

Realtime
publishers

The Shortcut Guide[™] To



Implementing Virtualization in the Small Environment

sponsored by



Greg Shields

Chapter 4: What Are the Best Practices in Implementing Small Environment Virtualization? 31

 The Life Cycle of a Virtualization Implementation 32

 The Importance of Education and Awareness 33

 Assessing Your Environment 33

 Measuring Candidate Performance with Native Tools..... 36

 Leveraging Easier Performance Toolsets..... 38

 The Importance of Assigned Virtual RAM..... 39

 Purchase & Implementation 40

 Converting Physical Servers to Virtual with P2V 40

 Extending Backup Capabilities 42

 Adding Affordable Disaster Recovery..... 44

 Implementing Virtualization Involves Technology and a Plan..... 46

Copyright Statement

© 2009 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 or 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 eBook was downloaded from Realtime Nexus—The Digital Library for IT Professionals. All leading technology eBooks and guides from Realtime Publishers can be found at <http://nexus.realtimepublishers.com>.]

Chapter 4: What Are the Best Practices in Implementing Small Environment Virtualization?

In three short chapters, this guide has discussed many of the questions you and your business should be asking about virtualization. You've learned about the technologies behind virtualization, with a focus on the underpinnings that make virtualization a play for small businesses. You've also come to know the financial benefits that virtualization brings to your environment.

In the previous chapter, you got a comprehensive look at the components you'll need to get started in virtualization. Chapter 3 took a technology-based focus on the pieces of technology you need in order to start your virtualization project. There, you learned about the types of storage available, the impact of networks and network cards, and the critical role of available memory in your physical hosts. This chapter will augment that information with the added knowledge and best practices you need to properly implement that virtual environment.

It goes without saying that any technology insertion must have a well-defined purpose if it is to be considered successful by the business. Yet there's a problem intrinsic to many IT projects in that they're often not well-scoped before the vendors' checks are written. In fact, Computer Associates reported in a 2007 survey (<http://www.continuitycentral.com/news03155.htm>) that 44% of respondents were unable to declare their server virtualization deployments a success. The primary reason was focused on the inability for respondents' businesses to quantify their project's ROI.

It is because of issues like this that this guide was developed. It was stated back in Chapter 1 that the central goal of this guide is to help you value the benefit virtualization provides to your small environment. That benefit is driven primarily through virtualization's low cost of entry and is augmented through its easy management.

Yet, actually seeing a benefit out of a virtualization implementation project means running that project correctly from the get go. Before you ever purchase hardware or software, or before you ever start clicking buttons in any virtual management interface, you must first identify what you want your end state to be. That process of defining the architecture ensures you don't end up with too many virtual servers attempting to use too few resources. It also ensures that you've purchased the right amount of hardware to support your virtual machines, their failover needs, and your expected level of expansion.

To Outsource or Not to Outsource

Going through the process of defining your architecture, validating any equipment you plan to purchase, and ultimately getting virtualization installed can be a big task for a small business or environment. Although actually clicking the buttons to install the software is easy, some of the surrounding tasks might require extra effort. As such, you may consider bringing in consultants or working with a Value Added Reseller (VAR) for some portion of your virtualization deployment. This can be a good thing, but it also adds expense.

Consider carefully the decision to bring in outside help. Virtualization—especially its no-cost products—are specifically designed to be easy to deploy, even for technology generalists. Your decision on outsourcing will most likely be affected by the size and complexity of your desired environment.

Since many virtualization solutions are mature technologies, enjoying many years of successful implementations in businesses of all sizes, a set of best practices has evolved that helps ensure success. This chapter will finish the discussion by helping you understand the right way to plan and implement the rollout of your virtual environment.

The Life Cycle of a Virtualization Implementation

It can be argued that every virtualization implementation project has a life cycle just like the software itself. Irrespective of size, virtualization implementations tend to go through a series of phases. Latter phases build upon what was learned and/or implemented in the previous phase towards an expanded level of use. Figure 4.1 shows one representation of that life cycle.

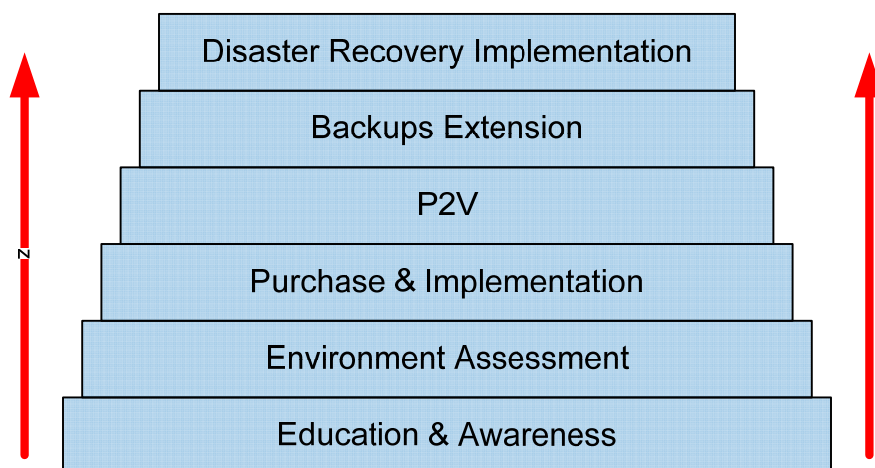


Figure 4.1: Six phases that are commonly associated with a virtualization implementation.

