

Product Brief

Tintri VMstore: Application-Aware Storage

Date: June 2014 Author: Terri McClure, Senior Analyst

Abstract: *Server virtualization brings tremendous value in terms of adding operational agility and significant cost savings through consolidation, increased utilization, and workload portability, but storage can be an inhibitor to realizing that value. Traditional storage systems were not designed to be shared between applications, and the gyrations required to do so, such as mapping LUNs, volumes, ports, and zones, introduce complexity without solving the entire storage challenge related to supporting virtual environments. This is because they cannot provide sufficient quality of service to ensure the right applications get the right amount of resources. What is needed is a solution that is application-aware in order to ensure application quality of service. That is the approach Tintri took with its VMstore.*

Overview

Server virtualization is a well-entrenched technology that has the potential to serve as the stepping stone for private cloud deployments, a capability on which more organizations are beginning to capitalize. Regardless of an organization's virtualization evolution, these deployments have obvious implications on IT environments, especially on storage. The level of virtualization maturity is reflected by the finding that most respondent organizations (55%) have virtualized more than 40% of the x86 servers that they believe are virtualization candidates. While there isn't a significant difference in terms of company size, midmarket organizations—with smaller and typically less complex environments—are more likely to have in excess of 60% of x86 servers virtualized than their enterprise counterparts.¹

In these environments, multi-hypervisor strategies are—and will continue to be—pervasive. Nearly two-thirds of the organizations ESG surveyed last year when we looked at virtualized environments reported using more than one hypervisor. In addition to providing pricing leverage, these strategies allow IT staffs to accommodate the specific needs of different applications by matching them with the most appropriate hypervisor, which is especially important to larger organizations with more diverse environments. The majority of organizations using multiple hypervisors plan to maintain this strategy moving forward.²

When it comes specifically to storage, in another survey, ESG asked IT managers responsible for the storage environment about virtual server environments, and found that more than one-third of organizations expect server virtualization to impact data storage over the next 12-18 months, while 43% of organizations cite the capital cost of new storage—whether incremental capacity or net-new systems—as a significant challenge related to server virtualization support (see Figure 1). Other top challenges include melding existing storage-based disaster recovery capabilities with those provided by virtualization technology (42%), as well as virtual server storage capacity planning (36%) and limited I/O bandwidth (29%).³ I/O bandwidth constraints in virtual server environments is a trend that ESG has seen in previous research, in which more than one-third of current solid-state storage users indicated that server virtualization I/O bottlenecks were the primary reason they deployed the technology.⁴ It is worth noting that *only 5% of respondents report not having encountered any storage-related challenges stemming from the support of server virtualization implementations.*

¹ Source: ESG Research Report, [Trends for Protecting Highly Virtualized and Private Cloud Environments](#), June 2013.

