Server-Based Storage Acceleration

Entering the Performance Race With a Server-Based Approach

By Jeff Boles

A little over two weeks ago, EMC made a major announcement of their offering in a product category that looks poised to reshape how enterprise storage works over the next few years. That product previously known as Lightning is now called VFCache, and is one of many emerging technologies in a category we more broadly call *storage accelerators*. EMC is not the first or sole vendor in this space – they are preceded by Marvell with Dragonfly, Fusion IO with IOTurbine, and others – but EMC is the biggest vendor to announce such a product. The exciting thing about this type of technology is the potential for changing how enterprise storage works, for the better. The promise is that the business will be able to easily deliver and scale performance without requiring wholesale alteration of existing storage infrastructures or the way we implement and manage storage today.

In this article, we'll recap the on-going battle for more storage performance, describe where this new generation of technology fits into the picture, and discuss what potential customers should have in mind when looking at making server-based storage accelerators (SSA's) a part of their storage performance strategy. In a nutshell, there is tremendous opportunity in these accelerators that may well influence and change storage architectures for years to come, but early adopters should be thinking strategically about how and why to implement these accelerators both today and in the future. Asking your vendor the right questions may help you make a decision that maximizes your current and future value from these highly affordable acceleration technologies.

Accelerating the battle for storage performance

Over the past two decades, delivering performance from data storage has become the bane of the enterprise storage administrator, and the challenges have recently seemed poised to severely constrain the capabilities of the infrastructure in the face of increasing virtualization, consolidation, and data analytics. While data density and compute horsepower have increased almost logarithmically over the years, storage performance improvements have been comparatively miniscule.

Storage accelerators can be server-based or network-based, but at the most basic level they sit in a server or in a network and cache or tier data in order to intercept and handle IO requests from high speed local storage (typically a combination of SSD, DRAM, and even spinning disk). The goal of these vendors is to transparently insert high-speed storage technology between the application and the storage, so that the traditional storage array has to deliver less IO. As a result, the totality of storage performance within an enterprise's walls can be scaled beyond the limits within the enterprise's storage arrays, and the cache that is closer to workloads than the storage controller can even improve latency over what is possible from an array or filer.

Server-based storage accelerators or SSAs, similar to VFCache, are the most recent addition to this market. Previously, acceleration looked like it may remain solely the domain of network appliance innovators like Alacritech with their ANX1500, Avere with their FXT appliances, and GridIron in

the FC fabric. The hardware component of these more recent server-based storage accelerators (SSA's) may take the form of integrated storage adapters (although no HBA or NIC enabled caches have yet been introduced), solid state PCIe devices, or even SAS or SATA SSD devices. In all cases so far, these devices must be coupled with a host-based software layer that is responsible for intercepting and altering storage interactions to intelligently cache or tier the most active data to the solid-state storage on these server-attached devices.

In some cases, redirected IO may be used purely for read-only access, with all write IOs requiring acknowledged access on the storage array (this is the case with VFCache). In other cases, local solid state media may act as a read-write cache, with writes being protected one way or another, and then batched together and synchronized to the array in a way that requires less IO and less processor utilization from the storage controller. In yet other cases, SSA's may also permanently and authoritatively tier data off of backend storage (either authoritatively locking it for an extended period of time or altogether relocating it).

It is worth observing that the biggest challenge resides in intercepting and tiering data. In this sense, the best innovators in the software layer and underlying data intelligence may well win the game.

Finding maturity in an evolving market

At this early stage in the game, vendors are racing for leadership in the market, and capabilities are advancing rapidly. What EMC has claimed as a significant performance advantage will soon be topped by the closest competitor, and so on; when it comes to solid state technologies, this is the norm. But SSA's will be about far more than the raw solid state technology that delivers this performance.

As this landscape evolves, given the highly effective pricing of server-based accelerators, customers with performance problems shouldn't shy away and wait for advancements when there are reasonable assurances of interoperability. But when making purchase decisions, focus on a few key areas will help customers understand where innovation and differentiation will likely surface, and consequently decide if their investing in the right vendor and technology over the long haul. Let's take a look.

Efficiency and integrated design, for now, makes a difference.

The current generation of server-based storage accelerators are going to intermediate the server-to-storage connection to intercept, redirect, and accelerate IO. The industry has indirectly been working on this for at least a decade (and arguably more when considering caching in an array, which is fundamentally little different). There is little about this intermediation that is simple, and it is only available today because of intersecting trends around operating systems, virtualization, consolidation, and processing power that have enhanced the ability of vendors to interact with storage IO paths. Today we have a much better IO stack in which to interact than ever before, irrespective of the operating system, application, or hypervisor under consideration. Moreover, we now have the processing performance and affordable accelerated storage substrate (solid state) with which to do it.

