



January 28, 2009

Why Isn't Server Virtualization Saving Us More?

A Few Small Changes May Dramatically Increase Your Efficiency

by **Galen Schreck**

with Stephanie Balaouras and Alex Crumb

EXECUTIVE SUMMARY

Companies have rapidly adopted server virtualization over the past few years, but there are big differences in their use of the technology. Most companies are able to reduce server hardware spending, but many don't realize an ongoing reduction in management costs. Furthermore, not everyone sees the same level of savings on hardware. If you're not saving as much money as you'd hoped from your virtualization and consolidation efforts, you can learn from the best practices of the companies that are realizing huge savings. These companies run at least five times more virtual machines per physical host than other companies and use streamlined IT processes.

YOU'RE VIRTUALIZING, BUT SAVINGS ARE LESS THAN YOU HAD HOPED

Consolidation of x86 servers through virtualization has reduced capital expenditures, lowered electrical costs, and freed up expensive data center space. These are big savings, but many of them are one-time windfalls that fail to alter your ongoing cost of operations. You can do much better — you may be leaving money on the table with your virtualization projects because you're:

- **Not aggressive enough with the ratio of virtual machines to physical host.** Some virtualization implementations fail to produce the expected savings because the cost per virtual machine (VM) is too high. This usually results from poor VM density and high hardware costs in shops that are too conservative with their deployments. For example, many companies break even with physical hardware after three VMs, but they only put a total of five onto a server capable of running 15 VMs. They could have provisioned 10 more VMs for no incremental cost. Instead, it will take two more virtualization hosts to accommodate the overflow.
- **Networked to an expensive Fibre Channel storage area network (FC SAN).** If you want the flexibility to move VMs between different physical hosts for high availability or for load balancing, your physical hosts must be networked to shared storage. In the past, VMware only supported FC SANs. So while you were saving money consolidating physical servers you were also spending thousands of dollars per FC port to network your servers to storage.
- **Slow to update legacy IT processes.** Once a server is virtualized, you have the option of managing it pretty much like a regular server or taking advantage of your new capabilities to streamline management tasks. Many companies decide not to rock the boat with overhauling their processes, and not surprisingly their ongoing management costs remain the same — with some one-time savings associated with the initial provisioning.

Increase Your Virtual Machine Consolidation Ratio

Satisfied with the knowledge that they've eliminated a large percentage of their physical servers, few IT organizations analyze the cost of their virtualized environments. Sure, you're saving money — but you're probably still over-provisioning the physical server environment. To bring down the cost per VM you need to increase virtualization density by:

- **Turning on dynamic resource scheduling.** You can bring down the cost per VM by running the maximum possible number of VMs per physical host. Many virtual infrastructure administrators prefer to allocate around one processor core per VM. This is excessive and limits you to just six to eight VMs per physical host. But you need to prevent applications from hitting the wall in the event of a spike in demand. Both VMware and Citrix offer workload management capabilities that optimize VM placement on a continuous basis. Similarly, tools like HP Insight Dynamics allow you to mathematically optimize the placement of VMs based on historical performance.
- **Piloting VMware's memory overcommit.** If you're a VMware shop, you can take additional steps to increase the number of VMs per physical host and possibly buy cheaper servers with fewer memory sockets. On a typical virtualization host, a large number of your memory pages will be the same, since each VM might be running a variant of the same Windows or Linux operating system or application components. Memory overcommit is a VMware feature that essentially shares these common memory pages among VMs, rather than storing duplicate copies — allowing you to free up memory for more VMs at the expense of some added CPU overhead. But your actual mileage may vary depending on how much duplication exists — we suggest piloting this feature and measuring the benefits and any performance impact in your own environment.

Network Your VMs To Shared Storage With Ethernet

One of the top complaints we hear from our clients trying to save money with server virtualization is the cost of the Fibre Channel SANs. This seems to be a holdover from earlier versions of VMware ESX that required a VM to be stored on Fibre Channel SANs. Although you're saving money by consolidating server hardware, your FC SAN is now eating a disproportionate part of your hardware budget. Excluding the storage array, the typical physical host requires an outlay of almost \$4,000 for FC networking gear. If you're ready for a storage upgrade or you're about to embark on a new virtualization and consolidation project, you should consider deploying an iSCSI SAN or network-attached storage (NAS) system. Ethernet networking gear such as iSCSI host bus adapters (HBAs) or standard network interface cards (NICs) are significantly less expensive than FC.¹ And you're not sacrificing functionality; VMware's advanced features like VMotion and VMware HA are fully supported with iSCSI and NAS. NAS storage systems are supported via the NFS file system protocol. iSCSI SANs and NAS systems are available from all the major storage vendors.

Update Your IT Management Processes

Lots of IT organizations are familiar with the hardware savings driven by virtualization and consolidation. But they often find that ongoing costs have remained pretty much the same because they're still performing all the same duties with all the same tools. Sure, you might get some incremental savings, but you must truly revamp system management if you want to see a significant cost reduction. What management tasks will you need to examine?

