Realtime
publishers
"Leading the Conversation"

The Shortcut Guide™ *To*

Selecting the Right Virtualization Solution

sponsored by

Parallels™

Greg Shields

Copyright Statement

© 2008 Realtimepublishers.com, Inc. All rights reserved. This site contains materials that have been created, developed, or commissioned by, and published with the permission of, Realtimepublishers.com, Inc. (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 Realtimepublishers.com, Inc or its web site sponsors. In no event shall Realtimepublishers.com, Inc. 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.

Realtimepublishers.com and the Realtimepublishers 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 Realtimepublishers.com, please contact us via e-mail at info@realtimepublishers.com.

[**Editor's Note:** This eBook was downloaded from Realtime Nexus—The Digital Library. All leading technology guides from Realtimepublishers can be found at http://nexus.realtimepublishers.com.]

Chapter 3: Best Practices in Implementing Virtualization

Purchasing a virtualization solution is only the first step. Properly implementing it is critical to gaining the greatest Return on Virtualization. So far in this guide, we've run through a comparison of virtualization architectures and aligned products with each architecture. For organizations looking to finding the best-fitting—right—virtualization solution, it is necessary to know the field of entrants. As both previous chapters explored, there are multiple types of virtualization architecture, each of which is served by a dominant product.

Chapter 1 focused heavily on the specific architectures being used in today's computing environments. The chapter talked about how *Hardware Virtualization* abstracts computers at the layer of their physical hardware, virtualizing their memory, processors, disks, and network cards. The chapter showed how *Hardware Virtualization* solutions scale individual computers horizontally atop a hypervisor layer. We contrasted that architecture with the idea of *OS Virtualization*. In *OS Virtualization*, we move the layer of abstraction to within the operating system (OS) itself. We showed that for certain configurations, *OS Virtualization* provides an enhanced benefit over *Hardware Virtualization* due to its near-native speeds in homogeneous environments. We also compared and contrasted these two architectures with those of *Paravirtualization* and *Application Virtualization*.

In Chapter 2, we moved away from architectures to talk about specific products. We looked at the feature sets associated with individual virtualization products and analyzed four popular products currently available in the market:

- VMware ESX & Virtual Infrastructure and Microsoft Virtual Server, both of which are excellent examples of Hardware Virtualization.
- Xen and specifically the Citrix XenSource distribution, which is the most prevalent manifestation of the Paravirtualization architecture currently being used.
- Parallels Virtuozzo Containers, a product that uses OS Virtualization to abstract virtual environments (also called *Containers*) as linked components of the virtual host.

Getting the right solution in-hand is only one step in the process. Understanding the ways in which that product will be implemented is critical to ensuring that the best solution is acquired. This chapter will attempt to explain some of the best practices for using all of these types of virtualization within your computing environment today.

Virtual Environments Are Different than Physical Environments

Virtual environments require a different level of care than do physical environments. Many organizations move to virtual environments because of their natural capabilities for easier management and the potential for greater availability and uptime. What is not readily understood is that virtual environments in some ways add their own set of risks that must be properly managed separately from pure physical environments. These risks align with three greater concepts—their improved density, their greater complexity, and their enhanced capabilities for automation.

Greater Density

First and foremost, most organizations incorporate virtualization into their networks because of the driving need to consolidate physical machines. This consolidation activity reduces the total number of physical machines in the environment while allowing for the same number of network services to operate as before. One concern associated with this increased level of density surrounds the capacity for more individual services to operate on the same physical hardware.

Physical hardware failure characteristics are common across devices. This is the case regardless of whether a single OS instance is installed directly on the physical hardware or a virtualization solution is implemented to consolidate multiple machine instances. In the case of virtualization, the difference is that a host failure can impact more than one service. An increase in density due to virtualization consolidation means a corresponding increase in outage exposure when a host failure occurs.

Consider an all-physical situation in which a single network service occupies a single physical server. In this example, the loss of a single server means the loss of a single service. Next consider an all-virtual situation in which 10 network services—each on an individual virtual machine—are all housed on a virtual host. The loss of a virtual host can mean the loss of 10 network services.

