IBM and HP 6-Gbps SAS RAID Controller Performance

Evaluation report prepared under contract with IBM Corporation

Introduction

With increasing demands on storage in popular application servers, the server manufacturers have begun to provide 6-Gbit/sec (Gbps) SAS RAID controllers as the standard RAID controllers. It is useful to understand how well these controllers work under sustained heavy loads, especially in real-world workload situations.

With this purpose in mind, IBM commissioned Demartek to test their standard 6-Gbps SAS RAID controller in one of their popular servers and compare its performance to the standard 6-Gbps SAS RAID controller provided by Hewlett Packard in one of their equivalent popular servers.

Evaluation Summary

In comparing these two 6-Gbps SAS RAID controller adapters, especially in real-world workloads, we found some similarities in their performance and some differences. These real-world workloads include a mixture of reads and writes in various block sizes. For example, typical web-server, file-server, workstation and database applications read and write 4K, 8K and 64K blocks in a random fashion, often with multiple reads and writes occurring at the same time in a "bursty" fashion. Other applications, such as media streaming or video-on-demand applications, read and write their data blocks sequentially with a consistent block size, such as 32K or 64K, in parallel as each camera feeds its data to the storage device, or as each video is served out at a relatively steady pace.

For the lighter real-world workloads (smaller queue depths), the performance of the two controllers was very similar. However, as the workloads increased (running higher queue depths), we began to see a pattern in all the RAID configurations (RAID0, RAID10, RAID5 and RAID6) where the IBM controller continued to scale upwards but the HP controller performance curve began to flatten. This suggests that under heavy load, the IBM controller will sustain its I/O rates but may begin to be a concern for the HP controller.

For the media server workloads, the IBM controller showed noticeably higher throughput rates when several simulated video streams (threads) were processing in parallel.

We believe that real-world applications are of most interest to customers. It is our conclusion that in most real-world applications, especially as the workloads increase to heavy levels, the IBM ServeRAID M5015 SAS/SATA Controller provides superior performance to the HP Smart Array P410i RAID controller.

Evaluation Environment

The evaluation was conducted at the Demartek lab facilities in Arvada, Colorado. The specifications for the test system are listed in the table below. The basic goal was to take equivalent versions of popular models of servers using the 6-Gbps SAS array controller that is provided with that server. The servers were configured so that an array of eight (8) disk drives could be thoroughly tested with that 6-Gbps SAS array controller. The external boot drive was accessed through a separate controller so that normal operating system activities did not interfere with the testing of the array controller and its associated disk drives.

	HP System	IBM System
Model	ProLiant ML370 G6	System x3650-M2
Processor	Dual Intel® Xeon® X5570	Dual Intel® Xeon® X5570
Total Cores	8	8
Logical Processors	16	16
Memory	12 GB	24 GB
System Interface	BIOS	UEFI
Boot Controller	Onboard SATA	LSI SAS3442E-R
Boot Drive	Seagate Barracuda 7200.11, 500 GB	Seagate Barracuda 7200.11, 500 GB
Operating System	Windows Server 2008 SP1 Enterprise x64	Windows Server 2008 SP1 Enterprise x64
Array Controller	Smart Array P410i	ServeRAID M5015*
Controller Type	6-Gbps SAS 2.0 3-Gbps SATA	6-Gbps SAS 2.0 3-Gbps SATA
Controller Interface	On motherboard	PCI-Express 2.0 x8 slot
Controller Cache	512 MB	512 MB
Controller Firmware	2.50	2.0.03-0694[1]
Controller Driver	6.18.0.64	4.17.2.64[1]
Array Drive form factor	2.5 inch	2.5 inch
Array Drive Quantity	8	8
Array Drive Model	HP-branded Seagate Savvio 10K.3	Seagate Savvio 10K.3
Drive part number	ST9146803SS	ST9146803SS
Array Drive Capacity	146 GB	146 GB
Array Drive RPM	10K	10K
Logical Volume Type	MBR	GPT

* The IBM ServeRAID M5015 SAS/SATA Controller is designed and supplied by LSI Corp.

The external boot drive for the IBM system was connected to a popular 3-Gbps SAS/SATA controller. The external boot drive for the HP system was connected via an eSATA connection to one of the motherboard SATA ports. We believe that the choice of controllers for the external boot drive did not materially affect the performance of the 6-Gbps SAS controller and the array of drives connected to it.

[1] General availability on 11/17/2009

Although there were differing amounts of system RAM in the two systems tested, the memory consumed by the I/O workload application was considerably less than the physical RAM in either system. We believe that this difference in RAM did not materially affect the outcome of the tests.

Test Process

The test scenarios included a full suite of tests using IOmeter, an open-source I/O workload generator. In addition to the standard "four corner" tests of 100% read and write with 100% random and sequential I/O, we ran additional tests designed to emulate several real-world workloads. Each workload ran through a series of queue depths that simulate multiple concurrent I/O requests. All I/O tests were run on "raw" volumes. These workloads included:

- File server
- Web server
- Workstation
- OLTP
- O.S. paging
- Media reader
- Media writer

The I/O profile and specifications for each of the real world workloads are shown below.

File server

This is an attempt to emulate a typical file server workload. All the I/O for this profile is 80% read and 100% random I/O.

- 10% 512 byte blocks
- 5% 1K byte blocks
- 5% 2K byte blocks
- 60% 4K byte blocks
- 2% 8K byte blocks
- 4% 16K byte blocks
- 4% 32K byte blocks
- 10% 64K byte blocks

Web server

This is an attempt to emulate a typical web server workload. All the I/O for this profile is 100% read and 100% random.

- 22% 512 byte blocks
- 15% 1K byte blocks
- 8% 2K byte blocks
- 23% 4K byte blocks
- 15% 8K byte blocks
- 2% 16K byte blocks
- 6% 32K byte blocks
- 7% 64K byte blocks
- 1% 128K byte blocks
- 1% 512K byte blocks

Workstation

This is an attempt to emulate a typical workstation I/O workload. All the I/O for this profile is 80% read and 80% random I/O.