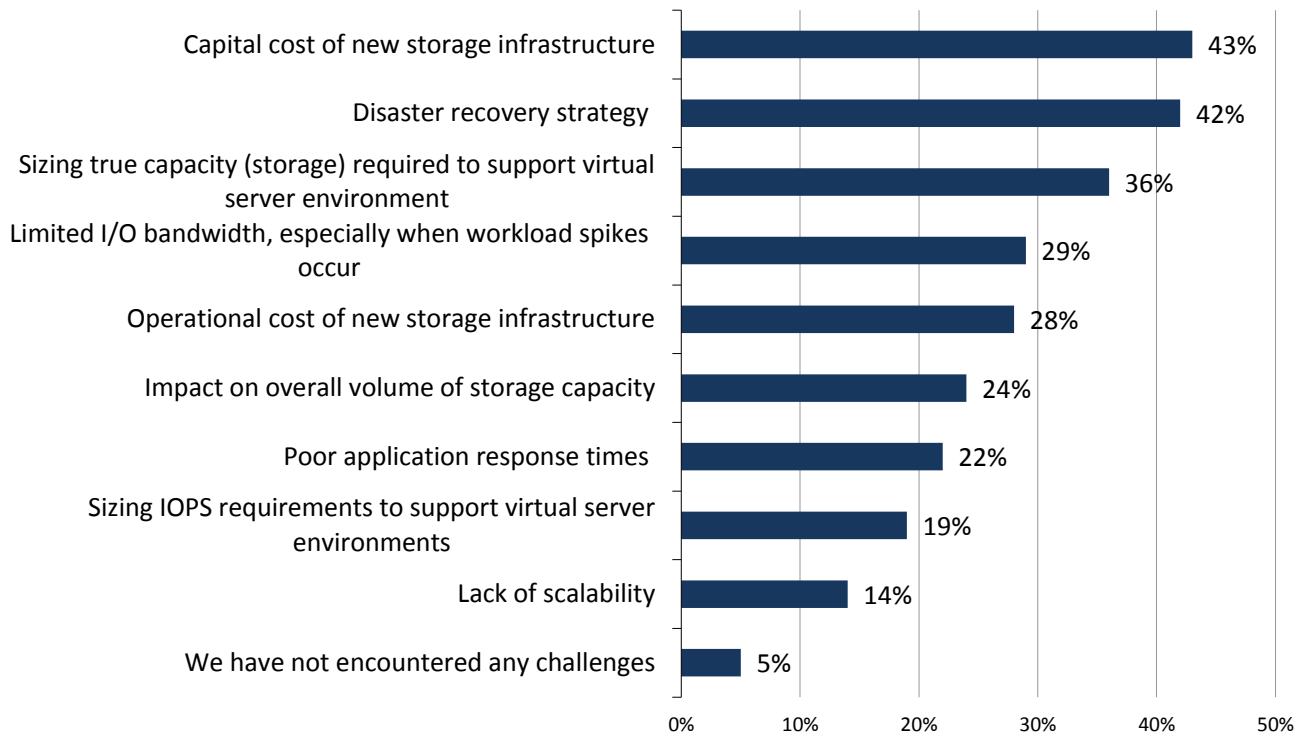
² Ibid.

³ Source: ESG Research Report, [2012 Storage Market Survey](#), November 2012.

⁴ Source: ESG Research Report, [Solid-state Storage Market Trends](#), November 2011.

Figure 1. Storage Challenges Stemming from Server Virtualization Usage

From a storage infrastructure perspective, which of the following would you consider to be significant challenges related to your organization’s server virtualization usage? (Percent of respondents, N=418, multiple responses accepted)



Source: Enterprise Strategy Group, 2014.

What’s driving all these challenges? Ninety-nine percent of storage systems in use today simply were not built to be "adaptive" to the changing requirements that server virtualization introduces. For example, the “I/O blender effect” is introduced when I/Os from multiple applications living on a single physical machine intermix I/Os. Storage was designed mainly as a one-size-fits-all function. It was designed for pre-virtualization workloads in which users built a system (physical in production, virtual only in the lab) that was "fixed" to an application, or a series of applications (i.e., workload[s]). Thus, users could test and deploy said workload and know *exactly* how it would perform, and performance was predictable because it was fixed. Because storage performance was known, I/Os per second were reliable. That’s fine in an unchanging workload environment, but when the workload suddenly changes, say a new VM is spun up or an application moves or multiple applications need to share storage resources, problems arise. That is because a fixed infrastructure is, by definition, not adaptable and users are stuck trying to figure out how to manually deal with performance spikes, bandwidth contention, and capacity planning.

It is not a fixed world today. Users build workloads (Exchange, Oracle, etc.) inside a VM and plop them (multiple workloads) on a (single) physical machine. We make sure the storage works. We initially treat it like it is "fixed." It works great! Then something happens. That workload moves elsewhere or another workload suddenly appears next to it, and then all of them start fighting for the I/O resources of the storage. The storage can only do so much since it is not application-aware and has no idea which applications are generating which I/O (because they are blended, hence the blender effect term), so it starts to arbitrate. Now the I/O performance on every workload starts to suffer.

