,239,917,056	FCB800000	4,232,052,736	FC4000000	3,761,766,400	E03800000	3,769,106,432	E0A800000

Note	LBA Counts for drive capacities greater than 8TB are calculated based upon the SFF-8447 standard publication. ftp://ftp.seagate.com/sff/SFF-8447.PDF
-------------	---

3.6 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.

3.7 Factory-installed options

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

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

4.0 Performance characteristics

This section provides detailed information concerning performance-related characteristics and features of Exos X18 drives.

4.1 Internal drive characteristics

Drive capacity	18TB	16TB	(formatted, rounded off value)
Read/write data heads	18	17	
Bytes/track	1,840,943	1,785,977	Bytes (average, rounded off values)
Bytes/surface	1,000,000	941,176	MB (unformatted, rounded off values)
Tracks/surface (total)	543,200	526,981	Tracks (user accessible)
Tracks/in	482,000	46,609	TPI (average)
Peak bits/in	2,470,000	2,396,251	BPI
Areal density	1146	1146	Gb/in ²
Internal data rate	2899	2812	Mb/s (max)
Disk rotation speed	7200	7200	RPM
Avg rotational latency	4.16	4.16	ms

4.1.1 Format command execution time

5xxE-byte sectors (minutes)	18TB models	16TB models
Maximum (with verify)	2930	2740
Maximum (without verify)	1465	1370

Execution time measured from receipt of the last byte of the Command Descriptor Block (CDB) to the request for a Status Byte Transfer to the Initiator (excluding connect/disconnect).

When changing sector sizes, the format times shown above may need to be increased by 30 minutes.

4.1.2 Fast Format

Drive sector size transition

- Single code to support sector sizes from 512E
- T10 fast format conversion between 512E configurations in the field.
 - Possible only if sector sizes are exact multiples of 8 & vice versa
- The selected sector size will take effect only after fast format or full format
- Drive default is 512E from the factory.
- 512E features set after Fast Format

T10 Fast Format

- Implements the fast format based on T10 Spec.
- To request Fast Format, the FFMT bits (Byte 4, Bits 1:0) should be set to 01b.
- A setting of 10b or 11b will return a check condition with 05/24 sense code (pointing to FFMT MSB in CDB).

Mode Select - Parameter list header

- Set Write buffer: // Set Block Descriptor Length = 0x08, Number of LBAs = 0xFFFFFFFF
- 00 00 00 00 00 00 08 FF FF FF FF 00 00 **02 00** // Set block size to 0512 (0x0200)
- Then Send Mode Select Command
- cdb: 55 01 00 00 00 00 00 10 00 // (SP bit = 1, Parameter list = 0x10)

Format Unit - Parameter list header

- Set Write buffer: // Set IMMED = 1
- 00 02 00 00
- Then Send Format Unit Command
- cdb: 04 14 00 00 01 00 // (FMTDAT = 1, DEFECT LIST FORMAT = 010b, FFMT = 01b)

4.1.3 General performance characteristics

Minimum sector interleave	1 to 1
Maximum Internal data rate*	2.7 Gb/s
Sustained transfer rate (for 18TB models)	103 to 258 MiB/s **
Sustained transfer rate (for 16TB models)	103 to 249 MiB/s **
SAS Interface maximum instantaneous transfer rate	1200MB/s* per port (dual port = 2400MB/s*)
Logical block sizes	
512 (default), 520 or 528. 4096, 4160 or 4224.	
Read/write consecutive sectors on a track	Yes
Flaw reallocation performance impact (for flaws reallocated at format time using the spare sectors per sparing zone reallocation scheme.)	Negligible
Average rotational latency	4.16ms

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

** MiB/s x 1.048 = MB/s

4.2 Start/stop time

Power-on to ready time is based on typical operating conditions, default full current spin-up profile, and clean shutdown prior to measurement. To ensure a clean shutdown issue a START STOP UNIT command with the START bit set to zero and the IMMED bit set to zero, so that the device will return status after the operation is completed.

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

If the drive receives a NOTIFY (ENABLE SPINUP) primitive through either port and has not received a START STOP UNIT command with the START bit equal to 0, the drive becomes ready for normal operations within 30 seconds (excluding the error recovery procedure).

If the drive receives a START STOP UNIT command with the START bit equal to 0 before receiving a NOTIFY (ENABLE SPINUP) primitive, the drive waits for a START STOP UNIT command with the START bit equal to 1. After receiving a START STOP UNIT command with the START bit equal to 1, the drive waits for a NOTIFY (ENABLE SPINUP) primitive. After receiving a NOTIFY (ENABLE SPINUP) primitive through either port, the drive becomes ready for normal operations within 30 seconds (excluding the error recovery procedure).

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

The START STOP UNIT command may be used to command the drive to stop the spindle. Stop time is 23 seconds (maximum) from removal of DC power. SCSI stop time is 23 seconds. There is no power control switch on the drive.

