

the *Availability Digest*

www.availabilitydigest.com
[@availabilitydig](https://twitter.com/availabilitydig)

VMTurbo – Managing Virtualization

April 2015

VMTurbo



VMTurbo is a monitoring and management tool that maintains a virtualized data center in the “Desired State.” VMTurbo defines the “Desired State” of a virtualized environment as that which eliminates inefficiencies without compromising application performance.



Maintaining the Desired State is typically a complex, time-consuming, and error-prone job requiring the skills of highly experienced engineers. The ever-changing mix of applications coupled with the requirements to meet SLAs, to provide the specified redundancy, and to comply with affinity rules require constant monitoring and adjustment of workloads across the virtualized infrastructure. (See our companion article in this issue entitled “Managing Your Private Cloud.”¹)

Traditional virtualization performance tools that monitor VMware or Hyper-V use threshold-based alerts. When a metric crosses a threshold, an alert is triggered and the operations team must decide whether the alert is ‘smoke’ or ‘fire’. VMTurbo is a proactive management tool that maintains the virtualized environment in the Desired State automatically. It *prevents* problems rather than issuing alerts indicating that there are problems.

Demand-Driven Control

Traditional virtualization management solutions leave it to the IT staff to monitor dashboards, alerts, and reports and to make the appropriate reallocations of virtualized resources to keep the virtual machines (VMs) in compliance with corporate policies.

VMTurbo continuously monitors resource usage and automatically reassigns VMs to compute clusters to maintain corporate policies. Each application is different and is governed by the requirements of its own SLA. The set of SLAs (for perhaps hundreds or even thousands of applications) and other configuration criteria represent the corporate IT policies for the virtualized environment and include:

- Application performance criteria, whether it be response times or batch-running times.
- The maximum allowable loading of host servers.
- The required redundancy for each application.
- Affinity requirements:
 - affinities - some applications should always run on the same host server for performance reasons.
 - anti-affinities – some applications should never run on the same host server for failure reasons.

¹ [Managing Your Private Cloud, Availability Digest, April 2015.](http://www.availabilitydigest.com/public_articles/1004/virtualized_data_center.pdf)
http://www.availabilitydigest.com/public_articles/1004/virtualized_data_center.pdf

- Load balancing to even out the load on virtualized resources.
- Recovery to a remote site for critical applications.
- Capacity monitoring to determine future resource expansion requirements.

VMTurbo will recommend placement, sizing, and start/stop actions across hosts, clusters, and datastores to prevent queuing, latency, and ballooning. As VMTurbo detects potential problems, it can automatically reconfigure the virtual environment to prevent those problems. Alternatively, it can provide suggestions to the IT staff to take to avoid problems. These suggestions are provided in an intuitive set of instructions with the ability for the IT staff to approve any of the suggestions, which then VMTurbo will execute:

To Do: (1 - 10 of 10 rows)					
Apply All Categories					
<input type="radio"/> Show All <input checked="" type="radio"/> Show Top					
Select	To Do	Target	From	To	Risk
<input type="checkbox"/>	Move Virtual Machine	ubuntu4	esxaut	esxauto3	Mem congestion on Physical M
<input type="checkbox"/>	Move Virtual Machine	yuval-test	vmturb	vmturbo	CPUProvisioned congestion or
<input type="checkbox"/>	Provision new Physical Machine	vmturboclo		vmturbo	Critical Mem congestion. Provi
<input type="checkbox"/>	Provision new Physical Machine	vmturboclo		vmturbo	Critical Mem congestion. Provi
<input type="checkbox"/>	Increase VMem capacity	DVS Contrc	1.43 G	2.43 GB	VMem congestion in Virtual Mi
<input type="checkbox"/>	Increase VCPU capacity	UbuntuLG9	1	2	VCPU congestion in Virtual Ma
<input type="checkbox"/>	Increase VCPU limit	UbuntuLG7	1071.0	2871.0 M	VCPU congestion in Virtual Ma
<input type="checkbox"/>	Increase VCPU limit	UBuntuLG4	446.0	2246.0 M	VCPU congestion in Virtual Ma
<input type="checkbox"/>	Increase VMem capacity	vCenter_VA	8.0 GB	10.0 GB	VMem congestion in Virtual Mi
<input type="checkbox"/>	Reduce VCPU limit	UbuntuLG6	714.0	0.0 MHz	VCPU congestion in Virtual Ma

VMTurbo Operations Manager

The VMTurbo Operations Manager is the basic module of VMTurbo. It works with the hypervisor to provide the following functionality:

Workload Multiplexing: Workloads with coincident load spikes are identified and placed on separate physical hosts or datastores to prevent contention and to enable high VM density.

Automatically prioritize mission-critical applications: VMs can be classified into criticality tiers. If resources become constrained, lower tier applications are throttled and higher level tiers are prioritized to assure mission-critical performance.

