Product Manual

Cheetah 15K.4 FC

ST3146854FC

ST373454FC

ST336754FC

Seagate

Product Manual

Cheetah 15K.4 FC

ST3146854FC

ST373454FC

ST336754FC

Seagate

Revision status summary sheet

Revision	Date	Sheets Affected
Rev. A	06/30/04	1/1, v-viii, 1-96. Initial release.
Rev. B	07/14/04	Page 9.
Rev. C	05/04/05	Pages 2,11, 25, 26, 36, 38, 50, 51, 57, and 75 - 79.

Contents

Cheetah 15K.4 FC Product Manual, Rev. C

List of Figures

1.0   Scope

This manual describes Seagate Technology® LLC, Cheetah® 15K.4 FC (Fibre Channel) disc drives.

Cheetah 15K.4 FC drives support the Fibre Channel Arbitrated Loop and SCSI Fibre Channel Protocol specifications to the extent described in this manual. The *Fibre Channel Interface Manual* (part number 77767496) describes the general Fibre Channel Arbitrated Loop characteristics of this and other Seagate Fibre Channel drives.

Figure 1.   **Cheetah 15K.4 FC family disc drive**

2.0   Applicable standards and reference documentation

The drive has been developed as a system peripheral to the highest standards of design and construction. The drive depends upon its host equipment to provide adequate power and environment in order to achieve optimum performance and compliance with applicable industry and governmental regulations. Special attention must be given in the areas of safety, power distribution, shielding, audible noise control, and temperature regulation. In particular, the drive must be securely mounted in order to guarantee the specified performance characteristics. Mounting by bottom holes must meet the requirements of Section 8.5.

2.1   Standards

The Cheetah 15K.4 FC family complies with Seagate standards as noted in the appropriate sections of this manual and the Seagate *Fibre Channel Interface Manual*, part number 77767496.

The Cheetah 15K.4 FC disc drive is a UL recognized component per UL1950, CSA certified to CAN/CSA C22.2 No. 950-95, and VDE certified to VDE 0805 and EN60950.

2.1.1   Electromagnetic compatibility

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

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

2.1.1.1   Electromagnetic susceptibility

As a component assembly, the drive is not required to meet any susceptibility performance requirements. It is the responsibility of those integrating the drive within their systems to perform those tests required and design their system to ensure that equipment operating in the same system as the drive or external to the system does not adversely affect the performance of the drive. See Table 2, DC power requirements.

2.1.2 Electromagnetic compliance

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

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

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

Electromagnetic compliance for the European Union

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

Australian C-Tick

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

Korean MIC

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

Taiwanese BSMI

If this model has the Chinese National Standard (CNS) 13438 marking, it complies with Chinese National Standard (CNS) 13438 and meets the Electromagnetic Compatibility (EMC) Framework requirements of the Taiwanese Bureau of Standards, Metrology, and Inspection (BSMI).

2.2   Reference documents

Cheetah 15K.4 FC Installation Guide   Seagate part number: 100220460

Fibre Channel Interface Manual   Seagate part number: 77767496

ANSI Fibre Channel Documents
X3.230-1994   FC Physical and Signaling Interface (FC-PH)
X3.297.1997   FC-PH-2 Fibre Channel Physical and Signaling Interface-2
X3.303.1998   FC-PH-3 Fibre Channel Physical and Signaling Interface-3
X3.272-1996   FC Arbitrated Loop (FC-AL)
X3.269-1996   Fibre Channel Protocol for SCSI (FCP)
NCITS TR-19   Private Loop SCSI Direct Attach (PLDA)
NCITS TR-20   Fabric Loop Attachment (FC-FLA)
SFF-8045   Specification for 40-pin SCA-2 Connector with Parallel Selection
SFF-8067   Specification for 40-pin SCA-2 Connector with Bidirectional
                                         Enclosure Services Interface

ANSI Small Computer System Interface (SCSI) Documents
X3.131-1994   (SCSI-2)
X3.270-1996   (SCSI-3) Architecture Model
NCITS 305-199X   (SCSI-3) Enclosure Services

Specification for Acoustic Test Requirement and Procedures
                                         Seagate part number: 30553-001

Package Test Specification   Seagate P/N 30190-001 (under 100 lb.)
Package Test Specification   Seagate P/N 30191-001 (over 100 lb.)

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

3.0   General description

Cheetah 15K.4 FC drives combine giant magnetoresistive (GMR) heads, partial response/maximum likelihood (PRML) read channel electronics, embedded servo technology, and a Fibre Channel interface to provide high performance, high capacity data storage for a variety of systems including engineering workstations, network servers, mainframes, and supercomputers. Cheetah 15K.4 FC drives also support 2-Gbit Fibre Channel which can transfer data at up to 200 Mbytes per second and 400 Mbytes per second in dual-loop configurations.

Cheetah 15K.4 FC drives support the Fibre Channel Arbitrated Loop (FC-AL) and SCSI Fibre Channel Protocol as described in the ANSI specifications, this document, and the *Fibre Channel Interface Manual* (part number 77767496) which describes the general interface characteristics of this drive. Cheetah 15K.4 FC drives are classified as intelligent peripherals and provide level 2 conformance (highest level) with the ANSI SCSI-1 standard.

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

Cheetah 15K.4 FC drives use a dedicated landing zone at the innermost radius of the media to eliminate the possibility of destroying or degrading data by landing in the data zone. The heads automatically go to the landing zone when power is removed from the drive.

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

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.

3.1 Standard features

Cheetah 15K.4 FC drives have the following standard features:

- 2-Gbit Fibre Channel interface. For 2 gigabit operation, the block size must be divisable by 8.
- Integrated dual port FC-AL controller
- Concurrent dual port transfers
- Support for FC arbitrated loop, private and public attachment
- Differential copper FC drivers and receivers
- Downloadable firmware using the FC-AL interface
- Supports SCSI enclosure services via interface connector
- 128-deep task set (queue)
- Supports up to 32 initiators
- Drive selection ID and configuration options are set on the FC-AL backpanel or through interface commands. Jumpers are not used on the drive.
- Supports SCSI Enclosure Services through the interface connector
- Fibre Channel worldwide name uniquely identifies the drive and each port
- User-selectable logical block size (512 to 528 bytes per logical block). For 2 gigabit operation, the block size must be divisable by 8.
- Selectable frame sizes from 256 to 2,112 bytes
- Industry standard 3.5-inch low profile form factor dimensions
- Programmable logical block reallocation scheme
- Flawed logical block reallocation at format time
- Programmable auto write and read reallocation
- Reed-Solomon error correction code
- Sealed head and disc assembly (HDA)
- No preventive maintenance or adjustments required
- Dedicated head landing zone
- Automatic shipping lock
- Embedded Grey Code track address to eliminate seek errors
- Self-diagnostics performed at power on
- 1:1 interleave
- Zone bit recording (ZBR)
- Vertical, horizontal, or top down mounting
- Dynamic spindle brake
- 8,192 Kbyte data buffer (see Section 4.5).
- Embedded servo design
- Reallocation of defects on command (Post Format)
- Fibre Channel interface transports SCSI protocol

3.2 Media description

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

3.3   Performance

- Programmable multi-segmentable cache buffer
- 200 Mbytes/sec maximum instantaneous data transfers per port
- 15k RPM spindle; average latency = 2.0 msec
- Command queuing of up to 128 commands
- Background processing of queue
- Supports start and stop commands (spindle stops spinning)
- Adaptive seek velocity; improved seek performance

3.4   Reliability

- 1,400,000 hour MTBF
- Increased LSI circuitry integration
- Self-Monitoring Analysis and Reporting Technology (S.M.A.R.T.)
- Dithering

3.5   Formatted capacities

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

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

1. Seagate designs specify capacity points at certain sector 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:

Sector Size	ST336754FC		ST373454FC		ST3146854FC	
	Decimal	Hex	Decimal	Hex	Decimal	Hex
512	71,687,372	445DCCC	143,374,744	88BB998	286,749,488	11177330
514	70,512,692	433F034	141,025,384	867E068	282,050,768	10CFC0D0
520	70,197,546	42F212A	140,395,092	85E4254	280,790,184	10BC84A8
522	68,914,176	41B8C00	139,463,602	8500BB2	279,041,740	10A1D6CC
524	68,766,592	4194B80	137,577,184	83342E0	275,154,368	106685C0
528						

2. Seagate drives also may be used at the maximum available capacity at a given sector size, but the excess capacity above the guaranteed level will vary between 10K and 15K families and from generation to generation, depending on how each sector size actually formats out for zone frequencies and splits over servo bursts. This added capacity potential may range from 0.1 to 1.3 percent above the guaranteed capacities listed above. Using the drives in this manner gives the absolute maximum capacity potential, but the user must determine if the extra capacity potential is useful, or whether their assurance of backward and forward compatibility takes precedence.

3.5.1   Programmable drive capacity

Using the Mode Select command, the drive can change its capacity to something less than maximum. See the Mode Select Parameter List table in the *SCSI Interface Product Manual*. Refer to the Parameter list block descriptor number of blocks field. A value of zero in the number of blocks field indicates that the drive shall not

change the capacity it is currently formatted to have. A number in the number of blocks field that is less than the maximum number of LBAs changes the total drive capacity to the value in the block descriptor number of blocks field. A value greater than the maximum number of LBAs is rounded down to the maximum capacity.

3.6   Factory-installed accessories

OEM standard drives are shipped with the *Cheetah 15K.4 FC Installation Guide* (part number 100220460) and the *Safety and Regulatory Agency Specifications*, part number 75789512 (unless otherwise specified).

3.7   Factory-installed options

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

• Other capacities can be ordered depending on sparing scheme and sector size requested.

• Single-unit shipping pack. The drive is normally shipped in bulk packaging to provide maximum protection against transit damage. Units shipped individually require additional protection as provided by the single unit shipping pack. Users planning single unit distribution should specify this option.

• The *Cheetah 15K.4 FC Installation Guide,* part number 100220460, is usually included with each standard OEM drive shipped, but extra copies may be ordered.

• The *Safety and Regulatory Agency Specifications*, part number 75789512, is usually included with each standard OEM drive shipped, but extra copies may be ordered.

3.8   User-installed accessories

The following accessories are available. All kits may be installed in the field.

• Evaluation kit, part number 73473641.
  This kit provides an adapter card ("T-card") to allow cable connections for two FC ports and DC power.

• Single-unit shipping pack.

4.0   Performance characteristics

This section provides detailed information concerning performance-related characteristics and features of Cheetah 15K.4 FC drives.

4.1   Internal drive characteristics

	ST3146854FC	ST373454FC	ST336754FC	
Drive capacity	146.8	73.4	36.7	Gbytes (formatted, rounded off value)
Read/write data heads	8	4	2	
Bytes per track	471,916	471,916	471,916	Bytes (average, rounded off values)
Bytes per surface	18,351	18,351	18,351	Mbytes (unformatted, rounded off value)
Tracks per surface (total)	50,864	50,864	50,864	Tracks (user accessible)
Tracks per inch	85,000	85,000	85,000	TPI
Peak bits per inch	628	628	628	KBPI
Internal data rate	685-1,142	685-1,142	685-1,142	Mbits/sec (variable with zone)
Disc rotation speed	15k	15k	15k	rpm
Avg rotational latency	2.0	2.0	2.0	msec

4.2   Seek performance characteristics

See Section 9.5, "FC-AL physical interface" on page 64 and the *Fibre Channel Interface Manual* (part number 77767496) for additional timing details.

4.2.1   Access time

		Not including controller overhead[1,2](msec)	
		Read	Write
Average	Typical	3.5	4.0
Single track	Typical	0.2	0.4
Full stroke	Typical	7.4	7.9

1. Typical access times are measured under nominal conditions of temperature, voltage, and horizontal orientation as measured on a representative sample of drives.
2. Access to data = **access time** + latency time.

4.2.2   Format command execution time (minutes)

	ST3146854FC	ST373454FC	ST336754FC
Maximum (with verify)	90	60	30
Maximum (without verify)	45	30	15

4.2.3   General performance characteristics

Data buffer to/from disc media (one 512-byte logical block)*

Minimum	685 Mbits/sec
Maximum	1,142 Mbits/sec

Fibre Channel Interface maximum instantaneous transfer rate	200 Mbytes/sec* per port (dual port = 400 Mbytes/sec*)

Logical block sizes

Default is 512-byte data blocks
Variable 512 to 704 bytes per sector in even number of bytes per sector. If n (number of bytes per sector) is odd, then n-1 will be used.

