

Technical White Paper

5 Reasons That Flash + VAS is Key to Enterprise Cloud



Contents

- Introduction..... 1
- Reason 1: Autonomic Operations and Orchestration 2
- Reason 2: QoS-driven Tiering per VM 3
- Reason 3: Real-time Analytics..... 3
- Reason 4: Scale out 4
- Reason 5: DevOps Benefits and Agile Development Support..... 4
- Tintri VAS Architecture 5
- Summary 5

Introduction:

The benefits of cloud computing are well-documented and have led to its widespread adoption across organizations. Implementation options comprise enterprise, public, hybrid, and private options. An enterprise cloud consists of private cloud in the datacenter and an option to connect to public cloud for some specific workloads. A public model leverages economies of scale across all clients, simplifies and (to a degree) offloads management, and typically provides access to a broader array of storage and compute options than a private cloud. Enterprise clouds, which are predominantly private clouds and inward focused provide workload tiering, portability, and cost optimization.

Enterprise clouds provide similar benefits to public clouds, such as multi-tenancy with chargeback, scalability, self-service, dynamic provision, and parallel computing options for Big Data workloads. But unlike public clouds, it provides support for more traditional/custom applications (that are not written or tuned for the public cloud), and direct user management. As such, an enterprise cloud is best for non-cloud-native, mission-critical, and/or very complex workloads that demand strong uptime, minimal latency, and tight management.

In addition, while public cloud offer a cost advantage for short-lived or new applications and infrastructures, enterprise clouds can provide more compelling financial returns for longer-term workloads due to the ability to amortize up-front infrastructure and ongoing management costs over a longer time period, avoiding the higher marginal costs of public cloud subscriptions.


Because of the performance demands of typical enterprise cloud workloads, such as online transaction processing, deploying solid-state flash storage becomes essential to meet service level agreements, especially when they are virtualized. But the unique requirements of workloads, such as strict performance, robustness, and flexible management, require more than just basic flash components.

There are areas where public clouds augment private clouds well which creates a well rounded enterprise cloud strategy. Those areas are remote access, offsiteing, and surge support. There are always risk benefits to co-locating data, and when that second location provides anywhere access, it further reduces the risk. Public clouds can also provide rollover capacity to cover temporary demand spikes, ensuring continuous performance, even during heavy loads. So even a near pure-play private cloud would benefit from on-demand capacity and storing snapshots in a 3rd party public cloud for redundancy, robustness, recovery purposes.

This paper reviews key reasons why the advanced software-based storage services of VM-aware storage (VAS) simplifies and optimizes the deployment and management of storage infrastructure in addition to all flash array deployment in enterprise clouds, allowing organizations to maximize its benefits across private and public cloud:



1) Autonomic Operations and Orchestration



2) QoS-driven Tiering



3) Real-time Analytics



4) Scale out



5) DevOps Benefits and Agile Development Support



Reason 1: Autonomic Operations and Orchestration

Managing an infrastructure in a private or enterprise cloud configuration can be a monumental challenge. Compute, storage, and networking must all work in concert or performance and reliability will suffer greatly. Manual management of unit of computing such as VM and its associated storage and network attributes' orchestration simply compounds that challenge

When traditional storage or standard flash becomes part of an enterprise cloud deployment handling virtual workloads, it brings tremendous management baggage due to legacy storage structure bindings. Traditional shared storage systems manage objects such as LUNs, volumes or tiers, which have no intrinsic meaning for VMs. Each new VM instance must be assigned a specific storage LUN or volume. When IO requirements and VM behavior are poorly understood, a painful trial-and-error process ensues to make sure the storage needs of each VM are met. Storage and VM administrators must coordinate to ensure each application not only has the capacity it needs, but also enough IO performance for the expected loads. Some specific challenges include:


- Virtualization combines applications and workloads to create extremely random IO, resulting in performance requirements that choke traditional storage. This storage bottleneck and complexity of scaling storage beyond a single system creates inefficiencies that limit VM densities within a LUN, precludes virtualizing IO-intensive applications, such as databases and critical business applications, and adds unnecessary management overhead. When one or more of the VMs in a LUN gets overly demanding or goes rogue, you've got a noisy neighbor. Noisy neighbors can starve other VMs in the same LUN of resources and severely complicate VM management in a private cloud.
- System management also becomes problematic. Management of a LUN- and volume-based storage back end and its associated functions simultaneously with compute and networks is a very manual process, especially when trying to optimize a dynamic virtual environment. Each request for a new VM triggers a manual configuration of storage and network attachments. Handling this by hand does not scale, which counteracts one of the main reasons for moving to a private cloud.


