

Tegile and Tintri Comparison

July 02, 2015

Jeremy Li

The storage technology advancement has proved to the world that high-end and midrange as well as NAS systems are equally good enough to compete with high-end SAN systems even in business-critical environments thanks for flash storage arrays either via all-flash (SSDs) or hybrid flash arrays widely deployed in production.

Tintri, founded in 2008, claims the top Fortune 15 are deploying Tintri systems in production environments, while none of the current disruptive storage vendors can achieve that level, even though Tintri only supports NFS protocol, although SMB 3.0 support is coming soon.

In June 2015, Tegile, founded in 2010, earned a spot in Gartner Magic Quadrant for Solid-State Arrays as a visionary for the first time ever; while Tintri does not provide an all flash array at this time because it has been achieving 100% writes to SSDs and 99% reads from SSDs in its hybrid array for the last four years.

In Nov. 2014, both Tegile and Tintri earned a spot in Gartner Magic Quadrant for General-Purpose Disk Arrays as Visionaries, and Tintri surpassed Tegile in both axial - Ability to Execute and Completeness of Vision in the Gartner Magic Quadrant. See the [Full Report](#).

Tintri has produced a system that can support 100TB effective storage, delivering 140K IOPS and supporting up to 3,500 VMs in a single 4U in the VMstore T880 system at a list price of \$250,000.00 due to its focus on both performance and capacity efficiency. It has the following unique features:

(1) Per VM Management; (2) Per VM QoS; (3) Per VM Replication; (4) Key Managed Encryption; and (5) Multi-hypervisor supports on a single 4U appliance.

Tegile has produced a system with [314TB RAW capacity](#) with a storage pool capability and [Tegile IntelliCare Flash 5 Guarantee](#).

Tegile supports both physical and virtual machines (VMs), while Tintri only supports VMs.

The definition of each Magic Quadrant:

Magic Quadrants provide a graphical competitive positioning of four types of technology providers, where market growth is high and provider differentiation is distinct:

- **Leaders** execute well against their current vision and are well positioned for tomorrow.
- **Visionaries** understand where the market is going or have a vision for changing market rules, but do not yet execute well.

Tegile and Tintri Comparison

- **Niche** Players focus successfully on a small segment, or are unfocused and do not out-innovate or outperform others.
- **Challengers** execute well today or may dominate a large segment, but do not demonstrate an understanding of market direction.

Below are a few key items between two arrays:

I. VVOLs Technology

Tintri has developed VMstore four years ago that equates to a VVOLs technology recently available in vSphere 6.0. The VMStore provides a VM-aware/App-aware array that has been used for more than 4 years in production among many enterprises customers.

Tegile lacks VVOLs technology and relies on LUN/Volume or a Mount Point that also lacks the granularity feature for storage configurations and Quality of Services (QoS). Below are a few items that must be taken cared prior to using VVOLs and new issues arisen from implementing VVOLs:

- Existing arrays with LUN limits can max out quickly – 30 VMs will be max out per array if a snapshot runs hourly for the retention period of 3 days and a half, meaning dramatically reducing the snapshots retention period
- Implementing VVOLs requires an array firmware upgrade
- VASA 2.0 upgrade is required via a storage vendor that will support bidirectional flow and VVOLs
- vSphere 6 is required
- Enabling VVOLs will break up vCloud Air, SRM and others
- VVOLs must live in a Storage Container (SC), meaning implementing VVOLs requires to move resources from its current location into a SC

Read a great article at <http://info.tintri.com/VVols-Taneja-Tech-Brief.html> (Source: TANEJA GROUP) and watch a great webinar titled “So you think you know VVols” via a link at <http://event.on24.com/wcc/r/1010699/7F757C4BA8A2BB976E68D9AE8AC7D386>

In summary, most enterprises will not be ready to implement VVOLs for a long time due to the above factors mentioned in the article.

II. Tegile Dual Active/Active Controller vs Tintri Active/Passive Architecture

Tegile supports two active/active controllers, meaning data access can be used simultaneously via two controllers among a storage pool, while Tintri only supports active/passive controller architecture.