How do most organizations handle it? They throw hardware at the problem and over provision across the board, undoing all the savings realized from server virtualization. Capital and operational costs of storage increase, DR becomes more challenging, and lots of manpower is expended fixing performance problems. How do traditional storage vendors handle it? They provide some QoS features, allowing storage administrators to tune performance by changing multiple parameters on the array in the hope of finding a combination of settings that will solve a performance issue, or perhaps intrusive application-specific agents or virtual machine settings that might enable better performance. The problem with

this is approach is the fact that storage still does not have a direct connection to the virtual machines that are the source and the victim of poor storage performance. QoS tuning on a traditional storage array tends to be a zero sum game, enhancing the performance of certain LUNs to the detriment of other LUNs. And because there is no application or VM awareness, any change in the environment (even as simple as adding a new VM) may negate the benefits of QoS tuning. A better answer is needed.

What IT needs today is a storage system(s) that can be adaptive (ideally in real time) to changing workload requirements. It needs to provide predictable performance under changing workload requirements in order to guarantee performance, not hope for performance. IT needs storage that can a) handle a high I/O load (much higher than what is needed for normal operations), b) delineate between workloads in terms of importance, and c) guarantee the important workloads get appropriate performance resources. IT needs QoS *at the application level*. QoS functionality evolves from simply exposing low-level tuning parameters to the storage administrator for constant manual adjustment to a more holistic, self-tuning approach where application-level awareness and smart algorithms provide the benefits of tuning without the need for manual intervention.

If the storage were smart and self-optimizing, using the right combination of intelligent caching, flash, and disk; could deliver QoS at a granular enough level, such as the application; and was able to be managed in terms of the virtual environment, rather than the physical (LUN-based) environment, then the storage and server worlds align and all the storage waste we see today would go away. There are post-virtualization storage products on the market designed in this manner from new entrants like Tintri, which is seeing market traction thanks to the efficiency that can be gained by taking a top-down view of storage requirements that is better aligned with a virtual server world.

Tintri VMstore: Application-Aware Storage

Tintri is a storage company formed in the post-virtualization world, so it has taken a different approach to storage than that taken by pre-virtualization companies. Its VMstore was designed from the outset to be used specifically in virtualized environments and managed in the context of virtual machines, not storage.

Traditional storage takes a bottom-up approach. It starts with disks, which are wrapped into LUNs, which are wrapped into volumes and then assigned to ports and data paths. And that is how they were managed: starting at the disk and mapping up to the server before getting assigned to the application. It had to be that way in a physical world with the technology available when traditional RAID arrays were invented. Running a traditional storage array in this world often means creating LUNs with different performance characteristics and assigning those to ports, paths, servers, and ultimately applications. Again, it is a fixed approach that does not respond well to change and, because it is done at the LUN level, there can be lots of stale or slow data that lives on high-performance LUNs.

Tintri takes a top-down approach, starting with the virtual machine-based application, and is managed in terms that a virtual administrator understands, such as VMs and virtual disks rather than mount points, LUNs, volumes, and arrays. All data management is done at VM level. Because of this approach, Tintri can take a fundamentally different tactic when it comes to providing storage for virtualized environments. Performance allocation can be provided at the application (a set of VMs) level, rather than the LUN or volume level, and resources are allocated end-to-end appropriately, even as the application grows and as additional workloads with varying performance (latency and throughput) needs are added to a VMstore. The storage system adapts to the dynamically changing application environment.

Because of this approach, management is streamlined and simplified. That means:

- Faster setup because LUN, volume, port, and zone mapping are eliminated and the administrator just needs to set up VMs or virtual disks.
- Easier scaling thanks to the same simplicity of provisioning realized during setup. Additional VMstore systems can be added as capacity or performance needs grow. Tintri Global Center, a centralized platform for managing multiple VMstore systems, can be used to realize the same level of efficiency in managing at VM level from a single pane of glass on a global scale.
- Protection policies are set up at the VM level, so they travel with the VM.

- Troubleshooting is simplified and expedited because the management console provides visibility end-to-end, from the hypervisor, through the network and storage layers at virtual disk and VM level.

