Realtime publishers
"Leading the Conversation"

The Shortcut Guide™ To

Selecting the Right Virtualization Solution

sponsored by

Parallels™

Greg Shields

Copyright Statement

Chapter 4: Managing Virtualization Environments

The total cost of ownership of any virtualization solution is driven by its recurring management costs. Greater automation within a virtualization system will reduce that recurring cost, ultimately reducing its overall cost of ownership. This guide has attempted to take a "soup to nuts" approach to helping you determine the best virtualization solution for your computing environment.

As we've discussed throughout this guide, there are a number of virtualization solutions available on the market. Each of those solutions has a best-fit within the various types of workloads needed by any particular IT department. Heterogeneous environments with support requirements for numerous operating systems (OSs) might benefit from the emulation capabilities of *Hardware Virtualization* and VMware Virtual Infrastructure. Alternatively, environments that support a single OS but desire the highest possible consolidation density might prefer *OS Virtualization* and Parallels Virtuozzo Containers. The intent of this guide has been to illustrate where your specific needs for virtualization fit best within the virtualization solutions available today.

Chapters 1 and 2 discussed a few of the architectures currently on the market. Chapter 1 focused heavily on the differences between architectures, while Chapter 2 featured a comparison between the products that enjoy those architectures. In that comparison, we discussed the similarities and differences between VMware Virtual Infrastructure's and Microsoft Virtual Server's *Hardware Virtualization* architecture with Citrix XenSource's *Paravirtualization* and Parallels Virtuozzo Containers's *OS Virtualization* architectures.

Chapter 3 focused on the best practices of implementing virtualization within the computing environment. In that chapter, we talked at length about potential uses of virtualization and which architectures tend to best support those uses. This discussion is necessary because obtaining the software is only the first step. Building the environment correctly from the beginning ensures the best possible initial return on your virtualization investment.

As is stated earlier, to get the best overall return, you must manage your chosen environment using best practices. Those best practices involve the use of automation to enhance the efficacy of available administrator personnel. Depending on the virtualization solution chosen for your environment, the management tools that are part of that solution will include automation support. This chapter will talk about the best ways to make use of those management capabilities to achieve the lowest overall cost of ownership and greatest return on your virtualization investment.

Key Concepts in Virtualization Management

We've talked before about how integrating virtualization into your computing environment adds a layer of complexity while at the same time enabling additional functionality. As virtualization segregates workloads between virtual machines and virtual hosts, two classes of servers now need to be considered. Each class is managed differently than the other.

Virtual hosts are managed as collections of resources. Those resources are to be distributed to residing virtual machines. Virtual machines are managed similarly to their all-physical analogs but with added functionality driven by the capabilities of the virtualization environment. Depending on the virtualization solution chosen, the *layer of abstraction* associated with virtualization can be felt all the way down to individual process, service, log, and network connection management. This occurs because virtualization enables administrators to work with groups of machines as a unit.

The key take-away from this section, and indeed this chapter, is that the addition of virtualization to an environment removes the traditional boundaries associated with individual *one server per chassis* management. This concept is represented visually in Figure 4.1. With virtualization and its associated management software, the hard lines between individual servers blur. Individual host machines become less like boxes that sit in the data center and more like collections of processor, memory, disk, and network resources that are granularly assigned to virtual machines as needed. Let's talk now about some of the elements necessary when considering these layers of management.

