

A Deep Dive into Tintri VM Scale-out

Contents

A Deep Dive into Tintri VM Scale-out	1
Tintri Technologies	1
VM-Level Performance Visualization	1
VM Scale-out	1
Tintri Analytics	2
Live Resource Usage and Organic Growth Trends	2
View the Live Resource Usage	2
View the Organic Growth Trends	3
Using Tintri Analytics, VM Scale-out and Other Tintri Capabilities Together	3
Example 1—Solving the Initial VM Placement Challenge	3
Example 2—A New VDI Project	4
Example 3—Deploying 100 New Database Development Environments	5
Example 4—Scaling out with Tintri Storage	6
Summary	7

A Deep Dive into Tintri VM Scale-out

LUNs, volumes, RAID, striping and more have nothing to do with the virtual machines and applications that run your business, and only result in storage performance issues and management complexity. Tintri lets you manage what matters: individual virtual machines. With the announcement of VM Scale-out and analytics, Tintri provides virtualization scale-out technology that makes it possible to focus on managing individual VMs instead of storage.

This paper illustrates how Tintri's per-VM management tools work together with predictive analytics and VM Scale-out technology to make it possible to scale out storage while utilizing the management simplicity that is designed into each Tintri VMstore.

Tintri Technologies

A deeper look into three core Tintri Technologies—Performance Visualization, VM Scale-out, and Tintri Analytics.

VM-Level Performance Visualization

Tintri has always provided you full visibility into your individual virtual machines. See the root cause of latency, changes in performance usage and much more—all actionable in real-time.

Tintri characterizes the throughput and latency performance profile of each VM running on the VMstore. This real-time data provides the basis for Tintri's automated performance isolation feature. Conventional storage cannot separate the I/O requirements of individual VMs that are forced to share LUNs, and so demanding (noisy) neighboring workloads can disrupt your mission critical applications. VMs running on Tintri are guaranteed separate I/O performance lanes which means that every VM receives low latency access to the performance reserves of the VMstore.

The performance data is a working history of all the VMs on the VMstore. This detailed real-time data is made visible in the VMstore for viewing the performance of individual VMs, the virtualization host, the network and the storage. Having this performance insight allows administrators to quickly diagnose and resolve performance issues with the entire virtualization infrastructure.

The benefits of this data multiply when accessed from Tintri Global Center (TGC). In addition to managing data protection policies for 100,000 VMs spread across 32 VMstores, TGC monitors 30 days of real-time performance data for each VM. The data is used by the VM Scale-out engine to optimize the placement of individual VMs.

Up to 3 years of virtual machine performance data is also used in Tintri Analytics to model storage and performance capacity, which allows the customer to proactively plan their performance and capacity reserves. Having accurate performance data for each and every VM is the only way for VM Scale-out and Tintri Analytics to provide accurate results.

VM Scale-out

Scale-out mines capacity and performance data to recommend the best placement for every VM.

When storage arrays run out of capacity or performance reserves, the entire virtualization infrastructure is affected. The workload on individual VMs is compromised—thin provisioned VMs experience write errors and stop working, backup jobs cannot run to completion, and end users complain about application response times. Yet maintaining the efficiency of your virtualized storage requires effort and consumes time.

What if you could scale-out your storage the same way you scale-out compute in your virtualized environment, by adding another server and letting the hypervisor manager optimize VMs across the pool? With Tintri, addressing capacity or performance limits is as simple as adding storage and letting Tintri optimize the resources.

The VM Scale-out technology in Tintri Global Center Advanced automatically optimizes the capacity and performance of your virtualized storage by optimizing the placement of individual VMs. TGC Advanced makes predictions that account for the next week of activity, considering the overall performance and capacity needs of each VM and the capabilities of each VMstore array, and selects VMs to move in order to optimize both performance and capacity. Tintri VM Scale-out technology makes it easy to grow your storage footprint without ever having to worry about the placement of individual VMs.

Tintri Analytics

Use Tintri real-time and predictive storage analytics to view organic growth trends and support unpredictable inorganic growth requests.

When storage sizing and trend analysis is based on rough estimates of the current VMs then the planning process that follows will be inaccurate. Tintri Analytics models capacity and performance with precision. Based on historical performance data from individual VMs, Tintri Analytics provides accurate analysis of organic storage capacity growth and performance usage. The trend graphs clearly show whether capacity or performance will run out first, and highlight the date in the future when this will happen. Analytics also makes it simple to locate VMs that are on the critical path.

