

Tintri Analytics

VM-Level Predictive Analytics for Improved
Resource Planning

Contents

Tintri Analytics – VM-Aware Predictive Analytics for Resource Planning	1
Understanding Tintri Analytics	2
Synthetic Metrics	2
Tintri Analytics Dashboard	3
Defining Applications	4
Using Tintri Analytics for Resource Planning	5
Determining whether or not your current footprint can support new workloads	5
Determining when you will need more storage—for organic or inorganic growth	7
Determining where capacity and performance changes are occurring	8
How To Get Started with Tintri Analytics	9

Tintri Analytics

VM-Aware Predictive Analytics for Resource Planning

IT teams rely on data and analytics to make planning decisions, but sometimes interpreting the available data and determining what action to take can be difficult. You may be able to tell what capacity a storage appliance is using today, and extrapolate what it will be consuming in the future if nothing changes, but the data center is a dynamic place. You need to be able to predict what the storage impact will be when you make changes, like doubling the number of instances of an application, or switching from standard IT to DevOps. Your planning should include not just capacity but also performance.

Most of the analytics tools available don't provide much help with these kinds of questions for two reasons:

- **They are focused on support.** A lot of storage analytics tools are designed primarily to predict hardware failures and facilitate the support process.
- **They have limited data.** Most vendors collect storage data at the LUN or volume level but are unable to correlate that to the virtual machine and application. You can see what's happening inside your storage, but you can't map it to your VMs and applications.

Tintri Analytics takes a different approach; it is designed from the ground up to support resource planning. And like everything that Tintri does, Tintri Analytics is VM-aware storage (VAS), providing the critical link between storage metrics, VMs, and applications. Having the ability to ask "what-if" questions and see answers in real time greatly simplifies resource planning.

This paper explains how you can use Tintri Analytics to answer important resource planning questions:

- Do I have the resources to add more instances of an application?
- When will my storage appliance(s) run out of capacity or performance?
- How is the resource usage of different applications changing over time?

Understanding Tintri Analytics

There are two key elements that distinguish Tintri Analytics from other analytics tools:

- **VM-aware data that facilitates resource planning.** Storage is the most complicated element of IT resource planning. Other storage analytics tools and application-level planning tools don't correlate storage information with application information. Only Tintri joins the two together. Tintri Analytics shows you what your applications need from storage. As a result, you can plan from the application down rather than the infrastructure up and make faster decisions based on actionable data.
- **A next-generation toolchain that enables real-time analysis.** Tintri Analytics runs as a service and is built on Elasticsearch, a "distributed, scalable, real-time search and analytics engine." Elasticsearch lets you sift and summarize large amounts of data very rapidly in response to ad hoc queries, delivering real-time synthesis and analysis of data. Other storage analytics tools use columnar databases, which are designed on the assumption that you know what types of queries you are going to run in advance. Tintri allows you to answer real questions about your footprint as they arise.

Real-time responsiveness matters. You interact with data much differently when you are able to drill down quickly and find things you didn't know you were looking for. If your ability to explore your data is complex or constrained, you're reduced to assumptions and guesswork.

Synthetic Metrics

The metrics displayed through Tintri Analytics are synthesized from raw storage data. Raw information can be displayed, but in most cases it is not useful by itself. Understanding the resource usage inside a storage system has become complex.

There are three primary metrics reported by Tintri Analytics:

- **Logical Capacity.** The size of the virtual machine (or defined group of virtual machines) after thin provisioning but before compression and deduplication. This also includes the size of snapshot data.
- **IO Performance.** An estimation of the IO and throughput resources required by a virtual machine (or defined group of virtual machines). This metric is a linear combination of reads, writes, IOPS, and throughput.
- **Working Set.** The amount of flash needed to store the "hot" data for a virtual machine (or defined group of virtual machines), after Tintri data reduction algorithms have been applied.

Tintri Analytics gives you information that is immediately actionable. Each metric is a synthesis of multiple elements of raw data, so you don't have to try and interpret results. Take performance as an example. If you look at raw IOPS for an application, what is the IO size? And what is the read/write mix? Does throughput matter? Providing IO Performance as a single synthetic metric gives you a single value that combines all these elements and that you can immediately understand in the context of other applications and the total performance available from a VMstore. Tintri Analytics is the only storage analytics tool currently on the market that determines and reports working set size.

