

Guide to Monster VM Management

Welcome, storage gamekeeper, to the Guide to Monster VM Management—Tintri’s primer on identifying and managing the Monster VMs hidden in your virtualized environment.



Let’s start with the fundamentals—what’s a Monster VM? Well, it’s any virtual machine residing in your data center. You might call them virtualized databases, virtual desktops or something else entirely. We call them monsters.

Now, is that really fair? Certainly. It doesn’t mean they are evil or malicious—far from it. They have all the right intentions, and when in their preferred habitats they are reasonable and harmless creatures. But when provoked or poorly placed, they all have the potential to be badly behaved.

The key to being a successful storage gamekeeper is (1) knowing exactly what Monsters you are dealing with, and more importantly, (2) managing to their very individual needs.

For instance, where should you house your Monster VMs—in all-flash or hybrid-flash environments? How can you help them grow with vSphere, Hyper-V, RHEV or OpenStack? Now’s the time to understand the care and feeding necessary and we’re here to help. Here’s how Tintri VM-aware storage can help you manage your monsters:

Gulper

This is a monster Oracle database, SQL server database or another very large and transactional application. These are all I/O- or throughput-intensive VMs that need performance isolation—to get their own “lane”—and Tintri can do just that to ensure they (a) get the performance resources they need, and (b) don’t disrupt other VMs.

Consequently, Tintri’s Hybrid-Flash platform is very capable of handling these VMs; but if yours come in larger sizes and/or quantities, Tintri also offers you the choice of All-Flash. Both platforms share a common operating system and real-time analytics so you can optimize the location of these particular monsters according to their need.



Conventional storage

The Gulper has a voracious appetite for I/O and needs a lot of room to roam. But conventional storage stuffs Gulpers into the same LUN habitat of other VMs. The Gulper steps on the VMs in its path as it fights for more than its share of resources, flattening their performance.



VM-aware storage

Give that Gulper all the I/O it can eat with guaranteed Quality of Service. And isolate it in its own yard, so if it makes a mess, no one else is going to step in it. No LUNs to confine it. No conflict over resources to anger it. You’ve turned your Gulper into a puppy.

Werewolf

Here we have an ERP system, a finance server or another application that tends to falter at predictable times. With Tintri VM-aware storage, though, you have VM-level Quality of Service, so you can set a minimum GoS threshold for each VM to guarantee I/O. Hybrid-Flash works great with one or two applications, but if you have many, All-Flash is recommended.



Conventional storage

Though usually mild-mannered, the Werewolf will “wolf out” on a predictable schedule. So if your storage isn’t prepared, it’s hard to shuffle him around or tune him out—and you don’t want to move a mad Wolf. The only thing to do on conventional storage? Wait out the storm. Not great.



VM-aware storage

Your Werewolf will keep working through the night, even during a full moon-like event. Forget the hairy behavior—you can dynamically assign your Werewolf the performance it needs to keep things operating smoothly. No need to destroy the furniture.

Swarm

Sound familiar? This is a VDI implementation. To prevent scans and storms from slowing down desktops, Tintri Hybrid-Flash serves 99%+ of I/O from flash. And if you have thousands of persistent VDI desktops, Tintri offers All-Flash storage.

Should a single user ever complain, you can troubleshoot in seconds, because you've got cross-infrastructure visibility. See truly real-time analytics (not correlations!) across host, network and storage in a single click to pinpoint the problem and prevent trouble tickets.



Conventional storage

When your Swarm's beset by a noisy neighbor or a storm, you can quickly lose control. And since each individual Swarm VM is stuffed in the same LUN, it's hard to tell from the outside which ones are still working, and which ones are just plain angry. So your first indication will have to be the trouble tickets that fill up your inbox.



VM-aware storage

When Tintri storage lets your Swarm move freely, your Swarm are consistent worker bees—regardless of how persistent they are. And because storage gamekeepers can keep track of each and every Swarm VM, if one starts to struggle, they can instantly see why. No wonder gamekeepers are all buzzing about VM-aware storage.

Beast

This might remind you of your dev and test VMs or private cloud implementation, which can grow, shrink and fuss without warning. Fortunately, Tintri lets you provision individual VMs (not LUNs), and clone and replicate with just a few mouse clicks (instead of hours).