When groups make requests to deploy hundreds of new VMs—inorganic growth requests—use Tintri Analytics to run “what if” scenarios based on existing or new application datasets. When asked whether the infrastructure can accommodate (for example) 25 more SQL servers, you can use exact behavior from current SQL servers to provide a definitive answer.

Live Resource Usage and Organic Growth Trends

Use Tintri Analytics to view live resource usage and organic growth trends.

View the Live Resource Usage

When logging into Tintri Analytics you are presented with a dashboard that provides an overview of the health of your storage resources. Graphs showing the overall resource usage means you can quickly gauge the health of the storage resources. Other graphs on the page highlight the amount and type of storage resources that are being consumed by each type of application.

The graphs and charts on the Tintri Analytics dashboard are designed for you to discover early-signs of excessive resource consumption, and to illuminate areas where capacity, performance, and working set resources are being wasted by VMs in the data center.

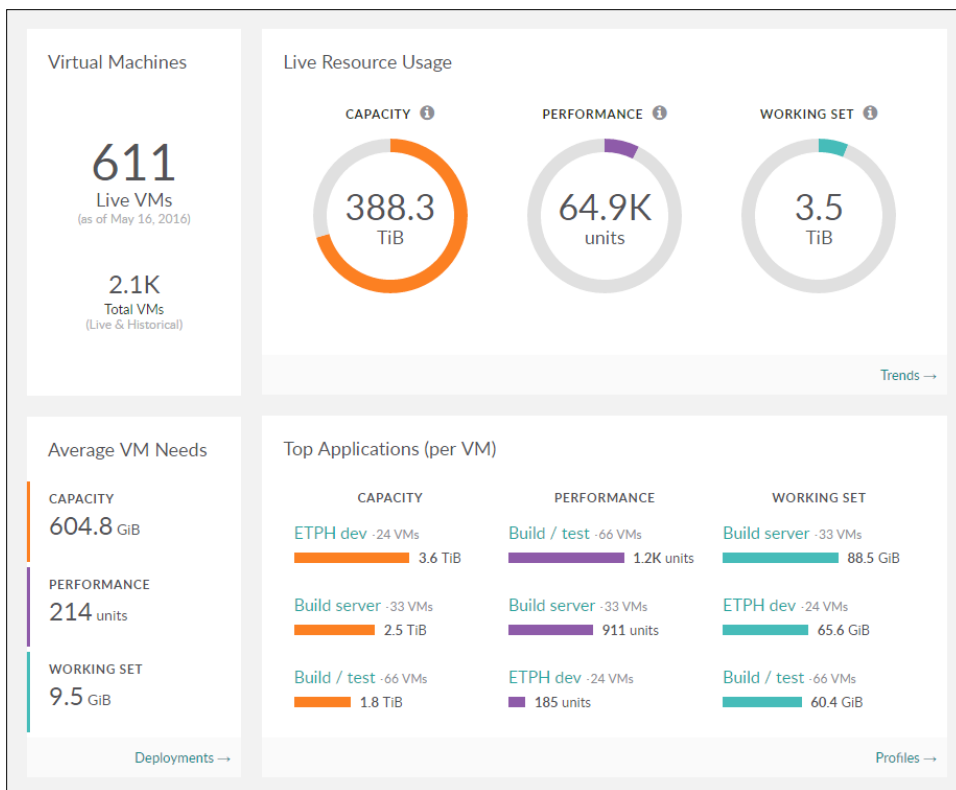


Figure 1—analytics dashboard showing the health of the environment

View the Organic Growth Trends

Another graph in Tintri Analytics shows the organic growth of the VMs in the data center. Based on detailed insights into storage, performance and working set size for individual VMs, the Analytics Insights tab displays both the actual organic growth and the projected organic growth trends. The scope of the analysis can include all VMs or be changed to focus on individual types of applications.

With all of this data available at the click of a mouse, you are now armed with storage and performance trend analysis that makes it simple to keep an eye on the organic growth of your data center.




