HP P6000 Enterprise Virtual Array performance

Technical white paper

Table of contents

Introduction

The HP P6000 Enterprise Virtual Array (EVA) product family offers high-performance, high-capacity, and high-availability storage solutions to midsize and enterprise-size customers to help them reduce their IT costs and complexity. The EVA family provides virtualized storage, enabling capacity pooling, simplified management, automatic performance load balancing, dynamic configuration, and reconfiguration. Continuing to build on this success are the HP P6300 EVA and P6500 EVA models.

This paper focuses on the read and write performance of the P6000 EVAs. They provide excellent performance under transaction-oriented workloads and sequential workloads. One purpose of this paper is to describe P6000 EVA performance quantitatively by showing the results of direct measurements made on the array using different RAID types and workloads.

To properly understand performance numbers the context must also be understood. Numbers without context may cause misunderstandings about achievable storage system performance in a production environment. Such misunderstandings can be expensive in terms of both equipment and human resource costs. So, the second purpose of this paper is to provide some context surrounding the performance numbers.

This white paper includes:

- A high-level discussion about performance numbers in general, including a discussion about the usefulness (or lack thereof) for capacity planning purposes of cache-only performance numbers versus end-to-end performance
- A summary of performance results for the P6300/P6500 EVAs

Sizing up performance numbers

The HP storage 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 new 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 they can repeat and verify at their site.

When reporting performance numbers for an array, HP has found that the numbers most relevant to real applications are those associated with workload, configuration, and other characteristics that tie directly to actual business conditions. Obtaining such numbers is an extremely difficult task because the performance of any disk array depends strongly on the following conditions:

1. The workload to which the array is subjected
2. The configuration of the array
3. The configuration of the SAN to which the array is connected

There are at least as many different workloads as there are users of storage, no single set of performance numbers can be used to describe what every user sees when an array is deployed in their environments. The configuration of the array is important, for example, because different RAID levels result in different performance characteristics for the same workload. Configuration of the SAN plays a role because some configurations create performance bottlenecks outside the array.

Despite all this complexity, experience garnered while working with a large base of storage customers has allowed HP performance engineers to identify a small set of workloads and configurations that expose some of the most important performance characteristics of a disk array.

These workloads and configurations are used to derive all of the performance numbers shown in this paper.

To make sure that the test results are relevant, HP disk array performance engineers design the test SAN configurations so that there are no host or infrastructure performance limitations that would affect the disk array performance measurements. This process verifies that the properties of only the array—and not something else, such as a server, a switch, or some host bus adapter—are determined. The performance measurements shown in the section entitled Performance summary were obtained with a test setup that satisfies this criterion.

Another important aspect of the measurement results shown in this paper is that the average throughputs and average response times reported correspond to steady-state behavior. Transient effects, such as those caused by empty write caches, which skews throughput numbers up and response times down dramatically, are not a part of any of the measurements. The steady-state condition implies that all measurement results are indicative of sustained performance, even regarding performance associated with 100% cache accesses.

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 connections to servers, cache, internal buses, and disk drives. The normal code path includes software or firmware running on an array’s controllers that moves user data between its disk drives and the servers that connect to the array.

End-to-end performance numbers

End-to-end performance numbers reflect typical data and code paths exercised in an array when that array is deployed in a customer environment. 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 accessing most, if not all, of the usable array capacity with a uniform access probability. Such measurements are more accurate than “short-stroking (de-stroking),” a configuration that allows each disk to access only a small fraction of its capacity, thus inflating throughputs artificially. 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 that it is meaningless to talk about array throughputs without specifying the average response time at which the throughput is measured. Response times are very important, sometimes even critical, to customers and their business applications. For this reason, HP publishes “end-to-end” throughputs at average response times of less than or equal to 30 milliseconds. The maximum average response time of 30 minutes is high for most customers under transaction-based workloads (15 minutes is often quoted as being high for many such applications), but it is used in case such high response times are acceptable to some users. Because cache numbers are not useful from a customer-deployment point of view, HP does not apply this response-time constraint to cache-only throughputs.

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

- Suitable for sizing in a real-life business condition
- Easy to verify at the customer site using a reasonable “black box approach” by removing 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 configurations make the best business, technical, and financial sense.

Cache performance numbers

Cache performance numbers register the largest throughputs possible with an array when no disk accesses occur. Therefore, they do not reflect the performance of a disk array under normal business conditions, rendering them useless for ascertaining whether a given array will meet a customer's performance requirements. Thus they are not appropriate for purposes of capacity planning.

Such cache-only workloads place a lighter load on an array's internal buses than would exist at a customer site. Because the cache-only data path is shorter than the normal data path, exercising only a part of an array's data and code paths does not reflect actual response times or latencies that normally occur. Cache-only data paths are used when 100% of all I/O requests are satisfied from the array's cache, with no disk accesses occurring. Cache performance numbers are unsustainable in day-to-day operations, although exceptions 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 highly likely consequence of using cache performance numbers for capacity planning, HP best practices do not recommend using such numbers.

