Tintri

Three Things you can ONLY do with VM-aware Storage

There's a disconnect in your data center. You've virtualized your applications and your teams are thinking and working in virtual machines. But your storage hasn't materially changed in decades—it still uses the same constructs (LUNs, volumes, RAID, striping) that were introduced for physical workloads. And that disconnect between virtualized applications and physical storage is costing your business time and money.

You might be exploring (or deploying) new storage technologies to address the pain, but it doesn't matter if those products are all-flash or hyperconverged—if they're not VM-aware they're not going to solve the root cause.

Only VM-aware storage (VAS) is specifically built for virtualized applications, stripping out the complexity of LUNs and volumes so you can manage only the VMs that matter. When you can take every storage action (replicate, clone, analyze, etc.) on individual virtual machines, you save time, money and sanity.

And that's exactly why VAS is increasingly visible and sought after. The percentage of data center professionals that have heard about VAS has leapt from 60% to 85% in just the past 12 months, and 79% say that VAS capabilities are extremely important to their next storage purchase.

It's also why storage companies are jumping on the VAS bandwagon. Many now tout their technology as VM-aware, VM-centric and other creative interpretations. Fortunately, there's a single question test you can use to determine whether or not storage is VM-aware:

Has the storage system removed all LUNs or volumes from its architecture?

YES.

Congratulations, that's VM-aware storage.

NO.

Unfortunately, it's not VM-aware storage.

And when storage is truly VM-aware, it can do many things that commodity all-flash and hyperconverged cannot entertain—the following three examples (with accompanying anecdotes) top the list:

1. See the source of latency across host, network and storage to pinpoint problems

Tintri's UI shows you every individual virtual machine stored on the platform. Search and scroll over any one, and a simple visualization shows you its latency broken down across host, network and storage. Never confuse symptoms for the root cause—the end-to-end visibility of VM-aware storage lets you solve problems in a few clicks.

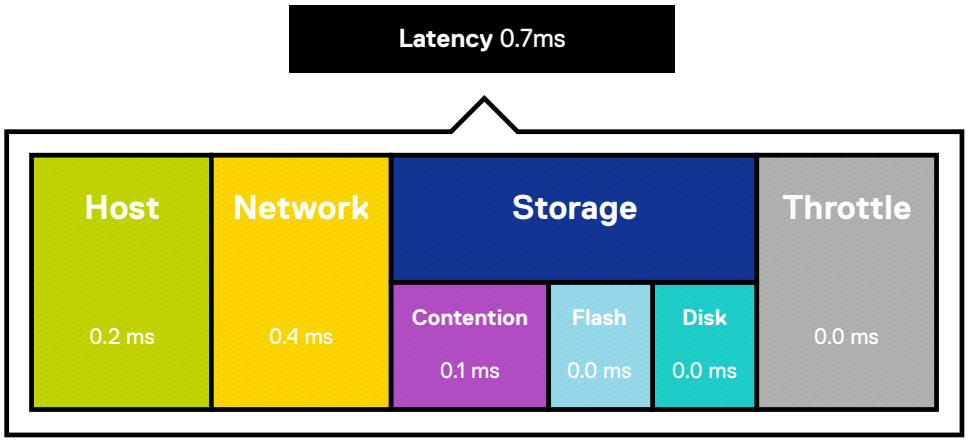

"Every Monday morning used to start with a finger-pointing meeting. As the storage admin, I would gather in a room with my counterparts from compute and networking. We'd work through a list of performance issues and we'd all point fingers at one another. Now I've got Tintri, and I can look at my UI and see the exact source of latency across compute, network and storage. My peers count on me for insight, and we've permanently canceled the finger pointing meeting."

Storage Admin, Fortune 100 Financial Services Company

"Tintri's VM level visibility is an unbeatable advantage. An example from this morning... we received a trouble ticket from an end user whose virtual desktop was performing poorly. We looked up the specific VM in Tintri and determined it was a CPU / Memory bottleneck. From user complaint to rectification in less than 10 minutes across a complex network infrastructure—that's why I love Tintri."

Geoff Grice, Head of IT, Coal Marketing Company

2. Eliminate noisy neighbors with VM-level Quality of Service (QoS)

VMs stuffed into a LUN share the performance resources and policies assigned to that LUN. When one of those VM's gets overly demanding or goes rogue, you've got a noisy neighbor. Tintri eliminates LUNs, instead assigning every VM its own lane. There's no conflict over resources or policies and therefore, no noisy neighbors.