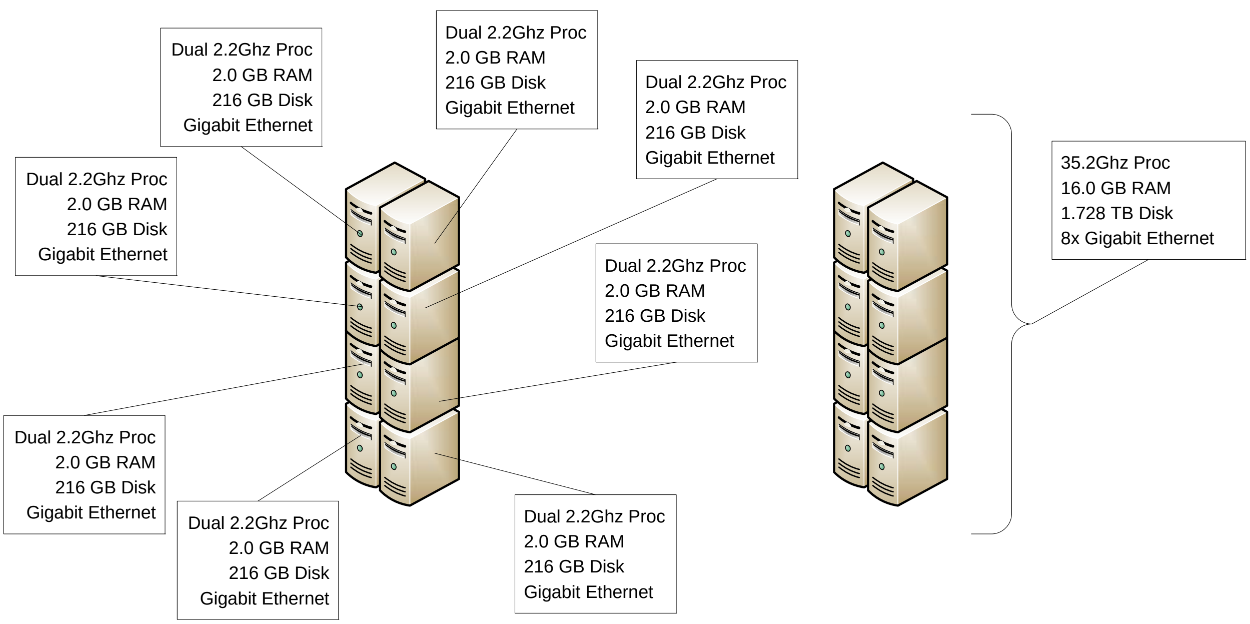

Figure 4.1: The figure on the left represents an all-physical environment, where one service is installed per chassis. Here, resources are managed on a per-server basis. The figure on the right represents an all-virtual environment. Here, resources are managed as collections across all machines and can be distributed on a just-in-time basis to needy network services.

Virtual Host Management

The servers that host the virtualization environment have special needs, different than traditional physical servers. Those needs are determined by the virtualization solution chosen. With environments that use VMware Virtual Infrastructure, the virtual host operating system is based on a Linux derivative. Thus, experience with this specialized operating system is necessary for troubleshooting and installation of any additional host software. For environments that use Parallels Virtuozzo Containers, the virtual host operating system is the same as its residing virtual machines. Thus, the same skills that are used in managing traditional workloads can be used in managing the virtual host.

Hosts in virtualization environments require other special attention as well. Consider the following special needs in managing your virtual hosts:

- *Virtual host configuration similarity*. Most virtual hosts include features that allow for the movement of virtual machines from one host to another. Because of this, the expected target host for a virtual machine migration should be as similar as possible to its source host. This helps to prevent conflicts or unexpected behavior on the part of the virtual machine that could occur due to a configuration mismatch.
- *Single service focus*. It is similarly important to consider virtual hosts as single-use servers. This is for a number of reasons. First, the responsibility of virtual hosts is to make resources available to virtual machines. Additional software running on the virtual host will consume resources that cannot be used by virtual machines. This is particularly the case with *OS Virtualization*. Due to the linkages between residing virtual machines and the virtual host with *OS Virtualization*, any software installed to a host in this environment can become a part of all residing machines, which can be an undesired result.
- *Backup and anti-virus software*. We'll discuss backup software in a minute, but be aware that backups in virtualization environments can be different than traditional backups. Backup software that supports both residing virtual machines as well as the virtual host can be necessary. Anti-virus software can be challenging within virtualization environments as well. The potential for conflicts with the host or residing virtual machines requires additional care to ensure full support. Additionally, the scanning processes associated with anti-virus software can be a resource drain. One of the benefits of *OS Virtualization's* focus on file sharing is an increase in the effectiveness and performance of antivirus software shared between the host and residing virtual machines.
- *Hardware similarity*. Depending on the virtualization solution chosen and feature sets required, hardware similarity may be a requirement. As an example, VMware Virtual Infrastructure's hot migration capability called VMotion can only function across certain processor classes. Processor architectures that are too dissimilar will negate the ability to successfully VMotion between hardware. Conversely, *OS Virtualization* solutions like Parallels Virtuozzo Containers virtualize at the OS level, which reduces its reliance on hardware.

Virtual Machine Management

Also necessary when managing a virtualization environment are the unique characteristics of virtual machines as well. Some traditional management and monitoring tools, such as Wake on LAN tools, or hardware vendor management tools may not behave in the same way once servers are virtualized. Additionally, the composition of simultaneously-running virtual machines on a single host can impact the user's experience when a single virtual machine consumes large quantities of physical resources.

