

Intel[®] RAID Performance

SAS 2.0 RAID Mainstream Controllers

March 2010
Version 1.0

Contents

Revision History	3
Disclaimers.....	3
Introduction.....	4
Architectural Benefits of SAS 2.0 Capable RAID.....	4
Performance Tests and Results.....	5
Intensive streaming applications.....	6
Transaction-oriented applications.....	8
Maximum Throughput.....	9
Maximum IOPs.....	10
SAS 2.0 Intel RAID: Best-in-Class Storage Portfolio	12
Summary.....	12
Hardware Configurations Tested.....	13

Revision History

Date	Revision Number	Modifications
March 2010	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.

Performance tests and ratings are measured using specific computer systems and/or components and reflect the approximate performance of Intel products as measured by those tests. Any difference in system hardware or software design or configuration may affect actual performance. Buyers should consult other sources of information to evaluate the performance of systems or components they are considering purchasing. For more information on performance tests and on the performance of Intel products, visit Intel Performance Benchmark Limitations at http://www.intel.com/performance/resources/benchmark_limitations.htm. 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 2010. Portions Copyright © LSI Corporation 2010.



www.intel.com/go/RAID

Introduction

Intel's SAS 2.0 RAID controllers, featuring the LSI SAS2108 SAS RAID-on-Chip (ROC) processor, offer significant streaming read and write performance enhancements for solutions architected with 3Gbps or 6Gbps drives. Their superior read/write performance makes them ideally suited for a broad range of application workloads, from compute intensive data center applications such as email, OLTP, and database to storage and content applications such as streaming video, and data archival.

Today's server applications demand storage technologies that deliver fast and reliable access to data. This paper will show how channel customers can make a low-cost investment in a SAS 2.0 Intel RAID Controller to significantly improve the performance of:

- Existing systems using 3Gb/s SATA and SAS hard drives
- End-to-end 6Gb/s SAS ecosystems
- Solid State Drive environments
- Throughput-intensive streaming applications
- Transaction-oriented server and database applications

In addition to outstanding performance, Intel's SAS 2.0 products such as the 6Gb/s-capable Intel RAID RS2BL080, RS2MB044, and RS2PI008 controllers support advanced features such as multi-pathing, VMWare ESXi virtualization, UEFI 2.0, and revertible hot-spare. Furthermore, they are highly flexible with support for high-performance SAS hard drives, high-capacity SATA hard drives, and energy-efficient solid state drives (SSDs).

Architectural Benefits of SAS 2.0 Capable RAID

Intel's RS2 generation of RAID products offer several architecture changes aimed to provide new or enhanced benefits. The new LSI SAS2108 chip combines the processing power of an 800MHz dedicated RAID processor with the added benefit of an integrated SAS 2.0, 6Gb/s controller. The faster clock speed of the controller coupled with dedicated DMA channels for every Phy results in a higher number of XOR calculations per second. This dramatically increases the RAID 5 and RAID 6 write capability of the controller as these write operations are heavily dependent on XOR calculations and direct memory access.

In addition, the integrated SAS 2.0, 6Gb/s capability of the controller allows for twice the SAS throughput of the previous generation with up to 600MB/s per controller port transfer speed. To support this increased bandwidth, Intel RAID SAS 2.0, 6Gb/s controllers include 512MB of on controller cache memory and a PCI Express 2.0 bus interface.

All of the hardware features come together with a RAID firmware stack that can support more than twice the number of I/Os per second (IOPs) of the previous generation. This is very important in order to support the latest generation of Solid State Drives currently on the market. Solid state drives are capable of performing more than 100 times the random activity of a standard SAS or SATA mechanical hard drive. The added IOPs capability of Intel 6Gb RAID products provides a cost effective means to support an SSD environment with high performance read and write operations.