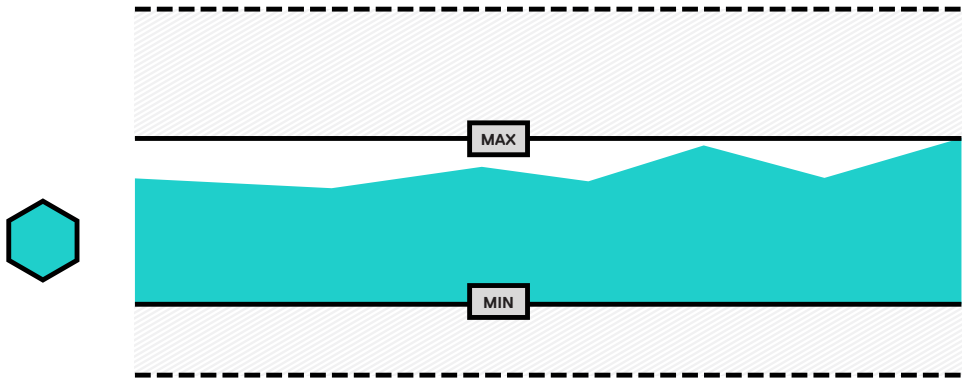

"On a few occasions Windows print servers get into a state where they generate IOs as fast as they can but do not go out of service. Since the service stays running the service Admins do not know there is an issue. That load would be problematic to other VM and be difficult to identify in the past. Now with the Tintri's they contain the noisy neighbor and allow us to identify the problem VM within 3 mouse clicks."

John Ward, Systems Architect, University of California Irvine

3. Apply VM-level analytics to identify issues with individual applications

As part of labeling themselves VM-aware, some storage providers claim they deliver VM level analytics. Under the covers, that usually means some basic math to average latency, throughput and more across all the VMs in a LUN (and so, it's still LUN level analytics). True VM level analytics allow you to dig into the performance characteristics of an individual VM with precision, not estimations or correlations.

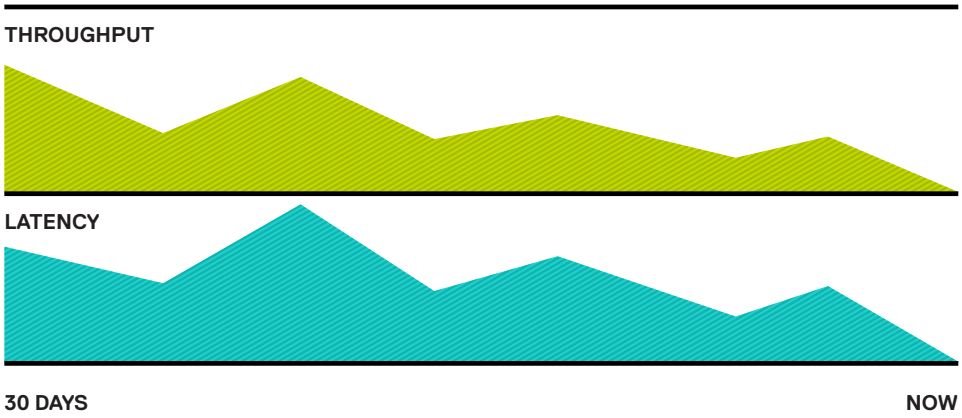

"We were experiencing slow performance writing to a database VM that was hosted on our old NetApp. Using vROps we tried to find the correct metrics to identify the bottleneck, but we could not pinpoint the issue. So, we moved the VM to our "Proof of Concept" Tintri and within a few minutes we were able to show the application developers that there was less than 1-millisecond latency between host and storage. They took a closer look at the application's coding, found the issue and optimized the process. Without Tintri they would have had us hunting down metrics for days."

Adam Way, Infrastructure Analyst, Garmin UK

Summary

While many storage products wear the VM-aware label, it doesn't take much to assess its authenticity. If the product's architecture is still built on a foundation of LUNs and volumes, it can never be fully VM-aware.

Truly VM-aware storage is built specifically for virtualized applications. It provides you visibility across your infrastructure, control over performance and complete analytics at the VM level. If your data center is experiencing pain—poor performance, complex management and cost over-runs—you can't solve it with all-flash or hyperconverged alone. Only VAS can resolve the disconnect and restore simplicity.

Tintri

Global HQ
303 Ravendale Dr.
Mountain View, CA 94043
United States
+1 650-810-8200
info@tintri.com

EMEA HQ
Fountain House 10th Fl
130 Fenchurch Street
London EC3M 5DJ
+44 (0) 203 053 0853
emea@tintri.com

APAC HQ
9 Temasek Boulevard
Suntec Tower 2, #09-01
Singapore 038989
+65 6407 1359
apac@tintri.com

Japan HQ
Level 6, Kishimoto Building
2-2-1 Marunouchi, Chiyoda-ku,
Tokyo 100-0005 Japan
+81 (3) 6213-5400
info.japan@tintri.com

www.tintri.com