VIOLIN CONCERTO 7000 ALL FLASH ARRAY

PERFORMANCE PACKED WITH DATA SERVICES

JUNE 2014

All Flash Arrays (AFAs) are plentiful in the market. At one level all AFAs deliver phenomenal performance compared to an HDD array. But comparing AFAs to an HDD-based system is like comparing a Ford Focus to a Lamborghini. The comparison has to be inter-AFAs and when one looks under the hood one finds the AFAs in the market vary in performance, resiliency, consistency of performance, density, scalability and almost every dimension one can think of.

An AFA has to be viewed as a business transformation technology. A well-designed AFA, applied to the right applications will not only speed up application performance but by doing so enable you to make fundamental changes to your business. It may enable you to offer new services to your customers. Or serve your current customers faster and better. Or improve internal procedures in a way that improves employee morale and productivity. To not view an AFA through the business lens would be missing the point.

In this Product Profile we describe all the major criteria that should be used to evaluate AFAs and then look at Violin's new entry, Concerto 7000 All Flash Array to see how it fares against these measures.

Impact of Storage on Application Performance

Enterprise computing is all about the application. Storage has been the bane of application performance as far back as one can remember. However, the situation has worsened because other infrastructure technologies, such as compute, networking, memory and busses have increased in speed and bandwidth at an exponential rate compared to storage. Yes, storage capacities have increased dramatically over this time period but storage performance has remained practically static, relatively speaking. As a result storage, without question, has become the single technology keeping application performance hostage. As an industry we have coped with this widening gap between storage and other elements of a computer through a variety of methods including overprovisioning and caching. These approaches have added cost, space, power, cooling and management penalties. For instance, to get higher IOPS, we add HDD spindles, since each HDD can only deliver a limited number of IOPS. We employ the highest performance HDDs we can find and yet only use 15% of their capacity just to get increased performance.

If something significant is not done immediately, the expense and inefficiency of disk-based storage will grind enterprise IT to a halt. Flash technology could not have come at a better time.

Flash is indeed a replacement for HDDs. But when application performance improves by a factor of 10 or 20 or 50, something else happens. You can do what was not possible before. You offer services that were impossible to dream of yesterday. You serve customers differently. You compete

87 Elm Street, Suite 900 · Hopkinton, MA 01748 · T: 508.435.2556 · F: 508.435.2557

differently. You perform your task differently. You get the point. Very rarely does a technology come to market that is able to transform who we are as people and how we relate to each other. It transforms societies. That is the power of flash.

Enterprise Demands on Storage

Enterprise Workload Demands

The probability that any enterprise storage array will be used only for one application is slim to none. Given all-flash array expense the probability is high that the user would deploy an AFA to handle a multitude of workloads. This is especially true since an AFA can enable a mission critical application to be brought under the umbrella of virtual servers when it was otherwise not feasible (due to performance criticality). This requires a well-designed AFA to deliver consistent performance across a multitude of workloads. As new workloads are added the existing performance must not suffer (assumption is maximum array performance is not reached by existing applications, of course). Neither should the latency go up for existing applications. AFAs that fail to make this grade are dangerous for your environment since they would impact exactly the applications that were sacred to the enterprise.

Enterprise Data Protection Demands

All enterprise caliber arrays today have a complement of data protection features built in to the arrays. Flash is a different media than an HDD. It deteriorates differently. It fails differently. It behaves differently under failures. Flash requires brand new thinking and it is crucial for the user to understand exactly how data is being protected inside an AFA. One should be prepared to learn entirely new concepts, given the variety of solutions available in the market today for flash. How is performance impacted by the loss of a cell/page/flash chip/flash shelf? How vulnerable is the array to another failure? For how long? These are all questions that become critical in an AFA.

Enterprise Cost Demands

When does cost not matter? Some AFA vendors have inline data deduplication and compression built in as a standard features. These vendors use a factor of 4 or 10 to 1 to reduce their cost per GB in order to show a near equivalence to HDD-based system prices. Yet many workloads do not lend themselves to data deduplication or compression. Some lend themselves to compression and not deduplication. Some benefit from deduplication. Some benefit from neither. Your mileage will vary when calculating cost by projecting the impact of data reduction technologies. Because of this uncertainty, making comparisons based on raw capacity is a more reasonable approach to comparing cost.