In purchasing hardware to support the greater density that occurs with virtualization and consolidation, take special care to acquire server hardware that is most resilient to physical failure. That resiliency may be manifested in the form of physical redundancy or higher-quality server-class components. These higher-quality components are usually available in the high-end server equipment needed to support virtualization.

> For example, the *Mean-Time Between Failure* (MTBF) metric for a hard drive is the same no matter what type of data that drive is storing. That metric does change, however, when the disk is used more often. In virtualization environments, the need for disk resources by multiple, simultaneous virtual machines mean more disk activity. This increase in disk activity can reduce the overall effective life of that disk. Because of this characteristic of greater use, server redundancy features are critically important within the virtualization environment.

Most high-end server-class equipment is also equipped with on-board hardware management and notification capabilities. These capabilities can notify a central network management system when failure—or even pre-failure—events have occurred on the system. In many cases, these capabilities are already a part of the hardware but may not be used. Enabling these features helps ensure better service resiliency in the environment. In the case of pre-failure warnings, leveraging your notification system gives the administrator time to relocate virtual machines to healthy equipment prior to a failure.

Greater Complexity

Many all-physical computing environments do not follow the practice of *service isolation*, often due to financial constraints. When an IT organization practices service isolation, they are using a single computer to support a single network service. Adding a new network service means adding a new server. Effective service isolation means that the loss of a server will result in the loss of only a single service. It also reduces the complexity of troubleshooting associated with identifying problems on a particular server.

Service isolation is operationally expensive in physical environments, especially in organizations with limited funding; the move to virtualization can lower many of these cost barriers. Creating a new server and its associated service is as easy as a “copy-and-paste.” It takes little time and effort to replicate services. Software licensing is financially friendly towards virtualization-hosted OSs.

There is one problem. An unintended consequence of this “easiness” can be a perfect storm of server and service expansion. An environment that moves to virtualization can experience a massive bloat in total server count, even as the number of physical machines stays the same.

This increase in total OS instances has a tendency to increase total environment complexity. More instances means more machines to patch, more applications to manage and monitor, and more services to maintain. Some virtualization solutions, such as those associated with *OS Virtualization*, provide native tools to assist with the management of virtual environment OSs and their installed applications. Other solutions may not natively support these features. Care towards the management of an increasingly complex environment must be taken such that diseconomies of scale do not appear.

Additionally, the ease of creating new virtual machines introduces new security-based issues into the environment. The horizontal increases in the total number of machines may drive the creation of more machines that are only rarely used. With some types of virtualization architectures, these rarely used computers can pose an additional risk to the operating environment associated with their configuration.

Let’s think for a minute about this problem. If a virtual machine is created, used for a period of time, and then shelved for future use, that machine is no longer continuously operational. Typical systems and patch management tools often don’t have the capability to power on the machine, patch it, and power it back down. Thus, the nature of these machines being powered off for long periods of time increases the risk that an already-patched exploit can infiltrate the machine. Because these machines are powered off during the typical patch cycle, they can “miss” certain critical patches.

In the case of *OS Virtualization*, the files of the powered-off virtual computer can still be linked to those on the virtual host. That virtual host typically remains powered on, which means its configuration is known and up-to-date. Thus, its resident virtual machines—even those that are powered off—are more likely to power on with a correct and fully patched configuration.

> The problems of powered down equipment are only now being recognized. These problems are not specific to virtual machines but are exacerbated by their nature. Products are only now being made available that assist in managing this security hole.

Greater Automation

One very positive characteristic of virtualized environments is their capacity for enhanced levels of automation. This automation is more extensible than with traditionally highly manageable physical hardware. Whereas individual OSs already include rich tools for programmatically modifying and managing configurations, the tools for managing physical hardware have traditionally been highly device-specific.

With virtual environments, the processes for machine-specific automation tasks such as powering on, powering off, backups, and bare-metal restoration are all similar for each virtual machine. The use of different physical hardware does not require different management tools to support this automation.

In addition, with most virtualization solutions scripting and programmatic exposure via published APIs adds management flexibility. These APIs are built into the virtualization framework and can be used no matter where the virtual machines may be housed. A good scripter or coder can make use of these APIs to easily create their own custom interfaces.

> For example, if you need to create scripts for mass rebooting, backing up at the level of the virtual machine, or even mass restoration activities, the process is much easier than with physical machines alone.

