Storage Performance for Virtualized Tier One Applications

John Webster
August, 2011

Technology Insight Series

Evaluator Group

enabling you to make the best storage decisions

The Importance of Storage Performance for Virtualized Applications

The Evaluator Group has observed a common pattern among IT administrators who are managing their growing VMware environments. The lower tier applications are the ones that get moved from physical to virtual first, once they have been tested. These are typically print and file serving, and other non-critical applications. Next comes the more challenging assignment: moving the more critical tier one applications into virtual machines so that as many applications as possible are managed under the VMware umbrella.

We have also seen that the more critical applications running on ESX servers are now placing new demands on the storage environment. Enter vSphere 5. This latest major release of vSphere contains much in the way of new storage-related functionality for VMware administrators. In fact, vSphere 5 contains even more storage-related streamlining standardizing, and automating features that can be integrated with capabilities in the storage array than have all the previous versions combined.

Assuring the Satisfaction of Business Users

When moving applications from physical to virtual, the worst possible outcome is to have user groups experience performance degradation after the migration and become less productive as a result. The experience of the critical business application users must be the same or better. Otherwise, user groups, especially those powerful enough, will demand a reverse migration back to the one-to-one, application-to-server deployment model that can now be seen as terribly inefficient. Additionally, future application migration from physical to virtual will become a "tough sell" when approaching other users groups and ultimately impact the efficiency of the business.

Therefore, the business user's experience at the outset must be at least the same or better than that which users currently experience. On the plus-side, a positive outcome from application server virtualization could represent an opportunity to change the nature of the relationship between centralized IT and particular user groups. Virtualization could be an avenue for establishing a centralized services delivery model backed up by Service Level Agreements[1] (SLAs) where these do not yet exist.

Application server virtualization can ultimately reduce time to market for new business applications and enhance data protection and disaster recovery capabilities. However, the addition of new applications, which is facilitated by VMware, for example, cannot impact the performance of existing applications on virtualized servers and should avoid impact to the CAPEX budget for as long as possible.

Underwriting virtualized application performance, business continuity, and ultimately business user satisfaction is the performance and scalability of the storage environment. And with the enhancements in vSphere 5, VMware now gives IT administrators the choice of running many storage related functions on the ESX host or offloading them to high-performance storage arrays using storage-resident APIs like

[1] IT organization's SLAs are moving towards a Cloud-like offering, where not only performance and uptime is measured, but also speed of delivery and provisioning of utility computing capabilities. Virtualization becomes essential to delivering to these new SLAs. For more information see "Data Center Transformation," by Russ Fellows.

VADP, VAAI, VASA, and others yet to be announced by VMware[2]. Doing so reserves ESX host processing cycles for applications. However, when accepting the offload, the array must be able to handle these functions without impact to I/O performance. In addition, variability in I/O access patterns over the course of a business day typically place greater emphasis the use of array-based load balancing to assure performance across all VMs.

IT Administrative Requirements

IT administrators will need to know what determines both good and bad storage performance. Being able to configure an array and then letting the array dynamically respond to the variations in the demand for storage performance is highly desired. They cannot be tasked with constant readjustment of the VMware storage environment in order to assure performance and application availability.

Evaluator Group has also found the IT administration and management of storage in a VMware environment is a new discipline and requires more collaboration among groups within the IT organization. Without IT organizational change, virtualizing application servers can actually add complexity—a problem that server virtualization is supposed to solve. Storage performance cannot therefore become a constantly reoccurring issue as application performance problems are resolved and new applications are assigned to VMs.

Finally, once having begun the migration of business-critical applications to VMware, IT administrators will require that their preferred storage vendors quickly support all of the storage performance enhancements delivered by VMware on an ongoing basis. We believe this statement will be particularly true with the release of vSphere 5 that adds significantly to the three storage-related performance and efficiency-enhancing APIs in vSphere 4.1, now referred to collectively as VMware vStorage APIs for Array Integration (VAAI).

In this paper, we look at the performance related enhancements delivered in vSphere 5, and look to how Hitachi has integrated its AMS and VSP storage array families with these new capabilities.

vSphere 5 Storage Essentials for Performance

VMware vSphere Storage APIs for Storage Awareness (VASA)

Both medium scale modular and large high performance storage arrays offer a range of high value features and functions. Depending on the application, the mix of the most relevant features required for a particular VM running a high value business application will vary. Array based data replication will be more important for some of these applications while predictable performance or scalability attributes may dominate others.

