

TOP CHALLENGES IN STORAGE VIRTUALIZATION

(AND HOW TO OVERCOME THEM)

TURBONOMIC WHITE PAPER

TOP CHALLENGES IN STORAGE VIRTUALIZATION

(AND HOW TO OVERCOME THEM)

Executive Summary	3
Storage Challenges in Virtualized Environments	4
The Turbonomic Approach	5
Mapping Workloads to the Storage Domain	6
Common Storage Problems	7
Conclusion	11

EXECUTIVE SUMMARY

One of the biggest challenges in virtualized infrastructure impacting application performance is the ability to manage resource contention in the underlying storage infrastructure. A common approach to reduce the impact of this challenge is to over-provision storage platforms with IO and disk capacity and manage mission-critical and other IO-intensive applications in silos. However, this is a trade-off that leads to a higher cost per gigabyte of storage allocated to every virtual machine, as well as significantly impacting operational costs associated with service assurance.

The adoption of so-called Solid State Drive (SSD) all-flash arrays tends to eliminate the IO bottleneck, as all-flash arrays offer 10x+ the IOPS capacity of Hard Disk Drive (HDD) arrays. However, the challenge of workload alignment, heterogeneous storage management, and efficient capacity planning remain when introducing SSD arrays into virtual and cloud environments.

Turbonomic's autonomic approach to virtualized storage management addresses these challenges by discovering the full stack of virtualized and cloud environments and enabling workloads to self-manage.

Using Turbonomic, enterprises and service providers can:

- Reduce ongoing storage infrastructure costs by 20-30%
- Significantly reduce operational costs by preventing complex storage problems and their impact on workloads, and, in turn, on end users
- Ensure application performance and reduce risk
- Enable IO-intensive applications to be virtualized reliably, bringing down the overall cost to deliver compute services to the business and customers
- Seamlessly implement heterogeneous SSD/HDD storage environments by determining which workloads require SSD access and when

This white paper explores the challenges of controlling the data center's virtualized storage infrastructure and the Turbonomic solution.

TOP CHALLENGES IN STORAGE VIRTUALIZATION

(AND HOW TO OVERCOME THEM)

Storage Challenges in Virtualized Environments

Virtualization hinges upon a “shared everything” infrastructure including shared storage, which has enabled many virtualization innovations such as live migration (vMotion in VMware environments), automated high availability, disaster recovery, and many others. Virtualization allows businesses to significantly lower the cost of compute services and ease management. Due to the shared nature of storage in a virtualized environment, multiple virtual machines must contend for underlying storage resources in order to satisfy workload demands. This includes consuming resources from the storage domain such as storage space, IOPS, and other resources consumed indirectly from the storage infrastructure, such as controller CPU load from the hosting storage arrays.

Numerous studies and surveys confirm that the storage underlying the virtualization layer continues to drive significant problems for administrators and operations teams. According to Gartner’s July 2014 Magic Quadrant for x86 Server Virtualization Infrastructure, average x86 server virtualization levels have surpassing the 70% mark.¹ A VMware report revealed that the percentage of tier-1 (i.e., business-critical applications), such as Microsoft Exchange, Oracle, and SAP are still trailing this,² with much of the reticence to virtualize stemming from the concern around virtualizing the most IO-intensive (and typically business-critical) applications.

The two most common reasons for this are storage costs and storage performance concerns. According to DataCore survey,³ the majority of respondents reported that storage accounted for more than 25% of their overall virtualization budget, and that storage costs are proving prohibitive to the virtualization of mission-critical (storage-intensive) applications. Ensuring storage IO performance to virtualized workloads is also a consistent top cause of virtualization performance issues, with many having learned first hand of application performance problems relating to storage issues, resulting in a lack of confidence in the ability to virtualize the most IO intensive applications. This has led to missed opportunities to significantly reduce compute service cost and increase agility to the business. The desire to solve for storage performance problems has also led to storage over-provisioning, further exacerbating the storage cost problem.

¹ Gartner, July 2014. Magic Quadrant for x86 Server Virtualization Infrastructure.