III. Tegile Supports Pooled Storage, While Tintri Does Not

Tintri does not support a storage pool. Therefore, Tegile has a better scale out capability than Tintri's.

Tegile and Tintri Comparison

A response from Tintri: That we feel is an incorrect statement. Tintri scales easier and endlessly. We actually provide more reliability for scale out than “Active-Active” can.

Because of our TGC, Tintri Global Center- we are able to scale endlessly under one pain of glass, with consistent IOPs as you add VMStores. We have a large percentage of business that is Service Provider, and Cloud oriented. They rely on us for our ability to easily scale.

IV. Tegile Multiprotocol Data Access Support vs. Tintri NFS Only Support

Tegile can provide multiprotocol data access support (FC, iSCSI, NFS, CIFS and SMB 3.0), while Tintri only supports NFS for simplicity in file system. Per Tintri, any customers only need to connect a 4U array to high performance, low-latency top of rack (ToR) switches (e.g., Cisco 5600 switches) by eliminating one extra hop (e.g., a fiber channel layer - two Cisco MDS 9509 FC switches) and will not see a performance difference by comparison to connecting a Tegile array directly to two FC switches.

It is worth of noting that Cisco is not going to make Fiber Channel Only Data Center switches in the future due to the following factors:

- An IP-based fabrics switch, a component of [Programmable Data Center Networks](#), along with Software Defined Network, and Storage Defined Solutions becomes an important element of the emerging Software Defined Data Center that can scale very well – Massive Scalability (e.g., from 3 servers to over 1,000 servers. Storage grows automatically as application servers are added)
- By segmenting the storage network from the existing IP Fabric using a 4-way Arista 7500E spine with 40Gb/s uplinks from dual top of rack switches per rack, the bandwidth and resiliency provided at 320Gbps/s per rack exceeds that of traditional FC SANs and takes their storage architecture to the next generation

Note 1: Cisco Data Center Switches only support 8Gbps in FC, while it can support 40Gbps/s now, and will support 100Gbps/s in the future for IP-based fabrics switches.

Note 2: The most recent [three month stock trend](#) from Brocade might also indicate that the FC fabric Data Center switches might not be able to grow constantly in a long term.

Read ARISTA AND EMC SCALEIO WHITE PAPER via <http://www.arista.com/assets/data/pdf/JointPapers/Arista-ScaleIO-High-Performance-Storage-Solutions.pdf> for the reason why FC Data Center switches will become obsolete.

Generally speaking, without FC support, enterprises may lose nothing. On the other hand, with multiprotocol support, they have flexibility.

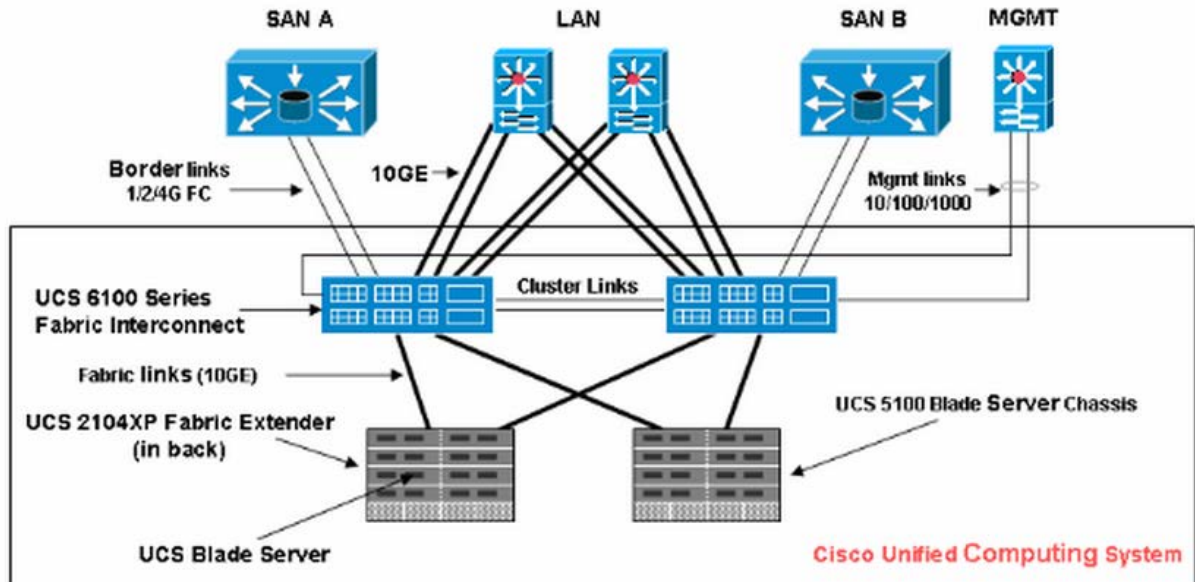
Traditionally, a FC network is a separate non-routable network, meaning it is more secure. On the other hand, using a non-routable IP network with a non-routable VLAN

