

THE ESSENTIAL GUIDE TO

VM-aware Storage (VAS)



Tintri

Table of contents



01 The disconnect,
defined

02 Enter VM-aware
storage (VAS)

03 The business
benefits of VAS

04 Words of wisdom
from VAS converts

05 An interview
with Kieran Harty

06 What to do next

A letter to the storage visionary

Fact: Organizations spend too much time and money on their storage.

There are plenty of proof points, and you're one of them—if your storage met your needs, you wouldn't be reading this.


Unfortunately, you've already invested too much time and money troubleshooting performance, provisioning new platforms, carving and tuning LUNs and volumes, and sitting on hold with storage vendors. You've lost evenings, weekends and sanity to storage that doesn't deliver what you need. And you're ready to accept that there's a better way.

If your organization is highly virtualized, or if you're planning a virtual-first strategy for your organization, you cannot meet your objectives with conventional LUN and volume-based storage. Conventional storage architectures that were built for physical workloads decades ago are still used by both legacy providers and storage newcomers today. Instead, you need storage specifically built for virtualization and cloud: VM-aware storage.

Tired of using physical storage for your virtualized applications? Want predictably fast—even guaranteed—performance? Ready to save your time and money? Good—let's get started.

01 The disconnect, defined

Virtualized applications drive value for your business —they are the backbone of your virtualized data center— but conventional storage doesn't let you interact with them on a 1:1 level. VM-aware storage does.



The disconnect in your data center

If your data center still runs on a conventional storage architecture built for physical workloads, you're not realizing the full value of virtualization.

DATA POINT:

Today, 77% of organizations have a virtualization-first strategy for workloads, and the percentage of virtualized applications has leapt from 2% to more than 70% in the past ten years.

Conventional storage was built to manage physical workloads via a 1:1 relationship. But virtualization forces conventional storage to manage virtualized applications via a 1:many relationship. To handle the “many”, conventional storage lumps those applications into LUNs or volumes. And those many applications get spread across different LUNs or volumes, with different characteristics and varying results.

What's the result? A tangled mess of mapping and what many call the “I/O Blender” where VM requests for I/O get all mixed up and are forced to fight for resources with no sense of prioritization.

How do organizations purport to clean up this mess? With lots of manual effort—and so it's little wonder that over the past decade, while capital expenses in the data center have grown 3x, operating expenses have grown 8x.

Let's break down that disconnect right now:

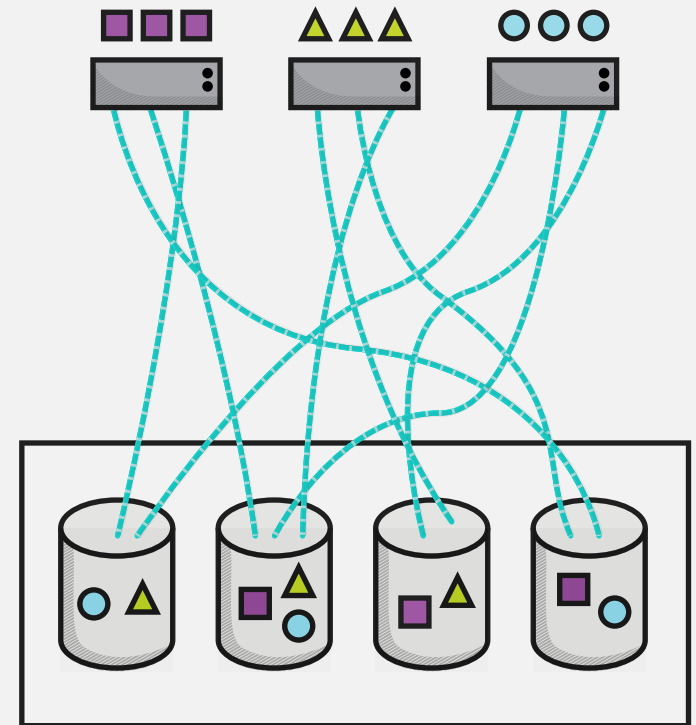
You:

An organization who's virtualized your applications and whose team members are all thinking in virtual machines (VMs).



Conventional storage:

Aging architecture made for physical workloads, with logical unit numbers (LUNs) and volumes that don't let you work with your individual VMs.



Conventional storage has very real costs that can slow down your business and increase unnecessary expenditure. As we've mentioned, it's costing you:

Time

You're getting your hands dirty—a lot. Your most talented people waste time on menial tasks: shuffling VMs, tweaking policies and troubleshooting tickets. Mission critical applications languish in IO queues, or compete with other VMs that want the same resources. And your data center challenges interrupt or even preclude your key business initiatives—from software development cycles to IT-as-a-service. Sure, you feel busy—but with all that opportunity cost, do you feel productive?

