WHY VMWARE STORAGE IS SO PAINFUL AND HOW TO FIX IT

Storage Switzerland, LLC

Tintri
Intelligent Infrastructure

Since the advent of the all-flash array, many vendors have declared VMware storage infrastructure problems solved. As a result, many vendors have moved on to other more "glamorous" projects like Artificial Intelligence, Machine Learning, and NoSQL workloads. While those projects are undoubtedly important, the VMware infrastructure is still the heart and soul of most organizations, and for many organizations, it is THE environment. The reality is that tuning, managing, and efficiently utilizing storage is still the most painful part of the VMware storage infrastructure. These issues are also potentially the most significant inhibitor to scaling the VMware Infrastructure.

WHY VMWARE STORAGE IS STILL A PROBLEM

Most VMware storage solutions attempt to fix storage pain points with a sledgehammer instead of a scalpel. While the sledgehammer approach does solve or at least mask some VMware storage problems, many still remain. More importantly, the life of the IT professional tasked with managing the VMware storage infrastructure the same way they always have didn't get any easier as a result.

THE ALL-FLASH SLEDGEHAMMER

All-flash arrays, especially as they become increasingly affordable, help IT professionals solve one of their biggest challenges, VMware's infamous IO blender. The IO blender is the result of multiple physical servers, each populated with potentially dozens of virtual machines continuously access the storage system, which becomes a choke point. Instead of prioritizing workload IO, the all-flash system resolves the issue by responding much quicker to IO demands than hard disk or hybrid systems can.

Like most sledgehammers, all-flash seems to solve the problem, but as the environment continues to increase virtual machine density and mix workload types, the IO blender problem creeps back in. IT professionals quickly learn that lower latency infrastructure is only part of the answer. Storage systems have to use all of their resources intelligently to provide balanced consistent performance. All-Flash arrays alleviate some of the IO blender problem because they reduce latency, not because VMware can tap into their raw IOPS capability. The lack of an intelligent storage system forces the organization to either not virtualize some workloads or to dedicate certain types of storage systems for each workload type.

The result is a management nightmare for IT professionals. They are having to constantly rebalance workloads across the various storage systems supporting the infrastructure and are in the dark about how the next new workload will impact the performance of the currently running virtual machines. For example if the organization decides to virtualize a bare metal MS-SQL cluster the VMware administrator not only doesn't know how much available resources they have they also can't measure the impact of virtualizing the new workload. The only "alert" available is when uses start complaining about performance.

THE HYPERCONVERGED SLEDGEHAMMER

Another attempt at addressing VMware's IO blender problem is Hyperconverged Infrastructure (HCI). HCI approaches vary but typically, they work by running a component of the storage software on the same hardware as the hypervisor and virtual machines (VM). They also keep a copy of each VM's data locally on the server that is hosting the VM, as well as a distributed copy for data protection and to facilitate VM mobility. Ideally, the local copy facilitates all read IO, which reduces network traffic and the impact of the IO blender. Additionally, an increasing number of HCI architectures are also all-flash, further reducing the creation of an IO blender effect.

HCI scales by adding additional physical servers, typically called nodes, to the hypervisor cluster. Each node includes compute, memory, networking and storage. The problem is that most data centers don't scale all of these components in lock step. Usually, each organization tends to need significantly more of one type of resource, than another type given the diversity of applications. For example, if the organization needs more storage, it buys a node with it to meet that need, but that node also comes with all the other resources as well, and those resources go unused. Additionally, each time IT adds another node to the hypervisor cluster, complexity increases, especially on the network.

HCI hits the VMware IO blender problem with an even bigger sledgehammer. It does leverage some intelligence by reducing the amount of read IO on the network but it increases the impact of write IO on the network. HCI sacrifices efficiency in its attempt to eliminate the IO blender.

THE INNOVATION THAT VMWARE STORAGE NEEDS

Addressing VMware challenges like the IO blender, increasing VM density and efficiently leveraging more powerful servers requires intelligent application of infrastructure resources end to end, a deep awareness of the VMware operating environment and efficient scaling of storage capacity and performance. Instead of a generic system designed for a multitude of workloads that might happen to include VMware, IT should consider a storage system purpose built for VMware.

Our next chapter details how VMware aware storage systems that operate at the VM, which is the atomic unique of a VMware environment, provide not only taking action on specific VMs but also empowers learning based on behavior of the VM and the VMs, visibility and control. The result is significantly more efficient systems that provide better overall performance and significantly reduces administration time.

