PeerSpot

Case Study

Tintri VMstore

Jeff Wilhelm

Founder & CEO at Infused Innovations, Inc.

Review by a Real User

Verified by PeerSpot

What is our primary use case?

We were using Tintri as a delivery mechanism, in a managed service type of environment.

How has it helped my organization?

At the time, a decade ago, storage vendors were talking about LUNs, fiber channel, worldwide names, and volumes—and all of the underlying framework of storage. Tintri was coming to the market with a different approach. They said, "You don't have to worry about any of that stuff. We're going to make it easy. If you're using VMware, we tak a novel approach and present disks that you can assign to VMs, and they're managed at the VM level." At the time, that was a unique value proposition and turned out to be true. It vastly simplified operations and the associated expense of delivering storage to customers.

Some customers would use it in a primary storage capacity, running virtual machines on it. Others would use it as a replication service. If they had Tintri in their own co-lo facilities or data centers, they would use it as a replication target for block-level replication. The ability to offer hundreds of terabytes of those services and not even require one full-time storage administrator on our side to do so, was frankly a win from an operating expenses perspective, but also a really novel offering when most of the other storage out there really took a lot to manage.

The OpEx improvements are from reduced administrative time, not needing multiple storage administrators dealing with storage day in and day out—the managing of volumes and LUNs, et cetera. That one area has completely

disappeared. We don't have to worry about that anymore. That allows our people to focus on more important, value-added functions. It's a bit cliche to say "set it and forget it." But if Tintri is deployed correctly, and the Tintri Global Center stack is configured to allow for machine-learning-based placement of virtual machines, it really is set it and forget it. That is something that has absolutely cut down on administrative time and allowed us to focus on customer value delivery.

Another major benefit accrues to our customers. Customers are looking for two things: for their storage to perform without people having to say constantly, "Let me fix that," or "Let me look into that"; and for protection models that match their organizational commitments and requirements. Some of those commitments are cyber-related and some are contractual. Our customers that have Tintri on-premises want to know that their environments are protected in a fully managed way with immutable recovery and low RTO and RPO. We're able to provide that. The customers that are hybrid or using us in a fully managed service want to know the same thing, but they want to know that we have our arms around that so that they can meet the commitments they're making to their customers, insurers, and their business owners. Tintri has allowed us to do both of those things.

In terms of business impact, the performance of an array is directly tied to

what workloads can be delivered how many can be delivered in a given storage ecosystem what the performance looks like to the customer. I'm certainly not saying Tintri is the only game in town, but when you look at the overall value proposition, from the perspective of a protection model, storage setup, operating expense, administrative overhead, ability to have deep granular visibility, and everything else, it's really hard to compare the Tintri solution unfavorably to anything else.

Tintri has been a real differentiator for our business that allows us to deliver what we say we will. It allows us not to spend time working on all of the nuts and bolts and underlying things that a lot of other storage vendors require and allows us to focus on the business. We can deliver good business objectives at both the virtual machine and application delivery levels, although the latter is what we really care about. It's not about the virtual machine, it's about applications that are delivering in a way that performs for the business. That has never been in question with Tintri. It allows us to abstract the storage choice away from the delivery of application reliability and performance.

The solution's autonomous operations have impacted IT department productivity positively. Knowing that Tintri Global Center allows for the intelligent placement of workloads, rather than requiring us to keep our hands on the levers, has been helpful. VMware has done this with DRS. The distributed resource scheduling allows for virtual machines to be placed on hosts, based on RAM and CPU and all the other requirements that a business

might have. Tintri allows for that to happen at the storage level, and it allows for business objectives to be mapped into that process. For example, we can create application groups in Tintri Global Center and specify what we care about for each group. It uses that, as well as the machine learning based on previous choices that we've made, to do things autonomously without somebody having to log in once a week or once a month. It can remember the seasonality of certain workloads to make these decisions in a way that we find to be accurate for the business requirements that we put in place.

We have customers with Tintri in their offices but they don't have the budget to be able to spin up an entire second instance of their office environments in a co-lo and buy multiple storage arrays to do it. They're small to medium businesses, but they still want to be able to have access to features such as immutable backup and recovery capabilities, replicas that exist in a disparate physical location, the ability to spin up machines if there is a problem and have them online with recovery points that are sub-one-hour. Those features are somewhat unheard of in the SMB market. Yes, enterprises have had those features for a long time, and they pay for them and for the staff to manage them. But the ability to say to small and mid-size businesses that they can have all those enterprise features as a managed service is unique.

Another huge aspect is that VMstore enables replication, snapshots, and setting QoS at the virtual machine level. That's especially true when organizations are trying to decide where to invest their dollars at the same time as cybersecurity insurance and Zero Trust security issues are happening around them. The ability to help them feel that they're not choosing between security and productivity, but rather, that the two can go hand in hand on Tintri, helps a lot. Our customers want to know that they have immutable off-site recovery points without having to understand the technical requirements for doing so. We're able to help them be more secure and protected while still delivering good quality of service and delivery of their applications. We're able to add more value for our customers by answering those questions in that way.