DATA POINT

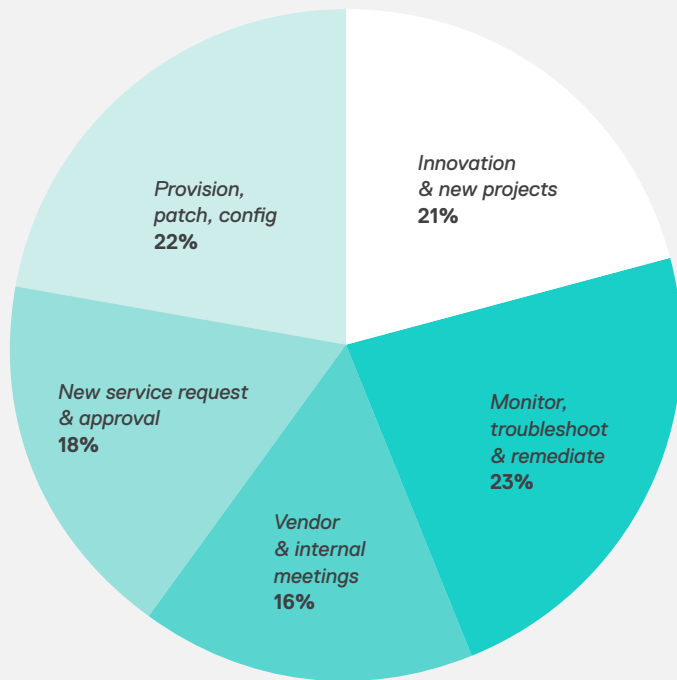
Today, IT staff spend far more time on troubleshooting, working through service requests and wasting away in meetings than on innovation and new projects. This isn't a formula for success; it's backwards. You need to free your people to work on what matters—and that starts with the virtualized applications that drive your business. SEE FIG. 1

Money

With conventional storage, you're blind. You never know how much performance you have in reserve, so you either wait until you hit a cliff, or over-provision to maintain a buffer. As your footprint gets larger, you need more power & cooling, and more third-party software to manage it all—expense after expense. It's a storage racket: buy more conventional storage to escape the very pain that it caused. You might as well buy protection from the local mafia don.


FIG. 1

How IT professionals spend their time



Source: IDC Survey, August 2014

02 Enter VM-aware storage (VAS)



This begs the question: if the disconnect is so obvious, why haven't storage vendors solved it?

Well, it's hard.

To eliminate LUNs, volumes and other remnants of physical storage, vendors have to rewrite their entire file system. For legacy vendors, we're talking a massive undertaking, and for emerging vendors moving fast to capture share, it's an unacceptable delay.

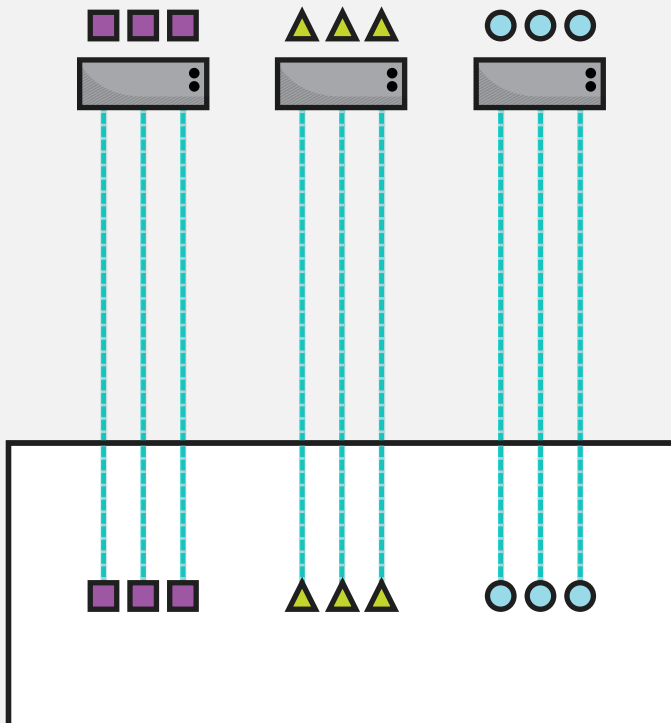
But over the past few years, some pioneers have stepped up to the plate. They've recognized—and vowed to eliminate—that disconnect between virtualized applications and physical-era storage. They've set out to reimagine storage from the ground up, building storage specifically for virtualized enterprises and cloud: VM-aware storage (VAS).

