



Global Knowledge®

Expert Reference Series of White Papers

Five Secrets for Successfully Virtualizing a Data Center

Five Secrets for Successfully Virtualizing a Data Center

John Hales, Global Knowledge VMware instructor,
A+, Network+, MCSE, MCDBA, MCT, VCP, VCAP, VCI, EMCSA

Introduction

This white paper is designed to provide some “secrets,” tips, and tricks for virtualizing your datacenter. We want to introduce some best practices for virtualization, while not being too biased towards one virtualization vendor or another. We’ll use some common examples of products and tools that work with VMware’s vSphere and Microsoft’s Hyper-V, but with an eye toward virtualization in general, and not the specifics of any of the capable platforms that could be used. We will assume, however, that bare metal hypervisors, in other words virtualization platforms where the hyper visor is the OS, will be used as opposed to running a hypervisor on top of an existing general-purpose operating system (which is great in a lab, but terrible for data center projects).

In this white paper, we’ll look at five basics that should be considered in any data center virtualization project: planning, balancing the various hardware components, sizing the storage properly, managing capacity, and automation.

Plan

The dreaded word in IT seems to be *Plan*. In years of training and consulting, I’ve found that few people in IT like it, fewer do it, and even fewer like to do it. That causes big headaches, as without planning, anything will do. Simply buy whatever you have budget for, plug it all together, and go to work. The problem with this is that without proper planning and requirements gathering, there is a great chance that you will buy less than is really needed to get the job done, so performance will be poor, requiring more money to go in and fix it later. On the other hand, on the off chance that you buy too much, you end up lowering the ROI and increasing the TCO, just the opposite of the goals of virtualization.

So what do you need to know to properly plan for a datacenter virtualization project? First and foremost, you need to know what you are consuming today. For example, you’ll need answers to questions like these (both on average as well at peak periods):

- How many GB of RAM is in use in each server? By this we mean, actually used, not simply installed.
- How many MHz (or GHz) are in use in each server? Again, it’s not what is installed, but what is in use.
- How many GBs (or TBs) of storage are used in each server?
- How many MB/s or I/O operations per second (IOPS) are being consumed today?
- What kind of storage is in use today? Is it Fibre Channel, Fibre Channel over Ethernet (FCoE), iSCSI, or NAS-based, or do you use local storage? If you use local storage today, will you use shared storage in the future? What local storage will be required, if any?

- How much bandwidth is required on the network?
- What kind of security is required for the data being stored?

We could go on and on, but you get the idea. Another important thing here is to figure out what is “normal” or average and when peak periods are, with a goal of not having VMs all peaking at the same time on the same server. Maybe they can be scheduled differently so the peaks are at different times (for example, when backups are run or antivirus scans are executed), or maybe they can be placed on different physical servers to spread the load and reduce the peak demand on any one server. If this is not possible, you’ll need to buy more capacity to handle the larger peaks that the combined load will cause.

This particular section is closely related to the next three – planning so that the appropriate hardware can be purchased that will work well together and will be sized properly, sizing storage to handle the load, and once the system is in place and operational, planning for the future based on real-world conditions as they evolve. In other words, on-going planning as opposed to the upfront planning discussed in this section. We’ll discuss each in the next few sections in more detail.

Balance Hardware Components

It is very important to balance all of the hardware components properly; in other words the goal is to keep all of the equipment roughly evenly loaded to minimize cost and maximize utilization (within reason).

For example, you don’t want a 10 Gb Ethernet network to handle all of your iSCSI needs paired with a low-end iSCSI device that can’t connect at 10 Gb, or an iSCSI array that can’t push data effectively at 10 Gb. In this case, it would be better to save some money on the networking equipment and put it into better storage.

Likewise, on the CPU and memory side, the goal is to be balanced as well; in other words enough RAM to run all of the applications and keep the CPUs fairly busy (averaging 60% to 75% is fairly normal). If there is a lot more RAM that the CPUs can effectively use (for example, CPU-intensive tasks that require modest amounts of memory), the extra RAM is wasted. On the other hand, if you have many machines that are not CPU-intensive and they all run on the same host, you may exhaust the RAM available, causing lots of swapping to disk, drastically reducing performance.

