

Realtime
publishers

Desktop Virtualization and Evolving Strategies for IT Service Delivery

Jaime Halscott

sponsored by



Introduction to Realtime Publishers

by **Don Jones, Series Editor**

For several years now, Realtime has produced dozens and dozens of high-quality books that just happen to be delivered in electronic format—at no cost to you, the reader. We’ve made this unique publishing model work through the generous support and cooperation of our sponsors, who agree to bear each book’s production expenses for the benefit of our readers.

Although we’ve always offered our publications to you for free, don’t think for a moment that quality is anything less than our top priority. My job is to make sure that our books are as good as—and in most cases better than—any printed book that would cost you \$40 or more. Our electronic publishing model offers several advantages over printed books: You receive chapters literally as fast as our authors produce them (hence the “realtime” aspect of our model), and we can update chapters to reflect the latest changes in technology.

I want to point out that our books are by no means paid advertisements or white papers. We’re an independent publishing company, and an important aspect of my job is to make sure that our authors are free to voice their expertise and opinions without reservation or restriction. We maintain complete editorial control of our publications, and I’m proud that we’ve produced so many quality books over the past years.

I want to extend an invitation to visit us at <http://nexus.realtimedpublishers.com>, especially if you’ve received this publication from a friend or colleague. We have a wide variety of additional books on a range of topics, and you’re sure to find something that’s of interest to you—and it won’t cost you a thing. We hope you’ll continue to come to Realtime for your educational needs far into the future.

Until then, enjoy.

Don Jones

Introduction to Realtime Publishers..... i

Chapter 1: Understanding the Challenges of Building VDI Now for Future End User Service Delivery..... 1

 Desktop Virtualization: Means or End? 1

 Planning..... 3

 Licensing..... 3

 The Goals of Desktop Virtualization..... 4

 Identifying How Users Interact with Desktops Is Critical..... 5

 Lesson Learned..... 5

 Reliability 6

 Speed..... 6

 Supporting Offline Users 7

 Understanding Internal Connectivity Needs 10

 Tackling External Connectivity Problems..... 10

 Metrics for Ensuring Network Effectiveness..... 11

 Where Do Physical Desktops Fit? 14

 Summary 15

Copyright Statement

© 2011 Realtime Publishers. All rights reserved. This site contains materials that have been created, developed, or commissioned by, and published with the permission of, Realtime Publishers (the “Materials”) and this site and any such Materials are protected by international copyright and trademark laws.

THE MATERIALS ARE PROVIDED “AS IS” WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, TITLE AND NON-INFRINGEMENT. The Materials are subject to change without notice and do not represent a commitment on the part of Realtime Publishers its web site sponsors. In no event shall Realtime Publishers or its web site sponsors be held liable for technical or editorial errors or omissions contained in the Materials, including without limitation, for any direct, indirect, incidental, special, exemplary or consequential damages whatsoever resulting from the use of any information contained in the Materials.

The Materials (including but not limited to the text, images, audio, and/or video) may not be copied, reproduced, republished, uploaded, posted, transmitted, or distributed in any way, in whole or in part, except that one copy may be downloaded for your personal, non-commercial use on a single computer. In connection with such use, you may not modify or obscure any copyright or other proprietary notice.

The Materials may contain trademarks, services marks and logos that are the property of third parties. You are not permitted to use these trademarks, services marks or logos without prior written consent of such third parties.

Realtime Publishers and the Realtime Publishers logo are registered in the US Patent & Trademark Office. All other product or service names are the property of their respective owners.

If you have any questions about these terms, or if you would like information about licensing materials from Realtime Publishers, please contact us via e-mail at info@realtimepublishers.com.

[**Editor's Note:** This book was downloaded from Realtime Nexus—The Digital Library for IT Professionals. All leading technology books from Realtime Publishers can be found at <http://nexus.realtimepublishers.com>.]

Chapter 1: Understanding the Challenges of Building VDI Now for Future End User Service Delivery

Our industry's trend towards virtualization quickly evolved from a buzz phrase to a storm. Server virtualization has been a great fit for a vast majority of companies, but can the same be said for desktop virtualization? There is no doubt that desktop virtualization and virtual desktop infrastructure (VDI) are emerging industry trends, but are they right for your exact business goals?

