The Impact of using All-Flash Arrays on Mission Critical Application Performance

December, 2013

Neuralytix™

Neuralytix™

White Paper

Sponsored by Pure Storage

The Impact of using All-Flash Arrays on Mission Critical Application Performance

Benjamin S. Woo

December, 2013

Overview

Applications that leverage structured data – the data that is typically stored in relational database management systems (RDBMS), also known as enterprise applications – are still the applications first seen by customers, and the applications with which most customers interact. Therefore, optimizing these applications is paramount.

One of the ways enterprises can optimize these applications is through performance. Increasing the number of transactions, or increasing the speed with which transactions complete translates directly to improved revenues and profits.

In this White Paper, Neuralytix examines the use of an all-flash storage system to optimize RDBMS operations while minimizing the impact to capital and operational expenses.

info@neuralytix.com | www.neuralytix.com | @Neuralytix
1129 Lexington Avenue, #3 | New York, NY 10075 | (212) 724.1188

Table of Contents

info@neuralytix.com | www.neuralytix.com | @Neuralytix

1129 Lexington Avenue, #3 | New York, NY 10075 | (212) 724.1188

Introduction

Enterprise applications – including Microsoft Exchange and SQL Server, Oracle, and SAP – represent an enterprise's most mission critical applications. It enables the enterprise to transact business and provide critical functions within the enterprise. Therefore, it naturally flows that ensuring optimal performance to these applications is paramount.

Whereas recently, there has been a lot of attention given to unstructured data, archive and *dark data*, enterprise applications still require attention to optimize performance in order to provide the enterprise with front-line competitive advantage.

Neuralytix research conclusively shows that roughly 5% of a given enterprise's data actively changes throughout any given time period. Large enterprises will want to focus on optimizing this 5%, while smaller and medium enterprises may opt to simply optimize the entire database environment. This can help enterprises of all sizes to minimize capital investments in storage, while significantly improving the performance of their enterprise applications.

Over the years, there have been many attempts to improve enterprise application performance. Most notably:

- Increase main memory (RAM); or
- Only writing on the outer tracks of high performance, high rotational speed, traditional, magnetic hard disk drives (HDDs), also known as "short stroking".

The problem with these approaches has primarily been cost.

Increase Main Memory (RAM)

Ultimately, DRAM is expensive. In some cases, costing 100x per unit of storage compared to HDD capacity. This makes increasing main memory in large database environments cost prohibitive.

Additionally, in very large environments, the server running the application may not support the amount of memory necessary to put the entire database into RAM. This in turn

info@neuralytix.com | www.neuralytix.com | @Neuralytix
1129 Lexington Avenue, #3 | New York, NY 10075 | (212) 724.1188

results in paging, bringing the problem back to an I/O constraint on the storage subsystem (ie. a disk issue).

RAM also has another problem. If used in a write-back mode, where data destined for another persistent storage device, a discontinuation of power could result in data becoming inconsistent or data corruption. Therefore, RAM is only useful for reads.

Only writing on the outer tracks of high performance HDDs (aka "short stroking")

Modern enterprise applications have easily identifiable "hot files." These files usually represent the indices, log files, and caching components of enterprise applications.

However, irrespective of whether one just focuses on the "hot files" or the entire database, in order to double performance, the enterprise still needs to double the number of disk drives.

High rotational speed HDDs provide 33-50% higher performance. To achieve the highest performance gain, storage system vendors can limit where data is written. By only writing data on the outer tracks of these disk drives, which translates to less movement by the read arms inside the mechanical HDD, the maximum performance, from the highest cost disk drives, is observed.

The problem here is also one of cost. To double the performance of an environment, an enterprise would have to change the mix of disk drives. Rather than 100% of slow disk drives, it would require roughly 58% of slow disk drives, and 42% of fast, higher cost disk drives.

Table 1 shows this effect:

	# of 7,200-RPM drives	# of 15,000-RPM drives	IOPS
Original Environment	100	0	10,000
Flash Environment	95	70	20,000

Assumptions: 7,200-RPM disk drives yield 100 IOPS, while 15,000-RPM disk drives yield 150 IOPS
Notes: This table is for relative performance comparisons only, and does not represent actual performance

Table 1: IOPS comparison between optimized and non-optimized (original) disk only environments (Neuralytix, 2013)

info@neuralytix.com | www.neuralytix.com | @Neuralytix
1129 Lexington Avenue, #3 | New York, NY 10075 | (212) 724.1188

Again, on a relative comparison basis, assuming slower 7,200-RPM drives cost $100 per unit, and 15,000-RPM drives cost 33% more, the total cost of the "optimized" environment would be 88% more expensive.