VAS in enterprise cloud replaces traditional storage architectures with one built around an individual VM or Instance as the base unit. This allows for granular management of VMs, providing dedicated lanes for each VM to eliminate the noisy neighbor problem and automatic, dynamic configuration of storage, compute, and network connections as VMs are spun up. This greatly simplifies the deployment and management of private cloud virtual environments.

Although enterprise clouds provide a platform to support scaling of storage as needed, managing that storage as it scales can burn a lot of time. Without automated virtualization orchestration that can provision, resize, protect and recover cloud storage for VMs and their environments, IT must intervene as employee use of the cloud evolves. Having to manually configure storage whenever needs change means responses to these needs are not real-time. It also means less time for IT to spend on high ROI projects.

For an organization with over 150,000 VMs, for example, aligning cloud storage needs as they develop is essentially impossible without some form of simplified management of cloud storage and automated orchestration to keep it current.

Building a private cloud virtual environment on VAS dramatically simplifies the orchestration of VMs at scale. VAS automatically aligns storage to changing needs without any IT intervention needed, providing one central management console that requires only minutes per day to manage.





Reason 2: QoS-driven Tiering per VM

While storage tiering capabilities have evolved considerably at a system level, managing storage placement of developer's individual VMs based on policy is still quite difficult in enterprise cloud. Essentially, since non-VAS storage systems define policies only for groups of VMs at the LUN or volume level, granular management and true optimization of storage in enterprise cloud for each tenant is untenable, especially at scale. In a 3-tier application, QoS for webserver, app server and database shouldn't be put together in LUN or volume as all tiers have their own requirements.

In addition to challenges around tiering, without VAS, it is practically impossible to support multi-tenancy or to offer performance guarantees to meet SLAs. Despite being able to craft policies that reflect best practices for your organization, applying them at a VM-level is also problematic. Allocating minimum and maximum IOPS for each VM while controlling for latency is essentially impossible without VAS.

VAS makes it possible for its customers to guarantee performance of applications automatically without administrative intervention through automated tiering based on QoS policies and needs at the VM level. VAS delivers granular performance and management of virtual environments to extract full value from an enterprise cloud.




Reason 3: Real-time Analytics

As noted above, most virtual environments, even in a private cloud, are tied to LUN-based storage architectures with volumes and tiers, so they suffer significant shortcomings in identifying VM-level sources of latency, creating barriers to real-time analytics. Traditional storage's resource pool structure has no intrinsic meaning for VMs, making granular real-time monitoring for visibility on a per-VM basis and visibility into cloud latency issues all but impossible without massive additional effort. Without VAS monitoring's visibility to support real-time cloud storage analytics, it becomes impossible to optimize a private cloud virtual environment and extract maximum value from a flash array.

VAS's detailed visibility driving real-time and predictive analytics across cloud-based hosts, networks, and storage addresses this shortcoming of traditional storage. For maximum value, these analytics must provide real-time actionable insights at a VM-level to rapidly identify factors that could contribute to increased cloud storage latency, degraded performance or even downtime. Rapid identification and resolution of these factors across hosts, networks, and storage requires analytics provide a real-time view of what is going on in the private cloud's virtual infrastructure.

In addition, leveraging time and condition data to generate predictions of future behavior can help identify trends and prevent problems in the cloud before they manifest themselves. The resulting end-to-end troubleshooting capability can instantly identify developing and potential cloud performance hot spots with comprehensive analytics visualization.





Reason 4: Scale out

One of the big benefits of the cloud is scalability. As needs grow, it is a relatively simple task to add more capacity. For most storage architectures, the solution is to add more capacity shelves (scale-up). But that puts increasingly heavy loads on storage controller nodes, leading to reduced performance. This pretty much precludes offering enterprise private cloud storage for mission-critical, tier-1 applications. This is a setback to cloud strategy pursued by most enterprises when they want to turn into cloud-first companies.