Consider the following best practices in managing your virtual machines:

- *Performance management.* We'll discuss performance in detail in the next section, but know that the move to virtualization places a heavy requirement on vigilant performance management for the environment.
- *Virtual machine instance bloat.* IT best practices suggest the use of a single server per network service. However, in many physical environments this is not operationally feasible due to costs. The movement to virtualization eliminates this barrier by easing the creation of new servers. An unintended result of this is often an expansion in the total number of machines in the environment. Care must be taken to ensure that the number of virtual machines under management does not grow unnecessarily.
- *Licensing and activation.* Along these same lines, as virtual machines are replicated to expand services, care must also be taken to ensure that the number of machines and installed applications do not expand beyond the point of available licenses. Along these lines, with some operating systems licenses require an activation to be fully recognized. That activation may be lost during the replication process due to its tie into hardware identifiers unique to each virtual machine instance. If this occurs, the activation will need to be reacquired. This can be a problem because some software vendors limit the number of reactivations allowed for any particular license.
- *Machine uniqueness.* Operating systems have identifiers unique to each instance that must be globally unique for them to properly work together. For example, in Microsoft Windows the Security ID (SID) must be unique across all systems in an Active Directory (AD) domain. All virtualization solutions allow for the easy replication of instances, though not all automatically randomize these uniqueness components as part of that replication. Diligence is necessary to ensure that these characteristics are changed upon replication.

Centralized Process, Service, Log, and Network Management

The management toolsets enjoyed by some virtualization solutions also provide a level of abstraction over virtual machine processes, services, logs, and network connections as well. In this case, individual processes, services, and logs can be managed as a unit across all virtual machine instances. This provides a significant value to the administrator, giving them the ability to globally manage all instances from a single location. Parallels Virtuozzo Containers is one tool that can provide this level of element management.

Some examples of how this benefits the environment are:

- *Processes.* When all processes across all virtual machines as well as their resource use can be viewed through a single interface, it is possible through a single view to see where bottlenecks exist in the environment. Further, it is easy to restart or kill processes as necessary when their activities oversubscribe available resources.
- *Services.* Unified service management provides the administrator with a central location whereby services can be enabled, disabled, and otherwise configured across all virtual machines. This enhances the security profile of the environment while reducing the management overhead of manually configuring these services to the necessary baseline.
- *Logs.* When logs across multiple systems can be analyzed within a single view, it is more easily possible to isolate problems within the environment. Error logs can be lined up by timestamp across multiple machines to provide a whole-network view of how machines are interconnecting. For problems that occur amongst multiple machines, this substantially improves the troubleshooting process.
- *Network connections.* Virtualization aggregates network connections among all virtual machines on the host. A layer of logical networks is placed on top of available physical connections. This means that complex networking arrangements can be easily created and managed between virtual machines without the need for costly physical re-cabling.

Performance Management

A variation on Moore's Law suggests that hardware processing power doubles every two years. This means that every two years, we can expect the load carrying capabilities of new server hardware to be twice that of two years before. At the same time, the hardware requirements of server software have not necessarily kept up with this increase in capability. As an example, between 2003 and 2008 Microsoft Windows has released only two new server operating systems. Over this six year span, the hardware requirements of Windows Server 2008 are not eight times greater than for the original release of Windows Server 2003. Over the same six years, Moore's Law would suggest that hardware has gotten eight times better.

This dissonance between hardware capabilities and software needs has driven downward the productive usage of non-virtualized server hardware resources. As discussed in Chapter 3, the average utilization of a Microsoft Windows server across all industries is around 5%. This means that 95% of the time non-virtualized server hardware is performing non-productive work. These economies of performance are the basis for virtualization and consolidation.

A trend associated with this situation has been the reduction and/or elimination of performance management needs in most physical environments. As a matter of resource economics, in an environment with ample supply and little demand, effective performance management is unnecessary. Conversely, when moving to a virtualization environment, productive resource use can grow to 80% or more.

The point of this argument is that the resource economics associated with virtualization mean that performance management regains prominence as a critical activity of daily management. Performance management is fundamentally critical in virtualization environments to ensure a good user experience. This being said, let's talk now about a few of the needs associated with doing it effectively.

Counters to Manage

For the new administrator, one of the hardest parts of performance management is the identification of the right counters to manage. A typical Microsoft Windows server can monitor on hundreds of counters. Finding just the right ones that correctly tell the tale of a server's health is critical to maintaining a healthy virtualization environment. Let's take a look now at a short set of critical performance counters that are relevant to all virtualization environments. No matter what virtualization solution you choose, consider these a minimum set of counters to watch for all systems.

- *% Processor time.* The counter for % Processor time gives a view of the overall processor use within the server. % Processor time relates to the amount of time the processor is performing useful work and is not in the System Idle Process. As virtual machines all work concurrently on the host, the sum total of all virtual machine processor needs plus the needs of the host itself equal this metric. When individual virtual machines oversubscribe the processor or when individual processes spike processor use, this metric can rise to nearly 100%.
- *Processor queue length.* An individual system processor can only process a single instruction at a time. Instructions are swapped in and out of the processor rapidly to allow multiple requests to be completed with the illusion of simultaneity. The number of requests that are in-queue and waiting for the processor is measured through this counter. When this counter is high, then onboard processors are unable to keep up with resource requests from the host and virtual machines. Additional or faster processors may be necessary, or a reduction of virtual machine instances hosted by this server.
- *Context switches per second & System threads.* Related to the above metric, the metrics for system threads illustrates the number of instruction sets that are currently in work. An individual thread can have multiple instructions which the processor will need to accomplish based on its scheduling. Context switches are what occurs when the processor changes which thread upon which it is currently processing. High metrics for either of these two counters indicates also that the processor is unable to keep up with the requirements of its residing virtual machines.