In this image, you see a number of phases, starting at the bottom with the all-too-critical education process. It is upon this platform of education that the other phases lie, and from which their data and information is based. Important also in this education phase is recognition of the level of hype that exists in the virtualization ecosystem. This recognition includes an understanding of the solutions and features that are necessary for your particular needs.

The Importance of Education and Awareness

Recognizing what you need in this first phase is critically important to ensuring the success of your project. It has been discussed a number of times in this guide that technologies and features exist in today's virtualization products that are justifiably exciting. These advanced features provide the highest levels of uptime for virtual machines. They enable fantastic levels of load balancing and failover in cases of server downtime while protecting running virtual machines from planned host outages.

Yet at the same time, these added technologies might not necessarily be needed by all environments or might not be affordable by all budgets. One of the necessary activities you must undergo during this phase is in understanding your requirements as well as your budget.

You can inexpensively implement a virtualization environment today that provides basic functionality and reasonable resiliency. Such an environment can be built using existing equipment and no-cost virtualization software, and will provide many of the financial benefits you've learned about in this guide. Once you're ready for the more high-level benefits and features, "upgrading" to for-cost solutions is often a simple process.

Assessing Your Environment

Once your implementation team fully understands the products and features available, the next step in the rollout is to fully understand your existing environment. In this way, the term "assessment" is multifaceted. It involves looking at your current inventory of hardware and OS installations with an eye towards which workloads will work into your planned virtual environment. Such an inventory is necessary because not all workloads work well when collocated with others atop a virtual host. As will be discussed shortly, a look at server performance is critically necessary to discover which physical servers make good virtual servers.

At a minimum, your virtualization assessment should include the following:

- **An inventory of the physical servers in your environment.** This inventory should include the characteristics of the physical servers being used, including installed memory, processors, and any attached peripherals. Information about networking and network domains to which servers are connected is also vitally important.
- **A quantitative performance analysis of each server.** This performance analysis should include performance metrics that make sense for the server in question as well as its workload. These performance metrics are possibly the most critical data you can collect, as they will document the prevailing behaviors on that server over a period of time. A listing of some sample counters as well as suggestions for how they are collected will be discussed in the following section to get you started.
- **An analysis of the hardware needs to support virtualized servers.** With the right level of performance information in hand, it will be possible to see what level of virtual machine collocation you can expect atop your physical hosts. In this case, and based on the performance information you have, you should be able to determine how many physical servers are required to support your virtual server needs. Although this step seems like a lot of work, it will assist with planning the hardware costs for your project.
- **An analysis of failover needs.** Completing the immediately previous step only gets you partway to your goal. Being successful with your virtualization implementation means also planning for failure states. For example, if you know that three virtual hosts can support ten virtual machines each, purchasing those three gives your environment the support for a total of thirty concurrent virtual machines. However, if any of those virtual hosts experiences a failure (see Figure 4.2), the failover of virtual machines will have a negative impact on the performance of your entire IT environment. This is the case because in such a failure state, a virtual host must now support more than its maximum of ten virtual machines. To prevent this, your virtualization implementation should include the added hardware capacity to support the loss of any single component.

