



Visualizing Intelligent Data In Motion

Tintri and Turbonomic offer solutions that improve virtualized machine performance and efficiency by providing Quality-of-Service control, storage provisioning efficiency with per-VM management and visibility that helps deliver the true essence of a Software Defined Datacenter.



Conventional disk-based storage systems spanning datacenters across the world were never designed to accommodate the requirements of virtualized and shared workloads, Quality of Service (QoS), and the 'minute' to 'massive' scale-out virtual and cloud computing infrastructures. This is an ideal solution for any customer looking to host business critical applications that require predictable performance and complete application and VM (VM) control across their infrastructure – host, network, and storage. In today's rapidly evolving world, no VMs are created equal – so you shouldn't have to host and manage all your applications in a single one-size fits all container.

Benefits and Capabilities

Tintri VMstore and Turbonomic Operations Manager help fine-tune your VM workload to assure performance while maximizing efficiency. Automatable decisions can help assure that business critical applications are not affected by resource contention, host, network, or storage latency and prioritize mitigations according to business and performance criteria.

- Automated performance-enhancing decisions at the hypervisor, network, and storage level for continuous workload assurance
- QoS custom controls allowing minimum and maximum IOP configurations for business critical VMs
- Concurrent and seamless support for best-of-breed hypervisors with VM-level management across the entire infrastructure
- Visualization of storage and capacity trends at the VM-level
- 100% VM and virtual disk visibility
- Get per-VM snapshots, cloning, and efficient replication
- Demand-driven control to assure performance while maximizing utilization

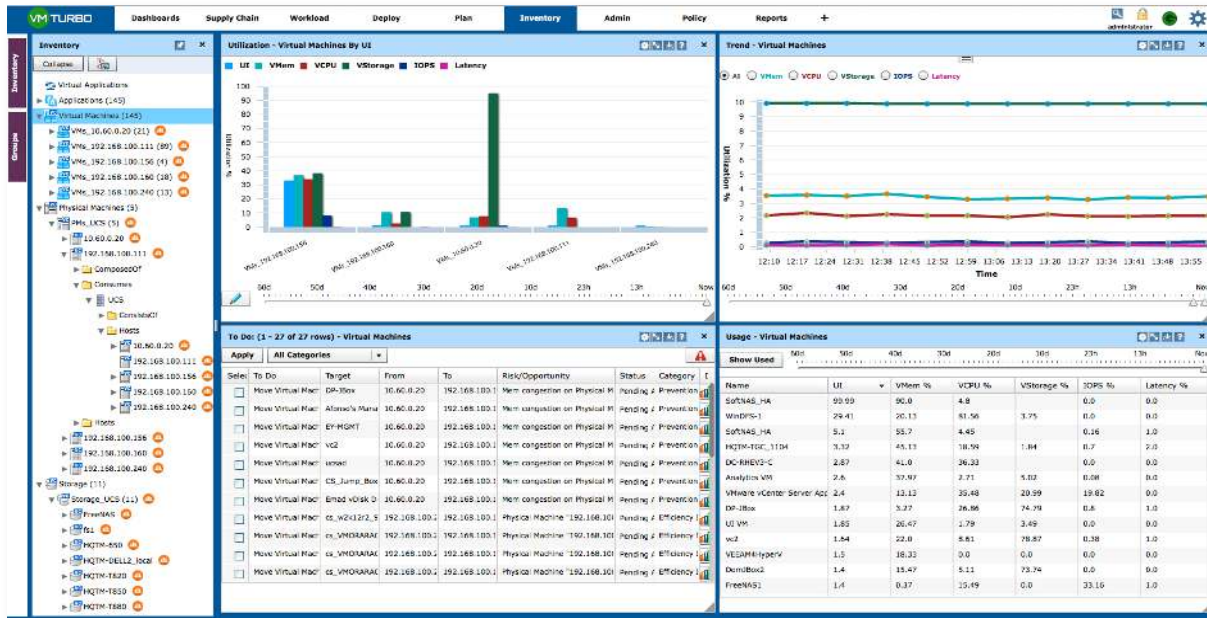


Figure 1 – Turbonomic Operations Manager dashboard and analytics

The Storage Management and Visibility Challenge

Current storage arrays were never designed for the unique challenges presented by large-scale, multi-tenant cloud and virtualized datacenters. Performance, QoS, availability, scale and cost requirements are different in yesterday's enterprise environments. Tintri's storage system is built from the ground up to specifically address these aspects for the workloads of today and tomorrow.

