



Global Knowledge®

Expert Reference Series of White Papers

Five Secrets for Successfully Virtualizing Desktops

Five Secrets for Successfully Virtualizing Desktops

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

Introduction

This is the second whitepaper in a series on “secrets” surrounding virtualization. In the first white paper, *Five Secrets for Successfully Virtualizing a Data Center*, we discussed general tips to successfully virtualize a data center. While not specifically discussed, the primary focus of the document was surrounding the virtualization of servers. In this white paper, we’ll build on the previous paper with specific guidance for virtualizing desktops (this is sometimes referred to as VDI – Virtual Desktop Infrastructure). While this white paper does not require an understanding of the first one, you will see similarities and amplifications of content from that one.

We want to introduce some best practices for desktop 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 View, which runs on top of vSphere as a common method of virtualization, but other vendors, including Microsoft and Citrix, have great solutions in this area as well. This paper will discuss desktop virtualization in general, and not the specifics of one platform or another (or any of the other capable platforms that could be used).

This white paper focuses on desktop virtualization implementation; if you are interested in the ROI for desktop virtualization, Global Knowledge has one on desktop virtualization that I have authored, *10 Ways to Save with Desktop Virtualization*.

We’ll look at five basics that should be considered in any desktop virtualization project:

- Planning
- Balancing the various hardware components
- Sizing the storage properly
- Minimizing the number of base images deployed
- Accepting a wide variety of end user devices

The first three are named the same as the data center virtualization paper previously mentioned, though they will be discussed here in relation to desktops, while the last two are unique to this white paper.

Plan

The dreaded word in IT seems to be Plan. Over the years of training and consulting, I have 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. This goes for the storage, servers, network, and end-user connectivity devices. While the solution may work, it is not likely to work well.

So what do you need to know to properly plan for a desktop virtualization project? First and foremost, you need to know what you are consuming today on your desktops. 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 desktop? By this we mean, actually used, not simply installed. In most scenarios, it is usually less than 2 GB, except for some power users.
- How many MHz (or GHz) are in use in each desktop? Again, it's not what is installed, but what is in use. Patterns of CPU utilization are particularly important to look for and understand to be sure that the environment works well at those peaks. Patterns to look for include:
 - Certain times of days that CPU utilization peaks, such as when people first get to work, just before leaving, and just before and right after lunch.
 - Times in the environment when the entire environment is busy. For example, virus scans or updates may be scheduled at certain times that affect everyone in the environment. These scheduled tasks can very quickly kill a VDI deployment. If the tasks must be executed, spread them out so they don't happen all at the same time. In other words, doing a virus scan at noon on every desktop in the company will drastically affect the environment; it would be better to schedule them throughout the day. And while on the topic, having them all run at midnight isn't much better.
 - Scheduled tasks at the OS level (especially the default ones) should be carefully analyzed and removed if not required.
- How many GBs of storage are used in each desktop? How much of that is for applications and the OS vs. user data?
- Is user data stored on each desktop or stored on central file servers?
- How many MB/s or I/O operations per second (IOPS) are being consumed today on each desktop? Many VDI projects assume 5 IOPS as average with heavy users at 15 IOPS. That is probably not accurate in your environment; it may be much lower or higher, so measure it so sound decisions can be made.
- What kind of storage is in use today on the server side? 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? Will the desktop virtualization project use the same storage (on different disks), or will new storage be required for the VDI project?
- How much bandwidth is required on the network by each desktop?
- What kind of security is required for the data being stored? It probably isn't a lot for the OS, but understanding security for the stored data may be a much bigger priority so appropriate security mechanisms can be planned for and deployed.

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 the desktop 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.

One of the primary drivers for this phase is to gather the knowledge necessary to figure out how many desktop VMs can be placed on each server – the more that can be placed on a single server, the greater the consolidation ratio and the lower the TCO.

This particular section is closely related to the next two – planning so that the appropriate hardware can be purchased that will work well together and will be sized properly, and sizing storage to handle the load. We'll discuss each in the next couple of sections in more detail.

Another primary objective of this phase is to determine how people access their data today, as well as how they would like to access it if they had a choice (and it made sense from a business perspective). These options will be discussed in more detail in the last of the five secrets discussed in this white paper, but it is obviously important and part of the planning process.

Balance Components

Next, 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 low-end storage solution or slow disks that will cause the virtual desktops to perform slowly, nor do you need a top-of-the-line storage system in most cases. The key is to find one that will provide the requisite performance.