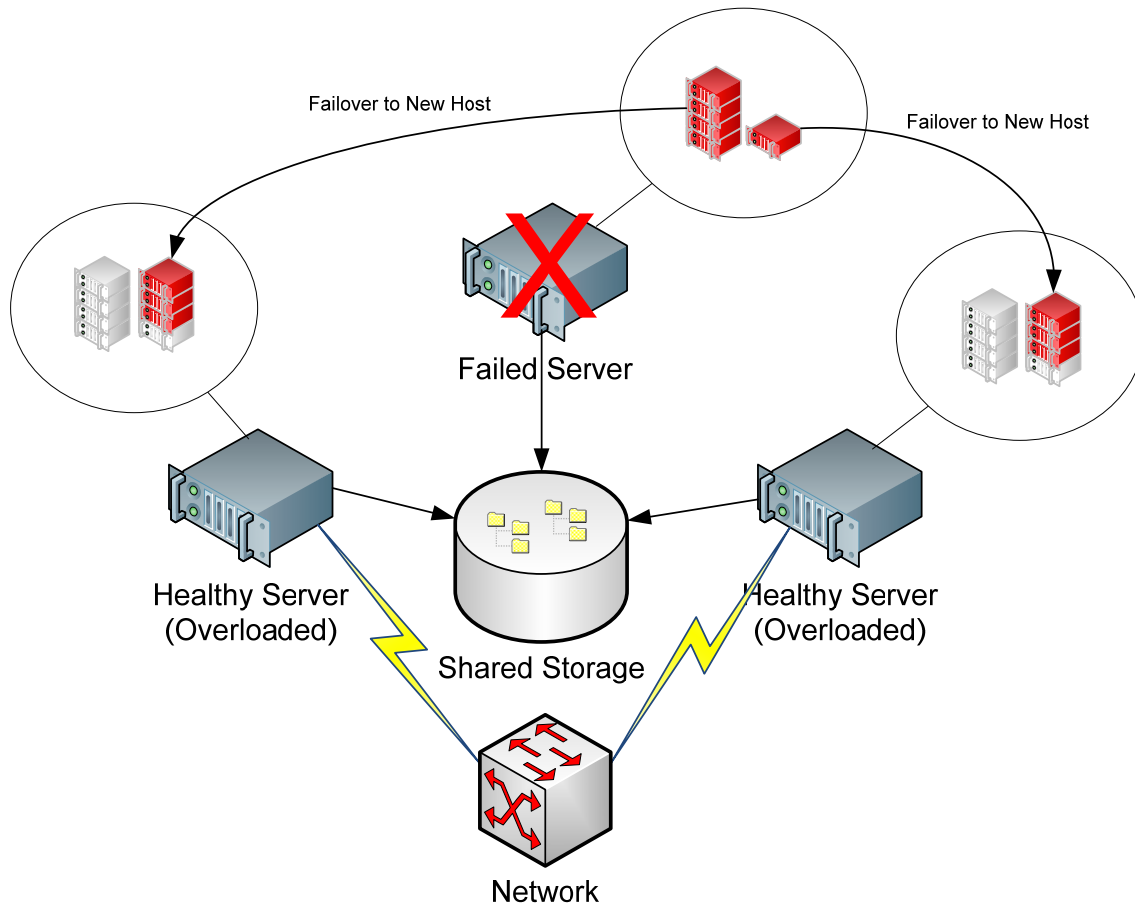


Figure 4.2: Environments that do not properly prepare for failover can cause healthy servers to be overloaded when one server fails.

- Planning for near-term expansion.** Environments that move to virtualization immediately gain the ability to rapidly spin up new servers and services when necessary. Yet this capability can be a double-edged sword when not properly managed. Environments that don't manage the expansion of virtual machines quickly find themselves awash with far too many. Your planning should include the extra hardware and software necessary to support some measure of near-term expansion. At the same time, you should consider planning to ensure that the expansion doesn't get out of control.
- A look at initial hosting and collocation for virtual machines.** This step is particularly important for environments that elect not to implement automated failover technologies with their initial implementation. Without the ability to fail over virtual machines from one host to another with no loss of service, environments must be careful about where virtual machines are positioned. Virtual machines with high resource requirements should not be collocated with others on the same host.

Your Level of Planning Will Vary

You needn't be overwhelmed by the items in this list. The suggestions here are intended to give you a comprehensive look at all the planning that can be done for larger-scale rollouts. If your environment takes a more measured approach, your level of pre-deployment planning will be dramatically simpler. No matter what your size of initial deployment, however, pay careful attention to the performance measurement information in the next two sections.

Measuring Candidate Performance with Native Tools

With each of these ideas in your plan, the next step is involved with measuring the performance of those servers you may virtualize. This process is critical because there are a limited amount of resources available on any virtual host. The virtual machines running atop that host must share its available resources for their processing. As such, servers with high resource requirements often don't make good candidates for virtualization when consolidating many servers onto few is your goal.

Back in the early days of virtualization, many IT professionals took an application-centric approach to virtualization candidacy. During those times, you would hear or read that *"Exchange servers shouldn't be virtualized,"* or that *"You should never virtualize a SQL box,"* or even that *"Citrix and Terminal Servers don't make good candidates."* However, this early focus on the applications that are installed to candidate servers was flawed. As an example, an Exchange server that serves an extremely large number of clients with massive mailbox sizes and extremely high network traffic might not be a good candidate. Yet the same server that services relatively few clients with small mailboxes and light network requirements might actually be a good fit. Thus, focusing on what is installed to your servers isn't the right approach.

In the end, that early application focus gave way to a more quantitative approach focused on actual performance measurements. Think about the servers you have in your environment today. Do you have performance counters enabled on those servers? If not, how do you know what types and level of processing they are doing during the day? Do you know whether they're particularly busy servers or relatively lightly-used servers? Without this information, it is difficult to understand the true usage of your applications and the physical resources you've assigned to the servers that host those applications.

In the Microsoft Windows operating system (OS) the primary mechanism for identifying the level of resource use on a server is the Performance Monitor (PerfMon) tool, shown in Figure 4.3. This service measures counters across a system for metrics such as processor and memory use. It digs deep into resource usage across extremely specific counters for all kinds of system and application measurements.