- **Backup of virtual environments.** If you want to achieve a much higher ratio of VMs-to-physical host, you can't back up individual VMs like physical servers. First, the cost would be too great because traditional backup vendors have still not adjusted their per-agent pricing models for the virtual world. Second, as you increase the ratio of VMs to physical host, it becomes almost impossible to complete backups during maintenance windows because the physical host simply doesn't have enough bandwidth. In VMware environments, consider backup applications that have native GUI integration with VMware Consolidated Backup (VCB), or look at alternatives such as storage-based snapshots and backup software designed for virtual environments; two examples are Vizioncore vRanger Pro and Veeam Backup.²
- **Configuration management and software distribution.** Managing operating system configuration, patching, and auditing absorbs a lot of administrators' time. You can simplify management tasks by standardizing on fewer variants of your base operating system and using virtual machine templates to their full advantage. Rather than building configurations from scratch, you concatenate multiple templates containing your base operating system, management agents, and applications to form a new virtual machine volume. In the future, these clones will be linked to their original templates, meaning that changes to the original template — like patches — will propagate to all derivative virtual machines. This technology is already in use for virtual desktops and greatly reduces the amount of storage required, since the desktop version uses a single shared copy of the bits.
- **Self-service provisioning.** How long do your users or developers wait for a server after putting in a request — weeks? In addition to speeding time-to-market, you can offload many trivial tasks from your administrative staff by providing a self-service portal for virtual resource requests. Surgient specializes in self-service provisioning tools, while VMware offers virtual lab management tools that allow developers to instantly provision test environments.
- **Lower-cost disaster recovery.** In the past, you had few options for disaster recovery: slow tape-based recovery or expensive clustering tools. Virtual servers allow you to offer mid-range disaster recovery and HA services to take the place of high-end capabilities that aren't really necessary. By using virtual machines for recovery infrastructure, you can tap into test and development hardware for emergency capacity and move some apps that don't quite warrant array-based replication to same-day recovery on virtual hardware.³

RECOMMENDATIONS

COMPANIES CAN INCREASE VMs PER SERVER WITH MINIMAL EFFORT AND RISK

To lower the cost of your virtual environment, you must scale up the number of VMs per physical host and decrease infrastructure costs (e.g., storage) to a minimum. For many companies, this is more a matter of risk aversion than a technical barrier. We believe that many IT shops running hypervisors like VMware's could easily move from an average of five VMs per server to 10 VMs per server with little risk and no additional technology. By enabling live migration and automated workload management, densities of 20-plus VMs per server should be well within reach. Best-practice users of x86 server virtualization often get around 25 VMs on a typical rack-optimized server such as the HP DL585.

Don't focus on achieving a specific number of VMs per physical host — there are too many variables, and your applications may be quite different from those of other organizations. Even if you're getting 20 VMs on some servers, you will probably have systems running just a few very large VMs. The key is to keep over-provisioning to a minimum by analyzing the amount of capacity required to keep your applications performing normally. Platforms like VMware allow you to reserve a specific quantity of system resources for an application — make sure these reservations are tuned to an appropriate size, rather than taking the easy route and massively over-provisioning them.

ENDNOTES

- ¹ Although a great deal of energy is expended discussing the performance differences of IP-based and Fibre Channel (FC)-based storage area networks (SAN), it is generally accepted that it is cheaper to deploy an iSCSI SAN. This report dissects the components of an IP-based SAN and compares the up-front and ongoing costs versus a FC SAN. The cost advantage of iSCSI is so significant in terms of server side and switch network costs that you owe it to your budget to determine which applications would be a good fit for this technology. Some implementations of IP-based storage networking can struggle with high intensity application workloads; others have matured significantly and are more capable of handling enterprise-class applications. Given the hardware acquisition cost benefits as well as synergies in network staff and infrastructure, significant benefits can be gained from the use of iSCSI-based SAN architectures. See the February 7, 2008, "Cost Comparison Of iSCSI Versus Fibre Channel SAN Components" report.
- ² Forrester regularly receives calls about how best to back up large VMware environments. What have we found? IT managers should default to VMware's Consolidated Backup (VCB) as the best option today, but don't hesitate to explore specialized tools and storage-specific functionality — both of which have come a long way to support VM-specific intelligence. To detail this further, here are our answers to the most common questions. See the July 14, 2008, "Q&A: Top Questions About VMware Backup" report.
- ³ According to a recent Forrester study, 49% of enterprises surveyed that are implementing or interested in x86 server virtualization indicate that improving disaster recovery/business continuity continues to be

a very important motivation for adoption. Server virtualization facilitates a rapid — or even automatic — restart of applications after an IT failure, and when used in conjunction with data replication between data centers, it can restart applications at a recovery site following a primary site failure. See the October 24, 2007, “X86 Server Virtualization For High Availability And Disaster Recovery” report.