



white paper

Best Practices for Implementing iSCSI Storage in a Virtual Server Environment

Server virtualization is becoming a no-brainer for any business that runs more than one application on servers. Nowadays, a low-end server is 64-bit capable and comes with at least 8GB of memory. Without virtualization, most such servers cruise along at 5 percent of CPU capacity with gigabytes of free memory and some I/O bandwidth to spare. Virtualization helps you better utilize these resources.

Not only do you make more effective use of the hardware while using less power and rack space, you get the ability to move virtual machines between physical servers for added flexibility and resilience. All told, you do more with less and gain flexibility and improve uptime. It's a big win.

It was not so easy on the storage side of server virtualization, at least until the advent of iSCSI SANs. A centralized storage pool shared by multiple physical servers is required to make sure you can quickly rehost virtual machines if the physical server they were guests on goes down. But the Fibre Channel technology first available to provide shared storage for virtualization was expensive to acquire and manage. Not so with iSCSI SANs, according to a VMware white paper¹, which says the benefits of iSCSI for server virtualization include "enterprise-class storage at a fraction of the price of a Fibre Channel SAN" and leverage of your existing Ethernet infrastructure "without the need for additional hardware or IT staff."

VMware goes on to note that IT departments "unable to afford a Fibre Channel SAN will now be able to take advantage of advanced VMware infrastructure functionality" including redundant storage paths for resiliency, centralized storage for greater space utilization, live migration of running virtual machines, plus dynamic allocation and balancing of computing capacity.

That's not to say that all iSCSI SANs are created equal. Or that you don't need to carefully plan and manage this high value storage environment.

This white paper covers many of the best practices for implementing iSCSI storage in order to optimize the resiliency, flexibility and performance of your virtual server environment.

business

¹"Configuring iSCSI in a VMware ESX Server Environment," VMware 2008.

Begin with an enterprise-class iSCSI array

When all of your mission critical application workloads and information assets are consolidated onto a storage platform, the need for high-performance, non-disruptive scalability, and continuous availability from that storage increases. A data-center-class, high availability SAN needs to be completely resilient through data protection provided by RAID or mirroring, redundancy of power supplies and cache batteries, and, most important, dual controllers.

A dual-controller design not only achieves high availability by providing redundant paths between the servers and the data they need to access, it doubles the storage processing power of the array.

One best practice is to begin with one controller active and the other standing by in case of failure. As the demands of your virtualized applications grow, you can make both controllers active to share the load and still enjoy the resiliency of failover. When using an active/passive iSCSI target, use a preferred path failover policy. This makes servers use the preferred path whenever it is available, fail over when it's not, and fail back to the preferred controller when it comes online again. When using an active/active iSCSI target, you should use MRU (most recently used) failover.

Note that VMware's virtual infrastructure is providing this multipathing function through its storage abstraction. Other necessary storage functions including storage snapshots and some that might provide benefits such as thin provisioning are also provided. (Thin-provisioned volumes aren't automatically striped over a large number of disk spindles, cutting I/O performance which is so critical in many virtual server environments.) However, VMware does not do it all as far as storage is concerned.

Go beyond VMware's storage management

Certainly VMware's VirtualCenter infrastructure management system does what it can to maintain desired performance levels across workloads of varying priorities running on shared servers. VMware's Distributed Resource Scheduler granularly controls sharing of the host server's CPU and memory, for example. But effectively managing disk access performance levels is beyond the standard virtual infrastructure system. To map appropriate levels of storage quality of service, your storage system must differentiate and manage performance across different virtual server workloads. "Performance must be managed in the storage controllers," says a report on VMware virtual environments from IMEX Research.² Ideally, the virtual infrastructure administrator can manage storage quality of service as easily and non-disruptively as she administers CPU and memory QoS for virtual servers. However, traditional arrays require managing performance at the physical disk RAID level, making it neither easy nor non-disruptive.