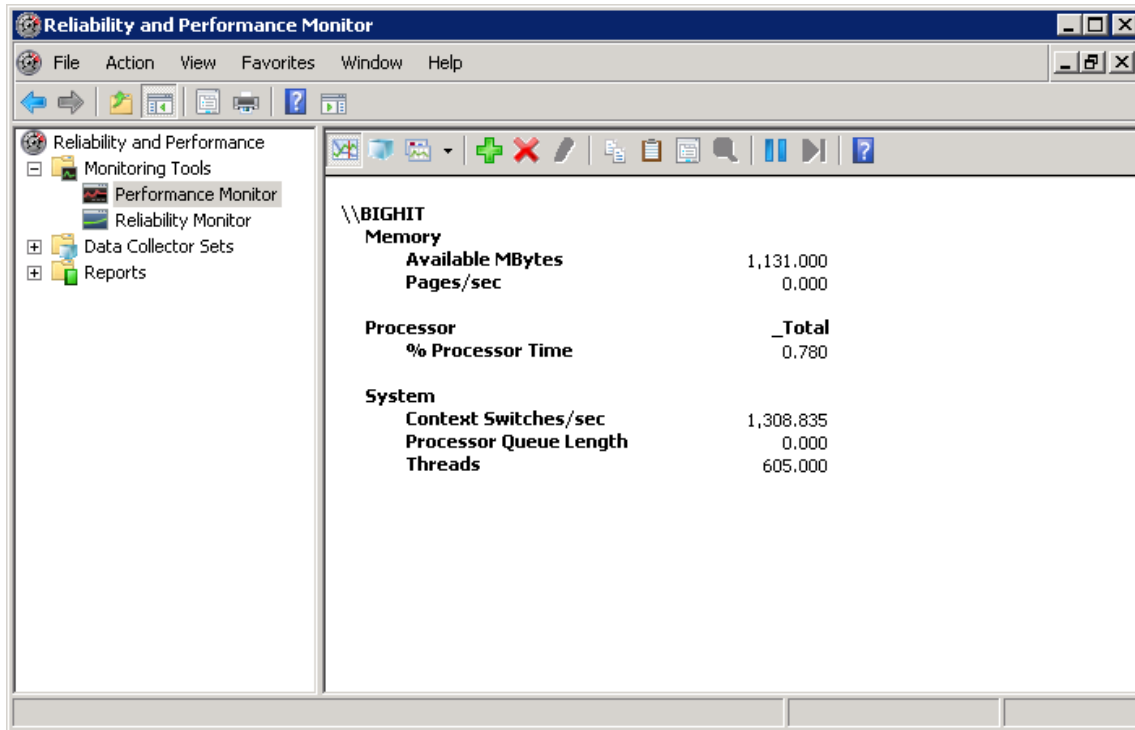


Figure 4.3: PerfMon is a tool in Microsoft Windows for measuring performance across a computer.

Yet, PerfMon's comprehensive list of available counters also brings a level of confusion to the uninitiated. How do you actually use this tool—or others like it—to measure server performance? With PerfMon's hundreds or thousands of counters, which ones are important to watch?

Table 4.1 shows an example of PerfMon counters that may help get you started. The thresholds in this table are not intended to be hard-and-fast rules for definitively identifying candidates and non-candidates. Instead, they are listed as example threshold values that should pique your attention towards servers that might not be the best candidates. These counters should be set up for a candidate system and run over an extended period of time—perhaps a few days or a week—to get a long-term sense of the level of resource use for a candidate server.

<i>Category</i>	<i>Performance Metric</i>	<i>Example Threshold</i>
Physical Disk	% Disk Time	> 50%
Memory	Available MBytes	Below Baseline
Memory	Pages / Sec	> 20
Paging File	% Usage	> 70%
Physical Disk	Current Disk Queue Length	> 2 * Number of Spindles
Processor	% Processor Time	> 40%
System	Processor Queue Length	> 2 * Number of Processors
System	Context Switches / sec	> 5000
System	Threads	> 2000

Table 4.1: A table of sample counters and example thresholds.

Leveraging Easier Performance Toolsets

With all this in mind, you can see how this process of measuring candidate performance is somewhat challenging using native tools such as PerfMon. You will need to set up performance monitoring individually across multiple servers, monitor that configuration over time, and ultimately pull the results into a useable format for later reporting. You will need to know what counters to watch and what results actually have meaning for your environment. Adding to the difficulty is that the native Windows PerfMon tool is notoriously difficult to use. It has a tendency to create multiple log files with each reboot that must be later reassembled. Its reporting capabilities are ridiculously challenging to use.

