HP StorageWorks 4400 Enterprise Virtual Array (EVA4400) performance white paper

Introduction

The HP StorageWorks Enterprise Virtual Array (EVA) product family offers midsize to enterprise size customers leading performance, high capacity, and high-availability storage solutions for reducing IT costs and complexity. The EVA provides virtualized storage enabling capacity pooling, simplified management, automatic performance load balancing, dynamic configuration, and re-configuration.

Continuing to build on this success is the HP StorageWorks 4400 Enterprise Virtual Array (EVA4400). Engineering advancements have enabled substantial real-world performance improvements over the HP StorageWorks 4100 Enterprise Virtual Array (EVA4100).

The EVA4400's increased levels of sequential performance provide customers with improved data transfer rates for applications such as data warehousing, streaming video, high-performance technical computing, and backups and restores.

In random workloads, the EVA4400 performance scales up to the maximum number of disk drives supported. Customers will have the predictability they need to choose configurations that best match actual business conditions. Increased random workload performance translates into better throughputs and response times for applications such as files systems, transaction-orientated databases, and email.

This white paper provides a high-level overview of the:

- Importance and differences related to cache and end-to-end performance numbers
- Performance summary associated with the EVA4400

Sizing up performance numbers

The HP StorageWorks research and development (R&D) philosophy embodies two important tenets:

- Create products that deliver value and quality to the customer. (This tenet guides HP decision-making for deploying cutting-edge technology.)
- Build trust by using testing procedures that, within normal engineering constraints, customers can verify and repeat at their place of business. (Many HP customers refer to HP as their "trusted advisor," and have "bet their business" on HP technology. HP enables customers to substantiate any claims by using test procedures that customers can repeat and verify at their site.)

When reporting performance numbers for an array, HP has found that the most useful numbers for customers are those associated with workload configuration and other characteristics that reflect actual business conditions.

For the purposes of this paper, HP refers to the performance numbers that exercise both the normal data path and the normal code path as "end-to-end" performance numbers. A normal data path includes the array's Fibre Channel connections to servers, cache, internal buses, and disk drives. The normal code path includes software or firmware that moves user data between its disk drives and the servers that connect to the array. "Cache" performance numbers are those obtained from 100% cache-only accesses.

End-to-end performance numbers

End-to-end performance numbers reflect typical data and code paths used in a customer situation. These numbers are an excellent choice for capacity planning when obtained under two additional engineering constraints.

First, when evaluating performance under random workloads, HP recommends using uniform probability to access most, if not all, usable array capacity. This measurement is more accurate than "short-stroking (destroking)," a configuration that allows the disk to access only a small fraction of the drives' capacity, thus enhancing performance numbers. Measuring performance by accessing all of an array's capacity leads to sizing estimates that will not become obsolete over time as the customer consumes more of the array's capacity.

The second constraint applies to the recognition from HP that response times are very important to customers and their business applications. HP publishes "end-to-end" throughputs at average response times of 30 milliseconds or less. Because cache numbers are not useful in this context, HP does not apply this constraint to cache-only measurements.

HP best practices recommend using end-to-end performance numbers for capacity planning because they are:

- Measurements that can be used for sizing in a real-life business condition
- Easy to verify at the customer site using a reasonable "black box approach" by eliminating bottlenecks outside the array and balancing the load across an array's controllers and host ports

End-to-end performance numbers provide customers with the best data to determine which array configuration makes the best business, technical, and financial sense.

Cache performance numbers

Cache performance numbers register the largest throughput numbers possible with an array but do not reflect use under normal business conditions, diluting their usefulness for capacity planning.

Such cache-only workloads place a lighter load on an array's internal buses than would exist at a customer site. Exercising only a part of an array's data and code paths does not reflect actual response times or latencies that normally occur because the cache-only data path is shorter than the normal data path. Typically, these are the paths used when 100% of all I/O requests are satisfied from the array's cache. Cache performance numbers are unsustainable in day-to-day operations, though exceptions will occur momentarily and show up as short-lived transients.

Cache serves as an intermediary between the servers and the disk drives. When viewed separately, cache performance numbers provide an artificial view of an array's performance capabilities. What is most important about cache performance is the interaction of the cache with a disk array's disk drives.