VMstore has other features such as deduplication, compression, space-efficient snapshots, zero-space clones, and WAN-optimized remote replication that are all managed at the VM level and tightly tied to the application; no more figuring out and managing the set of LUNs an application lives on or, worse, managing a LUN or volume that has dozens of VMs and virtual disks. VMstore is a hybrid array with a flash-first design, meaning it handles 100% of writes and most of the reads from flash to deliver 99% I/O from flash, eliminating latency and disk contention associated with supporting mixed workloads.

Tintri's To Do List

Tintri is still a small company relative to the storage behemoths on the market. It previously limited support to VMware environments, which covered a large portion of the market, but since so many environments are multi-hypervisor, it meant that a second array vendor was required for non-VMware applications. It is addressing that significant barrier to market momentum by delivering support for multiple hypervisors with VM-level data management: Tintri supports VMware now, and will support Red Hat Enterprise Virtualization in July and Microsoft Hyper-V by the end of 2014.

It is clear from Tintri's technology approach that it understands the requirements of the virtual era. But to further accelerate market traction and growth, it needs to:

- Market in terms that both the storage administrator and virtual server administrator can understand. Server virtualization is forcing some level of organizational change as more administrative responsibility is put into the virtual administrator's domain. But many companies can and should keep the storage administrator role in place (albeit with broader understanding of the overall environment). If the storage administrator starts offering storage in terms that the virtual admin can understand, that is a win across the board: Both lives are made easier and more value can be realized for each IT dollar spent.
- Marry product capabilities with assessment services. Storage administrators are typically conservative buyers because changing the storage approach is considered risky; it is considered safer to stick with the known than venture into the unknown. Because of this, storage is typically the last holdout of change in IT. Tintri needs to ensure it can make IT organizations comfortable with change by showing that although its approach may be different, it is safe and can drive significant capital and operational savings on the storage front as well as help remove barriers to realizing value from server virtualization initiatives (and drive further capital and operational savings on that front!).
- Educate sales and their channel partners about a two-pronged approach to selling to virtualization and storage administrators, and selling the value of application-aware QoS. It is a competitive differentiator that traditional storage arrays cannot provide.
- Invest in marketing and awareness around application-aware storage. It is not a topic that IT is familiar with because it is not something that has been available from vendors. It is still something that the big vendors have trouble doing, so prospects won't hear anything about it from the storage behemoths. It takes time and money to create this level off awareness, but the value associated with application-aware storage and the ability it drives to provide QoS at the VM level are tremendous, so the payoff would be worth the investment.

The Bigger Truth

In the fixed physical world, there hasn't been a real problem with I/O for the past couple of decades. Storage vendors knew workload profiles, and knew that if they provided an array with x number of disks, y number of channels, and some known amount of cache, they would get z performance and meet user needs. All storage systems were "good enough" in the fixed-workload world. But workload mobility changes everything. It allows IT to finally get full utilization of physical server resources by moving applications to where the CPU resources are, but that completely breaks the traditional storage paradigm. And IT cannot afford to overprovision storage just to meet unpredictable workload requirements.

For IT to realize the value of virtualization initiatives, storage needs to be smart, self-optimizing, and self-managing. It needs to use the right combination of intelligent caching, flash, and disk. It needs to allow policies to be set at the VM level, *not* the LUN level, to ensure performance and protection policies travel with the application as it moves. It needs to be application-aware, as offered by Tintri VMstore.

This ESG Brief was commissioned by Tintri and is distributed under license from ESG. All trademark names are property of their respective companies. Information contained in this publication has been obtained by sources The Enterprise Strategy Group (ESG) considers to be reliable but is not warranted by ESG. This publication may contain opinions of ESG, which are subject to change from time to time. This publication is copyrighted by The Enterprise Strategy Group, Inc. Any reproduction or redistribution of this publication, in whole or in part, whether in hard-copy format, electronically, or otherwise to persons not authorized to receive it, without the express consent of The Enterprise Strategy Group, Inc., is in violation of U.S. copyright law and will be subject to an action for civil damages and, if applicable, criminal prosecution. Should you have any questions, please contact ESG Client Relations at 508.482.0188.