Because this pre-implementation performance assessment is so critical to virtualization success—and so natively hard to accomplish—many virtual platform vendors now provide free tools that assist. These tools (see Figure 4.4) gather information from native services such as Windows PerfMon across multiple machines at once. They provide a more streamlined approach to performance metrics gathering and reporting. More importantly, these tools tend to be preconfigured with known performance characteristics of candidate servers as well as the underlying logic that provides you with quick answers regarding which servers might work well together once virtualized. The result is that these tools provide much easier heads-up reporting that assists you in understanding which candidate servers will and won't work well in your virtual environment.

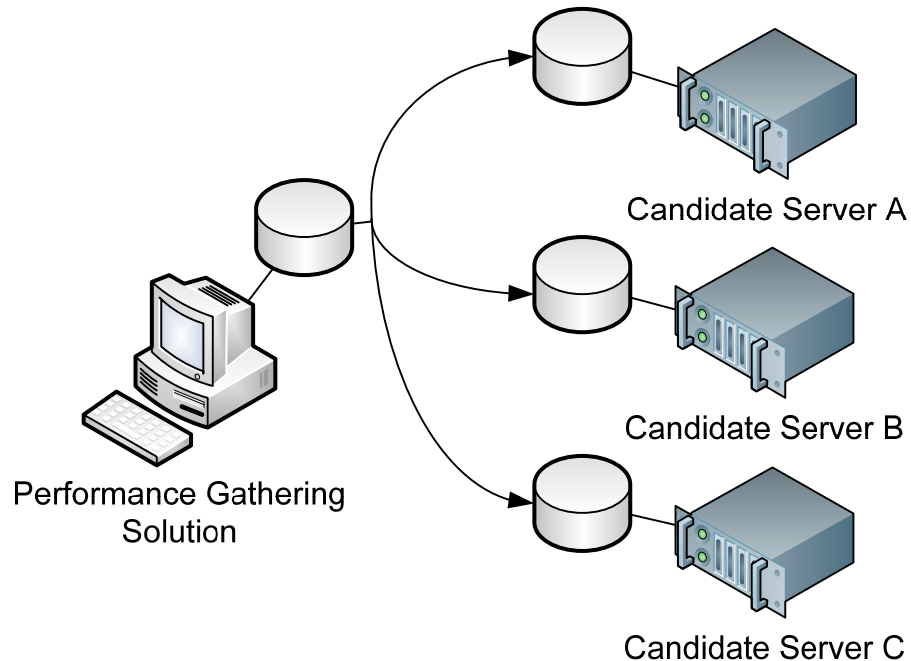


Figure 4.4: Some virtualization platforms include software that eases the challenge of gathering quantitative performance data.

Prior to starting any virtualization implementation, it is critical to understand how potential virtual machines may work with each other. The use of these tools ensures that your completed implementation isn't resource restricted, a situation that will reduce your overall satisfaction with the result.

The Importance of Assigned Virtual RAM

Another part of your assessment process is the determination of which resources should be assigned to a server once virtualized. These resources can relate to what kinds of external drives are attached or to which networks the server should be connected.

One very important determination here is the decision about how much virtual RAM to assign to the virtual machine. Virtual environments enjoy the flexibility to assign very granular levels of RAM to virtual machines, with the added ability to adjust that assigned level on the fly. This is much different than in physical machines where adding physical RAM requires a hardware purchase and can require a complete reconfiguration of the hardware itself. Virtual environments are quite different. In a virtual environment, it is possible to very specifically assign a particular amount of RAM to a virtual machine. Changing that level at a later point is as simple as a few clicks and a virtual machine reboot.

Yet history has a role to play in this decision. Administrators have long been forced to deal with challenges in augmenting the level of RAM in a physical machine. Changing the level of RAM in a physical machine is hard. As a result, many administrators have historically overloaded their physical servers with the maximum amount of RAM, notwithstanding how much RAM that server's applications might need. As a result, many physical servers today are working with vast quantities of RAM memory that are simply not being used.

Virtual environments are much different. With many virtual machines residing atop a host, administrators can no longer afford the luxury of assigning lots of RAM to each virtual machine irrespective of its actual needs. This is the case because the level of RAM on a virtual host is a fixed resource. Assigning too much RAM to one virtual machine reduces the amount that others on the same host can use. The net result is a reduction in the total number of virtual machines that a single host can support, and an overall reduction in the value you get from your virtualization investment.

As such, a best practice is to determine and assign the lowest acceptable level of virtual RAM that a server will require in the processing of its mission. Determining this level of RAM is another activity that can be done using native performance monitoring tools such as PerfMon, or through the performance-gathering tools provided by your virtual platform vendor. When considering the level of RAM you plan to assign to your virtual machines, look first at their resource use prior to being virtualized. Over a long-term measurement, memory counters such as the *Available MBytes* counter shown in Table 4.1 will help you identify the level of RAM that is actually being consumed by the system. Once you recognize the actual amount of RAM that system requires, consider this as a starting point for its assigned virtual RAM.

Purchase & Implementation

Once you've completed the assessment of your environment, you're ready to begin looking at the virtualization solutions available on the market today. As has been discussed in earlier chapters, that analysis will look at various solutions that may or may not involve the use of no-cost virtualization software. Most virtualization vendors today provide multiple tiers of software, with each tier providing added capabilities over those in the tier below.

