

# STORAGE SWITZERLAND REPORT

## FOR DATA CENTER VIRTUALIZATION, IOPS IS MORE IMPORTANT THAN AIR



George Crump, Senior Analyst

IOPS, or Input Output operations Per Second, are one of the key barometers to how well your storage is performing. It gives you a good idea of how quickly a storage request from an application or server can be fulfilled by the storage system. Typically the lower the IOPS the longer the requesting server has to wait. High IOPS is important for specific high transaction applications or concurrent/aggregate workloads like those found in the simulation and design fields.

Until the widespread adoption of server virtualization, high IOPS haven't typically been a requirement for the more typical data center 'workhorse' applications. The combination of desperate virtual servers creates a highly random I/O pattern from each host. Adding additional VMs to a host result in a exponential increase in random I/O operations per second. The result is that every server in the environment is a potential IOPS consumer meaning that storage system IOPS is more important than air.

### The Cause Of I/O Demands

In the past every application had its own physical hardware, which also meant its own network and storage

I/O channels. Even if the storage that the server was attached to was shared, a specific section of that storage was zoned off and allocated to that particular server. The disks themselves were typically allocated to only a few actual servers. While many data centers have a small number of their overall server populations that need high-end storage performance, the average IOPS required per server is generally low. In the traditional, single host - single application configuration it's also easy to identify and isolate the servers that need high storage throughput.

The challenge, of course, to single workload servers is low CPU utilization. This model becomes expensive, wasting resources which also wastes power, cooling and data center floor space. Server virtualization allows each application workload to run on its own virtual server, all on a single physical host server so that resource utilization can increase. While each of these virtual servers are protected from each other they do share the same pool of resources and while CPU resources and I/O bandwidth are plentiful, storage IOPS are not. The speed at which the physical storage system can respond to storage requests, the storage IOPS, becomes the choke point for performance in virtual server infrastructures.

The requestor for each individual storage I/O may be a simple virtual server with a moderate resource consuming application. When these workloads are multiplied by a dozen or more virtual machines (VM) the single physical host becomes a very demanding consumer of storage I/O. For example a common VMware rule of thumb is 50 IOPS per VM, meaning that a host with 15 VMs will need 750 IOPS. However, real world implementations have found the number of IOPS per VM to be closer to 150, bringing the I/O demand per physical host closer to 2,250 IOPS. Additionally more virtual infrastructures are pushing the density of VMs per host well beyond 15, 25 is not uncommon, bringing the IOPS per host up to 3,750. Assuming a medium sized virtual infrastructure of 50 physical hosts, the total IOPS for the virtual environment alone can reach 187,000.

Remember, this IOPS number is a *sustained* level of I/O demand and I/O capacity to handle an unpredictable *peak* 'I/O storm' must be factored into the design of the system. Also, add to this IOPS baseline those unique servers mentioned above that are not virtualized because of their own individual high IOPS needs. Since it's ideal to have all the hosts connected to a single storage system, this virtual environment, in which each VM runs applications with 'moderate' performance requirements just became a very demanding storage infrastructure.

IOPS is not only important when considering workload performance, it's also important when considering data protection performance. Copying VMs from primary storage to a secondary device will consume IOPS as well, since backups of VM images are often the equivalent of backing up hundreds of very large files.

#### Why Are IOPS Critical To Server Virtualization?

The level of IOPS that can be delivered from the storage system directly impacts the density that can be achieved per host and most virtualization hosts are not CPU bound,

they're I/O bound. The higher the density the more cost effective the server virtualization environment becomes. Since the available IOPS from the storage infrastructure directly impacts performance of the VM it directly impacts VM density. In 2010 and 2011 a key initiative in virtual environments is to drive up VM density per host, something that's difficult without enough IOPS on the storage platform.

With inadequate IOPS the other cost that the virtual environment forces on IT is time. The lack of IOPS can be somewhat worked around with very detailed resource planning and constant monitoring. The problem is that this takes time and tools to monitor and forecast where VMs should be placed, delaying VM deployment and detailed storage system design. When a rebalance of VMs on physical hosts needs to take place to optimize resource utilization the migration effort not only impacts the VM being moved but also other VMs sharing the same connections.

If a storage infrastructure could provide all the needed IOPS then VM densities would increase and management time would be lessened, providing further cost savings for the virtual infrastructure. The challenge is that traditional storage systems don't provide the needed level of IOPS.

#### Impact On Traditional Storage System

The response to IOPS shortcomings in the virtualized server environment is to configure traditional storage systems with faster drives and to increase drive spindle count or in some cases, even add solid state drives. While this will help improve IOPS at the mechanical drive level it does so at the added cost of power and data center floor space. Eventually though the real bottleneck becomes the storage controller itself. In a large virtual infrastructure the storage IOPS demand will soon outpace the ability of controllers and connecting I/O bandwidth of the storage system to deliver the needed throughput.

Eventually almost every storage system manufacturer is forced into (or looks forward to) selling the customer additional systems. While adding two or three total storage systems to the environment does add somewhat to the management burden, many mid-sized and larger environments need five, ten or even more individual storage systems to support the total IOPS demand caused by server virtualization. This not only impacts ROI since storage is a high percentage of the overall infrastructure cost, it also significantly increases storage management activities as it relates to the virtualized environment. For example, provisioning a new VM now means that every storage system and every LUN or volume within that system needs to be inspected to identify which is the best candidate for responding to that VM's IOPS needs. This also means that careful planning of the proposed VM must take place to understand what its exact requirements are going to be. An error in either calculation not only impacts the performance of that VM but also the performance of the VMs that share the storage and server environment it's on.

While resource planning in the virtualized environment certainly has value, it's not an exact science. There are simply too many variables for the administrators that manage the environment and perhaps too many different tools to help them keep track. The ideal solution would be to have a storage system that's pre-built to handle the bulk of the IOPS load upfront and then can be scaled as needed.

The good news is that the IOPs deficit can be closed. Storage systems today are available to deliver plenty of IOPs upfront and into the future. For example platforms like [BlueArc's](#) Titan and Mercury storage systems provide the ease of access for the virtual environment via NAS and do so within a single system that can deliver the IOPS demand that the virtual environment needs. They have consistently proven to be one of the leaders in SPECFS, a benchmark that compares IOPS performance and have achieved these ranking without requiring high spindle counts. Having a single system that can sustain the IOPS workload for a virtual infrastructure is critical to being able to broaden its use. With IOPS demands fulfilled VM densities can increase, more demanding stand alone servers can be virtualized and if the storage I/O can be delivered in a single or finite number of systems, management time can be reduced.

### **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>.**