Performance Tests and Results

In comparing the SAS 2.0 Intel RAID RS2PI008 with a leading competitor's 3Gb/s offering -- hence to be referred to as Competitor A -- Intel RAID outperformed Competitor A in a comprehensive list of typical SATA, SAS, and SSD applications. (Note: At the initial release date of this paper, Competitor A does not offer a SAS 2.0, 6Gbps product. Therefore, only comparison to their 3Gb/s product is possible.)

Using IOMeter, an open-source I/O workload generator, Intel used I/O workload profiles to demonstrate how real-world applications might perform. Streaming read and write benchmarks are featured for throughput intensive applications that are well suited for SATA hard drives performing heavily sequential reads and writes. Real-world benchmarks are featured for transaction-oriented applications that are well-suited for SAS hard drives and SSDs. These applications include:

- Streaming Read workloads represent contiguous read requests made to the disk. They are predominately sequential I/O and are typically seen in applications such as Media servers (e.g. video-on-demand) or Virtual Tape Libraries (VTL)
- Streaming Write workloads represent contiguous write requests made to the disk. They are predominately sequential I/O and are typically seen in applications such as Media Capture, (e.g. video surveillance, medical imaging, etc), VTL, and archival/backup
- Workstation workloads represent single users running multiple applications where they are a predominately handling random read requests, with a mix of sequential reads and writes.
- Web Server workloads typically service multiple, predominately random simultaneous requests.
- OS Drive workloads represent the I/O activity of an operating system drive. They are handling random I/O with a read/write ratio of 70:30.

- File server workloads represent a large number of users requesting access to unstructured reference data and files. They are handling random read/write requests with a ratio of 80:20.
- Email server workloads typically service multiple simultaneous requests to transport or store email. They are handling a fairly even mix of random reads and random writes.
- OLTP workloads represent transaction processing and database query I/O requests that require a read of the original value, followed by an update, write and verify operation. They handle predominately random requests with two reads for every write.

The following spreadsheet outlines the access pattern characteristics for each workload.

	% Reads	% Writes	% Sequential	% Random
Streaming Read	100%	0%	100%	0%
Streaming Write	0%	100%	100%	0%
Workstation	80%	20%	20%	80%
Web Server	100%	0%	0%	100%
OS Drive	70%	30%	0%	100%
File Server	80%	20%	0%	100%
Email Server	50%	50%	0%	100%
OLTP	67%	33%	0%	100%

Intensive streaming applications

Because of their excellent sustained media rates, applications with heavy sequential reads and writes are well-suited for SATA hard drives. These include applications such as video/audio streaming, nearline VTL and disk-based backup. The SAS 2.0, 6Gb/s Intel® RAID Controllers produce up to 32% better sequential writes and up to 30% better sequential reads over Competitor A's product using SATA hard drives in 3Gb environments. When compared in a 6Gb infrastructure, Intel is able to outperform Competitor A by two to four times.

Streaming applications are all about efficiency, storage capacity, and throughput. RAID 5 and RAID 6 are two of the most popular RAID levels for these applications because they deliver data protection, good performance, and the most efficient use of drive capacity. In Figure 1 and Figure 2, the Intel RAID RS2PI008 shows superior performance versus the Competitor A in the same 3Gb/s SATA environment with these two most popular RAID levels as well as other commonly used RAID levels. (Note that – although only Intel RAID Controller RS2PI008 was used for testing, identical performance from the RS2BL080 and the RS2MB044 is expected given the similar architecture of these products. The four port RS2BL040 controller performance may be similar for some workload profiles.)