Pay special attention to the information in Chapter 2 to help you identify which features you need and which you can save for later augmentation. Also, be aware of each vendor's upgrade path between tiers. Some virtual platforms can enable higher-level tiers without requiring a wholesale replacement of the existing infrastructure. If your environment is considering a measured approach to virtualization, starting small and working your way up, a trivial upgrade path between platform tiers will make the decision to expand much easier down the road.

Converting Physical Servers to Virtual with P2V

During virtualization's early years, once the virtual platform was up and running, administrators' next step was to rebuild servers in the virtual environment. This process was challenging and involved a large amount of time, as no automated mechanism existed for converting physical servers to virtual ones. Since that time, nearly every virtualization vendor has incorporated the tools necessary to make this automated conversion possible. The process to complete this automated conversion, explained first in Chapter 1, is commonly referred to as P2V.

From a high level, the P2V process is functionally similar to the process used to image and rapidly deploy workstations. In the workstation analogy, the imaging process analyzes a source computer and creates a file in which is stored the configuration of that computer. That file is later used by the same tool to deploy the source computer's configuration to one or more other computers.




Figure 4.5: The P2V process automatically converts a physical computer to a virtual machine.

The P2V process is similar to this process in that P2V tools look at a source computer to gather its configuration (see Figure 4.5). That configuration is, however, transferred not to another physical computer. It is instead transferred into a virtual disk file that resides within your virtual environment. Where the P2V process differs from the traditional imaging process is that during the conversion, the P2V tools also change the set of drivers installed into that computer. The existing drivers that work with the physical server are replaced by new drivers that work with the virtual environment. The result when complete is that the source computer now resides also in the virtual environment and is fully ready for use once the process completes.

Although the original physical server is unchanged during the P2V process, be aware that a conversion is a one-directional action using most toolsets. Thus, once a server has been virtualized using a P2V tool, it is a best practice to power down and decommission or repurpose the original physical server. Both copies of the server configuration should not run in your environment at the same time. As such, a good practice after P2Ving a server is to verify that the conversion completed successfully and that all the necessary data is correct on the server's virtual instance before proceeding.

Extending Backup Capabilities

Traditional server backups have always been a pain in the neck. Ensuring that each and every file and folder is correctly backed up among the tens of thousands on a server is a challenging job. Replicating that process across each of the servers in your data center virtually ensures that some are going to get missed in the process. Yet missing even a single file or folder can be a major problem if that file needs restoring. Missing that file during a server restore can fail the entire restoration process.

It is for this reason many organizations enjoy the benefits to backups that virtualization provides. As you can see in Figure 4.6, virtual machines on a virtual host have multiple ways in which they can be backed up.

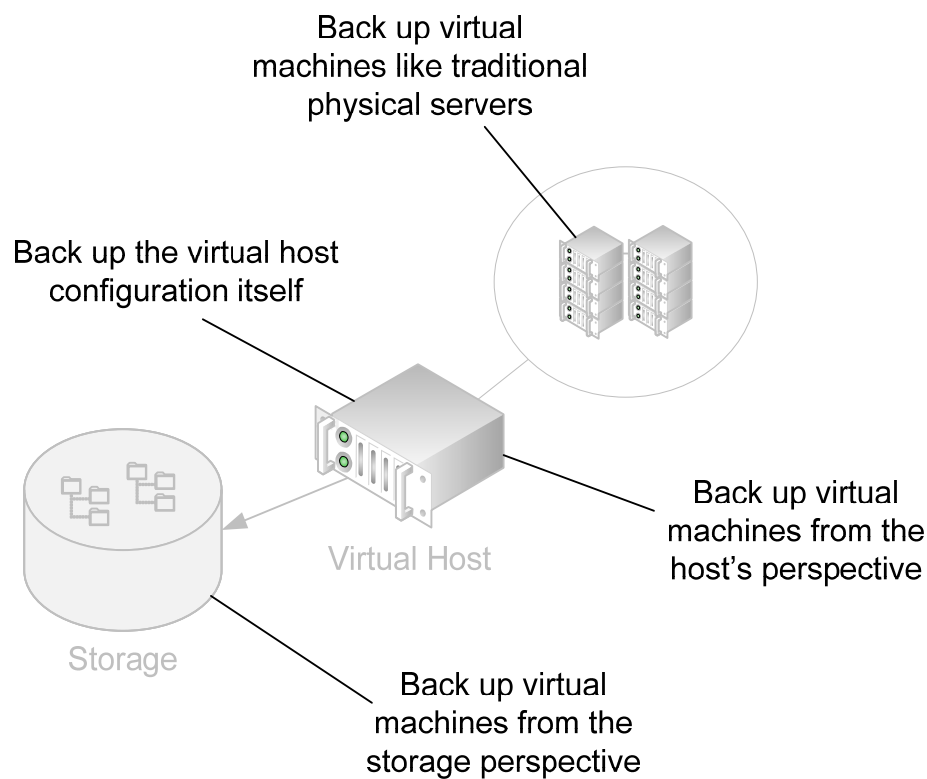


Figure 4.6: Backups in a virtual environment can occur from multiple perspectives.