However, in general, if data reduction comes at the expense of performance one has to question why one is buying the AFA in the first place. And secondly, does data efficiency impact consistency of performance. If it does, all bets are off.

Operating costs are another consideration. The density of an AFA will help determine the amount of operational costs such as power, cooling, and floor space. Any flash array will give you better density (except when substituting SSDs for HDDs in a legacy architecture array) today. But the real comparison has to be amongst AFA vendors. And here density varies all over the map. If one vendor can give you 70TB in 3U it may require 8Us from another vendor. If all else is equal one must opt for the highest density possible. In any case, in the world where space is limited in data centers, density matters.

87 Elm Street, Suite 900 · Hopkinton, MA 01748 · T: 508.435.2556 · F: 508.435.2557

Enterprise Scale Demands

The industry is rife with scale-up vs. scale-out discussions. Both have their ups and downs. Of course, the ideal architecture is one that provides the ability to grow as you add new applications, and to keep up with application growth. The ability to grow LUNs, add shelves and expand a single name space across physical devices can be very useful tool to manage growth.

Enterprise Performance Demands

Early AFA designs were notorious for giving you blazing performance when initially deployed. IOPs were strong and latency was in 10s or 100s of microseconds. And bandwidth was excellent. But after a few days of being in production all three parameters of performance dipped dramatically. Why did this happen? It happened because flash media writes differently than magnetic media. As long as the writes were to new cells (which were ready to be written into at the factory) performance was great. But then the writes were to existing cells, requiring surrounding existing data to be moved, and the cells to be "flashed" before new writes could be written. This process slowed the write performance dramatically and latency increased correspondingly.

Consistency of performance matters. And good design solves the problem by a variety of methods used to mitigate garbage collection and wear-leveling. Without this engineering, application performance remains unpredictable causing unexpected damage, depending on the type of application.

Enterprise Availability Demands

With the rise of global enterprises, and the ability to run 24 hour/day operations, the data center needs to be available all the time. Planned outages are not acceptable, and steps must be taken to prevent unplanned outages. All enterprise appropriate AFAs need to provide hot swapping of key components, non-disruptive upgrades for software and firmware, and HA configurations with redundancy so that there are no single points of failure. When the data center suffers a storage outage without these techniques, you are out of business.

Enterprise Data Services Demands

All enterprise caliber storage arrays today offer a variety of data services that assist the IT administrator to get the most out of the array. Data efficiency features like thin provisioning and thin clones. Simple provisioning. Data protection features like snapshots. DR features like synchronous and asynchronous replication. And so on. Several of the initial set of AFAs that came to market came without these features. However, the application performance pain level was so high that IT user had no choice but to accept them as such. This is no longer true and the user should expect a full complement of well-designed storage applications to accompany the AFA.

Clearly, not all storage applications are equally efficient. Snapshot technology from one vendor to the next can be vastly different. Similar concerns apply to other applications. Not only is it reasonable to expect the availability of these applications but one should expect them to be best of breed today.

87 Elm Street, Suite 900 · Hopkinton, MA 01748 · T: 508.435.2556 · F: 508.435.2557

All Flash Array Options for Enterprise Storage

SSD-based Enterprise All Flash Arrays

At this stage of AFA evolution we believe hardware design matters. The reasons are many-fold. One, using SSDs from another vendor relegates important aspects of wear leveling and garbage collection, along with several other functions, to the SSD vendor. These functions are therefore controlled at the SSD level beyond the control of the array vendor. This creates a situation where a critical component in system design (the SSD) cannot be tightly integrated with the design of the storage subsystem. If the array vendor chooses more than one SSD vendor (for supply reasons) they have to account for the differences that arise from different methodologies, as they relate to performance and other characteristics of the system. The array vendor gets what they get from the SSD vendor and they deal with it.

Just as in the case of density, all AFAs deliver on the promise of lower power and cooling. However, those that are denser will generally be more efficient in power usage. Control over flash chips, for instance, gives the designer the ability to use one controller layer rather than a controller per SSD plus an array controller, typically yielding better results in power and cooling. Vendors that control their hardware design and use flash chips generally achieve higher densities than those who use SSDs.