- *Memory working set.* Similar in function to the counter above, but related to system RAM memory, this counter discusses the overall use of system memory on the host. The memory requirements of all residing virtual machines plus the needs of the host combine here to be represented by this counter. Residing virtual machines can require substantial memory per concurrent machine or only a small amount. The virtualization architecture used will have much to do with this. Virtual machines within a *Hardware Virtualization* architecture tend to share fewer elements within core memory than do those within *OS Virtualization*. Thus, more memory can be required to support the same number of virtual machines. Like with the % Processor time counter, it is necessary to watch this counter to ensure that virtual machines and their applications have enough memory to properly operate with minimal swapping.
- *Available megabytes.* In many ways, the inverse of the above counter, Available megabytes illustrates the amount of free memory available on the system. When this counter is low, then large levels of onboard memory are being consumed.
- *Disk % Free Space.* Within many architectures, it is possible to configure virtual machine disk files to grow as needed. Because disk files need not necessarily be created at full size, this feature allows for the creation of a greater number of disk files, each with a smaller size. However, because each disk file assumes that it has enough space to fill its configured maximum size, it is possible for those files to grow beyond the physical limits of the available disk. When this occurs, it can cause corruption of the virtual machine's disk file. Thus, it is necessary to monitor for available free disk space to ensure that the drive does not fill to capacity.

> ✐ Network monitoring counters can also be useful for monitoring, especially in situations where virtual machines make heavy use of the network. Though in most cases, with multiple gigabit Ethernet network cards installed into most modern server hardware, network utilization oversubscription is rare.

Depending on the virtualization solution chosen, the monitoring of some or all of these counters may be automatically done as a component of its management interface. Virtualization solutions with rich management interfaces that support the automated collecting of these counters make this process easy. As this monitoring needs to be done at both at the level of the host and within the individual virtual machine, consider this capability a critical need for your chosen solution.

> Be aware that with some virtualization architectures the monitoring of counters within the virtual machine can show incorrect results. For example, in a *Hardware Virtualization* environment it is improper for the individual virtual machine to monitor its own counters. Resource scheduling on the part of the host will cause time dilation problems that will skew the results.
>
> For *OS Virtualization* environments, individual virtual machine processes are components of the host. Thus, their processes can be more accurately monitored from within each virtual machine instance.
>
> In either case, the best mechanism for monitoring counters across multiple machines is via the virtualization solution's management interface. Consider centralized counter management as a valuable tool when selecting a virtualization solution.

Dynamic Resource Control

Now that we understand the nature of some potential counters of interest, it is important that the virtualization environment be configured in such a way to granularly control the assignment of resources. Depending on the virtualization architecture chosen, different capabilities are possible associated with the assignment of resources to virtual machines. As we discussed in Chapter 3, *Hardware Virtualization* solutions like VMware Virtual Infrastructure use emulated device drivers that must be coded to support a specific level of resources such as maximum number of processors or maximum amount of memory. Thus, within this architecture there will be limits to the amount of resources that the emulated device drivers can support. Using *OS Virtualization* solutions like Parallels Virtuozzo Containers, virtual machine drivers are components of the host, which means that resources can scale to the level of those on the host.

Also important with this are the needs to set resource levels on virtual machines as appropriate to ensure poorly-behaving applications do not oversubscribe resources on the host. Let's take a look at the types of resource levels that can be configured:

- *Processor*. In addition to assigning a particular number of processors to a virtual machine, it is often necessary to set a guaranteed minimum level of processor resources as well. That guaranteed level of processor resources, measured in GHz, ensures that a server under high resource contention will always provide a necessary level of resources to the virtual machine. Conversely, it may also be beneficial to set a limit of processor resources that the virtual machine can request. By setting a maximum level, the virtual machine is less likely to impact the operations of other virtual machines on the same host. A combination of maximums and minimums can be necessary to maintain the environment in situations of resource contention.
- *Memory*. Virtual machines are assigned a static quantity of RAM memory for which they can make use. Due to the positioning of the virtualization layer, that quantity is typically much higher for *Hardware Virtualization* architectures. This is due to the need to support the OS in addition to applications. With *OS Virtualization*, a lower count of assigned memory can be used. This is because that quantity is typically used only for the applications specific to the virtual machine rather than its applications in addition to the OS itself.