Tegile and Tintri Comparison

by separating a corporate LAN network with a private storage network might mitigate the security concerns. But, it all depends on your environment.

Network Diagram

This is a typical topology used with the Cisco UCS:



Source: <http://www.cisco.com/c/en/us/support/docs/servers-unified-computing/ucs-manager/110434-fcoe-vsan-connectivity.html>

Note 1: UCS 6200 Series (e.g., 6248 or 6296) Fabric Interconnect (FI) is a brain of UCS and has replaced the first generation UCS 6100 series.

Note 2: Today, high-performance servers deployed in cloud will be able to support more virtual machines and workloads than ever before. The requirement for deploying new servers on demand places additional strain on the network fabric. The Cisco Nexus 5600 platform switches address this challenge by providing scalability and performance, making it a good platform for meeting current and future needs.

Note 3: The top-of-rack (ToR) Layer 2 and Layer 3, 10/40 Gbps switches (e.g., Cisco [Nexus 5600 series](#)) are not shown in the above picture.

V. Visibility - Telemetry Data Visibility for Each VM

Since Tintri uses VMstore technology referenced above, it can gain a granular visibility for any telemetry data via a single pane of glass without a third party VM tool to achieve a similar result.

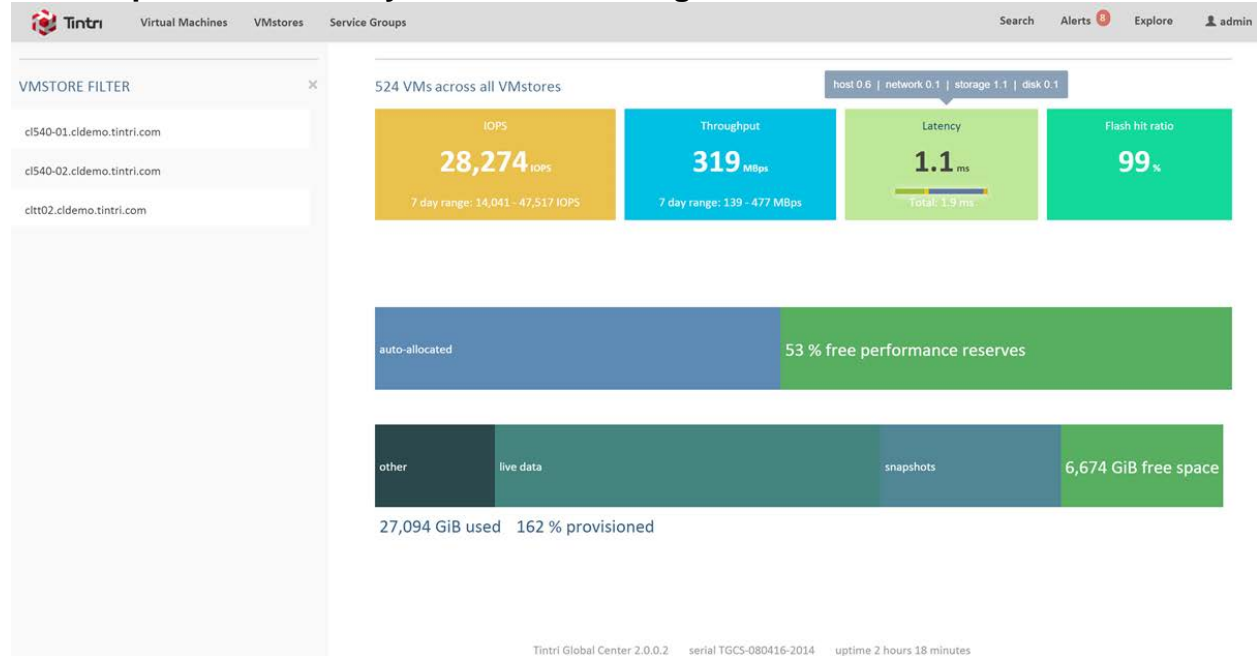
Tintri architecture does not use a conventional storage backbone that includes LUNs, volumes or mapping points. With VMstore capability, you can see a granular visibility or throughput, IOPS, and even latency across the three segments (host, network and