Performance Considerations. Let's step back for a moment to look at why the aggregation of so many virtual server loads into one centralized storage facility increases the need for high performance storage and the flexibility to tune that performance easily.

In addition to the compounding of I/O and bandwidth demands when many virtual server workloads on multiple physical servers access an iSCSI SAN, more performance stress comes from the scarcity of memory in many virtualized servers, which causes more paging to disk and further impact on application performance.

An essential best practice is to know what aggregated performance demand your virtualized servers will put on your storage. How to know you haven't exceeded the total bandwidth throughput and I/O capacity? It's best to measure throughput and I/O of each individual application when the applications are still in a physical environment, looking for both average and peak needs for several days prior to moving them to a virtualized environment.

The bottom line is that a high-performance platform is critical for supporting a number of virtualized servers. Look for storage arrays that provide more I/Os per second of processing power and are capable of driving the host network ports at their full line speed. As you move to higher speed network cores and more virtualized servers or additional applications using the SAN, you are also going to need a huge data pipe out the back end. Whether it's a number of 1-Gigabit Ethernet host data ports or a single 10-Gigabit port depends on the particulars of your virtualization load, but a big data pipe is critical.

Storage capacity is a secondary concern in most cases. Capacity will be sufficiently scalable in this class of SAN array for the most part, particularly if the array supports daisy chaining of additional disk cabinets.

How finely to manage performance. VMware automates setting up storage for virtual servers in several ways. The default for the VMware host ESX server is to work with your SAN's volumes as VMFS (Virtual Machine File System) Datastores, each of which will contain the virtual disks of multiple virtual machines. This is a completely automatic process left to VMware. Unfortunately, while pooling storage resources in this way may increase utilization and simplify management, it can lead to contention for storage that causes serious performance issues.

Best practices for VMFS datastores. You can certainly use a shared VMFS volume for virtual disks that have light I/O requirements, but the best practice with VMFS Datastores is to aggregate your performance measurements and map virtual servers to individual VMFS Datastores on volumes configured for that particular mix of applications. Tuning storage arrays for this approach is a matter of monitoring the specific performance statistics (such as I/O operations per second, blocks per second, and response time) and making sure to spread the workload across all the storage controllers.

The disks of any virtual machine with heavy I/O requirements should at the very least be on a dedicated VMFS volume to avoid excessive disk contention. To provide the storage controller granularity and control recommended by IMEX and other experts, you can use an RDM (raw device mapping) or directly connect to a volume on the SAN via an iSCSI initiator running in the guest O/S.

Advantages of SAN arrays that support virtual volumes. If you use the right iSCSI array environment, intelligence in the array can set the quality of service requirements as you create SAN volumes for each virtual machine. The system assigns RAID type, striping level and block size appropriate for the class of application that particular virtual machine will run. Traditional approaches require you to define volumes as sets of exactly matched proprietary drives and trap unused space in these inflexible array groups. You are stuck with volumes absolutely dedicated to one level of storage quality of service appropriate for one class of application. Such an approach is problematic even when letting VMware assign multiple VMFS Datastores to a volume.

Instead, look for systems that handle volumes in a way that allows each drive to participate flexibly in array groups. One impact of this “virtual volume” approach is that such systems can utilize unused capacity available in any drive. It also becomes possible to spread the volume at any point in time across as many drives as needed, tapping the I/O from each additional spindle until you’ve satisfied performance requirements. Most important, this flexible approach allows expansion, movement or reconfiguration of virtual volumes any time your virtual machines require. Storage therefore achieves the kind of flexibility virtualized server environments get from being able to move virtual machines on the fly: you can balance the storage load and expand overall capacity without shutting down mission-critical applications.

You are probably aware that setting up dozens of volumes for virtual servers in a traditional SAN environment takes considerable amounts of time. Therefore, if you go this route, make sure your array vendor has the capability to simplify and accelerate the volume creation process as just described.