Figure 2—insights dashboard showing the trend of capacity, performance, and working set sizes

Using Tintri Analytics, VM Scale-out and Other Tintri Capabilities Together

VM Scale-out eliminates the task of monitoring capacity and performance, while Tintri Analytics shows capacity and performance utilization trends. Other built in features such as GoS, clones, snapshots, SyncVM provide even more possibilities, but how do you take advantage of these capabilities to bring efficiency to your virtualized data center?

The following four examples illustrate how Tintri Analytics and VM Scale-out combine to bring simplicity to scale-out.

Example 1—Solving the Initial VM Placement Challenge

VM Scale-out technology ensures that new VMs are given the capacity and performance they need to guarantee their best performance.

When deploying new VMs administrators face the challenge of matching the capacity and performance requirements for these new VMs. Choosing the right storage location is a matter of considering the approximate storage capacity and performance reserves of the storage, as well as guessing the capacity and performance profile of the new VM. But why guess when Tintri knows the answer?

With Tintri there is no need to worry about the initial placement of VMs, they can be placed on any VMstore in storage pool managed by VM Scale-out. Once the VMs are created, Tintri analyzes VM-level footprint and performance statistics every ten minutes to identify the optimal location for all VMs, including the new VMs. When Tintri recommends a placement optimization, all you do is click “execute”. VM Scale-out also ensures that data protection policies, snapshots, and historical performance data follow the VM.

Example 2—A New VDI Project

Using Tintri Analytics to size a large scale deployment of virtual desktops.

Tintri Analytics provides tools that accurately size large deployments of any application, using sizing that is based on a single type of VM. This is extremely important in the case on VDI, where many virtual desktop VMs are deployed based on assumptions so any inaccuracy in the initial sizing data will be amplified in the deployment. This is where the histogram feature of the Tintri Analytics is useful.

Experiments

Once a Tintri Analytics application definition is created by selecting a group of currently deployed virtual desktops VMs (in this case, virtual desktops), it is ready to be used for this sizing exercise. Using the histogram chart (see Figure 3) to display the size of the example desktops allows you to discover specific information about the logical, physical, and provisioned capacity for the Windows VMs that will form the basis for this sizing exercise. In this case selecting a logical capacity from the upper range of values as the basis for the sizing exercise will cover a high percentage of the VDI desktops.

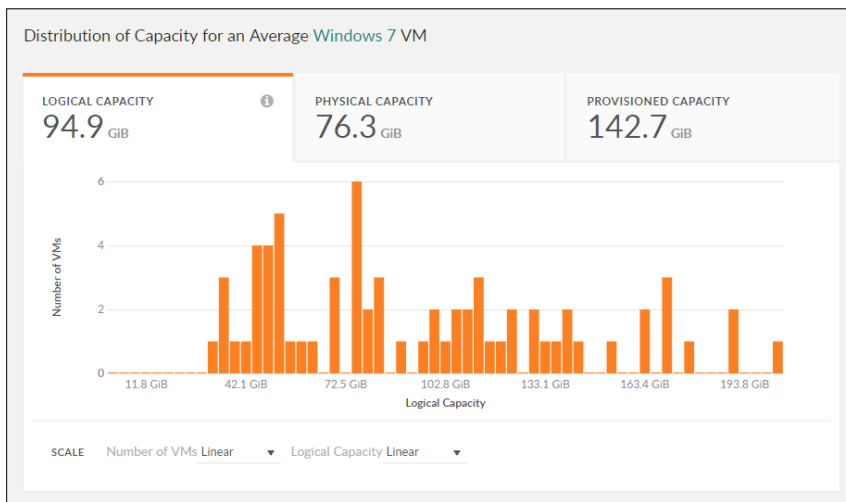


Figure 3—histogram showing the distribution of logical capacity for a Windows 7 VM

Next, use the Tintri Analytics Experiments page (see Figure 4) to test multiple “what if” scenarios to determine whether there is enough storage capacity and performance resources necessary to host the VDI deployment. Simply enter the number of instances (in this example, virtual desktops) that you intend to deploy and Tintri Analytics will show you—via the bar graphs on the right side of the screen—whether or not you have sufficient capacity, performance and working set. You can use the average requirements of existing VMs as a basis for modeling, or input other values according to the information you have available. The experiments can be run against all of the VMstores, a selected group of VMstores, or even against just one VMstore. The Experiments page will make it immediately clear if additional resources are needed.




Figure 4—sizing experiment based on a Windows VM capacity of 133GiB

Deployment

Use Tintri VM-level cloning to deploy the Virtual Desktops.