- *Disk quota.* Aligned with our discussion above about the complexities of expandable virtual disk files, monitoring and changing disk quotas as necessary is also necessary. One of the benefits of virtualization is the ability to easily expand virtual disks as required to support operations. Your virtualization solution's management tools should support the capability to change disk files on the fly as necessary to support the needs of your environment.
- *Processes and kernel level threads.* Limiting a virtual machine to a particular number of processes and kernel level threads protects the virtual host against excessive process creation or runaway processes. Setting limits here further protects the environment from resource overuse.
- *Network utilization and traffic shaping.* Lastly, in situations where multiple residing virtual machines consume high levels of network bandwidth, it may be necessary to limit their network usage. Network traffic shaping, also called network bandwidth management, is the idea of throttling the ability for certain network connections to make use of available bandwidth. Many virtualization solutions provide an ability to set limits on network usage through the management interface.

It is important to remember that any form of resource control can slow the performance of a virtual machine when under high load. When that virtual machine requires resources that are limited by the interface, it will take longer for it to accomplish its tasks. This is a necessary evil during times of high resource contention. Lacking these policies, an unmanaged virtual machine can take all available resources, leaving none for others collocated on the same host.

Right-Sizing Physical Hardware

Our discussion on resource management can't be complete without the consideration of right-sizing the virtual host. At first blush, many administrators might opt for "the best hardware money can buy" to support their virtualization environment. However, this "bigger is better" approach may do for resources in the virtualization environment as what was done in the previous physical environment. Namely, when all resources within the machine aren't *in balance*, the overall performance of the machine will be bound by one class of resources while others are underused.

Let's look at a simplistic example of this situation. Let's assume that an environment wants to move to virtualization as a solution for consolidating server instances. Currently they have 100 servers, all of which are dual processor 2.2 GHz servers with four gigabytes of RAM. A small number of those servers make heavy use of their processors and RAM, such as database servers, remote application servers, or email servers. A much larger number of servers, those such as AD domain controllers or patch management servers have relatively little resource use. The combination of these results is an average processor utilization of approximately 5% across all servers. Average memory utilization is approximately 1 gigabyte per server. The organization wants to consolidate all of these servers onto five virtual hosts. In order to do this, they purchase 5 dual-processor, dual-core 3.0GHz servers, each with 16 gigabytes of RAM.

In virtualizing their environment, they find that their resulting virtual servers are now memory-bound. This occurs because their processor utilization is well within their original limits. They were using 5% of 200 2.2GHz processors, which multiplies to a total of 22GHz of processor power. Their new environment supports 60GHz as the summation of all processors across all virtual hosts. However, their available memory is far short of the necessary amount of RAM. Their old environment required an average of 100GB of RAM, but their new environment only supports 80GB.

In the end, a better solution for this company would have been to purchase more RAM and less processing power to support the environment. Depending on the cost breakdown between dual-processor and quad-processor systems, it may have been cheaper overall to purchase a larger number of dual-processor systems, each with more RAM, than the composition they chose.

As was discussed in the beginning of this chapter, the movement to virtualization allows for the aggregation of resources across all available systems. Because of this, a pre-virtualization performance assessment of this example environment would have shown where the resulting hardware composition would be lacking.

> Some virtualization solutions include candidate performance monitoring and analysis tools natively within their interface. Others provide such tools as an add-on or third-party service. As a performance analysis is critically necessary prior to virtualizing any physical servers, consider these capabilities as a requirement when evaluating solutions.

Right-Sizing Storage and the Use of Server Templates

Another consideration when managing a virtualization environment are the needs of storage. Storage needs can be considered as much a management topic as an installation topic due to the ever-growing needs for storage within any environment. Across all industries, the need for additional storage means that storage augmentation projects are necessary at regular intervals to support the needs of the business.

The move to virtualization can significantly increase the amount of storage necessary for the computing environment. Depending on the solution chosen, the type of storage can change as well. Using the VMotion features of VMware Virtual Infrastructure’s Enterprise Edition, a form of shared storage is required such as iSCSI or Fibre Channel Storage Area Networks (SAN). These high-end storage devices can be costly in terms of dollars per managed gigabyte. The cost of interconnecting storage to servers is also high. Fibre Channel SAN’s require Host Bus Adapters to be installed into each storage target, usually in a redundant configuration. These costly HBA cards can increase the cost of a server by thousands of dollars. The special fibre cabling to connect server to storage is an additional cost, both in initial purchase as well as recurring maintenance as the environment changes over time. Though iSCSI SAN’s do not require special cabling, the cost of additional cabling runs in addition to the extra network cards required for iSCSI does increase storage costs over that of direct-attached storage.

One benefit of *OS Virtualization* solutions like Parallels Virtuozzo Containers is that the hosting of virtual machines does not necessarily require these more expensive classes of storage. With Virtuozzo Containers, the hot migration process can occur between servers equipped only with relatively inexpensive direct-attached storage.