- 100% 8K byte blocks

OLTP workloads

Three separate OLTP workloads were defined using 67% reads and 100% random I/O.

1. 4K byte blocks
2. 8K byte blocks
3. 64K byte blocks

O.S. drive

This workload is designed to emulate the I/O workload of an operating system volume. The I/O for this profile is 70% read and 100% random.

- 100% 8K byte blocks

O.S. paging

This workload is designed to emulate an operating system paging file I/O workload. The I/O for this profile is 90% read and 100% sequential.

- 100% 64K byte blocks

Media reader

This workload is designed to emulate multiple streams of media reading, such as a video-on-demand (VOD) service. This workload performs 100% sequential reads of 24 streams in parallel from the disk array

- 100% 64K byte blocks

Media writer

This workload is designed to emulate multiple streams of media writing, such as a video surveillance system with multiple cameras feeding data simultaneously. This workload performs 100% sequential writes of 24 streams in parallel to the disk array.

- 100% 64K byte blocks

Configuration Comments

The two RAID controllers had some features in common and some features that were unique.

Basic Array Configuration

Both controllers could be configured at boot time from the BIOS and using management software running on Microsoft Windows Server. However, at boot time, only one logical volume could be created using the HP controller using the BIOS interface, while several volumes could be created using the IBM controller through the BIOS interface.

The management software application for both controllers provided access to all the array configuration parameters and can be used to configure several logical volumes on the array. Both provide a graphical user interface that is straightforward to use, although each has a different interface. A scriptable command-line utility is also provided with each controller that can be used to configure the controller and volumes, as well as interrogate the status, firmware levels, etc. of the controller.

RAID5 and RAID6 Initialization

For RAID5 and RAID6 volumes, a parity initialization process creates the proper parity across volumes on the array. Although I/O can be performed to arrays that have not been initialized, the performance is sub-optimal for un-initialized RAID5 and RAID6 volumes. In all cases, we waited for the initialization process to complete before running any I/O workload tests. For the tests, a single volume was created using the entire available capacity of the array of eight 146GB, 10K RPM SAS disk drives for all RAID types.

For the IBM controller, this process consistently took 15 – 30 minutes for RAID5 and RAID6 volumes. When the IBM controller was configured through the BIOS interface, a progress bar was displayed and the time to completion could be easily estimated. When the IBM controller was configured through the management software, one of the menus showed the initialization progress.

For the HP controller, the parity initialization process took considerably longer. For RAID5 arrays, the process took up to 6 hours, and for RAID6 arrays, the initialization process took up to 24 hours. This variability in initialization time appeared to depend on the number of volumes in the array and may also depend on the stripe size used for the volumes on the array. After some initial experimentation, we settled on a single volume and used 128K stripe size for all volumes. With these settings the initialization took approximately 5 hours for RAID5 and approximately 6 hours for RAID6.

The difference in RAID5 and RAID6 parity initialization times was dramatic and not what we expected. This is a concern and should be taken into account when deploying these systems. These times may vary with different capacities and speeds of disk drives.

Stripe Size

The default stripe size for the IBM controller is 128K for RAID0, RAID10, RAID5 and RAID6. The default stripe size for the HP controller is 128K for RAID0 and RAID10, 64K for RAID5 and 16K for RAID6. In order to make the tests as similar as possible, the stripe size was set to 128K for all tests. We believe that if the default stripe size for the HP controller had been used for RAID5 and RAID6, the performance of the HP system would have been lower.

Read Ahead

The IBM controller has a "read ahead" feature, that can improve read performance. Read ahead can be set to "on", "off" or "adaptive". Adaptive read-ahead will turn read ahead on but will adapt the read ahead functions to the current I/O activity and will only use read ahead if it appears that there will be a benefit to using it. For all tests, we set the read ahead to "adaptive".

The HP controller did not have a read-ahead setting.

Cache Read-Write Mix

The HP controller's cache can be adjusted to specific read-write mixes, and are expressed as a mixture of percentages. The percentages can be set in 25% increments, ranging from 0% read and 100% write to 100% read and 0% write. For all tests, we set the cache mix to 50% read and 50% write, because we were running a variety of real world workloads.

The IBM controller did not have a cache read-write mix setting.

Disk Drive Cache

Both controllers allowed the disk drive cache to be enabled or disabled. In order to achieve the best performance, we enabled the disk drive cache for all tests for both systems. There is the potential for data loss in the event of a power failure when the disk drive cache is enabled. We connected both systems to a UPS for our tests and did not experience any electric power outage during our tests.

Summary of Controller Settings

	HP Controller	IBM Controller
RAID 0 Stripe Size	128K	128K
RAID 10 Stripe Size	128K	128K
RAID 5 Stripe Size	128K	128K
RAID 6 Stripe Size	128K	128K
Read Ahead	N/A	Adaptive
Read-Write Cache Mix	50% read, 50% write	N/A
Disk Drive Cache	Enabled	Enabled
RAID 5 initialization time	5 hours	20 minutes
RAID 6 initialization time	6 hours	30 minutes

Test Results - Overall Results

In the 4K OLTP workload, both controllers appear to scale evenly until the heaviest queue depths, where the IBM controller outperforms the HP controller.