Having confirmed that the resources exist to deploy the VDI project, now the actual deployment will make use of the Tintri cloning and fast file copy technology to rapidly deploy individual desktops. Tintri supports both the VMware vStorage APIs for Array Integration (VAAI) and Hyper-V Offloaded Data Transfer (ODX). This hypervisor integration offloads cloning and file copy tasks to the storage array.

Common VDI-related tasks that are offloaded through hypervisor integration include copying of VM templates, virtual disks, and cloning of entire VMs. Each VM deployment takes seconds to complete instead of minutes. An additional benefit is that there is no additional IO or compute overhead incurred on the hypervisor host, freeing up CPU and memory resources for running the VDI deployment.

Another feature of the VMstore is the tight integration with the hypervisor management tools. When VM clones are created on the VMstore, the new VMs are also populated into the hypervisor management tools. No further work is required to view and manage the new VMs.

Tintri VM clones are space efficient. The VM clones are based on Tintri VM-level snapshots, which consume zero storage. A cloned VM appears to be a complete and independent VM from the perspective of the hypervisor. However cloned VMs do not consume any storage until the end-user edits existing files in the VM, or creates new files. The amount of storage consumed by end user desktops is tracked by Tintri Global Center and in Tintri Analytics, the tools that watch over the storage and performance consumed by the VDI deployment.

Example 3—Deploying 100 New Database Development Environments

Using Tintri Analytics to size developer workstations.

Tintri Analytics will size deployments of groups of VMs. Consider the case where a group of VMs comprise a development environment for database developers. The group of VMs includes the production database server, a database client, a web server and several other supporting VMs. Each of these VMs has different capacity requirements, OS types, and performance requirements.

This is no problem for Tintri Analytics.

Experiments

Before running experiments it is necessary to create a Tintri Analytics application definition for the Database Developer and populate it by selecting the production database server, a sample database client, the production web server, and other selected VMs. These VMs can be located anywhere in the virtualization infrastructure. Using the experiments page, increase the number of deployment instances to 100 to find there is room to deploy these development systems using the existing storage and performance capacity.

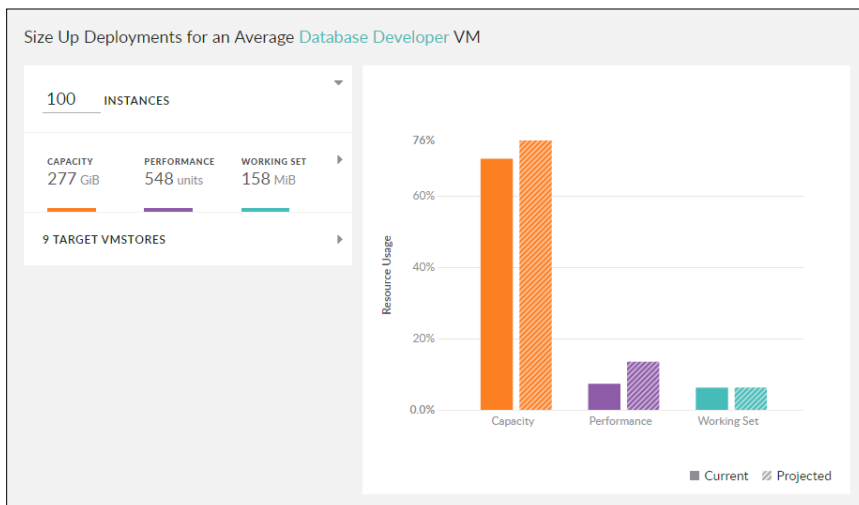


Figure 5—sizing experiment based on the VMs in the definition “Database Developer”

Deployment

Database development environments are perfectly suited to Tintri's powerful data management tools—Clone, QoS and SyncVM—that are used to create and manage databases running in VMs.

The development environment is created by cloning the production VMs, which can be done through the VMstore graphical user interface, through the hypervisor management GUI, or by using the Tintri automation toolkit for cloud-scale operations. The Tintri VMstore is engineered to store and manage VMs as objects—there are no LUNs to complicate things—which makes it possible to clone entire VMs in seconds. The Tintri cloning wizard also incorporates Sysprep tools that are used to customize VMs before being powered on for the first time.