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	2.0 msec

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

4.3   Start/stop time

If the Motor Start option is disabled, the drive becomes ready within 20 seconds after DC power is applied. If a recoverable error condition is detected during the start sequence, the drive executes a recovery procedure and the time to become ready may exceed 20 seconds. Stop time is 30 seconds (maximum) from removal of DC power.

If the Motor Start option is enabled, the internal controller accepts the commands listed in the *Fibre Channel Interface Manual* less than 3 seconds after DC power has been applied. After the Motor Start command has been received, the drive becomes ready for normal operations within 20 seconds (excluding the error recovery procedure). The Motor Start command can also be used to command the drive to stop the spindle.

There is no power control switch on the drive.

4.4   Prefetch/multi-segmented cache control

The drive provides a prefetch (read look-ahead) and multi-segmented cache control algorithms that in many cases can enhance system performance. Cache refers to the drive buffer storage space when it is used in cache operations. To select this feature, the host sends the Mode Select command with the proper values in

---

the applicable bytes in page 08h. Prefetch and cache operations are independent features from the standpoint that each is enabled and disabled independently using the Mode Select command; however, in actual operation, the prefetch feature overlaps cache operation somewhat as described in sections 4.5.1 and 4.5.2.

All default cache and prefetch mode parameter values (Mode Page 08h) for standard OEM versions of this drive family are given in Table 16.

4.5 Cache operation

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

Of the 8,192 kbytes physical buffer space in the drive, approximately 7,077 kbytes can be used as a cache. The buffer can be divided into logical segments (using Mode Select Page 08h, byte 13) 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 *Fibre Channel Interface Manual*), data requested by the host with a read command is retrieved from the buffer, if possible, before any disc access is initiated. If cache operation is not enabled, the buffer (still segmented with the required number of segments) is still used, but only as circular buffer segments during disc 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. On a cache miss, all data transfers to the host are in accordance with buffer-full ratio rules. On a cache hit, the drive ignores the buffer-full ratio rules. See the explanation provided with the information about Mode Page 02h (disconnect/reconnect control) in the *Fibre Channel Interface Manual*.

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

Case A—read command is received and the first logical block is already in the cache:

1. Drive transfers to the initiator the first logical block requested plus all subsequent contiguous logical blocks that are already in the cache. This data may be in multiple segments.

2. When a requested logical block is reached that is not in any cache segment, the drive fetches it and any remaining requested logical block addresses from the disc and puts them in a segment of the cache. The drive transfers the remaining requested logical blocks from the cache to the host in accordance with the Mode Select Disconnect/Reconnect parameters, page 02h.

3. If the prefetch feature is enabled, refer to section 4.5.2 for operation from this point.

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

1. The drive fetches the requested logical blocks from the disc 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 4.5.2 for operation from this point.

During a prefetch, the drive crosses a cylinder boundary to fetch data only if the Discontinuity (DISC) bit is set to 1 in bit 4 of byte 2 of the Mode Select parameters page 08h. Default is zero for bit 4.

Each cache segment is actually a self-contained circular buffer whose length is an integer number of logical blocks. The wrap-around capability of the individual segments greatly enhances the cache's overall performance, allowing a wide range of user-selectable configurations. The drive supports operation of any integer number of segments from 1 to 32. Divide the 7,077 Kbytes in the buffer by the number of segments to get the segment size in bytes; then divide by the sector size to get the number of sectors per segment, any partial sectors remaining are not used. Default is 3 segments.

Note. The size of each segment is not reported by Mode Sense command page 08h, bytes 14 and 15. The value 0XFFFF is always reported regardless of the actual size of the segment. Sending a size specification using the Mode Select command (bytes 14 and 15) does not set up a new segment

---

size. If the STRICT bit in Mode page 00h (byte 2, bit 1) is set to one, the drive responds as it does for any attempt to change an unchangeable parameter.

4.5.1   Caching write data

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

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

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

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

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

Table 16 shows the mode default settings for the drive.

4.5.2   Prefetch operation

If the Prefetch feature is enabled, data in contiguous logical blocks on the disc 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.

Since data that is prefetched replaces data already in some buffer segments, the host can limit the amount of prefetch data to optimize system performance. The Max Prefetch field (bytes 8 and 9) limits the amount of prefetch. The drive does not use the Prefetch Ceiling field (bytes 10 and 11).

During a prefetch operation, the drive crosses a cylinder boundary to fetch more data only if Mode parameters page 08h, byte 2, bit 4 is set to 1 (Discontinuity--DISC bit).

When prefetch (read look-ahead) is enabled (enabled by DRA = 0), it operates under the control of ARLA (Adaptive Read Look-Ahead). If the host uses software interleave, ARLA enables prefetch of contiguous blocks from the disc when it senses that a prefetch hit will likely occur, even if two consecutive read operations were not for physically contiguous blocks of data (e.g. "software interleave"). ARLA disables prefetch when it decides that a prefetch hit will not likely occur. If the host is not using software interleave, and if two sequential read operations are not for contiguous blocks of data, ARLA disables prefetch, but as long as sequential read operations request contiguous blocks of data, ARLA keeps prefetch enabled.

4.5.3 Optimizing cache performance for desktop and server applications

Desktop and server applications require different drive caching operations for optimal performance. This means it is difficult to provide a single configuration that meets both of these needs. In a desktop environment, you want to configure the cache to respond quickly to repetitive accesses of multiple small segments of data without taking the time to "look ahead" to the next contiguous segments of data. In a server environment, you want to configure the cache to provide large volumes of sequential data in a non-repetitive manner. In this case, the ability of the cache to "look ahead" to the next contiguous segments of sequential data is a good thing.

The Performance Mode (PM) bit controls the way the drive switches the cache buffer into different modes of segmentation. In "server mode" (PM bit = 0), the drive can dynamically change the number of cache buffer segments as needed to optimize the performance, based on the command stream from the host. In "desktop mode" (PM bit = 1), the number of segments is maintained at the value defined in Mode Page 8, Byte 13, at all times. For additional information about the PM bit, refer to the Unit Attention Parameters page (00h) of the Mode Sense command (1Ah) in the *Fibre Channel Interface Product Manual*, part number 77767496.

5.0   Reliability specifications

The following reliability specifications assume correct host and drive operational interface, including all inter-face 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 with minimum receive eye.
	Less than 1 error in 10^{14} bits transferred with typical receive eye.
	See Section 9.6.4, "Differential PECL input." on page 72
MTBF:	1,400,000 hours
Preventive maintenance:	None required

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

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 paragraph 6.2, "DC power requirements."

• Errors caused by host system failures are excluded from error rate computations.

• Assume random data.

• Default OEM error recovery settings are applied. This includes AWRE, ARRE, full read retries, full write retries and full retry time.

5.1.1   Recoverable Errors

Recovereable 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 recov-ered 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

Unrecoverable Data Errors (Sense Key = 03h) are specified at less than 1 sector in error per 10^{15} bits trans-ferred. Unrecoverable Data Errors resulting from the same cause are treated as 1 error for that block.

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. The total error rate for a loop of devices is the sum of the individual device error rates.

5.2   Reliability and service

You can enhance the reliability of Cheetah K15.4 disc 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 8.3 provides recommended air-flow information.

5.2.1   Mean time between failure (MTBF)

The production disc drive achieves an MTBF of 1,400,000 hours when operated in an environment that ensures the case temperatures specified in Section 6.4.1 are not exceeded. Short-term excursions up to the specification limits of the operating environment will not affect MTBF performance. Operating the drive at case temperatures above these values will adversely affect the drive's ability to meet specifications (see Section 6.4, "Environmental limits").

The MTBF target is specified as device power-on hours (POH) for all drives in service per failure.

The following expression defines MTBF:

$$\text{MTBF per measurement period} = \frac{\text{Estimated power-on operating hours in the period}}{\text{Number of drive failures in the period}}$$

Estimated power-on operating hours means power-on hours per disc drive times the total number of disc drives in service. Each disc drive must have accumulated at least nine months of operation. Data is calculated on a rolling average base for a minimum period of six months.

MTBF is based on the following assumptions:
* 8,760 power-on hours per year
* 250 average on/off cycles per year
* Operating at nominal voltages
* System provides adequate cooling to ensure the case temperatures specified in Section 6.4.1 are not exceeded.

Drive failure means any stoppage or failure to meet defined specifications caused by drive malfunction.

A S.M.A.R.T. predictive failure indicates that the drive is deteriorating to an imminent failure and is considered an MTBF hit.

5.2.2   Preventive maintenance

No routine scheduled preventive maintenance is required.

5.2.3   Hot plugging the drive

Inserting and removing the drive on the FC-AL will interrupt loop operation. The interruption occurs when the receiver of the next device in the loop must synchronize to a different input signal. FC error detection mechanisms, character sync, running disparity, word sync, and CRC are able to detect any error. Recovery is initiated based on the type of error.

The disc drive defaults to the FC-AL Monitoring state, Pass-through state, when it is powered-on by switching the power or hot plugged. The control line to an optional port bypass circuit (external to the drive), defaults to the Enable Bypass state. If the bypass circuit is present, the next device in the loop will continue to receive the output of the previous device to the newly inserted device. If the bypass circuit is not present, loop operation is temporarily disrupted until the next device starts receiving the output from the newly inserted device and regains synchronization to the new input.

The Pass-through state is disabled while the drive performs self test of the FC interface. The control line for an external port bypass circuit remains in the Enable Bypass state while self test is running. If the bypass circuit is present, loop operation may continue. If the bypass circuit is not present, loop operation will be halted while the self test of the FC interface runs.

When the self test completes successfully, the control line to the bypass circuit is disabled and the drive enters the FC-AL Initializing state. The receiver on the next device in the loop must synchronize to output of the newly inserted drive.

If the self-test fails, the control line to the bypass circuit remains in the Enable Bypass state.

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

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

5.2.4   S.M.A.R.T.

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

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

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

Controlling S.M.A.R.T.

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

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

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

Performance impact

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

Maximum processing delay

	On-line only delay DEXCPT = 0, PERF = 1	Fully-enabled delay DEXCPT = 0, PERF = 0
S.M.A.R.T. delay times	42 milliseconds	163 milliseconds

Reporting control

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

Determining rate

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

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

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

Predictive failures

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

5.2.5   Thermal monitor

Cheetah 15K.4 FC drives implement a temperature warning system which:

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

2. Signals the host if the temperature exceeds a user-specified value.

3. Saves a S.M.A.R.T. data frame on the drive which exceeds the threatening temperature value.

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

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

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

The current algorithm implements two temperature trip points. The first trip point is set at 65°C which is the maximum temperature limit according to the drive specification. The second trip point is user-selectable using the Log Select command. The reference temperature parameter in the temperature log page (see Table 1) can be used to set this trip point. The default value for this drive is 65°C, however, you can set it to any value in the range of 0 to 65°C. If you specify a temperature greater than 65°C in this field, the temperature is rounded down to 65°C. A sense code is sent to the host to indicate the rounding of the parameter field.

Table 1: Temperature Log Page (0Dh)

Parameter Code	Description
0000h	Primary Temperature
0001h	Reference Temperature

5.2.6 Drive Self Test (DST)

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

There are two test coverage options implemented in DST:

1. Extended test
2. Short text

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.

5.2.6.1 DST failure definition

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

The following conditions are considered DST failure conditions:

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

Recovered errors will not be reported as diagnostic failures.

5.2.6.2   Implementation

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

5.2.6.2.1   State of the drive prior to testing

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

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

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

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

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

5.2.6.2.2   Invoking DST

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

5.2.6.2.3   Short and extended tests

DST has two testing options:

1. short
2. extended

These testing options are described in the following two subsections.

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

Short test (Function Code: 001b)

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

Extended test (Function Code: 010b)

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

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

5.2.6.2.4 Log page entries

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

1. The Function Code field is set to the same value as sent in the DST command

2. The Self-Test Results Value field is set to Fh

3. The drive will store the log page to non-volatile memory

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

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

5.2.6.2.5 Abort

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

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

5.2.7 Product warranty

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

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

Shipping

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

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

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   AC power requirements

None.