What is most valuable?

Among the most valuable features are its granular replication ability to define asynchronous or synchronous replication, which gives us very definable RTOs and RPOs around that type of service granular quality-of-service configuration, which allows for cases where you've got multiple customers on a single Tintri, but you want to be able to offer strong quality-of-service metrics and KPIs. Those are really useful features.

Obviously, having reliable, performant, highly available storage are the table stakes, but they are important as well.

We've also always liked the visibility that

VMstore gives you at a virtual machine level. That's been a Tintri feature for a decade. It helps us help our customers understand workloads. One of our customers was on our managed platform that was built on Tintri and VMware for five years. Over those five years, they would run high-end analytical workloads and large ETL projects with multi-billion-record datasets. We were able to tailor the storage to their workload requirements. It also enabled us to answer questions about why something might perform in a certain way, or how they could improve on their side from an application delivery perspective, with actual data at the application and virtual machine and storage levels. That is unique, especially since these customers aren't $10 billion enterprise accounts; they are small and mid-market customers. They need those answers in real time to run their businesses better without paying a fortune to do so. That visibility is something we've really benefited from.

Also, you can launch the interface to look at either application groups, virtual machines groups, or at the individual level. The interface is really nice. Tintri Global Center, running in a virtual machine in a customer's or a managed environment, does the same thing with even more granularity. It has the ability to do predictive workload moves based on those parameters and characteristics that are being reported on.

The interface gives you the ability to look at IOPs: how much is hitting flash versus spinning disk, if it's a hybrid array; or if it's an all-flash array, what the characteristics of the workload look like. You can also look at it from an application perspective and see that SQL is reporting this while the VM is reporting that.

Those are really important features. In the past, we've deployed arrays across the Tegile family before it was bought by DDN, which owns Tintri, as well as Dell, HPE, and ZioTech, going way back. We're very familiar with what existed in the space. While the ability to offer granular information at the virtual machine level is becoming more commonplace now, Tintri does a really good job of doing it in a way that doesn't require you to pull the data out and manipulate it in a reporting package, and it's very visually robust.

And the T7000 Series that we're using has 40-gig network interfaces on them uplinked to 40-gig interfaces on our core. For our VMware environment, which is delivered on a Dell chassis, it is very performant. Some of our customer workloads are extremely demanding, running multi-billion-record merge-purge, ETL, and index search applications, and there have been no performance issues. In fact, when customers moved from a Dell array to the Tintri T7000s, we saw nearly a doubling of performance in some of their most demanding workloads, which was something the customers were very happy with as well.

What needs improvement?

The biggest area for improvement, and there has been some roadmap work in this area already, is cloud integration. It would be hard to find an organization that isn't actively embracing the cloud, at least in a hybrid way. Some workloads are very well suited for public compute on Azure, AWS, GCP, et cetera. That's important for us and for our customers, but it's also important from an OpEx perspective.

Tintri has been investing in this area and I'm sure will continue to, but cloud integration has been the biggest area that we've been crossing our fingers and hoping for quick development around.

We're a Microsoft Gold partner and a lot of our customers have heavily embraced the Microsoft ecosystem. To do modern work securely, public compute and productivity on the Microsoft Office 365 stack, Exchange Online, SharePoint, Teams, OneDrive, et cetera, are important for mid-market and smaller organizations. There has to be a deep integration point between the on-premises solution, a hybrid cloud solution, and a public cloud consumption model. That is becoming an increasingly important decision point in all of our vendor-alignment discussions, Tintri included.

For how long have I used the solution?

We have been working with Tintri since the end of 2013 or in 2014, when it was a standalone company. We had Tintri T540s at the time, and we were also part of the alpha and beta programs for some of the Tintri OS releases, Tintri Global Center, and Tintri Analytics. We are also part of some of the steering groups around product features for the T650 and T800 Series.

What do I think about the stability of the solution?

We've never had an issue with crashing, lagging, or downtime.

What do I think about the scalability of the solution?

The scalability is good.

One thing that is a little unique about Tintri is that, while you can present a large pool of storage, there are some idiosyncrasies around how that is seen by VMware. However, Tintri Global Center makes it appear as a single storage volume because it moves workloads predictively.

How are customer service and support?

I have probably opened a couple of dozen tickets over the decade that I've worked with Tintri. Like most support organizations, they acknowledge a ticket very quickly. If it's a high-

severity alert, remote hands can be dispatched very quickly. With us, Tintri has never breached its SLAs for remote hands or support.

Some of the things we work with them on our tickets related to R&D or beta programs that we're doing with them. Those involve different teams, but they are still really great engineering talent, and we've never had a bad experience working with them in the last five years.

How would you rate customer service and support?

Positive

Which solution did I use previously and why did I switch?

In the last decade we have used a vast selection of competitors to Tintri. Compellant, ZioTech, Dell, HPE, Pure Storage, and Tegile before and after its acquisition. Currently, we only recommend Tintri to our customers.