Should you ever get into a jam, restore individual VMs instantly from any recovery point without ever losing performance history. Hybrid-Flash is best for these inconsistent VMs, with Tintri providing automatic and dynamic tuning to meet whatever demands you experience today.



Conventional storage

You never know what your Beast wants. One day it'll want the highest QoS in all the land, and the next it'll suddenly need space. You can try to troubleshoot its needs, but all the PerfStats and diagnostics you'll need to run will only fuel your frustration. You want your Beast to be your best friend, but without VM-level insight, you just can't understand one another.



VM-aware storage

Tintri is the beauty to your Beast. It's VM-aware storage to the rescue, giving you immediate insight into your Beast's mood. You can monitor it without being obstructed by LUNs, and if its state of mind starts to swing, just swipe QoS ceilings up. Should your Beast begin to bark and bite, you've got per-VM backup and restore, so you can have a fresh start in seconds.

Blob

This mass of mostly cold data is easily solved with Tintri Hybrid-Flash, which divides workloads into active sets in flash and non-active sets in high-capacity disk. This is far more space and cost efficient for your mission critical VMs that have a sizeable amount of infrequently accessed data, like your Exchange server. Tintri's Hybrid-Flash arrays offer up to 100TB of effective capacity and can handle 3,500 VMs in just 4U. Keep colder data on Hybrid-Flash and save cash.



Conventional storage

The Blob demands space, not performance. But because the conventional storage gamekeeper can't see at the VM level, the Blob is often mistaken for its fellow big Monsters, such as the Gulper—meaning it gets housed in an expensive All-Flash Array. So even though the Blob consumes very little storage performance, it's consuming a lot of gamekeeper money ... when it'd be just as happy elsewhere.



VM-aware storage

Your Blob really isn't a monster on VM-aware storage, but he can still rob resources if you set him up in expensive all-flash accommodations. Tintri can help you with that—put the Blob in Hybrid-Flash, and it's just as happy. You'll never mistake it for other monsters, and should your Blob(s) ever require flash, Tintri Hybrid-Flash is designed to deliver 99% of I/O from Flash.

Managing your monsters

Managing monsters in the VM-aware era requires the right equipment. You wouldn't keep a stegosaurus in your shed—housing a Gulper without VM-level Quality of Service would be equally foolish. Here's our advice:

Create the right habitat

As you've learned, some Monster VMs are more comfortable in an All-Flash environment—your demanding database Gulpers and Swarms of persistent VDI desktops. But most Monster VMs live in harmony on Hybrid-Flash, with the VM-level visibility to identify and address their individual needs, and the design to serve them with flash as needed. Fortunately, Tintri offers both, with a common operating system and real time analytics to help you optimize their location across All-Flash and Hybrid-Flash.

Simplify care and feeding

You've got hundreds, perhaps even thousands of Monster VMs to manage. That's a lot of mouths to feed. So, you need the help of automation. Well, Tintri lets you manage 100,000 Monsters from a single control console. You can easily group like Monsters and set service policies that apply to all members—from data protection to QoS. As you move your monsters around, those policies stick with them, so they are always well cared for. Tintri simplifies monster management.

Keep watch

Storage won't drive business value, but your Monster VMs sure will... so you've got to keep watch. VM-aware storage is the only way to manage individual Monsters. You can monitor their behavior not just on your storage, but also across host and network—so if there's a problem, you know the root cause in seconds. You can do it all in real-time, or look at historical data to identify trends.

Expert storage gamekeepers understand that effective management requires the right equipment and support. To guarantee that each Monster VM gets what it needs to stay well-behaved and productive, you need VM-aware storage. And to see what VM-aware storage can do, you need to [try VM-aware storage](#). We've set up a hands-on Tintri demo for you to set QoS, replicate VMs, and try other features we've mentioned here.

See for yourself. With conventional storage, your monsters are too prickly to tame. With Tintri, you can put them to work—so you can work, too.



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

ANZ HQ

MLC Centre Level 56
19 Martin Place
Sydney NSW 2000
+61 2 9238 2128
anz@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