6.2   DC power requirements

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

Table 2:   DC power requirements for operation in 2Gb mode

		Notes	ST3146854FC		ST373454FC		ST336754FC	
			(Amps)	(Amps)	(Amps)	(Amps)	(Amps)	(Amps)
Voltage			+5V	+12V [2]	+5V	+12V [2]	+5V	+12V [2]
Regulation		[5]	±5%	±5% [2]	±5%	±5% [2]	±5%	±5% [2]
Avg idle current DC$\overline{X}$		[1] [7]	0.94	0.78	0.92	0.57	0.94	0.43
Maximum starting current								
(peak DC) DC	3σ	[3]	1.07	1.91	1.08	1.88	1.08	1.58
(peak AC) AC	3σ	[3]	1.36	3.87	1.44	3.98	1.37	3.30
Delayed motor start (max) DC	3σ	[1] [4]	0.60	0.03	0.59	0.03	0.59	0.03
Peak operating current:								
Typical DC$\overline{X}$		[1] [6]	1.03	1.18	1.01	1.00	0.99	0.88
Maximum DC	3σ	[1]	1.04	1.19	1.03	1.02	1.00	0.88
Maximum (peak) DC	3s		1.56	2.84	2.00	2.78	2.02	2.54

[1]   Measured with average reading DC ammeter. Instantaneous +12V current peaks will exceed these values. Power supply at nominal voltage. N (number of drives tested) = 6, 35 Degrees C ambient.

[2]   For +12 V, a −10% tolerance is allowed during initial spindle start but must return to ±5% before reaching 15000 RPM. The ±5% must be maintained after the drive signifies that its power-up sequence has been completed and that the drive is able to accept selection by the host initiator.

[3]   See +12V current profile in Figure 2.

[4]   This condition occurs when the Motor Start option is enabled and the drive has not yet received a Start Motor command.

[5]   See paragraph 6.2.1, "Conducted noise immunity." Specified voltage tolerance includes ripple, noise, and transient response.

[6]   Operating condition is defined as random 8 block reads at 309 I/Os per second for ST3146854FC and

314 I/Os per second for ST373454FC and ST336754FC models. Current and power specified at nominal voltages. Decreasing +5 volt supply by 5% increases 5 Volt current by 2.5%. Decreasing +12 volt supply by 5% increases 12 volt current by 3.0% (2% for ST373454FC and ST336754FC models).

[7]   During idle, the drive heads are relocated every 60 seconds to a random location within the band from three-quarters to maximum track.

General DC power requirement notes.

1.   Minimum current loading for each supply voltage is not less than 1.2% 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.

6.2.1   Conducted noise immunity

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

  +5 V   =   250 mV pp from 0 to 100 kHz to 20 MHz.
 +12 V   =   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.

6.2.2   Power sequencing

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

6.2.3   Current profiles

The +12V and +5V current profiles for the ST3146854FC, ST373454FC, and ST336754FC models are shown below in the following figures.

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

**Figure 2.**   **Typical ST3146854FC drive, 2 Gbit, +12V current profile**

**Figure 3.**   **Typical ST3146854FC drive, 2 Gbit, +5V current profile**

Figure 4. **Typical ST373454FC drive +12V current profile**



**Figure 5.** **Typical ST373454FC drive +5V current profile**

Figure 6.   **Typical ST336754FC drive +12V current profile**

Figure 7.   **Typical ST336754FC drive +5V current profile**

6.3   Power dissipation

ST3146854FC

Typical power dissipation under idle conditions in 2Gb operation is 13.96 watts (47.6 BTUs per hour).

To obtain operating power for typical random read operations, refer to the following I/O rate curve (see Figure 8). Locate the typical I/O rate for a drive in your system on the horizontal axis and read the corresponding +5 volt current, +12 volt current, and total watts on the vertical axis. To calculate BTUs per hour, multiply watts by 3.4123.


Figure 8.   ST3146854FC DC current and power vs. input/output operations per second

ST373454FC

Typical power dissipation under idle conditions in 2Gb operation is 11.49 watts (39.21 BTUs per hour).

To obtain operating power for typical random read operations, refer to the following I/O rate curve (see Figure 8). Locate the typical I/O rate for a drive in your system on the horizontal axis and read the corresponding +5 volt current, +12 volt current, and total watts on the vertical axis. To calculate BTUs per hour, multiply watts by 3.4123.
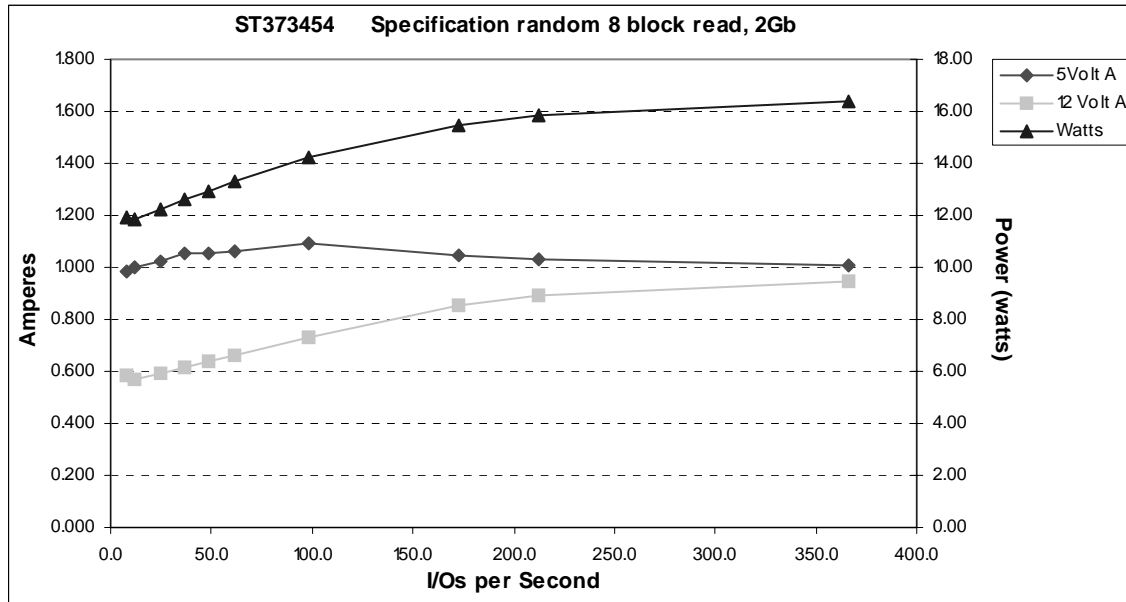

Figure 9. **ST373454FC DC current and power vs. input/output operations per second**

ST336754FC

Typical power dissipation under idle conditions in 2Gb operation is 9.88 watts (33.7 BTUs per hour).

To obtain operating power for typical random read operations, refer to the following I/O rate curve (see Figure 8). Locate the typical I/O rate for a drive in your system on the horizontal axis and read the corresponding +5 volt current, +12 volt current, and total watts on the vertical axis. To calculate BTUs per hour, multiply watts by 3.4123.
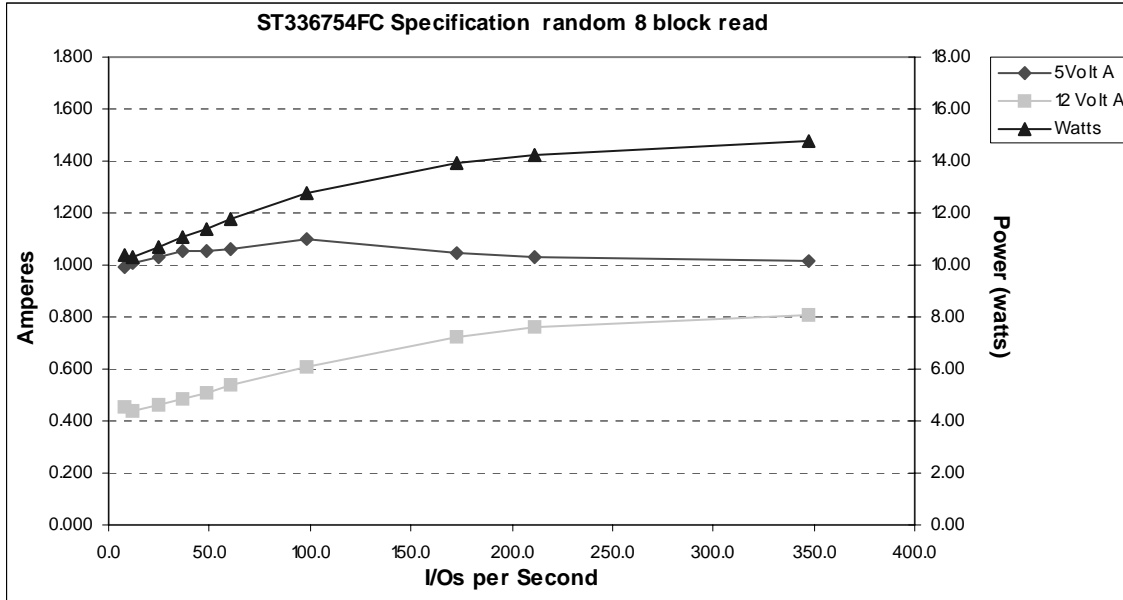

Figure 10.   ST336754FC DC current and power vs. input/output operations per second

6.4   Environmental limits

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

6.4.1   Temperature

a. Operating

The maximum allowable continuous or sustained HDA case temperature for the rated MTBF is 122°F (50°C) The maximum allowable HDA case temperature is 60°C. Occasional excursions of HDA case temperatures above 122°F (50°C) or below 41°F (5°C) may occur without impact to specified MTBF. Continual or sustained operation at HDA case temperatures outside these limits may degrade MTBF.

Provided the HDA case temperatures limits are met, the drive meets all specifications over a 41°F to 131°F (5°C to 55°C) drive ambient temperature range with a maximum temperature gradient of 36°F (20°C) per hour. Air flow may be needed in the drive enclosure to keep within this range (see Section 8.3). Operation at HDA case temperatures outside this range may adversely affect the drives ability to meet specifications. To confirm that the required cooling for the electronics and HDA case is provided, place the drive in its final mechanical configuration, perform random write/read operations and measure the HDA case temperature after it has stabilized.

b. Non-operating

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

Figure 11.   Locations of the HDA temperature check point

6.4.2   Relative humidity

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

a. Operating

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

b. Non-operating

5% to 95% non-condensing relative humidity.

6.4.3   Effective altitude (sea level)

a. Operating

–1,000 to +10,000 feet (–305 to +3,048 meters)

b. Non-operating

–1,000 to +40,000 feet (–305 to +12,210 meters)

6.4.4   Shock and vibration

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

The limits of shock and vibration defined within this document are specified with the drive mounted by any of the four methods shown in Figure 12, and in accordance with the restrictions of Section 8.5. Orientation of the side nearest the LED may be up or down.

6.4.4.1   Shock

a. Operating—normal

The drive, as installed for normal operation, shall operate error free while subjected to intermittent shock not exceeding 15 Gs at a maximum duration of 11 msec (half sinewave). The drive, as installed for normal operation, shall operate error free while subjected to intermittent shock not exceeding 65 Gs at a maximum duration of 2 msec (half sinewave). Shock may be applied in the X, Y, or Z axis.

---

b. Operating—abnormal

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

c. Non-operating

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

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

The drive subjected to nonrepetitive shock not exceeding 250 Gs at a maximum duration of 2 msec (half sinewave) does not exhibit device damage or performance degradation. Shock may be applied in the X, Y, or Z axis.

The drive subjected to nonrepetitve shock not exceeding 100 Gs at a maximum duration of 0.5 msec (half sinewave) does not exhibit device damage or performance degradation. Shock may be applied in the X, Y, or Z axis.

d. Packaged

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

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

Drives packaged in single or multipacks with a gross weight of 20 pounds (8.95 kg) or less by Seagate for general freight shipment shall withstand a drop test from 48 inches (1,070 mm) against a concrete floor or equivalent.

Z
Y
X

X
Z
Y

Figure 12.   Recommended mounting

6.4.4.2   Vibration

a. Operating—normal

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

 5-500 Hz @ 0.5 G (zero to peak)

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

Operating normal translational random flat profile

10 - 400 Hz   0.4 g RMS

b. Operating—abnormal

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

 15 minutes of duration at major resonant frequency

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

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

Operating abnormal translational random flat profile

10 - 400 Hz   1.2 g RMS

c. Non-operating

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

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

 5-22 Hz @ 0.040 inches (1.02 mm) displacement (zero to peak)
 22-500 Hz @ 2.00 G (zero to peak)

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

Non-operating translational random flat profile

10 - 400 Hz   1.2 g RMS

6.4.5   Air cleanliness

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

6.4.6   Corrosive environment

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

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

6.4.7　　　Acoustics

Sound power during idle mode shall be 3.6 bels typical when measured to ISO 7779 specification.

