

The Essential Guide to Storage for Virtualization



CONTENTS

- 3 Introduction
- 4 Constraints of Legacy Storage
- 7 Storage for Virtualization
- 9 Action
- 10 Automation
- 11 Analytics
- 13 Four Use Cases
- 14 Business Benefits: Time
- 15 Business Benefits: Money
- 16 Next Steps

Introduction

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. 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.

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.

The cons of constraints

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.



Conventional constraints

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 the 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.

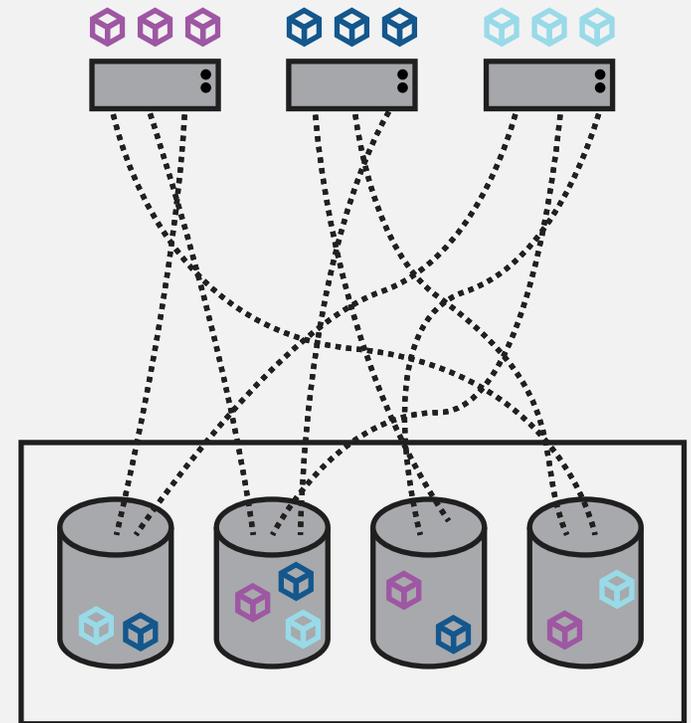


FIG. 1

How IT professionals spend their time

21%

- 23% Monitor, troubleshoot & remediate
- 22% Provision, patch, config
- 21% Innovation & new projects
- 18% New service request & approval
- 16% Vendor & internal meetings

Source: IDC Survey

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. What if you could flip the ratio and spend 80% of your time on innovation? 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.

Storage for virtualization



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

Well, it's hard.

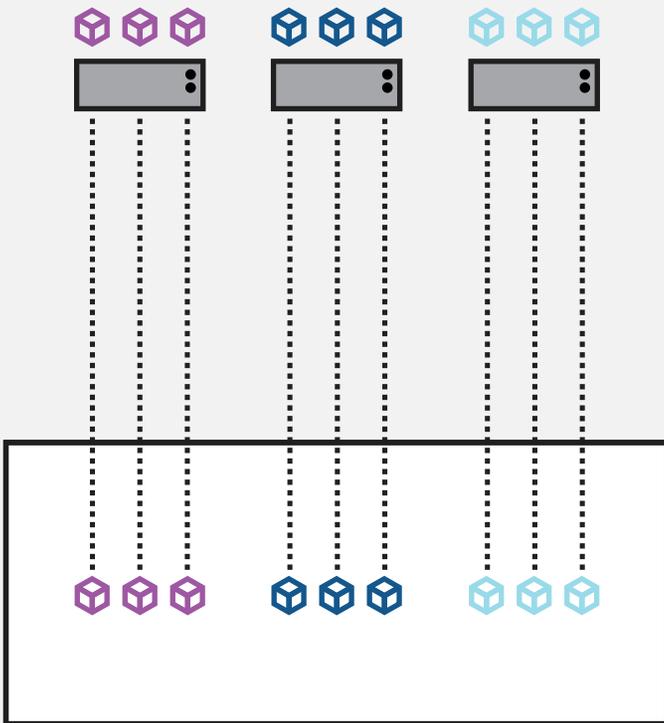
To eliminate LUNs, volumes and other remnants of physical storage, vendors have to rewrite their entire file system. For legacy vendors, we're talking a massive undertaking, and for emerging vendors moving fast to capture share, it's an unacceptable delay.