VMWARE STORAGE NEEDS INTELLIGENCE AND AWARENESS FOR A DIFFERENT EXPERIENCE

VMware continues to be at the heart of many data center infrastructures and will continue to be that heart for years, if not decades, to come. Many of these infrastructures are still struggling however with the most basic of data management and data protection functions. All-Flash arrays may have alleviated some of the infamous IO blender issue but there are still many more storage challenges to tackle. Two key challenges and areas for innovation are gaining insight into the storage IO demands and behaviors of each specific virtual machine as well as the need to better predict and plan for scale.

FROM VISIBILITY TO INSIGHT

Most storage systems that support VMware environments are block-based, which by default provide no visibility into the specific virtual machine (VM) IO activities. In 2015, VMware delivered a feature called VVOLS that provided increased VM visibility but with limited ability to respond rapidly to specific IO conditions. VVOLS is still volume based and requires the creation and management of volumes. An alternative is to use a file-system based storage architecture, which, because VMs are essentially files, provides visibility into each VM's IO profile.

Visibility into each VM's IO profile is an improvement over block-based storage, but to take full advantage of this granular view of VM storage, it requires more than just loading VMs on a NFS volume. The storage system needs to have intelligent software built-in that performs a continuous analysis of each VM's IO pattern, storage capacity consumption rate and provides predictive forecasting / modeling of future use. Armed with this insight, IT can easily respond to complaints about storage performance and either take corrective action or prove that storage is not the source of the bottleneck.

FROM INSIGHT TO LEARNING

Insight into the IO characteristics of a specific VM enables IT to more quickly and precisely intervene when problems arise. Intelligent infrastructure learns from that analytics captured, allowing the system to take corrective action on its own. Based on the analysis, it should be able to take corrective action either to mitigate outages or to meet changing performance demands. Applying machine learning to the data the storage system already collects enables organizations to avoid spending all day manually monitoring and managing storage.

DEALING WITH SCALE

One reality that almost every VMware administrator and the infrastructure admins that support them must deal with is scale. Either the current storage system will run out of storage capacity or it won't be able to keep up with storage IO demands. Scaling typically means adding another storage system and migrating workloads to the new system so that the old one can be retired. Several vendors have brought out scale-out storage solutions or scale-out hyperconverged solutions to address the scaling problem but these environments tend to start too large, don't scale granularly enough and put extra pressure on the storage network. A more intelligent approach is a system that can scale up by adding additional storage capacity and then scale-out by adding additional storage systems. The second storage system can start small and have capacity added to it as the need demands.

The typical problem with adding multiple storage systems is managing them and figuring out which VMs IT needs to migrate to the new system. Storage systems need innovation so that IT can forecast, by using the methods described above, when the need for a new storage system will occur. There is also a need for innovation in automating which and how VMs move to the new storage system since in most cases, the current system still has years of reliable service left. Trying perform a migration, at-scale, between systems that are not VM aware is much more difficult, time consuming and more than likely will impact production applications.

Armed with an intelligent scale-out capability IT can buy a new storage system that performs much better but initially has significantly less capacity then the current system. The storage software can then leverage the analytics information to move the most viable candidates automatically, to the new system. This automate process frees up capacity on the current system while improving the performance of VMs that need it.

CONCLUSION

In order to enable IT professionals to focus on tasks that more directly and positively impact the organization, they need technology that manages itself. The storage system and infrastructure are an excellent starting point. With proper intelligence software, a VM aware system can deliver valuable telemetry data that IT can use to manage storage better. The endgame though, is to have the storage system teach itself from this telemetry data and automatically take corrective measures, freeing IT to work on higher level tasks.

