



Hypervisor Economics

A Framework to Identify, Measure and Reduce the Cost of
Virtual Machines

By David Merrill and Michael Heffernan

November 2011

Table of Contents

Executive Summary and Introduction	3
A New IT Ecosystem	4
Growing Pains for Virtual Storage and Virtual Machines in the New Ecosystem	4
Consider the Operational Factors in Any Economic Assessment	5
An Economics Framework Approach to VM Cost Improvements	6
Identify Your VM Costs	7
Measure	8
Reduce	9
2011-2012 VM Cost Impact Areas	12
Conclusion	13

Executive Summary and Introduction

Global, national, local and business economic factors require IT departments to constantly evaluate and implement plans to reduce costs. As the size of IT continues to grow (in terms of storage capacity, servers and applications), business demands that IT costs hold flat, and that more must be done with less, which means with fewer resources. New applications and architectures are introducing new technologies that stretch both IT capability and IT budgets. Virtual machines (VMs) are one such technology that, since becoming more common in recent years, has been promoted from a lab-only capability into a staple for production computing. In this escalation of VM deployments (which we often call sprawl), a new set of views is required to understand and improve on VM architectures and VM costs.

If you are reading this paper, then you are probably not interested in the return on investment (ROI) of virtual machines; you made that financial and business decision years ago. At this point, the ROI is irrelevant, and the focus for continuous improvement has shifted to cost reduction, or unit cost reduction per VM. During the IT slowdown of 2007-2008, VM architects were relatively unaffected by capital expenditure (CAPEX) limitations and budget cuts; the VM revolution had just started, and seed money necessary to get started was small, if it existed at all. Today's economic times are different, and the VM segment of IT consumes a large percentage of total IT spending and operational budgets. VM planners and architects will be asked (in the next budget crunch) to do more with less and reduce the unit costs if VM quantities are to grow. Unit cost measurement and reduction are necessary approaches, as IT cannot always grow the VM budget or change VM demand. Measuring and taking steps to reduce the unit cost or total cost of ownership per VM (TCO-per-VM) is an essential business practice.

Hitachi Data Systems (HDS) brings 12 years of storage economics to the VM economics arena. Coupled with years of experience with mainframe costs and best practices, moving economic concepts into this hypervisor technology on x86 platforms is a natural fit. Thousands of proof points, plus measurement and mapping to solutions that reduce costs can be applied to help identify, measure and reduce the cost of VMs. The patterns and methods of storage economics are compatible with VM economics, so years of lessons and actionable tactics can and should be applied to controlling VM sprawl costs.

The authors understand operational and technical differences of the prevailing hypervisor products in the market today. This paper will not attempt to compare or contrast (economically or technically) these different systems, but we will outline a framework to reduce the unit and total cost, regardless of the vendor.

A New IT Ecosystem

The server and storage industry is a large ecosystem in which each technology development currently can have a compounded economic and technical effect on another technology or development. The fact that there has been a massive change in x86 server technology over the years has now provided other advances in technologies like server virtualization and storage virtualization. The combination of server and storage virtualization technologies can make a compounded economic impact to IT environments that has not been previously available.

Hypervisors have assisted in providing improvement from a single application on a single server where there have been major inefficiencies both with architecture and server utilization. Organizations today are adopting server virtualization in large capacities to address these inefficiencies and VMs are now considered the foundational platform for the data center. Server virtualization technology has also developed an ecosystem; however, this is quite unique as now there is interdependency between server and storage. An example of this interdependency is with VMware using SCSI commands to offload certain processes from the server to the storage. So what does this mean? New storage architectures can now leverage the true power and capability of both servers and storage to deliver a high-performance and robust infrastructure to manage data center transformations.

The popularity or sprawl of hypervisors has just begun, and, as a recent analyst¹ has observed:

- More than half of all workloads (51%) will be virtualized by the end of 2010 and over 2/3 (69%) by 2013, while only 12.8% of all physical servers were virtualized in 2009.
- VM densities continue to rise predictably; the current average was 6 VMs per physical server in 2009, which will increase to 8.5 in 2013.