The costs associated with this example is further made worse when taking into account the real estate costs associated with the number of racks necessary, as well as the power and the cooling costs!

These comparisons clearly show that while performance may become optimal, the cost of optimization is very expensive.

Introducing Flash

While flash solid-sate storage technology has been around for a long time, it has only been in recent years that the economic viability of flash has become recognized.

Flash for the enterprise comes in many form factors. Three of the more popular form factors are:

- Solid-State Disk (SSD);
- PCIe cards; and
- All-flash storage arrays.

SSDs and PCIe form factors have been popular for those environments that only seek to optimize one server or application. However, given the new economics around NAND flash, and the improvements in data efficiency technology, it is now not only possible, but also highly viable to leverage all-flash storage systems in place of or in addition to traditional HDD based storage systems.

Flash can provide 100X the performance of traditional HDDs. Flash can cost 10X more than high performance HDDs. However, data deduplication and compression algorithms can provide an average of 10:1 reduction in the actual amount of data stored, thereby, negating the increase in cost.

info@neuralytix.com | www.neuralytix.com | @Neuralytix
1129 Lexington Avenue, #3 | New York, NY 10075 | (212) 724.1188

Augmenting traditional HDD environments

In large complex database environments, with say 1PB of data, approximately 50TB of the data would require optimization. Using the above table again, but instead of 15,000-RPM disk drives, we replace these with flash capacity (without data deduplication), we can see how much more efficient flash is.

	# of 7,200-RPM drives	# of **flash** drives	IOPS
Original Environment	100	0	10,000
Flash Environment	100	1	20,000[1]

Assumptions: 7,200-RPM disk drives yield 100 IOPS, while flash drives yield 20,000 IOPS (assumes equal capacity)
Notes: This table is for relative performance comparisons only, and does not represent actual performance

Table 2: IOPS comparison between non-optimized (original) disk only and optimized disk+flash environments (Neuralytix, 2013)

This environment only incurs a small 10% increase in cost.

Replacing traditional HDD environments with flash only environments

Most mission critical applications leverage some form of relational database management system (RDBMS). RDBMS typically have files that are considered "hot." These files include log files, indices, caches, etc.

For many years, users have used short-stroking on HDDs. This is a technique where data is only written on the outer tracks of disk drives. This reduces the amount of movement required for the disk drive arms and heads to read and write data – optimizing the performance of the disk drive. However, this results in significant waste in capacity, because the inner tracks (i.e. the majority of the disk drive) are not used. Utilization rates with this technique can often be less than 30%, and in more extreme cases, less than 1%.

Since flash performance is essentially one speed – ultra-high speed, it does not require special techniques like short stroking to achieve high performance. As such, flash does not have the same overhead.

Take flash technology one-step further, and not only is there no loss of capacity, but capacity can actually be further optimized. Data can be compressed or deduplicated. Given

[1] This is a conservative estimate. Many flash drives can exhibit more than double this performance.

info@neuralytix.com | www.neuralytix.com | @Neuralytix
1129 Lexington Avenue, #3 | New York, NY 10075 | (212) 724.1188

the randomness of RDBMS data, compression of the data is the preferred and viable way of data efficiency.

For data not in a RDBMS or unstructured data, data deduplication is likely to prove optimal. So under whatever workload and furthermore, under multiple workloads, flash cannot only maximize performance, it can also optimize capacity through compression and deduplication.

Figure one shows the balance between compression and data deduplication. At lower rates of overall data efficiency, compression plays a commanding role. As data efficiency gets better and better, deduplication becomes a major factor in achieving these extraordinary rates of data efficiency.

Every enterprise's environment is going to be different. Even those having identical workloads. Nevertheless, one thing is consistent: compression and deduplication help to maximize capacity. Only flash has the ability to perform both of these data efficiency techniques at performance.

info@neuralytix.com | www.neuralytix.com | @Neuralytix
1129 Lexington Avenue, #3 | New York, NY 10075 | (212) 724.1188