These legacy systems are LUN-centric or volume-centric, and when it comes to snapshots, clones, and replication, it can be extremely wasteful and overly complex to manage. A VM consisting of many files (VMFS), residing on storage that is shared with other VMs is not a good thing. Here are a couple of reasons:


First, if there is a need to replicate a specific VM (either locally or remotely) then all the VMs within that LUN/Volume get replicated. There is no way to selectively choose a particular VM or application. Time, network bandwidth, and disk resources are unnecessarily over consumed when you need access to just a specific data set.

Second, from a storage array perspective, the LUN is the lowest level of granularity in terms of performance and QoS. As the storage array has no way to determine the individual contents of the LUN, it can't prioritize workload by VM. What this means is that *all* VMs, which are independently unique in size, workload, performance, and availability, must share a common configuration, originally configured for *one* specific workload or application.

To support the demands of modern day virtual and cloud infrastructure, QoS should allow mission critical workloads to be intertwined with less critical ones, without the impact of a 'noisy' neighbor stealing available storage performance. The problem is that most storage QoS implementations are not tuned for virtualization and can actually make the process of assuring mission critical performance for VMs more complicated than without QoS.

Manage vDisks,
not LUN's or Volumes

- Manage the *virtual*, not the physical
- *Application* needs, not infrastructure needs



The VM-Centric Solution

Tintri and Turbonomic solve this challenge, respectively, by making modern flash-based hybrid storage arrays VM-centric and providing automated demand-driven allocation for better VM performance and QoS. Unlike all other storage platforms, that are blind even to the concept of a VM, Tintri is application-aware, able to peek into the software environment to reveal and even call attention to issues. Tintri VMstore automatically surfaces information on VM behavior and yields insights about performance – including IO, storage growth, and utilization – for individual VMs.

Tintri's patented technology unlocks traditional boundaries of storage performance typically measured in Input/Output Operations per Second (IOPS) and capacity (GB). IT storage administrators can now provision, snapshot, clone and dynamically adjust each VM or virtual disk independently based on the unique demands of the individual application workloads. This level of isolation and flexibility solves the “noisy neighbor” problem that often plagues virtual workloads residing in traditional storage arrays where different workloads share a common LUN or volume resource configuration. Simply put, Tintri allows for automated management and protection of individual VMs than blindly applying policies to LUNS/Volumes as a whole.

We believe that performance measurement should be more than just a calculation of IOPS, but to also highlight on other factors including latency, throughput, network utilization, and the impact on vCPU. This joint solution outlines the ability to report on how each of these resources is being utilized in a visual fashion and provides immediate feedback on any changes that are made to those settings. You will realize why all the factors matter in terms of accurate performance monitoring and how they jointly enable IT administrators to make quick, knowledge-based decisions with new levels of confidence, efficiency, and ease.



Figure 2 – Tintri VMstore dashboard showing per-VM statistics and latency measurements

The Intelligent QoS Control

Predicting resource utilization and the performance of each individual VM can be an extremely daunting task, which can be challenging and time consuming, if not impossible, for any IT Administrator to figure out. By leveraging the multiple levels of custom controls from Tintri and Turbonomic Operations Manager, you can now eliminate any guesswork around Quality of Service, allowing you to guarantee and protect the performance of individual VMs.

With Tintri, system administrators can see the immediate impact of changes on VM-level as visualization spans the entire infrastructure, including latency stemming from host, network, storage and the QoS throttle. Admins may choose to toggle the QoS settings to limit

the resources rogue VMs can consume by reducing the maximum number of IOPS. Or they can guarantee the minimum number of IOPS to meet the high demands of a mission critical business application. Volumes and LUNs can contain hundreds of VMs running applications with unique service requirements and assigning the same QoS setting to all VMs offers very little value. This on-demand control allows you to set guidelines and tiers of service for each individual VM to meet and satisfy service-level agreement (SLAs). For internal customers or cloud service providers, the solution help ensures that an end-user who purchased a 'platinum' service do not experience a 'bronze' level service.




Figure 3 – Lower maximum IOPS for 'noisy' VMs




Figure 4 – Increase minimum IOPS for business critical VMs

Turbonomic enables you to define different layers of QoS levels for any of your business applications. Its Demand-Driven Control platform will make specific resource allocation decisions that adhere to the specific commitments of your SLAs. Conversations with application owners are elevated to a discussion about delivering on a specific QoS (e.g. response time, transactions per minute) after initial infrastructure sizing and resource configuration have been set.

Tintri's VMstore QoS visual dashboard detects latency within the storage appliance, and Turbonomic mitigates latency upstream of the environment. Let's take network 'east-to-west' network traffic latency as an example. Turbonomic's Network Aware Placement (NAP) control collocates 'chatty' VMs and workloads at the host level to reduce potential bandwidth bottlenecks without sacrificing compute and storage resource utilization needs. The solution will recommend placement, and start/stop actions across hosts, clusters, and datastores to prevent queuing, latency, ballooning and swapping. These actions can be executed manually ad-hoc or automated on a user-defined schedule.