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

6.4.8　　　Electromagnetic susceptibility

See Section 2.1.1.1.

6.5 Mechanical specifications

The following nominal dimensions are exclusive of the decorative front panel accessory. However, dimensions of the front panel are shown in figure below. Refer to Figure 13 for detailed mounting configuration dimensions. See Section 8.5, "Drive mounting."

Height:	1.000 in	25.4 mm
Width:	4.000 in	101.6 mm
Depth:	5.75 in	146.05 mm
Weight:	1.6 pounds	0.73 kilograms

Notes:

[1] Mounting holes are 6-32 UNC 2B, three on each side and four on the bottom. Max screw penetration into side of drive is 0.15 in. (3.81 mm). Max screw tightening torque is 6.0 in-lb (3.32 nm) with minimum full thread engagement of 0.12 in. (3.05 mm).

Dimension Table

	Inches	Millimeters
A	1.028 max	26.10 max
B	5.787 max	147.00 max
C	4.000 ± .010	101.60 ± .25
D	3.750 ± .010	95.25 ± .25
E	0.125 ± .010	3.18 ± .25
F	1.750 ± .010	44.45 ± .25
H	1.122 ± .020	28.50 ± .50
J	4.000 ± .010	101.60 ± .25
K	0.250 ± .010	6.35 ± .25
L	1.638 ± .010	41.60 ± .25
M	0.181	4.60
N	.040	1.02
P	1.625 ± .020	41.28 ± .50
R	1.618	41.10
S	0.276 ± .040	7.00 ± 1.02
T	.015 max	0.38 max
U	.015 max	0.38 max

Figure 13. Mounting configuration dimensions

7.0   Defect and error management

The drive, as delivered, complies with this product manual. The read error rates and specified storage capacities are not dependent upon use of defect management routines by the host (initiator).

Defect and error management in the SCSI protocol involves the drive internal defect/error management and FC-AL system error considerations (errors in communications between the initiator and the drive). Tools for use in designing a defect/error management plan are briefly outlined in this section. References to other sections are provided when necessary.

7.1   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 *Fibre Channel Interface Manual*. Also, more information on the drive Error Recovery philosophy is presented in the *Fibre Channel Interface Manual*.

7.2   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 *Fibre Channel Interface Manual*. The drive implements selectable error recovery time limits required in video applications.

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

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

Table 3 equates the read and write retry count with the maximum possible recovery time for read and write recovery of individual LBAs. The times given do not include time taken to perform reallocations. Reallocations are performed when the ARRE bit (for reads) or AWRE bit (for writes) is one, the RC bit is zero, and the recovery time limit for the command has not yet been met. Time needed to perform reallocation is not counted against the recovery time limit.

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

Table 3:   Read and write retry count maximum recovery times [1]

Read retry count	Maximum recovery time per LBA (cumulative, msec)		Write retry count	Maximum recovery time per LBA (cumulative, msec)
0	51.87		0	23.94
1	59.85		1	35.91
2	203.49		2	55.86
3	235.41		3	67.83
4	271.32		4	119.7
5	283.29		5 (default)	147.63
6	315.21			
7	395.01			
8	454.86			
9	486.78			
10	522.69			
11 (default)	1,264.83			

[1]   These values are subject to change.

Setting these retry counts to a value below the default setting could result in an increased unrecovered error rate which may exceed the value given in this product manual. A setting of zero (0) will result in the drive not performing error recovery.

For example, suppose the Read/Write Recovery page has the RC bit set to 0, read retry count set to 4, and the recovery time limit field (Mode Sense page 01, bytes 10 and 11) set to FF FF hex (maximum). A four LBA Read command is allowed to take up to 271.32 msec recovery time for each of the four LBAs in the command. If the recovery time limit is set to 00 C8 hex (200 msec decimal) a four LBA read command is allowed to take up to 200 msec for all error recovery within that command. The use of the Recovery Time Limit field allows finer granularity on control of the time spent in error recovery. The recovery time limit only starts counting when the drive is executing error recovery and it restarts on each command. Therefore, each command's total recovery time is subject to the recovery time limit. Note: A recovery time limit of 0 will use the drive's default value of FF FF. Minimum recovery time limit is achieved by setting the Recovery Time Limit field to 00 01.

7.3   FC-AL system errors

Information on the reporting of operational errors or faults across the interface is given in the *Fibre Channel Interface Manual*. The FCP 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 *Fibre Channel 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.

7.4   Background Media Scan

Background Media Scan (BGMS) is a self-initiated media scan. It performs sequential reads across the entire pack of the media while the drive is idle. In RAID arrays, BGMS allows hot spare drives to be scanned for defects prior to being put into service by the host system. On regular duty drives, if the host system makes use

of the BGMS Log Page, it can avoid placing data in suspect locations on the media. Unrecovered error sites found during BGMS will be logged and recovered error sites will be logged or reallocated per ARRE/AWRE settings.

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

Since the background scan functions are only done during idle periods, BGMS causes a negligible impact to system performance. Background media scan will begin after 500ms of idle time. Other features that normally use idle time to function will function normally because BGMS functions for bursts of 800ms and then suspends activity to allow other background functions to operate. Additionally, BGMS operates in segments of 400 block reads, checking for host commands upon completion of each of the block segments.

The BGMS function will complete any BGMS-initiated error recovery prior to returning to service host-initiated commands. Overhead associated with a return to host-servicing activity from BGMS only impacts the first command that interrupted BGMS. This results in a typical delay of about 4 ms.

7.4.1   Media Pre-Scan

Media Pre-Scan is an optional feature that allows the drive to handle unrecovered media errors that would otherwise have been found by the host system during critical data accesses early in the integration process at customer sites. This should reduce integration DPPM (defective parts per million) by identifying and repairing media errors that arose during shipping and handling. Media Pre-Scan replaces all host-initiated Write commands with Write-Verify commands when the drive is accessing areas of the media yet untouched by BGMS-initiated scanning activity.

The drive will attempt to run BGMS on the entire pack of the drive. This will result in increasingly large portions of the drive having been accessed, provided there is adequate idle time to allow BGMS to run. During the Pre-Scan period, all Write commands that are in areas not yet scanned by BGMS are converted to Write-Verify commands. To expedite the scan of the full pack and the subsequently exit from the Pre-Scan period, BGMS will begin immediately when the drive goes to idle during the Pre-Scan period, as opposed to waiting 500ms as is done under non-Pre-Scan conditions.

8.0   Installation

Cheetah 15K.4 FC disc drive installation is a plug-and-play process. There are no jumpers, switches, or terminators on the drive. Simply plug the drive into the host's 40-pin Fibre Channel backpanel connector (FC-SCA)—no cables are required. See Section 9.5 for additional information about this connector.

Use the FC-AL interface to select drive ID and all option configurations for devices on the loop.

If multiple devices are on the same FC-AL and physical addresses are used, set the device selection IDs (SEL IDs) on the backpanel so that no two devices have the same selection ID. This is called the hard assigned arbitrated loop physical address (AL_PA). There are 125 AL_PAs available (see Table 25). If you set the AL_PA on the backpanel to any value other than 0, the device plugged into the backpanel's SCA connector inherits this AL_PA. In the event you don't successfully assign unique hard addresses (and therefore have duplicate selection IDs assigned to two or more devices), the FC-AL generates a message indicating this condition. If you set the AL_PA on the backpanel to a value of 0, the system issues a unique soft-assigned physical address automatically.

Loop initialization is the process used to verify or obtain an address. The loop initialization process is performed when power is applied to the drive, when a device is added or removed from the Fibre Channel loop, or when a device times out attempting to win arbitration.

- Set all option selections in the connector prior to applying power to the drive. If you change options after applying power to the drive, recycle the drive power to activate the new settings.

- It is not necessary to low-level format this drive. The drive is shipped from the factory low-level formatted in 512-byte logical blocks. You need to reformat the drive only if you want to select a different logical block size.

8.1   Drive ID/option selection

All drive options are made through the interface connector (J1). Table provides the pin descriptions for the 40-pin Fibre Channel single connector (J1).

8.2 LED connections

A connector, J6, is provided on the printed circuit board assembly (PCBA) to provide port bypass, drive active, and drive fault LED connections (see Figure 15). See Sections 9.5.7, 9.5.8, and 9.5.9 for descriptions of LED functions.


Figure 14.   Physical interface

[1]   The drive has a 2.2K ohm resistor in series with this LED driver. Tie the minus side of an external high-efficiency LED (i.e., 2ma) to this pin. Connect the plus side of the LED to +5V.

[2]   An external current-limiting resistor is required when connecting an LED to this pin. The minus side of the resistor/LED combination is connected to this pin. Connect the plus side to +5V.

[3]   Jumper storage location (across pins 2 and 4).

Figure 15.   LED indicator connector

8.2.1 J6 connector requirements

Recommended mating connector part number: Berg receptacle, 6-position, Berg part number 690-006.

8.3 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.

8.4 Cooling

The host enclosure must dissipate heat from the drive. You should confirm that the host enclosure is designed to ensure that the drive operates within the temperature measurement guidelines described in Section 6.4.1. In some cases, forced airflow may be required to keep temperatures at or below the temperatures specified in Section 6.4.1.

If forced air is necessary, possible air-flow patterns are shown in Figure 16. The air-flow patterns are created by fans either forcing or drawing air as shown in the illustrations. Conduction, convection, or other forced airflow patterns are acceptable as long as the temperature measurement guidelines of Section 6.4.1 are met.

Above unit

Under unit

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

Above unit

Under unit

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

Figure 16.   Air flow

8.5 Drive mounting

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

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

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

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

Note. Before mounting the drive in any kind of 3.5-inch to 5.25-inch adapter frame, verify with Seagate Technology that the drive can meet the shock and vibration specifications given herein while mounted in such an adapter frame. Adapter frames that are available may not have a mechanical structure capable of mounting the drive so that it can meet the shock and vibration specifications listed in this manual.

8.6 Grounding

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

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

9.0   Interface requirements

This section partially describes the interface requirements as implemented on Cheetah 15K.4 FC drives. Additional information is provided in the *Fibre Channel Interface Manual* (part number 77767496).

9.1   FC-AL features

This section lists the Fibre Channel-specific features supported by Cheetah 15K.4 FC drives.

9.1.1   Fibre Channel link service frames

Table 4 lists the link services supported by Cheetah 15K.4 FC drives.

Table 4:   Link services supported

Type of frame	Link service
Basic link service frames	Abort Sequence (ABTS)
Basic link service reply frames	Basic_Accept (BA_ACC) Basic_Reject (BA_RJT)
Extended link service frames	N_Port Login (PLOGI) Fabric Login (FLOGI) Logout (LOGO) Process Login (PRLI) Process Logout (PRLO) Read Link Status (RLS) Fabric Address Notification (FAN) Port Discovery (PDISC) Address Discovery (ADISC) Third-party Process Logout (TRPLO)
Extended link service reply frames	Accept (ACC) Link Service Reject (LS_RJT)
Fibre Channel Services	Register FC-4 Types (RFT_ID)

9.1.2   Fibre Channel task management functions

Table 5 lists the Fibre Channel SCSI Fibre Channel Protocol (FC SCSI FCP) task management functions supported.

Table 5:   Fibre Channel SCSI FCP task management functions

Task name	Supported
Terminate task	No
Clear ACA	Yes
Target reset	Yes
Clear task set	Yes
Abort task set	Yes

9.1.3   Fibre Channel task management responses

Table 6 lists the FC SCSI FCP response codes returned for task management functions supported.

Table 6:   FC SCSI FCP response codes

Function name	Response code
Function complete	00
Function not supported	04
Function reject	05

9.1.4　　　Fibre Channel port login

Table 7 identifies the required content of the N_Port Login (PLOGI) payload from an initiator.