Tegile and Tintri Comparison

storage) for every virtual machine – an end-to-end visibility in a single pane of glass without needing a third party tool.

Most storage vendors, including Tegile, will not be able to see all telemetry data, as illustrated in the screenshot below without relying on a third party tool.

Most Important Telemetry Data Under A Single Glass of Pane



This will reduce the TCO dramatically in a virtualized data centers.

VI. Snapshots Comparison

Most importantly, (1) Tintri can actually snapshot or replicate a single VM at the storage level; (2) Tintri can understand the busy blocks inside an actual VMDK; and (3) Tintri can assign Quality of Service (QoS) to each individual V-Disk, while Tegile and most storage vendors will not be able to accomplish those tasks due to LUNs, volumes or mapping points used in their systems.

Tintri uses 8K block sizes on the back-end for Redirect on Write (ROW) snapshots, meaning Tintri ultimately consume less space than nearly all of its competitors for storing snaps.

Some vendors use large 4MB-512MB block sizes which results in massive snapshot overhead that in turn much less snapshots can be produced and stored on storage. As a result, it will affect recovery time objective (RTO) and recovery point objective (RPO).

VII. Replication Comparison

While Tegile only supports asynchronous replication today, Tintri has been supporting synchronous replication for a while. With Tintri's newest VMstore add-on feature called

Tegile and Tintri Comparison

“SyncVM”, SyncVM can let customers move back and forth between any recovery point without EVER losing any of their other snapshots or performance data.

SyncVM can update hundreds of “child” VMs from a single, refreshed “master.” It seems like a VDI environment, or update one golden image, all 100’s or 1000’s nodes will be updated automatically. Therefore, the update of all “child” VMs happens without physically moving data or reconfiguring storage.

What does it mean to you? A third party vendor might not be needed in order to achieve a VM to VM replication (e.g., Zerto - zerto.com).

VIII. Hybrid Flash and All Flash Array

Tegile offers both Hybrid Flash and All Flash Array, while Tintri offers only Hybrid Flash array. On the other hand, Tintri claims to achieve 100% writes to SSDs and 99% reads from SSDs in its hybrid array for the last four years. The top Fortune 15 deployed Tintri arrays in a production environment might be good enough to support its claim.

IX. Real-time (Inline) Data Reduction

Both Tegile and Tintri offer inline (real-time) compression and inline deduplication with ingest data always randomly written to DRAM, then, landed to SSDs right after inline compression and inline dedupe are complete.

X. Tegile IntelliCare Flash 5 Guarantee

- 5X faster performance
- 5 Year Flat Support Pricing
- 5:1 Data Reduction in highly virtualized environments
- 5 9’s Uptime
- 5X better in terms of \$/GB

Tintri can achieve a similar result.