Virtualized and cloud datacenters are modeled like a stock market where you have buyers who are the workloads, VMs, applications, and hosts. Numerous entities like vMemory, vCPU, network bandwidth, and disk space are the hot commodities that are traded, bought, and sold. Turbonomic provides response time guarantees and assures VM performance without affecting physical servers by delivering an algorithm designed to focus on fluctuating workload demands, such as the economic stock market. Through continuously brokering the supply and demand cycles, automation and recommendations on placement, sizing, and start/stop decisions fine-tune consumer-provider relationships to meet business application demands.

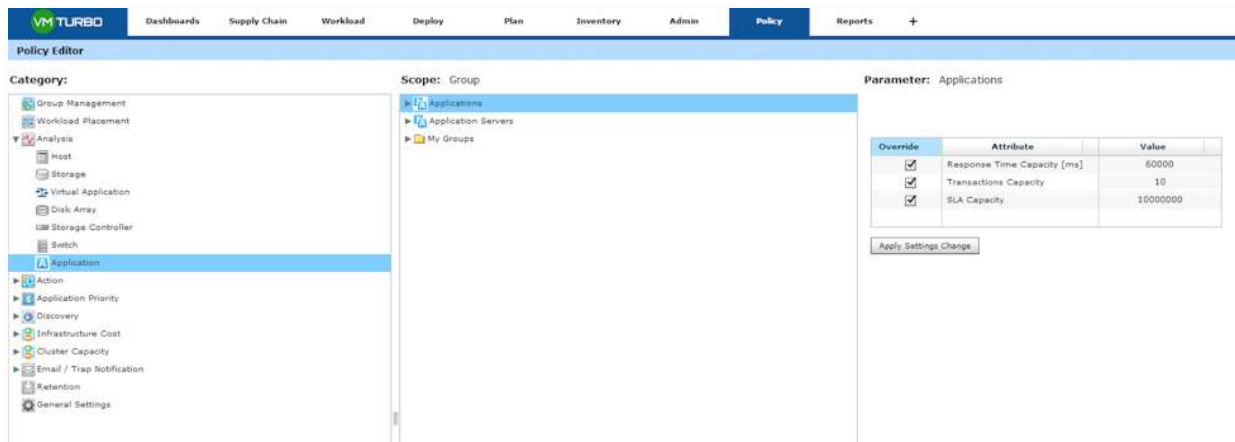


Figure 5 – Specify application QoS definitions for automated SLA adherence

Summary

The complementarity of the solution comes from Tintri's architecture and ability to do all things storage at the VM-level including performance guarantees and per-VM visibility, snapshots, cloning, and replication across different datacenters. Turbonomic provides the upstream level of VM response time guarantees and assures VM performance while maximizing utilization. It's an end-to-end solution for 100% VM control.

Turbonomic is widely adopted by customers utilizing heterogeneous storage. If what's being stored are VMs, Tintri may be the optimal solution to extend VM-level management from host through network and storage.



Tintri builds smart storage that sees, learns and adapts, enabling IT organizations to focus on virtualized applications and business services instead of managing storage infrastructure. Tintri application-aware storage eliminates planning and complex troubleshooting by providing VM-level visibility, control, insight and agility, with all flash performance for virtualized environment and the cloud.

Tintri powers hundreds of thousands of VMs running business critical databases, enterprise apps, desktops and mobile apps, and private cloud deployments. Tintri helps global enterprises such as AMD, F5 Networks, GE, NEC, NTT, MillerCoors and Time Warner maximize their virtualization and cloud investments. For more information, visit www.tintri.com and follow us on Twitter: [@Tintri](https://twitter.com/Tintri)

303 Ravendale Drive
Mountain View, CA 94043
Main: +1 (650) 810-8200
www.tintri.com



Turbonomic's autonomic platform is trusted by enterprises around the world to guarantee the performance of any application on any cloud or infrastructure. Turbonomic detects changes to application demand in real-time and adjusts cloud and infrastructure resources to maintain a perpetual state of health.

Launched in 2010, Turbonomic is one of the fastest growing technology companies on the market. Leveraging Turbonomic's autonomic platform, customers can confidently accelerate their adoption of cloud, virtual, and container deployments.

With Turbonomic, customers drive real-time performance, guarantee a Quality of Service, build confident agility, and minimize OpEx/CapEx spend. To learn more, visit turbonomic.com.

500 Boylston Street
Boston, MA 02116
Main: +1 (866) 634-5087
www.turbonomic.com