Table 7:　　N_Port login (PLOGI) payload

Bytes																	
0-15	03	00	00	00	09	09	BB	BB	CF	XX	FS	FS	XX	XX	XX	XX	Common
16-31	XX	XX	XX	XX	PN	PN	PN	PN	PN	PN	PN	PN	NN	NN	NN	NN	
32-35	NN	NN	NN	NN													
36-47					XX	XX	XX	XX	XX	XX	XX	XX	XX	XX	XX	XX	Class 1
48-51	XX	XX	XX	XX													
52-63					XX	XX	XX	XX	XX	XX	XX	XX	XX	XX	XX	XX	Class 2
64-67	XX	XX	XX	XX													
68-79					SO	SO	IC	IC	XX	XX	FS	FS	XX	CS	XX	XX	Class 3
80-83	OS	OS	XX	XX													
84-95					XX	XX	XX	XX	XX	XX	XX	XX	XX	XX	XX	XX	Reserved
96-99	XX	XX	XX	XX													
100-111					XX	XX	XX	XX	XX	XX	XX	XX	XX	XX	XX	XX	Vendor
112-115	XX	XX	XX	XX													Version

X　　Indicates a four-bit (hex) field is not checked.

x　　Indicates a single bit is not checked.

BB　　BB-Credit. This field is not checked. The FC-AL drive uses BB-Credit of zero (0).

CF　　Common features. This binary field selects the common features requested by the initiator login.

	MSB	Continuously increasing offset	Must = 1
		Random relative offset	Not checked. Port Login Accept will return a 0—not supported.
		Valid version level	x
		N_Port/F_Port	Must = 0, N_Port
		Alternate credit model	Must = 1
		Other bits reserved	xxx XX

FS　　Receive buffer field size. The FS field in the common and Class 3 parameters is checked for the range 128 < FS < 2,112 and a multiple of four bytes. For multiple frame sequences, all frames but the last frame of the sequence must be this size. Only the receive buffer field size in the Class 3 parameters is used.

PN　　Port name (initiator's)—saved with the login parameters. If a change of the port name/AL_PA address association is detected during a Port DISCovery, and implicit logout occurs and the initiator returns a LS_RJT.

NN　　Node name. The node name is not checked or saved by the drive.

SO　　Service options Class 3 only.

	MSB	Class valid	Must = 1
		Intermix	x
		Stacked connection req.	xx
		Sequential delivery	x
		Other bits reserved	xxx XX

IC　　Initiator control

	MSB	XID reassign	xx
		Proc Assc	10 or 11 causes the login to be rejected. Other values are accepted.
		Other bits	XXX

CS　　Concurrent sequences　　　　　　Must be a value greater than 0.

OS　　Open sequences per exchange　　　Must be a value greater than 0.

9.1.5   Fibre Channel port login accept

Table 8 identifies the N_Port Login access payload values.

Table 8:   N_Port Login Accept (ACC) payload

Bytes																	
0-15	02	00	00	00	09	09	00	00	88	00	FS	FS	00	FF	00	01	Common
16-31	00	00	01	F4	2P	00	CC	CC	CC	UI	UI	UI	20	00	CC	CC	
32-35	CC	UI	UI	UI													
36-47					00	00	00	00	00	00	00	00	00	00	00	00	Class 1
48-51	00	00	00	00													
52-63					00	00	00	00	00	00	00	00	00	00	00	00	Class 2
64-67	00	00	00	00													
68-79					80	00	00	00	00	00	FS	FS	00	FF	00	00	Class 3
80-83	00	01	00	00													
84-95					00	00	00	00	00	00	00	00	00	00	00	00	Reserved
96-99	00	00	00	00													
100-111					00	00	00	00	00	00	00	00	00	00	00	00	Vendor -
112-115	00	00	00	00													Version

CC	Bytes (22, 23 and 24) and (30, 31 and 32) are currently factory set to 00 04 CF respectively (subject to change).
FS	Receive buffer field size. The drive returns and uses the receive buffer size from the N_Port Login Class 3 receive buffer.
UI	Unique identifier. This 24-bit field is uniquely assigned to the drive. This same UI appears in the Port Name and Node Name fields.
P	Byte port identifier field.
	0	P_LOGI received on Node.	
	1	P_LOGI received on Port A.	
	2	P_LOGI received on Port B.	

9.1.6   Fibre Channel Process Login

Table 9 lists the process login payload data.

Table 9:   Process Login (PLRI) payload

Bytes																
0-15	20	10	00	14	08	00	20	00	XX	XX	XX	XX	XX	XX	XX	XX
16-19	00	00	00	22												

XX   Indicates fields that are not used.

9.1.7 Fibre Channel Process Login Accept

Table 10 lists Cheetah 15K.4 FC process login accept payload data.

Table 10: Process Login Accept (ACC) payload

Bytes																
0-15	02	10	00	14	08	00	21	00	00	00	00	00	00	00	00	00
16-31	00	00	00	12												

9.1.8 Fibre Channel fabric login

Table 11 lists the fabric login payload from the drive.

Table 11: Fabric Login (FLOGI) payload

Bytes																	
0-15	04	00	00	00	09	09	00	00	08	00	08	40	00	00	00	00	Common
16-31	00	00	02	F4	2P	00	CC	CC	CC	UI	UI	UI	02	00	CC	CC	
32-35	CC	UI	UI	UI													
36-47					00	00	00	00	00	00	00	00	00	00	00	00	Class 1
48-51	00	00	00	00													
52-63					00	00	00	00	00	00	00	00	00	00	00	00	Class 2
64-67	00	00	00	00													
68-79					80	00	00	00	00	00	08	40	00	00	00	00	Class 3
80-83	00	00	00	00													
84-95					00	00	00	00	00	00	00	00	00	00	00	00	Reserved
96-99	00	00	00	00													
100-111					00	00	00	00	00	00	00	00	00	00	00	00	Vendor -
112-115	00	00	00	00													Version

CC Bytes (22, 23 and 24) and (30, 31 and 32) are currently factory set to 00 04 CF respectively (subject to change).

UI Unique identifier. This 24-bit field is uniquely assigned to the drive. This same UI appears in the Port Name and Node Name fields.

P Port identifier field.

    1   FLOGI originated on Port A.

    2   FLOGI originated on Port B.

9.1.9 Fibre Channel fabric accept login

Table 12 lists the required content of the Fabric Login Accept (ACC) payload from the fabric.

Table 12: Fabric Login Accept (ACC) payload

Bytes																	
0-15	02	00	00	00	09	09	BB	BB	CF	XX	FS	FS	R_	A_	T0	V_	Common
16-31	E_	D_	T0	V_	PN	PN	PN	PN	PN	PN	PN	PN	NN	NN	NN	NN	
32-35	NN	NN	NN	NN													
36-47					XX	XX	XX	XX	XX	XX	XX	XX	XX	XX	XX	XX	Class 1
48-51	XX	XX	XX	XX													
52-63					XX	XX	XX	XX	XX	XX	XX	XX	XX	XX	XX	XX	Class 2
64-67	XX	XX	XX	XX													
68-79					SO	SO	xx	xx	XX	XX	FS	FS	XX	xx	XX	XX	Class 3
80-83	OS	OS	XX	XX													
84-95					XX	XX	XX	XX	XX	XX	XX	XX	XX	XX	XX	XX	Reserved
96-99	XX	XX	XX	XX													
100-111					XX	XX	XX	XX	XX	XX	XX	XX	XX	XX	XX	XX	Vendor -
112-115	XX	XX	XX	XX													Version

X Indicates a four-bit (hex) field is not checked.

x Indicates a single bit is not checked.

BB BB-Credit. This field is not checked. The FC-AL drive uses BB-Credit of zero (0).

CF Common features. This binary field selects the common features requested by the fabric login.

MSB	Continuously increasing offset	x
	Random relative offset	x
	Valid version level	x
	N_Port/F_Port	Must = 1, F_Port
	Alternate credit model	Must = 1
	Other bits reserved	xxx XX

FS Receive buffer field size. The FS field in the common and Class 3 parameters is checked for the range 128 < FS < 2,112 and a multiple of four bytes. The receive buffer field size in the Class 3 parameters is used. The drive uses the lower FS of Fabric Login Accept or N_Port Login when sending frames to an initiator.

PN Port Name. The fabric port name is saved with the login parameters. If a change of the port name is detected during a FAN, an implicit logout occurs and a LS_RJT is returned to the fabric.

NN Node Name. The drive does not check or save the node name.

SO Service Options—Class 3 only.

MSB	Class valid	Must = 1
	Intermix	x
	Stacked connection req.	xx
	Sequential delivery	Must = 1
	Other bits reserved	xxx XX

9.1.10 Fibre Channel Arbitrated Loop options

Table 13 lists the FC-AL options supported by Cheetah 15K.4 FC drives.

Table 13: FC-AL options supported

Option	Supported
OPEN Half Duplex	Accepted from another device.
OPEN Full Duplex	Sent to open another device. Accepted from another device.
Private Loop	Yes
Public Loop	Yes
Old Port State	No
Loop Position	Yes
Loop Position Report	Yes

9.2 Dual port support

Cheetah 15K.4 FC drives have two independent FC-AL ports. These ports may be connected on independent loops or on the same loop. Port A and Port B may be connected in any order or combination.

• If both ports are connected on independent loops and hard addressing is used, the drive interface address is selected through the interface connector, both ports will seek the same loop address. If no conflict, both ports will have the same loop address.

• If both ports are connected in the same loop and hard addressing is used, at least one port will attempt taking a soft address to prevent an address conflict.

Note. When a Cheetah 15K.4 FC drive is connected in loops with previous Seagate FC drive products:

Barracuda 4LP FC (ST32171FC, ST34371FC, and ST34571FC)
Barracuda 9FC (ST19171FC)
Cheetah 4LP FC (ST34501FC)
Cheetah 9FC (ST19101FC)

the connection of Port A and B for these products must follow the requirements in their product manuals.

Subject to buffer availability, the Cheetah 15K.4 FC drives support:

• Concurrent port transfers—The drive supports receiving transfers on both ports at the same time when the ports are on independent loops.

• Full duplex—The drive supports sending FCP_Data, FCP_RSP, FCP_XFR_RDY and ELS transfers while receiving frames on both ports.

9.3   SCSI commands supported

Table 14 lists the SCSI commands supported by Cheetah 15K.4 FC drives.

Table 14:   Supported commands

Command code	Supported [4]	Command name
00h	Y	Test unit ready
01h	Y	Rezero unit
03h	Y	Request sense
	Y	Extended sense
	Y	Field pointer bytes
	Y	Actual retry count bytes
04h	Y	Format unit [1]
07h	Y	Reassign blocks
08h	Y	Read
0Ah	Y	Write
0Bh	Y	Seek
12h	Y	Inquiry
	Y	Vital product data page (00h)
	Y	Unit serial number page (80h)
	Y	Implemented operating def. page (81h)
	Y	Device Identification page (83h)
	Y	Firmware numbers page (C0h)
	Y	Date code page (C1h)
	Y	Jumper settings page (C2h)
	Y	Device Behavior page (C3h)
15h	Y	Mode select (same pages as Mode Sense command shown below) [3]
16h	Y	Reserve
	N	3rd party reserved
	N	Extent reservation
17h	Y	Release
18h	N	Copy
1Ah	Y	Mode sense
	Y	Unit attention page (00h)
	Y	Error recovery page (01h)
	Y	Disconnect/reconnect control (page 02h)
	Y	Format page (03h)
	Y	Rigid disc drive geometry page (04h)

Table 14:   Supported commands *(continued)*

Command code	Supported [4]	Command name
	Y	Verify error recovery page (07h)
	Y	Caching parameters page (08h)
	Y	Control mode page (0Ah)
	Y	Fibre Channel Interface Control page (19h)
	Y	Power control page (1Ah)
	Y	Information exceptions control page (1Ch)
1Bh	Y	Start unit/stop unit
1Ch	Y	Receive diagnostic results
	Y	Supported diagnostics pages
	Y	Translate page
	Y	Enclosure services page
1Dh	Y	Send diagnostics page
	Y	Supported diagnostics pages
	Y	Translate page
25h	Y	Read capacity
28h	Y	Read extended
	Y	Disable page out
	Y	Force unit access
	N	Relative address
2Ah	Y	Write extended
	Y	Disable page out
	Y	Force unit access
	N	Relative address
2Bh	Y	Seek extended
2Eh	Y	Write and verify
	Y	Disable page out
	Y	Byte check
	N	Relative address
2Fh	Y	Verify
	Y	Disable page out
	Y	Byte check
	N	Relative address
30h	N	Search data high
31h	N	Search data equal
32h	N	Search data low