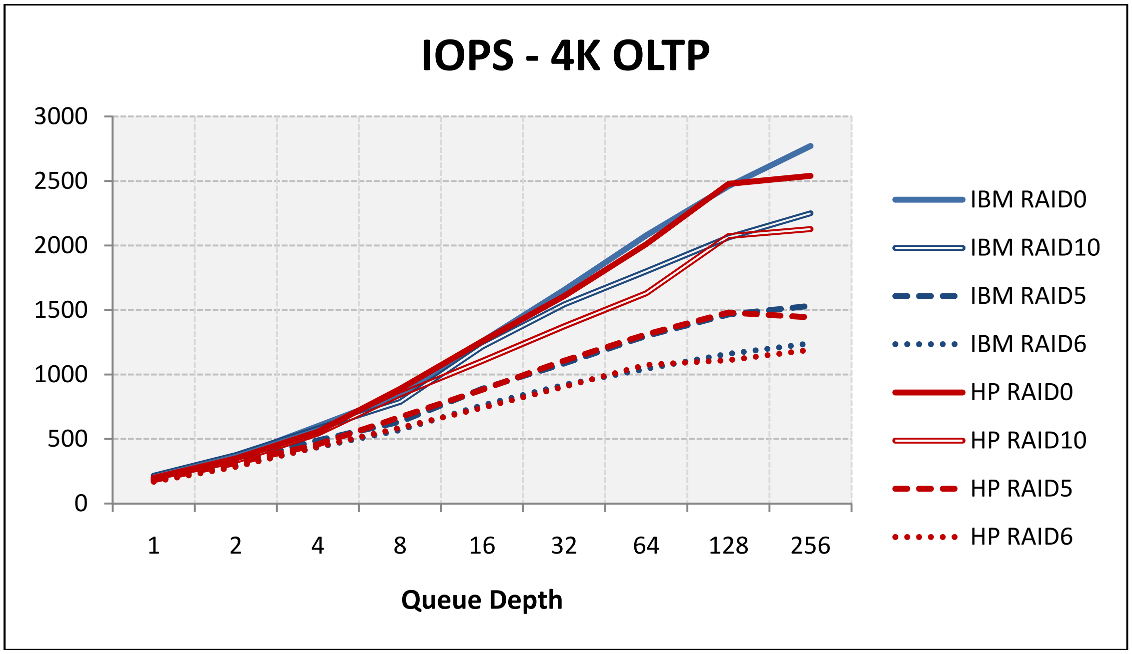

In the 8K OLTP workload, both controllers appear to scale evenly until the heaviest queue depths, where the IBM controller outperforms the HP controller.

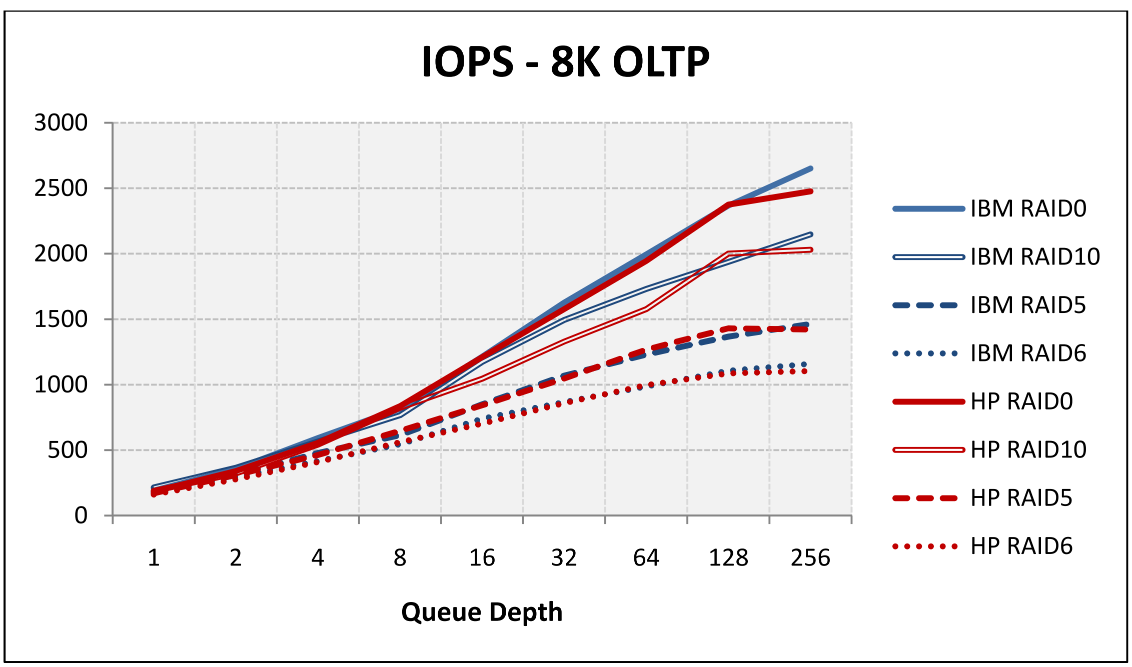

In the 64K OLTP workload, both controllers appear to scale evenly until the heaviest queue depths, where the IBM controller outperforms the HP controller, except for RAID6, where the HP controller performs better at the higher queue depths.