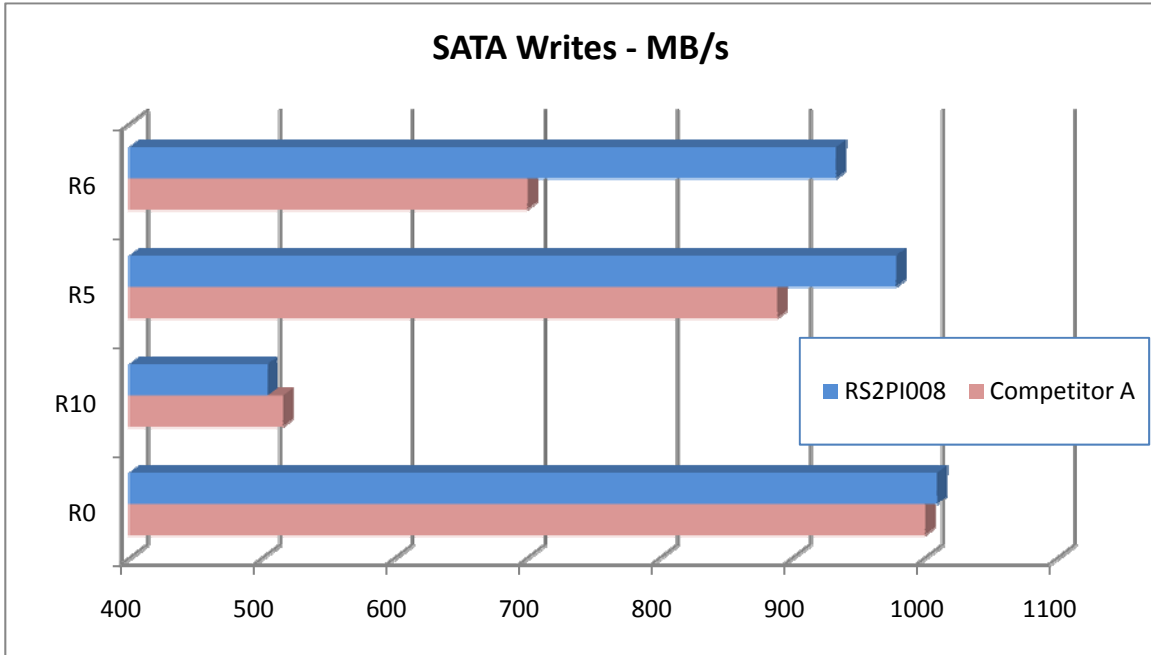


Figure 1 – Competitive performance analysis for controllers with 24 3Gb/s 7200 RPM SATA hard drives in write-intensive environments