That is the question this guide intends to answer. Although building VDI is a goal of many IT organizations, does the excitement surrounding VDI technology obfuscate the business problems it intends to resolve? This book will serve as a guidepost toward effectively determining the best alignment between technology and business need. That guidepost can serve as both a directional marker and a warning—and the intention is to accomplish both. The aim is to help the IT architect determine what business requirements drive the move toward desktop virtualization as well as choose the right path in desktop virtualization.

Desktop Virtualization: Means or End?

The IT industry finds itself in a quandary between technology and business requirements. Technology exists that dramatically improves the efficiency of the IT organization, but does the implementation of that technology in fact burden its users? Could an IT return on operational efficiency actually hinder the business in other ways? Determining whether desktop virtualization represents a solution or a means to that solution requires asking key questions:

- How does desktop virtualization align with end user service delivery?
- What additional benefits or enhancements to end user service delivery will come from the VDI project?
- Are those benefits cost justified?
- Do the desktop virtualization steps taken now align and integrate with future growth and other emerging technologies?

Timing the transition from legacy desktop infrastructure to a method for anytime, anywhere user workspace delivery is critical. As IT budgets continue to tighten, no IT architect can or should realistically ask executives to approve replacement of infrastructure that “isn’t broken.” Thus, it is necessary to frame the entire endeavor around the true goal of end user service delivery. However, legacy desktops may be meeting the goal of end user service delivery quite well.

So the big question is, Why replace them? Desktop virtualization isn’t an end in itself. It is part of a path toward the evolving delivery of services to end users. This journey is transformative and won’t happen overnight. There isn’t going to be a technical solution to all of the problems that will be encountered. Desktop virtualization doesn’t address them all at this point and as a technology set, may never do so. Vendors have made great strides towards building technical solutions that we will address in detail; however, understanding *where* to leverage desktop virtualization and *where not to* is of the utmost importance.

Desktop virtualizers face a whole new set of challenges from their server counterparts. Even seasoned IT architects familiar with the aspects of a successful server virtualization project can be caught off guard by the differences between server and desktop virtualization. Understanding these critical differences is crucial to a successful project.

Reference

Later chapters will analyze the differences between server and desktop virtualization with the goal of better understanding which technologies fit which future goals.

In addition, technologists must identify the reasons for the project in the first place and answer key questions: Why virtualize desktops? What is driving the VDI initiative? If the reasons and challenges aren’t thought through and addressed at the onset, the project will quickly fall short of the business goals. Technical challenges are just part of the problem. End user training, awareness, and expectations are vital elements to address throughout the process. Let’s start with a series of lessons designed to help you understand how desktop virtualization differs from server virtualization.

At the core, there are a lot of similarities. However, the implementation of desktop virtualization introduces challenges. We will analyze those challenges and come to an overall determination of whether desktop virtualization should even be a part of your overall IT plan. We will then look at how to determine how much of a fit it really is and where to leverage it most effectively. The first step in the process is to analyze the server virtualization model with the goal of understanding why it doesn’t work for desktop virtualization.

Lesson 1: Why the Server Virtualization Model Doesn't Work for VDI

Virtualization is a topic that just about every IT professional has had at least some basic exposure to at this point. Many businesses have adopted some form of virtual infrastructure, but this adoption is, at this time, mostly limited to server-based virtualization. That isn't to say that desktop virtualization is foreign, but it is emerging, and as such you may have very limited exposure to it.

Server virtualization brought with it a host of justifications for projects, including server consolidation, decreased overall hardware costs, decreased support costs, simplified backup, and many others. Do these hold true for desktop virtualization? We'll explore that answer as we compare server and desktop virtualization. Let's start with the planning phase.

Planning

I've seen a lot of server virtualization projects succeed even with the worst planning up front. The last server virtualization project I saw run into serious problems was the result of a complete lack of planning. Even that one was salvaged and showed an immediate return on investment. My personal experience has shown that desktop virtualization is less tolerant to planning missteps. This isn't a bad thing per se.

Every IT project should be planned from the start, and desktop virtualization is no different. Building a plan that takes into account as many variables, goals, and metrics as possible makes the project that much more likely to succeed.