> We'll talk more about some potential automation benefits later on in this chapter.

Potential Usage Scenarios and Best Practices

Depending on the virtualization solution chosen, there are several usage scenarios that fit well within that solution's architecture. In this section, we'll take a look at five potential ways of using virtualization to gain an accelerated return over with physical machines alone. For each of these five scenarios—Infrastructure Service Consolidation, Disaster Recovery & Business Continuity, Dynamic Workload Management, Code Development & Testing, and Virtual Desktop Infrastructure we'll also talk about best practices associated with their use.

Infrastructure Service Consolidation

As we discussed in Chapter 1, the average metric for server utilization is around 5% across all industries. This means that servers are on-average performing useful work 5% of the time. For the other 95% of their operational life cycle, they aren't adding any value to the business. This problem is particularly relevant for machines labeled "infrastructure servers." These typically low-use servers—such as DNS servers, Active Directory (AD) domain controllers, patch management servers, and reporting servers—are necessary for the operation of the environment. They typically cannot collocate their services with others, and they traditionally have the lowest usage of all servers in an environment.

> ✎ Infrastructure servers can be the lowest hanging fruit for virtualization.

Unlike mission-critical services, such as databases or services with complex configurations such as industry-specific software, the services on these servers are typically well-understood. The movement of these services to virtualization is usually supported by their vendors. Their movement can be done through a physical-to-virtual process as easily as a complete rebuild. They are typically redundant, so the loss of a single instance will not impact the environment as a whole.

Any of the virtualization architectures we've discussed in this guide can well support these types of infrastructure services. As low-resource services, they typically enjoy an excellent consolidation ratio. Most importantly, they typically do not involve large numbers of third-party applications where introduction into a virtualization environment may violate support agreements. Infrastructure services are often an easy win for virtualization-based consolidation.

Disaster Recovery & Business Continuity

Our conversation on infrastructure services dovetails perfectly into a discussion on the twin topics of Disaster Recovery and Business Continuity. Both are parts of the same whole—that of ensuring service continuance during and after an impacting event. But because Disaster Recovery and Business Continuity are subtly different, there are differences in the types of solutions necessary to compensate for them.

> ✎ *Business Continuity* typically associates with the need for the continuance of a service after the loss of a single service or service element. This involves adding compensating mechanisms to help protect against the situation in which a single server or service goes down.
>
> Conversely, *Disaster Recovery* usually involves the loss of an entire location or data center, often due to a natural or manmade disaster. As you can see, a disaster recovery event involves many more impacted systems than the single instance of a business continuity event. Thus, different tools and techniques are needed to prepare for a disaster recovery event.

As noted, incidences of business continuity typically involve the loss of a single service or service element. When that service goes down, the business is incapable of performing a critical operation. The standard solution for these sorts of events is to provide a redundant server that will fulfill the needs of the business once the primary server goes down. Traditionally, this “clustering” approach was costly to implement and challenging to maintain. Similar hardware and configurations were and are critical for clustering solutions so that the cluster can reliably move services from node to node.

With virtualization, the need for similar hardware is lessened somewhat. From a hardware perspective, the servers in the virtualization environment are mirrors of each other. Individual virtual hosts can relocate virtual server instances between them as necessary in preparation for a failure event. This further enhances the uptime capabilities of the redundant services. Many virtualization architectures support this “hot migration” capability, moving virtual servers from host to host without the loss of uptime. All virtualization architectures support the “cold migration” capability, moving virtual servers from host to host after the virtual machine has been powered off.

The time-to-restore for an individual failed server can be quite different based on the type of virtualization architecture chosen. As an example, with OS Virtualization, the startup process for a failed system can be significantly faster because the host server is typically already in a running state. Resources are already online and ready for use, which speeds the booting process. Contrast this with Hardware Virtualization where individual machines must complete a full boot cycle after a failure along with all the associated resource spin-up necessary to complete the boot process.

In the situation of disaster recovery, backups are critical. But backups are only one piece of the puzzle. A disaster event can be catastrophic to a business’ livelihood if services are not brought back online within a very short period of time. Mere tape-based backups may not be capable of bringing services back online in a timely enough fashion to support a business’ requirements. This is often due to the sheer number of individual files that make up a single computer, the loss or corruption of which—as part of the backup—can compromise the successful restoration of that server. For most organizations, after a disaster, speed is of the essence.

