

Seagate Constellation/Barracuda* ES.2 SAS Hard Drives and Intel® RAID SAS/SATA Adapters Powered by LSI MegaRAID*

White Paper

Ideal for Low-Cost Servers and Nearline Applications

Revision 1.0

July 2009

Enterprise Platforms and Services Division - Marketing

Revision History

Date	Revision Number	Modifications
July 2009	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.

The Intel® RAID Controllers may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

Except as permitted by license, no part of this document may be reproduced, stored in a retrieval system, or transmitted in any form or by any means without the express written consent of Intel Corporation.

Intel, Pentium, Itanium, and Xeon are trademarks or registered trademarks of Intel Corporation.

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

Copyright © Intel Corporation 2009. Portions Copyright © LSI Corporation 2009.

Table of Contents

- 1. Introduction 1**
 - 1.1 SAS and SATA Overview 1
- 2. Tiers of Storage 2**
 - 2.1 Storage Market Summary 3
- 3. Tier 2 SAS over Enterprise SATA and Desktop SATA 4**
 - 3.1 Tier 2 SAS and Enterprise SATA Differences 4
 - 3.2 Enterprise SATA, Tier 2 Nearline SAS, and Desktop SATA Differences..... 7
 - 3.3 Feature List Summary 8
- 4. Intel® RAID SAS/SATA adapters 9**

List of Figures

Figure 1. Tiers of Storage	2
Figure 2. Barracuda* ES.2 Drives	4
Figure 3. Tier 2 SAS Features for Nearline Applications	6
Figure 4. SAS/SATA Electronics.....	6
Figure 5. Intel® RAID SAS/SATA Adapters.....	9
Figure 6. Load Balancing	10

List of Tables

Table 1. Performance Comparison	6
Table 2. Feature List Summary.....	8

1. Introduction

With SCSI technology dwindling and Serial ATA (SATA) and Serial-Attached SCSI (SAS) technology gaining ground, many customers question which hard drive and RAID adapter are best for their specific application. This paper discusses the different tiers of storage and how Intel® RAID SAS/SATA adapters and Seagate Constellation and Barracuda ES.2 SAS drives are ideal for nearline and low-cost server environments. It compares the features and benefits of these new Tier 2 nearline SAS drives, traditional Enterprise SATA drives, and Desktop SATA drives.

The combination of Intel® RAID SAS/SATA adapters and Seagate Constellation and Barracuda ES.2 SAS drives offers a unique benefit to system builders and channel OEMs. Both product lines offer advanced SAS functionality at virtually the same cost as comparable SATA-only products with interposters. And with support for new features such as multipathing and load balancing, Intel® SAS adapters can take full advantage of dual-port, full-duplex, and concurrent active channels on the Tier 2 nearline SAS hard drives. This is something that no other vendors can offer today. Together, Intel and Seagate offer customers an unrivaled combination of performance, reliability, and affordability for nearline and low-cost server applications.

1.1 SAS and SATA Overview

For nearly two decades, parallel SCSI has dominated the high-performance, high-reliability enterprise storage hard drive market. But, with the need for increased bandwidth and flexibility, parallel SCSI is nearing obsolescence due to its fundamental limitations. In recent years, Serial ATA (SATA) and Serial-Attached SCSI (SAS) adoption rates have skyrocketed because of their ability to overcome the previous limitations of parallel ATA and parallel SCSI.

Users typically select SCSI for one of several reasons: scalability, reliability, and/or performance. With the debut of enterprise SATA drives, which brought high capacities, low cost per gigabyte, improved scalability, and some improved reliability features, only users requiring extreme performance made the transition from SCSI to SAS. Mission critical SAS drives (such as Seagate Savvio* 15k) are preferred in online and transactional applications due to their exceptional performance and reliability. Yet, many system builders and channel customers have found desktop and enterprise SATA drives to be “good enough” for many low-cost Server and Tier 2 Nearline applications.

