

# STORAGE SWITZERLAND

## HIGH PERFORMING VIRTUALIZED WORKLOADS REQUIRE A NEW STORAGE PROTOCOL



George Crump, Senior Analyst

Virtualization workloads are different, they have different storage demands in terms of performance, scalability and of course, cost. A new storage protocol may be required to address these three conflicting demands, one that can keep up with the highly random workload, scale to meet expanding capacity and port requirements while matching the cost reducing popularity of server virtualization. As data centers converge to Ethernet, the ATA over Ethernet (AoE) protocol provides a compelling new vehicle to meet these challenges. In the first article in this series we will look at the performance challenges that server virtualization workloads cause and how they can be addressed by AoE.

### Performance Challenges

Virtualized workloads create unique performance challenges caused by the consolidation of multiple servers to a single physical host. In the past each workload was on its own server with its own storage and network I/O channels to the shared storage device. In most data centers the overwhelming percentage of servers had only modest storage I/O requirements. At any given time there was no unusual stress placed on the network. When and if

there was a high level of storage I/O activity it could be easily identified and addressed for the particular workload.

In the virtualized server environment things change. First, the dozens (if not hundreds) of servers with modest I/O needs are consolidated and run as virtual machines on just a few physical hosts. The I/O demand of each of these now-virtual servers still exists but is contained in a single host with fewer ports and less I/O bandwidth. The result is a single physical server that can have very high I/O density. This new level of density changes the way storage protocols should be examined.

The second challenge is the random timing of demands these various virtual workloads make on storage that spikes I/O. It's very difficult to predict which virtual machines will need I/O at any given moment. As a result customers tend to overbuild their storage infrastructure around a worst-case scenario, using expensive high-end legacy storage arrays, or artificially limit the number of virtual machines that any given physical host will support. This practice introduces inefficiency back into the virtual infrastructure that server virtualization was supposed to drive out. It can also scale only so far given the cost impact of legacy storage infrastructures.

Some protocols have developed sophisticated tagging capabilities to identify a specific virtual machine and assign it priority access to storage I/O. This assumes that the virtual machines needing that I/O can actually be identified and given the correct priority. Once assigned, the prioritization of these various virtual machines is relatively static, it can't automatically be adjusted for workloads that peak at different times. To be most efficient, all workloads would need to be given a high priority at random times. The reality is in most cases virtual machines need I/O at different times; while a single or group of workloads can be singled out for priority access it's rarely the right workload at the right time. This type of prioritization is like an HOV lane on a highway that goes largely unused by the bulk of drivers. If you could simply build a highway that has ten times as many lanes for the same amount of taxpayer dollars, you would.

Storage protocols that utilize another protocol like iSCSI (IP) and NAS (IP) have an inherent weakness when it comes to dealing with performance. There's a 'translation' of sorts that has to take place when data is sent or received by both the host and the storage system. While they do have the advantage of being able to connect hosts to storage via a simple Ethernet infrastructure, there's significant latency and overhead because of this conversion that keeps them from reaching the full potential of the underlying network. Additionally, there is a "tax" on host CPU resources as the conversion from iSCSI to IP is made. At a high level, both iSCSI and NAS make a sacrifice of performance for the simplicity of Ethernet.

Fibre Channel (FC) and AoE are native block protocols that require much less conversion. They can easily exploit the connection type to its fullest extent with minimal host CPU resource cost. While FC requires a separate legacy network with different cabling, switches and management, AoE has the advantage of connecting through a standard Ethernet network. It simply uses Ethernet at the physical layer as the

connection methodology, delivering more raw speed with the simplicity of Ethernet cabling and management. It also takes advantage of the rapidly improving economics of 10Gb Ethernet.

In virtualization environments, performance has become a primary concern since every host on the SAN now has the potential to be very I/O demanding at any moment.

Because of performance limitations of traditional IP-based protocols in larger virtualized environments, hosts need to be limited on the number of virtual machines they will support. Additionally, in these environments more demanding tier-one applications may never be virtualized because of concerns over their performance and how they may impact the performance of neighboring virtual machines. While tagging of VMs can help assure service levels of particular virtual machines giving them priority access to storage I/O channels, it does add another layer of complexity and the potential for an inefficient use of storage I/O.

FC and AoE, because they are inherently higher performing without having an impact on host resources, become viable alternatives to the popular IP based protocols. If the network and storage is fast enough to meet any and multiple performance demands, the need to spend time managing the infrastructure is substantially reduced. This is one of the reasons that FC has been such a dominant storage protocol for virtual workloads. It provides fast performance, but once tuning is eventually or the topology needs to be changed, it becomes challenging for most organizations to implement because of its steep learning curve. Additionally there is a significant cost issue when it comes to FC. Both of these problems AoE attempts to address.

AoE leverages the inexpensive Ethernet infrastructure without the overhead of IP and has performance characteristics and resource utilization efficiencies that are similar to fibre. This makes it ideal for virtualized workloads. The protocol provides simple, fast performance that does not require constant fine-tuning or virtual machine inspection to try to set various levels of service quality. If performance does become a problem AoE can be easily scaled by adding ports for linear scaling - add another card and it's automatically bonded and servers can use all available I/O ports as if they are one. For example two 10GbE ports would appear to be a single 20GbE port. As a result when the eventual fine-tuning does occur it is simpler and easier to implement thanks to the Ethernet foundation and inherent simplicity of the protocol.

When dealing with virtual workloads most protocols can be tuned and monitored to provide acceptable performance, particularly if I/O needs are modest. If the environment is going to grow, especially if that growth will be rapid or sporadic, then the process of I/O management becomes too much of a burden on an IT staff that's probably stretched too thin already. The "just make it fast" capabilities of FC and AoE make them viable options. When the "keep it simple" characteristic of AoE is added, this new protocol becomes genuinely compelling. Finally, include the abilities to scale and remain cost effective, and AoE becomes a top storage protocol candidate for virtual infrastructures.

The next challenge in dealing with virtualized workflows is one of scalability. A scalable storage infrastructure is critical to meet the demands of virtual workloads, since these workloads can now be brought online almost instantly. The storage infrastructure has to respond in the same fashion.

### **About Storage Switzerland**

Storage Switzerland is an analyst firm focused on the virtualization and storage marketplaces. For more information please visit our web site: <http://www.storage-switzerland.com>

*Copyright © 2010 Storage Switzerland, Inc. - All rights reserved*