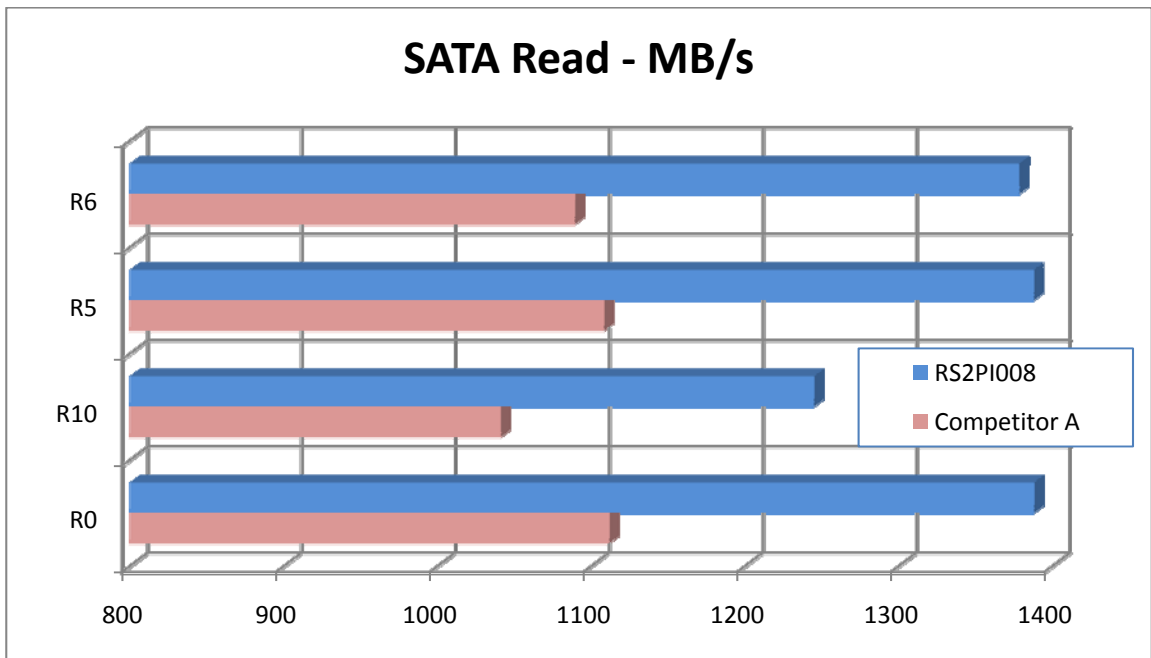


Figure 2 – Competitive performance analysis for controllers with 24 3Gb/s 7200 RPM SATA hard drives in read-intensive environments

The benefit of the SAS2.0 design of the Intel RS2PI008 controller is clearly defined in preceding figures (1 and 2). Even when operating in a 3Gb environment with 3Gb drives attached, the Intel controller performance is outstanding.

Transaction-oriented applications

Applications dominated by random I/O access patterns are well-suited for SAS hard drives or SSDs. These applications include email servers, web servers, databases, data warehousing and OLTP (Online Transaction Processing). In real-world benchmarks, the Intel RAID RS2PI008 consistently showed higher performance than the Competitor A. Whether customers are using 3Gb/s or 6Gb/s drives, investing in Intel RAID 6Gb/s SATA+SAS controllers can net performance results that are not possible with the Competitor A.

As shown in Figure 3, the 6Gb/s Intel RAID RS2PI008 provides superior performance over Competitor A even when connected to 3Gb/s drives. When the RS2PI008 is coupled with 6Gb/s drives, even greater performance can be achieved.




Figure 3 - Competitive real-world workload performance comparison for Intel RAID RS2PI008 controller and Competitor A using twenty-four 15k SAS drives in a RAID 5 array.

As customers look to increase the overall performance of their storage infrastructure, SSDs will continue to make inroads into the enterprise server market. SSDs are well-suited for server workloads requiring very high IOPs. Applications that fit this profile include online transaction systems with small, random reads and writes such as reservation and ecommerce systems.

The true performance advantage of Intel RAID 6Gb/s controllers are best demonstrated with SSDs. Only with SSDs can one see the full performance capabilities of the Intel RAID RS2PI008. Figure 4 shows up to an 82% performance advantage over the Competitor A.




Figure 4 – IOPs Competitive real-world workload performance comparison for Intel RAID RS2PI008 and Competitor A using eight SATA Solid State Drives (SSDs) in a RAID 0 array.

Maximum Throughput

Throughput of the controller is measured in Megabytes per second (MB/s). Approaching maximum throughput of the controller is achieved by utilizing a large number of hard disk drives, where drive sustained throughput is aggregated through the use of a SAS expander which will allow multiple SAS or SATA hard disk drives to be fanned out from each controller port. Enabling virtual drives controller cache helps to reduce drive transaction latency. Read throughput is virtually the same across RAID 0, RAID 5, and RAID 6 as the read operations are the same. However, write operations for RAID 5 and RAID 6 require that an XOR operation must be performed on the data being written before it is sent to the drive. These calculations create additional write latency, although this is mitigated through the use of a high performance ROC chip on and larger on controller memory on the Intel SAS2.0 generation of controllers. Figures 5 and 6 below shows the additional throughput capability of the Intel Controller compared to Competitor A.