VAS directly connects your file system and storage, using not LUNs and volumes, but rather individual VMs as the level of abstraction. With VAS, admins can manage, replicate, clone, snapshot and analyze on an individual VM level—no more tuning LUNs or keeping a sprawling spreadsheet of VMs. VAS bridges the disconnect between your storage and your virtualized applications, and in turn saves your business time and money.

VAS Capabilities

Since working on the VM level is completely unprecedented, it's tough to articulate the magnitude of the difference. So let's take a look at what we call the Three A's of VAS, or how VAS improves your daily storage workflow by simplifying VM-aware action, enabling powerful automation and offering granular analytics.

(You can also think of this as “How VAS makes me the smartest person in the room.”)



Action

Here's what it means to take every storage action on the VM level.

First, guarantee a critical application's performance with VM-level quality of service (QoS). No stress—just drag the MIN IOPS up to guarantee that the application will always have access to sufficient resources. SEE FIG. 2

Got a rogue application that is behaving badly? Drag its MAX IOPS down, and it won't consume more than its share. SEE FIG. 3

And hey, we know this sounds way too easy—and that makes sense.

The conventional storage you're used to just can't give you this level of granularity and control. With clumsy LUNs, you've had to change settings for the entire LUN, affecting all the (very different) VMs contained inside.

That's like giving all your friends the same gift, even if you know not everyone will like it.

When it comes to replication, conventional storage admins have to replicate entire LUNs, again affecting every VM within. With VAS, you just click the exact VM for replication and apply policies that affect it alone. That precision can reduce WAN usage by up to 95%.



Still sound too easy? Head to try.tintri.com to see (and use!) the Tintri VM-aware storage interface. Guess how many clicks it takes to replicate a VM. (Hint: you can count 'em on one hand.)

FIG. 2

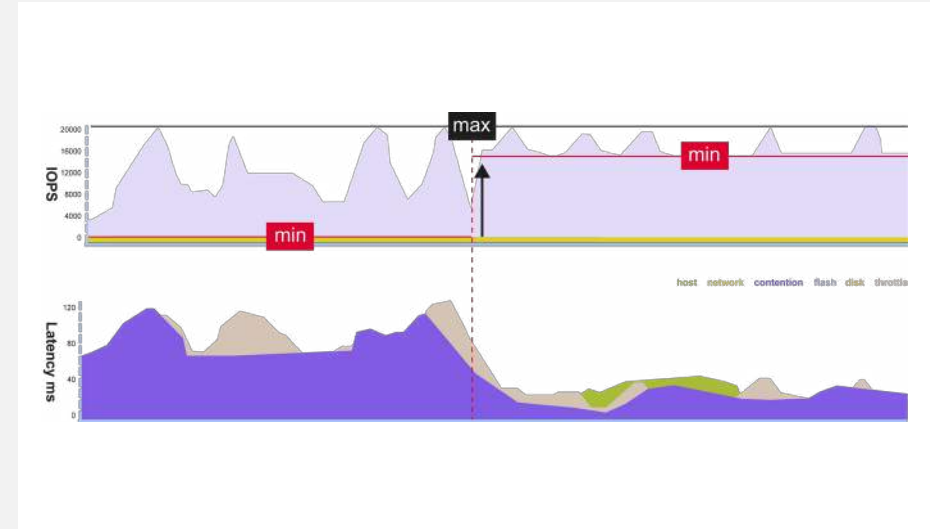
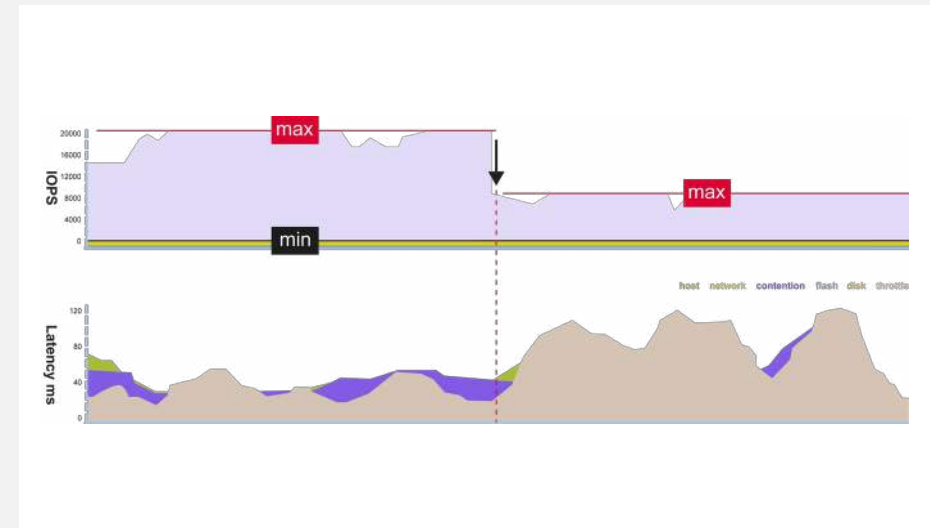