VMware administrators and storage administrators working in VMware environments typically assign and provision storage features and functions manually to VMs—a time consuming and error prone process. Manual management also requires continual monitoring and optimizing the storage

[2] With the introduction of vSphere 5, it is now clear that VMware intends to leverage through integration new storage technologies as they become available.

environment for each VM as the virtual server environment grows. As a result, a layer of management complexity gets added back that virtualization was supposed to do away with in the first place.

To eliminate these manual and error prone processes, vSphere 5 allows VMware administrators to automate these processes via VASA. VASA's APIs, when implemented by a storage vendor at the array level, can be used to communicate arrays available features and functions to vSphere. IT administrators can automate the assignment to VMs to storage that is optimized for each VM under established management policy. Continual monitoring is also automated. As time goes on, vSphere determines if the storage assigned to a VM is responding adequately to the usually growing needs of the VM. If not, the VM is migrated to available storage that does fit the required profile.

And, the fact that storage can be "profiled" from the standpoint of vSphere is an important concept in this discussion. IT administrators can automate the provisioning of storage to VMs using vSphere 5 Storage Profiles—characteristics of a storage device that is available to a given VM.

There are a growing number management processes that will benefit as a result of VMware's continuing commitment to VMware/storage integration:

Performance Management—A VM will remain on a given class of storage as defined in the profile as long as it meets predetermined performance requirements. If not steps are taken like moving the VM via vMotion to a different class of storage (see Storage DRS bellow) to bring storage for that VM back into compliance with performance requirements.

Storage Distributed Resource Scheduler (DRS)—Storage DRS determines the optimal placement of a VM using vSphere's knowledge (via VASA-derived information) of available disk space and current I/O load (IOPs and I/O latency). DRS determines initial placement of the VM, after which it continually monitors the VM's performance requirements against actual conditions. Other capabilities include the creation of storage clusters for the aggregation of multiple storage devices under Storage DRS as well as creating rules regarding the sharing of storage resources by VMs.

Enhancements to VAAI

VAAI APIs can allow IT administrators to offload processing for certain operations (e.g., provisioning a new VM from a template) from ESX hosts to storage arrays that support the VAAI APIs. Doing so has the immediate benefit of improving overall ESX host cluster performance and freeing up server CPU and memory to support additional VMs. Additionally, VAAI simplifies storage provisioning, which has been one of the greatest prohibiting factors for "virtualizing more" in the data center. VAAI currently supports *block storage arrays only*. VAAI for NAS is not supported with vSphere 5 but will likely be supported with release 5. However, Fibre Channel, iSCSI, and FCoE connections between the ESX server and the storage array are supported.

Specifically, VAAI addresses the following distinct storage performance and efficiency issues. The first three listed here were introduced in VMware vSphere 4.1. The fourth—TP-STUN is new in vSphere 5. These features are currently supported on the Hitachi Virtualized Storage Platform (VSP and VSP VM) and Adaptable Modular Storage (AMS) 2000 series.

> Full Copy leverages existing SCSI commands available in the SCSI command set and sent from the ESX host to the AMS 2000 storage array to clone, snapshot, and migrate data within and AMS 2000 array. These functions are normally used when creating new VMs and will be

particularly useful in Virtual Desktop Infrastructure (VDI) deployments. No ESX host server cycles are consumed in the performance of these copy functions, enhancing overall VMware-based application performance and greatly reducing the time required to create new VMs and virtual desktops.

Hardware-Assisted Locking supports block-level granularity when protecting VMFS metadata. Previously, SCSI reservations were used at the LUN level to satisfy the same requirement. While this API will enhance performance under certain conditions, the real benefit is a gain in the number of VMs that can be supported per storage array in a single volume. This feature makes it possible to support up to 2TB LUNs per data store, which has huge benefits from a performance, administrative and scalability standpoint.

Block Zeroing greatly speeds-up the standard disk initialization process in a VMware environment. This function is particularly useful in reducing the storage administration time required to create virtual disks in eager-zero thick (EZT) format from large SATA disks.

Thin Provisioning-STUN (TP-STUN) TP-STUN sends an alert to the VMware ESX host when a thin provisioning pool is full and temporarily pauses a virtual machine when disk space is exhausted. TP-STUN then allows for the allocation of additional space to the datastore or the migration of an existing virtual machine without resulting in the failure of the virtual machine.

UNMAP reclaims the physical storage blocks occupied by a virtual disk when that virtual disk is deleted. Under UNMAP, vSphere informs the array that a group of physical blocks is once again available after virtual disk deletion. Otherwise the array would continue to "see" these blocks as being used. TP-STUN allows the array to reclaim capacity, thereby enhancing performance and efficiency.