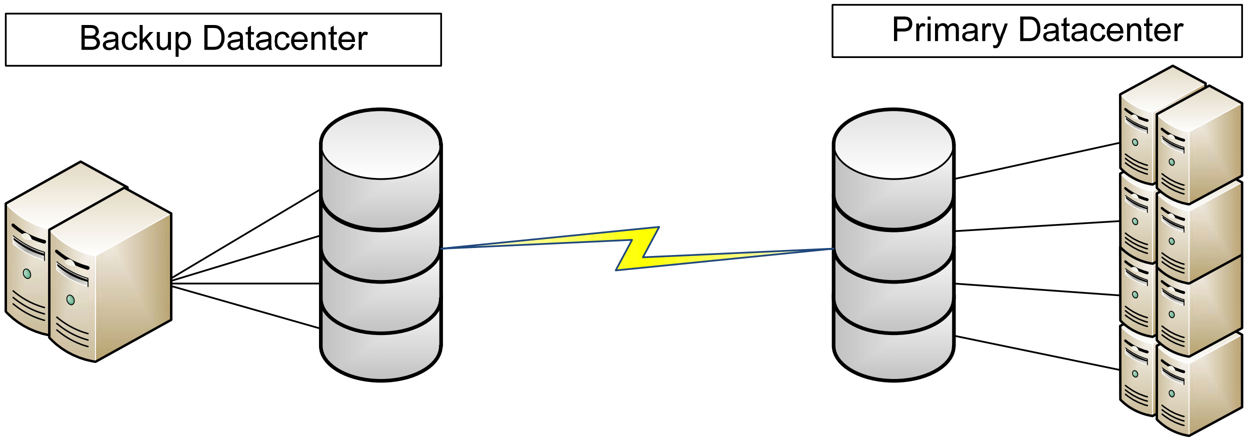

Figure 3.1: With data replication from primary site to backup site, virtual machine files can be replicated elsewhere. A major benefit with virtualization is that a 1:1 ratio of primary to backup machines is not necessarily required.

One solution for speeding this process is to use a tool that replicates the virtual machine backups to an alternative site in real or near-real time. That backup arrives at the backup site in a way that is easily restorable to a replacement host. As Figure 3.1 illustrates, the virtualization hosts in the primary site can automatically transfer their backups over the network to a data store at a backup site. Virtual machines at the alternative site can be rapidly provisioned from this backup data to replacement hardware, thereby greatly speeding the return-to-operations of necessary network services.

Enhancing this solution even more is the nature of post-disaster operations. In many industries, fewer users are using network services during a disaster; a 1:1 ratio of virtualization hosts may not be necessary at the backup site. More virtual machines can be consolidated onto fewer hosts because their resource needs are less. Fewer hosts at the backup site mean a lower cost to support that backup site.

✎ Since the beginning of data centers, the cost-effective ability to provide a full disaster recovery facility has been elusive. Early attempts involved the creation of a fully redundant set of physical servers at the alternative site. Configuration changes at the primary site needed to be made also at the redundant site. Human errors over time could misalign configurations, causing the backup site to no longer mirror the production site.

Many virtualization solutions include, either natively or as a third-party add-on, the capability to offload backup files to a remote site for rapid re-provisioning after a disaster event. This software-based approach is significantly more cost-effective than with previous solutions. When considering the right virtualization solution for your environment, look for one that can support your potential or real needs for supporting recovery in the case of a disaster.

Dynamic Workload Management

Related to the need for consolidation is a key desire to squeeze more useful processor cycles out of expensive server hardware. This is particularly useful in situations in which available funding or space to support additional hardware does not exist. Dynamic workload management is the virtualization concept that individual virtual machines and their workloads can be spread across multiple host systems. Depending on the virtualization product chosen, these migrations can occur automatically or manually and with or without an outage to the host system. Other products allow for the highly granular assignment of physical resources to individual virtual machines that can scale to fill all the resources that make up the host system. Some require a service outage to support resource changes and some do not.