² VMware, March 2012. Percentage of tier-1 instances that are virtualized: SAP (40%), Oracle database (35%), Microsoft Exchange (47%)

³ Datacore, The State of Virtualization and the Impact of Storage, March 2013.

TOP CHALLENGES IN STORAGE VIRTUALIZATION

(AND HOW TO OVERCOME THEM)

Tackling these issues has proved challenging for a number of reasons. In many cases, the technology teams and operational processes supporting virtualization and storage are fragmented, as are the tools being used to manage these domains. This makes it difficult to understand the key dependencies between them and control them in an optimally efficient and performant state.

The Turbonomic Approach

The Turbonomic solution controls the compute and underlying storage infrastructure in order to assure application performance, while maximizing efficiency in the underlying infrastructure.

Its patented economic model enables workloads to self-manage, ensuring they continuously get the resources they need to perform. The solution abstracts the virtualized IT stack into a service supply chain, or marketplace, of buyers and sellers of resources. Following economic principles of supply and demand, as utilization of a resource increases so does its virtual price. Workloads simply shop around for the best overall price for all the resources they need to perform. By allowing workloads to self-manage in this way, the workloads make decisions based on all the compute, storage, and network resources it needs—not just a single metric in isolation. As a result, the platform provides specific automatable workload placement, sizing, and provisioning decisions. This autonomic approach prevents performance issues from occurring in the first place.

With respect to storage, workloads are consuming storage space and IOPS from the underlying storage infrastructure. Turbonomic simply applies the principles of supply, demand and price, which enable workloads to make a broad set of resource allocation decisions including virtual machine placement (across data stores) and the sizing and creation of underlying storage entities including volumes/LUNs/NFS mounts/exports, aggregates (RAID groups) and storage controllers/arrays.

The following sections examine common storage problems and how the Turbonomic approach works to prevent these problems, optimizing disk space and IOPS usage. The resource allocation decisions apply equally well whether administrators and operations teams are deploying virtual machines for the first time or operating the environment on an ongoing basis.

TOP CHALLENGES IN STORAGE VIRTUALIZATION

(AND HOW TO OVERCOME THEM)



Mapping Virtualized Workloads to the Underlying Storage Domain


Fundamental to driving optimization across virtualization and storage domains is to understand the topology and connectivity between them. Virtual machines that map to data stores in the virtualization domain, map to a complex storage topology “underneath.” This includes LUNs (or NFS mounts), volumes, aggregates (or RAID groups) and, ultimately, physical disks (the pinning is less complex in all-flash arrays, for which aggregate and RAID provisioning do not exist). Mapping from “virtual machine to spindle” is a necessary pre-requisite to providing real-time control and optimization of both domains to obtain desired storage efficiency gains and workload storage performance assurance. Since teams, processes and traditional management tools are often restricted to a single domain in isolation, it can be extremely challenging and time consuming to assure performance as teams take readouts from separate tools and try to “join the dots.” These activities often occur in firefighting scenarios when application performance has already been impaired. Furthermore, given the dynamic nature of virtualization, workloads will frequently move to different compute and storage locations, which means that the relationships between both domains must be continually tracked in order to control and optimize the environment effectively.

Turbonomic’s full-stack understanding, includes the detailed storage domain and its interconnection with the virtualization domain. This allows administrators and IT operations to understand the key dependencies between entities and across layers of the stack, while the autonomic platform makes the decisions that assure performance, while maximizing efficiency.

Since teams, processes and traditional management tools are often restricted to a single domain in isolation, it can be extremely challenging and time consuming to assure performance as teams take readouts from separate tools and try to “join the dots.”

TOP CHALLENGES IN STORAGE VIRTUALIZATION

(AND HOW TO OVERCOME THEM)



The Turbonomic platform will make resource allocation decisions including virtual machine (data store) placement, as well as sizing/creation decisions on the underlying volumes, aggregates and arrays.