An unexpected power loss event, spin up at cold or hot temperature extremes may cause the drive to exceed the typical and max time to ready by 5 to 20 seconds. Extended time to ready is dependent on cache state and environmental conditions prior to the unexpected power loss and during the subsequent power on.

4.3 Prefetch/multi-segmented cache control

The drive provides a prefetch (read look-ahead) and multi-segmented cache control algorithms that in many cases can enhance system performance. Cache refers to the drive buffer storage space when it is used in cache operations. To select this feature, the host sends the Mode Select command with the proper values in the applicable bytes in page 08h. Prefetch and cache operations are independent features from the standpoint that each is enabled and disabled independently using the Mode Select command; however, in actual operation, the prefetch feature overlaps cache operation somewhat as described in sections [Table 4.4.1](#) and [Table 4.4.2](#).

All default cache and prefetch mode parameter values (Mode Page 08h) for standard OEM versions of this drive family are given in [Table 7](#).

4.4 Cache operation

Note	Refer to the <i>SAS Interface Manual</i> for more detail concerning the cache bits.
-------------	---

The buffer is divided into logical segments from which data is read and to which data is written.

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

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

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

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

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

1. The drive fetches the requested logical blocks from the disk and transfers them into a segment, and then from there to the host in accordance with the Mode Select Disconnect/Reconnect parameters, page 02h.
2. If the prefetch feature is enabled, refer to section [Table 4.4.2](#) for operation from this point.

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

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

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

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

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

The Synchronize Cache command may be used to force the drive to write all cached write data to the medium. Upon completion of a Synchronize Cache command, all data received from previous write commands will have been written to the medium. [Table 7](#) shows the mode default settings for the drive.

4.4.2 Prefetch operation

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

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

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

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

5.0 Reliability specifications

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

Seek error rate:	Less than 10 errors in 10^8 seeks
Read Error Rates ¹	
Recovered Data	Less than 10 errors in 10^{12} bits transferred (OEM default settings)
Unrecovered Data	Less than 1 sector in 10^{15} bits transferred
Miscorrected Data	Less than 1 sector in 10^{21} bits transferred
Interface error rate:	Less than 1 error in 10^{12} bits transferred
Mean Time Between Failure (MTBF):	2,500,000 hours
Annualized Failure Rate (AFR): ²	0.35%
Preventive maintenance:	None required

1. Error rate specified with automatic retries and data correction with ECC enabled and all flaws reallocated.
2. See [Section 5.2, "Reliability and service"](#) for rated MTBF device operating condition requirements.

5.1 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 [Section 6.4, 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.

5.1.1 Recoverable Errors

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

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

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

5.1.2 Unrecoverable Errors

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

5.1.3 Seek errors

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

5.1.4 Interface errors

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

5.2 Reliability and service

Users can enhance the reliability of Exos X18 disk drives by ensuring that the drive receives adequate cooling. [Section 6.0](#) provides temperature measurements and other information that may be used to enhance the service life of the drive. [Section 9.2](#) provides recommended air-flow information.

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

The production disk drive shall achieve an annualized failure-rate of 0.35% (MTBF of 2,500,000 hours) over a 5 year service life when used in Enterprise Storage field conditions as limited by the following:

- 8760 power-on hours per year.
- HDA temperature as reported by the drive $\leq 30^{\circ}\text{C}$
- Ambient wet bulb temp $\leq 26^{\circ}\text{C}$
- Typical workload
- The AFR (MTBF) is a population statistic not relevant to individual units
- ANSI/ISA S71.04-2013 G2 classification levels and dust contamination to ISO 14644-1 Class 8 standards (as measured at the device)

The MTBF specification for the drive assumes the operating environment is designed to maintain nominal drive temperature and humidity. Occasional excursions in operating conditions between the rated MTBF conditions and the maximum drive operating conditions may occur without significant impact to the rated MTBF. However continual or sustained operation beyond the rated MTBF conditions will degrade the drive MTBF and reduce product reliability.

Nonrecoverable read errors	1 per 10^{15} bits read, max
Load unload cycles	600,000 cycles
Maximum Rated Workload	Maximum rate of <550TB/year Workloads exceeding the annualized rate may degrade the drive MTBF and impact product reliability. The Annualized Workload Rate is in units of TB per year, or TB per 8760 power on hours. Workload Rate = TB transferred * (8760 / recorded power on hours).
Warranty	To determine the warranty for a specific drive, use a web browser to access the following web page: http://www.seagate.com/support/warranty-and-replacements/ . From this page, click on the "Is my Drive under Warranty" link. The following are required to be provided: the drive serial number, model number (or part number) and country of purchase. The system will display the warranty information for the drive.
Preventive maintenance	None required.

5.2.2 Hot plugging the drive

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

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

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

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

5.2.3 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 users 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.

Users 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.