Tintri Analytics Dashboard

Figure 1 shows the Tintri Analytics dashboard. VM data, such as the number of live VMs running on your VMstores and average VM needs, is displayed in the left two panes. The panes on the right show resource usage as well as which application profiles are the top consumers for each resource. You can see immediately if you're approaching a resource limit and drill down to find out what is consuming a resource.

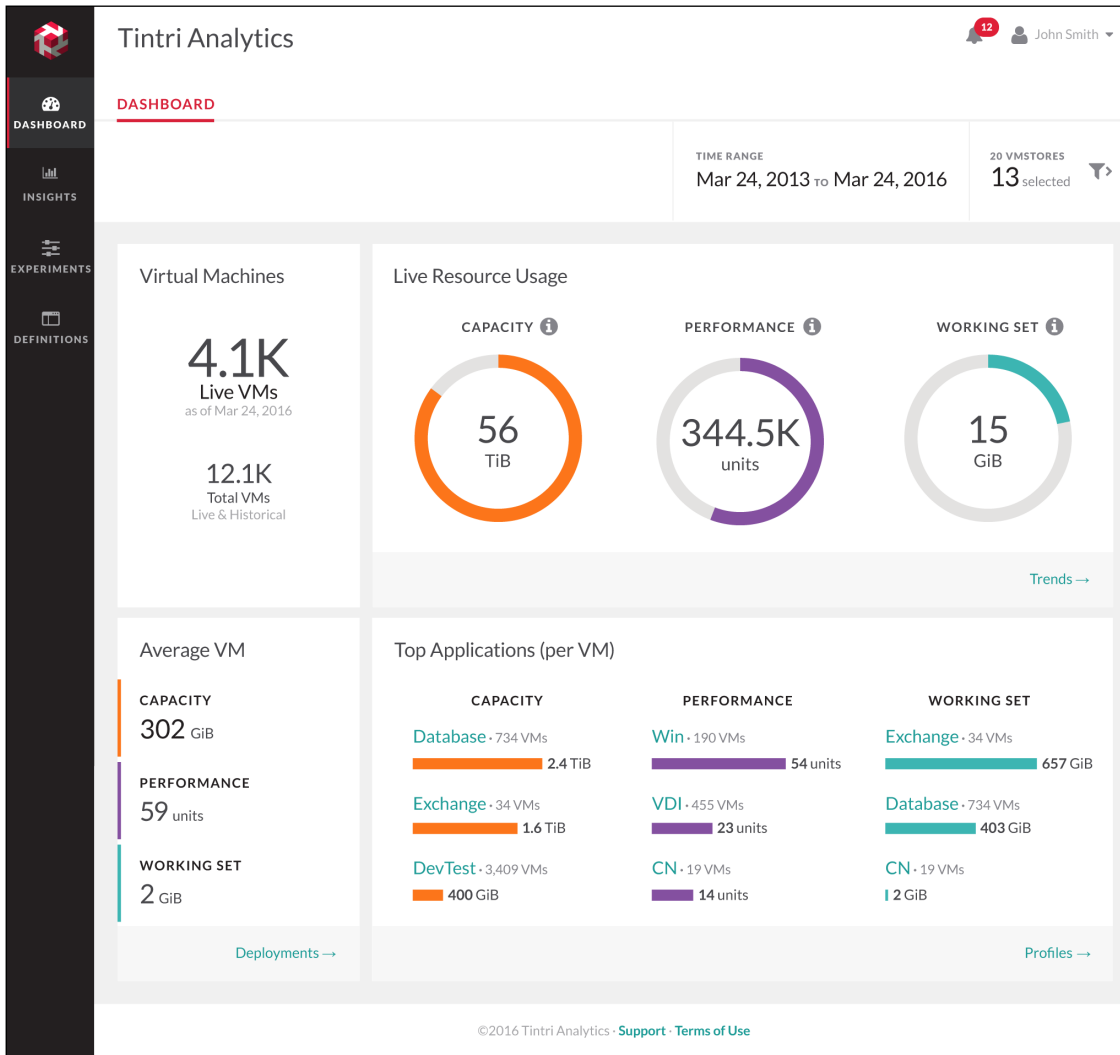


Figure 1 – Tintri Analytics Dashboard.