The challenges of trying to balance everything well, while at the same time leaving some resources available to handle outages (both planned and unplanned) and future growth can be somewhat daunting. This is why planning is so important.

The virtualization vendors have tools to help you determine how to best handle these challenges. For example, VMware has a tool called Capacity Planner (available from VMware partners only) that will gather all of the performance metrics over time (they recommend a minimum of 30 days) and then recommend a plan to put the components together. It also offers some what-if scenarios, such as more or faster CPUs, or more or less RAM in each server, recalculating the load, and creating a new plan in terms of resources required. Microsoft has a similar tool in their Microsoft Assessment and Planning (MAP) toolkit for Hyper-V. Many third-party companies, such

as server and storage vendors, have tools to analyze an environment and suggest hardware from that vendor that will handle the load. There are also vendor-agnostic tools from other third-parties that can help analyze the environment and suggest what would work best.

In this section, we have discussed tools that are available to help out in the planning phase; in other words, the initial sizing and balancing. The secret we are discussing here, though, is not a one-time event, but rather an ongoing (or at least periodic) process of reevaluating what is in use, how it is running currently, and what the projected needs and demands are a few months (at least) out so that equipment can be ordered, installed, and ready when it is needed. This is discussed in more detail in the Manage Capacity section, which we'll get to in just a little bit.

First, however, we need to think about the most commonly neglected area in sizing our virtual environment to run well, storage.

Size Storage Properly

This secret is really a part of the last one, but most people don't think of storage holistically and thus don't size it properly. This has probably killed more virtualization projects than any other area.

When sizing storage, most people simply count the TB of space required, but the requisite number of drives to provide that space (with a few extra for parity, hot spares, etc.) and consider the project complete. That is not adequate, however, as different disks have widely varying performance capabilities. A decade ago, it was all about capacity – the drives we had were small (by today's standards) – 9 GB or 18 GB were common. To get to 1 TB of space, one hundred or more drives would often be required. That provided a lot of performance relative to capacity. Today, 1 and 2 TB drives are common, but obviously replacing an entire 100 drive SAN with a single disk is not going to provide the same performance, even though the capacity is the same (or even better).

To help better understand this topic, a brief look at average performance values for different kinds of disks is helpful; this is shown in the table below.

Tier	Disk Type	Speed	IOPS
3	SATA	5,400	60 – 80
		7,200	75 – 100
2	Fibre Channel / SAS / SCSI	10,000	125 – 140
		15,000	160 – 200
1	SSD / EFD / Flash	N/A	4,000 – 10,000

Table 1 Drive types, speeds, and performance summary.

The values in Table 1 are approximate and vary depending on I/O size (larger I/Os will result in fewer IOPS), random vs. sequential access patterns (sequential is faster as there is no seek time involved for spinning drives;

this is not a factor for Solid State Disk (SSD) drives, and reads vs. writes (reads are usually faster, especially when the overhead of mirroring or parity are added for writes; note that SSD drives by their nature are typically much faster at reading than writing).

SSDs, Enterprise Flash Drives (EFDs, these are drives that are usually faster with more redundancies and other optimizations designed for use in servers not desktops and laptops), and flash drives all mean roughly the same thing. They are all memory-based, and not based on spinning platters.

Note in Table 1 that some vendors call things Tiers 1 – 3, while others prefer 0 – 2. The disk types in each tier are fairly universally agreed upon.

What this means in today's environment is that faster disks must be deployed and utilized in an environment. There are several ways this can be done; some vendors initially write to a fast tier and then migrate the data to slower tiers if it is not accessed frequently. Others use SSD drives as big caches that can handle the incoming I/O with the goal of pulling the most often accessed data from SSD drives instead of spinning disks. Others use various tiering techniques to move data over time. The point is that most vendors today offer various mechanisms to optimize the speed of storage and that they need to be carefully considered and implemented in most environments to get adequate performance from the consolidated environment.

It almost goes without saying, but to be clear, in all but the smallest of environments, shared storage of one type or another will be required to take advantage of all of the capabilities that modern virtualization platforms provide. Some vendors even offer simple, low-end solutions that transform local storage into shared storage so that these benefits can be realized; see for example VMware's VSA (vSphere Storage Appliance), which will take local storage on two or three servers and replicate it so there are two copies of the data (one each on two servers) to provide redundancy.

