Storage Switzerland, LLC

Fibre Forward - Why Storage Infrastructures Should Be Built With Fibre Channel

Prepared by: George Crump, Lead Analyst
Prepared: June 2014

Performance is a top demand of the storage infrastructure but it is not the top demand. While the speed at which data can be delivered to applications is critical for organizations looking to support more transactions, more virtual machines and more users, the ability of the storage network to maintain high availability and provide maximum, but simple scalability is paramount. After all, the performance of a storage network that is down is zero gigabits per second (Gb/s). Fibre Channel (FC) storage networks remain a logical choice not only for IT planners that are looking for maximum performance but also maximum scalability and resiliency.

While FC storage networks tend to enjoy a Gb/s bandwidth advantage over other commonly installed storage architectures, the rate at which data is transferred is only one small consideration when selecting a next generation storage network. In fact the line rate of the network is only a small part of the effective performance that the storage infrastructure will deliver. The ability to deliver data across a dedicated network infrastructure, without risk of transmission retries and with the ability to fine tune how the network is used, are potentially more important than the raw bandwidth.

The I/O Demand

There has been a dramatic increase in I/O demand in the past few years, driven by a number of factors. The universal cause of increased I/O demand is server and desktop virtualization. Instead of one host or desktop generating occasional I/O, a single host loaded with dozens of virtual machines or thousands of virtual desktops, means that each host attached to the storage network can potentially become consistent I/O generators. But the pace and size of that I/O is entirely random since each virtual machine will have its own I/O profile. Virtual environments have access to plenty of compute power and now thanks to flash, the storage media is prepared to instantly respond. Consequently, the storage network has to be upgraded to keep pace.

At the same time, the traditional I/O demand generators like OLTP/OLAP databases and trading environments are increasing their demand for I/O. Each of these environments has seemingly a never-ending demand to scale larger and respond faster. As is the case with virtualization, most of these environments have the compute power and storage media response time; they just need a fast, reliable storage network to transfer that data.

Another demand generator is what is commonly called big data analytics. In reality, big data analytics projects are actually millions of little data sets that are scanned and curated. This flat-file analysis needs a high speed network in order to deliver the results quickly enough that decisions can be made based on those analytics.

Real-time personalization is another area that an increasing number of organizations are looking to provide their customers. These organizations are essentially creating applications or web sites that adapt to the user as it learns about them and their surroundings. Real-time personalization involves being able to rapidly access a wide variety of information while the customer is interacting with the organization's services.

For example, if a customer is viewing a restaurant web site it might include directions to the closest location, along with all the menu options, reviews of specific menu items and available dining times. It could also provide a list of alternative franchise locations if the nearest location is fully booked and maybe the option to pre-order the meal online. This system might even learn that the viewer has certain dietary preferences and offer alternative menus.

Storage Is No Longer The I/O Laggard

The server layer of the data center is clearly able to generate more I/O than ever before thanks to increased processing power; virtualization environments are able to put that processing power to good use. In the past server processing power mostly went to waste, as did most upgrades to the storage network. This is because the actual storage media was too slow to respond to the increase in I/O demand, nor did it have ability to transmit I/O rapidly.

Times have changed. Now thanks to flash assisted and flash only storage devices, the storage system can typically respond instantly to I/O requests. Since these devices are solid state, there is very little latency introduced when the devices are accessed. Unlike HDDs, with SSD there are no platters to move into place to find data before each request can be served. Flash devices can handle a continuous stream of I/O, even when those streams are random in nature.

The result is that the compute layer, can generate massive amounts of random I/O. At the same time, the storage layer is able to respond to all that I/O instantly with almost no latency. The problem is that the network layer or the storage interconnect, falls behind creating an I/O deficit that if goes uncorrected, will “bankrupt” the investments made in the compute and storage layers.

The I/O Deficit

One of the single biggest challenges facing both the scaling of traditional workloads as well as the continued investment in new IT initiatives are the capabilities of the connection. As mentioned above, almost every data center has access to an unprecedented level of compute performance. From the storage media perspective, flash storage provides almost instantaneous response time and is rapidly replacing hard disk storage in many performance-demanding environments.

Legacy network connections, whether they be 2/4/8Gbps Fibre Channel or 1Gb/s Ethernet(1GbE), are simply too slow to deliver the I/O performance that these applications demand. A move to modern architectures like 10GbE or 16Gb Fibre Channel (16GFC), is a requirement for modern application environments to reach their full potential. The challenge for today’s IT planner is determining which of these infrastructures is most ideal for the next generation data center.

It’s About More than Bandwidth

When trying to decide between the various storage networking topologies and bandwidth speeds available, it is important to understand that selection should be about more than just performance. Certainly 16GFC Fibre Channel has an obvious numerical bandwidth advantage over 10GbE but as explained in the next section it also does a better job at sustaining that bandwidth in real-world situations.