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.

Forget constraints. Connect your file system and storage, using not LUNs and volumes, but rather individual VMs as the level of abstraction. Now, admins can manage, replicate, clone, snapshot and analyze on an individual VM level—no more tuning LUNs or keeping a sprawling spreadsheet of VMs. Storage that's purpose-built for virtualization saves your business time and money.

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 storage for virtualization.



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 VM-level operation, you just click the exact VM to replicate and apply policies that affect it alone. That precision can reduce WAN usage by up to 95%.

"I've worked in the storage industry for 25 years, and I've never put in a storage platform in 35 minutes, it just doesn't happen.' So we used Tintri for a while, seeing how it performed, expecting that it just can't be this simple. But it was."

 Southern Health
and Social Care Trust

FIG. 2

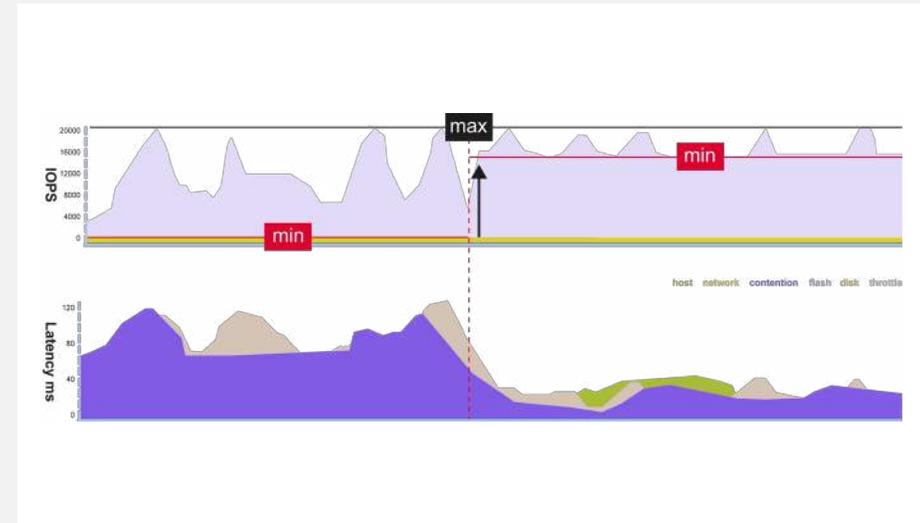
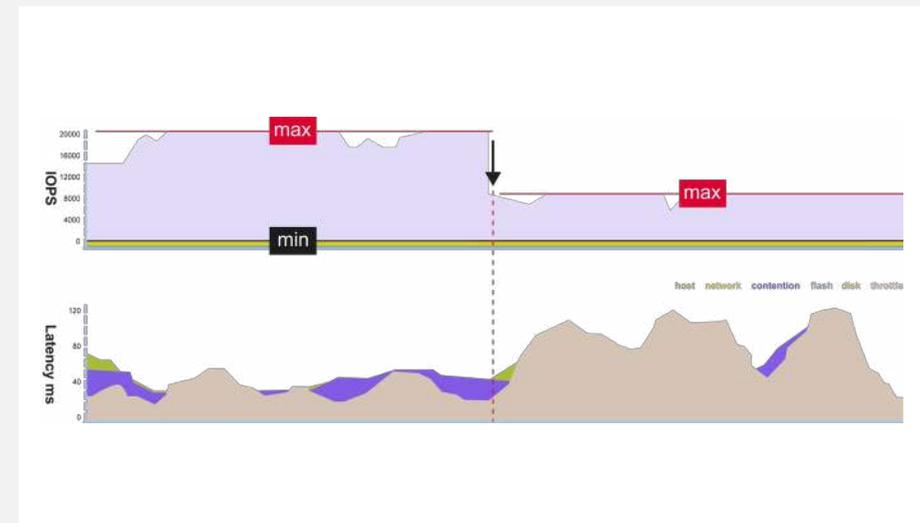
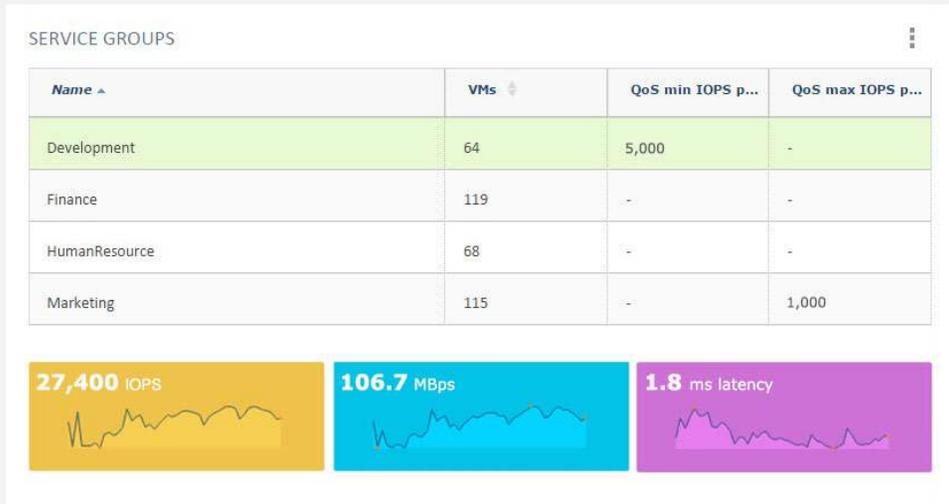