Another topic often not well thought out is the RAID level that will be used; many administrators choose RAID 5 because it provides a basic level of protection for the lowest price, but from a performance perspective, it is usually the slowest (or next to the slowest if the storage vendor supports RAID 6, which is the slowest) option. RAID 0 is the fastest option, but cannot tolerate any drive failures. Thus if performance and availability in the event of a drive failure are both important (and let's face it, that is the great majority of the time), RAID 10 (or 0+1, depending on the vendor's offerings) provides the best balance between the options.

Entire training courses exist from major vendors (such as VMware, NetApp, and EMC) on optimizing storage in a vSphere environment. We have touched on just a few of the major pieces relative to storage that must be considered when preparing for and running a virtualization solution; now let's turn our attention from planning for one to running one day-to-day.

Manage Capacity

Capacity management is the ongoing portion of a data center virtualization project, the part that involves understanding utilization over time and adapting to changing conditions as the environment grows, shrinks, new servers are virtualized, etc. Our fourth secret is the longest part of the process as it will continue in perpetuity as

you grow and upgrade your environment. In this phase you will want to look at the utilization of the four core components of any virtualization strategy, namely CPU, memory, network, and disk. You will want to make sure they remain balanced (as previously described) and rebalance as needed, for example, by moving VMs between servers, adding RAM, or by upgrading the network.

To handle the day-to-day variations in load, the automation tools (described in the last section) work great and should be employed. The discussion here, however, is not in the short term, but in the medium-term – months and quarters out. The idea is to anticipate when additional hardware will be needed so that it can be brought in and configured before it is needed to keep the environment running smoothly.

The question is, how do you know when that will be needed? There are a lot of good simulation tools from the virtualization vendors themselves; products in this group include:

- VMware's vCenter Operations Manager [vCops aka Ops Manager] or CapacityIQ.
- Microsoft's System Center Virtual Machine Manager [SCVMM].

In addition, an entire industry of third-party companies have brought their own insights and algorithms to play to help predict and plan for the future. Vendors in this category include:

- Veeam's One Reporting
- ManageIQ's Enterprise Virtualization Management (EVM)
- VKernel's vOPS family of products
- Quest's vFoglight

Note that the above list is not anywhere close to all inclusive, nor are any of the products recommended, but just to give an idea of some of the tools on the market. Some of the products only work with vSphere, some only Hyper-V, and some work with multiple vendors.

The point of this secret is one that is often lost on administrators – even those who planned up front – is that careful, ongoing analysis is required to keep the environment running smoothly for the long haul. It's not as simple as buying new physical servers every three years as was often done a decade ago. The server may not be the bottleneck – it could be the network or storage, and the solution could be as simple as installing a quad port NIC or upgrading to 10 Gb Ethernet.

While this ongoing analysis is important, it does not need to be time-consuming; checking once a week or once a month for anticipated capacity issues weeks or months away is often good enough, assuming there are no major changes in the environment in the interim.

Automate

The last secret we'll talk about is the power and benefit of using automation. Automation comes in many forms, from command lines for scripting (many of which are based on Microsoft's PowerShell framework) to manage-

ment platforms that do much of the management and load balancing between devices automatically. A brief history of automation is shown in the Figure 1.

