

Enhance VMware Performance: 3 Best Practices for Storage & Server Admins

(Or – how to keep the peace between teams)

WHITE PAPER BY BRYAN SEMPLE



Intro

It is no secret that friction many times exists between storage and virtualization teams. Cross discipline friction points are common throughout any enterprise and have sometimes existed between system administrators and the storage team. However, virtualization has increased the measure of complexity an order of magnitude and increased tensions accordingly. With enterprises now shifting from basic file and print virtualization to more advanced mission critical applications, solving this friction point is critical to expanding the virtualized datacenter. Not only must the friction point be solved, the team must unite more than ever to enhance VM performance. This white paper explores three steps to unify the teams based on VKernel's experience working with both system and storage administrators.

Friction Increases due to Virtualization

Prior to virtualization, a storage administrator was more closely involved with the system administrator in day to day operations, planning and deploying applications. With a one to one relationship between a LUN, a host, and an application, close coordination was required.

Virtualization changed all this. The virtualized server environment is now treated as an application itself. Storage teams simply allocate LUNs to a VMware or Hyper-V environment and move on to their next task. The problem is that the storage team is no longer involved in understanding specific application loads.

The disconnect between storage administrators and the underlying applications running in a virtualized environment is the root cause of much of the friction between the teams. Better uniting the teams requires returning the storage team to the planning process for virtualized environments.

Step One – Comprehensive Joint Planning for Application Loads

Today, the amount of joint storage and system administrator planning is mainly around future storage capacity requirements. The challenge here is that storage capacity in the form of GBs is generally not what causes performance problems. The dynamic nature of a virtualized deployment and its impact on data throughput is the primary driver of VM performance issues. This change comes about as:

- VMs are turned on and added to a data store
- VM application loads change
- Neighboring VMs change their load profile and start impacting other VMs
- VMDKs are moved between data stores

The starting point to preventing the dynamic nature of virtualized environments from impacting performance is a well balanced VM load across all resources including storage. To accomplish this, application owners, VM administrators, and storage administrators should review the expected load characteristics of an application prior to deployment with a goal of balancing loads across their

infrastructure. Application owners should characterize their expected application load with a standardized template so VM and storage teams can correctly allocate a VM and its associated data store to the right resource pools and corresponding LUNS.

What about DRS? The challenge with using DRS to solve the load balancing problem is that DRS does not take into account storage I/O bottlenecks. It may move a VM from a host that has plenty of I/O bandwidth to another host that is I/O constrained because of other VMs heavily utilizing the same datastore on a steady state or peak basis. Use of DRS can actually aggravate an already problematic storage situation.

Involving the storage team in application deployment, then returning periodically to review storage loading is the key to maintaining a balanced environment.

Step Two – Reduce Storage Consumption

The relative ease of deploying virtual machines means that storage consumption for virtual environments is actually higher than consumption for non-virtualized servers. This fact alone creates a point of contention for storage teams. Unless the enterprise has a chargeback system implemented, the storage team is forced to constantly ask for additional, potentially out of budget capital expenses for more storage. A higher consumption rate of storage draws scrutiny from the CFO's office leading to tension with the VM team.

To avoid this, a focus on controlling storage consumption is critical. Storage vendors offer de-duplication technology which is very effective in virtualized environments. Beyond de-duplication, zombie VMs, powered off VMs, abandoned images, old snapshots and templates should be proactively uncovered and eliminated to keep storage spending within budget.

Finally, the best way to reduce storage consumption and overall resource consumption is to implement some manner of chargeback or "showback" so the actual cost of VMs can be at least reported back to the business units requesting the resources. Actually charging back is problematic for a variety of systems and business reasons including a lack of information systems to support chargeback and no agreed upon costing for IT services. While chargeback may not be workable, arming the VP of Infrastructure for a budget meeting with information showing the actual virtualization cost a business unit is incurring does go a long way to controlling VM sprawl.

Step Three – Provide the Right Information for Troubleshooting

In the old, non-virtualized world, troubleshooting storage issues for an application was difficult enough. In the virtualized world, this complexity gets worse. The solution, however, is not to provide the VM admin with spindle level information. That level of detail is already present in a variety of tools offered to the storage team. The key is to restore the relationship between the storage admin and the application by detailing exactly which VMs and which datastores are suffering performance problems. Once this is established, sharing changes to the environment over the past 24 to 48 hours with respect

to new VM loads added to a resource pool, changes in VM load characteristics, and any vMotion for storage changes that have occurred will at least level set environment characteristics. Drilling down into VM performance characteristics and comparing information from the storage array should help to narrow down whether the storage issue is caused by storage hardware problems or capacity bottleneck issues at the HBA, fabric, array or spindle. A majority of these issues will be capacity bottlenecks caused by changes to virtualized loads and not connected to specific actions taken in the storage hardware.

Conclusion

Virtualized environments have introduced fresh points of contention between storage and server teams. With proper planning, a focus on containing storage consumption, and more proactive information sharing during troubleshooting, these contention points can be reduced and the full ROI for a virtualized environment can be achieved.