FIG. 3



Automation

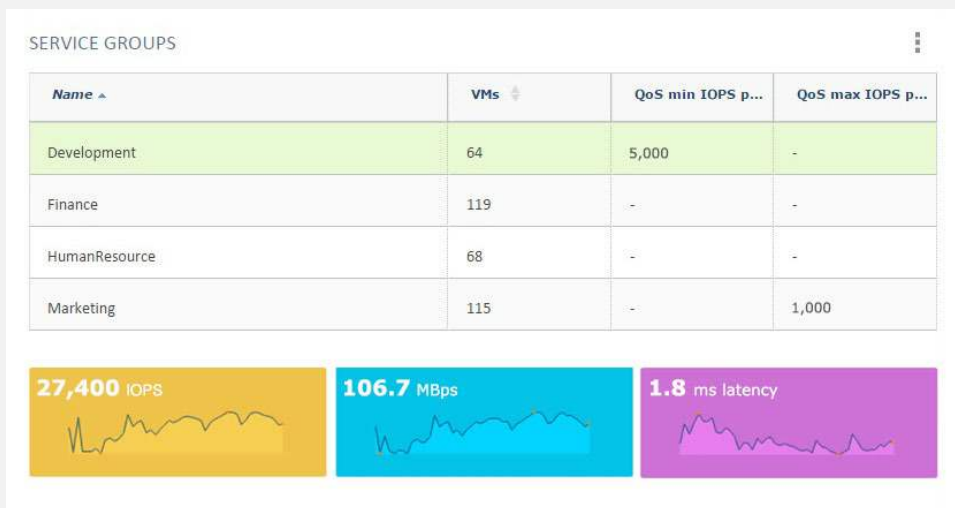
LUN storage and spreadsheets go hand in hand, but they're hardly a perfect pair. You might be used to mapping VMs to LUNs and tossing them in a never-ending .CSV, but all that scouring your spreadsheet for VMs when you need to shuffle them for performance just wastes your time.

VAS is completely different. Let's start with the absolute basics. Many admins burn hours manually realigning their VMs with their storage—a tedious task with no added value and a high opportunity cost. That's exactly the kind of action that Tintri automates, and that makes admins smile ear to ear.

You can easily automate even more advanced storage tasks.

Let's talk about service policies. With VAS, you can gather a set of VMs (based on their characteristics or other criteria) into a single service group. Then, you can apply policies to that service group—from QoS (as above) to replication. Now, as you move VMs from one service group to another, the original policies carry with them, or can be overwritten—automatically. That's the power of VAS automation. SEE FIG. 4

FIG. 4



If you're still a doubting Thomas, check out PART FOUR for customers and vExperts who can vouch for Tintri VM-aware storage. Or, check out this handy demo video page. We're as transparent as VAS.

Analytics

With conventional storage, you can't even see through a LUN to its resident VMs. With VAS, you have transparency, and thus unprecedented granularity in your storage analytics. For every single VM, you can see:

- Capacity
- Throughput
- Performance data (historical and real-time)
- Latency across host, network and storage SEE FIG. 5

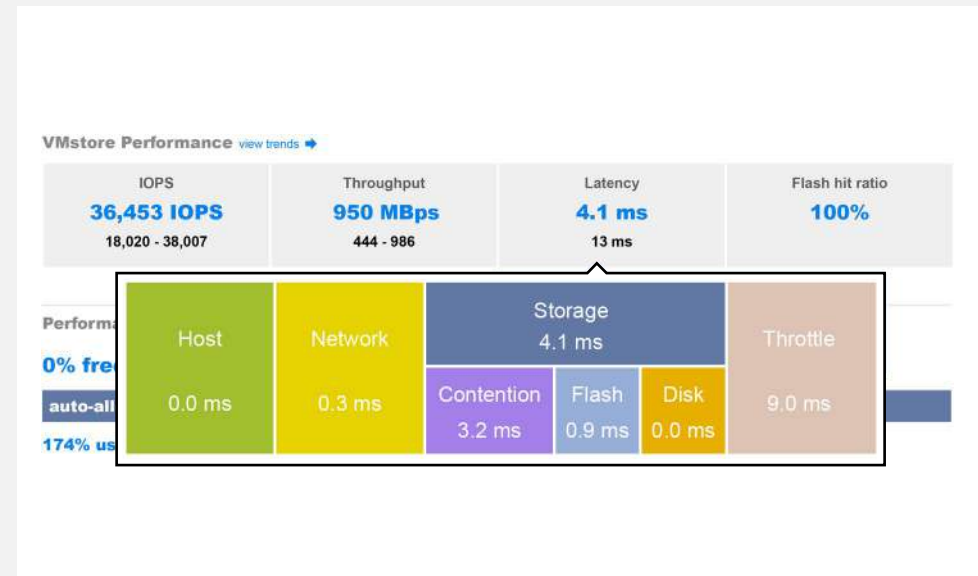
So when a latency issue arises, there's no finger-pointing or support calls. Check out the graphic above: scrolling over platform latency, or just the name of an individual VM, produces a pop-up graphic that breaks down latency by host, network, storage and throttle. You can immediately identify and address the root cause of any latency. Got LUNs instead? You can forget about accurate analytics.