Depending on the virtual platform you choose as well as the capabilities of your hardware, you may elect to choose one or a combination of these mechanisms for completing your backups:

- **Virtual machines can be backed up like traditional physical servers.** In this configuration, virtual machines are backed up just like their physical brethren. Accomplishing this task requires the same backup client to be installed to the virtual machine as is installed to traditional physical servers. Performing backups from this perspective maintains the same backup architecture and restore capabilities as with other physical computers but gains none of the added benefits enjoyed by virtualized computers. Additionally, the processing of the backup software can consume large levels of resources on the host, reducing the performance of other virtual machines during backups.
- **Virtual machines can be backed up from the perspective of the host.** In this configuration, the virtual disk file for a virtual machine is backed up at once. Here, the resulting backup file is a single file that represents the entire contents of the server's hard disk. As a result, restoring that entire virtual machine's configuration requires restoring only that single file. This dramatically improves the effectiveness of backups for entire-machine protection. Be aware, however, that this architecture may increase the effort required to complete an individual file restore, as that file will need to be extracted from the virtual disk. Using this architecture improves the performance of backups over traditional architecture because it requires no client to be installed into the virtual machine and its processing can be managed at the level of the virtual host.
- **Virtual machines can be backed up from the storage perspective.** Environments that leverage the use of segregated storage such as SAN and NAS enclosures add a perspective from which virtual machines can be backed up. In this configuration, the SAN or NAS enclosure itself can be configured to back up the virtual machine disk files on behalf of the virtual host. Offloading the processing of backups to the storage eliminates its performance impact on the virtual host itself.
- **The virtual host configuration itself can be backed up.** In addition to backing up the virtual machines on a host, that host's configuration can and should be backed up as well. Doing so ensures that individual files can be restored in the case of a loss, and provides a way to restore the virtual host in the case of a loss.

The Problem with Transactional Databases

Getting good backups of your servers' file systems is one thing. Capturing a successful backup of a running application is quite another. There is a known issue with some virtual backup solutions that back up virtual machines from the perspective of the host. These backup solutions do not have a mechanism for instructing transactional databases—such as SQL, Exchange, or Active Directory (AD) databases, for example—to correctly pause themselves for the purpose of the backup. Without the right pausing of databases during the backup, it is possible that the restore of a virtual disk may return the database in an inconsistent state.

The technical term for this pausing is called “quiescence” and is a necessary step for databases like the ones discussed earlier to be immediately available upon a virtual disk restore. For the Microsoft Windows OS, the most commonly used mechanism today for ensuring quiescence is the use of the Volume Shadow Copies (VSS) service in the virtual machine's OS. If you plan to virtualize workloads that make use of transactional databases and plan to use backups from the host's perspective, make sure that your backup solution supports the integration with the VSS service in any Windows-based virtual machines.

Adding Affordable Disaster Recovery

The final phase in this chapter's representation of the virtualization implementation life cycle is one that many businesses realize once they've made the move to virtualization. Specifically, smart organizations quickly realize that virtualization's augmentation to the architecture of backups readily enables new disaster recovery options.

To best explain how these disaster recovery options are made available through virtualization, consider first how disaster recovery and planning is done in physical environments. In these environments, the first step in creating a secondary site is to bring that site online. Bringing that site online requires high-speed network connectivity between primary and secondary sites. It also requires the purchase of hardware that is exactly the same between the primary and secondary site. On that hardware is needed a configuration that is exactly the same as well between the primary and backup sites. This precise level of similarity is necessary so that all the functionality and data from the primary site can still be used after the migration to post-disaster operations.

Yet, this requirement in an all-physical environment is operationally challenging. Creating two servers that are functionally equivalent requires double the hardware and the effort. Maintaining that exact similarity over time is an operational headache as well, with every configuration change needing replication to the alternate site. Software packages are available that assist in automatically replicating the configuration changes, but such a site still requires an effective doubling of hardware, software, and other infrastructure elements.

Virtualization changes the game significantly in part due to the way it handles backups of virtual machines. One of the options for backups mentioned in the previous section was for virtual disk files to be backed up from the perspective of the host. This single-file backup of an entire virtual machine quickly and easily captures the exact state of that virtual machine at the time of its backup. Also in that section it was stated that restoring such a virtual machine required little more than restoring its single file to a virtual host.

Now, frame these same capabilities around the restoration of that virtual machine instead at a secondary site. Once a business has their backup solution in place and virtual disk files are being backed up in easy-to-restore formats, the only remaining action left to enable disaster recovery is the replication of those disk files to an alternate location. Figure 4.7 shows a graphical representation of how this might work.