Desktop virtualization requires careful planning, testing, and piloting because it is so closely tied to end user experience. Users will not tolerate slowdowns or peripherals that don't work. *Remember that you aren't delivering just a desktop in desktop virtualization. You are delivering a whole user experience.*

Server virtualization gives more wiggle room, if you will, in the project. Taking a single physical server and repurposing it to support multiple virtual servers doesn't require a complete restructuring of the network or additional hardware. Existing metrics for processor, memory, disk I/O, and such should already be in place on the physical server. These performance counters do a good job of showing approximate workloads for virtual servers over time. This data makes it relatively easy to pilot virtual server workloads. Desktops are different. Some users are more resource intensive than others and have less tolerance for slowdowns. In general, the task of finding out this information isn't easy.

Licensing

Licensing is also different. In many cases, server virtualization doesn't require any additional operating system (OS) licensing. Most IT staff can rather quickly articulate what the Windows Server virtualization licensing rights are for the Standard, Enterprise, and Datacenter versions. The maturity of the server virtualization landscape has made this information widely known in the course of a few years.

That said, I'm willing to bet that hardly as many readers will be able to explain without extensive research how many instances of Windows 7, for example, can be installed on a particular VDI platform and how many can legally be run concurrently. There are alternative licensing schemes designed specifically for virtual desktops. Finding out how much a virtual license costs and the legal rights associated with that license remains a difficult task. This is part of a normal growing pain process with any technology. Vendors like Microsoft have come up with several plans for desktop virtualization licensing, which are meant to simplify and keep costs down—however, these plans still aren't fully fleshed out.

But before you can jump into licensing needs or even determining the key differences between end users' needs, you need to get back to the goals of the project. Let's start by analyzing the goals of desktop virtualization.

The Goals of Desktop Virtualization

You've probably started down this VDI path with the belief that the goals of desktop virtualization are the same as those for server virtualization. Trying to start with a familiar knowledge base in server virtualization shouldn't hurt a desktop virtualization project, should it? I think it might. Although some of the server virtualization goals apply directly to desktop virtualization, for most part, they don't share the main business drivers. Understanding the differences is critical.

Server virtualization is largely focused on driving up efficiency and driving down costs. It tends to achieve these goals well. IT architects that take on server virtualization can draw models that show cost savings in a variety of ways:

- Greater utilization of existing server hardware
- Space savings in the data center
- Reduced server maintenance costs
- Reduced power and cooling costs
- Software licensing savings
- Reduced server deployment time
- Application isolation
- Multiple platforms on the same hardware

Desktop virtualization provides some of these benefits but not nearly as directly or in a manner that is as easy to quantify as server virtualization does. A desktop virtualization solution achieves a host of soft cost savings that requires a different analysis. These cost savings are going to be more difficult to quantify at the onset of the project than they are with server virtualization. It doesn't mean they aren't just as beneficial, but rather that they aren't going to be immediately obvious.

In particular, desktop virtualization reduces support costs associated with deployment and management of the PC infrastructure by streamlining the process of provisioning new “desktops,” deploying new applications, simplifying imaging because of uniform virtual hardware, and offering a longer refresh cycle for the existing physical desktop infrastructure.

Identifying How Users Interact with Desktops Is Critical

Servers are back office products with which users don’t typically interact. Such is not the case for desktops. User experience is critical to the success or failure of a VDI project. And the goals of a desktop virtualization implementation should account for user expectations for performance and seamless access.

Do users perform separate roles when running on a physical desktop than on a virtual desktop? The answer is obviously no. Users aren’t all of the sudden using different applications or working with different data after a desktop virtualization implementation. Users are simply trading one form of service delivery for another form of service delivery. The applications and data that a user accesses will be no different whether interacted with directly on a physical piece of hardware (a traditional desktop) or delivered from a data center (a virtual desktop).

This user workspace concept will be important going forward because we are going to frame the use of VDI and other desktop and application virtualization efforts in this context. The concept of the user workspace is almost completely foreign to the server virtualization model. Why? The answer is a simple one: end users don’t interact with a server from an experience standpoint. The server is serving applications and data to the desktop, laptop, tablet, mobile device, or some other presentation method.

The most often cited reason for virtualizing desktops centers on the challenges and costs of managing PCs. Managing physical PCs is expensive and time consuming. The end user workspace environment can include tens to hundreds of various applications, all of which have their own supporting files, necessary dependencies, file conflicts, and so on. This results in conflict and a lack of standardization between desktops.