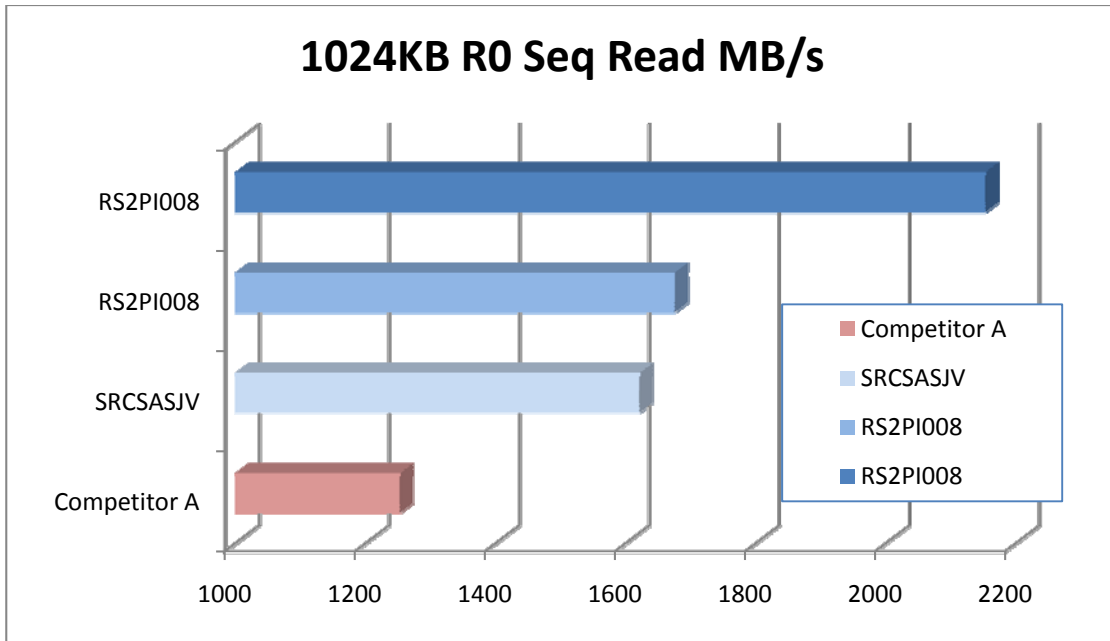


Figure 5 – RAID 0 Sequential Read with Read Ahead Cache Enabled using 24 SAS Drives