The fact that x86 servers with their multithreaded cores and high-performance cache can inadvertently support higher densities of VMs is providing opportunities to consolidate the data center even further. However, the challenge now becomes the integration between the server and the storage with the hypervisor. There are 2 main ways in which a storage system can help ensure that virtualized server environments are able to provide promised benefits with minimal undue dedication of storage resources:

- Enabling flexible and efficient storage allocation and utilization
- Providing features that will help overcome performance degradation that occurs in environments that experience a high degree of random I/O traffic and server memory overcommitment

Growing Pains for Virtual Storage and Virtual Machines in the New Ecosystem

Hitachi Data Systems is observing that customers are experiencing many management issues in their VM environments because of the complex configurations that have evolved over the years.

¹ IDC, "Server Virtualization 2009"

There are:

- Many data stores on different RAID groups and storage systems
- Many different architecture configurations for server and storage
- VM clone and backup issues
- Single points of failure for hosts running large numbers of VMs
- Examples of VMware vMotion and Microsoft live migrations that cannot move the VMs easily around in the virtual server infrastructure
- No disaster recovery plans in place
- Islands of storage systems with no use of storage virtualization

IT architects and planners are starting to see powerful impact and results from converged storage and server virtualization:

- Storage virtualization is just being discovered for these new x86 adopters.
- VMs now need enterprise-class storage with integration.
- VM infrastructure needs 100% guarantees for no downtime (risk is now a factor).
- NFS is now no longer suitable for large, consolidated, heavy-I/O VM environments.
- API integrations between the hypervisor and storage system are offloading more processing into the storage system.

Consider the Operational Factors in Any Economic Assessment

One cannot start an economic analysis or cost-reduction plan without understanding the operational and business needs of the VM installation. Cost areas and developing cost-reduction plans have to be correlated and balanced with existing (and often unchangeable) infrastructure elements. Some of the basic factors that will determine or exclude certain VM cost-reduction activities include:

- Data types and workload types — certain configurations of memory, spindle count, network connection are often dictated by the I/O or workload.
- Availability, reliability and recoverability add costs to any IT infrastructure, and we often see VM costs rise as reliability, availability and serviceability (RAS) qualities have to be provided to these designs.
- Geographies — many cost elements (such as power and labor) can have varying cost structures.
- Data growth rates — typical organizations are seeing rapid growth of virtual machines, I/O and capacity.
- Vertical market segments have different cost sensitivities, and need to be customized to the local business need.
- Capacity — to accommodate growth rates, organizations need to increase capacity of VM, storage, applications and memory.
- Age of existing infrastructure — if the VM licenses or servers are very old, the options to limit or reduce costs may be curtailed.

- Sunk costs always burden new initiatives, even if the initiatives can reduce costs. The finance department will always want to sweat the assets in order to maximize the return on assets (ROA). Being forced to reduce costs while using older, existing assets can be problematic.

Before a TCO-per-VM or cost reduction strategy is started, IT and financial planners will need to understand and agree on the above factors that may impact options to measure and reduce costs of the VM.

An Economics Framework Approach to VM Cost Improvements

New server and storage ecosystems and the projections of VM sprawl will demand better control and a measurement system to provide improvement and justifications for these growth areas. Storage and hypervisor economics can be an objective approach to measure and act on improvement options of these new systems.

There is an old saying that "When performance is measured, performance improves. When performance is measured and reported back, the rate of improvement accelerates." In this case, we are interested in the improvement of unit costs (being reduced) for VM instances. A proven framework that includes measurements and reports is essential to deliver on cost-reduction actions.

Step 1 is to identify the costs associated with hypervisor TCO. There are as many as 24 different types of costs²; some are CAPEX costs but most are operational costs. Focusing on the purchase price alone may only account for 15% to 20% of the total cost of a VM, so focusing on the best price deals from your vendors will only get part of the way to a cost reduction strategy.