XI. I/O Blender Effect

Q: I/O Blender effect happens inside the ESXi physical host, while VM-Aware storage resolved the I/O Blender effect only after it occurs via I/O dedicated lanes at a VM level, as illustrated in the screenshot below. In other works, Tintri cannot prevent the original I/O blender effect. Is it right?

A: The I/O blender effect happens in the LUN or Volume on the storage array. The storage array schedules all of its IO at the LUN or volume level. Let’s consider an example of a LUN with 10 VMs: 1 SQL Server database and 9 VDI desktops. At this point in time, the database is processing a massive transaction while the 9 VDI users are trying to move their cursor 1 pixel. Because the array does not know which VM the IO belongs to, the IOs for these VMs get

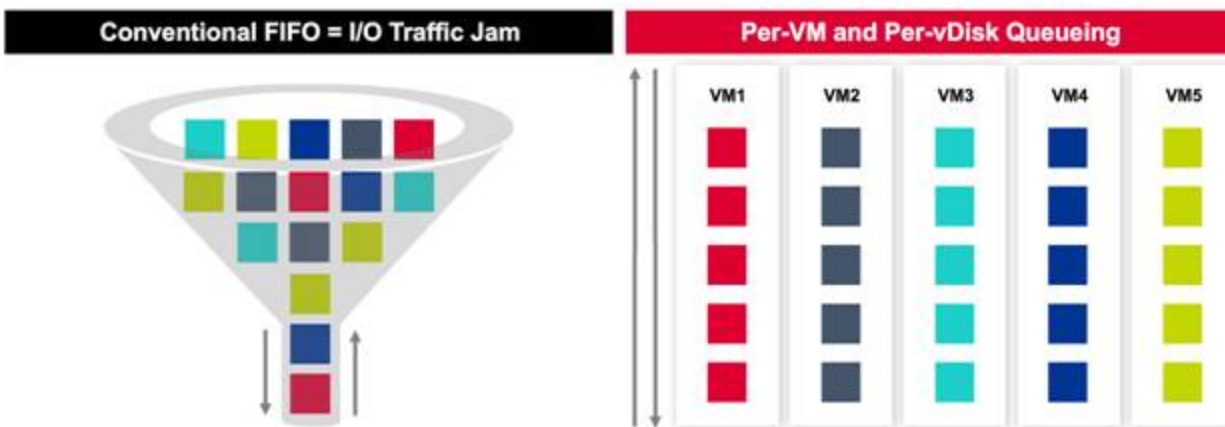
Tegile and Tintri Comparison

blended, so the VDI VMs may get stuck behind the large database transaction—VM users suffer. In the next minute, an antivirus scan may occur on the VDI machines causing lots of IO. Now the database suffers.

With Tintri, there are no LUNs or volumes. We know which IO belongs to which VM so we can schedule the IO fairly and give each VM the IO it requests in an isolated fashion. This is effectively what would happen if you could give each VM its own LUN, but it happens dynamically with Tintri and without requiring storage admins to do anything.

2 VM-Level QUALITY OF SERVICE

Rather than handle workloads sequentially (like conventional storage), Tintri VMstore assigns every VM its own lane to optimize performance



Additional response from Tintri:

The IO blender effect is happening when your applications are all trying to access storage resources at the same time. The IO Blender issue you have with your SQL, is likely a result of using traditional LUNS that result in contention for the same resources. This will be a problem in any virtualized environment, regardless of storage...*IF you aren't using Tintri. EVERYBODY else that is using a LUN or a Volume has contention issues Contention=IO Blender.

If there is another “blender” issue happening in the actual host (as your guru has suggested), that would be host based contention that no storage company can help with. That Host-Based contention is an issue you would likely take up with your Server Manufacturer, and/or Hypervisor provider.

That is much more uncommon issue than folks trying to use traditional storage for virtualization, and having problems caused by the LUN:VM disconnect and the contention associated at that level.

XII. What Can VM-level Quality of Service Do for You

Tegile and Tintri Comparison

Due to Tintri's unique VMSTORE technology, it can provide a VM-level QoS that in turn resolves or minimize a performance issue arising from a ubiquitous I/O Blender effect, as well as LUN/Volume legacy configuration's limitation referenced in a great article via <http://info.tintri.com/VVols-Taneja-Tech-Brief.html>.