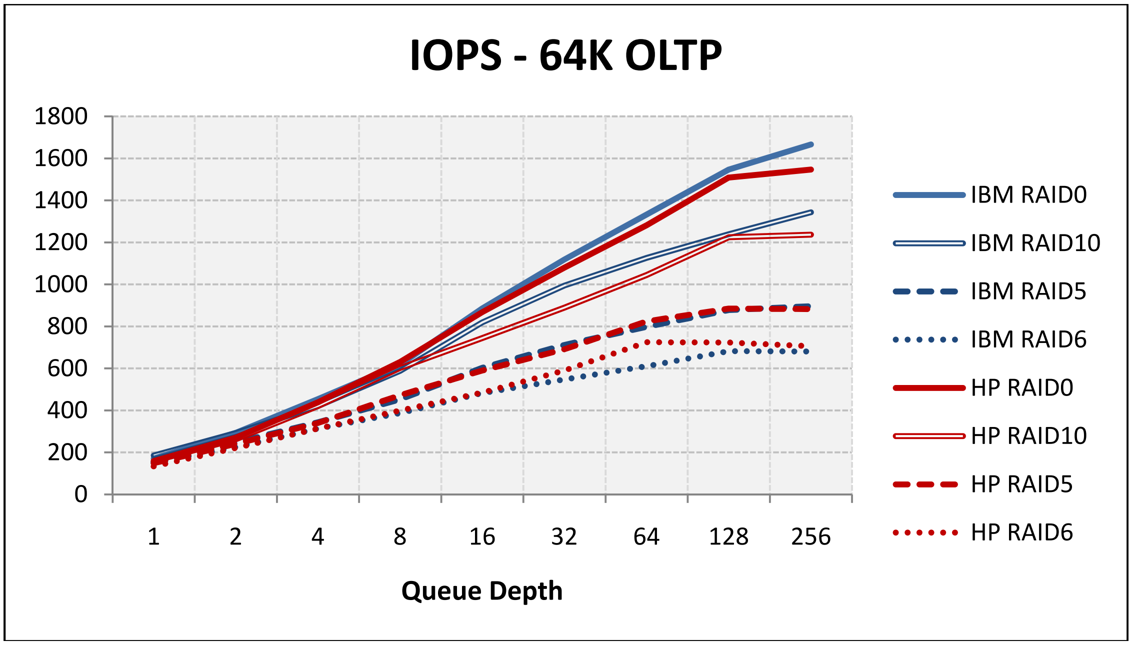

In the email server workload, both controllers appear to scale evenly at the lower queue depths. For RAID0 and RAID10, the HP controller has an interesting peak at queue depth of 128. Note that the HP controller performs better in RAID5 at the heavier queue depths.

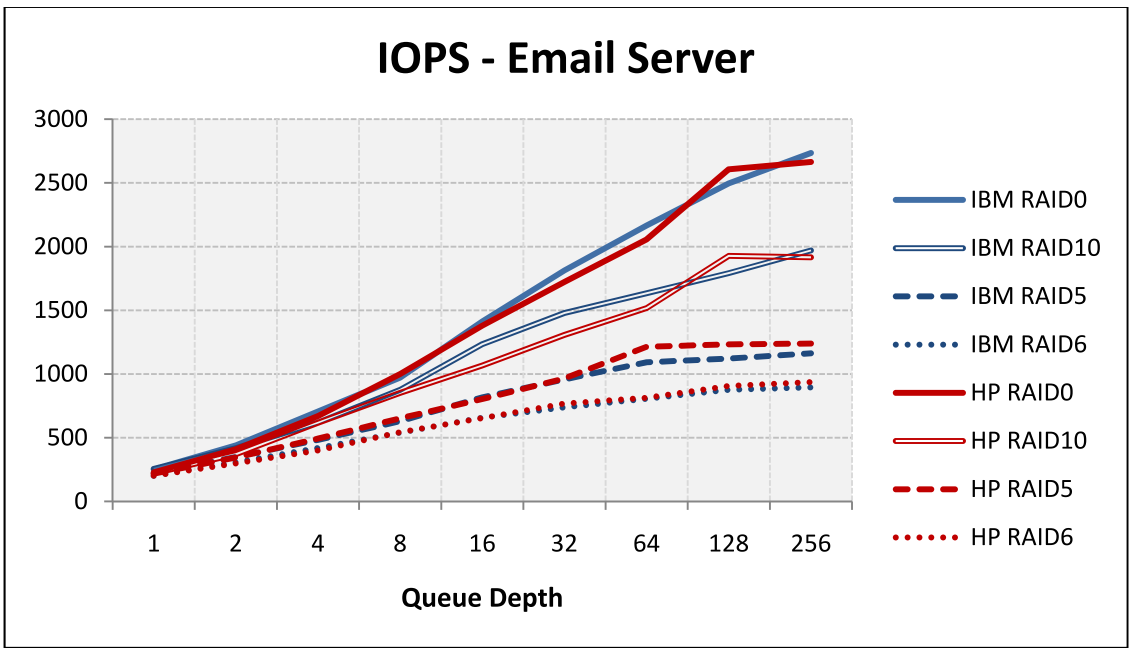

In the file server workload, both controllers appear to scale evenly until the heaviest queue depths, where the IBM controller outperforms the HP controller.