There are several navigational elements on the dashboard that are consistent across all pages:

- The **Navigation Bar** on the left side lets you switch quickly between the dashboard and other pages.
- The **Filters** above the informational panes, let you adjust the time range and select a set of VMstores to analyze. You can look at VMstores across your entire organization or a single appliance.
- The **Alerts Icon** at the top right has a pull-down menu with actionable alerts. Alerts flag VMstores that are likely to run out of a resource and other issues.

Defining Applications

In Tintri Analytics, an application is a group of virtual machines that all share a common type of workload. You have the flexibility to define applications according to your needs and existing practices. Figure 2 shows the Definitions page used to view and define the applications used in your environment. You can add application profiles and/or drill down from this page to find out more about each application.

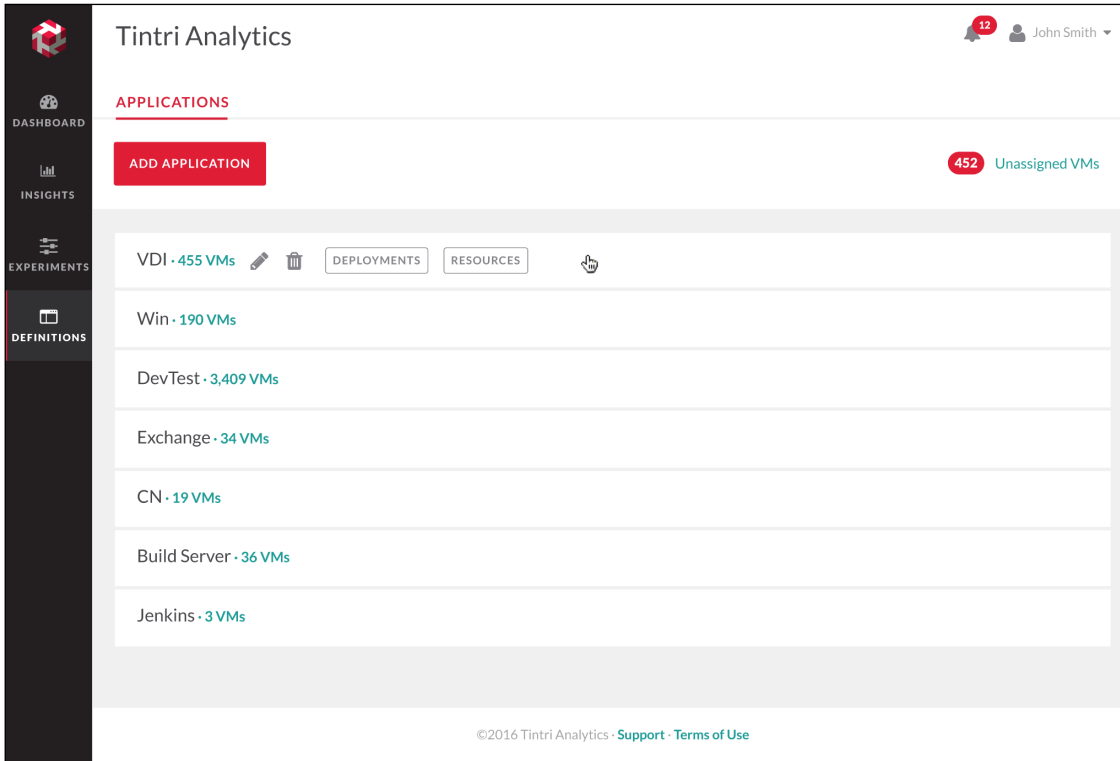


Figure 2 – Tintri Analytics Definitions page lets you define the applications used in your environment.

For example—in Figure 2, you can see an application called VDI, which includes all VMs supporting virtual desktops. Figures in later sections show data for this VDI application.

Using Tintri Analytics for Resource Planning

Tintri Analytics can help you better address a number of common resource planning tasks. You can model a variety of “what-if” scenarios to identify the best course of action. You can also perform trending analysis based on key metrics and look at changes in the rate of resource consumption.

Determining whether or not your current footprint can support new workloads

There are many situations where you need answers to what-if questions. For instance, you might need to know if your storage can support 10 more webserver VMs or 25 more VDI seats. Before Tintri Analytics, it was very difficult to answer these questions, especially for metrics other than capacity. Because of deduplication and compression, even capacity has become a complicated metric.