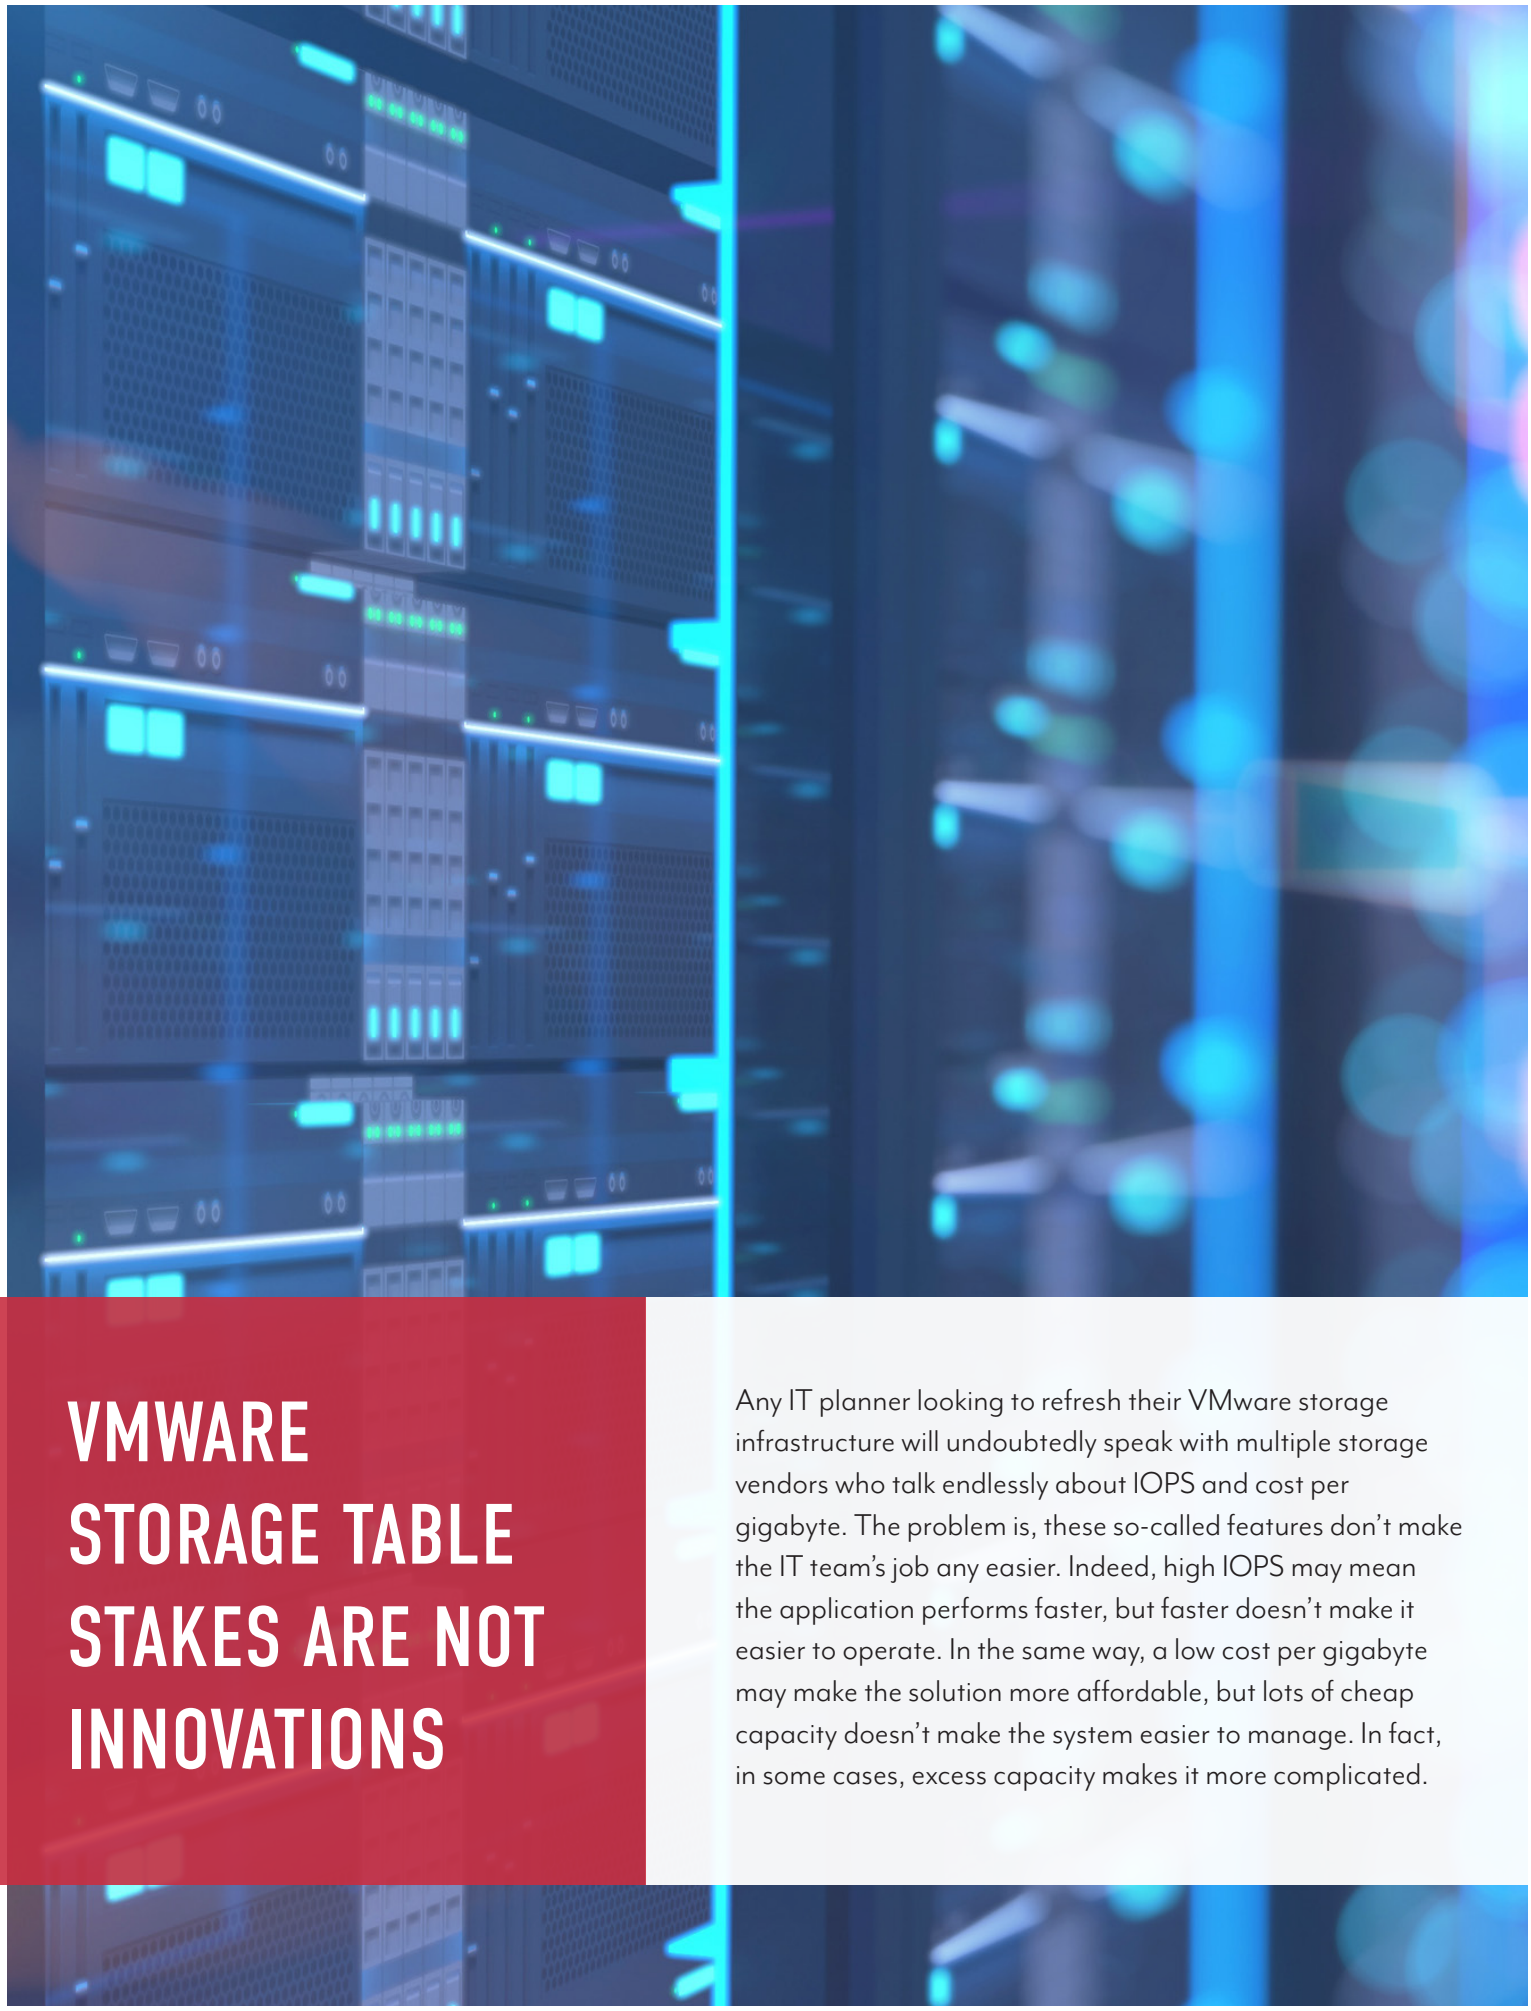

VMWARE STORAGE TABLE STAKES ARE NOT INNOVATIONS

Any IT planner looking to refresh their VMware storage infrastructure will undoubtedly speak with multiple storage vendors who talk endlessly about IOPS and cost per gigabyte. The problem is, these so-called features don't make the IT team's job any easier. Indeed, high IOPS may mean the application performs faster, but faster doesn't make it easier to operate. In the same way, a low cost per gigabyte may make the solution more affordable, but lots of cheap capacity doesn't make the system easier to manage. In fact, in some cases, excess capacity makes it more complicated.