Backplane-based Enterprise All Flash Arrays

On the other hand hardware-centric designs use flash chips and manage the functions of garbage collection and wear-leveling at the array level. They are also not restricted by the SSD form factor, which was designed to make flash look like a disk drive, and is not very space-efficient. At least at this stage of development, greater control is possible over garbage collection and wear-leveling when these functions are controlled at the array level. These impact resiliency and longevity of the array and allow the extraction of maximum possible performance out of the flash being used.

Keeping in mind that performance was the number one reason for evaluating an AFA, hardware design matters. Given that the AFA will likely be used for the most mission critical applications in the enterprise, resiliency and longevity of media matters.

Let's take a look at one of the early all-flash array vendor's latest offerings in the backplane-based All Flash Array space.

Violin Memory's Concerto 7000 All Flash Array

Concerto Hardware Design

Now that we have laid out criteria for evaluating an AFA let's look at how Violin's new product Concerto 7000 fares in these measures, across the board. We will discuss the details of the array in each dimension. First, a look at the hardware. Violin Concerto 7000 is a 4U box HA controller with up to four storage shelves. Each storage shelf is 3U and has 64 Violin Integrated Memory Modules (VIMMs) instead of SSDs, four vRAID controllers, two array controllers and two I/O controllers. The two array controllers are configured in an A/A manner so a failure of a controller does not affect availability. Up to four 3U storage shelves, each with up to 70TB of raw flash each can be added, for a total raw flash capacity of 280TB. The system delivers over 500K IOPs at a 70/30 read/write ratio of random data when fully loaded with mixed and multiple workloads. The design is fully

87 Elm Street, Suite 900 · Hopkinton, MA 01748 · T: 508.435.2556 · F: 508.435.2557

redundant, with hot swap power supplies, fans, vRAID modules, array controller modules and VIMM modules. There are no single points of failure in the design, which is enhanced with hot swap hardware components and non-disruptive upgrades for software and firmware.

Violin chose to design hardware using flash chips and not use SSDs as many vendors do. This control of flash at the array level yields gains in performance, density, and resilience. Garbage Collection is managed intelligently yielding consistent performance.

The Violin Flash Fabric Architecture

For NAND flash to be viable in data center applications, it requires a very different set of attributes compared to a PC or consumer device. Sustained and predictable performance is required for data centers. Violin Memory solved the inherent "behavioral issues" of NAND flash through innovative design and chip-to-chassis integration of hardware components known as the Flash Fabric Architecture and software called the Violin Memory Operating System. Violin's Flash Fabric Architecture in particular enables thousands of flash devices to operate efficiently together, masks the chip level issues, and delivers reliable and sustained system performance.

The architecture is designed for fast random writes, erases that are simple, fast and hidden from the user, reads that are accelerated through parallel access to billions of pages and other features designed to extract maximum performance without compromising reliability.

Violin's Flash Fabric Architecture uses multiple layers of hardware technologies, and patented flash optimization algorithms implemented in hardware, operating at line rate.

- At the system's core lies a resilient, highly available deep mesh of thousands of flash dies that work in concert to continuously optimize performance, latency, and longevity.
- Violin Intelligent Memory Modules (VIMM) organizes this mesh of individual dies into intelligent flash management units. VIMMs provide a hardware-based Flash Translation Layer with Garbage Collection, wear leveling, and error/fault management.
- VIMMs and the switched fabric layer work in conjunction with vRAID, Violin's hardware-based RAID algorithm specifically designed to increase reliability and reduce latency.

Collectively, the mesh of flash dies organized into VIMMs integrated into the switched fabric and overlain by vRAID make up the Violin Flash Fabric Architecture.

Violin's Flash Fabric Architecture works with the Violin Memory Operating System to enable reliable, highly available storage that offers multiple benefits:

- **Spike-free Low Latency** - The Flash Fabric Architecture delivers spike-free and predictable latency that is 95% lower than HDD and substantially lower than SSD and PCIe card solutions.
- **High Bandwidth** - A single Violin All Flash Array support over 4000 flash devices and 500 independent flash interfaces. This provides the bandwidth needed for outstanding flash performance with four times greater bandwidth than HDD storage systems and substantially greater than most SSDs.
- **Extreme Reliability** – All active components of the Flash Fabric Architecture are hot-swappable for enterprise grade reliability and serviceability.