Below are snapshots what a VM-level Quality of Service can do for your environment:

VM-level Quality of Service

Tintri is the only storage provider that lets you guarantee (and protect) the performance of individual VMs. In the process we've eliminated the guesswork around Quality of Service. Here's how — in words and pictures — Tintri gives you **unmatched control** and **visibility of performance**.

1 Do you have a rogue VM that is sucking up IOPS?
Simply drag the maximum toggle down to an acceptable level. Watch the "throttle" latency pop right away, but know that you've protected your other VMs from its bad behavior.

2 Do you have a mission critical VM with high performance demands?
Tintri helps you visualize contention, so you can raise MIN IOPS to meet a VM's needs.

3 Do you need to establish tiers of service?
Whether you're setting guidelines for internal customers, or you're a service provider with a growing customer base, Tintri makes it easy to set service tiers and offer chargeback.

4 If there's latency... where does it come from?
Tintri lays the root cause of latency bare with one simple visual. We show you the source, across your entire infrastructure, so you can address it and move on.

VMstore Performance [view trends](#)

IOPS	Throughput	Latency	Flash hit ratio
36,453 IOPS 18,020 - 38,007	950 MBps 444 - 986	4.1 ms 13 ms	100%

Performance Reserves

0% free	auto-allocated	174% used
---------	----------------	-----------

Latency Reserves: Host (0.8 ms), Network (0.2 ms), Storage (4.1 ms), Controller (5.2 ms), Flash (0.8 ms), Disk (0.8 ms), Throttle (0.8 ms)

Let's continue the conversation: Twitter: @Tintri & @Tintri_EMEA | Email: info@tintri.com
www.tintri.com



XIII. VVols vs. Tintri VMSTORE

VVols supports only VMware technology, while Tintri VMSTORE supports multi-hypervisor, including VMware, Microsoft, Redhat, Citrix and Openstack in a single 4U appliance that will support up to 3,500 VMs.

XIV. eMLC vs. cMLC

Tegile uses eMLC, while Tintri uses cMLC that is much lower cost.

Note: Both Tintri and Tegile claim that ingest data always randomly writes to DRAM, then, lands to SSDs only after inline compression and inline dedupe occur.

Tegile and Tintri Comparison

XV. 3,500 VMs per Tintri 4U Appliance vs. Tegile 4U Appliance

Tintri has produced a system that can support 100TB effective storage, delivering 140K IOPS and supporting up to 3,500 VMs in a single 4U in the VMstore T880 system at a list price of \$250,000.00 due to its focus on both performance and capacity efficiency

Tegile can deliver...to be continued...

XVI. Single Pane of Glass Management – A Built-in Management Tool

1. Tintri Global Console (TGC) allows for an administrator to manage all VMs globally from a single pane of glass, and can give global environment statistics as an aggregation or combined statistics. It allows an administrator to dive into a single VMstore or VM, regardless of how many VMs, or where a VM lives.
2. Tintri Array allows for a consistent performance from the 1st TB to the last, and is 100% useable.
3. Enterprise(s) will always have consistent IOPs across VMstores as you add them into a system.

XVII. Variable Block-level Dedupe

Tegile supports variable block sizes. All writes are fully striped with coalesced writes to SSD, similar to Nimble Storage's coalesced layout architecture.