Meanwhile, though most conventional storage can estimate when you'll run out of capacity, VAS knows exactly when you'll run out of capacity, **throughput** and performance. And it can guide you to the right mix of hybrid and all-flash storage, ensuring that you don't over-provision or overpay for what you really need.

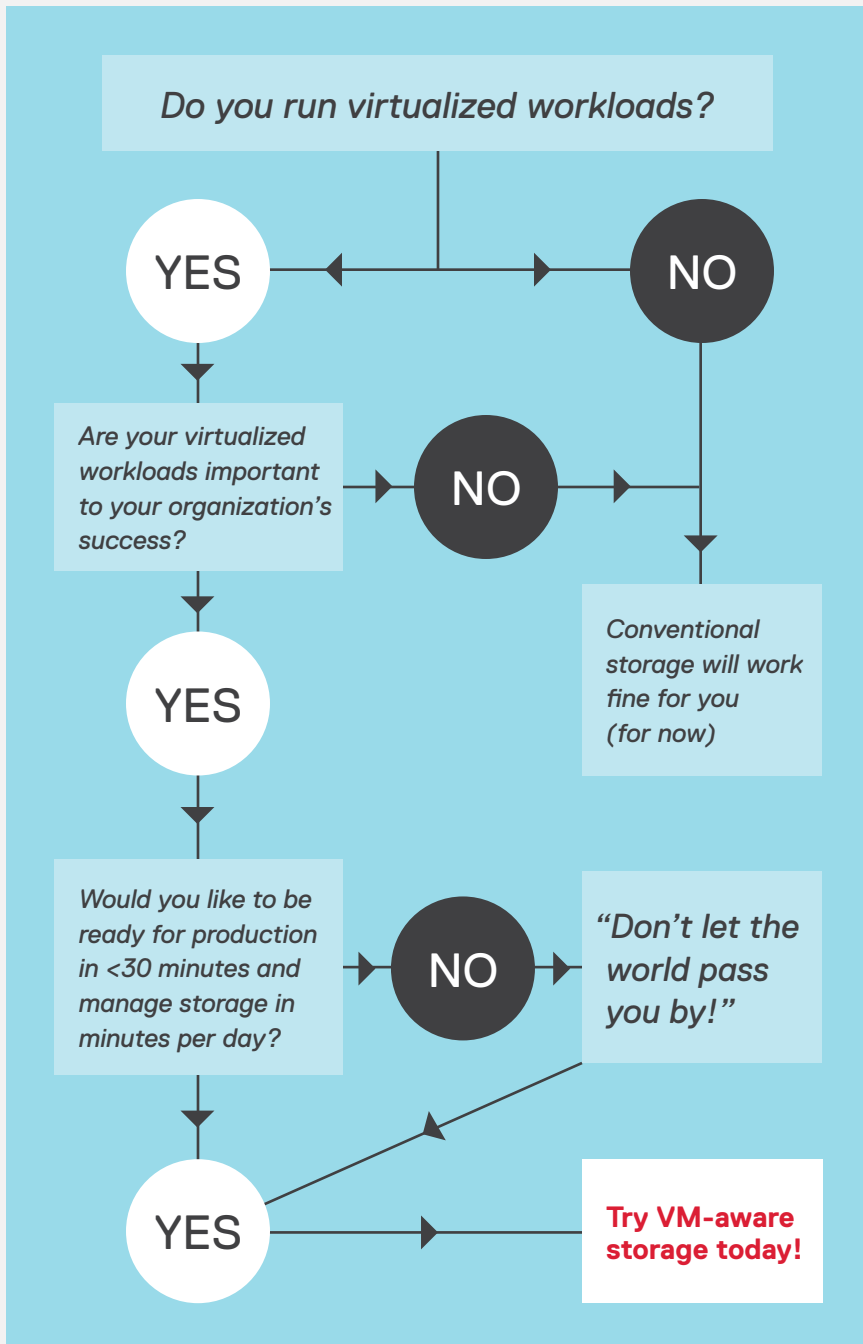


Yeah, I know some other guys offer "VM-level analytics." But here's a brain teaser: if their storage operates in LUNs and volumes, how can their analytics crunch in VMs? Averages and correlations, that's how. They'll get a VM's "performance" by averaging all the VMs in its LUN. But what's the average net worth of everyone at a McDonald's when Bill Gates walks in? Yep, averages don't help much—real VM-level analytics do.