Because disk overload is a possible consequence of using cache performance numbers for capacity planning, HP best practices do not recommend this measurement.

EVA4400 performance summary

The HP EVA4400 offers real-world performance at a higher level than the EVA4100 and is more affordable. With the EVA4400, HP engineers aimed to lower the customer's total cost of ownership (TCO) while delivering more capacity and better performance—all within the context of the existing EVA enterprise-class functionality.

The vast majority of customers use EVA arrays for two classes of workloads: large block sequential and small block random. The EVA4400 achieves improved sequential performance levels and at a reduced cost and complexity. With respect to random workloads, the performance of the EVA4400 scales with the number of its disk drives, allowing customers to choose configurations that best meet their needs.

The elevated levels of sequential performance provides customers improved data transfer rates for applications such as data warehousing, streaming video, high-performance technical computing, and backups and restores. Increased random workload performance translates into better throughputs and response times for applications such as filesystems, transaction-orientated databases, and email.

The following two tables display the difference between the EVA4400 and its predecessor, the EVA4100. HP measured each array with the maximum number of drives supported (56 drives on the EVA4100 and 96 drives on the EVA4400). The random workload performance of the EVA4400 is approximately twice that because of the increased number of disk drives coupled with increased controller capability. The gain in internal controller bandwidth of the EVA4400 accounts for the marked improvement under the large block sequential workloads.

EVA4400 and EVA4100 differentiation

Table 1. End-to-end performance @ <= 30 ms average response time

Workload	VRAID level	EVA4100 (56 drives)	EVA4400 (96 drives)
Random Reads IOPs	Any	13,800	26,000
Random Writes IOPS	RAID 1	7,300	12,000
Random Writes IOPs	RAID 5	4,000	7,200
Sequential Reads MB/s	Any	350	775
RAID 1 Sequential Writes MB/s	RAID 1	168	375
RAID 5 Sequential Writes MB/s	RAID 5	264	550

Table 2. Cache performance*

Workload	EVA4100	EVA4400
Small Reads (IOPs)	> 154,000	> 140,000
Large Reads (MB/s)	1,385	1550

*Cache performance numbers register the largest throughput numbers possible with an array but do not reflect use under normal business conditions, diluting their usefulness for capacity planning.

Experimental setup description

The following information provides a high-level overview of the measurement criteria, equipment, and benchmarking procedures used in the testing of the EVA4400. The equipment chosen ensured that the only potential bottleneck to performance was the EVA tested. The switches used in the test for convenience did not affect any measurements.

Measurement criteria

Server:

- HP rp8400 server
- Eight PA8700 750-MHz CPUs
- HP-UX 11.11 equipped with 32 GB of main memory
- Four AB378A Fibre Channel host bus adapters

Tested array:

- EVA4400 HSV300—96, 146.5-GB disk drives
- Drive rotational frequency—15K RPM
- 1 GB of main memory per controller
- HP StorageWorks Brocade 4Gb SAN Switch with 32 ports
- One-to-one mapping between the AB378A and LE HSV300 front-end ports

Benchmarking procedures:

- RAID 1—16, 381818-MB Virtual Disk
- RAID 5 and 16, 610910-MB Virtual Disk
- Wrote to all of the capacity before making the measurements
- Load balanced across host port when using four or more I/O processes and across controls when using two processes

Measurement results

Figures 1 through 6 show the relationships between the average response time and the throughput delivered by the EVA4400 for several workloads and RAID levels. Customers can use these curves to consider the tradeoff between increased throughput and the resulting increased average response time.

In Figures 1 through 4, the 60% read workload is what HP calls the OLTP workload.

Figure 1 shows curves for random workloads directed at RAID 1 storage on the EVA4400 for three workloads:

- Read requests
- Write requests
- 60% reads and 40% writes (labeled as OLTP)

All workloads used 4-KB request lengths. HP Server benchmark teams used the 60% read workload to predict the servers and storage performance under the TPC-C benchmark.

Figure 2 is similar to Figure 1 except that all workloads used a uniform request size of 8 KB rather than 4 KB.

Figures 3 and 4 are similar to Figures 1 and 2, respectively, except that RAID 5 storage replaced RAID 1 storage.