Challenges:

- *“Highly virtualized midsize enterprises with fewer than 200 virtual machines should opt for hyperconverged infrastructure.”* — Mike Cisek, Gartner Research Director and report author. This will make midsize enterprises to consider hyperconverged infrastructure
- VMware Virtual SAN - Cluster server disks to create radically simple shared storage for virtual machines, enabling a resilient, high-performance scale-out architecture that reduces TCO as much as 50 percent. No LUNs or RAID to deal with, a new approach to storage. (Source: <http://www.vmware.com/products/datacenter-virtualization/>)
- [VMware EVO:RAIL](#) - Create VMware compute, networking, and storage resources in this hyper-converged infrastructure appliance, an all-in-one solution offered by qualified partners
- Converged infrastructure also gained 57% year-over-year growth rate in 2014

Tegile and Tintri Comparison

- Many big storage vendors now have either all flash or hybrid flash in their array products. As a result, their customers can boost the SQL performance dramatically by moving **tempdb** and **transaction log(s)** into a flash tier prior to considering the following method – “Flash Cache” installed from each host
- Big server and storage vendors also introduced more intelligence flash storage pool cluster (e.g., Dell [Fluid Flash Cache](#)) at the server access layer via clustering technology. Dell Fluid cache uses write-back caching technology to protect data as it moves from cache pool servers to the SAN. As a result, it can increase the server to server performance to address the ever increasing traffic from east-to-west in modern data centers by eliminating many other bottlenecks from hosts to a storage array (see a diagram of 12 hops via <http://www.virtualizationadmin.com/articles-tutorials/general-virtualization-articles/all-flash-performance-without-buying-all-flash-array.html>). The solution, also known as decoupled storage architecture, produces faster application response time while generating extraordinary I/O performance (e.g., 5 million I/Os) that in turn increases applications performance such as OLTP and VDI in a scalable, low-latency environment. As a result, the many I/O intensive applications' performance will be increased by 3X to 4X. See [Dell Fluid Cache for SAN: Accelerate application performance](#) or [Dell Fluid Cache for SAN](#) for details. What does it mean to you? Enterprises can extend their investment by just adding a hot-swappable PCIe-based flash drive (e.g., Dell 350GB PCIe SSD) into a node to form a cluster to increase the performance dramatically via a flash storage pool cluster with both read and write cache techniques.

"In our SQL environment, a database query that previously ran for over an hour was completed in less than 15 minutes using Fluid Cache." – *Peter Lonis, LeaseWe*"

"Remove typical IO bottlenecks for your OLTP applications. "The Dell Fluid Cache for SAN solution allowed us to scale the database infrastructure to accommodate 4.4 times more transactions per second and 4 times more concurrent users ... with a 90 percent reduction of the average database query response time." – *Kai Yu, Senior Principal Engineering and Oracle ACE Director, Dell Global Solutions Engineering*"

(Source:<http://www.dell.com/learn/us/en/555/solutions/fluid-cache-licensing?c=us&l=en&s=biz&cs=555> and <http://www.dell.com/learn/us/en/555/videos~en/documents~fluidcache-for-san-tech-video.aspx?c=us&l=en&s=biz&cs=555>)

Dell new innovative fluid flash cache technology has addressed its customers' concerns by providing investment protection plan from currently deployed architectures.

- Many enterprises are looking for a hybrid cloud such as [AWS](#), Microsoft Azure and etc. for Platform as a Service (PaaS). AWS is able to achieve a “one click” of button to bring a live VM in production. See AWS keynote at <https://www.youtube.com/user/AmazonWebServices/Cloud> for details. Many

Tegile and Tintri Comparison

cloud providers, as a matter of fact, can provide more secure services than on-premises due to the following certifications enforced at a public cloud such as AWS, while most private cloud may not have all certifications in house.



Source: AWS Summit Keynote July 09, 2015

As a result, over 1,700 government agencies and over 5,000 educational and academic institutions rely on AWS cloud as of today (see below for your reference):



Source: AWS Summit Keynote July 09, 2015

- Tegile has Cisco UCS reference architecture with a 314TB clustering storage pool capability

Note: Tintri claims that its storage has received Cisco UCS IVT certification. See <https://www.tintri.com/partners/alliances/cisco> for details.

Conclusion:

Tintri has produced a system that can support 100TB effective storage, delivering 140K IOPS and supporting up to 3,500 VMs in a single 4U in its VMstore T880 system at a list price of \$250,000.00 due to its focus on both performance and capacity efficiency, while delivering a consistent performance from the 1st TB to the last, and is 100% useable.