But the multi-generational development of this opportunity can be easily undone by the smallest amount of performance inefficiency or overhead. Moreover, because this server-based accelerator technology stands to move much of the IO intervention to the edge, the smallest amounts of processing inefficiency could be devastating at scale. Efficiency innovation will be fought over in a very similar way that it was fought for with storage virtualization, where IBM with SVC and HDS with USP eventually used efficiency and latency to dominate much of the market.

Today, the battleground has moved much closer to the server, and the technology will be applied for a slightly different purpose. But latency will be even more important here, especially because the same solid state and processing technologies in the storage array are letting vendors deliver more and more performance with much faster levels of response. In turn, SSAs will always have to be faster and more performant to have broadly realizable benefit.

In our view, building low latency solutions takes both a tightly coupled IO interaction layer and an innovative hardware layer, and at this stage of the game, the best of those hardware layers will likely always be built around a special purpose ASIC. Right now, EMC claims to deliver highly efficient IO interaction as a result of bringing intellectual property from their very mature PowerPath software stack to the challenge of IO interaction. Based on current results in Taneja Group lab exercises, others will soon follow them and demonstrate similar efficiency. But as a word of warning, today, we see vendors who have designed products primarily as localized storage (effectively direct attached storage) struggling with IO stack redirection. Such vendors appear to be far from delivering the processing efficiency necessary to make their IO caching strategies effective for heavily consolidated enterprises. Look before you leap.

Integrated with Enterprise Storage

Vendors entering this market are universally focused on transparency, with the realization that the business customer has no place for a widely distributed technology that interferes with existing storage practices. But delivering an effective solution that is transparent in the right ways takes a clear understanding of business storage systems and storage practices.

SSAs in particular take on enormous complexity. Not only must they deliver efficiency with a small form factor, but they also have to avoid harming data, while moving considerable data handling functionality out into a server that may be far less protected and reliable. Running a key list of storage concerns by any given vendor will likely reveal whether a particular product has long term legs, or is simply a convenient but temporary hardware Band-Aid.

The most significant data handling issues at this early stage of the market will revolve around data control and data protection. How will your vendor protect against device failures while still delivering the speed boost really expected from an accelerator? How will the vendor let you manage the protection of data while data is being cached or tiered on these distributed devices? How deeply can you control what is cached to make sure you are achieving optimal performance and avoiding problems like flooding the cache with useless data? How is your vendor thinking about someday providing acceleration for multiple servers accessing the same storage volumes?

Over the longer term, and at larger scale, distributed technologies like SSAs could create significant complexity in other areas like service level management and capacity planning. How

will a vendor help you someday look at the performance acceleration delivered by these products as part of your entire storage service level, architect your infrastructure for maximum optimization and minimum risk, and plan to handle failures in a way that does not compromise the performance necessary to meet your service levels?

Today, expect equal part answers and equal part vision, because the right vendors are working on these problems, but the market is still in early stages. But vendor answers to these questions will reveal those who have the right storage pedigree to really tackle the challenges.

Integrated for Flexibility

With the announcement of any fundamentally new technology by a major vendor, many customers immediately begin considering the risks of vendor lock-in. With the solid state technologies behind server-based accelerators, lock-in risks may be greater than usual. Any given solution may not advance as rapidly as other acceleration alternatives in a highly competitive landscape. Moreover, these devices may create storage-to-server dependencies that restrict cross-generation compatibility and make maintaining the storage infrastructure more complex. With those risks, it seems hard to believe that Taneja Group thinks these technologies could become pervasive. Once again, it is entirely up to innovation.

We're optimistic that vendors will focus on innovating in the host layer as well as the hardware, but that they will also decouple these two things. This separation will help lead to an ecosystem of interoperable solutions, where the best vendors from both sides of the aisle can bring innovation to the table. For certain, EMC is painting a vision of an ambitious innovation that is rooted in key parts of their software and IO management expertise – both PowerPath and FAST-VP. Standardizing how these software layers might interact with a rich system of potential IO accelerators (server-based and network-based) may be a truly disruptive underlying innovation that is far different than the packaged accelerator device being offered to the market in this first generation of product.

The natural way to delineate where these functions reside would be with a standardized API approach across vendors – a task that seems particularly idealistic given the history of the storage industry and undertakings like SMI-S.

The road is paved with APIs

Such APIs are not just for SSAs. Network-based accelerators, or even tiering within an array, could make efficient use of very similar APIs as well.

What are these APIs? The most apparent opportunities for API innovation shake out in a few different areas:

Synchronization of distributed caches so centralized management isn't broken.