Lesson Learned

So rather than attempt to build a desktop virtualization model with concepts borrowed from server virtualization, a better plan is to start with a new model from scratch. Doing so reduces the temptation to pull cost saving justifications, planning strategies, and benefit/goal analyses from server virtualization models. The following list highlights factors that should be taken into account for a desktop virtualization cost model:

- Basic infrastructure upgrade costs (that is, more/faster servers)
- OS and application licensing
- Migration from physical to virtual desktops
- End user training
- Physical and virtual desktop management costs

You will find that these factors are difficult to quantify. How can you determine costs before you build a pilot and test? The trick is to start construction of a cost model that takes into account all of the factors you reasonably believe are going to affect the overall expense, build a proof of concept, then develop a pilot that is small enough to represent the desired goal. The results of the proof of concept and the following pilot will test the accuracy of the cost model and expose any weaknesses that can be corrected in the model before proceeding with a full-scale project.

Lesson 2: How to Overcome Difficulties Supporting Distant and Offline Users

Distant and offline users present one of the greatest technical challenges to the delivery of user workspaces using VDI. Much like server virtualization, desktop virtualization relies on a back-end and appropriate connectivity to that back-end to serve or present the workspace. This setup is fine as long as there is an appropriately robust mechanism for delivery of the user workspace end-to-end. Any level of latency, dropped connections, or general instability will seriously hamper the productivity of the end user.

We have already determined that the goal is end user service delivery. Interruptions to that delivery will critically impair the overall productivity of the user. These interruptions can result from many factors, such as problems with network reliability. To fully understand the challenges of service delivery to remote users and further down the spectrum, offline users, you have to look at the factors that affect both reliability *and* usability. Reliability influences whether you can deliver the service in a consistent manner. Usability is whether you can deliver it with a quality that won't adversely impact the user's productivity.

Reliability

Let's look at network reliability first. Every network connectivity technology is going to have varying levels of reliability. This doesn't come from the technology used as much as the way it is implemented and how well it is supported. Speed also plays an important role.

Speed

Network speed is crucial for a variety of reasons, and is more complex than just looking at the maximum theoretical speed of the connection. When you think about network speed, most often you think about bandwidth. However, bandwidth is only going to be part of the problem. Much more important is the latency introduced by the network. In a tightly controlled LAN, this isn't going to be a big problem.

Gigabit Ethernet is a high-bandwidth and low-latency technology. As a result, it is ideal for service delivery and VDI. However, it is uncommon to have end-to-end Gigabit Ethernet to all users all over the planet. The reality is that we have to live with some levels of reduced bandwidth and latency. Desktop virtualization implementation must be conscious of the way users interact with the virtual desktop infrastructure—inside the LAN on Gigabit Ethernet or faster as well as outside the LAN where the network speed can be much less.

Supporting Offline Users

Offline users present a slightly different set of challenges than remote users do. An offline user is one that spends a lot of time with little or no connectivity. A good example is the field sales users that spend a vast majority of their time in airports and hotels. The offline user may be remote most of the time and have connection back to the data center for access to the VDI.

Traditionally, offline or remote users have made use of a solution to address connectivity interruptions and allow for productivity while disconnected. Technologies a decade old allow for offline files. Users have also had traditional hardware at their location that would keep resident copies of needed applications, associated settings, and preferences.

In the world of the virtual desktop, the desktop session or published application has been presented remotely to the end user. The challenges to accomplishing this with limited to no connectivity back to the data center are well known. Quite simply, if you can't reach your data center, you can't get your "desktop" or applications. In order to address this limited or no connectivity problem, vendors have created several unique solutions.

Supporting offline users can take the form of an offline desktop, essentially a local virtual copy of all the OS and application needs of the user. However, we still have to manage that desktop and the computer on which it runs. The technologies for delivering an offline desktop vary, but for the most part involve some sort of checkout and version control mechanism for the virtual desktop. This setup allows the user to work with a local copy of the desktop and synchronize the changes back to the data center as connectivity permits. The problem is that in order to have a seamless user workspace, the user must understand the processes associated with check-ins, checkouts, version control, and synchronization. Each vendor will have slightly different process sets.

Reference

Chapter 2 will examine the different types of access hardware that will greatly affect the user workspace presentation. New types of devices may not support offline desktop function and might be a clunky fit for VDI.