Performance summary

In this section, a summary of the performance of the HP P6000 EVAs are provided. The summary is based on actual measurements made for a number of select workloads and RAID levels. Since the vast majority of EVA customers use EVAs to service large-block sequential workloads or small-block random workloads, we focus on the performance of the P6000 EVAs under these workloads.

The large-block sequential performance information is most relevant to customers who care about data transfer rates (MB/s), which is often the case when their applications perform data warehousing, streaming video, high-performance technical computing, backups, or restore functions. The small-block performance information is most relevant to customers who care about throughputs (IO/s or IOPS), which is often the case when their applications are file systems, transaction-orientated databases, or email applications.

Figures 1 through 12 show the relationship between the average response time and the throughput delivered by the P6000 EVAs for several workloads and RAID levels. These curves are useful when considering the tradeoff between increased throughput and the resulting increased average response time.

In Figures 1 through 6, the curve labeled OLTP is what HP uses to model an On Line Transaction Processing workload. It is composed of 60% reads and 40% writes.

Figures 1 and 2 show curves for random workloads directed at RAID 1 storage on the P6300 EVA and P6500 for three workloads:

- Read requests
- Write requests
- 60% reads and 40% writes (labeled as OLTP 60-40)
- 70% reads and 30% writes (labeled as OLTP 70-30)

All random workloads used 8 KB request lengths.