Most obviously, vendors will need a way to trigger synchronization of data back to centralized storage if snapshots and replication are going to be executed against a consistent and recoverable set of data. Standardizing on a single API for this can make sure all business data is handled and protected well, irrespective of the devices used, and help pave the way to prolific adoption.

Storage side and flash side collaboration over what is cached or tiered.
It may be important to inform storage accelerators in a systematic, centrally managed way about how to cache data and what data should be cached. There may be mission critical data, clustered storage systems, or unusual sets of data (such as large block sequential) that should never be cached. But informing the array about what is being tiered or cached is just as important – if a LUN or file is tiered to a server-based storage accelerator, it may be the case that is should not be consuming precious storage array cache. As EMC has already hinted, a server-based accelerator may someday become just another tier with Fully Automated Storage Tiering (FAST). The right-APIs can make storage-to-accelerator cooperation a reality, and can help storage administrators broadly apply large numbers of server-based storage accelerators to the right problems.

Virtual infrastructure integration
Ultimately, the industry is already approaching a point where this type of API integration will be more possible and practical than ever before, because server virtualization has largely taken us here already. With the virtual server hypervisor sitting in the middle of the infrastructure, APIs like VMware's vSphere Storage API's (VASA and VAAI) may enable the storage accelerator to provide storage-complimentary new features not currently available in a business's centralized storage array. As an example, an accelerator could enable enterprise-storage-like features, such as XCopy or Zero Block/Write Same even when those features don't exist on the storage array. No vendor has proposed specifics here, but the opportunities seem nearly limitless.

Visibility
Finally, server-based accelerators also stand poised to provide a whole new level of insight into what is really happening with stored data, and could employ standardized APIs to not only alert storage managers upon failure, but also inform array auto-tiering policies to re-optimize storage to maintain service levels. Hopefully, vendors will be proactive this go around, rather than leaving visibility as the after thought process that has been too often the storage industry norm.

Innovation driven adoption
The Storage IO challenge is an enormous problem that has for far too long been unaddressed and storage accelerators as an easy, cost-effective, and non-disruptive solution to this problem stand to achieve a significant footprint in every data center. But it shouldn't be missed that the potential reaches far beyond just delivering fast IO, and the resulting products will quite probably change how we architect and work with storage across the enterprise. But it takes a storage focus that extends well beyond just developing clever PCIe storage cards – instead it takes a focus on the entire storage infrastructure and how all the pieces fit together. Fortunately, we're quite confident that bright minds are fully engaged and innovating, with just this focus in mind.

One vendor that has remained out of the recent hubbub is Marvell with their Dragonfly accelerator. One of the reasons we

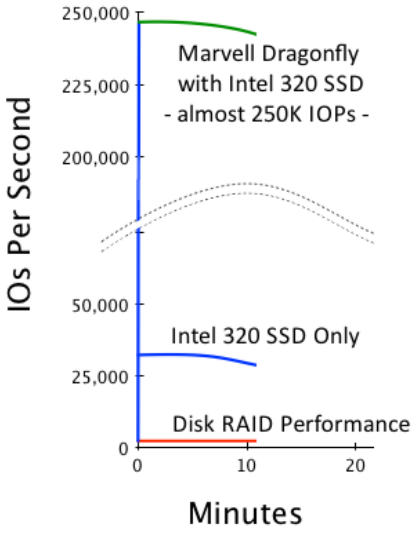

Figure 1: An upcoming Taneja Group Technology Validation Report will discuss hands-on findings with the Marvell's Dragonfly storage accelerator, including this initial preview performance sample from a 4K random write FIO benchmark, but also performance in a real world virtual desktop infrastructure. Look for the published test for more discussion on performance over time; how SSAs may solve SSD write-cliff problems; and how SSAs can be a tremendous boon for virtualization.

are proponents for this market is because we've recently spent extensive hands-on time characterizing this accelerator in both a formal lab test and through on-going use in a Taneja Group lab environment in Phoenix, Arizona. In another article on this topic we'll focus on what emerging acceleration technologies mean for the design of the storage infrastructure, and look at a few Marvell test results as an indication of how impactful this technology may be. For aggressive workloads, expect improvements that will need more than the fingers on two hands to measure.

As I've said before, 2012 will be the year of storage performance, and the emerging performance technologies may change the implementation and operation of storage inside the enterprise walls. There are now more choices for extending your investment dollar by optimizing for capacity and performance beyond the limits of a single array. This year's performance breakthroughs will completely reinvent the storage TCO equation for businesses of all sizes, and will as a consequence merit adoption sooner than any previously introduced storage technology. Now is the time to start paying attention to where and how centralized storage, virtual infrastructures, and acceleration technology intersect.