An important characteristic of the storage requirements of virtualization architectures has to do with the level of file sharing between virtual machines and virtual hosts. In *Hardware Virtualization* architectures, each individual virtual machine is fully atomic in terms of the files and folders that make up its instance. This means that the hosting of 100 instances of a ten gigabyte operating system will consume a full terabyte in storage.

Using *OS Virtualization*, the use of server templates links each residing virtual machine to its host. This has the tendency to reduce each virtual machine's storage significantly. As an example, a freshly installed Virtuozzo Containers virtual machine residing on a Microsoft Windows host consumes only .17 gigabytes. Thus, the hosting of 100 of these instances would consume only 17 gigabytes

> This number will obviously get larger as each Virtuozzo Containers-hosted virtual machine goes through its operational lifecycle, is installed with applications, and begins logging and other disk-based operations. But it is an excellent example of the starting position one experiences with each virtualization solution.

Converting Machines

Once the installation of the environment is completed, your next step is usually getting physical machines converted over to the virtual environment. This process, commonly called *P2V* or "physical to virtual", can be an extensive process. With some virtualization solutions, the full use of P2V tools can be limited by software packages, licensing agreements, or editions. Other platforms include P2V tools natively within the management interface. As the process of migrating machines from the physical world to the virtual one requires these tools, inclusion of these tools with your solution purchase is usually necessary.

In addition to the P2V process, in this section we'll also talk about the V2V and V2P processes. In the case of V2V, you are effectively cloning an existing virtual machine to another instance. With V2P, the migration process is reversed, taking a virtual machine and converting it back to a physical instance.

> When considering a virtualization product, look for those that include tools for monitoring and analyzing existing machines with an eye towards importing them into the virtual environment in the best way possible. These tools help eliminate the guesswork of positioning new virtual machines.

P2V

As the P2V process is commonly done during the early stages of a virtualization implementation, these tools are critical to ensuring an initial success. Typically, once the initial batch of physical machines is converted, P2V tools grow less used. Though, the tools come in handy later if additional servers need to be added to the environment.

For most tools, the process of completing a P2V activity is very simple. With most, the process involves pointing the tool to the source server to be virtualized, pointing it also to the target virtual host, and starting the process. The P2V tool copies the files, registry keys, and all configurations—in much the same way as a server imaging operation—to the target virtual host. Different from a server imaging process, the P2V tool also completes a series of internal conversions of the physical machine to enable its support for the environment. Those internal conversions can include changing device drivers, converting files to function with server templates, and changing its file system to work with the virtualization solution's file system.

> Lacking P2V tools, new machine instances with similar configurations and data would need to be created from scratch. Though the P2V process, depending on the tool used, can take many hours to complete, its process is highly automated. Completing a P2V usually takes much less time than rebuilding the server completely.

Be aware that P2V operations can involve a heavy use of resources on the target virtual server. Because of this resource use, the process of completing the P2V operation can impact the operations of virtual servers already on the host.

V2V

We've talked in previous chapters about the "cut-and paste" properties of virtual machines. The ability to clone or otherwise copy a virtual machine in whole is a major benefit associated with the move to virtualization. As this process is regularly used during the nominal administration of the virtual environment, be sure to consider management toolsets that include rich support for copying virtual machines.

V2P

Occasionally the need arrives to de-virtualize a server. Often this occurs because a server was virtualized that does not work well with others in the environment. Perhaps that server uses too many resources in comparison with others. Or, software vendors require the server to operate within a physical instance to be fully supported. In any of these cases, as with the P2V process, it is usually easier for an automated toolset to revert the server back to a physical platform rather than rebuild the server from scratch.

V2P tools are not typically included natively within most virtualization solutions. Though for many they are available as third-party add-on toolsets. These tools perform much of the same processes as P2V tools, but have the added requirement of discovering and installing the correct physical device drivers to the migrated machine appropriate for the target physical server.

> One excellent use of V2P tools is in the case where a virtualized server is not supported by an application vendor. Think about the situation where an organization wants to keep the server virtualized, but also wants vendor support when problems occur. Here, the organization can use V2P tools to create a temporary physical instance of the server for situations where vendor troubleshooting support is needed. Once the problem is identified and fixed on the temporary physical instance, its resolution can be replicated to the production virtual instance.

Server Backups and Disaster Recovery

The move to virtualization also changes elements within traditional backups. Due to how virtualization encapsulates physical machines into virtual machines, it provides new ways to backup and restore those servers. At the same time, depending on the virtualization solution chosen, it can also complicate the architecture of backups.

Along with this expansion of backup capabilities comes a set of new disaster recovery capabilities as well. Historically, disaster recovery operations have been operationally unfeasible for all but the largest and most critical of environments due to the need to host duplicate physical servers in multiple locations. Managing the configuration of those servers over time is difficult, as any changes to production instances would need to be replicated to the redundant ones. Let's take a look at the good and bad of backups and disaster recovery associated with virtualization.