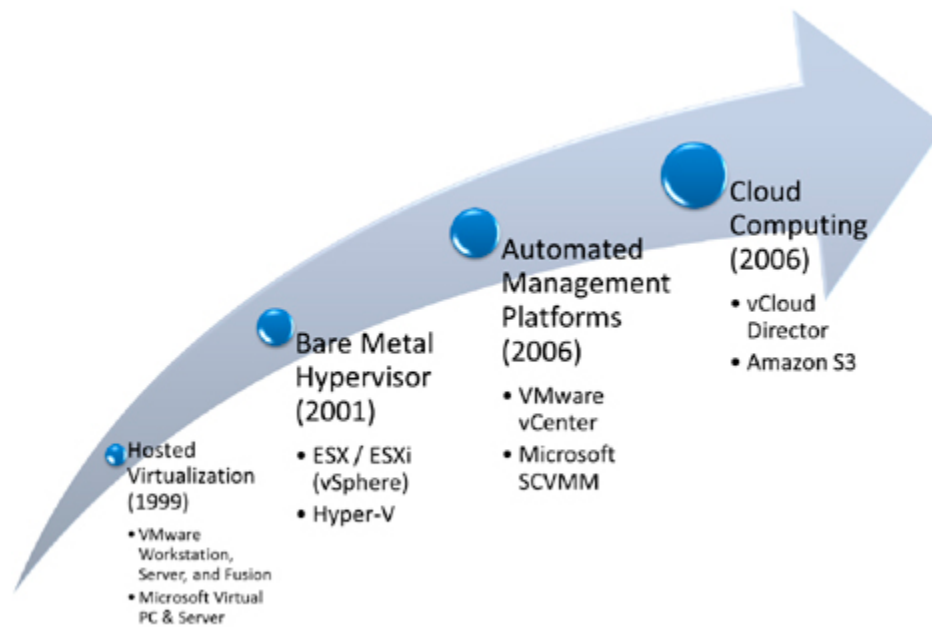


Figure 1 A brief history of virtualization.

Note that virtualization has gone from the early days of hosted and bare metal hypervisors where consolidation was achieved, but the administrator still had to determine which VM ran on which physical computer, to the automated platforms first introduced with vCenter, to today's self-service almost fully automated cloud platforms (at least from a user perspective; administrators are still required to manage the cloud platform).

A recent study of VMware customers and partners found that 92% used vMotion (which allows VMs to be moved from one physical computer to another without any downtime to the VM), 87% used the High Availability (HA) feature (which automatically restarts VMs after either a VM or physical server crash), and 68% used Storage vMotion (which allows a VM to be relocated to a different storage location with no downtime to the VM).

A new feature in vSphere 5 is Storage DRS that can automatically migrate VMs from one storage location to another based on capacity and latency of the underlying datastores. Given these great automatic tools, why aren't the values closer to 100%? This is the kind of automation that can easily be leveraged to help solve storage, CPU, and memory utilization issues. In talking with many students and customers, many don't because they think they are better. To be blunt – you are not smarter than an optimized algorithm and definitely are not as vigilant in looking for these optimizations. The system will analyze utilization 24x7x365 and make appropriate changes; you the administrator don't have the time, or even the ability, to do that kind of monitoring and management. So don't; leverage software that will do it better and focus on the things you need to do in the rest of the environment, such as capacity planning, helping users, and planning for upgrades.

In addition to using these automated tools, leverage third-party applications, scripts, and other tools to make your life easier. You have enough to do, so let the system help with what it does best.

Summary

So there you have it, five strategies, five secrets, to virtualizing your data center successfully. Plan carefully, both up front and on an ongoing basis, balance the components to minimize your TCO and maximize your ROI, and be especially careful to setup storage properly.

Learn More

Learn more about how you can improve productivity, enhance efficiency, and sharpen your competitive edge through training.

[VMware vSphere: Fast Track \[V5.0\]](#)

[VMware vSphere: Install, Configure, Manage \[V5.0\]](#)

[VMware vSphere: What's New \[V5.0\]](#)

[Citrix Desktop Virtualization Fast Track](#)

[Citrix XenApp 6.5 Advanced Administration](#)

[Citrix XenApp 6.5 Basic Administration](#)

[Implementing and Managing Microsoft Server Virtualization](#)

[Microsoft Cloud Enterprise Solutions](#)

Visit www.globalknowledge.com or call **1-800-COURSES (1-800-268-7737)** to speak with a Global Knowledge training advisor.

About the Author

John Hales, VCP, VCAP, VCI, is a VMware instructor at Global Knowledge, teaching all of the vSphere and vCloud Director classes that Global Knowledge offers, including the new View classes. John is also the author of many books, from involved technical books from Sybex to exam preparation books, to many quick reference guides from BarCharts, in addition to custom courseware for individual customers. His latest book on vSphere is entitled *Administering vSphere 5: Planning, Implementing and Troubleshooting*. John has various certifications, including the VMware VCP, VCAP-DCA, and VCI, the Microsoft MCSE, MCDBA, MOUS, and MCT, the EMC EMCSA (Storage Administrator for EMC Clariion SANs), and the CompTIA A+, Network+, and CTT+. John lives with his wife and children in Sunrise, Florida.