FIG. 5



Do you need VM-aware storage?:



Now, that applies to growth of your current workloads, but also to modeling new workloads. VAS lets you use the historical behavior of applications (up to three years) to define storage needs six months into the future. That's a far longer storage horizon that lets you to plan better and procure better.

DATA POINT

In 2013, fewer than 50% of organizations were familiar with VAS. In 2015, 85% of companies stated they were familiar with VM-level storage management, and 79% cited VAS capabilities as central to their storage purchases over the next 12 months.

Source: IDC Storage for Virtualization, July 2015



Can VAS analytics predict stock returns? No. Can VAS analytics predict the outcome of the Super Bowl? Maybe. Can they help you make better decisions about how and when you spend your storage dollars? Definitely.


Since engineers built VAS specifically for virtualization and cloud, VAS especially enables these four use cases:

Server Virtualization



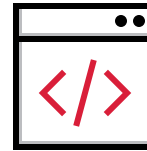
VAS gives each server VM its own lane. We're talking Exchange servers and flash-hungry Oracle databases, each avoiding I/O traffic jams. Plus, since they're not on LUNs, you can drill into an individual server VM's performance to see the source of any latency, set QoS and replicate it in seconds.

VDI



VAS is tightly integrated with the ecosystem—including VMware Horizon (with View) and Citrix XenDesktop—so you can guarantee performance and flip storage from VDI anchor to enabler. By isolating each VM in its own lane, you eliminate the contention that stems from boot storms and virus scans. And since VAS offers unprecedented VM density—thousands of VMs on one platform—you can use less space, sip less power and spend less money.

Test and Development



VAS enables you to deploy updates to developers in seconds, and roll out and tear down environments of thousands of VMs in minutes. Plus, with VM-level snapshots, you can forego days of restoring entire LUNs and volumes from scratch, restoring VM snapshots in just seconds. And it's so simple that the test and development team can own their storage footprint.

Private Cloud



VAS lets you deploy the storage backbone for your private cloud in less than 30 minutes. It isolates every VM in its own lane to eliminate noisy neighbors. You can group those VMs and apply policies that follow them as they move between devices and/or geographies. And you can manage a massive footprint—hundreds of thousands of VMs—from a single, centralized user interface.

03

The business benefits of VAS

You've learned about VAS capabilities—now let's explore how they will benefit your business.



VAS benefits

Time


What's the scarcest resource in your data center? Time.

It's hard to measure the opportunity cost and value of hours in the data center and/or in DevOps. Here's how we've seen VAS save organizations TIME.

People priorities. Every hour your data center team spends maintaining and managing storage is an hour they are not spending on strategic projects like advancing hybrid cloud or IT-as-a-service. Make sure you know how much effort your storage will take to install, configure, manage and troubleshoot day-to-day.

Specialists vs. generalists. Conventional storage requires an admin be versed in LUNs, volumes, queue depths and more. VAS requires an admin understand VMs—that's it. And that means that you don't need to hire (or can re-purpose the expertise of your) storage specialists; IT generalists can manage VAS in a small fraction of their day.

Business agility. Your test and development team may need to spin up VMs for testing exercises and then tear them back down again. With conventional storage, those actions can take days or weeks... and test and development has to wait. VAS takes minutes, so storage never slows the activities that generate value for your business.



“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.”

Fortune 100 Financial Services Company

VAS benefits

Money

Saving time certainly saves you money, but VAS will impact your bottom-line in other ways too. In our experience, VAS will directly save dollars in the following ways:

Capital expense. Since VAS is built specifically for virtualization and cloud, you can use capacity and performance more efficiently. You can pack VMs up to 10 times more densely to reduce your physical footprint, and you never have to over-provision to maintain a performance buffer.

Ecosystem integration. The design of VAS allows for tight integration with the virtualization ecosystem. For example, you can run multiple hypervisors (and even multiple workloads) on one storage device, rather than buying to separate nodes for separate hypervisors. And you won't need expensive software or custom development to make your storage sing.

Scale. Chances are your virtual footprint is growing—how well equipped is your storage to grow with your needs? Consider hyperconverged; it's all the rage, but forces you to grow server and storage in parallel when business is rarely that symmetrical. Conversely, you can add VAS as needed—standalone or converged with best-in-class server and networking.