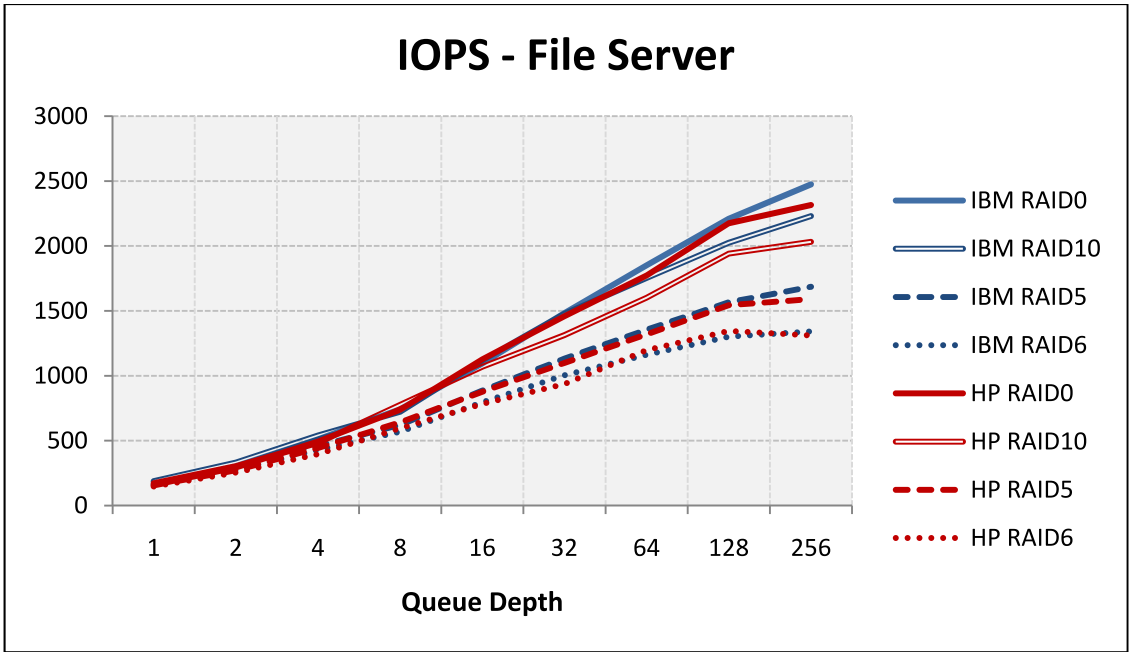

In the web server workload, the performance is much more even until the heaviest queue depths, where the IBM controller outperforms the HP controller. RAID10 performance is higher for both controllers than the other RAID types.

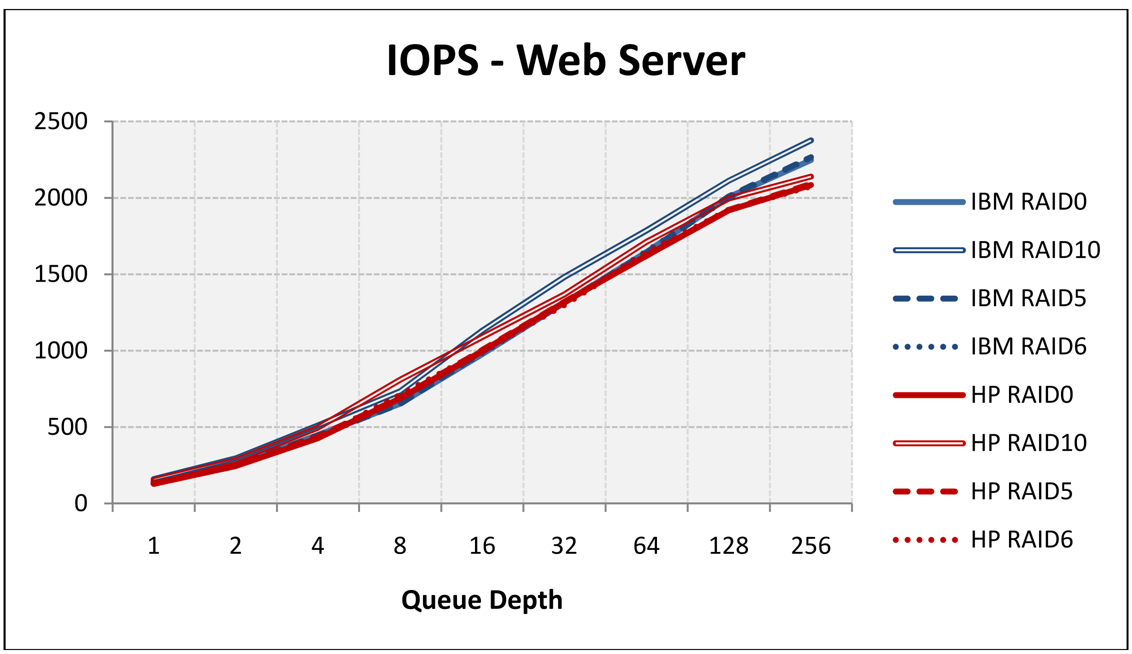

In the workstation workload, both controllers appear to scale evenly until the heaviest queue depths, where the IBM controller outperforms the HP controller.

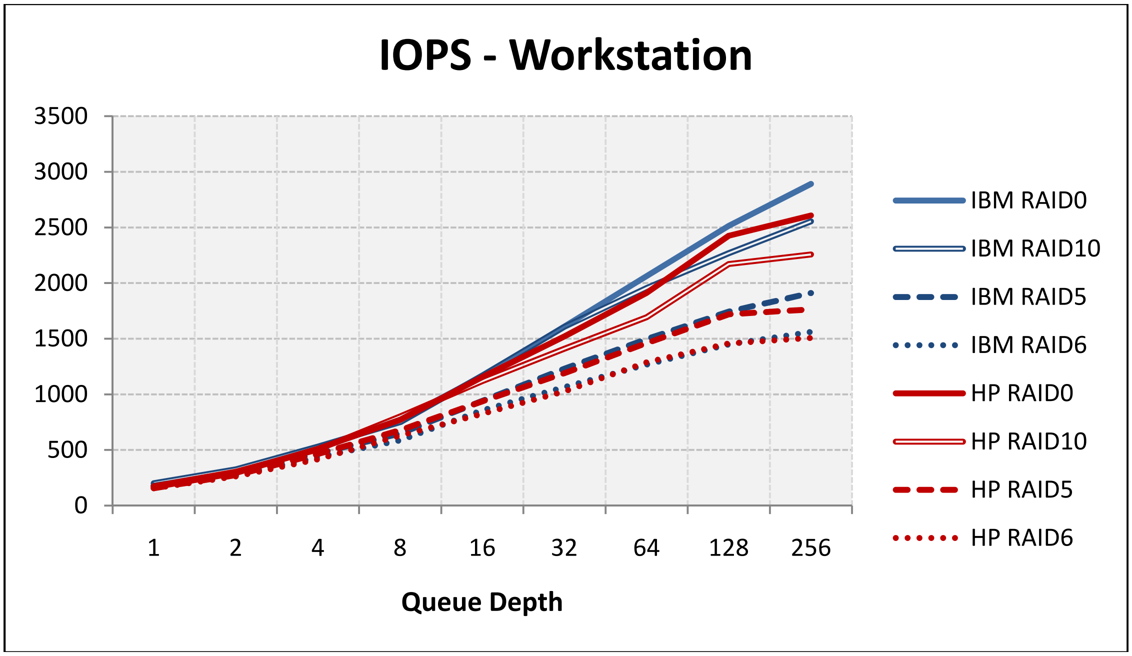

In the operating system drive workload, both controllers appear to scale evenly for RAID0 and RAID10 until the heaviest queue depths, where the IBM controller outperforms the HP controller. For RAID5 and RAID6, the performance is much more even.