Step 2 requires a measure of the costs. Taking the costs that you isolate as candidates for your VM TCO will then require that parametric values be applied for the total cost calculation. Several cost metrics are possible and popular with hypervisors:

- Unit cost measurement of dollars per VM per year
- Performance cost measurement of dollars per I/O per VM

Other metrics can include the number of VM per square unit of floor space or per rack, etc. Defining the costs and measurement correctly will allow you to find the cost problem areas in your TCO, and that will lead to choosing the right correction action or investment to reduce the total cost.

Step 3 is defining the actions to reduce the costs. Identification and measurement of costs allows for prioritization of the costs, and choosing technology, operation changes or business processes that will reduce the costs. There are direct and correlated investments that will reduce certain costs of a VM. This paper will outline a few of these options, but it is more important to know your costs before a rapid set of investments are made to reduce the costs of a VM.

² Please refer to this blog entry for a short summary of the costs that make up a hypervisor TCO.
<http://blogs.hds.com/david/2011/06/hypervisors-welcome-to-storage-economics.htm>

Identify Your VM Costs

Hitachi Data Systems has extrapolated (from Storage Economics) 24 unique types of cost that can be considered for the total cost of a VM, and can be used in VM economic calculations. These are summarized below into 4 different families:

Purchase (acquisition costs, usually CAPEX in nature and occurring only once)

1. Hardware depreciation (or lease) expense
2. Software purchase and depreciation
3. Cost related to the procurement process
4. CIFS- or NFS-related infrastructure

Operational or Run-rate Costs (recurring costs that last the lifetime of the asset)

5. Hardware maintenance
6. Software maintenance
7. General administration or management, labor
8. Power consumption and cooling
9. Monitoring
10. Data center floor space
11. Cost of growth
12. Cost of performance
13. Data protection or backup infrastructure
14. Disaster protection infrastructure
15. Security, encryption

Migration, Movement-related expenses

16. Migration, remastering
17. Data or workload mobility effort and time
18. Storage area networking
19. Provisioning time

Risk-related Costs

20. Cost of scheduled outage
 21. Cost of unscheduled outage (machine)
 22. Cost of unscheduled outage (people or process)
 23. Cost of disaster risk, business resumption
 24. Reduction of hazardous waste
-

Of these 24 VM operating expenditure (OPEX) and CAPEX costs, there are perhaps 8 to 10 costs that are used most frequently for VM total cost models. Each IT department is different and may choose different combinations of costs that reflect their local business environment, operational demands and IT processing needs.

1. Server and storage hardware, including memory, network adapter cards and the resulting SAN or LAN infrastructure
2. Operating systems (hypervisors) and all other software installed on the VM; this cost may or may not include the application or database software costs
3. Additional licensing costs for hypervisor software advanced features (rapid cloning, de-duplication, etc.)
4. All hardware and software maintenance that is charged after the warranty period is over
5. Labor and administration costs associated with managing the VM and storage infrastructure
6. Power, cooling, data center floor space
7. Network (SAN and IP) costs for local and long distance data movement or migrations
8. Migration costs (time and effort) related to VM workload movement
9. Data protection or disaster recovery protection costs related to redundant server clusters, network connections and staff
10. Outage risks and costs, whether planned or unplanned