Assure redundancy: N + K redundancy (K spares in addition to the minimum number of hosts, N, needed by the application) is guaranteed at all times to meet the SLAs of applications, even in dynamic environments.

Future Planning: Incorporate growth and high availability together before the environment is out of compliance.

Test Fault Tolerance: Ensure that a host failure or a disaster recovery event is properly handled with detailed workload placement and resource reconfiguration decisions.

Disaster Recovery: Automatically manage the disaster recovery environment to a remote data center in the event of a disaster.

Workload Reservations: Reserve space in the infrastructure for upcoming application deployments. Issue alerts if the infrastructure can no longer support the anticipated load. Upon deployment of new applications, identify on which targets (clusters, hosts, and datastores) each new VM should be deployed.

Self-service provisioning: Integrates with cloud management platforms including vRealize Automation, OpenStack, CloudStack, System Center Orchestrator and VMM to place new VMs directly in the right cluster, datastore, and host depending upon the anticipated load so that existing workloads are not impacted. New VMs can be based on custom or pre-defined templates.

VMTurbo Control Module Extensions

The VMTurbo Control Modules extend the functions of the Operations Manager from the hypervisor into each layer of the data center to provide tighter control of the Desired State and to further ensure application performance and resource utilization efficiency.

- The **Application Control Module** discovers and manages applications as a single entity, enabling policy-based prioritization for tiers of service. It understands the real-time performance characteristics of the application layer to make resourcing decisions that ensure performance across the board.
- The **Network Control Module** manages the tradeoffs between compute load, storage load, and network traffic flow to utilize the network resources as efficiently as possible.
- The **VDI (Virtualized Desktop Infrastructure) Control Module** controls the demand of I/O intensive events such as reboots and patch updates to maintain the desired end-user experience without having to run VDI deployments on a dedicated or overprovisioned infrastructure.
- The **Container Control Module** extends the Operations Manager to manage containers, ensuring the performance of containerized applications while utilizing the underlying infrastructure as efficiently as possible.
- The **Storage Control Module** manages NetApp, HP 3PAR, EMC, and Pure Storage to extend the control created for virtual servers to the storage layer.
- The **Fabric Control Module** provides the complete visibility and control of Cisco's UCS network fabric and converged infrastructure offerings such as VCE Vblock and NetApp FlexPod.
- The **Hybrid Cloud Control Module** provides the intelligence to match applications workload demands with the right resources, whether these resources reside in the data center private cloud, the public cloud, or a hybrid combination.

Free VMTurbo Services

VMTurbo Virtual Health Monitor

The VMTurbo Virtual Health Monitor is a free and unlimited virtualization monitoring and reporting tool. It monitors compute, storage and network performance in a single dashboard. Its monitoring includes heterogeneous environments using VMware, Hyper-V, and Xen.

VM metrics include vCPU, vMem, vStorage, I/Os, and latency. Host metrics include CPU, memory, I/Os, network, swap, and ballooning. Storage metrics include I/Os, latency, capacity, and utilization.

Business Impact Assessment

In addition, VMTurbo also provides a free “Business Impact Assessment so that an organization can clearly understand how effective it would be in mitigating the risks to which it is exposed in the event of failures in key components of its virtualized infrastructure, including complete data center failures. The assessment involves the following steps:

Scoping: The scope of the Business Impact Assessment is determined.

Software Deployment: VMTurbo’s assessment software is deployed.

Data Collection: VMTurbo’s software automatically analyzes the target environment over a one-to two-week period.

Data Analysis: VMTurbo’s engineers perform a number of “what-if” simulations to determine the risk to which IT operations might be exposed in the event of different failure scenarios.

Executive Report: VMTurbo prepares a report detailing the results including technical and financial analyses.

VMTurbo – The Company

Founded in 2009, VMTurbo’s charter is to transform IT operations in the cloud and virtualized environments from a complex, labor-intensive, and error-prone process to one that is simple, automated, and predictable. Its products ensure that applications get the resources required to meet corporate IT policies while maximizing the utilization of IT assets.

With offices in New York, California, the United Kingdom, and Israel, VMTurbo is used at over 9,000 enterprises worldwide.

Summary

Virtualization monitoring tools alert you when you have a problem and allow you to investigate. VMTurbo actively manages your environment to prevent alerts and problems from occurring.

VMTurbo can be deployed in about thirty minutes, meaning that the operations staff and engineering teams can quickly test and verify how VMTurbo can control their virtualized environment. Within one hour of deployment, VMTurbo will identify begin identifying performance issues and provide the specific actions to take to remedy the issues. It then continuously evaluates the changing demands of the organization’s applications against the infrastructure availability to deliver and maintain the Desired State of the virtual environment.

VMTurbo claims that its customers typically increase VM density by 20% to 40% without risking degradation in their environment. They further claim that many VMTurbo customers have reduced their infrastructure and licensing costs by 40% to 70%.