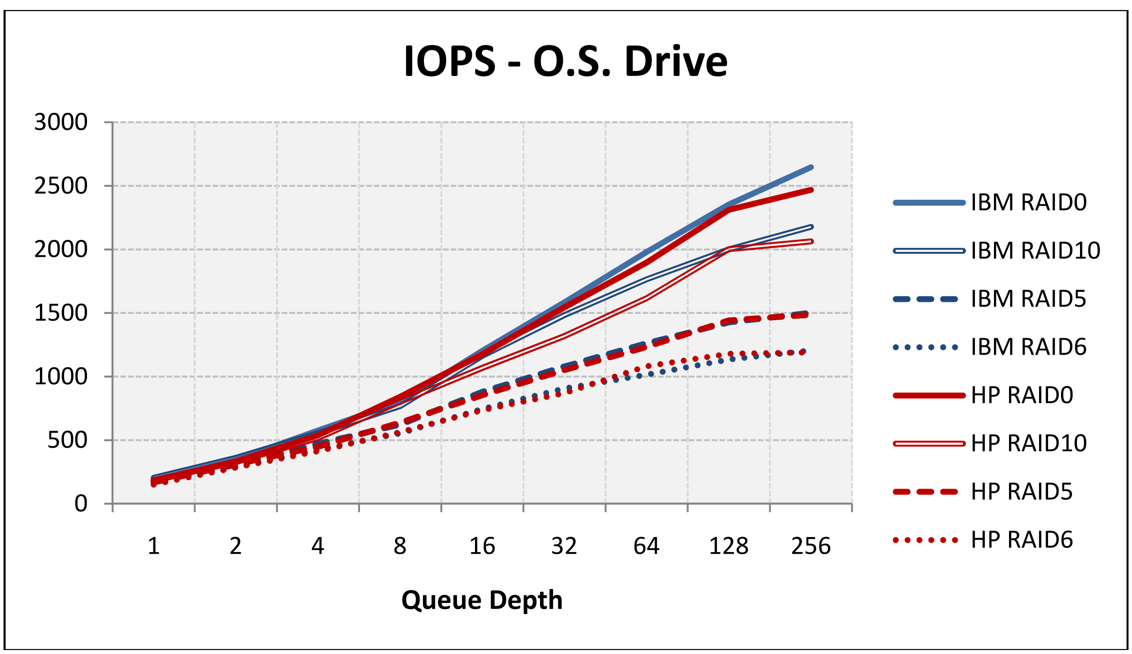

The operating system paging performance is quite different than the other workloads. In this case, the IBM controller has a clear advantage at the lower queue depths, but the HP controller has the advantage in the highest queue depths.

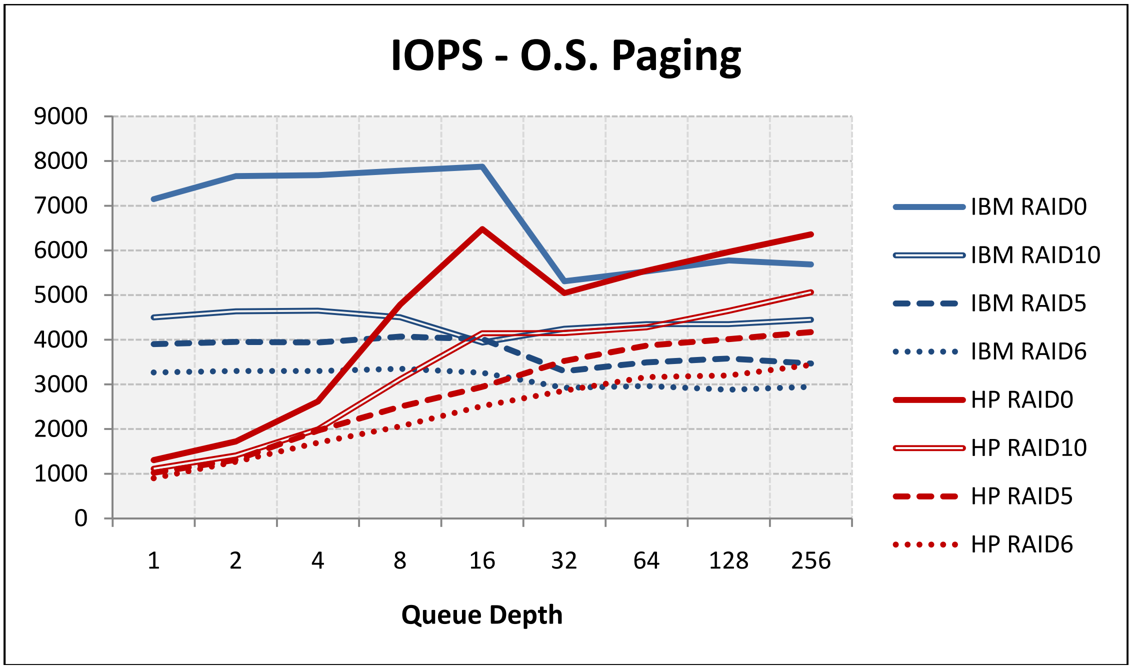

In the media reader workload, the IBM controller performs significantly better up to the range of 12-19 simultaneous threads (video streams), depending on RAID type.