Users 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 the user to control when S.M.A.R.T. interruptions occur. Forcing S.M.A.R.T. with the RTZ command resets the timer.

Performance impact

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

Maximum processing delay

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). An example, if the MRIE is set to one, the firmware will issue to the host an 01-5D00 sense code. The FRU field contains the type of predictive failure that occurred. The error code is preserved through bus resets and power cycles.

Determining rate

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

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

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

Predictive failures

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

5.2.4 Thermal monitor

Exos X18 drives implement a temperature warning system which:

1. Signals the host if the temperature exceeds a value which would threaten the drive.
2. 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.

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).

5.2.5 Drive Self Test (DST)

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

There are two test coverage options implemented in DST:

1. Extended test
2. Short test

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

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

5.2.5.1 DST failure definition

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

The following conditions are considered DST failure conditions:

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

Recovered errors will not be reported as diagnostic failures.

5.2.5.2 Implementation

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

5.2.5.2.1 State of the drive prior to testing

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

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

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

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

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

5.2.5.2.2 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.

5.2.5.2.3 Short and extended tests

DST has two testing options:

1. short
2. extended

These testing options are described in the following two subsections.

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

Short test (Function Code: 001b)

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

Extended test (Function Code: 010b)

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

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

5.2.5.2.4 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.

5.2.5.2.5 Abort

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

Users 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).

5.2.6 Product warranty

See [“Seagate® Technology Support Services” on page 5](#) for warranty contact 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 the authorized Seagate distributor to purchase additional boxes. Seagate recommends shipping by an air-ride carrier experienced in handling computer equipment.

Storage

Maximum storage periods are 180 days within original unopened Seagate shipping package or 60 days unpackaged within the defined non-operating limits (refer to environmental section in this manual). Storage can be extended to 1 year packaged or unpackaged under optimal environmental conditions (25°C, <40% relative humidity non-condensing, and non-corrosive environment). During any storage period the drive non-operational temperature, humidity, wet bulb, atmospheric conditions, shock, vibration, magnetic and electrical field specifications should be followed.

Product repair and return information

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

6.0 Physical/electrical specifications

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

6.1 PowerChoice™ power management

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

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

PowerChoice modes

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

PowerChoice can be invoked using one of these two methods:

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

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

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

6.1.1 PowerChoice reporting methods

PowerChoice provides these reporting methods for tracking purposes:

Request Sense command reports

- Current power condition
- Method of entry

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

Mode Sense command reports (mode page 0x1A)

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

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

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

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

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

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

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

6.2 Power Balance

- Mode page 01h byte 6 bits 0 & 1 define the Active Level
- Active Levels - 00b Default, 11b Lowest active power level

6.3 AC power requirements

None.

6.4 DC power requirements

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

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

Table 1 DC power requirements (18TB and 16TB)

	Notes	12.0Gb mode		
		(Amps)	(Amps)	(Watts)
Voltage		+5V	+12V [[2]]	
Regulation	[5]	± 5%	± 10% [2]	
Avg idle current DCX	[1] [6]	0.39	0.30	5.58
Advanced idle current				
Idle A		0.39	0.30	5.59
Idle B		0.33	0.17	3.75
Idle C		0.33	0.12	3.06
Standby		0.32	0.01	1.73
Maximum starting current				
(peak DC) DC	[3]	1.12	1.77	
(peak AC) AC	[3]	1.21	2.86	
Delayed motor start (max) DC	[3] [4]	0.60	0.11	
Peak operating current (random read 4K16Q)				
Typical DCX	[1]	0.47	0.62	9.79
Maximum DC	[1]	0.49	0.63	
Maximum (peak) DC		1.26	2.30	
Peak operating current (random write 4K16Q)				
Typical DCX	[1]	0.56	0.35	6.95
Maximum DC	[1]	0.57	0.35	
Maximum (peak) DC		1.28	2.38	
Peak operating current (sequential read 64K16Q)				
Typical DCX	[1]	0.95	0.30	8.32
Maximum DC	[1]	0.96	0.30	
Maximum (peak) DC		1.37	0.64	
Peak operating current (sequential write 64K16Q)				
Typical DCX	[1]	1.06	0.30	8.88
Maximum DC	[1]	1.08	0.30	
Maximum (peak) DC		1.42	0.64	

- [1] Measured with average reading DC ammeter. Instantaneous +12V current peaks will exceed these values. Power supply at nominal voltage. N (number of drives tested) = 6, 35°C ambient.
- [2] For +12 V, a -10% tolerance is allowed during initial spindle start but must return to $\pm 10\%$ before reaching 7200 RPM. The $\pm 10\%$ 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 1](#).
- [4] This condition occurs after OOB and Speed Negotiation completes but before the drive has received the Notify Spinup primitive.
- [5] See [6.4.1, "Conducted noise immunity."](#) Specified voltage tolerance includes ripple, noise, and transient response.
- [6] During idle, the drive heads are relocated every 60 seconds to a random location within the band from three-quarters to maximum track.