Hitachi's Performance Assurance Capabilities for VMware VMs

In addition to integrating vStorage APIs Hitachi also delivers intelligent array controller technologies to assure application performance under varying workloads and changing user demands. IT administrators ideally want an array to have the "intelligence" built into a storage array to automatically respond to changing I/O load conditions in real time. This becomes a very real requirement when some VMs become more active than others throughout the course of the business day. IT administrators cannot be constantly monitoring, tweaking, and re-tweaking to respond to changing conditions on a daily basis.

Hitachi will support the storage related enhancements in vSphere 5 upon general availability of vSphere 5 and are currently delivering storage systems with support for the original VAAI primitives. However, it should be noted that a virtualized storage system such as the Hitachi VSP that can attach a heterogeneous group of storage arrays and present them as a pooled storage resource to vSphere can extend VASA and VAAI support across the entire storage resource pool. As a result, the VSP has the ability to extend vSphere 5 benefits to its attached arrays, even if those attached arrays don't support APIs like VASA or VAAI natively.

From a VMware administrative perspective, and because all three of the APIs are supported on the Hitachi VSP, these benefits include:

- Increased efficiency when cloning VMs and creating new VMs from templates

- Increased migration performance when using VMware Storage vMotion
- Decreased need to worry about the size of VMware snapshots
- Increased ability to support critical applications on the vSphere platform.

For VSP, Hitachi also offers automated data tiering (Hitachi Dynamic Tiering) that can be integrated with VMware environments. HDT operates at the sub-LUN level (block) or object level (file) using SSD (SLC Flash), SAS, SATA, and virtualized third party storage as separately definable tiers. A maximum of three tiers (tiers 0, 1, and 2) can be configured. Tiers can also be defined by RAID group and drive speed. Volumes may be migrated among tiers without disruption as well as well using Hitachi data migration tools.

To assure performance, data is written to the highest performance tier first as defined by the user. As it becomes less active it is successively demoted to the lower performance tier(s). If activity increases, it will be promoted back to successively higher performance tiers. Movement can be controlled by time of day in cycles of between one and twenty-four hours.

For the AMS 2000 series, Hitachi will also support the vSphere 5 enhancements for storage with general availability of vSphere 5. It offers a number of features that support automated assurance of VMware application performance. These include:

Hitachi Symmetric Active-Active Load Balancing Controllers automatically distribute workloads and I/O across all data paths between VMs and storage to assure predictable performance under continuously variable load conditions and is complimentary to DRS. Symmetric active-active controllers allow any VM to connect to an assigned LUN through any host port. This reduces and possibly eliminates storage access-related performance bottlenecks and application outages caused by I/O path thrashing. Dynamic Load Balancing functions across up to 32 back-end SAS paths without requiring administrators to configure preferred paths between VMs and disk.

Hitachi Dynamic Provisioning (HDP) creates a pool of storage from which VMs can draw capacity; actual capacity is assigned only when a VM writes data to the LUN. HDP enables organizations to maximize the use of the 2TB LUNs VAAI enables by wide-striping the volume across spindles and is a prerequisite for TP-STUN support.

Conclusion

During the last few years, VMware has signaled that it wants to make efficient use of performance, data management and data protection services embedded within intelligent and virtualized storage arrays. The announcement of the first three VMware VAAI APIs allowed VMware administrators to offload processes from the ESX server to the storage array, enhancing overall performance and efficiency. The delivery of storage related enhancements in vSphere 5 such as VASA shows VMware doubling down on this effort. Support for VAAI has quickly become a baseline requirement for storage arrays and storage array vendors supporting VMware ESX servers among VMware administrators. We expect the same with vSphere 5. Hitachi currently supports VMware VAAI on it VSP and AMS 2000 arrays and will support vSphere 5 upon general availability.

Storage performance and the ability of IT administrators to predictably deliver an increasing number of I/Os to VMs has also become an absolute requirement, especially for tier one applications. However, this new core competency should not necessarily add to the complexity of managing VMware. Nor should it require highly trained storage administrators to deliver. Evaluator Group believes that Hitachi's VSP and AMS intelligent block storage subsystems which support VMware's storage-related enhancements and APIs have the ability to deliver on both of these requirements.

About Evaluator Group

Evaluator Group Inc. is dedicated to helping IT professionals and vendors create and implement strategies that make the most of the value of their storage and digital information. Evaluator Group services deliver in-depth, unbiased analysis on storage architectures, infrastructures and management for IT professionals. Since 1997 Evaluator Group has provided services for thousands of end users and vendor professionals through product and market evaluations, competitive analysis and education. www.evaluatorgroup.com Follow us on Twitter @evaluator_group