Of course, we do have customers that have multi-vendor solutions, and we're able to support them, but our current recommendation for storage is Tintri. It's been like that for a few years, and I expect it will continue like that into the future.

How was the initial setup?

The deployment model depends on the customer. We have Tintri deployed in customer data centers and, for really small customers, in their offices. We have it deployed in our own data center in a hybrid, co-location model, as well as a replication model for customers. We have about 500 terabytes of Tintri deployed at a Tier 4 data center in Massachusetts. And the cloud replication model is something that we're using as well.

Our most recent customer deployment was for a small to medium business that does about $100 million in revenue in the New England market. We deployed three new server hosts for them, some switching, and Tintri on the storage layer. It was a T Series deployment and from unboxing to rack-and-stack to IP addresses set and the storage being mapped into VMware, it took less than an hour.

For a small to medium business, the implementation strategy is to get the storage deployed in the customer's environment and map it into the VMware environment. If it's an implementation that requires a migration from another storage vendor, then there are storage vMotions that take place. Once the environment is up and running, we usually provide performance metrics to the customer. And we set up the RTO and RPO requirements in our data center. If the customer is taking advantage of the Tintri block-level or VM-level replication, we would match that to their business objectives related to application location delivery—tier one, tier two, or tier three— and what dependencies exist in their environment.

But the implementation and deployment model is really straightforward because as soon as the IP addresses are assigned to the storage and any VLAN requirements are completed, it's just a matter of moving the virtual machines to the storage and then focusing on the business rather than the storage itself. We usually require one engineer for deployment.

In terms of maintenance, these are disks or SSDs or storage modules and there is predictive, and at times reactive, maintenance. With enterprise storage, there's a four-hour response and 24/7 hands-on if that's what is needed. That means there is nothing required from our team, but sometimes we want our team to be the face of the account to the customer. In those cases, one of our engineers may go onsite to meet the Tintri support person. A lot of our customers have no visibility into any maintenance because we're offering this as a service within our co-location environment. So generally speaking, there isn't really any management that has to happen. And if there is some alert that requires field service, we're contacted, sometimes proactively and sometimes reactively, by Tintri support.

What was our ROI?

The return on investment for our organization has been that we're able to deliver what we say we will to a customer. We're able to offset the projected OpEx expenses of other vendors. The ROI for us is based on customer value and customer satisfaction. Storage has never been a detractor on our side in terms of customer value. Storage has never been the root cause or smoking gun for any problems that a customer may have asked us to look into.

The ROI for us is clear. We're able to deliver performance at a predictable price point with a predictable administrative and operating model. We really haven't looked into any other vendor in the last couple of years because of that.

We're able to offer competitive managed services and innovation services to our customers. Some of that is built on Tintri, and some of it is knowing the predictability of Tintri and the scaling model. It allows us to then invest in other areas of our business as well. If we had a less predictable scaling or OpEx model, it would introduce a lot of questions for us about whether we should be investing in these areas or putting these dollars towards storage vendor X or management person Y. The predictability, reliability, and performance of Tintri have enabled us to invest in other areas of our business without concern.

What's my experience with pricing, setup cost, and licensing?

We haven't found Tintri's pricing to be a barrier to entry for any of our customers. Tintri works through a number of distributors and when we've had to purchase it for our customers we

have found the pricing to be competitive.

We've worked with customers that have to get a three-quote comparison. A recent example was a customer that was looking at IBM storage, Tintri, and Dell. Tintri was competitive on price-for-features, as well as licensing and the support model. We didn't have a reason to push the customer in any other direction than Tintri. We felt that the feature set, performance, pricing, and total cost of ownership, inclusive of the operating expenses and administration, made Tintri a no-brainer.

What other advice do I have?

Tintri has gone from older models that were hybrid flash with a mix of spinning disks and SSDs through improvements to all-SSD arrays. Now there are NVMe benefits, which have lower latency and higher performance. As the product line has matured and improved and new generations of storage have come out, performance has followed suit. It's not just the hardware itself that adds performance benefits, but for organizations that have multiple arrays, Tintri Global Center allows for the intelligent positioning of virtual machines to maximize the performance. The combination of a heuristic, machine-learning-based placement model and improved hardware performance is part of the secret sauce of Tintri.

As for corporate stability, we were working with Tintri before the bankruptcy and before the DDN acquisition, and we continued to work with Tintri after those events. Stability is important, but OpEx performance and the ability to deliver are a part of the conversation. I found that Tintri was delivering innovation a decade ago when I first started evaluating them. I will admit that we did look at other vendors in the 2017-2018 time frame when Tintri was going through some of its challenges. But we are comfortable now, especially with DDN backing Tintri. There is good financial stability, good operating stability, a good roadmap into the future, and it's still a great product.

My advice is to talk to customers that are already using Tintri. Every vendor you're considering is going to tell a good vendor story, and every vendor is going put things in front of you that say why they're better than any other vendor. But talk to somebody who's actually using it and let them tell you why it was the choice that they made.

Read 7 reviews of Tintri VMstore

See All Reviews