Figure 1: P6300 EVA—Random IO performance for Vraid1 8 KB transfers

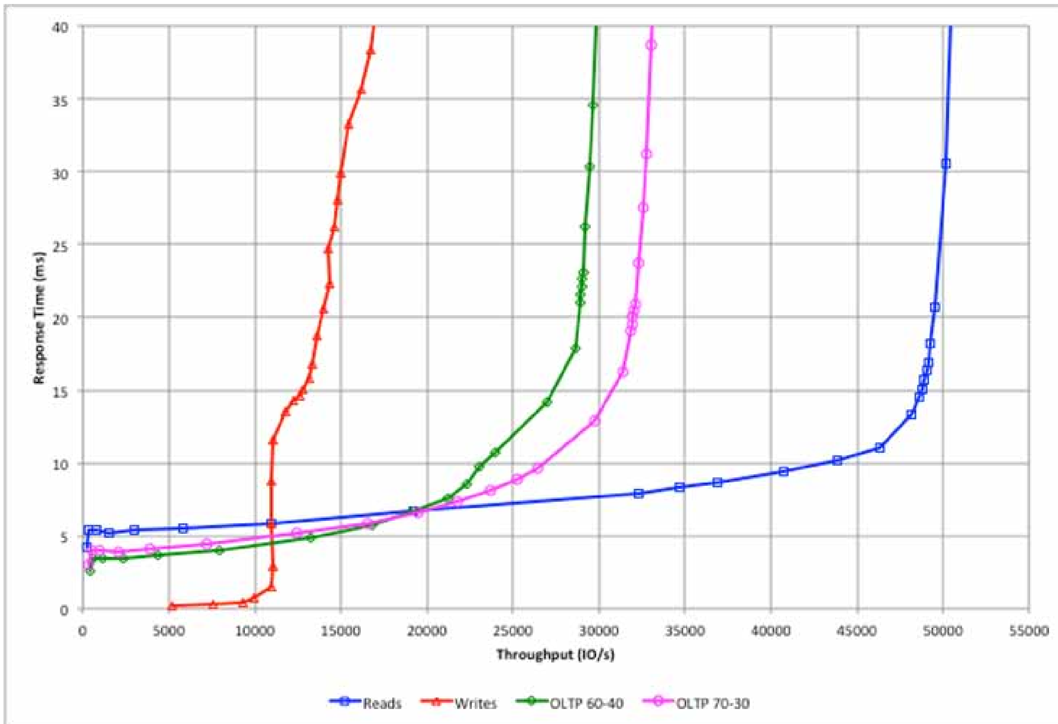

Figure 2: P6500 EVA—Random IO performance for Vraid1 8 KB transfers

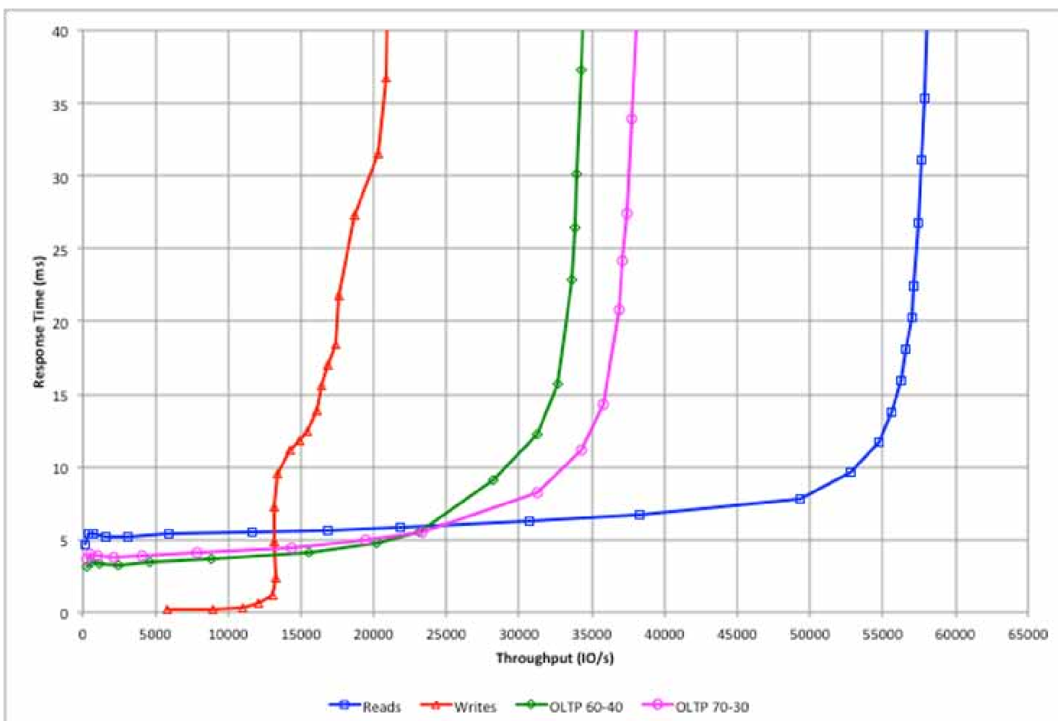

Figure 3: P6300 EVA—Random IO performance for Vraid5 8 KB transfers

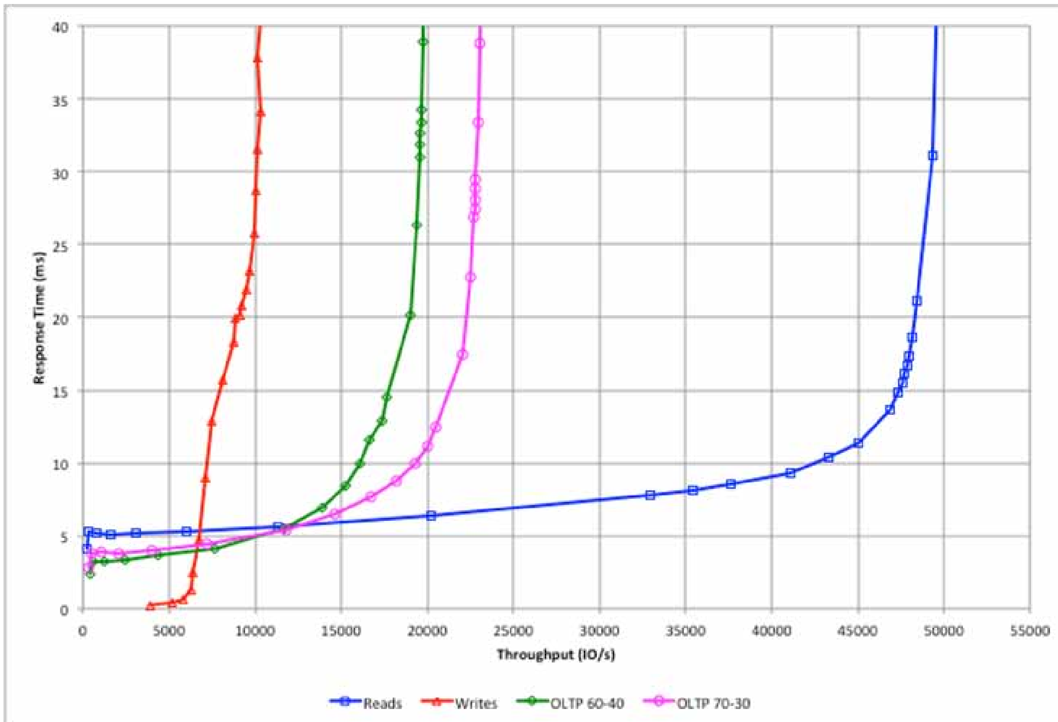

Figure 4: P6500 EVA—Random IO performance for Vraid5 8 KB transfers

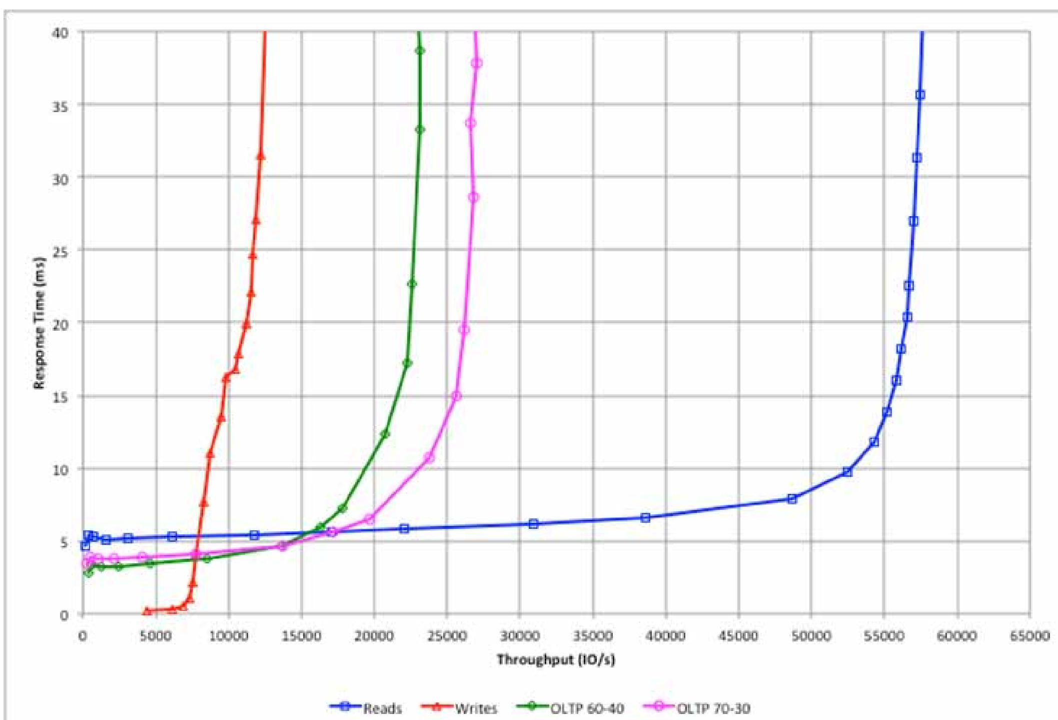

Figure 5: P6300 EVA—Random IO performance for Vraid6 8 KB transfers

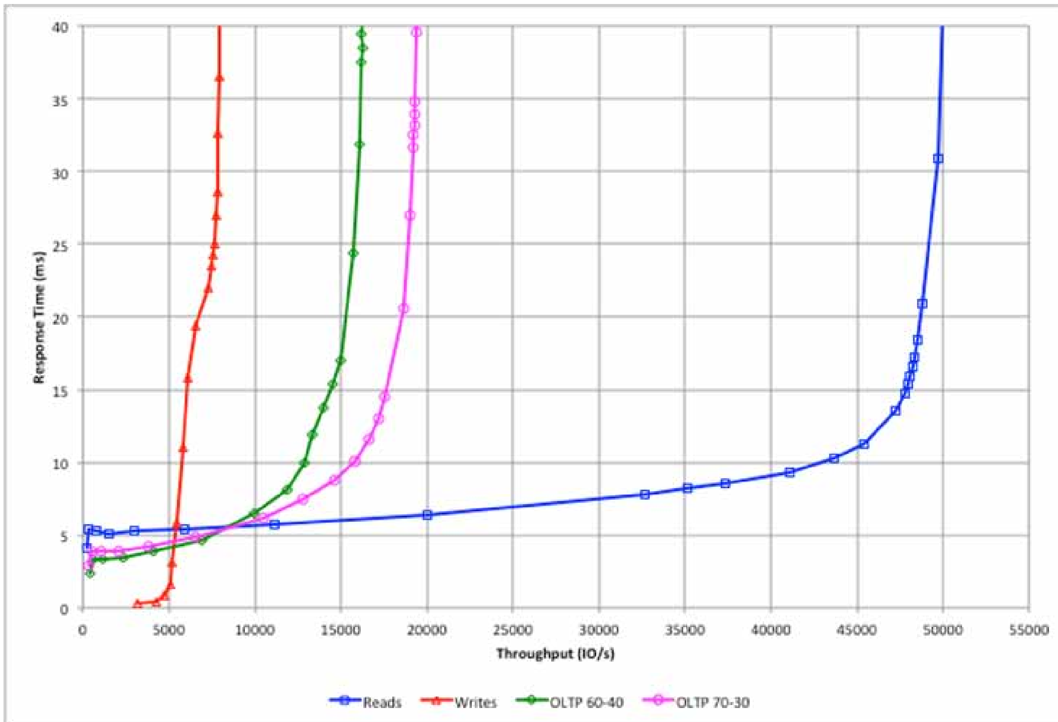

Figure 6: P6500 EVA—Random IO performance for Vraid6 8 KB transfers

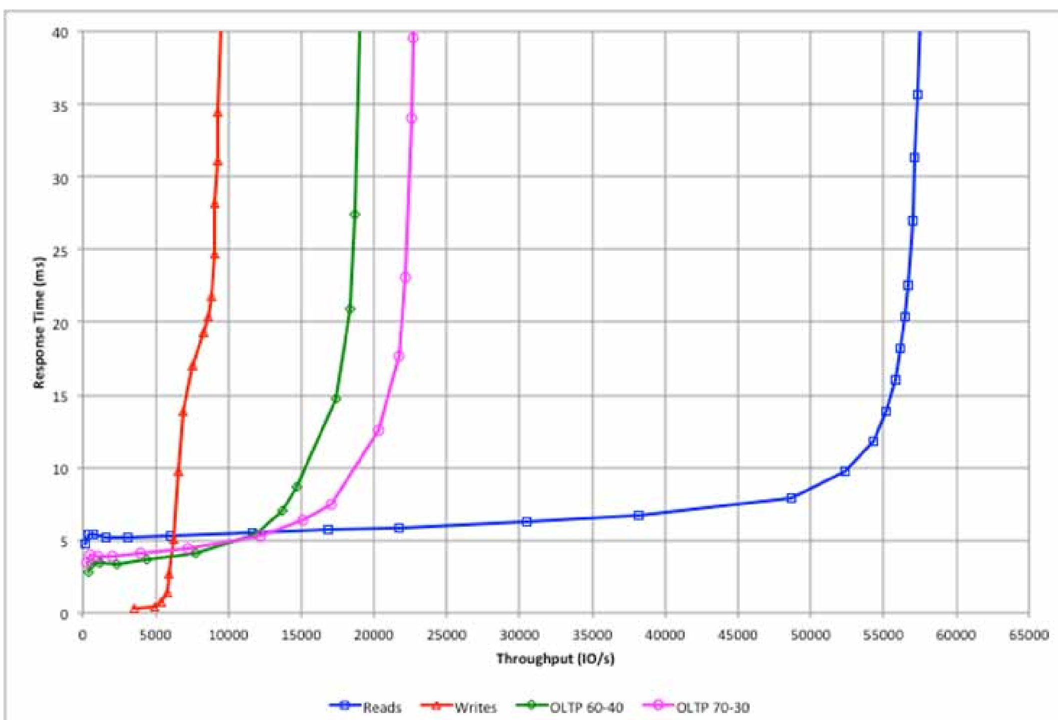

Figure 7: P6300—Sequential IO performance for Vraid1 128 KB transfers

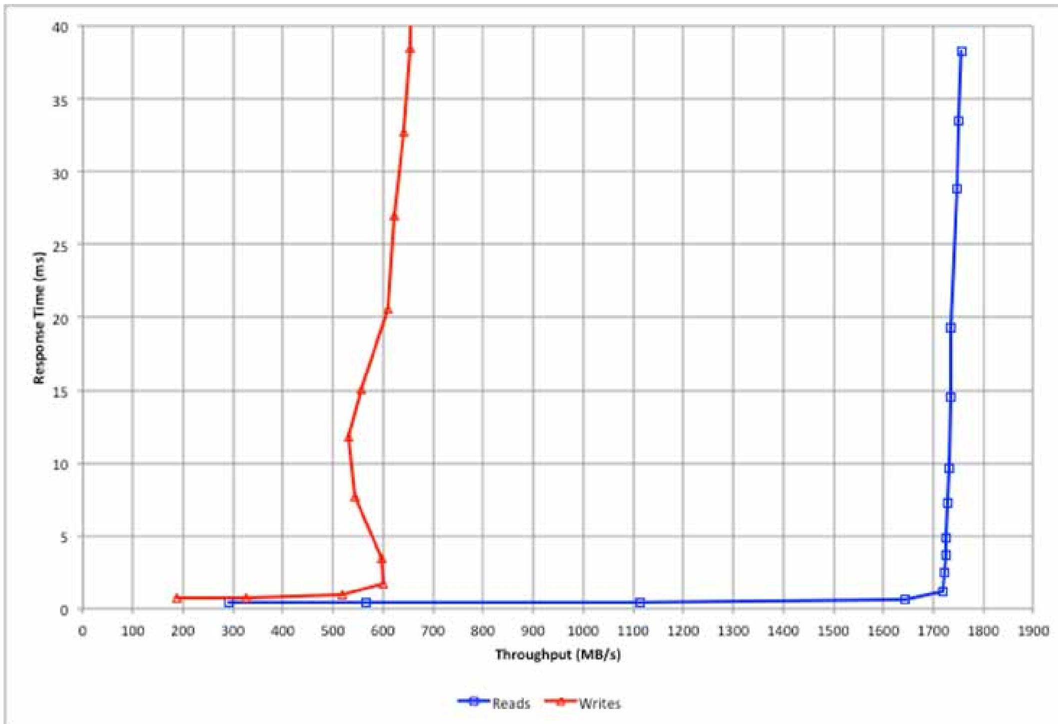