VDI's technologies are challenged to meet end user service delivery goals when it comes to offline and remote users. The limitations imposed by required connectivity are too much to overcome for many businesses. Vendors have made measured strides to achieving success here; however, preserving a satisfactory user experience remains a challenge with VDI and this user class.

Lesson 3: How to Support Heavy Users

Heavy users represent a unique class of users. The heavy user typically is more tied to specialized hardware than the average knowledge worker. For example, an architect that relies on high-resolution graphics and larger or multiple monitors will not be easily supported in a desktop virtualization scenario. I typically refer to this type of user as a power user, but because Windows has a built-in user group also called “Power Users,” I’ve started referring to these users as heavy users. Really, the terms are interchangeable.

Technical solutions have evolved to help meet the increased demands that the power user places on a VDI system, but again you must focus on the reason for desktop virtualization in the first place and, more importantly, the service delivery aspect. Is the heavy user class one that should be a part of a desktop virtualization scenario or is this simply not the way to meet the goal of end user service delivery?

Let’s start this analysis by looking at the way the typical heavy user interacts with technology. Using this information, you can determine what or how much of that work fits with VDI and other virtualization efforts.

What does an example architectural or engineering heavy user employ in a traditional non-virtualized desktop environment for service delivery? This user will have a CAD software package that has processor and memory demands that are greater than typical business productivity software packages. This software is needed in addition to other more common business applications such as office productivity, time tracking and billing, and various other supporting point solutions.

The second consideration is hardware. At the core, the processor, memory, and storage demands are higher for the heavy user than for the general business knowledge worker class of users. These demands are placed by the requirements of the specialized software packages themselves. These extra resources are expensive.

The costs per user that must be supported on the desktop virtualization platform will be higher for the heavy user than for the knowledge worker. Desktop virtualization solutions are starting to help with the specialized hardware and even now provide support for USB device redirection.

What are the criteria for when a heavy user can and cannot be supported in a virtual desktop? Here are some of the factors to consider:

- Directly connected peripherals
- Specialized input devices
- Multiple monitor support
- Complex applications

Heavy users often have requirements for 3D rendering and full motion video that are met quite well with physical desktops. The ability to offload these functions to a dedicated graphics system in a physical desktop is going to be easy and inexpensive, but this isn't the case with a virtual desktop. Only recently, companies have begun to offer 3D and full motion support as part of the technologies to improve the visual capabilities of virtual desktops. There has been some level of support for 3D in many vendor offerings, but the most recent generation has allowed it to approach the experience that heavy users expect from a traditional desktop.

Microsoft calls this RemoteFX. It is a great start, but it has requirements for Windows 7 end-to-end, both on the client and on the virtual desktop. Also, the graphics processing on the server hosting the virtual desktop has to be increased in order to support the more intensive graphics display required from RemoteFX. Servers that have thus far been used to support only server virtualization may not work. These servers will have to have expansion abilities to take professional workstation-level graphics cards.

Microsoft isn't the only provider that is keenly aware of the rich user experience that virtual desktop users demand. Citrix has built the HDX protocol for use with their virtualization solution, which enables what they term as the "high definition" desktop virtualization user experience. If you are unable to support the end-to-end Windows 7 requirements or upgrade the server hardware to support these technologies, the user experience will diminish accordingly.

In addition, monitor spanning was introduced a few years ago to assist users with multiple monitors to have remote desktop sessions take advantage of more than one monitor. It does have limitations that the heavy user may not be able to live with. Maximum resolution is limited and not all virtual desktop platforms will support all of the available monitors.

The possibility that heavy users can be supported isn't out of the question. However, you will have to address the factors mentioned to determine how to best support them. In doing so, you might find that the needs of your organization's heavy users make virtual desktops unfeasible for them.

Lesson 4: Making Sure the Network Infrastructure Is Up to the Task

The lynchpin in successful VDI delivery of user workspaces is its network infrastructure. In particular, network latency and reliability are paramount to success. So how do you determine the appropriate amount of bandwidth and thresholds for reliability?

One approach involves seeking the solution that fits best along the axes of good, fast, and cheap. As the adage goes, one can only expect the ability to prioritize two of these; however finding the combination that meets your needs is fundamentally important.

Network and infrastructure reliability are paramount, so consider the category of *good* to be a constant. *Good* means that the network connectivity and infrastructure are going to guarantee as much uptime as possible. In order to achieve this end, we are going to focus on two aspects, internal and external connectivity.