The Lossless Value

Transmitting storage traffic is not the same as sending an email message. Storage is generally much larger and more sustained; also storage traffic does not typically need to be transferred across a country or an ocean. Instead, the server causing the traffic is in the same data center, usually a few meters away.

As a result, the storage network can be and needs to be lossless with low latencies so that communication is guaranteed, reliable and fast. A lossless transmission is one where each data packet is deterministic, there are no dropped packets or retries required. A lossless, low latency transmission does mean that the physical storage needs to be closer to the attaching servers and within the same data center. As mentioned above, that is typically the case.

The advantage of a lossless infrastructure is that it allows for a low latency connection that provides the predictable performance that modern data center applications require. Without a lossless feature, the network may need to re-try transmissions repeatedly causing unpredictable application performance especially when under load. The unpredictability when multiplied across thousands of users and hundreds of virtual machines could result in user or customer dissatisfaction and even loss of revenue as customers leave your site to go to another. While there are ways to ensure storage over IP is lossless, Fibre Channel is innately lossless.

High Bandwidth Needs Controls

One of the challenges with designing a high bandwidth architecture is making sure the right applications get priority access to I/O resources. It is also important that those applications don’t consume so many resources that the other applications in the infrastructure are “starved out”.

For a well-designed storage architecture to assure consistent application performance it has to have enough raw bandwidth to be able to be provisioned and prioritized. In other words, the “bandwidth pie” has to be large enough to slice up before it makes sense to provision it.

Once there is enough bandwidth, the network then has to be prioritized or provisioned as such that critical applications get the performance they need whenever they need it. This capability is commonly called, Quality of Service (QoS). QoS allows for critical applications to participate in a shared infrastructure since they can be guaranteed a minimum level of performance.

Server and desktop virtualization makes guaranteeing application performance via QoS more challenging since the application is abstracted from the physical system. Some IP infrastructures are limited to just provisioning bandwidth at the physical host layer or they require special interface cards, software drivers or switches to provision bandwidth. Often these different components can’t automatically coordinate the provisioning policies between each other, especially if there is a mixture of vendor hardware being utilized. This adds to expenses and decreased flexibility in vendor selection. What is needed is a standardized way for the storage network to identify specific virtual machines and provision bandwidth to them as needed.

FC architectures can provision bandwidth in an end-to- end fashion so that bandwidth can be provisioned or guaranteed at the virtual machine level. The FC network then coordinates that bandwidth utilization throughout the infrastructure regardless of the HBA adapter, switch and storage system combinations in use. This includes the ability to have that provisioning guarantee follow the virtual machine even when it is migrated to another host.

High Bandwidth Needs Full Utilization

As bandwidth increases the links connecting hosts to switches and switches to storage need to be used effectively. Ideally this is done via a flat simple network that can be easily scaled. In all generations of FC, all paths or network connections are active without any sophisticated storage network design required. The ability to use all paths and do so simply is a critical factor and advantage that Fibre Channel has when a high performance storage network is implemented.

Fibre Channel - The Purpose Built Storage Network

The majority of storage networks are either Fibre Channel (4GFC/8GFC) or 1GbE based on iSCSI or NFS. As customers begin to evaluate options for the next generation of applications and storage, they need a network that is high performance but delivers that performance in a reliable and consistent manner. The latest generation of FC, is ideal for these modern workloads. It is a dedicated network, purpose built for storage. Other storage networks are more general purpose in nature, servicing not only storage traffic but other traffic like messaging that can consume bandwidth and cause performance unpredictability. IP networks were originally designed for very small packet transfers across long distances. Storage traffic is obviously different.

Deploying End-To-End Fibre Channel Connectivity Today

Other storage networks are typically a series of conversions between the layers of the data center. Fibre Channel on the other hand is based on SCSI end-to-end. This starts with FC host bus adapters (HBAs) that provide connectivity to FC switches and then to native FC arrays. This allows the creation of an end-to-end SCSI connection based on the Fibre Channel protocol. It is seamless and efficient requiring no conversion, transmission retries or routing. This creates an ideal infrastructure for flash based storage systems allowing them to reach their full potential.

Enterprise Proven, Backward/Forward Compatibility Assured

Fibre Channel is already the primary storage infrastructure for the majority of enterprise workloads. VMware continually reports that FC remains the dominant environment in its installed base. As with all generations of Fibre Channel, the latest FC protocol- 16GFC is backward compatible two prior generations of FC (4GFC and 8GFC).

This allows 16GFC to be implemented for those performance demanding use cases now, while integrating seamlessly into the existing storage network. Unlike other storage protocols, the FC standard has always guaranteed backward/forward compatibility to assure investment protection and compatibility across the datacenter with streamlined, common management tools enabling multiple generations of FC products to work together.