Table 14: Supported commands *(continued)*

Command code	Supported [4]	Command name
33h	N	Set limits
34h	N	Prefetch
35h	Y	Synchronize cache
36h	N	Lock-unlock-cache
37h	Y	Read defect data
39h	N	Compare
3Ah	N	Copy and verify
3Bh	Y	Write buffer
	Y	Write combined header and data mode (0)
	Y	Write data mode (2)
	N	Download microcode mode (4)
	Y	Download microcode and save modes (5)
	N	Download microcode with offsets mode (6)
	Y	Download microcode with offsets and save mode (7)
	Y	Firmware download option [2]
3Ch	Y	Read buffer
	Y	Read combined header and data mode (0)
	Y	Read data mode (2)
	Y	Read descriptor mode (3)
3Eh	Y	Read long
3Fh	Y	Write long
40h	N	Change definition
41h	Y	Write same
	N	PBdata
	N	LBdata
42-4Bh	N	Not used
4Ch	Y	Log Select
4Dh	Y	Log Sense
	Y	Support Log page (00h)
	Y	Write Error Counter page (02h)
	Y	Read Error Counter page (03h)
	N	Read Reverse Error Counter page (04h)
	Y	Verify Error Counter page (05h)
	Y	Non-medium Error Counter page (06h)
	Y	Temperature page (0Dh)

Table 14:   Supported commands *(continued)*

Command code	Supported [4]	Command name
	N	Application Client page (0Fh)
	Y	Self Test Results page (10h)
	Y	Cache Statistics Counter page (37h)
	Y	Factory Log page (3Eh)
4E-4Fh	N	Not used
50h	N	XD write
51h	N	XP write
52h	N	XD read
53-54h	N	Not used
55h	Y	Mode Select (10) [3]
56h	Y	Reserved (10)
	Y	3rd party reserve
	N	Extent reservation
57h	Y	Released (10)
58-59h	N	Not used
5Ah	Y	Mode Sense (10) [3]
5B-5Dh	N	Not used
5E	A	Persistent reserve in
5F	A	Persistent reserve out
60-7Fh	N	Not used
80h	N	XD write extended
81h	N	Rebuild
82h	N	Regenerate
83-8Fh	N	Not used
A0h	Y	Report LUNS
C0-DFh	N	Not used
EO-FFh	N	Not used

[1]   Cheetah 15K.4 FC drives can format to any multiple of four bytes per logical block in the range 512 to 704 bytes.

[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.

9.3.1 Inquiry data

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

Table 15: Cheetah 15K.4 FC inquiry data

Bytes	Data (hex)																
0-15	00	00	03**	12	8B	00	PP	0A	53	45	41	47	41	54	45	20	Vendor ID
16-31	[53	54	33	31	34	36	38	35	34	46	43][1]	20	20	20	20	20	Product ID
32-47	R#	R#	R#	R#	S#	S#	S#	S#	S#	S#	S#	S#	00	00	00	00	
48-63	00	00	00	00	00	00	00	00	00	00	00	00	00	00	00	00	
64-79	00	00	00	00	00	00	00	00	00	00	00	00	00	00	00	00	
80-95	00	00	00	00	00	00	00	00	00	00	00	00	00	00	00	00	
96-111	00	43	6F	70	79	72	69	67	68	74	20	28	63	29	20	32*	*Copyright
112-127	30*	30*	32*	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).

**   02 = SCSI-2 implemented with some SCSI-3 features (default).

      03 = SCSI-3 implemented.

PP   50 = Inquiry data for an Inquiry command received on Port A.

      70 = 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 18 through 26 reflect model of drive. The table above shows the hex values for Model ST3146854FC.

      Refer to the values below for the values of bytes 18 through 26 of you particular model:

      ST373454FC   53 54 33 37 33 34 35 34 46 43

      ST336754FC   53 54 33 33 36 37 35 34 46 43

9.3.2 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 16, refer to Mode page 81, in the row entitled "CHG." These are hex numbers representing the changeable values for Mode page 81. Note in columns 5 and 6 (bytes 04 and 05), there is 00h which indicates that in bytes 04 and 05 none of the bits are changeable. Note also that bytes 06, 07, 09, 10, and 11 are not changeable, because those fields are all zeros. In byte 02, hex value FF equates to the binary pattern 11111111. If there is a zero in any bit position in the field, it means that bit is not changeable. Since all of the bits in byte 02 are ones, all of these bits are changeable.

The changeable values list can only be changed by downloading new firmware 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 *Fibre Channel Interface Manual*).

Definitions:

SAV = Current saved value.
DEF = Default value. Standard OEM drives are shipped configured this way.
CHG = Changeable bits; indicates if default value is changeable.

Table 16:   Mode Sense data saved, default and changeable values for ST3146854FC drives

```
MODE SENSE (10 BYTE) HEADER DATA
00 ae 00 10 00 00 00 08 11 17 73 30 00 00 02 00


MODE SENSE PAGES DATA
SAV 81 0a c0 0b ff 00 00 00 05 00 ff ff
DEF 81 0a c0 0b ff 00 00 00 05 00 ff ff
CHG 81 0a ff ff 00 00 00 00 ff 00 ff ff


SAV 82 0e 80 80 00 00 00 00 00 00 05 bc 00 00 00 00
DEF 82 0e 80 80 00 00 00 00 00 00 05 bc 00 00 00 00
CHG 82 0e ff ff 00 00 00 00 00 00 ff ff 00 00 00 00


SAV 83 16 56 20 00 00 00 30 00 00 02 c2 02 00 00 01 00 aa 00 91 40 00 00 00
DEF 83 16 56 20 00 00 00 30 00 00 02 c2 02 00 00 01 00 aa 00 91 40 00 00 00
CHG 83 16 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00


SAV 84 16 00 c6 88 08 00 00 00 00 00 00 00 00 00 00 00 00 00 3a a7 00 00
DEF 84 16 00 c6 88 08 00 00 00 00 00 00 00 00 00 00 00 00 00 3a a7 00 00
CHG 84 16 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00


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


SAV 88 12 14 00 ff ff 00 00 ff ff ff ff 00 20 00 00 00 00 00 00
DEF 88 12 14 00 ff ff 00 00 ff ff ff ff 00 20 00 00 00 00 00 00
CHG 88 12 b5 00 00 00 ff ff ff ff 00 00 a0 ff 00 00 00 00 00 00


SAV 8a 0a 02 00 00 00 00 00 00 00 08 4d
DEF 8a 0a 02 00 00 00 00 00 00 00 08 4d
CHG 8a 0a 03 f1 00 00 00 00 00 00 00 00


SAV 99 06 00 00 00 00 00 00
DEF 99 06 00 00 00 00 00 00
CHG 99 06 00 ff 00 00 00 00


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


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


SAV 80 06 00 00 0f 00 00 00
DEF 80 06 00 00 0f 00 00 00
CHG 80 06 b7 40 0f 00 00 00


READ CAPACITY DATA
11 17 73 2f 00 00 02 00
```

Table 17:   Mode Sense data default and changeable values for ST373454FC drives

```
MODE SENSE (10 BYTE) HEADER DATA
00 ae 00 10 00 00 00 08 08 8b b9 98 00 00 02 00


MODE SENSE PAGES DATA
SAV 81 0a c0 0b ff 00 00 00 05 00 ff ff
DEF 81 0a c0 0b ff 00 00 00 05 00 ff ff
CHG 81 0a ff ff 00 00 00 00 ff 00 ff ff


SAV 82 0e 80 80 00 00 00 00 00 00 01 3c 00 00 00 00
DEF 82 0e 80 80 00 00 00 00 00 00 01 3c 00 00 00 00
CHG 82 0e ff ff 00 00 00 00 00 00 ff ff 00 00 00 00


SAV 83 16 2b 10 00 00 00 18 00 00 02 c2 02 00 00 01 00 aa 00 91 40 00 00 00
DEF 83 16 2b 10 00 00 00 18 00 00 02 c2 02 00 00 01 00 aa 00 91 40 00 00 00
CHG 83 16 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00


SAV 84 16 00 c6 88 04 00 00 00 00 00 00 00 00 00 00 00 00 00 00 3a a7 00 00
DEF 84 16 00 c6 88 04 00 00 00 00 00 00 00 00 00 00 00 00 00 00 3a a7 00 00
CHG 84 16 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00


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


SAV 88 12 14 00 ff ff 00 00 ff ff ff ff 00 20 00 00 00 00 00 00
DEF 88 12 14 00 ff ff 00 00 ff ff ff ff 00 20 00 00 00 00 00 00
CHG 88 12 b5 00 00 00 ff ff ff ff 00 00 a0 ff 00 00 00 00 00 00


SAV 8a 0a 02 00 00 00 00 00 00 00 04 7e
DEF 8a 0a 02 00 00 00 00 00 00 00 04 7e
CHG 8a 0a 03 f1 00 00 00 00 00 00 00 00


SAV 99 06 00 00 00 00 00 00
DEF 99 06 00 00 00 00 00 00
CHG 99 06 00 ff 00 00 00 00


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


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


SAV 80 06 00 00 0f 00 00 00
DEF 80 06 00 00 0f 00 00 00
CHG 80 06 b7 40 0f 00 00 00


READ CAPACITY DATA
08 8b b9 97 00 00 02 00
```

Cheetah 15K.4 FC Product Manual, Rev. C   61

Table 18:   Mode Sense data default and changeable values for ST336754FC drives

```
MODE SENSE (10 BYTE) HEADER DATA
00 ae 00 10 00 00 00 08 04 45 dc cc 00 00 02 00


MODE SENSE PAGES DATA
SAV 81 0a c0 0b ff 00 00 00 05 00 ff ff
DEF 81 0a c0 0b ff 00 00 00 05 00 ff ff
CHG 81 0a ff ff 00 00 00 00 ff 00 ff ff


SAV 82 0e 80 80 00 00 00 00 00 00 01 3c 00 00 00 00
DEF 82 0e 80 80 00 00 00 00 00 00 01 3c 00 00 00 00
CHG 82 0e ff ff 00 00 00 00 00 00 ff ff 00 00 00 00


SAV 83 16 15 88 00 00 00 0c 00 00 02 c2 02 00 00 01 00 f0 00 91 40 00 00 00
DEF 83 16 15 88 00 00 00 0c 00 00 02 c2 02 00 00 01 00 f0 00 91 40 00 00 00
CHG 83 16 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00


SAV 84 16 00 c6 88 02 00 00 00 00 00 00 00 00 00 00 00 00 00 00 3a a7 00 00
DEF 84 16 00 c6 88 02 00 00 00 00 00 00 00 00 00 00 00 00 00 00 3a a7 00 00
CHG 84 16 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00


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


SAV 88 12 14 00 ff ff 00 00 ff ff ff ff 00 20 00 00 00 00 00 00
DEF 88 12 14 00 ff ff 00 00 ff ff ff ff 00 20 00 00 00 00 00 00
CHG 88 12 b5 00 00 00 ff ff ff ff 00 00 a0 ff 00 00 00 00 00 00


SAV 8a 0a 02 00 00 00 00 00 00 00 02 5b
DEF 8a 0a 02 00 00 00 00 00 00 00 02 5b
CHG 8a 0a 03 f1 00 00 00 00 00 00 00 00


SAV 99 06 00 00 00 00 00 00
DEF 99 06 00 00 00 00 00 00
CHG 99 06 00 ff 00 00 00 00


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


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


SAV 80 06 00 00 0f 00 00 00
DEF 80 06 00 00 0f 00 00 00
CHG 80 06 b7 40 0f 00 00 00


READ CAPACITY DATA
04 45 dc cb 00 00 02 00
```

9.4   Miscellaneous operating features and conditions

Table 19 lists various features and conditions. A "Y" in the support column indicates the feature or condition is supported. An "N" in the support column indicates the feature or condition is not supported.

Table 19:   Miscellaneous features

Supported	Feature or condition
Y	FC-AL selective reset
Y	Automatic contingent allegiance
N	Asynchronous event notification
N	Synchronized (locked) spindle operation
Y	Segmented caching
N	Zero latency read
Y	Queue tagging (up to 128 queue tags supported)
Y	Deferred error handling
Y	Parameter rounding (controlled by Round bit in Mode Select page 0)
Y	Reporting actual retry count in Extended Sense bytes 15, 16, and 17
N	Adaptive caching
Y	SMP = 1 in Mode Select command needed to save RPL and rotational offset bytes

Table 20:   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
Y	ACA active
Y	ACA active, faulted initiator

9.5 FC-AL physical interface

Figure 17 shows the location of the J1 Fibre Channel single connection attachment (FC-SCA) and the J6 connector. Figure 19 provides the dimensions of the FC-SCA connector. Figure 20 provides the dimensions of the J6 connector.