Common Storage Problems and the Turbonomic Solution

Common storage domain problems when deploying and operating virtualized workloads can be further broken down:

Storage IO and Workload Performance

As stated above, assuring the storage IO performance of virtual workloads is critical to enabling the successful virtualization of mission-critical IO-intensive applications. The storage IO demands of virtual machines must be satisfied while efficiently leveraging the underlying storage infrastructure. The Turbonomic platform will make resource allocation decisions including virtual machine (data store) placement, as well as sizing/creation decisions on the underlying volumes, aggregates and arrays. This goes beyond what is possible by examining the virtualization domain in isolation. For example, when looking at the virtualization domain alone, it may appear desirable to Storage vMotion a virtual machine from one data store to another to relieve latency/IO—even though, upon closer examination of the storage infrastructure, these data stores may share the same underlying aggregates (RAID groups) and disks, which means that the IO contention will not be alleviated by the move.

Turbonomic discovers the types and IO capacities of the underlying disks used by the arrays under management, and calculates the IO capacity of the underlying storage entities (volumes, aggregates). As IOPS usage driven by workloads becomes high and, therefore, more expensive in the Turbonomic marketplace, resource allocation decisions to resize/create volumes and/or aggregates may be required in addition to workload placement to satisfy workload demands and keep the environment healthy.

TOP CHALLENGES IN STORAGE VIRTUALIZATION

(AND HOW TO OVERCOME THEM)

Controller CPU Load

Turbonomic treats controller CPU as a commodity in the economic marketplace and, therefore, prices this commodity according to supply and demand. Resource allocation decisions may include upgrading the controllers in an array or potentially adding more controllers via adding more storage arrays to the infrastructure.

Storage Space Obfuscation

As discussed previously, there are multiple layers within the virtualization and storage domains that must be constantly mapped and related, to drive storage IO and space efficiencies while satisfying the demands of hosted virtual workloads. There are a number of storage domain technologies that can further obfuscate the underlying storage space supply and demand, making it difficult and time consuming for operations teams to understand the actual available storage space in the supporting storage infrastructure.

Thin Provisioning

Thin provisioning is a technique to oversubscribe available storage resources, so that storage consumers believe there is more space available than advertised from storage suppliers. This increases usage efficiency of the underlying storage. Thin provisioning may be implemented at the virtualization (workload) layer as well as at the storage (volume) layer. It is also possible to enable thin provisioning at the virtualization and storage layers concurrently, sometimes referred to as a “thin on thin” configuration. These configurations can make it notoriously difficult to answer the question “how much storage space is this virtual machine actually using?” In addition, it makes it challenging for administrators and operations teams to understand how much storage is actually being consumed and understand the risk/efficiency trade-off to hosted applications.

Turbonomic understands the relationships and dependencies of thin provisioning in both virtualization and storage domains. The Turbonomic marketplace treats thin provisioned storage allocation as a commodity. As storage space is more aggressively oversubscribed via aggressive thin provisioning, its price increases in the marketplace, making those storage volumes less attractive to prospective virtual machines. This guides resource allocation decisions including virtual machine placement, as well as sizing and creation of volumes, aggregates and arrays to optimize storage space efficiency and risk.

TOP CHALLENGES IN STORAGE VIRTUALIZATION

(AND HOW TO OVERCOME THEM)

Deduplication and Compression

While thin provisioning acts to make storage space more scarce through over-provisioning, deduplication and compression technologies essentially work in the opposite fashion by attempting to make storage more plentiful through storage conservation. This further obfuscates storage space decisions since administrators and operators must account for all of these technologies together in order to understand the net available storage, which can be a time-consuming and complex process, particularly in a dynamically-changing virtual estate.

Turbonomic understands how much space is being saved by deduplication and compression by talking to the storage arrays and respective volumes. The storage space savings essentially work to make the price of storage cheaper in the Turbonomic marketplace, making the respective data stores and volumes more attractive to host-shopping virtual machines. So, even in cases where aggressive thin provisioning has oversubscribed an underlying aggregate, deduplication and compression could be working to cancel this out by saving storage space.