To support dynamic workload management, *Hardware Virtualization* products such as VMware Virtual Infrastructure use a concept called VMotion. VMotion is a process whereby the ownership and processing of virtual machines is transferred from one physical host to another without a corresponding loss of service. The virtual machine need not be powered off to support this capability. Adding this capability to a virtualization solution means that VMware-hosted environments can load balance virtual machines as necessary across multiple hosts. This ensures that resources are best distributed to servers that need them and no single physical host is overloaded.

To the virtual machine-specific assignment of resources, VMware Virtual Infrastructure can supply varying levels of resources such as disk, RAM, and processor power to virtual machines. For many of these changes, a reboot of the virtual machine is required in order for that virtual machine to recognize and use the resources.

OS Virtualization tools such as Parallels Virtuozzo Containers have similar tools for a virtual machine “hot migration” from host to host. Depending on the software chosen as well as the virtual machine OS this process may or may not require a short outage of the system, typically just the time to restart a virtual environment, which is much less because the OS is already running. *OS Virtualization* operates at a different layer than the hypervisor-centric *Hardware Virtualization*, so the process for migrating machines is slightly different so it does not require a SAN or dedicated storage However, the end result is the same—the capability to load balance resources across multiple physical hosts.

✎ A critical point to recognize here is the ability with these advanced features to further abstract computers from the hardware they reside upon. Once virtual machine migration capabilities associated with dynamic workload management are introduced into a virtualization environment, the administrator no longer needs to consider each host as an individual computer chassis. Rather, the host can now be considered a “set of resources” of which any virtual machine can make use. This further abstraction of resources makes possible the best and most efficient use of available resources.

Another feature of *OS Virtualization* tools like is the ability to dynamically alter resources assigned to virtual machines on the fly. Because *OS Virtualization* tools do not use emulated driver sets, they are not limited to predetermined and code-limited quantities of resources such as number of processors, hard drive or amount of RAM. Virtual machines hosted on OS Virtualization products can make use of any and all resources that make up the physical host with no limit on maximum size or use.

✎ Combining automation with the ability to dynamically adjust resources on the fly allows machines to be given just the amount of resources necessary to perform their function. No longer do machines need to be "over spec'ed" with resources they may not even use.

Code Development & Testing

Virtualization grants some very specific benefits for code development and testing environments as well. These types of environments are typified by high turnover rates, meaning a repeated need to rebuild environments to support new tests or new code versions. Often, due to the typical code development process, multiple, simultaneous environments are necessary to support overlapping code, unit test, qualification test, and staging environments. Depending on the testing schedule, multiple environments for each of these stages may be required. In an all-physical situation, these multiple environments can be prohibitively expensive to purchase and maintain.

Virtualization and its management tools provide an easy interface for speeding the rapid re-creation of these environments. The "snapshotting" feature of many virtualization tools freezes a machine's configuration. It makes trivial the steps to roll back the configuration to that snapshot after the test is complete. A test can be run over and over rapidly and repeatedly, with the only requirement being that the tester reverts to the snapshot between each test. The reversion process is simple and automatic and significantly reduces the amount of time necessary to reset the environment in preparation for another activity.

More so, most virtualization solutions provide the capability to rapidly deploy new servers. Depending on the solution chosen, this rapid deployment can support the stand-up of templatized machines only or templatized machines along with necessary software packages.

Tools such as VMware Virtual Infrastructure support the creation and subsequent deployment of virtual machine templates. These templates can be augmented with automation that automatically adds them to the necessary network locations and Windows Active Directory (AD) domains.

Other tools such as Parallels Virtuozzo Containers also support machine templates. As Virtuozzo Containers is an *OS Virtualization* tool, it also supports the addition of software packages to deployed templates as necessary. This native addition of software packaging further aids in the rapid provisioning of necessary servers by speeding the resolution of custom requirements. This makes Virtuozzo Containers a great tool to provision large quantities of servers for stress testing environments.

✎ Another valuable benefit here is the reduction in number of templates required overall. This is possible by removing individual software packages from the core image and packaging them separately. A "blank" template can be created and later augmented with software packages.

Virtual Desktop Infrastructure

Our final usage scenario involves the removal of the desktop altogether. The management and maintenance of individual desktops in an organization can be one of the most expensive operational costs for an IT department. When technicians need to individually go to desktops to solve problems, the environment may not be serviced as quickly as necessary due to schedule conflicts or available technician resources.