Scale-out storage, which provides both storage controller and capacity expansion when needed, is a better option for private clouds. But trying to manage conventional scale-out storage for virtual environments can be exceedingly complex. Here are some of the extensive storage management challenges with conventional scale-up storage.

- LUN- and volume-level data leads to bad guesses and poor decisions about optimal VM placement
- Poor migration recommendations based on incomplete data, and no visibility in to the impact on performance or the time required to complete a migration
- Problem VMs constantly getting bounced back and forth between arrays reacting to, but never resolving, performance issues
- Having to throw out your old storage, migrate to a scale-out architecture, and hire a team of Storage PhDs to run it

VAS is built on a scale-out platform where performance scales directly with the number of storage nodes, reducing both management complexity and time. Storage infrastructure within enterprise cloud should optimize the entire VM and associated storage footprint under the cover without causing any downtime or even interruption in VM performance. Scale-out and adding cloud capacity and performance should automatically trigger entire footprint re-balance to keep cloud performance above SLAs. That's the only way to truly scale cloud.

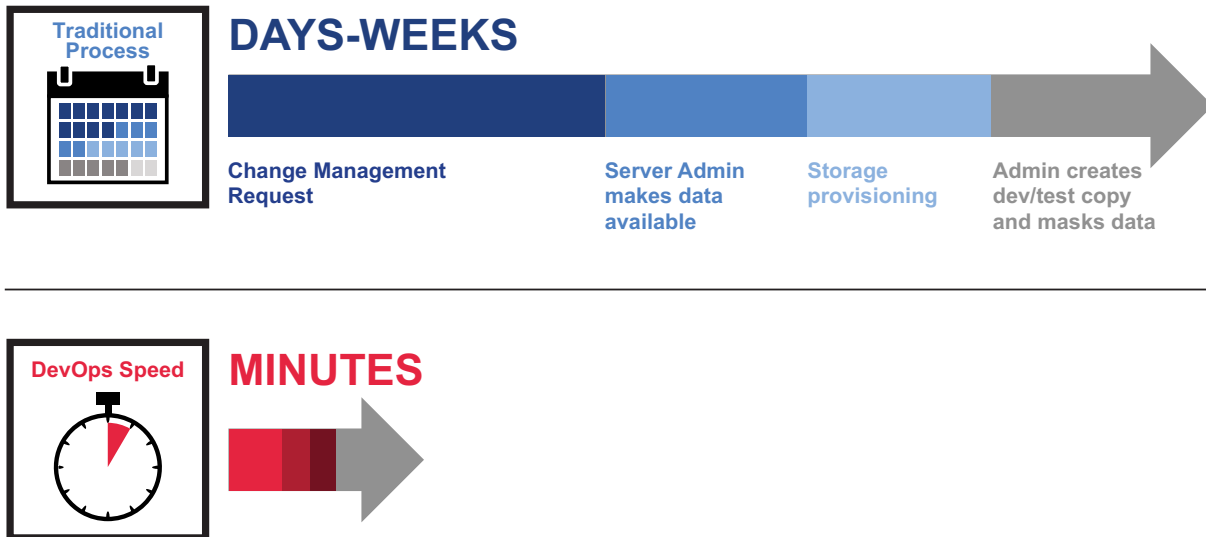


Reason 5: DevOps Benefits and Agile Development Support

A major driver of innovation and competitiveness within an organization is DevOps. For DevOps to be successful, it should have access to real-time copies of production data and the ability to ensure that changes in a master VM are instantly reflected in all of the copies that have been generated. This is especially necessary for teams that rely on agile development, which is built around a steady stream of smaller development efforts that drastically increase the changes that an infrastructure needs to accommodate. Having widespread access in an enterprise cloud to always current data without impacting the production environment provides a major boost in productivity to an organization's DevOps teams.

But providing this real-time access and consistency among VM copies is a herculean effort without the proper tools, especially in the cloud. Trying to manage this with the limited services offered by mainstream storage is not really possible due to the time and understanding needed of detailed dynamics of an infrastructure. In addition, with the growing popularity of containers within the developer community, it is becoming even more difficult for IT administrators to meet seamlessly the needs of its DevOps teams with traditional storage solutions in a private cloud environment.