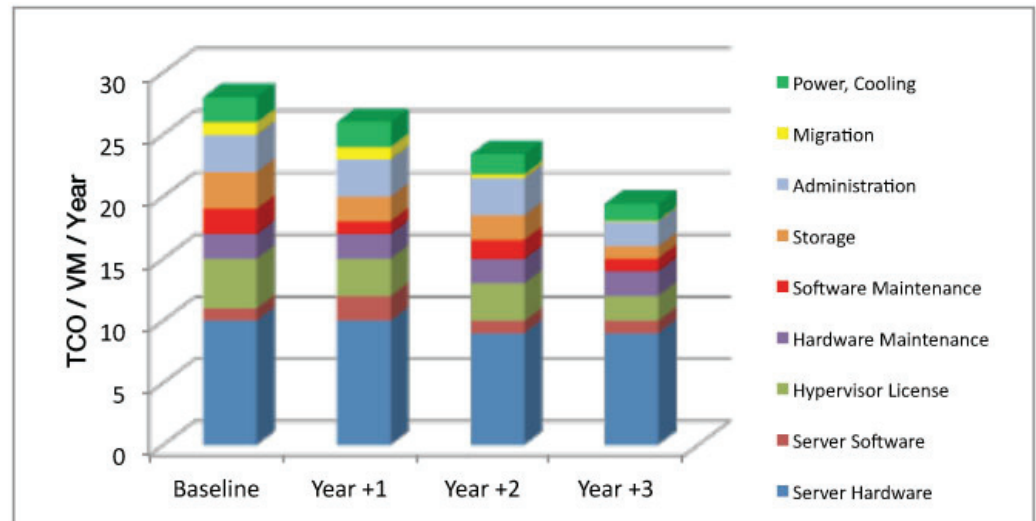
Each organization must prioritize the costs to be used in VM TCO measurements as well as differentiate costs between hard and soft or direct and indirect. Since all costs are not equal within the organization, coordination of the cost models must be determined in a cooperative approach with IT finance, procurement and operations.

Measure

It has been said that you cannot improve what you cannot measure. This is true for any VM cost reduction effort; a baseline measurement of the identified costs is essential to reduce costs. Using the cost categories (from the abovementioned 24 types), isolating and calculating each element provides a simple metric to determine TCO-per-VM. The measurements and metrics used to show continuous improvement (over time) need to be clear and unambiguous. As a metric, these measurements need to invoke action and planning in order to produce results for the IT department. Many individuals and groups will be required to cooperatively act and prioritize plans based on these measurements and metrics.

Over time, an annual TCO measurement of \$/VM can be valuable to show progress (see Figure 1) and to set priorities and expectations moving forward. Continuous improvement will require continuous measurements.

Figure 1. Representative VM TCO Measurement



Total costs per I/O or TCO/VM/year are the most relevant measurements for hypervisors. The "measure twice, cut once" rule is true for VM economics. Good data that is verifiable and believable is needed to justify strategic investments and directions to improve VM architectures. Measurement and re-measurement are essential to provide continuous improvements. Because each company will choose different costs when defining economic metrics, you may not be able to simply compare your VM TCO to other companies (although you will be tempted to do so). Rather, it is more valuable to measure your TCO costs each year to show management improvements in costs, as well as the actions or investments that produce those cost savings. Managers love metrics and dashboards, so start creating and reporting on VM economics with econometrics or quarterly reports.

Reduce

Cost-reduction tactics need to be tangible and actionable. Hope is not a strategy to reduce costs. A systematic approach to employ proven tactics to reduce the costs of a VM yields the best results. When roadmaps and IT strategies are aligned with actionable cost-reduction plans, reduction in OPEX and CAPEX can be more predictable.

There are several levers available to IT architects and planners to reduce the costs. Some levers can be purchased from a vendor; others require an internal review and change to operational and business processes. These levers are summarized below.

Technology Levers

Hypervisor Selection: This is one of the fundamental choices a customer must make for their VM environment. Note that "the bigger the ecosystem the more flexibility for choice."

- **Scalability.** Hypervisors technology has had a massive effect on cost. Memory management, compression, page sharing between Xen, VMware, kernel-based virtual machine (KVM) and

Microsoft Hyper-V are all different. For example light workloads with memory oversubscription can lead to high levels of consolidation.

- **Interoperability.** Ensuring that all components are completely certified reduces the risk for downtime and makes it easier to troubleshoot when issues arise. API programs also add a massive benefit, providing standardization for hypervisor to hardware integration.
- **Operational efficiencies.** These will reduce the need for skilled engineers (skills required to manage and maintain the environment), thereby reducing salary and training costs.
- **Ecosystem flexibility.** This is fundamental for the selection of the hypervisor and for 3rd-party backup products, disaster recovery and management software. If the hypervisor has built an API ecosystem for any partner to integrate, this then provides choice.
- **Hardware.** Storage features provide advanced integration that allows for higher levels of I/O and pooled resources. This includes storage technologies like thin storage, externalized virtualized storage and APIs like VMware's VAAI, etc.