The centralization of desktops was first pioneered through remote application tools such as Microsoft Terminal Services and Citrix Presentation Server. These tools, still commonly used today, provide an excellent mechanism for aggregating users and their applications onto server-based hardware. However, in some situations these tools cannot provide the necessary level of user separation. Some applications may not function properly on multi-session servers such as Microsoft Terminal Services or Citrix Presentation Server. Most importantly, users may want or need their own individual desktop that is not shared with other users. In any of these situations, it may be necessary to “host” the user’s desktop as a virtual machine.

The process of hosting a desktop is, in concept, relatively simple. The desktop is virtualized and hosted on a virtualization host. Users access their desktops through a network interface (for Microsoft Windows environments, this is most commonly Microsoft’s RDP protocol or Citrix’s ICA protocol). Quickly depreciating desktop hardware can be replaced with longer-depreciating thin client equipment. Users gain the ability to use their desktops from anywhere with a secured network connection.

All virtualization solutions provide the capability of virtualizing desktops. Some also add management components and easy user interfaces for connecting to those desktops. One important difference, however, between the various architectures is involved with the horizontal scaling of resources. With *Paravirtualization* and *Hardware Virtualization* tools such as Citrix XenSource and VMware Virtual Infrastructure, all files and other resources that make up the individual desktop machine must be replicated for each machine to be hosted. Thus, if an example desktop consumes 20GB of disk space, hosting 100 desktops will require at least 2TB of online storage to support the environment.

A major benefit associated with the architecture of *OS Virtualization* and tools such as Parallels Virtuozzo Containers is that individual virtual machine files can be shared across multiple virtual machines. Consider the case in which 20GB of disk space are required per desktop to support its hosting, but only 1GB is information that is different between desktops. If we attempt to virtualize 100 desktops, the total disk space necessary is only 120GB, a difference in size of nearly an order of magnitude. This is made up of the 20GB that is shared amongst the virtual machines plus the gigabyte of different data on each of the 100 desktops. With the cost for storage increasing geometrically as data size increases, this reduction in required storage space can be a major savings in overall cost.

✐ Another benefit with *OS Virtualization* is one discussed in our previous chapter. Hosting desktops does not eliminate the need to manage their security profile and configuration. With *OS Virtualization*, patching all hosted desktops means patching only the host. This reduces the total number of potential open points on a network from a security perspective and eases their management burden.

Obtaining Maximum Return on Virtualization

With all virtualization solutions, there will be an upfront cost to make the jump. With some solutions, especially those that require high-end iSCSI or Fibre Channel SAN storage, that upfront cost can be substantial. Thus, obtaining the best return on your virtualization dollar is important for validating the purchase.

The old adage "never trust an ROI that you didn't create yourself" still holds true today. But one very major difference between typical software ROI numbers and those associated with virtualization is that most virtualization ROI metrics are based on hard dollar return. Take a look back at the dollar calculations in Chapter 1. The typical metrics for RoV expenditures relate the cost of software and hardware to the return on reduced power and cooling costs, reduced provisioning costs, and reduction in overall data center footprint. Each of these is a hard, quantifiable metric.

That being said, the average payback for an investment in virtualization is around 6 months. This metric means that an investment in virtualization will provide positive benefit back to the business within a short 6-month period. In this section, we'll take a look at some of the ways an organization can maximize that payback, specifically associated with hardware, infrastructure improvements, power and cooling issues, and the virtualization software itself.

Hardware

First and foremost, virtualization's drive towards consolidation will naturally drive an organization towards larger and more powerful server hardware. Additionally, the aggregation of multiple services onto a single host will further drive higher-quality components that incorporate higher levels of redundancy. Thus, the servers of yesteryear are not likely to be well-suited for incorporation into a production virtualization environment.

Or maybe that's just what the consultants want you to hear. One major skill required for a successful virtualization deployment is a deep understanding of systems performance management. An administrator who has a deep understanding of the required and available performance of the systems in the environment will be well-suited for managing that environment with only as many resources as absolutely necessary.