While IOPS and affordability are an essential aspect of any storage refresh, most vendors can get within a range of each other in these two categories. IT planners should instead be looking for a storage system that increases VMware's usability while at the same time improving IT operations efficiency. A storage system can improve VMware usability by proactively managing IO so that more virtual machines (VM) can run per physical server, and more physical hosts can connect to the same storage system. Both of these capabilities, though, require a storage system that can analyze and intuitively adapt to changing conditions. A storage system can improve IT efficiency by providing clear, real-time insight into storage system telemetry data and automating mundane tasks.

IT STARTS WITH A FILE-SYSTEM

Block storage systems, long the mainstay of VMware storage infrastructure, are challenging to manage and don't provide the level of insight that IT operations need to maintain a busy VMware environment effectively. Organizations, to keep some simplicity, can place all VMs on a single logical unit number (LUN). The problem is that per VM visibility is lost, making management more difficult. Alternatively, IT can create a separate LUN for every VM.

Features like VMware's VVOLs simplify the VM per LUN process a bit but still don't provide the level of granularity that most storage managers want. File-systems, on the other hand, do ensure a high level of insight without impacting performance. File-Systems store every VM on a single volume, and each VMs datastore is accessible and visible to monitoring tools. The file-system can't be just any file-system though. It needs to be specifically designed for the VMware use case so that performance and other integrations are optimized. A purpose-built file system for VMware enables deep integration with the entire ecosystem. More importantly, it needs to leverage the per-VM visibility and via machine learning, proactively manage the storage infrastructure

QUALITY OF SERVICE

When trying to maximize the VMware investment, IT planners want to achieve the highest virtual machine density possible, and they want the storage system to support as many physical servers in the ESX cluster as possible. At a base level, all-flash arrays help organizations reach these goals, but even an all-flash array, under these highly-dense conditions can hit a performance wall if it is not optimized.

Organizations looking to maximize server host and VM density need quality of service (QoS) to make sure the most critical workloads get priority access to the storage system's performance. However, the system should go beyond just letting administrators preset QoS levels and also leverage machine learning to determine normal performance parameters and make sure the VM remains inside those boundaries. The proactive QoS capability ensures that a runaway VM doesn't take performance away from a more mission-critical VM.

CAPACITY PREDICTIONS

Another challenge facing highly optimized VMware environments is capacity planning. These environments tend to grow continuously and are adding new VMs multiple times a day. Of course, each of these VMs need to store data and can cause the storage system to run out of capacity if IT does not carefully monitor its consumption.

The problem is administrators of highly optimized VMware environments may be too busy to monitor their storage utilization constantly. Instead, the storage system needs to provide administrators with proactive analytics to help them determine in advance when the storage system will run out of capacity. This advanced knowledge means that organizations don't have to over-provision their storage systems, paying in advance for storage they won't use for months or even years.

DATA PROTECTION BUILT-IN

Backup should be part of every data protection strategy, but if IT is recovering from backup, it means that making recovery point and recovery time objectives are at risk. The VMware storage system needs built-in protection. It should take advantage of the VMware granularity, and be able to set different snapshots and replication schedules per VM.

CONCLUSION

Innovation in VMware storage is no longer IOPS or low price. Those are table stakes. VMware storage focuses on reducing IT operations overhead by proactively responding to VMware environmental conditions based on machine learning.

Storage Switzerland, LLC

Storage Switzerland is the leading storage analyst firm focused on the emerging storage categories of memory-based storage (Flash), Big Data, virtualization, and cloud computing. The firm is widely recognized for its blogs, white papers and videos on current approaches such as all-flash arrays, deduplication, SSD's, software-defined storage, backup appliances and storage networking. The name "Storage Switzerland" indicates a pledge to provide neutral analysis of the storage marketplace, rather than focusing on a single vendor approach.

Tintri
Intelligent Infrastructure

Based in Silicon Valley, Tintri is a wholly owned subsidiary of DataDirect Networks (DDN), the data-at-scale powerhouse and world's largest privately held storage company. Tintri delivers unique outcomes in Enterprise data centers. Tintri's AI-enabled intelligent infrastructure learns your environment to drive automation. Analytical insights help you simplify and accelerate your operations and empower data-driven business insights. Thousands of Tintri customers have saved millions of management hours using Tintri. Choose differently, the choice is yours. Where will you invest your resources to increase the scale and value of your business?

Learn more about the Tintri portfolio of solutions at https://www.tintri.com/products.