



Enterprise-class versus Desktop-class Hard Drives

Revision 1.0

April, 2008

Enterprise Platforms and Services Division - Marketing

Revision History

Date	Revision Number	Modifications
April, 2008	1.0	Initial Release

Disclaimers

Information in this document is provided in connection with Intel® products. No license, express or implied, by estoppel or otherwise, to any intellectual property rights is granted by this document. Except as provided in Intel's Terms and Conditions of Sale for such products, Intel assumes no liability whatsoever, and Intel disclaims any express or implied warranty, relating to sale and/or use of Intel products including liability or warranties relating to fitness for a particular purpose, merchantability, or infringement of any patent, copyright or other intellectual property right. Intel products are not intended for use in medical, life saving, or life sustaining applications. Intel may make changes to specifications and product descriptions at any time, without notice.

Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined." Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them.

Intel, Pentium, Celeron, and Xeon are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Copyright © Intel Corporation 2008.

*Other names and brands may be claimed as the property of others.

Introduction

In today's economic landscape server system integrators are challenged to balance data and system reliability with cost. It is important to balance the type, capacity, and number of hard drives in a system with the requirements of data integrity and availability.

This paper provides a technical explanation of the differences between enterprise-class and desktop-class hard drives to provide system integrators with the information they need to determine their enterprise hard drive requirements.

Enterprise versus Desktop Systems

Desktop systems are generally focused on running client applications with disk I/O operations supporting operating system load and runtime requirements as well as program/application load and runtime requirements. In many of these implementations client data is not stored locally, but is accessed across a network where data is stored on a file or storage server.

Enterprise-class server systems not only rely on operating system and application load and runtime support, but also may provide application and/or storage services to the network, which requires large capacity data storage that is always available and reliable.

The usage models for desktop systems and enterprise systems results in very different requirements for hard drives, which in turn influences the robustness of hard drive design, the implementation of features to enhance reliability, and the related costs. The table below provides an outline of the requirement differences for desktop and enterprise class systems. Each of these requirements is described in the following section.

Requirement	Desktop	Enterprise
Operational Availability	8 hours/day - 5 days/week	24 hours/day - 7 days/week
Work Load	10 - 20%	100%
Cost Sensitivity	Very sensitive to low cost	Moderately sensitive, balanced with requirements for reliability, availability, and data integrity.
Performance	Low to Moderate	High
Reliability	Moderate: <ol style="list-style-type: none"> 1. Outage affects only one user <ol style="list-style-type: none"> a. Critical data is not usually stored locally b. Higher Tolerance for long error recovery timeout 2. Lower Mean Time Between Failure Acceptable	High: <ol style="list-style-type: none"> 1. Outage affects multiple users 2. Higher Mean Time Between Failure 3. Intolerance for long Error Timeout
Data Integrity	Moderately desirable (a bit corruption may result in system lockup or critical data loss on one system)	Highly Desirable (a bit corruption may result in catastrophic critical data loss to multiple clients)

Enterprise versus Desktop Hard Drives

Hard Drive Features

To meet operational requirements, some vendors and drive models have significant differences in disk drive design for desktop class hard drives compared to enterprise class hard drives. This document examines these differences and describes how the drive design and feature set meet the needs of the two usage models. This document will contrast drive class differences in general terms, and will give a general description of features and methods used by drive vendors to meet these differing needs. This document will not specify vendor specific features or design implementations. Actual drive models may contain a mixture of the features that are described. For example, a particular desktop class drive model may include some enterprise class drive features, and some enterprise class drive models may not include all enterprise class features.

1.1.1 Operational Availability and Work Load

Desktop class systems normally have one hard disk drive, or a mirrored set of drives that are designed to work in a less demanding environment. The drive is only accessed to provide program or application access, access the swap file, or complete a limited data save or data retrieval for the operating system or running application. The system is usually shut down during the evening and weekend hours, or left idle for long periods of time.

An enterprise system may control multiple drives that add both capacity and redundancy to a storage subsystem. For example, an enterprise system may use four drives in a RAID 10 for the operating system and may use many drives for configuration in multiple RAID 5 or RAID 6 arrays to hold client data. The enterprise system will not only perform operating system and application tasks locally, but will also be expected to support client requests 100% of the time. During off peak time the enterprise system may be involved in patrolling the hard drives for defects or errors, system backup, and other maintenance tasks. Enterprise workloads create greater wear on bearings, motors, actuators, and platter media; which generates additional heat and additional vibration. Enterprise class drives are designed with heavy duty components and drive firmware programming to meet the rigors of this environment.