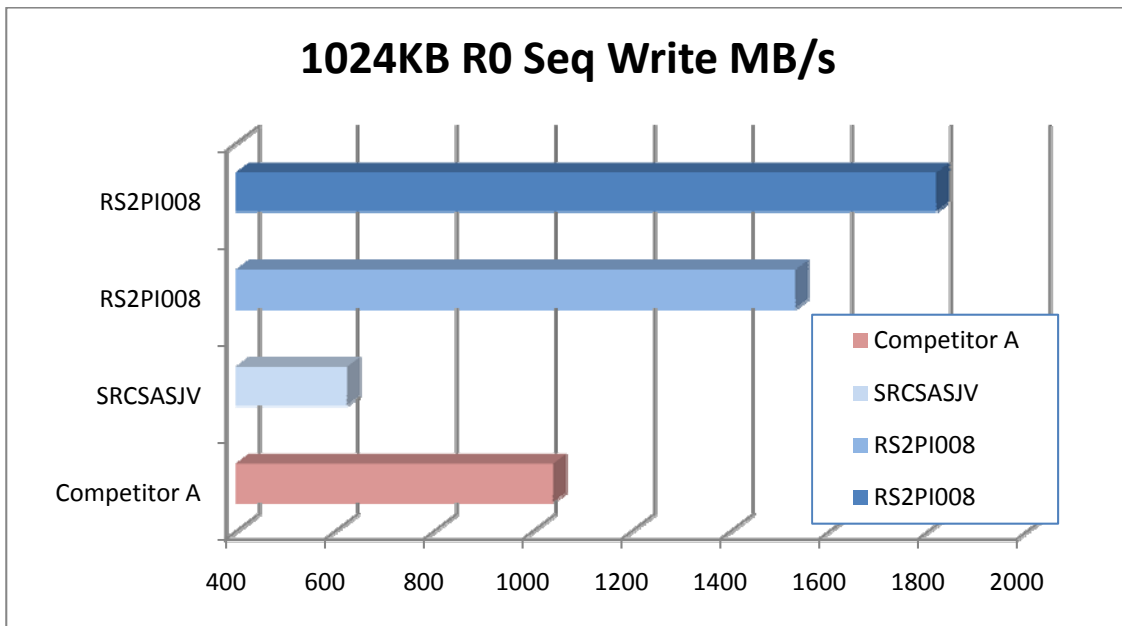


Figure 6 – RAID 0 Sequential Write Workload with Write Back Cache Enabled using 24 SAS Drives

Maximum IOPs

I/O per Second (IOPs) are very important when the usage model includes small random transfers. IOPs is best demonstrated using Solid State Drives (SSD), this is because SSD's are capable of more the 100 times greater IOPs than Hard Disk Drives. Like hard drives,

SSD's can read faster than they can write, but write capability of SSDs is comparatively high. The industry leading IOPs performance of the Intel SAS2.0 Raid controller is demonstrated in figures 7 and 8 below, the figures include the SAS 1.0 Intel SRCASJV controller for comparative purposes.

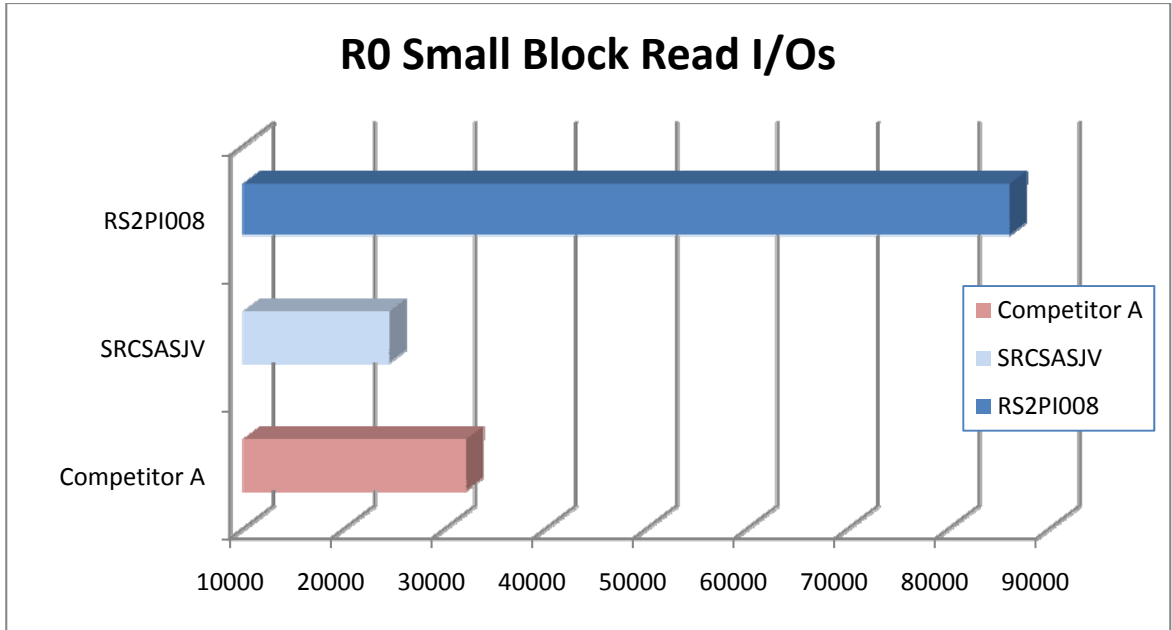


Figure 7 - Small Block Random Read I/Os per Second using 8 SSD drives




Figure 8 - Small Block Random Write I/Os per Second using 8 SSD Drives

SAS 2.0 Intel RAID: Best-in-Class Storage Portfolio

Intel is committed to providing our channel partners with a broad portfolio of best-in-class storage solutions. By being first to market with a complete lineup of high-performance SAS 2.0 storage adapters, Intel is enabling channel partners to build optimal storage solutions for a wide variety of applications.

