


STORAGE SWITZERLAND

ADDRESSING STORAGE PERFORMANCE CHALLENGES IN THE VIRTUAL SERVER ENVIRONMENT



Eric Slack, Senior Analyst

Server virtualization is causing performance problems for users, especially within the storage infrastructure. The complexity and abstraction with the way server virtualization deals with storage has created a new resource problem, storage IOPS, and rendered the existing management tools less effective. In lieu of simply adding more storage, we need to be smarter with the way we address storage performance challenges, which requires cross domain information and better tools to monitor and manage this dynamic environment.

How Did We Get Here?

Companies in the open systems era started with direct attached storage, essentially a 'shared nothing' environment. As IT needed scalability and flexibility to cope with data growth, SANs allowed users to 'share capacity' from a central point to multiple servers. But the LUN masking and switch zoning that SANs employed basically divided up a large pile of storage capacity into smaller chunks (LUNs) connected to a single, dedicated host. From an administration perspective, storage was easy to manage, as the environments were relatively static. Storage capacity was virtualized, but storage I/O wasn't,

as SANs essentially 'hard allocated' storage to host servers.

Shared Everything

Server virtualization changed all of that. While virtualization environments do use SANs to consolidate storage resources, there are major differences around managing storage in a virtual server environment compared to a traditional SAN environment of physical servers. Virtual machines (VMs) share a host's access to storage IOPS and storage network bandwidth. Physical hosts typically share access to the same storage area on the same storage network and, thanks to VM migration, they can at times share VMs.


To improve utilization, virtual server platforms use an overcommit strategy when leveraging resources. They allow the allocation of more CPU, memory, storage capacity and storage network bandwidth than physically available in order to facilitate the sharing of these resources among VMs. Similar to the thin provisioning rationale, the assumption is made that not all VMs will need the maximum amount of storage, CPU or memory that's been 'reserved' for them at the same time.

While storage does tend to show a slow steady growth that's predictable, a VM may spike the demand for any or all of those resources (CPU, memory, storage) at any moment. And, it's hard to predict these events, at least without some help. The overcommit strategy dramatically improves asset utilization, while it improves flexibility and drives down data center costs. The key is to manage problems caused by these spikes before they degrade performance.

Abstraction

Server virtualization also adds another layer of abstraction to the storage infrastructure beyond the traditional SAN environment. The hypervisor has a virtual storage stack and a virtual network stack between the guest OS on each VM and the LUNs. To be sure, this abstraction is the technology enabling server virtualization to work, but it's also the root of the complexity behind many of the storage performance challenges users have. These virtual NICs, virtual HBAs and virtual switches that live inside each hypervisor can combine to make it very difficult to map VMs to their storage resources and track their activity.

To make matters worse, VM sprawl 'turns up the volume' on the management issues outlined above and increases the chances that a problem will occur. As the results from a recent Storage Switzerland Survey show performance management and troubleshooting were almost 50% of the storage challenges indicated by users.



IOPS

In the traditional SANs, servers would be added to the network and storage capacity would be allocated to them. The required network bandwidth was consistent and predictable since LUNs dedicated to hosts across specific ports. Likewise, the load on the storage, measured in IOPS, was fairly consistent because the number of LUNs and connected hosts were generally static once all the storage was assigned out.

In virtualized environments, the combination of overcommitment, abstraction, and VM sprawl create the 'perfect storm' for hard to diagnose IOPS spikes and contention caused by the dynamic nature of virtualization.

Although the simple solution for performance problems might appear to be to add more storage, throwing money at the problem may not be an option and it can only be done for so long before you run out of data center floor space or budget dollars. Also, adding storage can compound the problem by contributing to VM sprawl and increase traffic on already overloaded networks.

The alternative is to improve efficiency and 'work smarter' to meet performance demands. But working smarter may be easier said than done, since it requires more information about the infrastructure and most of the vendor supplied tools are not up to the task. Instead of focusing on allocation, a tool must focus on *optimization*, with insight into the sources and consumers of those resources and their interrelationships. Also, both real time and historical information is needed to manage dynamic resources, like IOPS and storage network bandwidth.

End-to-End, Top-to-Bottom Visibility

Comprehensive real time information is required to untangle the abstraction of the virtual server environment. You can't examine wiring like you could in the physical data center, and instead need tools that can easily trace connectivity from VMs to storage and back.

What's needed is end-to-end visibility, from VMs, to hosts, to network elements, to storage. Also needed is 'top to bottom' visibility within the virtualization host, from guest OS to VM, to virtual NICs, HBAs and switches, to the physical host server. This removes the abstraction of the virtual storage and virtual network stacks from the management environment.

Solutions like [Storage Profiler from SolarWinds](#) can provide this comprehensive view of storage performance in real time and help resolve performance problems. If an application is running slow, for example, an end-to-end view of storage from the virtualization layer to physical disk would be one place to start. In this scenario such a view would show the VM running the slow application, the datastore supporting it, the host servers and VMs sharing that datastore, the LUNs that make up the datastore and the physical storage systems those LUNs live on. This kind of logical mapping would help you to identify the cause of the performance problem, virtual storage or physical. It could be too many VMs or hosts sharing the datastore or a datastore that's low on storage capacity. You could also drill down to view total IOPS at the LUN or RAID Group level to see if there is LUN contention or if physical storage is the bottleneck.

These end-to-end solutions can provide capacity analysis at both the physical and virtual levels. On the virtual side, they can show capacity and % utilization of storage volumes, hosts and VMs to determine the number of additional VMs the current infrastructure could support, for example. On the physical side, they can show raw storage and usable RAID capacity by the array, LUN and file system, plus the % of storage available to identify a need for additional capacity. And, unlike array-specific tools, they can show all storage systems that are supporting the virtual environment in a single pane of glass.

Trending and Predictive Analysis

Playing 'catch up' with a resource optimization problem isn't wise and tools must show trends and allow managers to prevent resource problems before they occur. The capacity and performance metrics described above can be captured and used to show how resource usage is changing over time, or used to drive a forecasting engine showing what resource levels will be required in the future. This can enable managers to provision physical storage before it's needed and load balance virtual resources to eliminate performance potential performance issues.

To improve resource utilization, LUNs which continually show lower IOPS levels could be good candidates for migration to a slower (and cheaper) storage tier. This ability to see performance utilization of a particular subset of your storage can help you take capacity a step further and indicate what kind of storage you need, not just how much, and where you can reduce costs. Similarly, monitoring file types for access history or age and ownership can help keep unused data on appropriate tiers of storage, or facilitate its deduplication or deletion.

Addressing storage performance challenges in the virtual server environment starts with timely, accurate information from multiple domains. Unfortunately, the tools used to monitor physical servers in SAN environments may not be much help. Server virtualization abstracts and overcommits physical resources to create a more complex environment in which storage IOPS are potentially the most critical storage resource. Working 'smarter' by optimizing resources is a key concept to managing this more challenging environment - but it will require a new tool set, like SolarWinds Storage Profiler, which provides a comprehensive, end-to-end view of storage and other resources.

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 © 2011 Storage Switzerland, Inc. - All rights reserved