Existing hardware should be measured prior to a production deployment to validate the quantity of virtual machines that can potentially run on this hardware. If this metric is acceptable to the business, it is likely more cost-effective to retain existing hardware for use within the virtualization environment until that hardware's effective end-of-life rather than an immediate purchase.

As virtual machines can be easily, and in some cases uninterruptedly, relocated from physical hardware to physical hardware, the upgrade of a piece of physical hardware is usually a trivial event. Virtual machines are moved from the virtualization host. That host is then powered down and replaced with a new piece of hardware. Virtual machines are then relocated back to the new hardware.

Depending on the virtualization architecture selected, licensing costs may be per processor, per socket, per hosted virtual machine, or combinations of these. That licensing model can impact the decision to keep existing hardware or purchase new. Either way, virtualization licenses for many products typically easily can be relocated from host to host. Thus, the need for an initial hardware purchase to support the environment can be easily delayed. This has the tendency to reduce the early costs for the move to virtualization.

Infrastructure

Depending on the virtualization product chosen, the introduction of virtualization into the computing environment can impact areas of its infrastructure. Network ports are one example. If an environment requires 50 servers to run the business and each server utilizes dual-redundant connections to the network, a total of 100 network ports are required plus the additional ports to interconnect the necessary switches and routers.

Making the move to virtualization, the total number of network ports can be significantly reduced. Let's assume that the virtualization solution itself requires only dual-redundant ports for a physical host. Assuming a 10:1 consolidation ratio is realized, the total number of ports can be reduced from 100 to merely 10 to support the 5 physical hosts running the virtual machines.

> ✎ This benefit is realized in even greater ways when one understands that a reduction in port count means a move towards smaller switches and less network infrastructure. These smaller network components are typically much less expensive to purchase and maintain.

For environments making use of keyboard, video, and mouse (KVM) technology for remote access to server consoles, this reduction in total port count can be even greater. Typical virtualization management consoles have the tendency to eliminate or significantly reduce the need for administrators to perform actions on the physical host's console itself. Individual virtual machine consoles are made available to administrators through a network-based interface. This eliminates the costly doubling of network infrastructure associated with remotely accessing server KVM elements or passing them over the network.

Related to security, there is a requirement in some environments for non-IT individuals to maintain access to the data center. This may be a result of their need to work on server consoles due to a software limitation. Allowing these quasi-trusted individuals into the data center can be a security concern. With most virtualization products, the capability exists to remotely access the console itself over the network. When this occurs, these individuals no longer will need access into the secured data center to perform their daily activities.

Power & Cooling

We talked about power and cooling enhancements at length in Chapter 1. But one specific area in which these two elements are especially improved through the application of virtualization is where a wholesale infrastructure expansion is required. The cost to augment existing data center equipment with greater wattage and tons of chilling—or to expand the data center itself—in most cases far out-costs a move to virtualization. Costs for both power and cooling often add little additional value to the business, so any reduction there eliminates a liability upon the business that is difficult to link to assets.

Software

Certain software and licensing costs can be similarly recouped. Specific to the costs for OS licensing, Microsoft has announced a bonus for licenses used within virtualization environments. Specifically for certain versions of its Windows Server OS, namely Enterprise Edition and Datacenter Edition, the purchase of a single physical server license bestows additional virtual server licenses for no added cost. The rules are as follows:

- *Windows Server Enterprise Edition*—Each Enterprise Edition server license for a physical server instance also bestows four additional licenses to be installed on virtual server instances. Thus, in a virtualization environment, the purchase of a single Enterprise Edition license nets a total of five total Windows Server licenses.
- *Windows Server Datacenter Edition*—Each Datacenter Edition server license for a physical server instance bestows unlimited licenses to be installed on virtual server instances. Thus, in a virtualization environment, the purchase of a single Datacenter Edition license nets an unlimited number of Windows Server licenses.

> ✎ There are two important notes about these rules: First, the purchase of a single physical license only bestows the additional virtual licenses onto a single server. Thus, the unlimited virtual licenses gained through the purchase of Datacenter Edition only licenses those unlimited virtual instances onto a single server.
>
> Second, be aware that the licensing rules state that this limitation is for *running instances*. You can host as many *non-running instances* as you want. That being said, if more non-running instances are present on the system, in order to fulfill your licensing requirement, a running instance must be powered off in order to power on a non-running instance.