2. Tiers of Storage

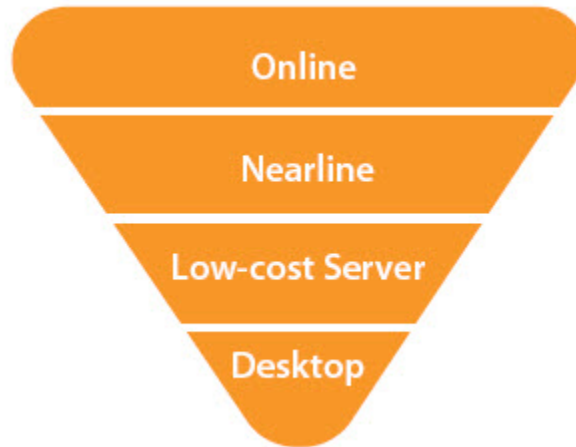


Figure 1. Tiers of Storage

There are several tiers that make up today's storage market: From entry-level (Tier 3) to Enterprise (Tier 1), there is Desktop, Low-Cost Server, Nearline, and Online storage.

Desktop applications represent the least demanding storage environment, typically supporting a single user using email, word processing applications, web browsers, and so on. For these applications, desktop SATA drives are ideal as they are designed to provide appropriate performance and reliability for single user, low-duty cycle environments at minimal cost.

Low-Cost Servers are generally used by departmental workgroups or a relatively small group of users for applications such as file/print serving and small web or e-mail servers. Although Low-Cost Servers are typically powered on 24x7, they are utilized mostly during normal business hours. Consequently, these servers have long durations of low duty cycle demands mixed with short durations of high duty cycle demands, when hard drives are heavily relied on. Equipment cost and support is also a major consideration when purchasing for this segment.

Nearline storage fills the gap between offline (archival) and online (primary) enterprise storage. These systems are designed to provide economical storage and quick access to reference data for organizations trying to improve existing tape backup, archiving, and overall data protection. Nearline applications such as medical imaging, multimedia, and supercomputing require effective management of large amounts of data that is not as critical as online data, but still needs to be accessed quickly and easily. With total costs lower than primary storage and overall performance better than desktop, Nearline storage helps organizations store and manage massive amounts of data more efficiently and reliably.

Online storage, also known as primary storage, is best suited for transactional applications that require constant and immediate access to data. These applications, ranging from highly critical databases to e-mail servers, rely on exceptional performance, reliability and availability. Online storage is powered on 24x7, handles mostly random requests, delivers high I/Os per second (IOPS), and provides high duty cycles. For that reason, mission critical SAS drives such as Seagate Savvio* or Cheetah* drives are the logical choice for online storage.

With the features and capabilities of SATA hard drives encroaching into the SAS hard drive space and vice versa, it is easy to see why the lines have blurred for when and where to use desktop SATA, enterprise SATA, enterprise SAS, and now Tier 2 nearline SAS hard drives. Traditionally, the difference in price/gigabyte between SATA and SAS hard drives has been a barrier to entry for those contemplating using SAS. With desktop drives at \$0.15/GB, enterprise SATA drives at \$0.25/GB, and mission critical SAS drives at nearly \$2.00/GB, many buyers opt for one of the SATA solutions. The dollars per gigabyte of SATA outweighed the performance and reliability advantages of SAS in some environments. That is why SAS has historically been reserved for the Tier 1, mission-critical applications where performance and rock-solid reliability are key and cost is not the primary consideration.

Nevertheless, the disparity between SATA and SAS adoption rates has not gone unnoticed by some vendors. In early 2007, Intel reduced pricing on the SAS/SATA value line adapters to match competitor's SATA-only RAID adapter pricing. This provided customers with the flexibility to create tiered storage with SATA and/or SAS hard drives with no cost premium, while also protecting their investment in storage.

With the introduction of the Seagate Barracuda* ES.2 product line, Seagate is offering the first SAS hard drive with 1 Terabyte (TB) capacity. Barracuda* ES.2 SAS drives use a SATA drive platform (7200 rpm) coupled with enterprise SAS electronics (similar to those of the Cheetah* 15k family). The result is a hard drive with SATA capacities, but with an enterprise-class SAS interface. With only pennies per gigabyte separating their Enterprise SATA and these Tier 2 nearline SAS drives, it is clear that Seagate also recognizes the desire for SAS features at enterprise SATA pricing.