Virtualization's Effects on Backup and Restore

Virtualization's impact on backups provides a much improved way to create highly-atomic snapshots of individual servers. This snapshot includes the entire configuration of the server in such a way that makes it very mobile and ultimately very restorable. Depending on the architecture chosen, those backups can take many forms.

With *Hardware Virtualization*, the files that make up an individual virtual machine are all part of a single file. That file contains the sum total of the virtual machine's configuration. Backing up this file is relatively easy when done from the virtual host. The host quiesces the virtual machine file, creates a snapshot, and backs up the snapshot as the virtual machine continues to run. The end result of this is a single file backup of that machine at that point in time.

The problem with this approach is that although entire machines can be restored easily through the restoration of a single file, restoring individual files grows more complicated. In order to restore individual files, the entire virtual machine file must be temporarily restored, opened via a management interface, and then the individual file obtained. This process can be time-consuming. Because of this problem it is often recommended that backups for *Hardware Virtualization* environments be done the traditional way within the virtual machine itself as well as through entire-machine backups as described above. This has the effect of doubling the amount of backups needed for an environment.

Contrast this situation to how backups are handled with *OS Virtualization* tools. Here, in order to backup an entire virtual machine, the management tools must back up the contents of that virtual machine's container on the host. Because the container has awareness of its presence on the host, quiescing features native to the operating system (such as the Volume Shadow Copy service in Microsoft Windows) can be used to perform the snapshotting.

This means that entire-machine backups for virtual servers are not single files. Instead they more resemble traditional file backups that are made up of the individual files that comprise the virtual machine. Though this means that an entire-machine backup is not encapsulated into a single file with *OS Virtualization*, the process of restoring entire virtual machines is the same as restoring individual files. The doubling of backups is unnecessary and incremental backups and restores are possible. *OS Virtualization* solutions like Parallels Virtuozzo Containers leverage server templates to provide the ability to restore a virtual machine to a different host if necessary.

Quiescence and Crash Consistency

The quiescence process for both of these architectures is critically important to what we call the *crash consistency* of the operating system and its applications. When any backup solution quiesces a virtual machine, it freezes that machine's contents in place for long enough to create a differencing snapshot. That differencing snapshot is then used during the subsequent backup to continue operating the server. At the same time the backup tool is backing up what is now a stable and unchanging OS file system. This is required to prevent corruption from occurring in the backup file should files change over the period of backup. This concept is often called *OS crash consistency*, and is a common feature of all virtualization architectures.

What OS crash consistency provides is the assurance that a restored server OS itself will operate as it did at the exact moment the backup was taken. But there is a problem with this level of crash consistency with some virtualization solutions. For some, the quiescence process does not properly "quiet" transactional databases that operate within the machine. When the quiescence occurs, the effect to any installed transactional databases is the same as a "hard stop" to the operating system. Transactional databases like Microsoft AD or Microsoft Exchange, when restored in this manner, can be restored showing an improper shutdown. Additional work will be required to bring the database back to proper health following a restore when using some *Hardware Virtualization* backup toolsets and entire-server backups. Some data could additionally be lost.

This issue is less problematic in *OS Virtualization* due again to the positioning of the layer of virtualization. Because the residing virtual machine's OS is aware of the management interfaces' attempt to quiesce it, it can do the same with its onboard transactional databases. This means that these databases are more likely to return from a restore in a healthy state.

> The problem of transactional databases is another reason why double backups are suggested for many *Hardware Virtualization* environments. Plan accordingly.

Replication

In any virtualization architecture, once backups are configured to support the storage of entire servers at once, the processes necessary to create a full disaster recovery solution become simple. Required are the movement of those backups to an off-site location and the positioning of hardware at that location to host virtual machines. In the case of a disaster, bringing the destroyed data center back on-line involves the restoration of backed-up virtual servers to new hardware and powering them on.

Three elements drive the costs associated with realizing this capability: The network connection between the primary and backup site, the quantity of time desired to return to operations after a disaster, and the amount of hardware pre-staged at the backup site:

- Network connection. The network connection between the primary and backup site must be of bandwidth large enough to support the regular transfer of backups to the backup site. Additionally, it must be sized appropriately so that it can remotely support the needs of the business after a disaster.
- Recovery Time Objective. The Recovery Time Objective (RTO) is the quantity of acceptable time computing operations must be restored after a disaster. Organizations that have a low tolerance for outages will have a low RTO and will require higher levels of automation in order to achieve the needed restoration target. Organizations with a high tolerance are more willing to afford a longer outage while administrators manually bring servers back on-line after a disaster event. Lower RTO metrics typically involve more cost to support the necessary automation.
- Hardware at backup site. Backup sites are typically considered either hot, cold, or warm sites, depending on the hardware pre-staged at the site. Hot sites include enough hardware to run the entire infrastructure and are configured to automatically come on-line immediately after a disaster. Warm sites include pre-staged hardware, but are not necessarily configured for an immediate failover after a disaster. Cold sites are those that have no pre-staged servers in place. Instead they typically involve leased accesses set in place for administrators to manually install hardware and restore services after a disaster.