FIG. 3



Automation

FIG. 4



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 wastes your time.

Storage that's built for virtualization 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 VM-level operation, 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 (at left) 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 storage automation. SEE FIG. 4

"Tintri 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 can accommodate twice the IOPS at less than a third of the latency in 1/4 the footprint."



Analytics

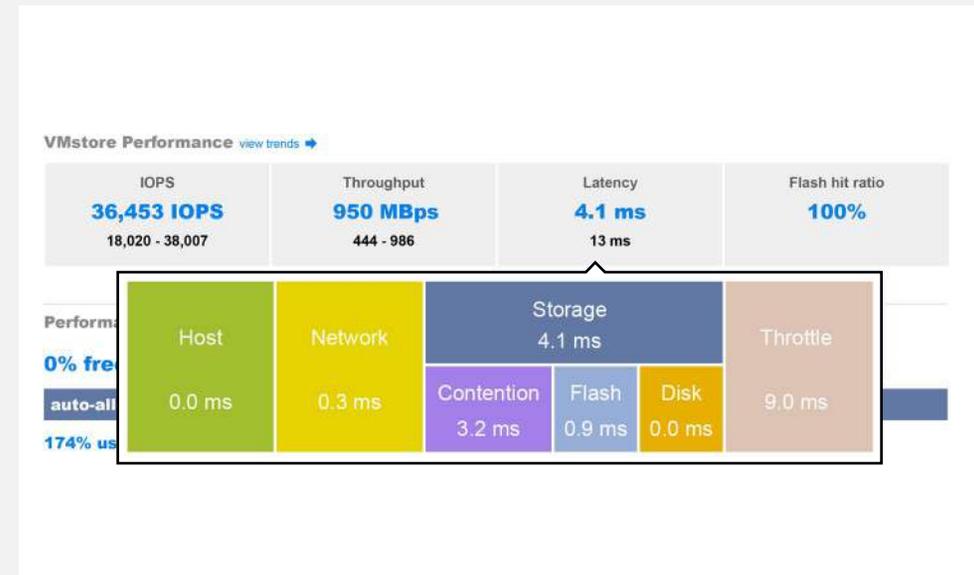
With conventional storage, you can't even see through a LUN to its resident VMs. With storage built for virtualization, you have transparency, and thus unprecedented granularity in your storage analytics. For every single VM, you can see:

- 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 at right: 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.