Understanding Internal Connectivity Needs

Internal connectivity will encompass all interconnects and networking components such as switches and routers. Current LAN technology is mature when connecting internal hosts. Vendors and the industry in general already have inexpensive solutions for fault tolerance and redundancy. Entry barriers for providing highly reliable internal infrastructure are far fewer than they were even 5 years ago. This is welcome news for an IT architect that must balance budget constraints with the business goals of a highly available virtualization solution.

Any supporting network infrastructure must support Gigabit Ethernet at a minimum, and 10 Gigabit Ethernet is quickly becoming the new standard. This standard has evolved due to ever-increasing network throughput creating areas of congestion in gigabit-speed networks. That congestion is effectively eliminated with the upgrade to 10 Gigabit Ethernet. With *good* as a constant, *fast* and *cheap* are the variables that remain. *Cheap* can be accomplished with internal connectivity. 10 Gigabit Ethernet prices have dropped significantly, giving even limited budgets the ability to invest in core infrastructure at this speed.

Fast internal connectivity also isn't much of an issue. It is nearly impossible to find anything slower than Gigabit Ethernet in the server space. Technologies inherent to servers allow for both load balancing and failover, which is great for getting a boost in speed and reliability, essentially combining *fast* and *cheap* in one package.

Tackling External Connectivity Problems

Where the challenges become most evident is with external connectivity. Choices for Internet access are numerous. For the purpose of analysis, we are going to break down this concept into two categories: those with guaranteed service levels and those without. Certain technologies are more reliable and cost more. It isn't that the second class of external connectivity is necessarily unreliable, but it is designed for speed over stability.

Internet connectivity to the remote users' endpoint represents the most likely hindrance to productivity. This issue has become much more complicated as users access their workspaces over cellular-enabled mobile devices. Are these users remote or offline? The answer depends on where they are, and the capabilities of the device, among a host of other factors.

Later, this book will discuss the types of devices and the emergence of new workspace access approaches that are pushing these limits. For now though, let's establish that access can be on so-called 3G and 4G networks, WiFi, WiMAX, or even legacy cellular networks. The IT architect will have to take into account all of these possibilities and set realistic expectations for the user experience for each of these types of networks and devices.

Metrics for Ensuring Network Effectiveness

What are the metrics that need to be monitored in order to ensure that the network infrastructure is up to the task? Internally, there should be metrics for total network utilization and per host/server utilization. Latency isn't going to play an important role in the determination of adequate resources on a high-speed network, but it will play a huge role in determining the adequacy of Internet connections. Externally, the metrics are much more important because they set the stage for both data center connectivity and remote user connectivity.

Employ a representative user during the pilot to determine the bandwidth requirements overall. The class of user won't play much of a role in determining bandwidth requirements with the obvious exception of video and multimedia-intensive users. The particular protocol used by the desktop virtualization solution handles screen updates and will affect the bandwidth more than the class of user.

For the purpose of determining bandwidth usage, let's select a single access method at first and look at how different vendor protocols will affect the total bandwidth consumed per user. By access method, I mean desktop virtualization using protocols such as RDP and Citrix ICA instead of application virtualization. These will take more bandwidth than application virtualization in general. The reason is that the display of the entire desktop and the graphical interface associated with it is going to consume more bandwidth than the display of only a virtual application.

Incremental user experience optimizations have been added as well, such as protocol optimizations and compression algorithms that decrease the bandwidth required to display a virtual desktop. Virtual infrastructure vendors have varying-level solutions to make the most out of limited bandwidth and the implication is that as time progresses this is going to improve the user experiences when running over low-bandwidth links.

End users expect a virtual desktop experience that closely approximates the one they are used to with their physical desktops. Slowdowns in display of graphics and videos will not be acceptable. If bandwidth cannot be provided that will support these users' expectations for display of video, users will complain about slowdowns and the entire project can be jeopardized. The importance of determining the representative user bandwidth requirements cannot be overemphasized. If these needs aren't determined up front, the bandwidth costs can balloon, jeopardizing the entire project.

Once you establish the baseline requirement for a representative user, you can build a formula by multiplying the number of each that are expected to be concurrently accessing their workspaces by the requirements for each class that you have determined. Take this number and add a peak usage factor to compensate for spikes without a degradation of service. At the same time, factor in anticipated growth.