General DC power requirement notes.

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

6.4.1 Conducted noise immunity

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

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

6.4.2 Power sequencing

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

6.4.3 Current profiles

The +12V (top) and +5V (bottom) current profiles for the Exos X18 drives are shown below.




Figure 1. 18TB and 16TB model current profiles.

Note All times and currents are typical. See [Table 1](#) for maximum current requirements.

6.5 Power dissipation

18TB and 16TB models in 12Gb operation

Please refer to [Table 1](#) for power dissipation numbers.

To obtain operating power for typical random read operations, refer to the following I/O rate curve (see [Figure 2](#)). Locate the typical I/O rate for a drive in the system on the horizontal axis and read the corresponding +5 volt current, +12 volt current, and total watts on the vertical axis. To calculate BTUs per hour, multiply watts by 3.4123.




Figure 2. 18TB and 16TB models (12Gb) DC current and power vs. input/output operations per second

6.6 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).

Note

To maintain optimal performance drives should be run at nominal drive temperatures and humidity. See [Section 5.2, "Reliability and service"](#) for rated MTBF device operating condition requirements.

6.6.1 Temperature

a. Operating

41°F to 140°F (5°C to 60°C) temperature range with a maximum temperature gradient of 36°F (20°C) per hour as reported by the drive.

The maximum allowable drive reported temperature is 140°F (60°C).

Air flow may be required to achieve consistent nominal drive temperature values (see [Section 9.2](#)). To confirm that the required cooling is provided for the electronics and HDA, place the drive in its final mechanical configuration, and perform random write/read operations. After the temperatures stabilize, monitor the current drive temperature using the Temperature log page (0Dh), with PARAMETER CODE 0000H TEMPERATURE. The TEMPERATURE field (byte 9) indicates the temperature of the SCSI target device in degrees Celsius at the time the LOG SENSE command is performed.

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.

6.6.2 Humidity

The values below assume that no condensation on the drive occurs. Maximum wet bulb temperature is 84.2°F (29°C).

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.

6.6.3 Effective altitude (sea level)

a. Operating

–1000 to +10,000 feet (–304.8 to +3048 meters)

b. Non-operating

–1000 to +40,000 feet (–304.8 to +12,192 meters)

6.6.4 Shock and Vibration

Shock and vibration measurements specified in this document are made directly on the drive itself and applied in the X, Y, and Z axis at the drive mounting point locations.

6.6.4.1 Shock

a. Operating

The drive will operate without error while subjected to intermittent shock pulses not exceeding 50g at a duration of 2ms.

b. Non-operating

The drive will operate without non-recoverable errors after being subjected to shock pulses not exceeding 200g at a duration of 2ms.

6.6.4.2 Vibration

a. Linear Random Operating Vibration

The drive will operate without non-recoverable errors while being subjected to the random power spectral density noise specified below.

PSD of 5-500 Hz random noise at 0.70 g rms					
Frequency (Hz)	5	20	200	250	500
G ² /Hz	0.00025	0.00210	0.00210	0.00020	0.00020

b. Random Rotary Operating Vibration

The drive will exhibit greater than 90% throughput for sequential and random write operations while subjected to the shaped random power spectral density noise specified below.

PSD Profile 20-1500 Hz at 12.5 rad/sec ²				
Frequency (Hz)	20	200	800	1500
(rad/sec ²) ² /Hz	5.53E-02	5.53E-02	3.49E-01	6.14E-04

c. Linear Random Non-Operating Vibration

The drive will not incur physical damage or have non-recoverable errors after being subjected to the power spectral density noise specified below.

PSD Profile 2-500 Hz at 2.27 g rms				
Frequency (Hz)	2	4	100	500
G ² /Hz	0.001	0.030	0.030	0.001

6.6.5 Acoustics

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

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

There will not be any discrete tones more than 9 dB above the masking noise when measured according to Seagate specification 30553-001.

6.6.6 Air cleanliness

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

6.6.7 Corrosive environment

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

Users should use caution exposing any electronic components to uncontrolled chemical pollutants and corrosive chemicals as electronic drive component reliability can be affected by the installation environment. The silver, copper, nickel and gold films used in hard disk drives are especially sensitive to the presence of sulfide, chloride, and nitrate contaminants. Sulfur is found to be the most damaging. 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.

Seagate recommends that data centers be kept clean by monitoring and controlling the dust and gaseous contamination. Gaseous contamination should be within ANSI/ISA S71.04-2013 G2 classification levels (as measured on copper and silver coupons), and dust contamination to ISO 14644-1 Class 8 standards, and MTBF rated conditions as defined in the Annualized Failure Rate (AFR) and Mean Time Between Failure (MTBF) section.

6.7 Mechanical specifications

Refer to **Figure 3** for detailed mounting configuration dimensions. See **Section 9.3**, "Drive mounting."