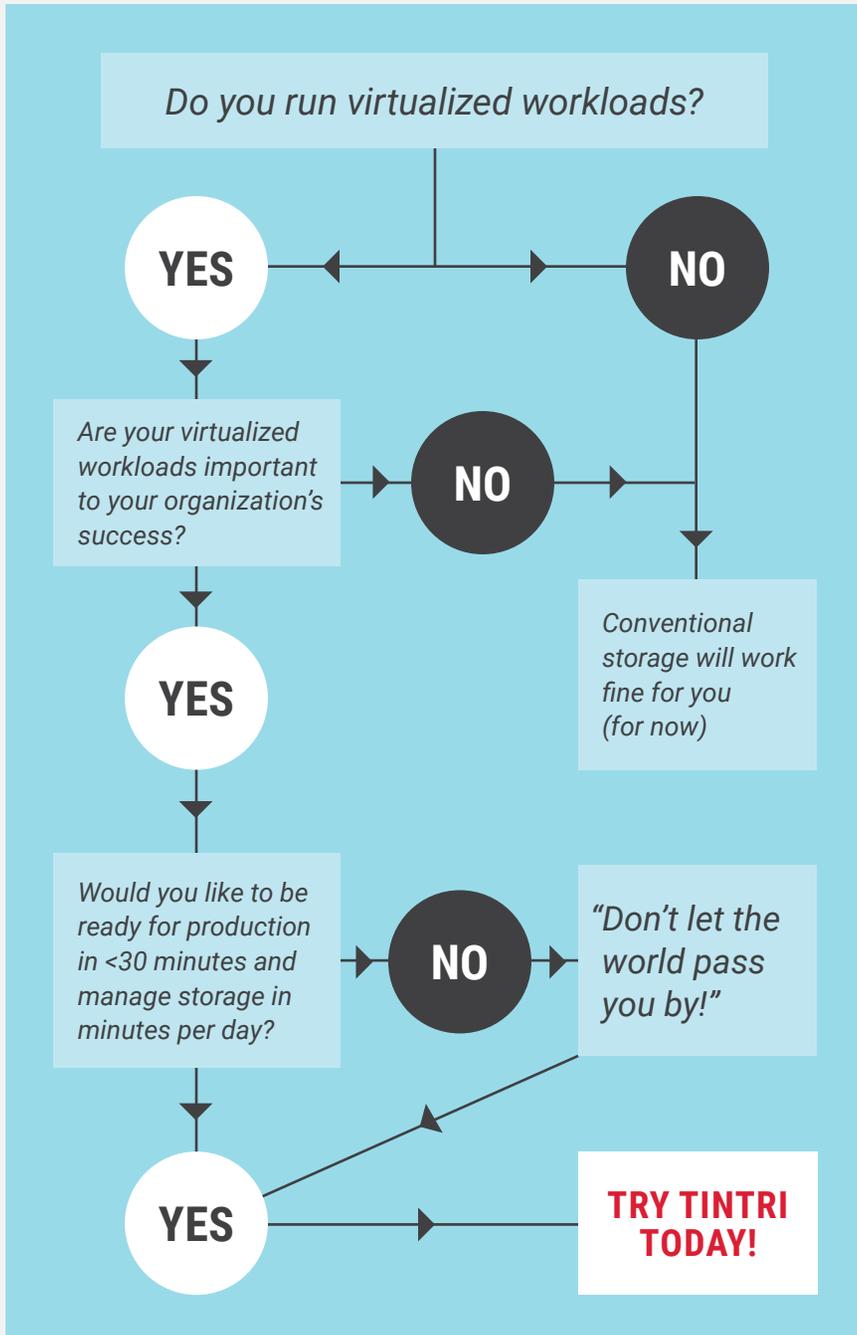
Meanwhile, though most conventional storage can estimate when you'll run out of capacity, Tintri all-flash storage knows exactly when you'll run out of capacity, **throughput** and performance. And it can guide you to the right mix of hybrid and all-flash storage, ensuring that you don't over-provision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

FIG. 5



"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."

Do you need storage that's built for virtualization?