Storage Architecture and Storage System Selection: This is one of the critical technology choices an organization must make for its VM environment. Storage capacity is vital to the operation of the VM, but it is also integral to overall performance and VM movement activities.

- If the external storage supports VAAI, then you can see a 300% to 400% improvement in provisioning VMs, and SCSI reserves can be reduced by 80%.
- If the storage supports thin provisioning, then you can recover deleted file space and reduce storage costs (capacity waste or reserve) by 40%.
- If you want to support more than 8 VM servers, then you may not be able to do it effectively with modular storage. These storage systems were designed to support workstations, not multiprocessor servers with high I/O.
- Modular dual controller storage is not high availability; when one controller is down you need to stop and fix it before you go back to work. The cost of outage increases with the number of VMs.
- If a storage system does not have a separate pool of global processors, then the system will not scale as you add more VMs.

Providing an integrated IT platform to support this evolution is the first step to address these items of maturity. Hitachi has built and put together an integrated IT platform for highly efficient operation in virtual server environment with VMware, Hyper-V and Citrix. In this stack, providing a Hitachi Adaptable Modular Storage (AMS) 2000 family system or Hitachi Virtual Storage Platform can definitely assist in overcoming performance degradation. Such degradation may occur in environments experiencing a high degree of "random I/O" traffic and server memory overcommitment within the hypervisor.

Hitachi Dynamic Provisioning and Hitachi Dynamic Tiering software, specifically, are suitable. They provide hardware thin provisioning technologies and wide striping capabilities for virtualized environments. This is a valuable combination since a virtualized environment will have multiple VMs placing many I/O requests in a single storage volume (at any given moment). This means the I/O requests are more likely to be random, reducing the usefulness of using storage cache to improve storage performance.

VM density is an important element for servers and storage Infrastructure in a unified virtual architecture. The density has a massive impact on the amount of infrastructure that must be purchased. Today there are many different hypervisors and the density is different on each of these because of their technology. This will determine the number of physical machines needed and lead to management software and higher licensing costs.

Organization Levers

- As IT systems converge, more staff and organization convergence is possible. Re-designing the server and storage team can provide more cross-trained and effective team members, resulting in fewer divisions of labor. It is no longer practical to have silos of teams to handle servers, storage, network and basic operating system management.
- The consolidation and cross-pollination of a single virtual team will improve staffing levels and total labor costs.

Business Levers

- Implement a chargeback or show-back system in order to assess the true cost per VM to the business unit.
- Use a server or storage services catalog to communicate VM offerings, configurations and pricing structures (total cost or cost of goods) to the requesting organization. Outline service level agreement (SLA) and operating level agreement (OLA) arrangements by each class of service. Outline incentives and penalties for noncompliance to company standards. Reducing the configuration variability will reduce long-term management and support costs.

Procurement Levers

Several considerations are fundamental to purchasing a hypervisor infrastructure. Some people call these a stack (server + storage + hypervisor), and these stack options can impact OPEX, CAPEX, enterprise agility (i.e. robustness of the solution) and ease of use or management.

- Simplicity of procurement
- Warranty and support (maintenance contracts and ongoing support)
- A single versus multiple points to support this stack (Who does the customer call?)
- Simplification of the stack's lifecycle (deploying, maintaining, refreshing)
- Elasticity of the stack (Are you forced to buy more than you need?)

With the cost areas above defined (and verified by your finance or management team), an alignment or mapping process can begin to prioritize the technical, operational or business levers that can be applied to produce the best impact to VM cost. Not all investments have the same impact to each type of cost; Table 1 demonstrates the relative cost impact for each possible combination. The priority of cost factors (for example power and cooling) will have a direct and specific impact as to what strategy to enact first in order to provide the fastest positive impact.