Figure 3 illustrates how a "what-if" query works. You simply go to the Experiments page, choose the desired application in the Filter Bar (in this case VDI) and select the number of instances (in this case 25). There is also a dropdown that allows you to assign the Target VMstores for this analysis, in case you want to model the impact on a subset of your total footprint. Tintri Analytics shows you the results in real time.




Figure 3 – Using the Tintri Experiments page to see the impact of adding 25 additional VDI instances.

You can see that in this environment, adding 25 VDI VMs would overload the 7 target VMstores in terms of performance. It would also be dangerously close to the limit for the working set.

Fine Tuning

Tintri Analytics gives you a number of options for fine tuning the results:

- **Adjust per-instance values.** You can use the average values for your existing VMs (in the example above the average capacity used by a virtual desktop in your VDI application profile is 12 GiB). Or, you can set arbitrary values, if for example you know that the next VMs to rollout will consume significantly more or less resources (e.g. bump the 12 GiB up to 20 GiB).
- **Target more (or fewer) VMstores.** If you have multiple VMstore systems in your environment, you probably have a preferred target for different application types. However, if your preferred VMstore(s) can't support the added load, you can see what happens if you spread the load across more appliances.
- **Adjust the number of instances.** If your storage can't support 25 new instances, you can quickly determine how many it actually can support by adjusting the number of instances up and down.

Because results are always updated immediately, this process is quick and painless, and allows you to zero in on the best course of action.

Determining when you will need more storage—for organic or inorganic growth

You have to monitor storage resource consumption on a regular basis so that you can take action before you run out of resources. Tintri Analytics makes this analysis fast and simple. The Insights page graphs resource usage over time and provides trend lines as shown in Figure 4. (Be sure and select “All VMs” from the drop down menu on the Filter Bar to see results for your full data set.) With this view, you can immediately identify what the critical resource is and when it is likely to run out.




Figure 4 – Trends in resource consumption help you determine when to add additional storage hardware.

In this example, the critical resource is capacity, which is projected to run out in about 3 months, meaning it's time to take action.

Fine Tuning

Once again, there are a number of available actions to fine tune your results:

- **Look at a subset of VMstores.** Click the funnel icon on the Filter Bar to view results for a particular subset of your VMstore appliances or a single VMstore.
- **View results by application.** Once you've reviewed the results for "All VMs" you can view the same data on a per-application basis to find out which applications are consuming resources and how application consumption is trending over time.
- **Change the timeframe.** If you suspect that resource consumption is accelerating (or decelerating), it may be helpful to shorten the timeframe so that earlier data is excluded. This can give you a more accurate projection of when resources will run out.

Determining where capacity and performance changes are occurring

In a virtualized environment, capacity requirements tend to grow over time. As you patch VMs and make other changes over time, VMs tend to accrete data. A question that can be hard to answer is, “am I adding more VMs or are my VMs getting bigger?”

The Profiles tab on the Insights page lets you see how the resource needs of an application profile change over time. Figure 5 shows this view for VDI.




Figure 5 – By adjusting the Time Range, you can see how the capacity usage and other metrics for an application VM are changing.

Using the Time Range filter, you can see how the resource needs of an application VM have changed over time. If capacity is growing, that may be a sign that your VMs are accreting data. By understanding where growth is happening in your environment—inside VMs or due to added VMs—you can take control of your footprint.

From the Profiles tab you can also drill down to see the raw data behind each metric or visualize the distribution of VMs in the application for each metric to identify outliers.

How To Get Started with Tintri Analytics

Tintri Analytics is a storage analytics tool unlike anything else on the market. These examples are a starting point, but you will almost certainly find additional uses for the wealth of actionable data that Tintri Analytics provides.

Tintri Analytics is available to Tintri customers that have a valid maintenance contract. It's a central element in Tintri's broader scale-out architecture, working in concert with VM Scale-out software to help organizations optimize their virtual footprint for both cost and performance.

As you grow your storage footprint, so grows the value of the data it generates. That's why the right analytics can have a profound impact on the way you make decisions that affect your data center and your business. Tintri Analytics helps you plan for resource needs in a dynamic, changing environment. And its VM-aware storage (VAS) foundation offers you complete visibility of your entire storage footprint. With Tintri Analytics, you'll make the right decision.

Tintri believes the information in this document is accurate as of its publication date. The information in this publication is provided as is and is subject to change without notice. Tintri makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims any implied warranties of merchantability or fitness for a particular purpose.