Figure 8: P6500—Sequential IO performance for Vraid1 128 KB transfers

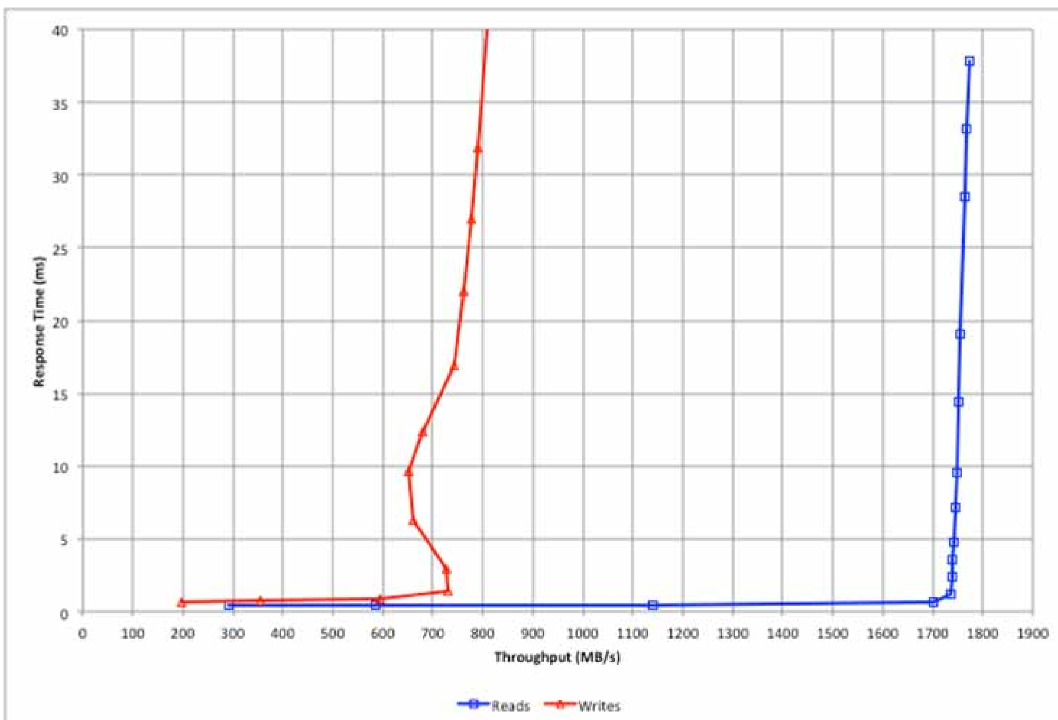

Figure 9: P6300—Sequential IO performance for Vraid5 128 KB transfers

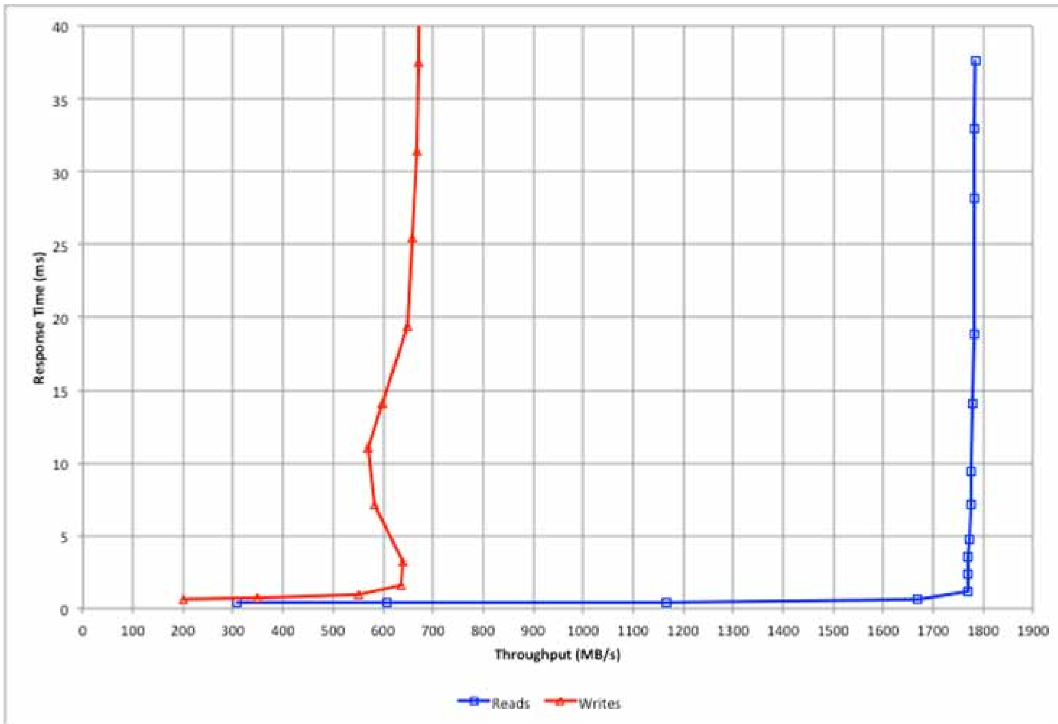

Figure 10: P6500—Sequential IO performance for Vraid5 128 KB transfers

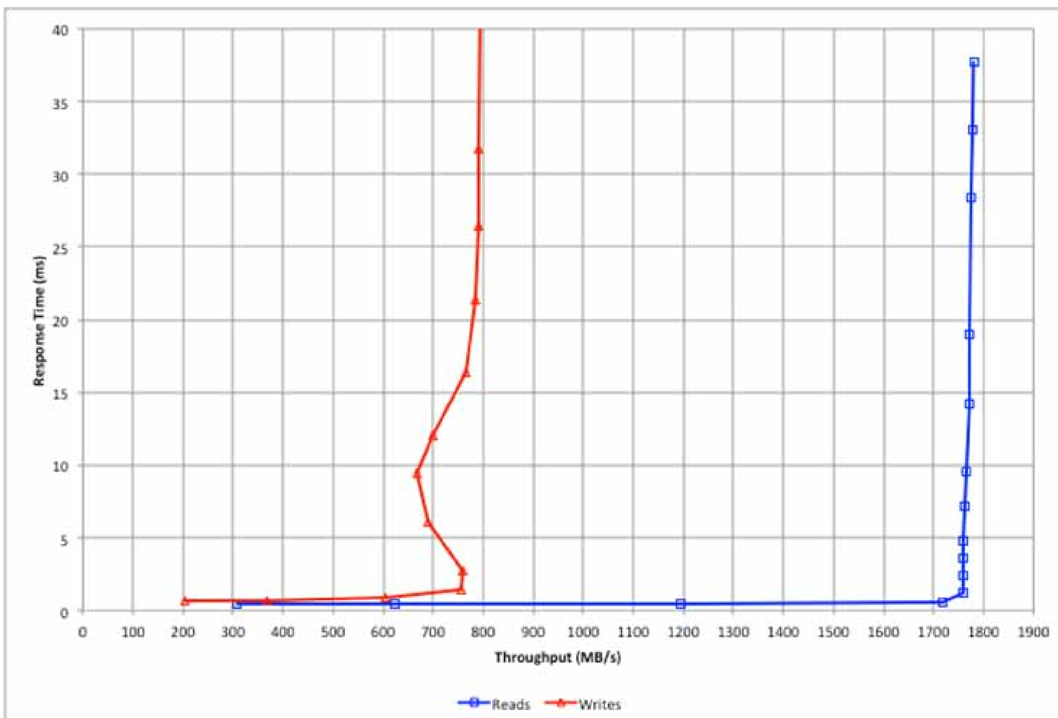

Figure 11: P6300—Sequential IO performance for Vraid6 128 KB transfers

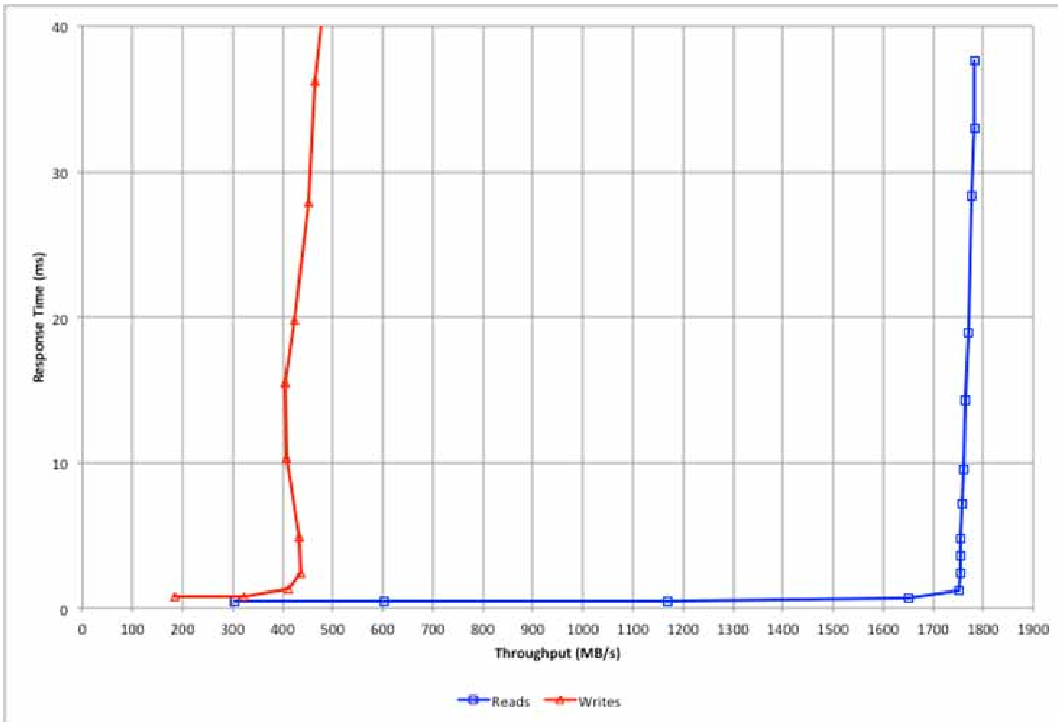

Figure 12: P6500—Sequential IO performance for Vraid6 128 KB transfers

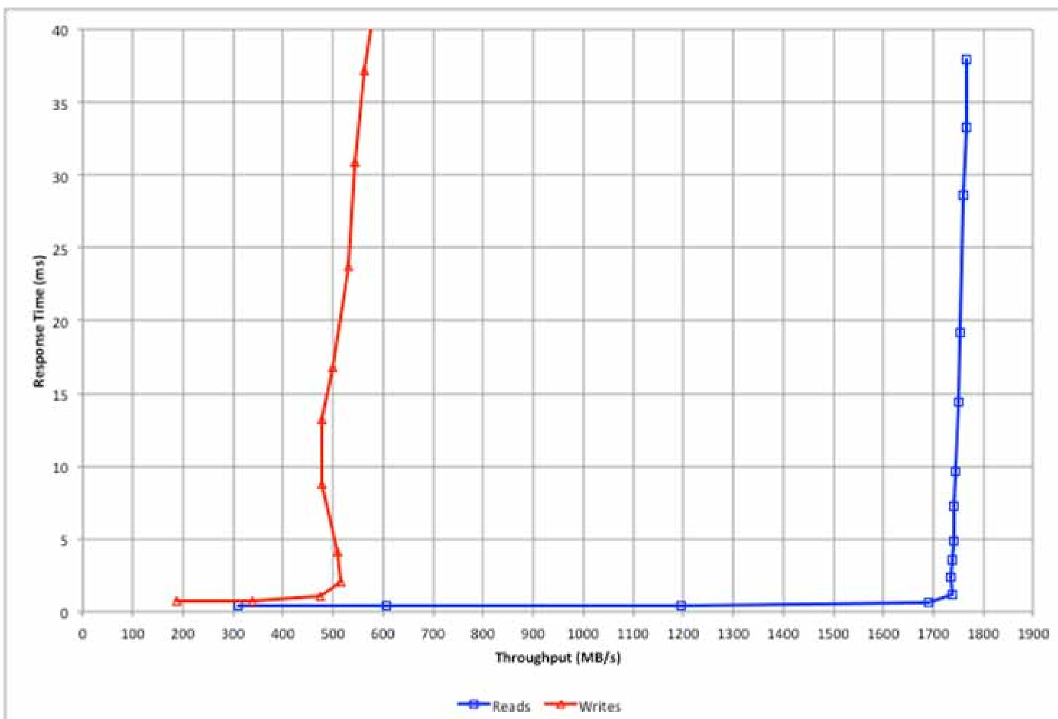

The data presented above reflects end-to-end performance of the P6000 EVAs. For completeness, we include the cache-only throughputs for small-block and large-block workloads as