2.1 Storage Market Summary

Online

- Transactional applications
- Constant and immediate access to data
- Critical databases, Large email servers
- Powered on 24x7
- High IOPs and high duty cycles

Nearline

- Economical storage for reference data
- Archival, medical imaging, multimedia, supercomputing
- Stores massive amounts of data
- Access to data must be quick and easy

Low-cost Server

- Departmental workgroups/small group of users
- File/print serving, small web or email servers
- Powered on 24x7, utilized more 8x5
- Mix low duty cycle/high duty cycle demands
- Equipment cost/support major consideration

Desktop

- Least demanding
- Single user
- Email, word processing, web browsing
- Powered on 8x5

3. Tier 2 SAS over Enterprise SATA and Desktop SATA

This section explains the reasons for choosing Tier 2 nearline SAS drives over:

- Enterprise SATA for Nearline Applications
- Desktop SATA for Low-Cost Server Applications

3.1 Tier 2 SAS and Enterprise SATA Differences

Constellation and Barracuda ES.2 drives offer the benefits of SAS technology at Enterprise SATA pricing.



Figure 2. Constellation and Barracuda* ES.2 Drive

With the release of the Seagate Constellation and Barracuda ES.2 SAS drives, Nearline applications can now enjoy enterprise peace of mind without the enterprise cost. Until now, organizations rightly chose to use Enterprise SATA drives due to their combination of affordability and increased reliability over desktop drives. However, with the price of Seagate's Tier 2 nearline SAS drives at pennies more per gigabyte than the comparable Enterprise SATA with interposer solutions offering, users no longer have to settle for "good enough".

Seagate Tier 2 nearline SAS drives offer many of the features and functionality of the Tier 1 SAS drives, such as Seagate's Savvio* and Cheetah*. These features include dual-port, full-duplex, concurrent active channels, robust command queuing, and superior error detection and correction.

The combination of full-duplex, dual-port, concurrent active channels and SAS command queuing, give SAS drives a large performance advantage over Enterprise SATA drives. Full duplex supports transfers in both directions simultaneously, which greatly increases throughput. This inherently provides better performance than Enterprise SATA, which only supports half duplex, or the transfer of data only in one direction at a time.

Dual port provides two redundant paths to every hard drive. This increases availability and reliability by supplying each drive with an alternate path of communication in case one path fails. SATA, which is a single port, only offers one path to every hard drive. Therefore, if this single path of communication fails, there may be data loss and system downtime.

SAS hard drives also provide two active (concurrent) channels, which allow transfers to and from the disk drive on two separate channels. With SATA, there is only one active channel, which means you can only have one read or one write at a time. According to Seagate benchmarks, the combination of improved optimization algorithms in the queue management, two processor architecture, and hardware acceleration features on the Seagate Barracuda ES.2 SAS hard drives yields an average performance increase of 135% over SATA.

Another feature of SAS that trumps Enterprise SATA is SAS command queuing. Seagate's Constellation and Barracuda* ES.2 SAS have a clear advantage with two CPUs (versus one with SATA). Although both SATA and SAS offer command queuing, the second CPU on the SAS drive provides smarter and more robust command queuing, thus improving efficiency and performance.

Perhaps more important than the performance benefits of SAS hard drives are the enhanced error detection and error correction capabilities that they offer. In a Nearline environment, where large amounts of data are used for reference or archival purposes, it is crucial that the data be correct when it is ultimately needed.

With SAS drives, users can be confident that they have full error detection (IOEDC) and error correction (IOECC) on both reads and writes. Also, SAS drives provide data protection while data is in flight, which ensures the integrity of data during reads, writes, and data en route for complete peace of mind. SAS drive ASICs also include IOECC and IOEDC on all internal processor instruction and data memory structures, while SATA drives typically do not.