Tegile and Tintri Comparison

Tintri has ease of use in configuration standard point of view because Tintri does not rely on LUNs/Volumes with RAID level legacy configuration. As a result, managing a massive spreadsheet (which LUN or Volume is associated with a VM) might not be necessary.

Tintri provides unique features: (1) Per VM Management; (2) Per VM QoS; (3) Per VM Replication; (4) Key Managed Encryption; (5) Multi-hypervisor supports on a single 4U appliance; and (6) 100TB effective storage and 3,500 VMs in 4U appliance.

Tegile does not charge any additional cost for software, meaning all features are bundled in the platform once the appliance(s) is purchased. It has IntelliFlash flash optimized software architecture that claims to support 1,000 VDI nodes per 4U system and can

- increase the SQL performance up to 10X
- reduce storage footprint by 50%
- eliminate disk I/O latency
- Achieve SQL enterprise-like features with SQL standard, thus, reducing SQL licenses for massive cost reduction

Both Tegile and Tintri arrays can bring down the CAPEX from 50% to 90% for less storage or initial acquisition costs due to their real-time data reduction, as well as OPEX that includes reduced footprint, power and cooling, in addition to less switches ports required to support the same storage capacity,

In summary, it all depends on what your environment need.

Recommendation:

It really depends on each environment. For example, Tintri might be a good choice if your environment does not need more than 100TB usable disk storage for a long time, while Tegile might be a good choice if your environment does need more than 100TB usable storage, but under 314TB RAW disk space.

Please keep in mind that Tintri does not rely on LUNs/Volumes, while most storage vendors must rely on. Tintri's solution has less complexity due to no LUNs and Volumes.

On the other hand, Tegile can provide an active/active storage pool, while Tintri cannot. Therefore, Tegile has a better scale out capability than Tintri.

Any organizations should perform a detailed analysis. Otherwise, they might be likely to make wrong decisions and may enter into inappropriate contracts.

At the end, it all depends on your environment.

Below is a recommendation from Arun Taneja, founder at Taneja Group and a renowned consulting analyst:

Tegile and Tintri Comparison

“As I can see it, you are dealing with three good products here. If Hyper-V is all want, and VM centricity is important, I think Gridstore will give you the best performance. If VM centricity is important but so is multi-hypervisor capability (and performance is important but not the driving factor), Tintri is the way to go. In my view, Tegile is traditional storage built on top of ZFS file system. As a result, it enjoys low cost and decent characteristics all around but I do not see it as VM centric (please confirm this with them) or blazing performance that comes from being a filter driver (Gridstore). Now it is up to you to decide what is important and make a selection.

Overall, you are dealing with choosing between three good players. “

Buyers may refer to the following to make a decision:

Tegile is Best-in-Class in 2015 DCIG Buyer’s Guide for Midrange Hybrid Storage Array?

Source: <http://www.tegile.com/company/awards/>

Tintri: [InfoWord.com 2015 Awards – Technology of the Year!](http://www.infoword.com/2015/05/20/tintri-technology-of-the-year/)

Recommended Reading:

1. [Magic Quadrant for Solid-State Arrays - 23 JUNE 2015 ID:G00270206](#)
2. [Magic Quadrant for General Purpose Disk Arrays – 20 November 2014 ID:G00263357](#)
3. [Simplify the Midmarket Data Center With Hyperconverged Infrastructure Solutions - 24 June 2015 ID:G00275308](#)
4. [Erasure coding vs Raid as a data protection method](#)
5. [L.A. CTO: We are actively upgrading city’s IT in the cloud](#)

Acknowledgement

Thanks for the TANEJA GROUP & Tintri webinar dated June 30, 2015 and a great article at <http://info.tintri.com/VVols-Taneja-Tech-Brief.html>.

Thanks for James Rudden and Tyler Nasseth at Tintri to provide a Q&A session on June 29, 2015.

Thanks for Nick Stocking and Megan Wardell at P1 Technologies for hosting an one-hour Tintri Technology Overview on July 1, 2015.