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 virtual desktops and keep the CPUs fairly busy (averaging 60% to 75% is fairly normal). If there is a lot more RAM than 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. Most VDI deployments look to place 80 – 100 desktops per server; even at only 1 GB of RAM each, you'll probably need at least 96 GB of RAM. If the users are not CPU intensive, such as call center workers who utilize a browser to do their jobs, more may be possible. The key is to have enough memory so you don't need to swap to disk.

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. This is especially true in VDI implementations – the failure of a single host may have 100 desktops that need to be restarted somewhere, and this can cause a large load on the hardware components that must be planned for.

While we are on the subject of balance, some people are tempted to mix and match their virtual servers and their virtual desktops on the same servers and/or storage. This is not a recommended practice in all but the smallest of environments. It is best to dedicate servers just for server VMs and other servers for desktop VMs; likewise separate disks should be used for desktops and servers, if not entirely separate SAN or NAS arrays. The I/O characteristics are very different, the memory loads and CPU loads are different, etc. While they can be deployed on the same servers and storage, performance will be sub-optimal (and maybe much worse).

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 desktop 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, such as those from Liquidware Labs and Lakeside Software. These tools are specifically for VDI projects, whereas many vendor tools are more generic.

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. Nowhere is this truer than with VDI projects. Adding a server or some RAM is relatively easy and inexpensive, but a major upgrade to storage can be very time-consuming, not to mention expensive.

When sizing storage, most people simply count the TB of space required, calculate 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 the above table 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, 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 the table above 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 a VDI 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.

This cannot be overemphasized: users have used PCs for years with their own storage devices – even the slow SATA ones provide for 80 or so IOPS. In a virtual environment, it is not reasonable to think that every two or three users will get their own 15K SAS or Fibre Channel drive; they are simply too expensive. For more and more companies, VDI on SSD is the most cost-effective method, especially if technologies such as Linked Clones in VMware View are used to minimize the disk space requirements. It is all about the IOPS, and SSD drives shine in this area. You may still use spinning disks for user data or you may wish to offload that to dedicated file servers, but it is important to plan for enough I/Os to work well.

This is even more an issue after a server fails – 100 VMs starting all at once (known as a boot storm) can make the storage extremely slow, and the down time for users to reconnect may be unacceptably large. Hence, SSD drives may be the best choice as they can handle a large number of I/Os and still perform acceptably.

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. This scenario requires very careful planning to ensure sufficient IOPS are available to handle the load.

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, especially when spinning disks are used. This may be less of an issue in SSD drives due to their high I/O characteristics and relatively small sizes, so many administrators choose RAID 5 in that scenario.

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 to another area that requires careful planning: minimizing the number of base images that need to be deployed.

Minimize Base Images

One of the advantages of virtualizing your desktops is that the possible devices and variations of software and hardware that need to be supported can be minimized. You no longer need to support dozens of models of Dell, HP, Lenovo, and ASUS notebooks and desktops; you have a single standard hardware platform. While the platform is common, you may need a few variants of virtual desktops for different use cases, such as a standard 1-CPU, 1-GB RAM, 32-bit Windows 7 desktop for most users, a 1-CPU, 2-GB RAM, 32-bit Windows 7 desktop for some power users, and a 2-CPU, 4-GB RAM, 64-bit Windows 7 for developers and multimedia content developers. You decide these basic configurations based on the data collected during the planning phase and the needs and requirements of the users. Some basic guiding principles in creating various standard VM sizes include:

- Use 1 CPU whenever possible; there is extra overhead and lower density per physical server when two or more CPUs are used.
- Use enough RAM so that you don't need to swap to disk, but provisioning more RAM than necessary simply decreases the consolidation ratio and/or causes a lot more RAM to be purchased that would otherwise be necessary.
- Use a 32-bit OS instead of a 64-bit OS when possible, as 64-bit operating systems have more overhead to them, in part due to the larger memory space they can address.
- Minimize the number of variations to make it simpler to patch and maintain the virtual desktops.

All of the above really describe how to minimize hardware variations, but what about software? There are a lot more software combinations required in an enterprise than software ones in most cases. There are two solutions to this problem, namely:

- Create separate VMs for each unique combination of software and "hardware" configuration required. This tends to have the net result of many possible base images, all of which need to be patched and updated separately, greatly increasing the load on the IT staff and raising the TCO.
- Use application virtualization techniques, such as VMware's ThinApp or Microsoft's App-V. This is the next logical step in virtualization and separates the application from the OS in much the same way that server or desktop virtualization separates the OS from the physical hardware. This technique allows the same application to be packaged and run on multiple operating systems without any modification.