> Obviously, the more temperature associated with the site, the more cost as well.

Environment-Specific Considerations

Lastly, there are some specific management considerations associated with special virtualization workloads. In this section, we'll talk about four of these special cases that are important to keep in mind.

Microsoft Terminal Services and Citrix Presentation Server

It is often heard in certain circles that "Terminal Services and/or Citrix don't play well with virtualization." Though some attempts at virtualizing remote application tools like these have resulted in failure, in many cases the problem lies not with the applications. Instead, it lies with the management of their performance. Terminal Services and Citrix are unlike nearly all other applications within the data center in that they involve the hosting of numerous users concurrently on server hardware. Because user actions are less predictable than server actions where no users are present, it is more difficult to fully understand the result on overall performance.

The problem is particularly difficult in *Hardware Virtualization* solutions like VMware Virtual Infrastructure. Due to these applications requirements for high levels of memory, *Hardware Virtualization* solutions can be insufficiently scaled to support the needs of users. This contrasts with *OS Virtualization* where memory sharing is done at the level of the host. Additional benefits are gained through centralized caching as well as the near-native throughput associated with *OS Virtualization* driver sets.

Notwithstanding the virtualization solution chosen, the most critical component to ensuring a successful integration of application and server virtualization is the vigilant management of performance. Ensuring that users are unable to perform actions on the server that tend to lead to high resource use, reducing the available touch points for users to specific applications, and monitoring of resource use to watch for spike conditions are all best practices.

Hosted Desktop

Hosted desktop situations can be in many ways very similar to the application virtualization situations described above. Rather than using Microsoft Terminal Services or Citrix Presentation Server to present applications to users, individual OS instances are created and connected to via network protocols. Like discussed above, this situation involves users directly accessing desktops within the data center. In order to prevent a single user from oversubscribing resources, resource maximums must be set in place. These maximums will ensure that the actions on any particular desktop will not negatively impact others on the same virtual host.

Where *OS Virtualization* excels in this situation is within its sharing of files on the server and its native driver support. Because hosted desktop environments are often very similar to each other, the individual personalization elements between desktops can be minimal. When personalization is kept to a minimum, then individual virtual machines are very small. This is different in comparison with *Hardware Virtualization* where a high level of file replication is necessary to support hosted desktops.

> Due to this resource sharing and near-native driver functionality, hosted desktops can be a best fit for *OS Virtualization*.

Virtual Appliances

The idea of a virtual appliance encompasses a ready-to-deploy virtual machine with a specific function. Typically that function is narrow in scope and involves very little setup to bring it to operation. Though virtual appliances are more often found in *Hardware Virtualization* environments due to the single file nature of its disk files, all virtualization architectures can support the creation of virtual appliances.

The process to create a virtual appliance involves creating a server template that includes the operating system, installed applications, and configurations that support the use of the appliance. Additional configuration is usually done to ensure that a quick appliance customization can be done within the target environment to support its use.

Although all architectures can support the creation of virtual appliances, one benefit of using *OS Virtualization* is the smaller file size of such virtual machines once created. As appliances are usually based off of server templates available in all environments, only the specific customizations need to be part of the appliance itself.

64-bit Environments

Lastly is the consideration of 64-bit environments for both virtualization hosts as well as residing virtual machines. 64-bit hardware has the capability of supporting higher levels of processor throughput as well as higher quantities of RAM memory. Because of this, 64-bit environments have the capability to support a higher consolidation ratio, with more virtual machines running on each virtual host.

Notwithstanding the processor architecture of the virtual host, individual virtual machines still require memory in order to operate. The move towards 64-bit hardware means that more virtual machines can be supported concurrently, though it is important for the virtual host to be configured with enough memory to support their concurrent operation. When considering the use of 64-bit hardware for supporting virtual environments of any architecture, be sure to include extra memory to support the higher level of concurrency.

Summary

Getting the right virtualization software in-house and getting it installed are only the first components of a successful virtualization rollout. Managing it effectively over time ensures continuous return on that initial investment. No matter what virtualization solution you choose, it will come with its own set of benefits and detractors. What is important is to find the solution that fits best within the workloads you plan to operate within your computing environment.

Throughout the last four chapters, we've attempted to unravel some of the hype surrounding virtualization in the market today. At the same time, you've had the chance to educate yourself on alternate options that may drive the needs your particular servers better than others. The decision is up to you to find the solutions that work best for you.

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.