Weight: 1.477 lb 670 g

Note These dimensions conform to the Small Form Factor Standard documented in SFF-8301 and SFF-8323, found at www.snia.org/technology-communities/sff/specifications.




Figure 3. Mounting configuration dimensions

Note The image is for mechanical dimension reference only and may not represent the actual drive.

7.0 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 2.1](#).

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

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

- Security Protocol Out
- Security Protocol In

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

7.1 Data encryption

Encrypting drives use one inline encryption engine for each port, employing AES-256 bit data encryption in AES-XTS 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 and cannot be disabled.

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 32 data bands (see [Section 7.5](#)).

7.2 Controlled access

The drive has two security providers (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.

7.2.1 Admin SP

The Admin SP allows the drive’s owner to enable or disable firmware download operations (see [Section 7.4](#)). Access to the Admin SP is available using the SID (Secure ID) password or the MSID (Manufacturers Secure ID) password.

7.2.2 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 32 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 31).

7.2.3 Default password

When the drive is shipped from the factory, all passwords are set to the value of MSID. This 32-byte random value can only be read by the host electronically over the interface. 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.

7.3 Random number generator (RNG)

The drive has a 32-byte hardware RNG that it uses to generate 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.

7.4 Drive locking

In addition to changing the passwords, as described in [Section 7.2.3](#), 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. In addition “ReadLockEnabled” and “WriteLockEnabled” must be set to true in the locking table in order for the bands “LockOnReset” setting of “PowerCycle” to actually lock access to the band when a “PowerCycle” event occurs. 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.

7.5 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 30 Data Bands may be defined in a similar way (Band2 through Band31) 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 7.6](#)) or the password when required. The bands should be aligned to 4K LBA boundaries.

7.6 Cryptographic erase

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

7.7 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.

7.8 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 6.4](#) for power requirements on the standard (non-SED) drive models.

7.9 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 5](#):

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

7.10 Sanitize - CRYPTOGRAPHIC ERASE

This command cryptographically erases all user data on the drive by destroying the current data encryption key and replacing it with a new data encryption key randomly generated by the drive. Sanitize CRYPTOGRAPHIC ERASE is a SCSI CDB Op code 48h and selecting the service action code 3 (CRYPTOGRAPHIC ERASE)

7.11 RevertSP

SED models will support the RevertSP feature which 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. In order to execute the RevertSP method the unique PSID (Physical Secure ID) printed on the drive label must be provided. PSID is not electronically accessible and can only be manually read from the drive label or scanned in via the 2D barcode.

8.0 Defect and error management

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

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

- Deferred Auto-Reallocation (see [Section 8.4](#))
- Idle Read After Write (see [Section 8.5](#))

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

8.1 Drive internal defects/errors

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

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

8.2 Drive error recovery procedures

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

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

The drive firmware error recovery algorithms consist of 12 levels for read recoveries and five levels for write. Each level may consist of multiple steps, where a step is defined as a recovery function involving a single re-read or re-write attempt. The maximum level used by the drive in LBA recovery is determined by the read and write retry counts.

Table 2 equates the read and write retry count with the maximum possible recovery time for read and write recovery of individual LBAs. The times given do not include time taken to perform reallocations. Reallocations are performed when the ARRE bit (for reads) or AWRE bit (for writes) is one, the RC bit is zero, and the recovery time limit for the command has not yet been met. Time needed to perform reallocation is not counted against the recovery time limit.

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

Table 2 Read and write retry count maximum recovery time

Read retry count*	Maximum recovery time per LBA (cumulative, ms)	Write retry count*	Maximum recovery time per LBA (cumulative, ms)
		0	35.94
1	124.32	1	53.91
5	621.62	2	79.89
10	1243.23	3	97.86
15	1864.85	4	175.85
20 (default)	2486.47	5 (default)	421.79

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

e.g. 1 ~ 5%

5 ~ 25%

20 ~ 100%

Setting these retry counts to a value below the default setting could result in degradation of the unrecovered error rate. For example, suppose the read/write recovery page has the RC bit = 0 and if the read retry count is set to 5, this means ~ 25% of error recovery will be executed which consumes 621.62 ms (please refer to the table above). If the limit is reached and a LBA has not yet been recovered (i.e. requires retries beyond 621.62 ms), the command will end with Check Condition status report and unrecoverable read error will be reported.

8.3 SAS system errors

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

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

8.4 Deferred Auto-Reallocation

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

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

8.5 Idle Read After Write

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

8.6 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 sectors that store 512 bytes of customer data with 8 bytes of Protection Information appended to it. The advantage of PI is that the Protection Information bits can be managed at the HBA and HBA driver level. Allowing a system that typically does not support 520 LBA formats to integrate this level of protection.

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

8.6.1 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 (e.g.0xFFFF) 8 bytes of Protection Information to be stored on the media, but the 8 bytes will not be transferred to the host during a read command.