1.1.2 Cost Sensitivity

Desktop systems are very price sensitive, driven in part because the ratio of clients to servers is much higher. Clients run user-oriented applications that place a lighter and intermittent load on the local drive(s). Enterprise systems are not as price sensitive because they are typically running multiple enterprise applications and must handle a higher work load with increased reliability and availability. These enterprise-class systems usually include a higher drive count. The features required to support the additional work load add cost to the drive system design.

1.1.3 Performance

Enterprise-class drives generally incorporate internal mechanisms that allow faster data access and retrieval. These features include heavier actuator magnets, faster spindle speeds, denser magnetic media, and faster drive electronic components with more cache memory and faster hard drive micro-processor speeds.

1.1.4 Reliability

Reliability is influenced by the factors described below.

1.1.4.1 Bad Sector Recovery

Desktop drives perform heroic efforts to recover data in a bad sector. Usually there is one hard drive in a desktop system and a bad sector can result in a catastrophic failure of the operating system or an application failure. Typically desktop systems do not provide an online backup copy of the sector. To recover data contained in a bad sector, desktop class drives are designed to attempt to re-read the sector many times before responding with an “un-recoverable read error”. During this recovery effort, the drive may become unresponsive and may ignore bus resets. When the drive disappears from the bus for an extended period of time, the operating system, the application, and the user are required to wait for the system to respond. A typical desktop drive command timeout can take many minutes and no disk access is allowed while the system attempts to retry the command.

Long drive recovery timeouts are not acceptable in an enterprise environment because multiple users can be affected, and because RAID systems, which are prevalent in enterprise systems, do not tolerate an unresponsive drive. A feature of enterprise-class hard drives is a short command timeout value. When a drive has a difficult time reading a sector and the short timeout is exceeded, the drive will respond by attempting to recover missing data from sector checksum if available. If that attempt fails, the drive will notify the controller and the controller will attempt to recover the sector from redundant data on adjacent disks and remap bad sectors related to the error. The short timeout allows these recovery efforts to take place while the system drives continue to support system disk access requests by the operating system. Typical timeout for an enterprise class drive is 7 to 15 seconds and retries are limited to a few attempts.

Desktop-class drives with timeout values exceeding 30 seconds should not be used in an enterprise-class system. They could cause drives to be marked as offline more often, and may increase operating system crashes, kernel panics, or blue screens.

1.1.4.2 Rotational Vibration

Rotational vibration is measured in Radians per second per second, which refers to vibration in the same plane that the drive spins. Vibration in other planes is frequently lumped in with this term though. Drive manufacturer specifications for vibration are usually limited to measurements within the rotational plane and high vibration frequency ranges may not be addressed by manufacturer specifications.

The source of vibration that can affect drive operation can come from moving components within a system, such as fans and neighboring hard drives. Vibration from fans can be transmitted through the system chassis to hard drives, and vibration from a drive can be transferred to another drive, or from the drive to the system and then reflected back to the originating drive. While data is read from or written to a disk, these vibrations can push the read/write head out of alignment with the data track. Failure to compensate for vibration induced misalignment can result in data that is written or read off track. This increases the possibility that an off-track write could corrupt data in an adjoining track, or that an off-track read could result in incorrect data or data that cannot be located or read.

Enterprise class drives provide a more sophisticated compensation for vibration by sensing the vibration motion of the drive, and by sensing head position and track alignment. The drive can then react with additional actuator strength or wait for the spindle motor to spin the target media location under the head again so that it can re-attempt access. Some forms of misalignment compensation can have a dramatic impact on performance due to the time required to spin the data target location back under the read/write head. Performance loss of up to 90% has been observed in some tests. The amount of performance loss is dependent on the frequency and strength of the vibration.