The Intel RAID SAS 2.0 Mainstream product family currently includes 4-port and 8-port controllers that address high-performance storage requirements for both inside and outside the server. This line of controllers offers the perfect mix of value and performance for internal SATA and SAS systems. In mid-2010, Intel will launch a series of SAS 2.0 I/O modules that add value above standard add-in cards. In Q3 2010, Intel will also launch several Scalable Performance SAS 2.0 products, with both internal and external connectivity options, to provide the ultimate in performance for a large number of drives and scalability beyond the confines of the server chassis.

The following table provides an overview of Intel's Mainstream SAS 2.0 Add-In Card features:

Intel RAID Controller	Internal Ports	Internal Connectors	External Ports	External Connectors	Processor	Cache Memory	Host Interface
Intel RAID RS2PI008	0	•	8	SFF8088	LSI2108 ROC (800MHz)	512MB DDRII SDRAM	x8 PCI-E
Intel RAID RS2MB044	4	SFF8087	4	SFF8088	LSI2108 ROC (800MHz)	512MB DDRII SDRAM	x8 PCI-E
Intel RAID RS2BL080*	8	SFF8087	0	•	LSI2108 ROC (800MHz)	512MB DDRII SDRAM	x8 PCI-E
Intel RAID RS2BL040	4	SFF8087	0	•	LSI2108 ROC (800MHz)	512MB DDRII SDRAM	x8 PCI-E

Summary

When architecting a RAID solution, careful consideration should be given to factors that will allow for the best performance of the intended application. This document has described several "Real World" scenarios and benchmarks for consideration as good representatives of performance within these scenarios. Using the Intel® RAID Controller RS2PI008 for benchmarks and comparison, this white paper conveys the message that

Intel® RAID Controllers powered by LSI MegaRAID* technology offer exceptional performance.

Intel's new generation of SAS2.0 controllers provide world class performance in a comprehensive set of usage models. These controllers excel in both large block sequential transfers and in small block random I/O. Real world workloads include a combination of both of these transfer and I/O operations, making the Intel controller an excellent choice for use in Video Surveillance, File Services, Email services, and on line transaction processing (OLTP) transactions.

Furthermore, Intel offers a wide range of RAID products making it possible to architect a highly optimized and scalable solution specifically for your environment. By utilizing advanced technology, world-class features, and a SAS design that allows for SAS, SATA and SSD drives, Intel RAID allows you to balance performance, reliability and cost requirements.

Hardware Configurations Tested

Intel Server Board S5500HCV

BIOS 38

BMC 40

Intel Chassis SR5600

HSC 1.12

Intel RAID Controller RS2PI008**

FW Version 2.60.03-0778

Driver Version 4.23.0.64

Intel RAID Controller SRCASJLV

FW Version 1.40.42-0615

Driver Version 4.23.0.64

Competitor A RAID Controller

FW Version – Latest public release at the date of this testing

Driver Version – Latest public release at the date of this testing

Drive Enclosures

LSI620J – 6G 2.5" SAS drive tests

LSI3600 – 3G 3.5" SATA drive tests

Direct connect – 3G 2.5" SATA SSD drive tests

Drives

ST3500320NS 3Gb/s 7.2K SATA drive

ST3146356SS 3Gb/s 15K SAS drive

ST9146852SS 6Gb/s 15K SAS drive

SSDSA2SH032G1GN 3Gb/s SATA Intel SSD

Test Suite

Windows Server 2008 x64

IOMeter Version 2006.07.27.win32.i386

****Note:** Although testing was performed with the RS2PI008 controller, similar performance from the RS2BL080 and the RS2MB044 is expected based on the similar architecture of these products. The four port RS2BL040 controller may yield comparable performance for some workload profiles