Type 3 - Seagate products do not support Type 3.

8.6.2 Setting and determining the current Type Level

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

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

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

8.6.3 Identifying a Protection Information drive

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

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

9.0 Installation

Exos X18 disk drive installation is a plug-and-play process. There are no jumpers, switches, or terminators on the drive.

SAS drives are designed to be used in a host system that provides a SAS-compatible backplane with bays designed to accommodate the drive. In such systems, the host system typically provides a carrier or tray into which users need to mount the drive. Mount the drive to the carrier or tray provided by the host system only using 6-32 UNC mounting screws. The screws should be inserted no more than 0.140 in (3.56mm) into the bottom or side mounting holes. When tightening the screws, do not overtighten use a maximum torque of 6 in-lb. Users can mount the drive in any orientation.

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

Slide the carrier or tray into the appropriate bay in 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 10.4.1](#) for additional information about these connectors.

Power is supplied through the SAS connector.

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




Figure 4. Physical interface

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

9.1 Drive orientation

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

9.2 Cooling

Cabinet cooling must be designed by the customer so that the ambient temperature immediately surrounding the drive will not exceed temperature conditions specified in [Section 6.6.1, "Temperature."](#)

The rack, cabinet, or drawer environment for the drive must provide heat removal from the electronics and head and disk assembly (HDA). Users should confirm that adequate heat removal is provided using the temperature measurement guidelines described in [Section 6.6.1](#).

Forced air flow may be required to keep temperatures at or below the temperatures specified in [Section 6.6.1](#) in which case the drive should be oriented, or air flow directed, so that the least amount of air flow resistance is created while providing air flow to the electronics and HDA. Also, the shortest possible path between the air inlet and exit should be chosen to minimize the travel length of air heated by the drive and other heat sources within the rack, cabinet, or drawer environment.

If forced air is determined to be necessary, possible air-flow patterns are shown in [Figure 5](#). 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 6.6.1](#) are met.





Figure 5. Air flow

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

9.3 Drive mounting

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

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

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

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

9.4 Grounding

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

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

10.0 Interface requirements

This section partially describes the interface requirements as implemented on Exos X18 drives. Additional information is provided in the *SAS Interface Manual* (part number 100293071).

10.1 SAS features

This section lists the SAS-specific features supported by Exos X18 drives.

10.1.1 task management functions

Table 3 lists the SAS task management functions supported.

Table 3 SAS task management functions supported

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

10.1.2 task management responses

Table 4 lists the SAS response codes returned for task management functions supported.

Table 4 Task management response codes

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

10.2 Dual port support

Exos X18 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 12Gb/s or the first port can run at 6Gb/s while the second port runs at 3Gb/s.) The supported link rates are 3.0, 6.0, or 12.0 Gb/s.

Subject to buffer availability, the Exos X18 drives support:

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

10.3 SCSI commands supported

Table 5 lists the SCSI commands supported by Exos X18 drives.

Table 5 Supported commands

Command name	Command code	Supported ^[4]
Change Definition	40h	N
Compare	39h	N
Copy	18h	N
Copy and Verify	3Ah	N
Format Unit ^{[1] [5]}	04h	Y
DCRT bit supported		Y
DPRY bit supported		N
DSP bit supported		Y
IMMED bit supported		Y
IP bit supported		Y
SI (Security Initialize) bit supported		N
STPF bit supported		Y
VS (vendor specific)		N
Get Physical Element Status	9Eh/17h	Y
Inquiry	12h	Y
Block Device Characteristics (B1h)		Y
Block Limits (B0h)		Y
Date Code page (C1h)		Y
Device Behavior page (C3h)		Y
Device Identification (83h)		Y
Extended Inquiry Data (86h)		Y
Firmware Numbers page (C0h)		Y
Jumper Settings page (C2h)		Y
Logical Block Provisioning (B2h)		Y
Mode Page Policy (87h)		Y
Power Condition (8Ah)		Y
Protocol Specific Logical Unit Information (90h)		Y
Protocol Specific Port Information (91h)		Y
SCSI Ports (88h)		Y
Supported Vital Product Data page (00h)		Y
Unit Serial Number page (80h)		Y
Lock-unlock cache	36h	N
Log Select	4Ch	Y
PCR bit		Y
DU bit		N
DS bit		Y
TSD bit		Y
ETC bit		N