Only VAS can provide the real-time storage management at the VM level that gives DevOps the granularity and data consistency it needs to streamline development efforts and accelerate the release of new applications. Having the access to current data and the ability to promulgate it across VM copies provided by VAS allows DevOps to experiment with different modifications of applications before rolling them out to production environments. In addition, providing an option for rollover or temporary capacity expansion via the public cloud as needed further enhances VAS benefits to DevOps.



VAS Architecture

To make the most of your virtual environment in a private cloud deployment, you must build that infrastructure around VAS. And Tintri's VAS platform is ideally suited to deliver the full value of enterprise cloud.

Tintri for Automated Operations and Orchestration

Tintri VAS eliminates LUNs, instead assigning every VM its own lane. This greatly simplifies and provides broad automation of resource allocation in your private cloud environment. With Tintri, there's no conflict over resources or policies and therefore, no noisy neighbors.

The ability to automate storage tasks has long been central to Tintri's differentiation. Tintri allows customers to manage up to 160,000 VMs from one central console, using automated policies and Tintri software to substantially reduce management effort.

For example, Tintri offers its VMware vRealize Orchestrator plugin, which automates many storage operations, such as snapshot, clone, replicate and copy data management at the VM-level. Customers can integrate vRealize Orchestrator into VMware cloud management and automation systems such as vRealize Automation and vCloud Director.

Tintri for QoS and Analytics

With Tintri's advanced QoS capabilities, administrators can now allocate exact maximum and minimum IOPS to each individual VM, not just to a LUN or volume. Unlike conventional QoS, which requires administrators to predict the right IOPS values, Tintri provides visual guidance on the QoS values to specify, removing guesswork. Cloud administrators can see the immediate impact of throttle changes on VM-level latency instead of waiting for end user feedback. The visualization spans the entire cloud infrastructure—including latency stemming from host, network, storage contention and QoS throttle.

Tintri's granular QoS policies help you manage mixed workloads for different customers, with different service level requirements, while hosting multiple types of hypervisors on the same Tintri VMstore in an enterprise cloud.

Tintri's analytics provides customers with cloud-based real time and predictable analytics to improve data center planning and operations., Tintri can model virtualized applications' need for storage capacity and performance. Tintri optimizes your VMs based on a complete picture of their storage capacity and performance needs

Tintri for Scale-out

The Tintri VAS platform is built around delivering performance as capacity scales. Tintri's modern scale-out architecture allows you to start small, with as few as 17TB and a couple hundred VMs, all the way to 10 PB and 160,000 with the same storage team you have today.

Tintri's VAS platform uses a million data points collected every 10 seconds from thousands of VMs to optimize VM distribution across multiple pools of storage within your private cloud. With Tintri, you can:

- Start small and scale big: Start with 1 or 2 VM stores and a couple of hundred VMs and scale up to 10 PB and 160,000 VMs or more.
- Scale smart: Deep workload analysis identifies VMs that are growing rapidly and applies sophisticated algorithms to model the growth ahead and avoid resource constraints within a pool.

It gives you least-cost recommendations, saving you time, bandwidth and capacity, to maintain optimal VM distribution across your enterprise cloud. You can review and edit the recommendations, and see the outcome before committing. Tintri learns every time you edit its recommendations and allows you to opt certain VMs out of migration.

Tintri for DevOps

Tintri makes providing and managing production copies at a VM level to DevOps seamless and simple. Tintri allows the user to move back and forth between snapshots of an individual VM without losing other snapshots or performance history. Administrators can also use this capability to update hundreds of "child" VMs from a refreshed "master" VM without physically moving data or reconfiguring the VM and/or storage. They can even automate the process with Tintri PowerShell or REST APIs.

Connectivity to Public Cloud

Further, Tintri's S3 API support opens up extended avenues for both protecting data and processing it under extraordinary loads thus completing your enterprise cloud strategy. Now you can use the public cloud in conjunction with the private cloud for all cloud-first workloads.

Summary

For a successful enterprise cloud, organizations have to start building it on top of latest advances in hardware such as all flash array. But to realize the full potential of their cloud strategy, organization need to consider storage solutions such as VAS, which focus on the unit of arbitration in cloud such as VMs or instances and remove any associated bindings to legacy definition of storage, which is waste of time, effort and budget.