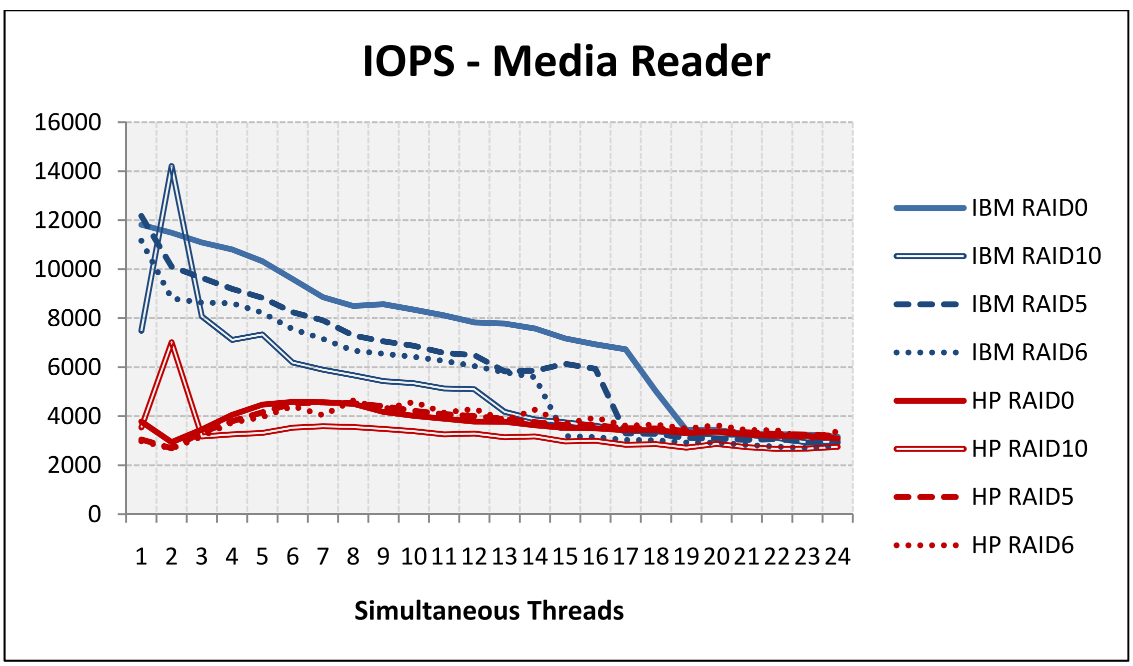

In the media writer workload, the IBM controller performs significantly better when higher numbers of simultaneous threads (video streams) are active.

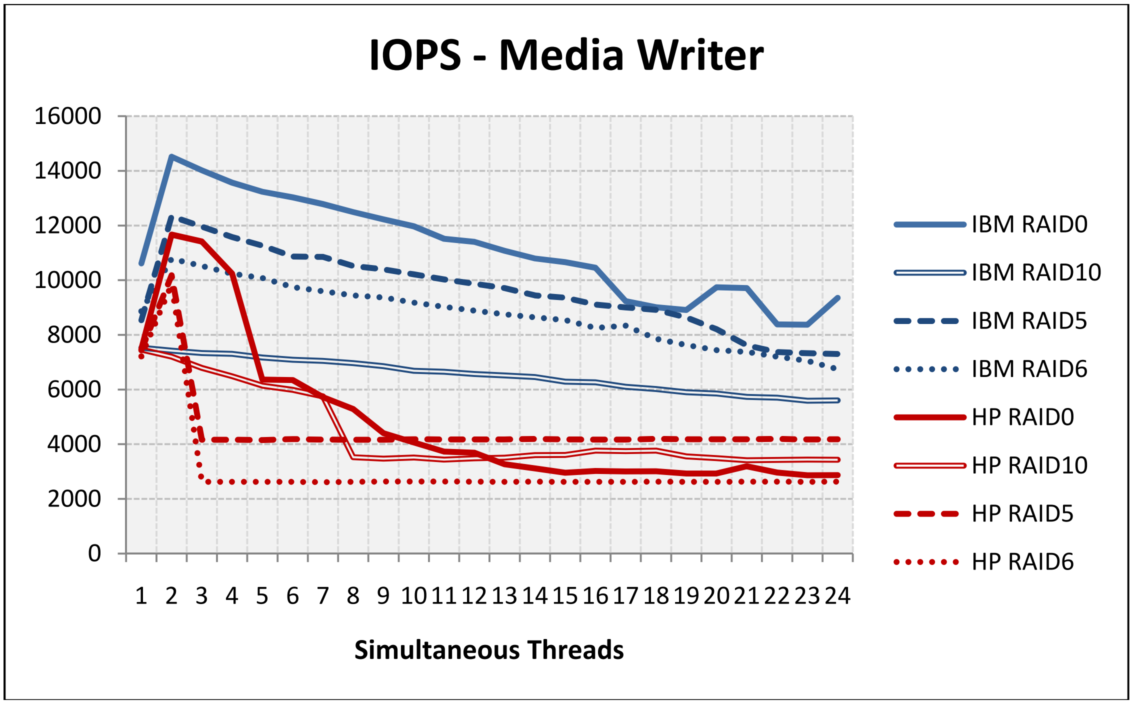

Test Results - Selected Data Points

This graph shows the comparison of all of the heavy real-world workloads tested at a queue depth of 256 for RAID0. The other RAID types showed similar patterns, but only the RAID0 data is included here to keep the graph relatively simple. This data shows a subset of data from the other graphs previously shown in this report.

In all but one workload, the IBM controller provides better performance at the heaviest queue depth.