shown in Table 1. Note that no RAID level is specified since there are no disk accesses done when these numbers are obtained.

Table 1: Cache performance*

Workload	P6300 EVA	P6500 EVA
Small Reads (IOPS)	>140,000	>180,000
Large Reads (MB/s)	1780	1780

* Cache performance numbers register the largest throughput numbers possible with an array, but have not relation to what will be observed by a customer when the array is deployed under a customer's application.

The following information provides a high level overview of the measurement criteria, equipment, and benchmarking procedures used in the testing of the P6000 EVAs. The equipment was selected to make sure that the only potential performance bottleneck was the EVA itself. The switches used for convenience in the test did not affect any measurements.

Server:

- HP Integrity rx8640 Server
- 4 Dual-Core Intel® Itanium® 2 9000 series (1.6 GHz, 18 MB) processors
- HP-UX 11.31 equipped with 16 GB of main memory
- Four 4 Gb Fibre Channel Host Bus Adapters (AB379A)

Tested arrays:

- HP P6300 HSV340 with 250 146 GB disk drives
- Drive rotational frequency—15k rpm
- 4 GB of cache memory per controller pair
- HP StorageWorks B-series 4 Gb SAN switch with 32 ports
- One-to-one mapping between the Fibre Channel HBAs and P6300 HSV340 front-end ports
- HP P6500 HSV360 with 450 146 GB disk drives
- Drive rotational frequency—15k rpm
- 8 GB of cache memory per controller pair
- HP StorageWorks B-series 4 Gb SAN switch with 32 ports