Although SATA drives do error correction on reads, if the errors exceed the ECC capability, SATA drives will continue to re-read the data for a long time before giving up. On enterprise SATA drives there is a deterministic limit to prevent the drive from being marked as bad by a RAID controller. SAS (and all previous versions of SCSI) have a maximum command time parameter that prevents this from occurring. Because SATA is half-duplex and single port, there is no way to talk to the drive to see if it is still functional while it is attempting to re-read the data. This can result in the drive erroneously being marked as "bad" or "missing" and taken offline, even if the drive was fully functional.

In applications where low-cost storage is a primary consideration, Seagate Constellation and Barracuda ES.2 SAS drives provide the necessary performance, availability, and reliability all at pennies more per gigabyte than comparable Enterprise SATA drives with interposer solutions. With significant feature and performance improvements over Enterprise SATA drives, these high-capacity Tier 2 nearline SAS drives help cut costs and reduce downtime for Low-Cost Server and Nearline applications. Like other enterprise-class drives, users can be sure that the data is available, quickly and accurately, when they need it.

Table 1. Tier 2 Nearline HDD Performance Comparison

Performance Comparing Seagate Barracuda* ES.2 SAS drives to Seagate Barracuda* ES.2 SATA drives (write cache off)
135% Average Performance Boost over SATA
✓ 110% Improvement in sequential reads
✓ 390% Improvement in sequential writes
✓ 20% Improvement in random reads
✓ 23% Improvement in random writes

Full Duplex

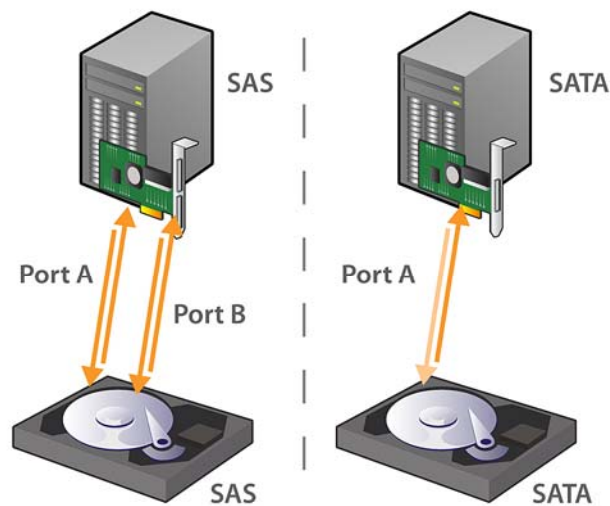
- Transmits in both directions at one time (orange arrows)

Dual Port

- Two redundant paths to each hard drive (Channel 1 and 2)

Two Active Channels

- Hard drive can communicate via two channels for concurrent transfers.
 - 2 concurrent Writes or
 - 2 concurrent Reads or
 - 1 Write + 1 Read



Half Duplex

- Transmits in one direction at a time

Single Port

- Single path to each hard drive (Channel 1 and 2)

One Active Channel

- Hard Drive can communicate through only one channel

Figure 3. Tier 2 SAS Features for Nearline Applications




Figure 4. SAS/SATA Electronics

3.2 Enterprise SATA, Tier 2 Nearline SAS, and Desktop SATA Differences

Today, many users choose Desktop SATA drives for Low-Cost Server environments due to the large capacities and low cost. However, Desktop SATA drives were not designed for the rigors of Low-Cost Server applications and using them in these environments can be very detrimental to server reliability and availability. In the end, using Desktop SATA in Low-Cost Servers can actually cost more than using Enterprise SATA or Tier 2 nearline SAS drives due to server downtime and replacement/maintenance expenses.

There are significant differences between Desktop SATA and Enterprise SATA drives, including those related to reliability, command queuing, error recovery, staggered spin up, workload, and rotational vibration tolerance.

From a reliability standpoint, Enterprise SATA drives have nearly twice the mean time between failure (MTBF) of Desktop SATA drives. While this number already seems large, keep in mind that the test environment for the Desktop drives is significantly less stressful in terms of duty cycle than that of Enterprise drives. This means that the MTBF for Desktop drives would be even lower (worse) when used in more stressful environments, like that of Low-Cost Servers.