Most applications can be virtualized and will run the same way that they do in the physical world. The virtualized applications can then be associated with the appropriate desktops using the vendors' native tools (Such as VMware's View Administrator or Citrix XenApp) or via third party application distribution mechanisms (such as AD Group Policies, Microsoft's SCCM – System Center Configuration Manager, or Altiris). This technique minimizes the number of base image variations that must be maintained.

More information on application virtualization can be found in the [10 Ways to Save with Desktop Virtualization](#) white paper previously mentioned.

The key point in this section is that the greater the number of images that must be maintained in your environment, the greater the management cost and time in keeping each patched and maintained. Thus techniques that minimize them reduce the TCO.

Finally, we'll turn our attention to one of the major benefits of desktop virtualization: access to your data via your own (virtual) desktop from any device at anytime from anywhere.

Accept a Wide Variety of End-User Devices

One of the big changes going on in IT at the moment is that users want access to their data from many different devices, and they want that access at all times. When planning a VDI deployment, consider the methods that your users will be able to use to access their desktop. Options include:

- Laptops with a locally cached copy of their virtual desktop (at least a copy that can be cached as needed) for those who, in some cases, may not have Internet access to get to their desktop, such as those who travel and sales people. They can use their online desktop most of the time, and yet "check it out" when they will be out of the office.
- Existing PCs, desktops, or laptops running Windows, Linux, and/or Macintosh OS variants, can be used. This does not offer many of the savings that VDI promises initially, but may save money up front, with these devices being replaced by other options as they wear out, come up for replacement, etc. These PCs may also be employee or contractor devices they use to access their desktops from home.
- Thin clients or zero clients, which are basically dumb terminals that have enough intelligence to connect to the VDI infrastructure to access the virtual desktops. The difference between a thin client and a zero client has to do with the OS on the device. A thin client has a small OS (such as Windows XP Embedded or Windows 7 Embedded), while a zero client has no OS, just a BIOS with API calls to connect to the appropriate virtual desktop. This is the cheapest option, for the device itself, as well as for the power to run it (some can even be powered by power supplied over Ethernet cabling, reducing the cables required), and for the maintenance to keep it going (there are usually no moving parts). These devices have an average life span of eight years, as there is nothing to upgrade and no moving parts to fail.
- Tablets, such as Android tablets or iPads that are becoming more and more popular. They provide simple, easy, convenient access from anywhere with a network connection, and many have 3G or 4G capabilities built in, allowing access virtually any time.
- Smart phones, for those times that none of the above are handy and access is required. Approximately two-thirds of phone owners today have smart phones and most want access to information via them.

Planning for these devices, including security issues (such as how to remotely wipe them, and how to securely provide access to the data stored at the company, as well as to the virtual desktop itself) and support issues when users run into trouble are all big concerns that must be planned for and managed. By virtualizing the desktop and possibly even the applications, support personnel have fewer issues to deal with (though the variety of devices they use to connect may provide an equally large challenge).

This trend of access anytime, anywhere with any device is being accelerated by the Bring Your Own PC (BYOPC) trend where users bring their own devices to connect to the corporate network and resources, including their desktops. Often, companies will provide a stipend every few years for the employee to purchase a device of their choice. Repair and maintenance of the device then becomes the employee's responsibility instead of the company's.

Summary

In conclusion, the five "secrets" espoused here will greatly increase the chance of a successfully VDI implementation. It is very important to plan for this, to balance the hardware environment to handle not just average, but peak load, to properly size storage, not just for capacity, but for performance as well, to minimize the number of base images so that the costs of maintaining each can also be minimized, and to accept and embrace the fact that people will connect with a wide variety of devices, and to create a plan to accommodate as many of these devices as is feasible at the lowest cost.

Planning and implementing is a challenge, but the results of doing so is cost savings, greater productivity, and a happier work force that has the flexibility to work when, where, and with what device they wish to use.

Learn More

To learn more about how you can improve productivity, enhance efficiency, and sharpen your competitive edge, check out the following Global Knowledge courses:

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

[Citrix Desktop Virtualization Fast Track](#)

[Citrix XenDesktop 5 Administration](#)

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.