- One-to-one mapping between the Fibre Channel HBAs and P6300 HSV340 front-end ports

Benchmarking procedures:

- P6300 HSV340
  – Vraid1: 8, 2040117 MB virtual disks
  – Vraid5: 8, 3264187 MB virtual disks
  – Vraid6: 8, 2720156 MB virtual disks
- P6500 HSV360
  – Vraid1: 8, 3699577 MB virtual disks
  – Vraid5: 8, 5919323 MB virtual disks
  – Vraid6: 8, 4932769 MB virtual disks
- 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 controllers when using two processes

Conclusion

The HP P6000 EVA product family provides virtual storage that enables capacity pooling, simplified management, automated performance load balancing, dynamic configuration, and reconfiguration. The HP P6300/P6500 EVAs offer customers all those features in an affordable array with excellent performance under random and sequential workloads.

HP best practices recommend reviewing the total system performance characteristics of the HP P6300 for the most accurate data. End-to-end performance numbers provide the most stable and accurate representation of the storage environment. Cache performance numbers register the largest throughput numbers possible with an array but do not take into account actual business conditions, rendering them useless for capacity planning purposes.

To learn more about HP P6000 Enterprise Virtual Array, please visit www.hp.com/go/P6000. To get more information on HP StorageWorks Enterprise Virtual Arrays, visit www.hp.com/go/eva and on HP StorageWorks products and solutions visit www.hp.com/go/storage.

4AA3-3832ENW, Created June 2011