On Desktop SATA drives, Native Command Queuing (NCQ) has traditionally been and in some cases is still an optional feature, while it is standard on Enterprise SATA drives. NCQ is a disk technology designed to increase performance by allowing the drive to order commands intelligently. Desktop drives that do not support NCQ must perform commands in the order that they were received. Enterprise drives optimize the order in which it performs commands, which greatly reduces seek times and improves overall system performance.

Enterprise SATA drives provide improved error recovery capabilities, which prevent drives from being taken offline prematurely. Desktop SATA drives stop communicating while they attempt to finish a problematic command. Enterprise SATA drives notify the system that “I have a problem and I am working on it”. This prevents any unnecessary downtime.

Staggered spin-up is an optional feature on some desktop SATA hard drives. Without staggered spin-up, hard drives spin up when power is applied. In systems with a large number of hard drives, this can cause a significant power surge thereby reducing the MTBF of the system power supply. Enterprise SATA hard drives do not spin up their platters when power is applied. Instead, they wait until told to do so. With staggered spin-up, Intel® RAID adapters can spin up the drives one at a time. This provides greater overall reliability and prevents power surges.

Another significant difference between Desktop SATA drives and Enterprise SATA drives is the way that they handle workload. Designed for far lower I/O processing (IOP) workloads than Enterprise SATA or SAS drives, Desktop SATA drives do not have safeguards to protect against extreme workloads. Higher duty cycles and higher system temperatures can ultimately work a desktop drive to premature death whereas Enterprise SATA drives are built for 24x7 workloads from the ground up.

Lastly, Enterprise SATA drives have a greater ability to tolerate vibration caused by adjacent disk drives in the chassis. When the rotational vibration exceeds a drives' acceptable limits, the head may move off the data track during read/write operations. This causes the drive to retry the read/write to ensure the integrity of the data, but it also results in overall performance degradations. Desktop drives are designed for use in low-temperature, single or dual-drive environments where read/write duty cycles are low. Enterprise SATA drives are built for multi-drive environments where temperatures are higher and read/write duty cycles are very high and occurring 24 hours a day, 7 days a week.

Although the need for inexpensive, high-capacity storage is essential for Low-Cost Server environments, it should not come at the cost of reliability, availability, and performance. Desktop SATA may have been "good enough" in the past, but today Enterprise SATA drives provide incomparable enhancements. And because Tier 2 nearline SAS drives, like the Seagate Constellation and Barracuda ES.2, offer benefits above and beyond Enterprise SATA drives at a very small cost delta (see previous section), it is logical that these would be the ideal solution for Low-Cost Server environments where reliability, availability and performance are just as important as cost.

3.3 Feature List Summary

There are multiple criteria for choosing Enterprise SATA drives over Desktop SATA drives including:

- Reliability
- Command queuing
- Error recovery
- Staggered spin up
- Workload
- Rotational vibration tolerance

Table 2. Feature List Summary

Feature List	Desktop SATA	Enterprise SATA
Reliability	750,000 Hrs	1,200,000 Hrs
Native Command Queuing (NCQ)	✓ (optional) ¹	✓
Error Recovery	Drive can go AWOL	Timeout set by host
Staggered Spin Up	✓ (optional) ¹	✓
Workload	Buyer Beware (vulnerable to overuse) 8 x 5 typical	Managed (drive protects itself) 24 x 7 typical
Rotation Vibration Tolerance (Tolerance to other drives in the chassis)	8.5 radians/s 10 to 300 Hz	12.5 radians/s 20 to 1500 Hz
Data Path Error Detection		✓

¹ Optional features on some desktop drives not including Seagate Barracuda 7200.

4. Intel® RAID SAS/SATA adapters

This section explains why Intel® RAID SAS/SATA adapters are ideal for pairing with Seagate Constellation and Barracuda ES.2 SAS drives in Low-Cost Server and Nearline Applications.