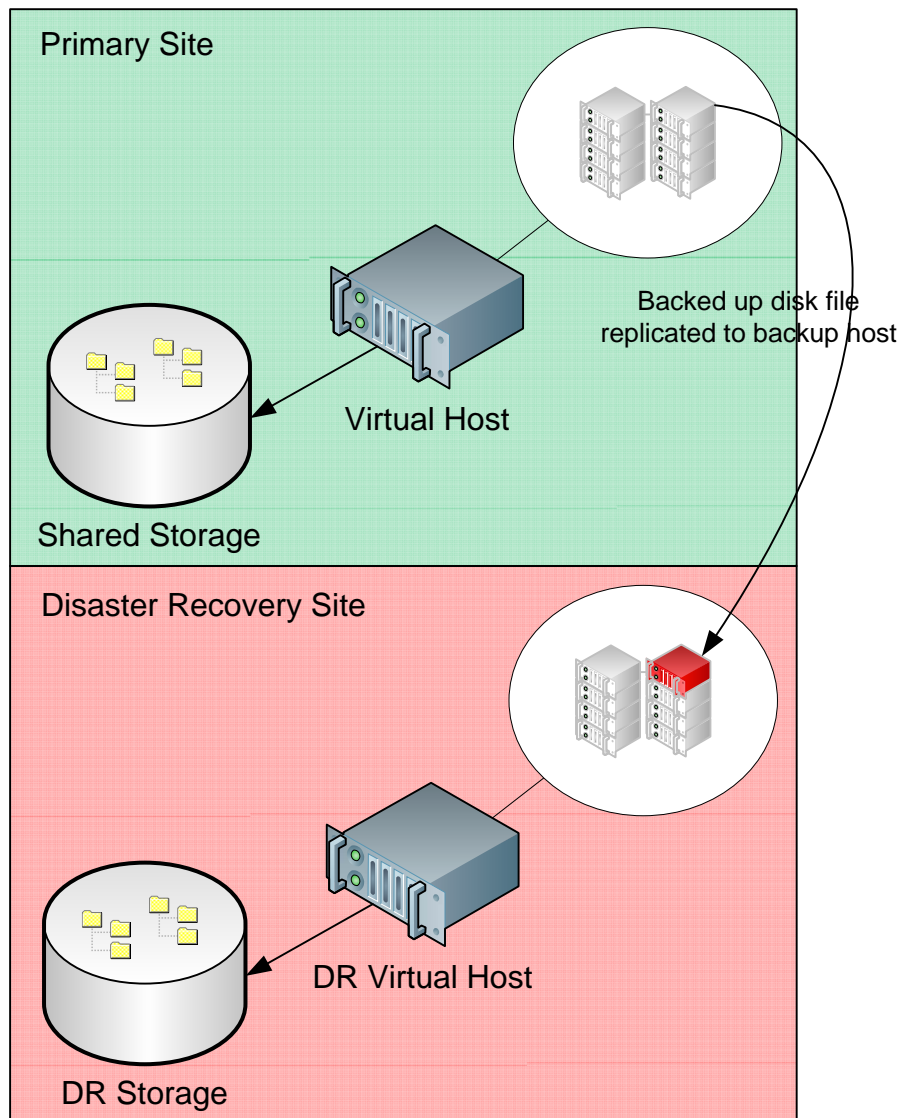


Figure 4.7: Replicating virtual machine disk files to an alternate location provides a cost-effective disaster recovery solution.

The process to replicate these backups from a primary to a disaster recovery site can occur through an automated mechanism—either as a component of the backup or through a dedicated replication solution—or can be manually transferred as necessary. Your level of automation required in the replication is directly related to your requirements for returning to full operations after a disaster event. Also a factor in this architecture is the level of granularity you require for your backed up data. Businesses with little tolerance for lost data will want solutions that back up virtual machines more-or-less continuously. Other businesses with greater levels of tolerance can leverage solutions that transfer backups on a scheduled or irregular basis.

Implementing Virtualization Involves Technology and a Plan

Again, all the technical information in this and the previous chapter are written to give you a holistic view of the options available in making the jump to virtualization. Your individual needs may require only a single server and a few virtual machines, a situation that reduces your level of required planning to begin an implementation. Projects with larger scopes and more complex needs will want to include each of the planning steps and best practices outlined in this chapter.

No matter your needs, this guide has attempted to show you how the jump to a virtualized infrastructure is a good thing for small businesses and environments. It increases your level of IT agility while at the same time reducing your overall operating costs. It provides levels of hard benefit to your business not seen in most IT projects. Above all, virtualization enables dramatically greater levels of flexibility for the processing needs of your business. With the right information in hand, your job now is to determine whether and where it fits in your overall business plans.

Download Additional eBooks 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 eBook to be informative, we encourage you to download more of our industry-leading technology eBooks and video guides at Realtime Nexus. Please visit <http://nexus.realtimepublishers.com>.