Figures 5 and 6 show the relationship between average throughputs and response times for large block sequential workloads. Figure 5 uses RAID 1 storage, and Figure 6 uses RAID 5 storage.

In Figures 5 and 6, the cache algorithms used by the EVA4400 results in very high data transfer rates at very low response times. This illustrates the importance of a disk array's cache and the algorithms used for managing it. Although we are seeing the effects of cache in the measurements shown, these effects are included as part of the end-to-end performance measurements.

Figure 1.

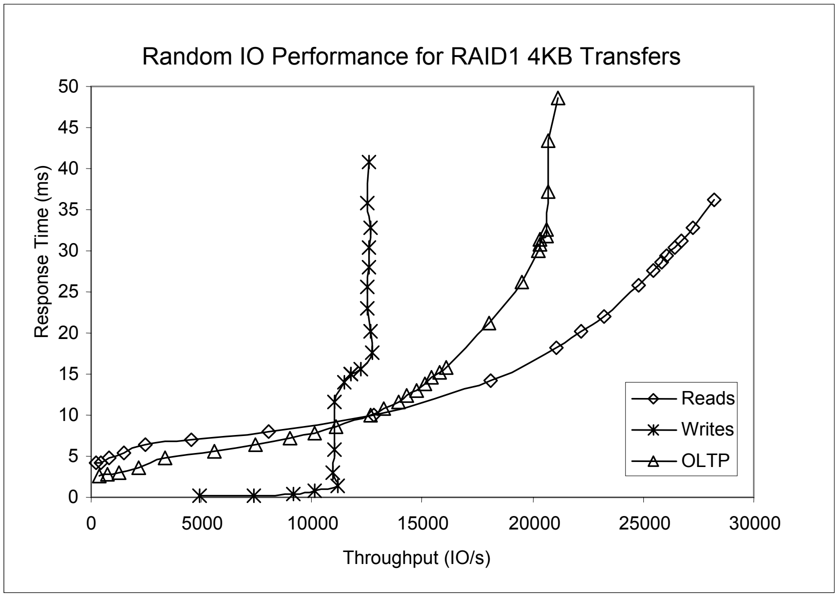

Figure 2.

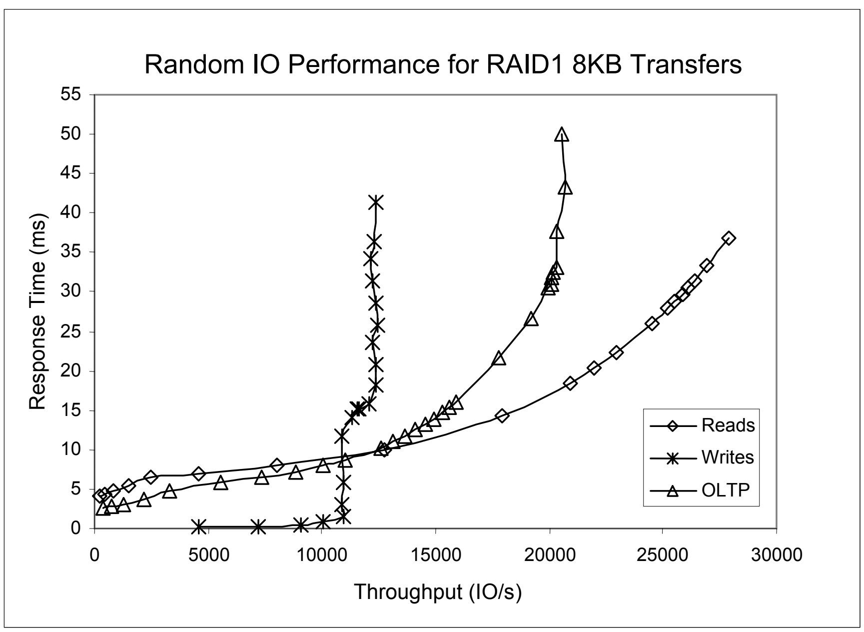

Figure 3.