Figure 5. Intel® RAID SAS/SATA Adapters

Intel® RAID SAS/SATA adapters offer the added fault tolerance, data security, availability and performance required by Nearline and Low-Cost Server environments. With a diverse portfolio of entry, mainstream and scalable performance RAID controllers, Intel offers a solution for any storage requirement.

The combination of Intel® RAID SAS/SATA adapters and Seagate Constellation and Barracuda ES.2 SAS drives offers a unique benefit to system builders and channel OEMs. With current Intel® RAID software releases, most Intel® RAID controller products now support single-controller Multipathing (failover) and Load Balancing. Together, these two new features take full advantage of the dual-port, full-duplex, and concurrent active channels on the Seagate Tier2 nearline SAS hard drives.

Multipathing takes advantage of dual-port drives by detecting and using multiple paths of communication. This increases reliability and fault tolerance in cases where the cabling becomes faulty or disconnected, or where one port on a drive becomes inactive. In this instance, Multipathing allows the Intel adapter to access the hard drive via an alternate path, thereby maintaining communication and eliminating downtime.

Load balancing, working in tandem with Multipathing, is a powerful feature for improving application performance. While Multipathing allows for use of multiple paths, it is Load Balancing that intelligently distributes read/write requests to the multiple paths. This maximizes throughput between the Intel adapter and the hard drive, removes bottlenecks, and improves

overall system performance. In fact, Intel has seen Load Balancing improve large block sequential reads by up to 40% and modestly improve random writes as well.

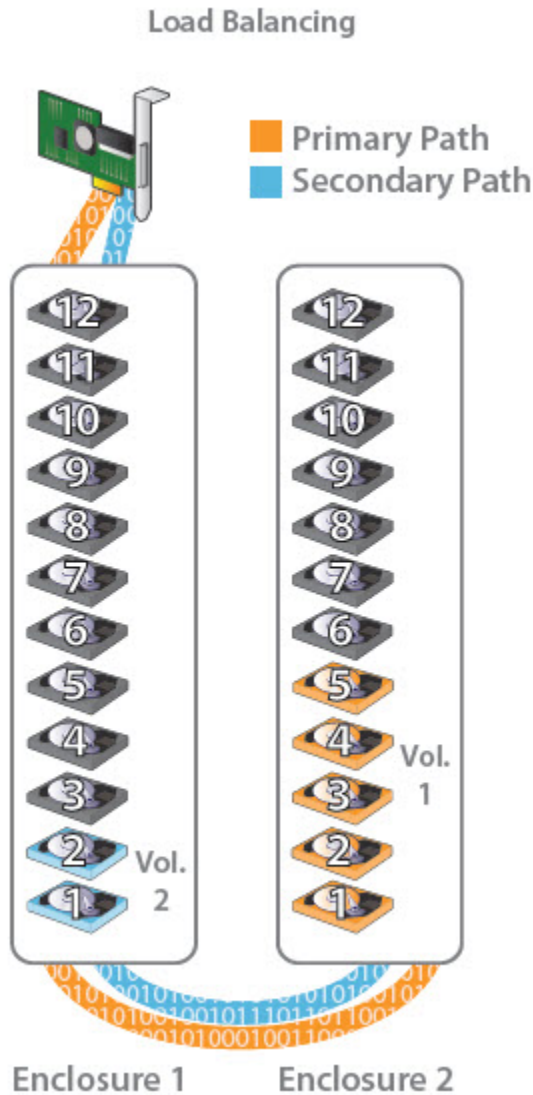


Figure 6. Load Balancing

Intel® RAID SAS/SATA adapters used with Seagate Constellation and Barracuda ES.2 SAS hard drives provide the most ideal storage solution for Low-Cost Server and Nearline applications. Not only are system builders and other channel customers able to purchase these SAS building blocks at SATA pricing, but they can also take full advantage of the SAS performance and reliability features.

For more information on Seagate Constellation or Barracuda ES.2 SAS drives or Intel® RAID SAS/SATA adapters, please visit:

www.seagate.com

www.intel.com/go/RAID