Concerto Data Services Software

For many enterprise customers, performance is often not enough to move to the all-flash data center. Because applications need to always be available, their storage must always be available.

87 Elm Street, Suite 900 · Hopkinton, MA 01748 · T: 508.435.2556 · F: 508.435.2557
www.tanejagroup.com

Storage must be easy to manage. It must be scalable. It must be protected. It must be efficient. The Concerto Data Services surround high performance with enterprise caliber business continuance, safety, scale and efficiency to enable the all-flash data center.

Concerto Data Services can provide several benefits:

- Deliver flexible and powerful business continuity with asynchronous and synchronous replication capabilities. Synchronous replication enables local or stretch clusters for zero RPO/RTO. WAN optimization technologies are used to reduce data transferred across a WAN to minimize bandwidth needs and reduce costs
- Data efficiency technologies such as thin provisioning and snapshots to reduce the amount of storage required. Snapshots can be crash consistent or application consistent and are writable. Thin clones are supported as well
- Data protection features such as clones, mirroring and encryption to protect valuable enterprise data. Consistency Group can be defined to support applications using multiple LUNs. Continuous data protection is supported
- Scalability to four storage shelves for a maximum of 280TB of raw capacity. Capacity and LUNs can be expanded online, without disruption to the application. Capacity is pooled across shelves for a single namespace

The Concerto 7000 system is designed to deliver the power and capabilities that allow enterprises to create an all-flash data center. The industry is just now coming to grips with the phenomenal benefits that accrue from an all-flash data center. And now that such a dream is possible CFOs can enjoy cost savings (power, cooling, space), and business users can enjoy transformational improvements in application performance and offer additional revenue-producing services to their customers. And for once the CIOs can stop spending a large majority of their time on mundane storage management and start focusing on value producing activities.

Taneja Group Opinion

We believe the time for AFAs is here and there are plenty of choices in the market. However, not all AFAs are born alike and it is important for users to understand the differences and choose the one that makes the most sense for their infrastructure and application needs. We believe there is no enterprise that is not a viable candidate for an AFA. 65% to 70% of all workloads have already been virtualized. By definition then those servers are running multiple VMs, often with vastly varying workload patterns. In these situations, it has become clear over the past five years that traditional storage array architectures do not fare well. In fact, they collapse. Flash technology could not have come at a better time. Whereas hybrid arrays will certainly fill the bill in many of those cases, the most mission critical of your applications will require an AFA. Given the right AFA you may even be able to bring the next set of most mission critical applications under the purview of virtualization and start enjoying the benefits that come with it. In many cases, these latter applications are still sitting on physical servers because that was the only way you could fine-tune the application performance. With an AFA all that changes.

But as you have seen from above not all AFAs are alike and the range of efficiencies and deficiencies is long, since the market is relatively new. Violin was a pioneer of AFAs, being the first vendor to announce and ship such an offering. They have several years of experience over the others and they are in the fourth generation of product. We believe Violin has eliminated the one issue that had prevailed before Violin Concerto 7000's entry: enterprise caliber data services. This issue is now off the table and Concerto 7000 comes with a vast array of these services. The architecture has already proven itself in the field. In our view the product stands out in the areas of performance, density,

87 Elm Street, Suite 900 · Hopkinton, MA 01748 · T: 508.435.2556 · F: 508.435.2557

consistency of performance and several other dimensions, as stated above. If you were looking at alternative AFA solutions because Violin was missing the enterprise data services, your search not need to go any further. Check out the Concerto 7000 and see if you like what you see.

Regardless, we encourage the user to look at the AFA as a business transformational tool, not simply a "faster than HDD array." Given the 5X or 10X or 20X (or more) application performance acceleration you are likely to see, you have to ask yourself questions like, "What can I do now that I couldn't do before? How can I serve my customers better? What new revenue-producing services/products can I offer that were not possible before? How can I improve my internal efficiencies?" This is the magic of AFAs. Violin may have lost a bit of the mojo after essentially creating the AFA market.

With Concerto 7000 Violin gets its mojo back!