Snapshots

Array-based snapshot technology allows for low-impact user-recoverable backup of files for a given volume. If left unchecked, these files can consume significant storage over time, filling up a given volume and putting hosted virtual machines and applications at risk of outages. Snapshot storage consumption is, therefore, a key factor for operations teams to take into account when deploying and operating virtual machines to ensure that storage space demands can be met.

Turbonomic accounts for array-based snapshot space in its economic marketplace and resulting resource allocation decisions. Storage space used by snapshots will help to raise the market price of storage space for virtual machines shopping for data stores and volumes, helping to achieve equilibrium in virtual machine storage space demand and minimizing the risk of snapshot overruns.

TOP CHALLENGES IN STORAGE VIRTUALIZATION

(AND HOW TO OVERCOME THEM)

Policy Constraints

The Turbonomic marketplace naturally seeks to optimize workload storage performance while making efficient use of the underlying infrastructure. Turbonomic also allows policy constraints to be imposed. Some examples include:

- Constraining virtual machines/images to a given data store that is serving as an image/OS repository in a virtual environment (or VDI deployment) to maximize underlying deduplication of similar OS images on that volume.
- Turbonomic discovers the type (HDD, SSD, etc.) and speed of the underlying physical disks in the storage infrastructure. While Turbonomic will naturally optimize virtual machine placement on data stores and the underlying storage based on workload IO demand, high IO (e.g., database) workloads may be constrained to particular data stores and storage infrastructure when making deployment or operational decisions, if desired.

Controlling Heterogeneous Storage Environments

Many enterprises practice an incremental flash-adoption strategy, introducing all-flash arrays into their IT environments in phases or stages spanning several business cycles. Similarly, other organizations choose to invest only partially in flash, retaining their legacy HDD storage for several hardware refresh cycles.


In either case, heterogeneous storage environments pose a simple yet challenging question to any organization choosing to build one: Which workloads should be placed on flash, and when, to maximize the performance of those workloads and the ROI of my flash purchase?

Turbonomic dynamically places workloads on the right storage, based on real-time workload demand, enabling the most responsible migration and optimal utilization of a flash storage investment alongside legacy arrays.

Turbonomic Storage Integrations

Turbonomic takes actions on the virtual storage layer (data store) including Storage vMotion, data store resizing, and data store provisioning/decommissioning. It can also discover and act upon EMC VNX, EMC VMAX, EMC XtremIO, NetApp, HPE 3PAR, Pure Storage, Dell Compellent and Nutanix at the aggregate and volume level, fully aware of all constraints and data reduction policies in place.

CONCLUSION



Turbonomic's autonomic solution controls the virtualization and underlying storage infrastructure, leading to significant improvements in assuring the storage performance of virtual workloads.

This paper examined some of the most common storage problems that are prevalent in virtualization deployments. Turbonomic's autonomic solution controls the virtualization and underlying storage infrastructure, leading to significant improvements in assuring the storage performance of virtual workloads while driving infrastructure efficiency and operational cost savings.

Using Turbonomic, enterprises and service providers can:

- Reduce ongoing storage infrastructure costs by 20-30%.
- Significantly reduce operational costs by preventing complex storage problems and their impact on workloads, and, in turn, on end users.
- Ensure application performance and reduce risk.
- Enable IO-intensive applications to be virtualized reliably, bringing down the overall cost to deliver compute services to the business and customer.
- Seamlessly implement heterogeneous SSD/HDD storage environments by determining which workloads require SSD access and when.

TOP CHALLENGES IN STORAGE VIRTUALIZATION

(AND HOW TO OVERCOME THEM)

About Turbonomic

Turbonomic delivers an autonomic platform where virtual and cloud environments self-manage in real-time to assure application performance. Turbonomic's patented decision engine dynamically analyzes application demand and allocates shared resources to maintain a continuous state of application health.

Launched in 2010, Turbonomic is one of the fastest growing technology companies in the virtualization and cloud space. Turbonomic's autonomic platform is trusted by thousands of enterprises to accelerate their adoption of virtual, cloud, and container deployments for all mission critical applications.