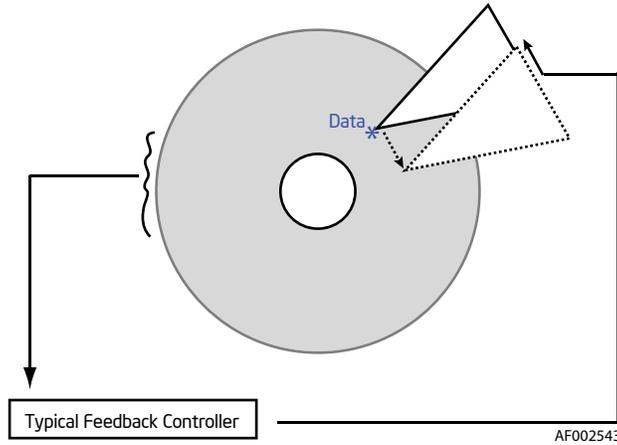
1.1.4.3 Misalignment detection

Most drives incorporate multiple “servo wedges” in the track. Drive firmware constantly monitors these wedges to determine the location of the head in relation to the track. If the drive firmware detects a misalignment, it will hold the read or write and wait for the target location to spin under the head again. Desktop class drives have fewer “servo wedges”, or a single combined servo/data path processor. With a single processor, the drive is unable to servo the head while writing to media. Most desktop drives do not have dedicated servo and data path processors or the firmware compensation algorithms that make them more susceptible to RV errors, which results in poor performance due to vibration. Most enterprise class drives have dedicated servo and data path processors with servo algorithms in the drive firmware to enforce this compensation.

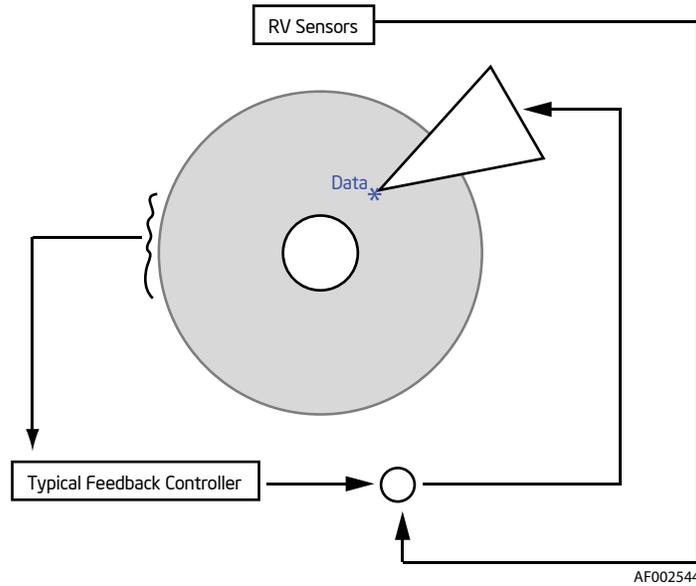
1.1.4.4 Vibration Sensors

Most enterprise drives implement a vibration sensor circuit on the drive’s electronic board. These sensors can detect the movement of the drive when the servo is moving the heads, which provides a reliable method of positioning the heads and determining when it is safe to read or write data. However, this can be less beneficial if there is a single, combined servo/data path processor. Enterprise class drive designs include a closed loop feedback system between the magnetic head and the spindle(s) to sense vibration anomalies and react accordingly.

Figure 1: RV Compensation Feedback Mechanisms



No RV Compensation (Desktop-class drive)



RV Compensation (Enterprise-class drive)