Bandwidth requirements vary wildly, not just between classes of users but also depending on the way that you define a class of user and their particular role. Optimizations in the particular virtualization platform, compression technologies, and many other unknowns will affect bandwidth. Requirements as low as 20Kbps and as high as 180Kbps are not unreasonable.

Remember too that how you present the user workspace impacts the overall network requirements. Using existing technologies, this presentation can be an entire desktop or a series of presented virtual applications. Testing either scenario is critical to success. I alluded to it earlier, but there will be variation, however slight vendors might claim, in the different desktop virtualization solutions.

Once you have tackled the calculations for data center bandwidth and factored it into your costs, you need to turn to the more difficult calculation of the minimum requirements for the client side. As noted earlier, the client can be accessing the user workspace from anything ranging from a mobile device with limited bandwidth to a traditional thin client or even a full desktop.

Lesson 5: Preparing for Unforeseen Costs

The cost models for server virtualization don't translate exactly for desktop virtualization. We saw this in our earlier analysis. The reasons for this are many—as we have looked at already. Even the best-planned project is going to be susceptible to some sort of cost overrun. Of course, the best way to ensure a successful budgeting for the project is to analyze as many factors as possible. Even then, there may be some that cannot be calculated with complete certainty. Let's start with the known quantities that are shared between both virtual and physical infrastructure design, then look at the factors that are exclusively part of a desktop and, in many cases, a server virtualization build-out.

If we start the cost model analysis with the factors that will be most comparable between physical and virtual desktops, the first item that comes to mind is the licensing costs associated with the OS and applications. Virtual desktop and virtual application software licensing hasn't yet matured to the level that it has in the server virtualization arena. The answer to the question of how many copies of a particular desktop OS or application can be run virtually isn't as easy to answer as the server question.

Hardware needs are certainly going to factor into the overall cost for the virtualization effort. It goes without saying that there are certain baseline needs for just about every form of virtualization technology. These include processor support for the platform at a minimum. Almost every server shipped in the past few years has the requirements for basic virtualization support, so there will be the drive to save costs by reusing and repurpose these resources. This is a great idea, but it can't be taken at face value.

The ability to present user workspaces in a consistent manner will depend on the creation of a backend infrastructure that is seamless to the users working in it. Existing high-availability technologies such as storage area networks (SANs), RAID arrays, redundant power supplies, and other technologies are going to be reusable when creating a desktop and application virtualization system. A single point of failure is going to interrupt the service delivery for more than a single desktop, so all design work has to focus on creating as reliable a system as possible within the allotted budget.

Most organizations are able to tolerate downtime for a single physical desktop system. Desktops are considered a commodity by many businesses. If one fails, it gets replaced and business continues. Things change up a little bit once we start virtualizing. A single server can handle several virtual desktops. The same principle may apply to a virtual desktop that does to a physical one, meaning that a single virtual desktop can fail without massive business interruption. However, you can't have a server fail because it will lead to several virtual desktops being unavailable.

High availability comes at a cost. Swapping a commodity desktop containing an inexpensive hard drive for a virtual desktop that uses high-end SAN storage is going to be several orders of magnitude more expensive. This additional cost will factor into the decision of reduced management costs versus increased infrastructure costs.

Perhaps the most difficult task is determining the specialized skills required to implement the project and support the solution going forward. An IT architect is going to have to understand the current technical landscape and anticipate the future needs of both the infrastructure as it will stand and any additional changes that might occur over the life of the infrastructure. The following list highlights factors that must be considered:

- The experience of the IT staff in general
- Previous staff experience with virtualization technologies
- Staff experience supporting remote users
- Implementation and support offerings from vendors
- Ability of the staff to embrace innovations
- Costs to outsource the project and support to a third party

There will be more than this to consider, but all too often, the IT staff isn't prepared to undertake a project of this magnitude without additional training, additional personnel, or third-party assistance. Performing an initial analysis of the specialized skills required for various virtualization efforts will greatly assist in providing a cost model that is accurate.

The most important takeaway is that changing from physical to virtual desktops may require more upfront costs than you might expect. Planning carefully for those costs is crucial to avoiding overruns.

