

# STORAGE SWITZERLAND

## STORAGE I/O MANAGEMENT IN VIRTUALIZED ENVIRONMENTS



George Crump, Senior Analyst

In essence, server virtualization takes a group of servers that have modest storage I/O requirements and then consolidates them into a few virtualized hosts. Those servers, because of their minimal requirements, were allowing CPU, Memory, Storage I/O and Network I/O resources go unused, which made them a prime target for server virtualization in the first place. The challenge however, is that the I/O those modest servers generated, when combined onto a single host, often created a ‘resource consuming monster’. While the new CPU and memory problems that result can be easily identified, I/O, and in particular storage I/O, can become a management nightmare. Consequently, storage I/O management in virtualized environments is turning out to be a key skill for IT personnel to develop.

### **Troubleshoot first**

Taming the storage I/O monster in virtualized environments requires leveraging the appropriate tools to develop a storage I/O management skill set. While the “by the book” first step is typically to plan out what the storage I/O requirements of the virtual environment will be, now and in

the future, this is not reality, nor is it practical. First, planning, especially in a virtual environment, is hard because there are so many variables and so many layers of abstraction. Planning is more practical and more accurate if it can be done after an initial baseline of storage demands are measured.

Almost predictably though, a storage I/O problem occurs while the baseline is being developed; and fixing the problem generally takes priority over developing a great storage I/O plan. In most environments storage systems and their I/O requirements don’t get attention until they become a problem. The thought of planning for storage I/O rarely warrants serious discussion until after a problem has been resolved - basically a ‘how do we make sure that doesn’t happen again’ approach.

Storage I/O problems usually present themselves when something in the environment, typically an application, suddenly becomes slower. Usually storage performance is assumed to be the culprit, and often is. As a result, troubleshooting is frequently the first reason that IT professionals start looking for assistance to help with the diagnosis.

Before server virtualization this process often meant touching each server, logging in and running built-in system utilities that would help the storage manager to find the cause of the problem. While not necessarily the best use of time, it was feasible to diagnose storage problems this way. The server typically had a one dimensional relationship with storage, like a single server being attached to a single storage LUN. Even in a shared storage environment where several servers could be attached to the same LUN, the population was usually small enough that the storage manager could get the job done, although more difficult and not recommended.

The virtualized environment makes manual diagnosis an impossibility. Not only are multiple virtual servers' I/O demands coming from a single physical server, the LUN supporting that host server is often shared across dozens of physical servers, each with their own set of virtual machines. Storage I/O problems can come from many different places including:

- Too many VMs could be accessing the same physical disk area
- A single VM may be too performance demanding, stealing I/O resources
- Storage switch interconnectivity
- Storage controllers being overwhelmed by I/O

The first step in diagnosing a storage problem is to isolate exactly what the problem is. This is going to require end-to-end visibility and the use of modern tools like [SolarWinds Profiler](#) products to present the storage manager with various views, and allow for faster diagnosis and correction. The environment needs to be seen from a storage view 'up to the server' or a physical system or virtual machine view 'down to storage'. The goal is to be able to see which virtual machines storage requests are going through, which network ports and which LUNs they're assigned to - and even the physical disks.

With a "disk up" view for example, the storage manager could see which virtual machines have data stores on that specific disk array and then monitor just those virtual machines. They could then re-allocate those virtual machines to other physical systems or move their virtual images to another LUN. In both cases the load would be lightened. If the LUN is not the problem the VMs could be monitored for the amount of I/O bandwidth they are consuming coming out of the physical server. If it's too high, the virtual machines with the heaviest I/O could be moved to separate physical hosts or the decision could be made to upgrade the I/O capabilities of that server.

## Monitoring

With the storage I/O problems under control the tool selected to help with the diagnosis should now be able to shift into a monitoring mode. A tool that is either so expensive or intrusive that it can only be justified in case of emergency is probably not the best solution. I/O management software should be easy enough to use and provide dashboard-like information so that it's effective at monitoring the environment as well as troubleshooting it.

The value of a good monitoring platform is that it allows for the prediction of problems before they occur, or at least before they impact business processes. Virtual machines can be monitored to track and trend increases in I/O need as their use evolves. Physical hosts can likewise be tracked and trended to see the overall I/O consumption rate as new virtual machines are added to the environment. This provides a macro host view as well as a specific virtual machine view. Both variables need to be known in order to maximize the environment. Storage can be monitored and trended to see when the capacity or I/O capabilities of a particular LUN reach its maximum.

Even the storage controller itself can and should be monitored to see if it's being overwhelmed by a growth in virtual machine activity. The storage controller is becoming the new bottleneck in storage performance for many environments, but virtualization products like VMware provide the tools to rebalance that storage workload. Storage VMotion allows the live movement of the virtual machine data store to an entirely separate storage system and controller. Leveraging that capability requires knowing which VMs to move when the first storage system becomes overwhelmed with storage traffic.

One of the keys to effective monitoring is finding a tool that does not try to squeeze all the information onto a single screen. Ideally it should be "dashboarded" in such a way that a quick visual inspection will provide a 'red - yellow - green', view of the environment; allowing a drill down on yellow alerts, for example, before they become red.

## **Planning**

With the troubleshooting work of storage I/O problems out of the way and the capability to predict upcoming problems handled with monitoring, the final step is to create a plan that will allow for the forecasting of storage I/O and capacity needs. To help with planning, the storage I/O management software application needs to be able to collect historical data, trend that data and then project where growth is heading. Most importantly, that data has to be presented in a clear and concise manner that doesn't require hours of spreadsheet manipulation to present the data in its final form to management.

By putting planning last, the storage manager is capturing a more accurate view of storage I/O in the virtualized environment. The baseline data is real, not an estimate, and as new virtual machines are added they can be factored into the existing baseline. Planning is never perfect, especially in the virtual environment, and this is the reason why a planning-only tool does not make a lot of sense. A troubleshoot - monitor - plan cycle is then replaced by a plan - monitor - troubleshoot (less often) cycle. A tool that can provide assistance in each of these aspects is critical to obtaining maximum storage I/O performance with minimal risk.

### **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>.**