VAS saves you time and money, sure, but it also saves you sanity. It's just not sane to keep using conventional storage for virtualized applications, and I for one like sane storage.

“We achieved a significant reduction in OPEX by moving to Tintri,” noted Gugliucci. “The Tintri platform is so easy to use, we were able to eliminate the specialized training needed to manage our previous storage environment. We can now reassign our storage admins to more strategic IT projects.”

Infocert

04 Words of wisdom from VAS converts

The promise of VAS is borderline unbelievable.
But don't take our word for it—here are VAS users raving
and sharing their favorite stories.



Shrink your storage footprint

“What used to take 16 rack units to accommodate our performance needs now takes only three rack units, considerably reducing power and cooling needs for our infrastructure.”

University of Arkansas

“[VAS] provides the performance we need with a small footprint – saving space, power and cooling in our lab and allowing us to scale efficiently. Compared to our previous storage, Tintri VMstore can accommodate twice the IOPS at less than a third of the latency in one fourth the footprint.”

F5 Networks

Provision VMs like a boss

“It used to take as long as two weeks to roll out thousands of VMs for new customer events. With Tintri [VM-aware storage], it only takes minutes. Plus, the additional VMs don’t strain the system’s ability to deliver high performance once provisioned. Tintri’s speed in provisioning is a huge strategic advantage.”

Department of Defense Cyber Range

DATA POINT

50% of storage professionals say latency is their sharpest storage pain. 41% say capital expense, 40% say it’s scaling to manage growth and 39% say it’s manageability.

Source: *State of Storage 2015*

Predictably fast performance

“[The customer] sent me an email that said, ‘I don't know what you did over the weekend, but the process ran 300% faster. The time to complete went from two hours down to 28 minutes.’ That time savings translates directly into significant business value for this customer.”

Envision Technology Advisors

“Before Tintri, our customer’s applications exhibited an average latency of 25ms, which would spike to 120ms under load. Tintri is delivering sub-microsecond latencies on average. The biggest spikes we are observing only reach 2ms during peak workloads. The 60:1 reduction in latency was quite refreshing.”

Amatis Networks

DATA POINT

2 out of 3 storage professionals have onboarded a new storage vendor in the last 24 months. 56% are evaluating hybrid storage vendors, and 42% are evaluating storage that is VM-aware.

Source: State of Storage 2015

Keep IT simple stupid

“[I said,] ‘It just can’t be this easy. I’ve worked in the storage industry for 25 years, and I’ve never put a storage platform in in 35 minutes, it just doesn’t happen.’ So we used [VM-aware storage] for a while, seeing how it performed, expecting that it just can’t be this simple. But it was.”

South Eastern Health and Social Care Trust

“Our business has grown 20x over the past few years, but Tintri allows us to manage our storage footprint in a fraction of the time we used to spend. Our VMware admins manage our Tintri VMstores with very little effort.”

Greencloud

DATA POINT

On average, storage professionals currently spend twice as much time tuning LUNs and volumes than they do on strategic projects.

Source: State of Storage 2015

05 Interview with Kieran Harty



Kieran Harty is CTO of Tintri and the godfather of VM-aware storage. In his own words, here's how it all started... and what's next.

Q: Where did the idea for Tintri—and VAS—originate?

A: From 1999 to 2007, I led engineering for VMware with a team that helped define virtualization. It became apparent that conventional storage was ill-equipped to support the virtualization wave crashing through the data center—their products were all built for physical workloads. And so a small group of experts came together to build storage specifically for virtualized applications.

Q: Tintri picks on LUNs a lot. Why do you see LUNs as so evil?

A: LUNs make a lot of sense when it comes to managing physical workloads—but do nothing but complicate virtualized workloads. One obvious piece of evidence is the spreadsheets that storage admins often maintain—showing how hundreds of VMs are mapped to different LUNs. When an admin places tens or hundreds of VMs into a LUN, they lose visibility of their individual needs and performance. VM-aware storage removes all that pain and complexity, letting the admin see and take action on each individual VM, so they can manage storage in a fraction of the time.

Q: If VAS eliminates so much work for the storage admin, doesn't it put the admin's job at risk?

A: Not at all—every storage admin we've met has a massive queue of projects, but they have to spend their time fighting fires and addressing trouble tickets. Tintri frees them so they can spend their time on the strategic projects that are more interesting to them and more valuable to the organization.

Q: When explaining VAS to the uninitiated, what is the “a-ha” moment when they understand it is different from conventional storage?