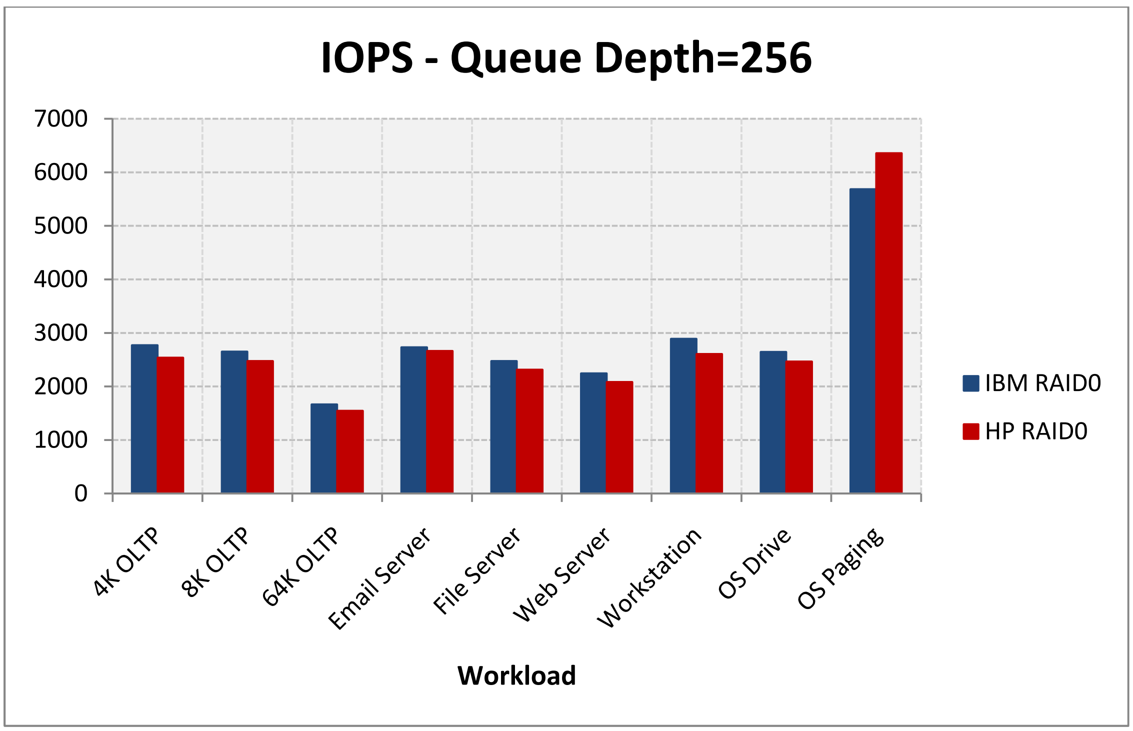

This graph shows the media reader and writer results for twelve (12) simultaneous threads, the midpoint of the media I/O workloads tested. The other RAID types showed similar patterns, but only the RAID0 data is included here to keep the graph relatively simple. This data shows a subset of data from the other graphs previously shown in this report.

The IBM controller produces noticeably better performance for these types of parallel streaming, sequential I/O.

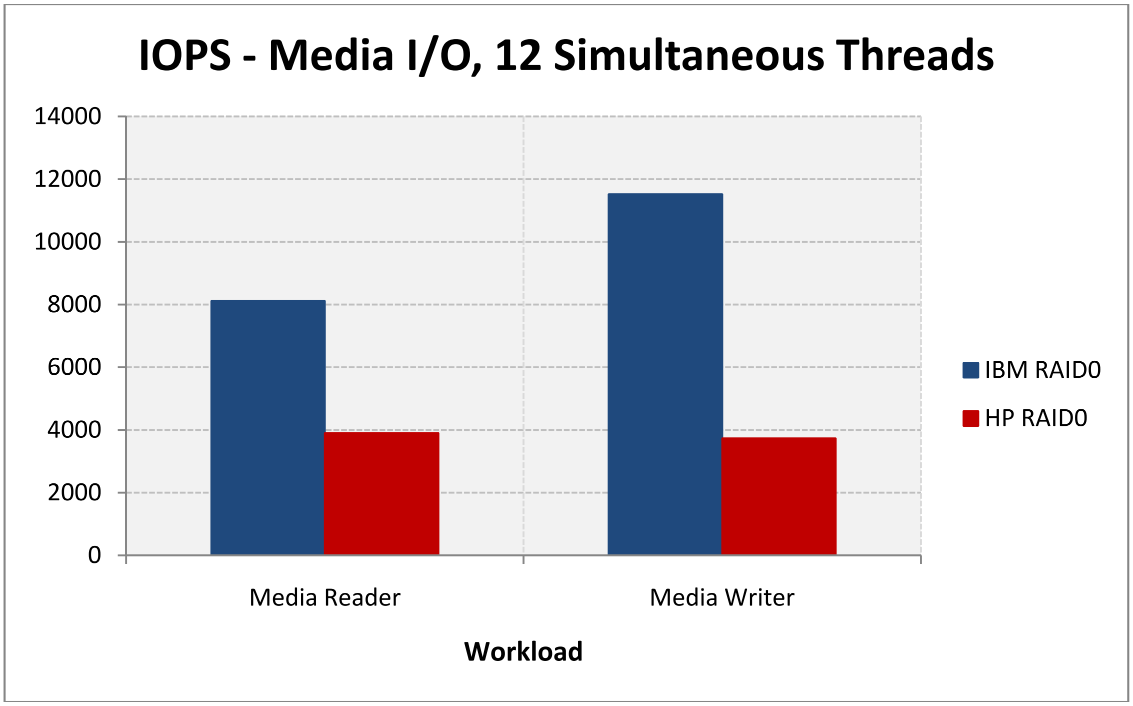

Conclusion

Our test results show that under heavy workload conditions, the IBM ServeRAID M5015 SAS/SATA Controller provides superior performance to the HP Smart Array P410i RAID controller. These tests were conducted using eight, 10,000 RPM SAS disk drives that have a 6-Gbps SAS interface. We believe that if 15,000 RPM disk drives had been used, the performance would have been higher for all tests, but the general pattern of results would have approximately the same.

These tests were run with eight disk drives in the array. We can only speculate on the results with a significantly larger number of disk drives in the array, but based on what we have seen, we believe that the results would be similar.

These tests show that although both RAID controllers provide similar performance on lighter workloads, the IBM controller is able to sustain good performance on the heaviest workloads.