The Use Cases For The Latest Generation of Fibre Channel

Fibre Channel For Flash Deployments

Fibre Channel should be seriously considered for any storage network that is implementing flash arrays; especially All-Flash Arrays that count on very low latency network connection. With 16 GFC performance, the All-Flash Array can service not only storage data but it can provide virtual memory resources thanks to the high performance low latency nature of the lossless network. This can lower the cost of the server DRAM investment, provide DRAM allocation beyond what is actually available in the server and potentially reduce or eliminate the need for server side flash.

Fibre Channel For Virtualization

16 GFC performance should be considered for any virtual server or desktop environment that is trying to achieve new levels of virtual machine (VM) density. The more virtual servers and desktops that server hosts can support, the lower the cost of the virtual environment. In some cases, host investment can be reduced by as much as 2/3's. But hosts, densely packed with VMs, generate a significant amount of random I/O and they do so continuously.

The random I/O of densely packed virtual hosts places a significant load on every component of the storage network from the HBAs to the switches to the host. However, 16GFC provides the lossless connection and guaranteed performance needed to respond to this demand. The lossless nature of FC allows that response to be done with a minimal amount of latency so that none of the available bandwidth goes to waste. The fabric based, “all paths active” approach of Fibre Channel allows for each investment in host bus adapters and ports to be fully utilized. Finally, the ability to use standardized capabilities like NPIV to fine tune how that bandwidth is utilized, allows for the safe virtualization of performance demanding mission critical workloads.

Fibre Channel For Scalable Applications

Applications need to scale to support more data and more users all the while providing faster response. When an application can’t scale because of storage networking bottlenecks, it needs to be somehow divided up either through clustering or sharing. This adds to the complexity of the application design and the design of the supporting architecture. It also adds cost because additional servers have to be purchased.

From a storage perspective, if flash is not leveraged, large application environments will typically have to deploy hundreds of hard drives across multiples storage systems to provide the performance needed for these applications. These environments will also typically implement server-side flash resources to augment performance. In both cases, these storage workarounds add complexity and cost.

16Gb FC allows mission critical applications to scale further in much the same way it enables the deployment of dense VM environments. Application end-users are in many respects the same as Virtual Machines. There are many of them, they all generate random workloads in the form of queries and file modifications, and they have the expectation of instant response times. Unlike a VM, however, they may have the option of using a competitive service. In essence, there may be a greater risk of revenue loss. With 16Gb FC, the network now has the ability to respond at the speed of flash allowing for far more scalable database and application servers, as well as more scalable storage systems.

Improved Big Data

As mentioned above, Big Data isn't really big data; it is often the scanning of billions of little pieces of data to come up with an answer. Whether that is a decision support system or a real-time personalization initiative, the faster the requests for answers can be satisfied, the better. These clusters are generating millions of requests across billions of files; unlike virtualization and applications, these requests tend to be a continuous sequential I/O pattern. In this situation, raw bandwidth is key. But the servers that manage these requests can be handling dozens of simultaneous requests, leading to a randomized, sequential workload. 16Gb FC with its 16Gb/s transfer capabilities and 100% active connections is able to meet this unique challenge.

Roadmap to 128GFC and Beyond

Storage infrastructures need to continue to evolve and keep up with advances in the compute and storage layers. There is little question that the compute layer will continue to become more powerful each year. At the same time, storage suppliers are already talking about the next generation of flash storage that will use a memory interface instead of a PCIe or SAS interface. Furthermore, memory based storage will eventually become available and will feature DRAM like speeds but with the non-volatility of flash. These advances will lead to a whole new generation of applications that are still in the imagination of developers, some yet unborn.

The next great application, however, will need an advancing network that can respond to applications and computers that are so fast they seem to think for themselves. FC is demonstrating the ability to make sure the advances in compute can be connected to the advances in storage. Already Gen 6 FC -this name was created to move away from speed-based naming to technology generation-based naming. Gen 6 Fibre Channel is based on the 6th generation (1, 2, 4, 8, 16 Gbit/s) with 32GFC being the sixth generation- and 128GFC (four lanes of serial 32GFC) on the horizon and the FCIA (Fibre Channel Industry Association) FC roadmap establishes speeds for 64GFC speeds and beyond.

Conclusion

The modern data center has evolved. Now thanks to virtualization and highly scalable database applications, the full potential of the compute layer can be tapped. At the same time, in terms of performance potential, the storage layer has gone from worst to first (thanks to flash based storage devices). The remaining holdout has been the slow to evolve storage infrastructure. But that infrastructure needs to do more than become faster, it needs to become faster with greater efficiency so that the full potential of the compute and storage layers can now be realized. The latest generation of Fibre Channel is an ideal infrastructure for this modernizing data center and should be seriously considered so that new levels of density and performance can be realized.