Choosing a Vendor

IT professionals looking to assemble and manage virtual storage architectures are looking for products that deliver the performance, flexibility, scalability, cost-efficiency and security necessary to make server virtualization deliver the promised benefits. IT departments considering solutions from suppliers such as HP and Dell have to wrestle with some tough choices. For budget-constrained buyers, the older systems these vendors offer lack the performance and scalability necessary to work well in a virtual setting. To satisfy performance requirements, newer systems from Dell and HP are pricey and require significant management resources.

D-Link’s new DSN-5000 Series of SAN solutions are a good fit for IT organizations looking to virtualize their storage with performance, scalability, fault-tolerance and manageability. . . .all in an affordable package.

The DSN-5000 provides advanced volume virtualization that increases virtual machine performance in a number of ways including increasing spindle count to hit desired performance, teaming ports without denying bandwidth to virtual servers, and providing industry-leading bandwidth with 850 MB/Sec of bandwidth in its 8x1GbE host port configuration and 1160 MB/Sec in its 10GbE host port configuration. This amounts to approximately 4 times the performance of competitively priced systems.

For situations where D-Link’s DSN-5000 product line will be deployed, HP and Dell offer two very different classes of solutions. At the low end of its portfolio, Dell offers its legacy system, the MD3000i, while it also offers a higher-performance solution from EqualLogic, a company Dell acquired in 2008. HP employs a similar approach to its solutions, offering the MSA-2012i for price-sensitive customers but newer solutions from LeftHand Networks, a company HP recently acquired.

Dell and HP are saddled with difficult challenges. Their low- and high-end solutions are based upon different, incompatible architectures, making migration a difficult and potentially costly task for customers who may buy in at a low price point but look to upgrade their functionality and performance with newer offerings. D-Link, by contrast, has designed in ease of migration by basing its 5000 family offerings on the same architecture, controller technology and industry standards.

Dell and HP's lack of a consistent storage solutions architecture present some tough choices for IT decision-makers. Scalability and migration become much more difficult and costly, and it suggests the probability that IT staffs will have to manage mixed-architecture storage deployments for their virtual infrastructure. Sooner or later, the legacy platforms for Dell and HP will have to give way to the EqualLogic and LeftHand Networks solutions, which offer better performance and more flexible design than the legacy systems, but are significantly more expensive than D-Link's offerings for what is often less performance. For instance, HP's MSA-2012i is a solution based on technology originally developed more than 15 years ago that has passed through two successive acquisitions by HP.

A side-by-side features comparison clearly points out the advantage enjoyed by D-Link's 5000 family: support for non-proprietary disk drives, more host ports, more controller memory, greater bandwidth, significantly greater host connections, greater flexibility through volume virtualization, support for more hard drives and greater overall capacity – all at comparable pricing.

When you dive into the virtualization process, the legacy approaches of Dell and HP quickly fall away. That's because older architectures like these simply don't scale as well as do D-Link's current designs; supporting virtual servers chokes network performance in these older architectures. So, compared with Dell and HP's legacy solutions, D-Link's DSN-5000 family avoids the performance and scalability hurdles.

When compared with the newer configurations of HP and Dell, D-Link's solutions offer comparable or better functionality, but at considerably lower price points. D-Link still offers advantages over the competitive products on host ports, controller memory, bandwidth, and its compelling ability to virtualize volume creation to ease scaling and tuning arrays for the quality of storage service required, all at price points generally 30 to 60 percent lower than those of HP and Dell.

Finally, D-Link iSCSI-based SANs have been tested and verified to work with the leading virtualization software platforms, including VMware and Citrix XenServer. This gives IT organizations confidence that the underlying storage platform in their virtual environments will work seamlessly and reliably while providing the processing performance needed from your virtual infrastructure.

Call or visit the web address below for more information
on the complete line of business products:

www.dlink.com | www.dlink.ca
888.331.8686 | Email: getmore@dlink.com

business