Table 5 Supported commands

Command name	Command code	Supported ^[4]
TMC bit		N
LP bit		N
Log Sense	4Dh	Y
Application Client Log page (0Fh)		Y
Background Scan Results page (15h) (SBC-3)		Y
Buffer Over-run/Under-run page (01h)		N
Cache Statistics page (37h)		Y
Factory Log page (3Eh)		Y
Format Status page (08h)		Y
Information Exceptions Log page (2Fh)		Y
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
Power Condition Transition page (1Ah)		Y
Protocol Specific Port page (18h)		Y
Read Error Counter page (03h)		Y
Read Reverse Error Counter page (04h)		N
Self-test Results page (10h)		Y
Start-stop Cycle Counter page (0Eh)		Y
Temperature page (0Dh)		Y
Verify Error Counter page (05h)		Y
Write error counter page (02h)		Y
Mode Select (same pages as Mode Sense 1Ah)	15h	Y ^[2]
Mode Select (10) (same pages as Mode Sense 1Ah)	55h	Y
Mode Sense	1Ah	Y ^[2]
Caching Parameters page (08h)		Y
Control Mode page (0Ah)		Y
Disconnect/Reconnect (02h)		Y
Error Recovery page (01h)		Y
Format page (03h)		Y
Information Exceptions Control page (1Ch/01h)		Y
Background Scan mode subpage (01h)		Y
Notch and Partition Page (0Ch)		N
Protocol-Specific Logical Uit (18h)		Y
Protocol-Specific Port page (19h)		Y
Power Condition page (1Ah)		Y
Rigid disk Drive Geometry page (04h)		Y
Unit Attention page (00h)		Y
Verify Error Recovery page (07h)		Y
Xor Control page (10h)		N

Table 5 Supported commands

Command name	Command code	Supported ^[4]
Mode Sense (10) (same pages as Mode Sense 1Ah)	5Ah	Y
Persistent Reserve In	5Eh	Y
Persistent Reserve Out	5Fh	Y
Prefetch (10)	34h	N
Prefetch (16)	90h	N.A.
Read (6)	08h	Y
Read (10)	28h	Y
DPO bit supported		Y
FUA bit supported		Y
Read (12)	A8h	N
Read (16)	88h	Y
Read (32)	7Fh/0009h	Y
Read Buffer (modes 0, 1c, 2, 3, Ah and Bh supported)	3Ch	Y (non-SED drives only)
Read Buffer (16) (modes 0, 1c, 2, 3, Ah and Bh supported)	9Bh	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
Remove Element And Truncate	9Eh/18h	Y
Report Identifying Information	A3h/05h	Y
Report LUNs	A0h	Y
Report Supported Operation Codes	A3h/0Ch	Y
Report Supported Task Management Functions	A3h/0Dh	Y
Report Timestamp	A3h/0Fh	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

Table 5 Supported commands

Command name	Command code	Supported ^[4]
Extent Reservation		N
Restore Elements and Rebuild	9Eh/19h	Y
Rezero Unit (6)	01h	Y
Sanitize	48h	Y
Block Erase	48h / 02h	N
Cryptographic Erase	48h / 03h	Y (SED only)
Overwrite	48h / 01h	Y
Sanitize Exit	48h / 1Fh	Y
Search Data Equal	31h	N
Search Data High	30h	N
Search Data Low	32h	N
Security Protocol In	A2h	Y (SED models only)
Security Protocol Out	B5h	Y (SED models only)
Seek (6)	0Bh	Y
Seek (10)	2Bh	Y
Send Diagnostics	1Dh	Y
Supported Diagnostics pages (00h)		Y
Translate page (40h)		Y
Set Identifying Information	A4h/06h	Y
Set Limits	33h	N
Set Timestamp	A4h/0Fh	Y
Start Unit/Stop Unit (spindle ceases rotating)	1Bh	Y
Synchronize Cache (10)	35h	Y
Synchronize Cache (16)	91h	Y
Test Unit Ready	00h	Y
Verify (10)	2Fh	Y
BYTCHK bit		Y
Verify (12)	AFh	N
Verify (16)	8Fh	Y
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

Table 5 Supported commands

Command name	Command code	Supported ^[4]
Write and Verify (32)	7Fh/000Ch	Y
Write Buffer (modes 0, 1A, 1C, 2, 6, A, D, E, F supported)	3Bh	Y (non-SED drives only)
Firmware Download option (modes 4, 5, 7, Ah and E) ^[3]		Y (non-SED drives only)
Firmware Download option (modes 4, 5, 7)		Y (SED drives only)
Write Long (10)	3Fh	Y
Write Long (16)	9Fh/11h	Y
Write Same (10) ^[5]	41h	Y
PBdata		N
LBdata		N
Write Same (16) ^[5]	93h	Y
Write Same (32)	7Fh/000Dh	Y
XDRead	52h	N
XDWrite	50h	N
XPWrite	51h	N

[1] Exos X18 drives can format to 512, 520 or 528 bytes per logical block.

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

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

[4] Y = Yes. Command is supported.

N = No. Command is not supported.

A = Support is available on special request.

[5] Approximately 1.5 increase in time to complete this command for a SED drive versus a non-SED drive of the same capacity.

10.3.1 Inquiry data

Table 6 lists the Inquiry command data that the drive should return to the initiator per the format given in the *SAS Interface Manual*.