SyncVM for Databases

Frequently, database development teams want to refresh the data in the development databases. This involves synchronizing data from the production database with the development databases. Without Tintri this is a time consuming task that involves creating backups of the production database and restoring the backups to each of the development database VMs, one by one. With Tintri the development database VMs are refreshed with production data using only the SyncVM data management tool.

Here's how it works: SyncVM clones the virtual disks attached to a production database and attaches them to the development database. Simple, isn't it? SyncVM runs quickly and requires no data movement.

Performance Isolation and QoS

Automatic performance isolation technology—part of Tintri's VM-aware storage filesystem—ensures that VMs with different workload profiles run together without individual VMs grabbing performance resources at the expense of other VMs. The VMstore applies latency to a VM that has broken out of its historical performance envelope. If a VM continues to request a new higher level of performance, one that exceeds the performance reserves, then VM Scale-out will suggest a way to optimize the VMs to allow for this new performance requirement. After a short period of time TGC will request permission to move VMs in order to optimize the performance resources on the storage.

If these VMs are development systems, consider using Tintri QoS to set performance limits and TGC won't make requests to move the VMs. For example, DBAs working on development databases will re-index the database, or run SQL queries that read every record in the database. The workloads generated by these tasks can consume storage resources that are better used by production VMs. In this case it makes sense to use QoS to limit the performance available to this group of VMs.

File Level Restores

Restoring individual files is a time consuming management task that Tintri simplifies. File Level Restore allows users to find and copy their own files. A VM administrator doesn't need to know what files they want to recover, just the drive where they were originally stored. Simply use Tintri File Level Restore to attach a snapshot of the virtual disk(s) from a previous minute, hour, day or week to the VM and then the user can retrieve the files they need.

Example 4—Scaling out with Tintri Storage

Adding VMstores doesn't add complexity.

Virtual infrastructure footprints are constantly changing, but VM Scale-out constantly monitors the environment for changes in workload, storage capacity, and performance reserves. Based on these changes, Tintri Global Center continues to make recommendations and optimize the storage to ensure that VMs are getting the capacity and performance they require.

As the number of VMs increases, and more storage capacity is needed, scaling out the storage is a simple matter of adding new VMstores. The VM Scale-out engine in Tintri Global Center Advanced recognizes when new VMstores are available and starts making recommendations to re-optimize the location of VMs to make the best use of the newly added storage capacity.

Once again, there is no need to worry about the capacity and performance profile of VMs. Tintri analyzes VM-level performance statistics every ten minutes to identify the optimal location of all VMs. Because of Tintri's VM-aware storage filesystem, VM Scale-out can consider not just the capacity requirements of each VM, but also its flash working set, VM activity, thin provisioning, deduplication, snapshots, etc. Tintri recommends a least-cost solution for the optimal placement for all VMs, and all you do is click "Execute". Confirming these actions means that VM Scale-out will move the VMs and optimize the resources available within the storage pool. VM Scale-out also ensures that data protection policies, snapshots, and historical performance data follow the VM.

Summary

Tintri storage has always provided detailed performance profiles for individual VMs. The detailed, per-VM performance counters make it simple to quickly determine the source of performance bottlenecks and latency in the virtualization stack. Now this per-VM performance history is leveraged in the Tintri VM Scale-out and Tintri Analytics tools.

This paper included four deployment examples to illustrate the value of Tintri's VM Scale-out and Analytics products when applied to real-world scale-out tasks:

- Tintri Analytics experiments accurately size deployments of hundreds or thousands of VMs.
- Real-time performance data—summarized by Tintri Analytics—reveals insights into the storage and performance capacity consumed by the VMs.
- VM Scale-out monitors storage capacity and performance reserves so there is no need to worry about the initial placement of VMs; let VM Scale-out manage that task for you.

Day-to-day management of VMs running on Tintri VMstores is performed through the use of VM management tools: snapshots, clones, and the SyncVM data recovery tool. These tools clone, protect, and restore entire VMs, or even individual disks within a VM. Applying the clone tool to VDI deployments results in fast and space efficient desktop clones. Applying the SyncVM tool to database VMs makes synchronizing data between production and development systems simple, fast, and repeatable!

Adding Tintri's VM management tools to Tintri Analytics and VM Scale-out results in a storage scale-out solution that scales capacity and performance without losing the management simplicity that is built in to every Tintri VMstore.