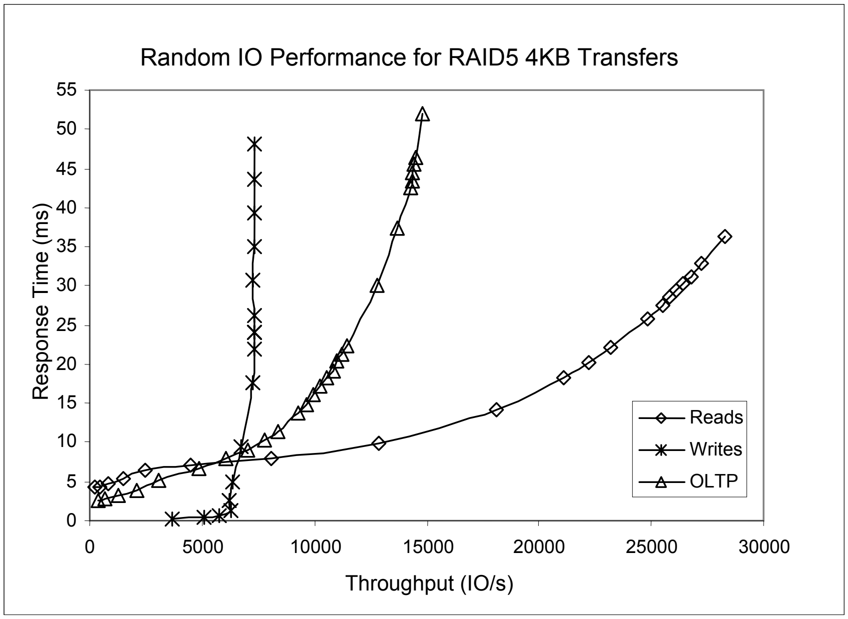

Figure 4.

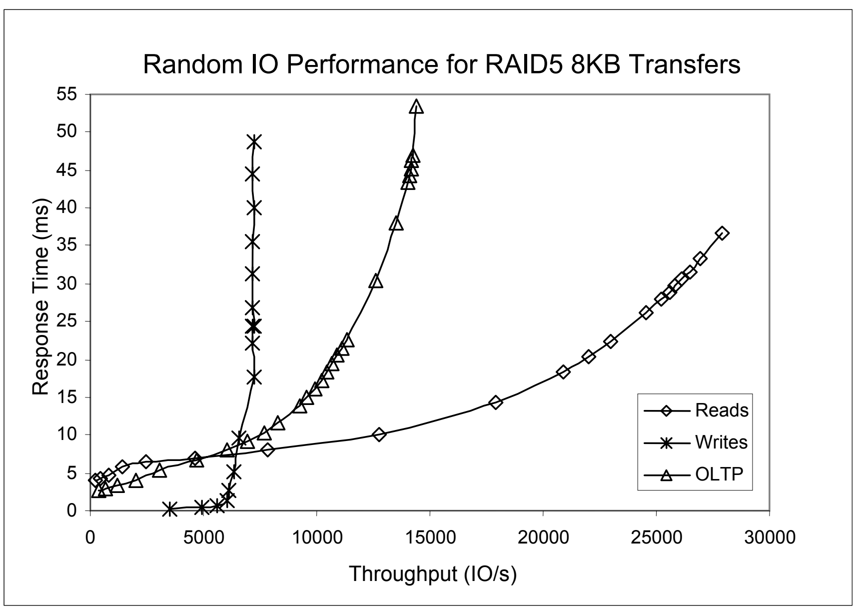

Figure 5.