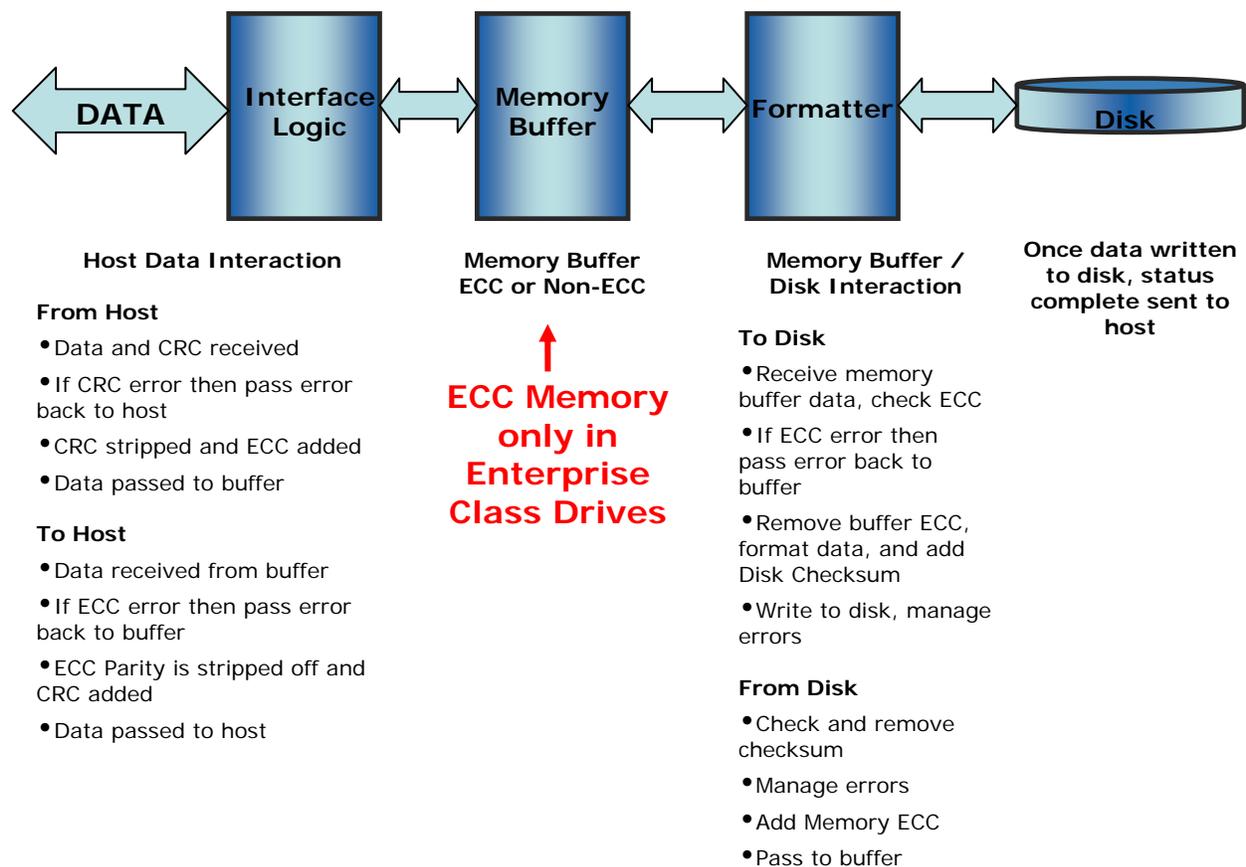
Desktop class drives have less sophisticated mechanisms to compensate for vibration induced errors, which causes greater performance loss and a higher error rate. Enterprise class drives are able to more effectively detect and compensate for vibration.

Symptoms of vibration related errors include lower drive performance, a larger number of logged medium errors, and an increased frequency of drives marked offline by the I/O controller. In addition, data corruption has been observed in high vibration environments with some drive models.

1.1.4.5 Data Integrity

One feature of an enterprise-class system is the implementation of “end-to-end” error detection. Transmitted data is accompanied by parity or checksum at every stage of transmission within the system. This allows data transmission errors to be detected, and in some cases corrected or retransmitted. Although desktop systems do have error detection in some subsystems, they do not usually provide this end-to-end data protection. For example desktop systems do not usually incorporate Error Correction Code (ECC) in system memory or drive memory buffers. Enterprise class systems use error detection at every data transmission stage within the system, including ECC support in system memory and drive memory buffers. Refer to the illustration below for more information.

Figure 2: Data Protection Within the Hard Drive Electronics



Enterprise class drives use ECC for data passing through drive memory and may use additional error detection methods for data transmitted within the drive electronics. The form of this error detection and correction capability is usually proprietary to the drive vendor. Desktop drives do not have this data protection feature within drive memory, If a data error occurs in memory it will not be detected. The error will be transmitted to the next stage in the drive electronics, and the error will be propagated to the drive media or to the host device. This type of error becomes very difficult to detect and can affect operating system stability or client data reliability.

1.1.4.6 Variable Sector Size

Many enterprise class hard drives (especially SCSI and SAS hard drives) have the ability to vary the sector size within limits. These drives utilize a 528 byte sector and allow the I/O controller (usually a RAID controller) to set the data portion of the sector to 512 bytes and to use the remaining sector capacity for a sector checksum. This allows the I/O controller to verify the sector data against the checksum and to recover data on the fly from the checksum. The I/O controller can track this error and remap bad areas of the drive in I/O controller memory. The drives have the ability to track and remap errors via sector parity.

Desktop (IDE and SATA) drives do not utilize a variable sector size. The sector is locked at 512 bytes, although there is a small area reserved for parity at the end of each sector. This parity is sufficient to detect an error, but it does not contain sufficient data to rebuild the sector or recover data if the sector is unreadable. It can only be used as part of the drives internal error detection mechanism.