Now, that applies to growth of your current workloads, but also to modeling new workloads. Tintri all-flash storage and software lets you use the historical behavior of applications (up to three years) to define storage needs 18 months into the future. That's a far longer storage horizon that lets you plan better and procure better.

DATA POINT

Four years ago, fewer than 50% of organizations were familiar with storage built for virtualization.

Currently, 85% of companies stated they were familiar with VM-level storage management, and 79% cited these capabilities as central to their storage purchases over the next 12 months.

Source: IDC Storage for Virtualization

"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 systems with very little effort."

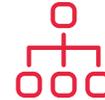


Four use cases



Servers and Databases

Give each server and database 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 VM's performance to see the source of any latency, set QoS and replicate it in seconds.



DevOps

Deliver 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 DevOps team can own their storage footprint.



VDI

Storage that is tightly integrated with your ecosystem—including VMware Horizon (with View) and Citrix XenDesktop—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 the right storage offers unprecedented VM density—thousands of VMs on one platform—so you can use less space, sip less power and spend less money.



Private Cloud

Deploy the storage backbone for your private cloud in less than 30 minutes. Isolate every VM in its own lane to eliminate noisy neighbors. You can group 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.

The business benefits

You've learned about VM-level operations—now let's explore how they will benefit your business.



Business benefits