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

Figure 17.   Physical interface

9.5.1   Physical characteristics

This section defines physical interface connector.

9.5.1.1   Physical description

FIbre Channel drives may be connected in a loop together or with other compatible FC-AL devices. A maximum of 127 devices may have addresses; however, one of the addresses is reserved for a fabric port switch device. This means 126 addresses are available for FC-AL devices. More FC-AL compatible devices may physically reside on the loop, but they will not be functional because they would not be able to obtain valid addresses.

Port bypass circuits (PBCs) allow devices to be inserted into unpopulated locations or removed from the loop with loop operation recovery after a brief interruption. These PBCs are located external to the FC-AL device. Figure 18 shows the relationship between the PBC and FC-AL device.
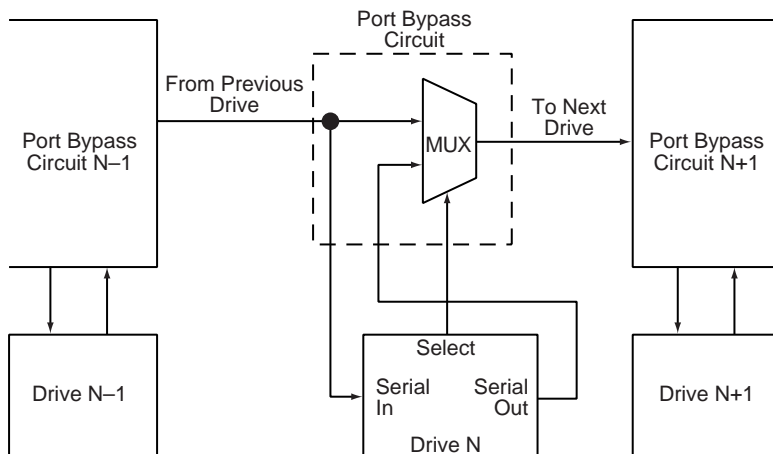
Figure 18.   Port bypass circuit physical interconnect

9.5.2   Connector requirements

Table 21:   Recommended mating SCA part numbers

Part description	Positions	Part number	Features
AMP Vertical (SCA sequence)	40	787317-1	With polarization
Berg	40	71781	With polarization
Methode	40	512-220-91-101N	With polarization
Molex	40	717431040	With polarization

The FC-AL SCA device connector is illustrated in Figure 19. The J6 connector is illustrated in Figure 20.

Figure 19.   FC-AL SCA device connector dimensions

Figure 20.   J6 connector dimensions

9.5.3 Electrical description

Fibre Channel drives use the FC-SCA connector for:

• DC power
• FC-AL interface
• Drive select (device identification)
• Option selection
• Enclosure Services interface

This 40-pin connector is designed to plug directly into a backpanel. External cables are not required.

9.5.4 Pin descriptions

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

Table 22: FC-SCA pin descriptions

Pin	Signal name	Signal type	Pin	Signal name	Signal type
1*	-EN bypass port A	Low Voltage TTL output	21	12 Volts charge	
2*	12 Volts		22	Ground	
3*	12 Volts		23	Ground	
4*	12 Volts		24*	+Port A_in	Diff. PECL input pair
5*	-Parallel ESI		25*	-Port A_in	
6*	Ground[1]		26	Ground	
7*	Active LED out	Open collector out	27*	+Port B_in	Diff. PECL input pair
8*	Reserved		28*	-Port B_in	
9*	Start_1[2]	TTL input	29	Ground	
10*	Start_2[2]	TTL input	30*	+Port A_out	Diff PECL output pair
11*	-EN bypass port B	Low Voltage TTL output	31*	-Port A_out	
12*	SEL_6	TTL input/output	32	Ground	
13*	SEL_5	TTL input/output	33*	+Port B_out	Diff PECL output pair
14*	SEL_4	TTL input	34*	-Port B_out	
15*	SEL_3	TTL input/output	35	Ground	
16*	Fault LED out	Open collector out	36	SEL_2	TTL input/output
17*	DEV_CTRL_CODE_2[2]	TTL input	37	SEL_1	TTL input/output
18*	DEV_CTRL_CODE_1[2]	TTL input	38	SEL_0	TTL input/output
19*	5 Volts		39	DEV_CTRL_CODE_0[2	TTL input
20*	5 Volts		40	5 Volts charge	

*Short pins in mating backpanel connector.

[1]   This pin may be connected to external logic to detect the presence of the drive. The drive connects this pin to the common ground.

[2]   Pins 9, 10, 17, 18, and 39 are option select pins and are tied high by the drive circuitry. The preferred electrical connection at the backplane is either open or grounded (open for the '1' setting, grounded for the '0' setting). Alternatively, these pins may be driven by a 3.3V logic device, pulled up to 3.3V through a pull-up resistor (recommended size of 10k ohm), or grounded through some other means.

9.5.5 FC-AL transmitters and receivers

A typical FC-AL differential copper transmitter and receiver pair is shown in Figure 21. The receiver is required to provide the AC coupling to eliminate ground shift noise.

Figure 21.   FC-AL transmitters and receivers

9.5.6 Power

Power is supplied through the FC-SCA with support for +5 volts and +12 volts. All of the voltage pins in the drive connector are the same length.

Four 12 volt pins provide +12 volt power to the drive. 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. The maximum current typically occurs while the drive motor is starting.

Three 5 volt pins provide logic power to the drive. The current return for the +5 volt power supply is through the common ground pins. Distribute supply and return current as evenly as possible among the voltage and ground pins.

The mating connector pins use shorter contacts to achieve power surge reductions and to aid in "hot plugging" the drives. There are longer voltage contacts in the connector to enable the drive filter capacitors to charge. Current to the drive through the long charge pins is limited by the system in which the drive operates. Three of the +12 volt pins are shorter to allow capacitive pre-charging through the longer +12 volt charge pin. Two of the +5 volt pins are shorter to allow capacitive precharging through the longer +5 volt charge pin.

9.5.7 Fault LED Out

The Fault LED Out signal is driven by the drive when:
• the drive detects failure of both ports
• the drive detects an internal failure
• the drive receives the appropriate fault LED command from the host

The Fault LED Out signal is designed to pull down the cathode of an LED. The anode is attached to the proper +5 volt supply through an appropriate current-limiting resistor. The LED and the current-limiting resistor are external to the drive.

9.5.8 Active LED Out

The Active LED Out signal is driven by the drive as indicated in Table 23.

Table 23:   Active LED Out conditions

Normal command activity	LED status
Spun down and no activity	Slow blink (20% on and 80% off a 2 sec cycle)
Spun down and activity (command executing)	On
Spun up and no activity	On

Table 23:   Active LED Out conditions

Normal command activity	LED status
Spun up and activity (command executing)	Off
Spinning up or down	Blinks steadily (50% on and 50% off)
Format in progress, each cylinder change	Toggles on/off

The Active LED Out signal is designed to pull down the cathode of an LED. The anode is attached to the proper +5 volt supply through an appropriate current limiting resistor. The LED and the current limiting resistor are external to the drive.

9.5.9   Enable port bypass signals

The – Enable Bypass Port A (– EN BYP Port A) and – Enable Bypass Port B (– EN BYP Port B) signals control the port bypass circuits (PBC) located external to the disc drive. The PBC allows a loop to remain functional in the event of a drive failure or removal. When these signals are active, low, the PBC bypasses the drive on the associated port. When an Enable Bypass signal is active, the corresponding Port Bypass LED signal in connector J1 is driven low by the disc drive. A pull down resistor, 1K, located with the PBC should be used to insure the bypass is enabled if the disc drive is not installed.

The Enable Bypass signal is active under failing conditions within the drive, on detection of the Loop Port Bypass primitive sequence, or on removal of the drive. In the bypass state the drive continues to receive on the inbound fibre. Enable Bypass may be deactivated by detection of a Loop Port Enable primitive sequence if the drive has completed self-test and a hardware failure is not present.

Failure modes detected by the disc drive that will enable bypass include:

• Transmitter/receiver wrap test failure
• Loss of receive clock
• Loss of transmission clock
• Drive interface hardware error

9.5.10   Motor start controls

The drive's motor is started according to the Start_1 and Start_2 signals described in Table 24. The state of these signals can be wired into the backplane socket or driven by logic on the backplane.

Table 24:   Motor start control signals

Case	Start_2	Start_1	Motor spin function
1	Low	Low	Motor spins up at DC power on.
2	High	Low	Motor spins up only when SCSI Start command is received.
3	Low	High	Motor spins up after a delay of 12 seconds times the modulo 8 value of the numeric SEL ID of the drive from DC power on.
4	High	High	The drive will not spin up.

9.5.11   SEL_6 through SEL_0 ID lines

The SEL_6 through SEL_0 ID lines determine drive address, and, optionally, for an Enclosure Services Interface. When the Parallel ESI line is high, the enclosure backpanel must provide address information on the SEL line. Refer to table 25 for a mapping of SEL to FC-AL physical addresses (AL_PA). You can think of the SEL

lines as the equivalent of a backpanel logic plug. The drives does not provide pull up resistors on these lines. The backpanel is required to provide high and low inputs to the SEL_ID lines per the specifications in table 27 on page 71.

Note. Table 25 gives AL_PA values for each SEL value. The first entry in the table is SEL_ID 00. The last entry is SEL_ID 7D. SEL_ID 7E is AL_PA 00 which is not valid for an NL_Port, so is not included in the table. Also, SEL_ID 7Fh does map to a valid AL_PA; however, this value signals the drive that physical addresses are not being assigned using the SEL lines and that a "soft" address will be determined by FC-AL loop initialization.

When the Parallel ESI line is low, the enclosure backpanel logic switches to ESI mode if supported. There are two modes of ESI, seven bits of enclosure status and a bidirectional mode. ESI support and the mode are determined by the drive using a discovery process. Refer to the *Fibre Channel Interface Manual* for a description of ESI operation.

9.5.11.1   Parallel Enclosure Services Interface (ESI)

The parallel ESI line is an output from the drive. This line provides the enclosure with an indication of the present function of the SEL lines. A high level, the default state, indicates the drive requires address information on the SEL lines. A low level indicates the drive is attempting an ESI transfer. The enclosure may not support ESI on any or all drive locations. It may only support the address function. Support of ESI is discovered by the drive. Refer to the *Fibre Channel Interface Manual* for a description of ESI operations.

Table 25:   Arbitrated loop physical address (AL_PA) values

AL_PA (hex)	SEL ID (hex)	Setting (dec)	AL_PA (hex)	SEL ID (hex)	Setting (dec)	AL_PA (hex)	SEL ID (hex)	Setting (dec)
EF	00	00	A3	2B	43	4D	56	86
E8	01	01	9F	2C	44	4C	57	87
E4	02	02	9E	2D	45	4B	58	88
E2	03	03	9D	2E	46	4A	59	89
E1	04	04	9B	2F	47	49	5A	90
E0	05	05	98	30	48	47	5B	91
DC	06	06	97	31	49	46	5C	92
DA	07	07	90	32	50	45	5D	93
D9	08	08	8F	33	51	43	5E	94
D6	09	09	88	34	52	3c	5F	95
D5	0A	10	84	35	53	3A	60	96
D4	0B	11	82	36	54	39	61	97
D3	0C	12	81	37	55	36	62	98
D2	0D	13	80	38	56	35	63	99
D1	0E	14	7C	39	57	34	64	100
CE	0F	15	7A	3A	58	33	65	101
CD	10	16	79	3B	59	32	66	102
CC	11	17	76	3C	60	31	67	103
CB	12	18	75	3D	61	2E	68	104
CA	13	19	74	3E	62	2D	69	105
C9	14	20	73	3F	63	2C	6A	106
C7	15	21	72	40	64	2B	6B	107
C6	16	22	71	41	65	2A	6C	108
C5	17	23	6E	42	66	29	6D	109
C3	18	24	6D	43	67	27	6E	110
BC	19	25	6C	44	68	26	6F	111
BA	1A	26	6B	45	69	25	70	112
B9	1B	27	6A	46	70	23	71	113
B6	1C	28	69	47	71	1F	72	114
B5	1D	29	67	48	72	1E	73	115
B4	1E	30	66	49	73	1D	74	116
B3	1F	31	65	4A	74	1B	75	117
B2	20	32	63	4B	75	18	76	118
B1	21	33	5C	4C	76	17	77	119
AE	22	34	5A	4D	77	10	78	120
AD	23	35	59	4E	78	0F	79	121
AC	24	36	56	4F	79	08	7A	122
AB	25	37	55	50	80	04	7B	123
AA	26	38	54	51	81	02	7C	124
A9	27	39	53	52	82	01	7D	125
A7	28	40	52	53	83			
A6	29	41	51	54	84			
A5	2A	42	4E	55	85			

9.5.12 Device control codes

The drive inputs a Device Control Code on the DEV_CTRL_CODE lines at power up to determine the link rate on the Fibre Channel ports. Both ports run at the same rate. If the backpanel does not connect to these lines, the drive has 10K ohm pull up resistors that default the device control code to 7 (1.0625 GHz). Table lists the supported codes.