Lesson 6: Making Sure the Reasons for the Project Are Well-Defined

Thus far, we have looked at cost models and how savings, while a part of the overall reason for a VDI project, are not the main business driver. Security is also an important factor but is secondary in many cases. When all factors are considered, the real reason for desktop virtualization is to obtain a greater level of consistent service delivery. That level of service delivery ties directly to the stated goal of providing users with a workspace that allows them to perform their job role regardless the how the workspace is accessed.

Previous experiences with server virtualization may lead seasoned virtualization veterans to gloss over the need to clearly define goals and scope. Even IT architects with limited server virtualization experience can fall into this trap. *The wealth of information that surrounds server virtualization efforts may tempt you to simply adapt its goals and head blindly into a desktop virtualization effort.* This method will certainly not address the appropriate delivery of user workspaces because its requirements are beyond those found in the server virtualization concept.

Vendor relationships are another potential pitfall. All too often, a vendor has a product, or even an entire solution, set to sell. Although these might meet goals and align with end user service delivery, every element of the VDI project must be critically analyzed against the well-defined reasons for the project. All too often an executive reads an article that espouses the virtues of desktop virtualization. This leads them to believe that it alone will solve all of the problems of PC management and usher in a new era of productivity. Vendors are trusted partners but can sometimes be tempted to pounce on these opportunities a little too enthusiastically. Tempering vendor relationships will give you a more accurate idea of what can and cannot be delivered from a desktop virtualization solution.

Where Do Physical Desktops Fit?

Reusing and repurposing physical desktops to support virtual desktops is a common approach. The first step to leveraging legacy desktops in a plan for future end user service delivery is to determine the remaining usable life for the desktops in use. If the focus is on creating effective user workspaces leveraging a virtualization solution, it doesn't make much business sense to replace desktops and then relegate them to nothing more than serving up a desktop session hosted on a server.

Once you've determined that the desktop meets the requirements for delivery of a virtual desktop, it can be repurposed to do so. However, having to manage a physical desktop OS and then a virtual desktop OS greatly increases the support requirements for that desktop. The best solution is to set up the desktop as a thin client. Doing so reduces costs by allowing existing hardware to be reused.

Vendors are starting to make great strides in the ability to repurpose physical desktops to become much more like traditional thin clients. The concept of the desktop as a thin client isn't really anything new. Vendors have been building solutions for years to take advantage of terminal services by converting PCs into thin clients. Recently, Microsoft announced Microsoft Windows Thin PC (WinTPC). Essentially this is a smaller footprint, locked down version of Windows 7 that lets businesses reuse legacy desktops in a VDI deployment. It has a lot of advantages in security and deployment over reusing the same Windows OS desktop and provides a means for future transition to thin clients.

This sort of approach appears to be the trend because of the added flexibility and reduced costs. It also allows your organization to “roll back” to traditional desktops should the need arise. This software solution will prove popular over the next few years because of the ability to have a desktop or thin client as needed.

Summary

There are a tremendous number of differences between server and desktop virtualization. The business drivers are different for each. The user experience is a major concern in desktop virtualization, but almost not at all in server virtualization. The nature of the users accessing the infrastructure is a large determinant in whether to virtualize desktops and, if so, how much to virtualize desktops. The cost savings come in through reduced support for desktop virtualization, more so than efficiency as in server virtualization.

You will need to navigate the current physical infrastructure that you have so that you can build a plan for what to reuse and what to remove. You will have to understand the challenges to service delivery and where the shortcomings in current technologies will impact those challenges. This isn't something to take lightly—end user experience is critical to success.

The next chapter will explore how the landscape of desktop virtualization is evolving outside the traditional desktop as a means for accessing virtual desktops. We'll also explore different service delivery methods and emerging technologies that will have to become a part of the plan for VDI. The cloud will be a key concern. There are many approaches to embracing the cloud to support desktop virtualization, and we will look at what works and what doesn't. I want you to be able to anticipate how VDI is going to be a part of an emerging strategy for end user service delivery. Chapter 2 will provide critical insights in how to avoid pitfalls that will either require you to scrap VDI completely or won't integrate your VDI well with future changes.

Download Additional Books from Realtime Nexus!

Realtime Nexus—The Digital Library provides world-class expert resources that IT professionals depend on to learn about the newest technologies. If you found this book to be informative, we encourage you to download more of our industry-leading technology books and video guides at Realtime Nexus. Please visit <http://nexus.realtimepublishers.com>.