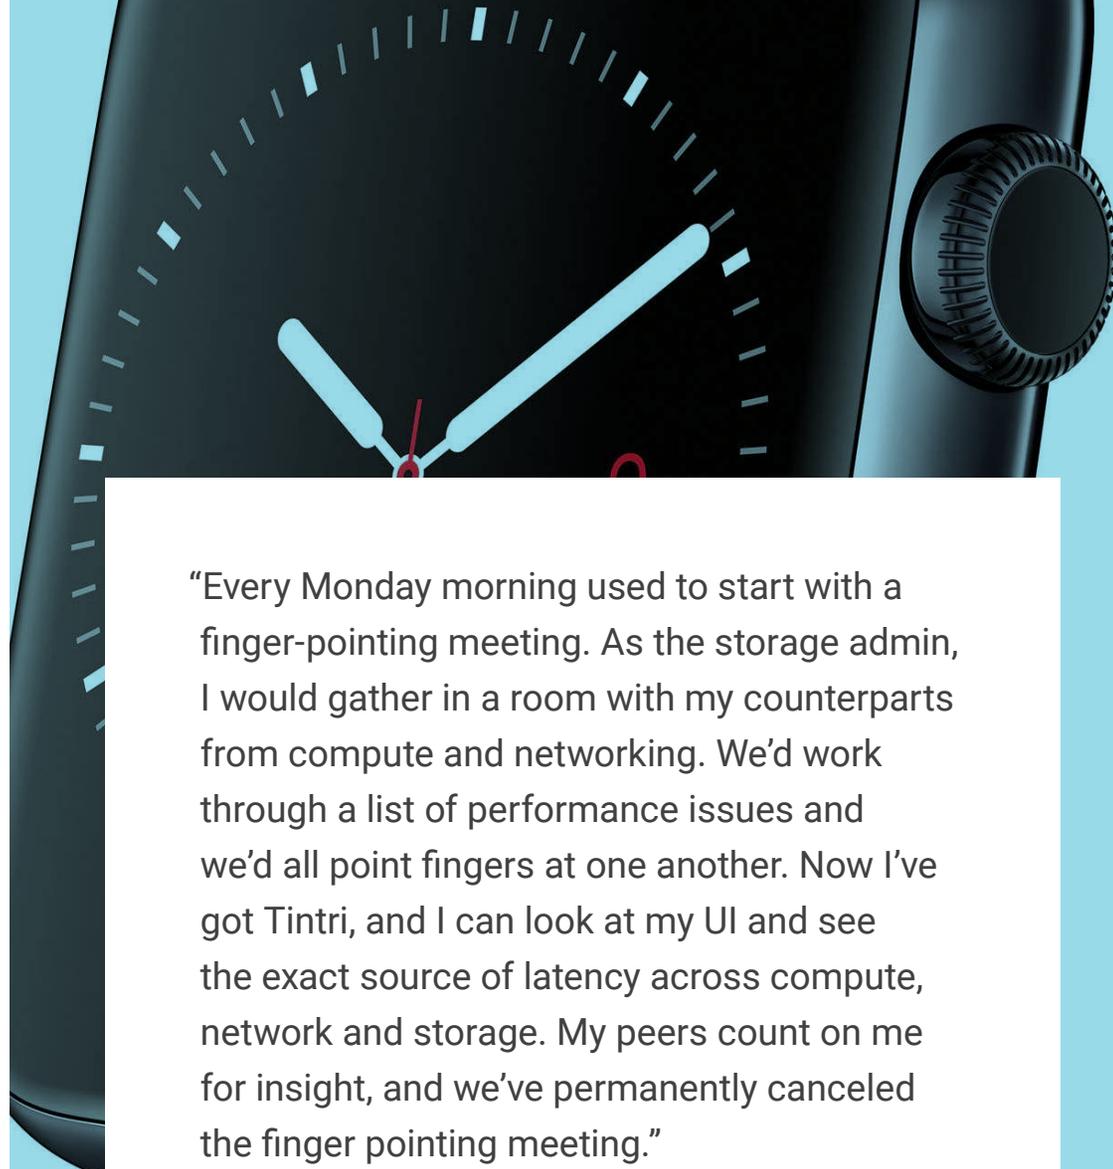
Time

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 Tintri 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 private 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 that an admin be versed in LUNs, volumes, queue depths and more. Tintri 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 storage 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. Tintri 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

Business benefits

Money

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

Capital expense. Since Tintri 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 Tintri 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 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 Tintri all-flash storage as needed—standalone or converged with best-in-class server and networking.

“We achieved a significant reduction in OPEX by moving to Tintri. 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

Next steps

Try it yourself—and test your skills

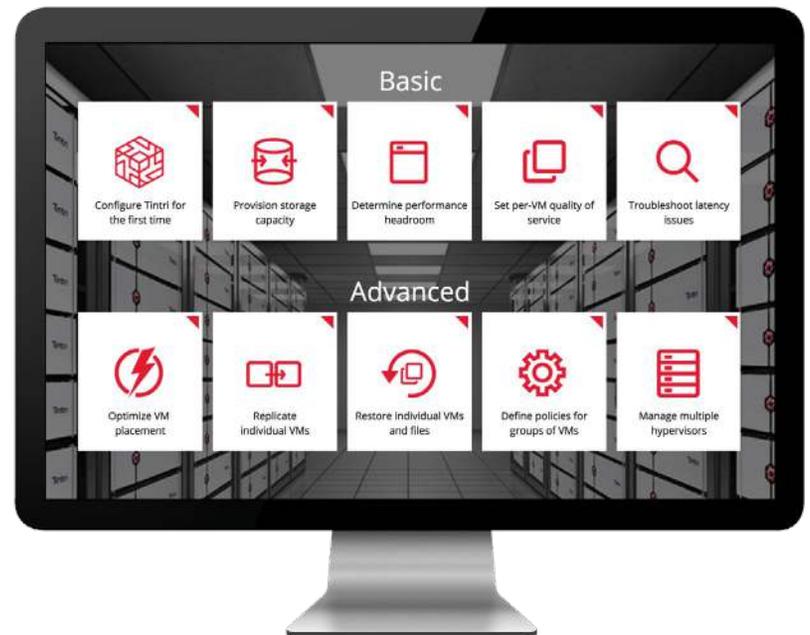
Now it's time to get hands-on. We've created a mock-up of the Tintri UI, so you can see how easy it is to guarantee performance, scale-out, replicate and more. Just visit:

explore.tintri.com

Contact us for a demo

Want to know exactly how Tintri can save you time and money? Then schedule a 30-minute demo to get your questions answered and see us in action. Just complete the form at:

tintri.com/contact-us





Tintri maximizes performance for applications and the people who manage them.

With all-flash storage and software for virtualized workloads, Tintri automatically manages each application, so you don't have to. That means you're free from decades-old storage constraints, so you can spend your time on high-impact projects.

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