Table 26: Device control code values

2 (pin 17)	**1** (pin 18)	**0** (pin 39)	**Definition**
0	0	0	Reserved for power failure warning.
0	0	1	Reserved for auto negotiation of link rate.
0	1	0	Reserved.
0	1	1	Reserved.
1	0	0	Reserved.
1	0	1	Reserved.
1	1	0	2.125 GHz operation on both ports.
1	1	1	1.0625 GHz operation on both ports.

9.6 Signal characteristics

This section describes the electrical signal characteristics of the drive's input and output signals. See Table 22 on page 66 for signal type and signal name information.

9.6.1 TTL input characteristics

Table 27 provides the TTL characteristics.

Table 27: TTL characteristics

State	**Voltage**	**Current**
Input high	$1.9 < V_{IH} < 5.5V$	$I_{IH} = \pm500nA$ max.
Input low	$-0.5V < V_{IL} < 0.9V$	$I_{OL} = \pm500nA$ max.
Output high (-EN Bypass A, B)	$2.4 < V_{OH} < 5.25V$	$I_{OH} < -3mA$
Output low (-EN Bypass A, B)	$V_{OL} < 0.5V$	$I_{OL} < 3mA$
Output high (-Parallel ESI)	$2.4 < V_{OH} < 0.9\ V_{CC}$ $V_{OH} > 0.9V_{CC}$	$I_{OH} < -2.4mA$ $I_{OH} < -500\mu A$
Output low (-Parallel ESI)	$0 < V_{OL} < .45V$	$I_{OL} < 2.4mA$
Output high (all other outputs)	$2.4 < V_{OH} < 0.9\ V_{CC}$ $V_{OH} > 0.9V_{CC}$	$I_{OH} < -1.6mA$ $I_{OH} < -500\mu A$
Output low (all other outputs)	$0 < V_{OL} < .45V$	$I_{OL} < 1.6mA$

9.6.2 LED driver signals

Fault and Active LED signals are located in the FC-SCA connector (J1) and through the indicator connector (J6). See Table 28 for the output characteristics of the LED drive signals.

Table 28: LED drive signal

State	Current drive available	Output voltage
LED off, high	$0 < I_{OH} < 100\mu A$	
LED on, low	$I_{OL} < -30$ mA	$0 < V_{OL} < 0.8V$

9.6.3 Differential PECL output

The serial PECL output signal voltage characteristics are provided in Table 29. The outputs are not AC coupled in order to deliver maximum signal without rise and fall time degradation. You must AC couple the receiver to isolate potentially different DC characteristics of the outputs and the receiver.

Table 29: Differential PECL output characteristics

Description	Parameter	Notes
Serial output voltage swing	$600 < V_{out} < 1300$ mV	Centered at 1.32V

Figure 22 provides the data output valid eye diagram relative to the bit cell time.

Figure 22.   Transmit eye diagram

9.6.4 Differential PECL input

The serial PECL input signal voltage characteristics are provided in Table 30.

Table 30: Differential PECL input characteristics

Description	Parameter	Notes
Serial input voltage swing	$200 < V_{in} < 1.300$ mV	AC coupled

Figure 23 provides the data valid eye diagram for typical and minimum requirements to recover data at the specified interface error rate. The inputs are AC coupled on the drive.

Figure 23.   **Receive eye diagram**

10.0   Seagate Technology support services

Online services

Internet
www.seagate.com for information about Seagate products and services. Worldwide support is available 24 hours daily by e-mail for your questions.

Presales Support: www.seagate.com/support/email/email_presales.html *or* Presales@Seagate.com
Technical Support: www.seagate.com/support/email/email_disc_support.html *or* DiscSupport@Seagate.com

mySeagate
my.seagate.com is the industry's first Web portal designed specifically for OEMs and distributors. It provides self-service access to critical applications, personalized content and the tools that allow our partners to manage their Seagate account functions. Submit pricing requests, orders and returns through a single, password-protected Web interface—anytime, anywhere in the world.

reseller.seagate.com
reseller.seagate.com supports Seagate resellers with product information, program benefits and sales tools. You may register for customized communications that are not available on the web. These communications contain product launch, EOL, pricing, promotions and other channel-related information. To learn more about the benefits or to register, go to reseller.seagate.com, any time, from anywhere in the world.

Automated phone services

SeaFONE® (1-800-SEAGATE) is the Seagate toll-free number (1-800-732-4283) to access our automated directory assistance for Seagate Service Center support options. International callers can reach this service by dialing +1-405-324-4770.

Seagate Service Centers

Presales Support
Our Presales Support staff can help you determine which Seagate products are best suited for your specific application or computer system, as well as drive availability and compatibility.

Technical Support
If you need help installing your drive, consult your system's documentation or contact the dealer's support services department for assistance specific to your system. Seagate technical support is also available to assist you online at support.seagate.com or through one of our call centers. Have your system configuration information and your drive's "ST" model number available.

SeaTDD™ (+1-405-324-3655) is a telecommunications device for the deaf (TDD). You can send questions or comments 24 hours daily and exchange messages with a technical support specialist during normal business hours for the call center in your region.

Customer Service Operations

Warranty Service
Seagate offers worldwide customer support for Seagate drives. Seagate distributors, OEMs and other direct customers should contact their Seagate Customer Service Operations (CSO) representative for warranty-related issues. Resellers or end users of drive products should contact their place of purchase or one of the Seagate CSO warranty centers for assistance. Have your drive's "ST" model number and serial number available.

Data Recovery Services
Seagate offers data recovery services for all formats and all brands of storage media. Our Data Recovery Services labs are currently located in North America. To get a free quick quote or speak with a case management representative, call 1-800-475-0143. Additional information, including an online request form and data loss prevention resources, is available at www.datarecovery.seagate.com.

Authorized Service Centers
In some locations outside the US, you can contact an Authorized Service Center for service.

USA/Canada/Latin America support services

Seagate Service Centers

Presales Support

Call center	Toll-free	Direct dial	FAX
Americas	1-877-271-3285[1]	+1-405-324-4730[1]	+1-405-324-4704

Technical Support

Call center	Toll-free	Direct dial	FAX
Americas	1-800-SEAGATE[2]	+1-405-324-4700[2]	+1-405-324-3339

Customer Service Operations

Warranty Service

Call center	Toll-free	Direct dial	FAX / E-mail
USA, Canada, Mexico and Latin America	1-800-468-3472[3]	—	+1-956-664-4725
Brazil			
Jabil Industrial Do Brasil LTDA[4]	—	+55-11-4191-4761	+55-11-4191-5084 SeagateRMA.br@jabil.com

Data Recovery Services

Call center	Toll-free	Direct dial	FAX
USA, Canada, and Mexico	1-800-475-01435	+1-905-474-2162	1-800-475-0158 +1-905-474-2459

---

[1]Hours of operation are 8:00 A.M. to 11:45 A.M. and 1:00 P.M. to 6:00 P.M., Monday through Friday (Central time).
[2]Hours of operation are 8:00 A.M. to 8:00 P.M., Monday through Friday (Central time).
[3]Hours of operation are 8:00 A.M. to 5:00 P.M., Monday through Friday (Central time).
[4]Authorized Service Center
[5]Hours of operation are 8:00 A.M. to 8:00 P.M., Monday through Friday, and 9:00 A.M. to 5:00 P.M., Saturday (Eastern time).

European support services

For presales and technical support in Europe, dial the Seagate Service Center toll-free number for your specific location. If your location is not listed here, dial our presales and technical support call center at +1-405-324-4714 from 8:00 A.M. to 11:45 A.M. and 1:00 P.M. to 5:00 P.M. (Central Europe time) Monday through Friday. The presales and technical support call center is located in Oklahoma City, USA.

For European warranty service, dial the toll-free number for your specific location. If your location is not listed here, dial our European CSO warranty center at +31-20-316-7222 from 8:30 A.M. to 5:00 P.M. (Central Europe time) Monday through Friday. The CSO warranty center is located in Amsterdam, The Netherlands.

Seagate Service Centers

Toll-free support numbers

Call center	Presales and Technical Support	Warranty Service
Austria	—	00 800-47324289
Belgium	00 800-47324283 (00 800-4SEAGATE)	00 800-47324289
Denmark	00 800-47324283	00 800-47324289
France	00 800-47324283	00 800-47324289
Germany	00 800-47324283	00 800-47324289
Ireland	00 800-47324283	00 800-47324289
Italy	00 800-47324283	00 800-47324289
Netherlands	00 800-47324283	00 800-47324289
Norway	00 800-47324283	00 800-47324289
Poland	00 800-311 12 38	00 800-311 12 38
Spain	00 800-47324283	00 800-47324289
Sweden	00 800-47324283	00 800-47324289
Switzerland	00 800-47324283	00 800-47324289
Turkey	00 800-31 92 91 40	00 800-31 92 91 40
United Kingdom	00 800-47324283	00 800-47324289

FAX services—All Europe (toll call)

Technical Support	+1-405-324-3339
Warranty Service	+31-20-653-3513

Africa/Middle East support services

For presales and technical support in Africa and the Middle East, dial our presales and technical support call center at +1-405-324-4714 from 8:00 A.M. to 11:45 A.M. and 1:00 P.M. to 5:00 P.M. (Central Europe time) Monday through Friday. The presales and technical support call center is located in Oklahoma City, USA.

For warranty service in Africa and the Middle East, dial our European CSO warranty center at +31-20-316-7222 from 8:30 A.M. to 5:00 P.M. (Central Europe time) Monday through Friday, or send a FAX to +31-20-653-3513. The CSO warranty center is located in Amsterdam, The Netherlands.

Asia/Pacific support services

For Asia/Pacific presales and technical support, dial the toll-free number for your specific location. The Asia/Pacific toll-free numbers are available from 6:00 A.M. to 10:45 A.M. and 12:00 P.M. to 6:00 P.M. (Australian Eastern time) Monday through Friday, except as noted. If your location is not listed here, direct dial one of our technical support locations.

Warranty service is available from 9:00 A.M. to 6:00 P.M. April through October, and 10:00 A.M. to 7:00 P.M. November through March (Australian Eastern time) Monday through Friday.

Seagate Service Centers

Call center	Toll-free	Direct dial	FAX
Australia	1800-14-7201	—	—
China (Mandarin)[1, 4]	800-810-9668	+86-10-6225-5336	—
Hong Kong	800-90-0474	—	—
Hong Kong (Cantonese)[1, 4]	001-800-0830-1730	—	—
India[2, 4]	1-600-180-1104	—	—
Indonesia	001-803-1-003-2165	—	—
Japan[3, 4]	0034 800 400 554	—	—
Korea[3, 4]	007 98 8521 7635	—	—
Malaysia	1-800-80-2335	—	—
New Zealand	0800-443988	—	—
Singapore	800-1101-150	—	+65-6488-7525
Taiwan (Mandarin)[1, 4]	00-800-0830-1730	—	—
Thailand	001-800-11-0032165	—	—

Customer Service Operations

Warranty Service

Call center	Toll-free	Direct dial	FAX
Asia/Pacific	—	+65-6485-3595	+65-6485-4860
Australia	1800-12-9277	—	—
India[4]	—	+91-44-821-6164	+91-44-827-2461

---

[1]Hours of operation are 8:30 A.M. to 5:30 P.M., Monday through Friday (Australian Western time).
[2]Hours of operation are 9:00 A.M. to 6:00 P.M., Monday through Saturday.
[3]Hours of operation are 9:30 A.M. to 6:30 P.M., Monday through Friday.
[4]Authorized Service Center

Publication feedback survey

We are interested in your comments and suggestions regarding this publication. Please take a few minutes to participate in our survey at the following URL:

http://survey.seagate.com/survey/techpubs.nsf

Thank you for your time and comments.

Index

Numerics

A

B

C