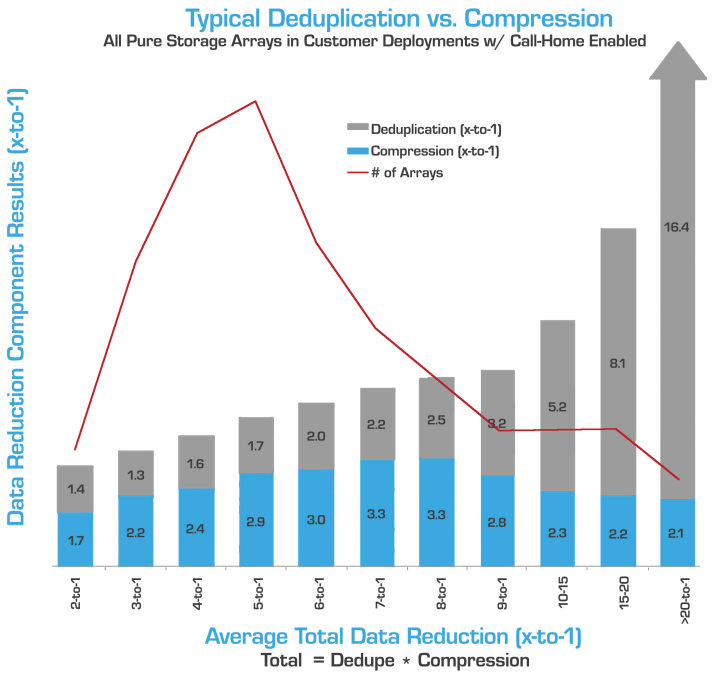

Figure 1: Typical Deduplication vs. Compression (PureStorage, 2013)

Using the assumptions in Table 2 above and augmenting that with data deduplication (where flash now has 10X the amount of capacity, thereby neutralizing the increase in the cost of flash), it is actually possible to double performance, and reduce the cost of storage by 99%! However, if the same amount of capacity is required, then some of the cost advantaged is absorbed.

Part of the movement towards a flash-only storage subsystem is simplicity. Having a singular storage system reduces management, improves performance, and creates a repeatable, predictable, scalable storage subsystem; allowing system administrators to focus on other areas of the IT environment.

The application owners also benefit. They know that their application data for the most part have the highest performance storage available to them. This gives them confidence. It provides them with a predictable level of performance. It allows them to scale with near linear performance – avoiding the political and tenuous tension that invariably arises when applications demand more storage, but system administrators are reluctant to over

info@neuralytix.com | www.neuralytix.com | @Neuralytix
1129 Lexington Avenue, #3 | New York, NY 10075 | (212) 724.1188

provision. By using an all-flash storage system, application owners no longer need storage to be over-provisioned in order to attain optimal performance.

Neuraspective™

Enterprise application performance has been a "hot" topic since enterprises began depending on their IT infrastructure to deliver competitive advantage.

The transition from traditional magnetic HDD reliant storage systems to all-flash storage systems is evolutionary. It is not a "rip-and-replace" approach. It gives enterprises opportunities to address point performance issues (particularly as it relates to 'hot' files, log files, and indices for enterprise applications), and eventually look to migrating the entire database onto flash storage.

The primary drivers are not only performance, but also cost. The combination of advanced data efficiency technologies augmented with ever-falling flash prices means that there are often little to no difference in capital cost between traditional HDD or hybrid storage systems and all-flash storage systems.

Additional drivers also come in terms of ancillary costs, such as reduced power and cooling, minimal rack and floor space. These cost factors can be a large contributor to the overall cost of delivering the necessary performance and capacity for enterprise applications.

Finally, there is also simplicity. An all-flash approach allows enterprises not to have to manage multiple storage systems or multiple storage media, while still benefiting from low cost capacity and optimal performance.

Predictability, reliability, repeatability and scalability are four cornerstones on which all-flash arrays have been designed and architected. For any IT infrastructure, predictability has always carried the price of management. By simplifying and unifying the storage system to a singular media type, it exponentially increases the predictability of the performance (and chargeback for capacity) and therefore decreases management complexity. Multi-controller and high availability designs ensure that all-flash storage systems are highly reliable. IT administrators can then design storage "templates" that can

info@neuralytix.com | www.neuralytix.com | @Neuralytix
1129 Lexington Avenue, #3 | New York, NY 10075 | (212) 724.1188

be repeated for each application, emphasizing the predictability aspect. Finally, all-flash arrays are designed to scale in capacity (and performance) with an enterprise.

Neuralytix research predicts that many organizations will look to all-flash arrays to first solve point performance issues related to enterprise applications, and that a material number of enterprises will see the benefit and simplicity of all-flash arrays and adopt this approach for their entire enterprise.

About Neuralytix™

Neuralytix™ is the global leader in contemporary and relevant IT market research and consulting firm. We have a holistic and forward-looking approach to research, which makes it unique and the most relevant research in the IT industry today.

Visit http://www.neuralytix.com to learn more.

info@neuralytix.com | www.neuralytix.com | @Neuralytix
1129 Lexington Avenue, #3 | New York, NY 10075 | (212) 724.1188