Table 6 Exos X18 inquiry data

Bytes	Data (hex)																
0-15	00	00	xx**	12	8B	00	30	02	53	45	41	47	41	54	45	20	Vendor ID
16-31	[53	54	31	38	30	30	30	4E	4D	30	30	34	4A]	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*	32*	30*	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 28 reflect model of drive. The table above shows the hex values for Model ST18000NM004J.

Refer to the values below for the values of bytes 16 through 28 of a particular model:

ST18000NM005J	53	54	31	38	30	30	30	4E	4D	30	30	35	4A
ST18000NM007J	53	54	31	38	30	30	30	4E	4D	30	30	37	4A
ST16000NM004J	53	54	31	36	30	30	30	4E	4D	30	30	34	4A
ST16000NM005J	53	54	31	36	30	30	30	4E	4D	30	30	35	4A
ST16000NM007J	53	54	31	36	30	30	30	4E	4D	30	30	37	4A

10.3.2 Mode Sense data

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

1. Default values

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

2. Saved values

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

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

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

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

3. Current values

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

4. Changeable values

Changeable values form a bit mask, stored in nonvolatile memory, that dictates which of the current values and saved values can be changed by a Mode Select command. A one (1) indicates the value can be changed. A zero (0) indicates the value is not changeable. For example, in [Table 7](#), refer to Mode page 81, in the row entitled "CHG." These are hex numbers representing the changeable values for Mode page 81. Note in columns 5 and 6 (bytes 04 and 05), there is 00h which indicates that in bytes 04 and 05 none of the bits are changeable. Note also that bytes 06, 07, 09, 10, and 11 are not changeable, because those fields are all zeros. In byte 02, hex value FF equates to the binary pattern 11111111. If there is a zero in any bit position in the field, it means that bit is not changeable. Since all of the bits in byte 02 are ones, all of these bits are changeable.

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

Note

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

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

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

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

10.4 Miscellaneous operating features and conditions

Table 8 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 8 *Miscellaneous features*

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

Table 9 *Miscellaneous status*

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

10.4.1 SAS physical interface

Figure 6 shows the location of the SAS device connector J1. **Figures 7** and **8** 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 6. Physical interface




Figure 8. SAS device plug dimensions (detail)

10.4.2 Physical characteristics

This section defines physical interface connector.

10.4.3 Connector requirements

Contact the 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 7](#) and [8](#).

10.4.4 Electrical description

SAS drives use the device connector for:

- DC power
- SAS interface
- Activity LED

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

10.4.5 Pin descriptions

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

Table 10 SAS pin descriptions

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

* Short pin to support hot plugging

** Power Disable (T10 Industry Standard) for remote management of the end device. Allows power cycling / power saving to be controlled by the host via interface pin 3.

† P1 & P2 tied for visible host detection.

10.4.6 SAS transmitters and receivers

A typical SAS differential copper transmitter and receiver pair is shown in **Figure 9**. The receiver is AC coupling to eliminate ground shift noise.




Figure 9. SAS transmitters and receivers

10.4.7 Power

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

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

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

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

10.5 Signal characteristics

This section describes the electrical signal characteristics of the drive’s input and output signals. See **Table 10** for signal type and signal name information.

10.5.1 Ready LED Out

The Ready LED Out signal is driven by the drive as indicated in **Table 11**.

Table 11 Ready LED Out conditions

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

The Ready LED Out signal is designed to pull down the cathode of an LED. The anode is attached to the proper +3.3 volt supply through an appropriate current limiting resistor. The LED and the current limiting resistor are external to the drive. See **Table 12** for the output characteristics of the LED drive signals.

Table 12 LED drive signal

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

10.5.2 Differential signals

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

Table 13 defines the general interface characteristics.

Table 13 General interface characteristics

Characteristic	Units	1.5Gb/s	3.0Gb/s	6.0Gb/s	12 Gbps
Bit rate (nominal)	Mbaud	1,500	3,000	6,000	12000
Unit interval (UI) (nominal)	ps	666.6	333.3	166.6	83.3
Impedance (nominal, differential)	ohm	100	100	100	100
Transmitter transients, maximum	V	± 1.2	± 1.2	± 1.2	± 1.2
Receiver transients, maximum	V	± 1.2	± 1.2	± 1.2	± 1.2

10.6 SAS-3 Specification Compliance

Seagate SAS-3 compatible drives are compliant with the latest SAS-3 Specification (T10/BSR INCITS 519 rev. 06).

The main difference from SAS-2 is the Tx and Rx training that allows the host and drive to adjust the amplitude and emphasis values to the channel. The receiver still employs Decision Feedback Equalizer (DFE) and Feed Forward Equalizer (FFE) circuitry to accomplish this training.

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.

10.7 Additional information

Please contact the 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.



Seagate Technology LLC

AMERICAS Seagate Technology LLC 47488 Kato Road, Fremont, California 94538, United States, 510-661-1000

Publication Number: 100865853, Rev. A
August 2020