TABLE 1. MAP VM OPTIONS TO COSTS

Cost Types	Virtual Machine Cost-reducing Strategies and the <i>Relative Impact</i> to VM Unit Costs						
	Hypervisor Type	Thin Volumes	VM or Storage Integration	Enterprise Storage	Advanced Provisioning	VM Density	Common Catalog
Hardware Purchase	High	High	High	High	Med	High	Med
Software License	High		Med		Med		Med
Hardware Maintenance		High			Med	High	Med
Software Maintenance			High		Med		Med
Administrative Labor	Med		High	Med			High
Power, Cooling		High		Low			
Floor Space		High	Low			Med	
Data, Disaster Recovery, Protection	Med		Med	Med	Low		Low
Migration	High		Med	Med	High		
Outage	Medium		High				Low

The time of cost reduction is a secondary view that needs to be factored into the cost-reducing roadmap. It will be obvious to choose first the investment actions that are easiest to implement and have the biggest cost impact to VM TCO.

2011-2012 VM Cost Impact Areas

This paper does not provide the quantitative details of the above investments to the cost categories; however, follow-on work and modeling services from Hitachi Data Systems will provide these types of details to assist in prioritization and quantification of the investments. There are observed activities that have been shown to directly impact VM costs as these technologies start the sprawl from testing and development environments into the production environment. There are 3 technology and 3 operational recommendations that can be mentioned here:

Technology Recommendations

1. Tune storage systems to provide the right level of I/O performance. This can be a function of cache allocation, thin provisioning, hot spot prediction and storage-system-level tuning and monitoring.
2. Use converged systems that are pre-built (cabling, hardware and software) and can support a commodity offering (or standardization), with integration done by the vendor.
3. Utilize data de-duplication and/or capacity optimization.

Operational Recommendations

1. Reduce the number of backups necessary; current VM backup schemes are often overengineered to protect the data through backups.
2. Employ automated provisioning methods to reduce the administration time and effort.
3. Develop a standard VM offering and configuration (or catalog) to reduce the variability of systems and simplify deployments.

Conclusion

Hoping that VM costs will come down is a dangerous strategy to adopt. Structured planning and identification and measurements of VM costs are essential. The current rates of VM growth or sprawl are unsustainable, as are the current cost rates, so specific action and architecture considerations will be required to effectively reduce the unit of a virtual machine over time.

This cost-reduction process will take time, and some investment areas will produce faster results. The determination of activities and cost areas to attack will depend on the current state of the VM environment, capacities, application and data types, and production workload.

New metrics are essential to show progress in cost reduction and to secure funding for new tactical improvements that will add to unit cost improvements. The metrics will support the actionable plans needed to deliver continuous cost improvements.

In summary, Hitachi is focused on building solid servers, storage and integration for all types of hypervisor architectures. Ongoing development and research into enhancing storage integration is fundamental to reduce VM costs. As VMs become denser, so, too, does the server and storage infrastructure. The dependency between all these components is critical for the future of data center transformation and data center cost efficiencies.

 **Hitachi Data Systems**

Corporate Headquarters

750 Central Expressway
Santa Clara, California 95050-2627 USA
www.HDS.com

Regional Contact Information

Americas: +1 408 970 1000 or info@hds.com
Europe, Middle East and Africa: +44 (0) 1753 618000 or info.emea@hds.com
Asia Pacific: +852 3189 7900 or hds.marketing.apac@hds.com

Hitachi is a registered trademark of Hitachi, Ltd., in the United States and other countries. Hitachi Data Systems is a registered trademark and service mark of Hitachi, Ltd., in the United States and other countries.

All other trademarks, service marks and company names in this document or website are properties of their respective owners.

Notice: This document is for informational purposes only, and does not set forth any warranty, expressed or implied, concerning any equipment or service offered or to be offered by Hitachi Data Systems Corporation.

© Hitachi Data Systems Corporation 2011. All Rights Reserved. WP-412-B DG November 2011