A: Seeing is believing and the “a-ha” moment happens as soon as they see the Tintri dashboard. The home page has a simple bar chart that shows performance reserves—something that’s simply not possible with conventional storage. Then they see a list of the individual VMs that are using the most capacity and performance. That level of visibility is a game-changer in the data center, and it’s something everyone understands—not just the storage admin.

Q: What is your favorite customer quote or anecdote?

A: There are so many, but one jumps to mind. Tintri was in a bake-off with another emerging storage vendor. The customer first deployed our competitor’s product and immediately a few workloads began performing erratically. The customer assumed the storage was to blame. They called us and we suggested they install and configure their Tintri right away. Within 30 minutes they had Tintri fully ready for production, and had moved over the workloads in question. We looked at the results together and were able to show them that our competitor’s storage product was working just fine—the problem was a networking issue. They loved that Tintri was able to instantly spot the root cause, and so even though we proved the competition worked great, it was obvious the visibility VAS offered was a differentiator. Needless to say we won the bake-off.

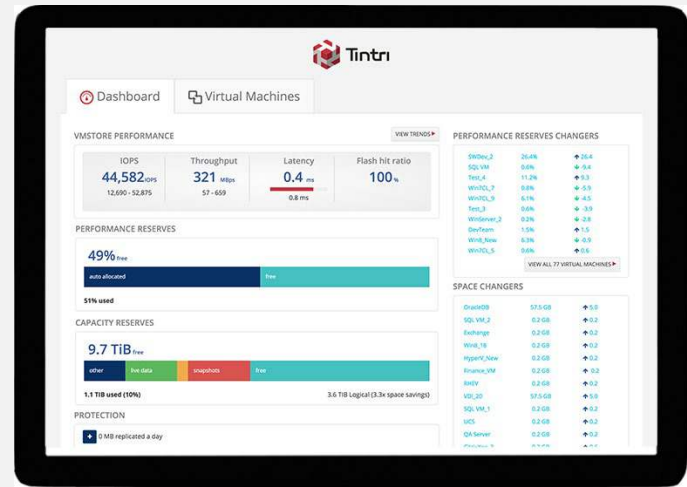
You've now been initiated. You know a little more about VAS, but if you want to get more serious, what should you do?

Try it yourself—and test your skills

This microsite shows how easy it is to use VAS. Work through four simple tasks in less than 90 seconds.



Don't worry, I'll be there to guide you through it every step of the way. My personal record for completing the four challenges is 21 seconds... top that, hot shot!



Watch an expert in action

Tintri's engineers, evangelists and technical marketers have built out a large and growing library of 2-to-5 minute product demo videos. We'd record longer videos, but with VAS it's just not possible. Greatest hits include:



Set QoS for an individual virtual machine



Tintri for Hyper V: Create and manage SMB shares



Automatically refreshing VMs with production data



Create VM-consistent snapshots in seconds

You've now been initiated. You know a little more about VAS, but if you want to get more serious, what should you do?

Read more—add to your knowledge

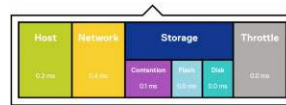
We recommend the following short and sweet documents:

INFOGRAPHIC



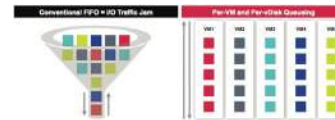
VM-aware storage
101 infographic

WHITE PAPER



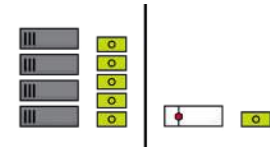
3 things you can
only do when your
all-flash is VAS

DATA SHEET



Tintri VMstore
Overview

WHITE PAPER



The Six Factors that
Determine Your
Storage TCO

Schedule a VAS demo

Want to know exactly how VAS can save you time and money? Then schedule a 30-minute demo to get your questions answered and see VAS in action. It's easy to set-up time, you can:



Submit your request at
www.tintri.com/contact-us



Email us at
info@tintri.com



Reach out to
[@Tintri](https://twitter.com/Tintri) on Twitter

Thanks for reading!



We hope The Essential Guide to VM-aware Storage got you thinking. Now it's time to make those wheels turn even faster—get hands-on with storage built specifically for virtualization and cloud. We've created a mock-up of the Tintri UI, so you can see how easy it is to guarantee performance, scale-out, replicate and more. Just visit:

Try.tintri.com

Tintri VM-aware storage is the simplest for virtualized applications and cloud. Organizations including United Healthcare, NASA, Toyota and 7 of the Fortune 15 have said “No to LUNs”. With Tintri they manage only virtual machines, in a fraction of the footprint and at far lower cost than conventional storage.



For more information, visit www.tintri.com and follow us on Twitter @Tintri