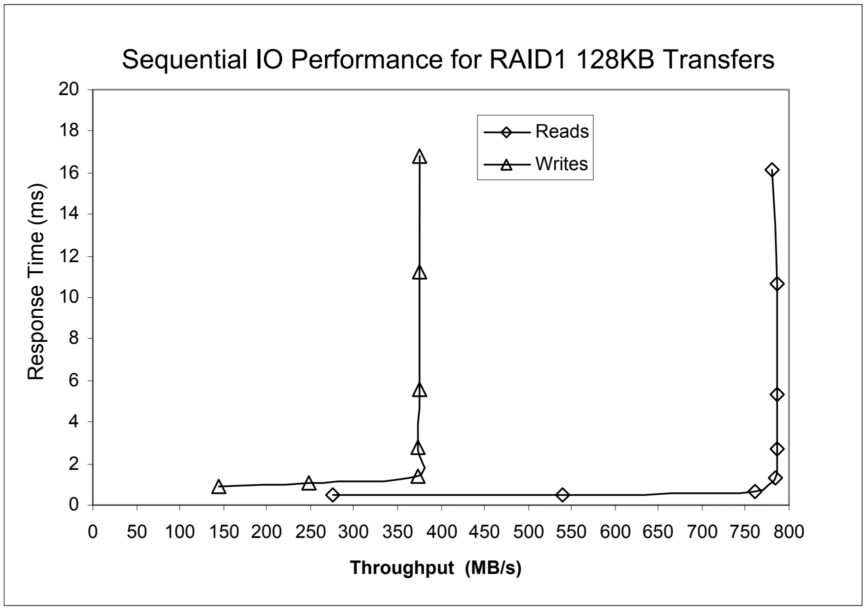

Figure 6.

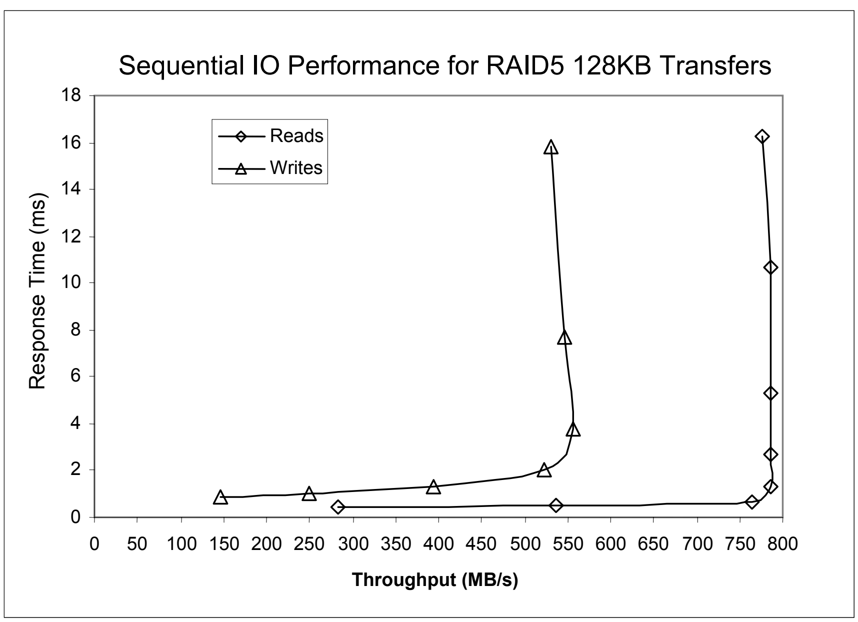

Conclusion

The EVA product family provides virtual storage that enables capacity pooling, simplified management, automated performance load balancing, dynamic configuration, and re-configuration. The new HP EVA4400 offers customers all those features in an affordable array with substantially improved performance in random and/or sequential workloads.

In random workloads, the EVA4400 performance scales up to the maximum number of disk drives supported. Customers will have the predictability they need to choose configurations that best match actual business conditions. The EVA4400 achieves improves sequential performance levels at a reduced cost and complexity.

By delivering twice the throughput as its predecessor (EVA4100), the EVA4400 doubles the transfer rates for a myriad of applications ranging from backups and restores to streaming video. Data scans can take half the time.

HP best practices recommend that customers review the total system performance characteristics of the HP EVA4400 for the most accurate data. End-to-end performance numbers provide the most stable and accurate representation of the customer's environment. Cache performance numbers register the largest throughput numbers possible with an array but do not take into account actual business conditions, diluting their usefulness for capacity planning.

For more information

- HP StorageWorks 4400 Enterprise Virtual Array (EVA4400)
  www.hp.com/go/eva4400
- EVA Product Family
  www.hp.com/go/eva
- EVA cache white paper
  http://h71028.www7.hp.com/ERC/downloads/4AA1-7945ENW.pdf
- HP StorageWorks
  www.hp.com/go/storage

4AA1-8473ENW, March 2008