Some SATA drive vendors are working to implement variable sector size capability in their enterprise class SATA drives. Larger sector sizes are forecasted for future implementation.

1.1.4.7 Feature Review

The following table provides a comparison of feature differences between enterprise-class and desktop-class hard drives. Vendor implementation of these features vary between drive models. Generally the high end of the feature spectrum includes enterprise-class SAS hard drives, and the low end includes desktop-class SATA drives. Enterprise-class SATA drives fall somewhere in between.

Table 1: Feature Reference

Drive Comparison Table

Feature	Enterprise	Desktop
Spindle Motor	Higher RPM	Moderate to lower RPM
	Tighter run-out (spindle end movement)	Lower specification for run-out
	Spindle anchored at both ends	Spindle anchored at one end
Media	Full media cert	Lower media specification and density
Head Stack Assembly	Structural rigidity	Lighter weight design
	Lower inertial design	Higher inertial design
Actuator Mechanics	Larger magnets	Smaller magnets
	Air turbulence controls	No air turbulence compensation
	RV sensors and closed loop RV suppression	No RV sensors or suppression - limited to servo wedge track alignment
Electronics	Dual processors (dedicated servo and data path processors)	Single processor
	Performance optimization	No performance optimization
	Advanced error handling	Standard error handling
	Advanced firmware algorithms	Standard firmware algorithms

Drive Comparison Table

Feature	Enterprise	Desktop
Performance		
Latency and Seek Time	5.7 msec @ 15K rpm	13 msec @ 7200rpm (or slower)
Command Queuing and Reordering	Full	Limited
Rotational Vibration Tolerance	Up to 21 rads/sec/sec	Up to 5 to 12 rads/sec/sec
Typical I/Os per sec/drive (no RV)	319	77
Typical I/Os per sec/drive (20 rad/sec/sec RV)	310	<7
Duplex Operation	Full	Half
Customization		
FW Code	Extensive	Limited
Variable Sector Sizes	Yes (SCSI/SAS only)	No
LEDs	Yes	No
Reliability		
MTBF	1.2M hours at 45 degrees C, 24X7 100% duty cycle.	700K hours at 25 degrees C and 8X5, 10%-20% duty cycle
Internal Data Integrity Checks	End to End	Limited, none in memory buffer
Maximum Operating Temperature	~60 degrees C	~40 degrees C
Warranty	~5 years	~1 to 3 years

Note: Typical MTBF, operating temperature, and operating rotational vibration vary depending on the manufacturer.

Conclusion

As hard drives technology changes, customers have a choice of products to use in their enterprise or desktop environment. It is important to properly balance the system requirements with hard drive features and choose a hard drive to meet usage requirements. Hard drive usage, reliability, and functionality can be optimized by choosing drives that meet the design specifications demanded by their target implementation.

Hard drives are designed with a wide variety of features that can impact data integrity and system uptime. Some hard drive manufacturers may differentiate enterprise from desktop drives by not testing certain enterprise-class features, validate the drives with different test criteria, or disable enterprise-class features on a desktop class hard drives so they can market and price them accordingly. Other manufacturers have different architectures and designs for the two drive classes. It can be difficult to get detailed information and specifications on different drive modes.

Intel does not recommend specific hard disk drive vendors or technology and feature implementations, but we encourage you to contact the hard drive vendor of your choice to discuss the hard drive and features best suited for your deployment application.

Check your applicable Intel® Server product THOL (*Tested Hardware and Operating system List*) for drive compatibility with Intel® Server Products before making a hard drive purchase. Intel® Server Product THOLs can be found at: <http://www.intel.com/support/motherboards/server/>

For more information, please contact your Intel representative.

References

Intel® Server and Intel® RAID products:

- <http://www.intel.com/products/server/chassis/index.htm>
- http://www.intel.com/products/server/storage/index.htm?iid=ncdnav2+stor_sys
- http://www.intel.com/products/server/raid/index.htm?iid=serv_body+raid

Enterprise-class hard drive vendors:

- <http://www.fujitsu.com/global/services/computing/storage/hdd/ehdd/>
- <http://www.hgst.com/portal/site/en/menuitem.fe06f94c22c4099056fb11f0aac4f0a0/>
- <http://www.seagate.com/www/en-us/products/servers/>
- <http://www.wdc.com/en/products/index.asp?cat=2>