Different virtualization architectures have differing capabilities of hosting this number of virtual machines simultaneously. Due to the sharing benefits of *OS Virtualization*, there is a higher potential for server density than with *Hardware Virtualization*. Thus for any particular set of hardware, when using an *OS Virtualization* product, there is a higher likelihood of financial gain associated with an upgrade to Datacenter Edition than with other architectures.

Best Practices in Systems Automation

Aligned with the scenarios discussed earlier are some elements of systems automation that become very easy when moving to virtualization. Due to the interface commonalities for scripting and programmatic automation, the creation of scripts that work across the entire enterprise is easy and convenient.

Also possible is an enhanced ability to empower the individual server user to perform many tasks without the need for IT administrator management. Adding the concept of user self-management into typical server management offloads many mundane tasks typically relegated to the systems administrator instead of the service owner. Going a step further, it may be possible to assign per-use accounting and chargeback models to the environment, ensuring that costly virtualization equipment is paid for by the entities that are using it.

Automated Provisioning

No discussion of automated provisioning is complete without linking the provisioning activity to server templates. The concept of server templates or “golden images” has been around in IT since the introduction of rapid deployment tools such as Norton Ghost and ImageCast many years ago. In fact, much of the common knowledge associated with templatizing servers has developed initially through physical direct-to-hardware rapid deployments.

With virtualization, the process of templatizing servers is in many ways no different than doing so for an all-physical environment. Typically, the process is completed with the following steps:

1. Create a virtual machine and install a desired OS to the instance.
2. Install necessary applications, patches, and configurations. Ensure that any installed code is that which should be common to all systems that will be deployed vis-à-vis this image. Code that is custom to a particular server should typically be installed separately from the image.
3. Generalize the server to a template. This process eliminates any name or networking references to already-present servers. This is done to prevent conflicts when the template server is later powered on for provisioning.
4. (Optionally) Install an automation component to the server, such as Sysprep for Microsoft Windows machines. This automation component will personalize the server immediately after its first post-deployment boot.
5. Copy the template to a template location and secure it against inappropriate changes.
6. Later, as desired, deploy the template using either manual or automated methods from within the virtualization interface. These automated mechanisms typically will copy the template from a secured location to a server of lowest current load, boot the template, and begin the personalization process.
7. (Optionally) Install any custom applications and/or OS customizations to the server to enable it for serving in the desired capacity.

Depending on the virtualization solution chosen, some or all of these processes can be automated. With VMware Virtual Infrastructure, an add-on tool called VMware Lab Manager can be used to automate the process. Native to the VMware Virtual Infrastructure interface are limited capabilities for doing this as well. For Parallels Virtuozzo Containers, many of these customization and templatizing tools are built-in to the management interface, making this process relatively trivial. As stated before, the Virtuozzo Containers interface also includes application packaging tools to ease the administrative burden of applying post-deployment application customizations.

Delegation

The administrative activities associated with server management have traditionally been left to the trusted systems administrators. Tasks such as patching, powering on and shutting down, application installation and management, and the creation of new server instances could only be done by a systems administrator. This bottleneck had the tendency to cause the instantiation of necessary services to schedule slip when administrators were tied up with other activities. When IT organizations didn’t make use of automation, these processes consumed even more time.

One component of virtualization management is the capability of assigning server resources to alternative individuals. Depending on the virtualization architecture chosen, this process may be subtly different. However, the end result is that individuals outside the typical systems administrator can be delegated certain responsibilities for server administration. This is possible due to virtualization’s granting of console access over the network combined with the management interface’s level of granularity.

Consider the following situation: A business decides to purchase a virtualization solution. That solution is co-purchased by both the IT department and a code development team. In a traditional all-physical environment, the process to stand up and configure the environment would typically be done by systems administrators. The developers typically would need to wait for the environment to be completed before they can begin their work.

With virtualization and automation components built-in to the interface, it is possible to assign certain levels of processor, RAM, and other server resources into a pool usable by the developer team. Thus, if in our example the developer team paid for half of the environment, they can be assigned 50% of the available resources to do with as they will. If they want to create one “big” machine out of their resources, they can. If they want to create dozens of “little” machines, that is similarly possible. Their resources are theirs to do with as they see fit.

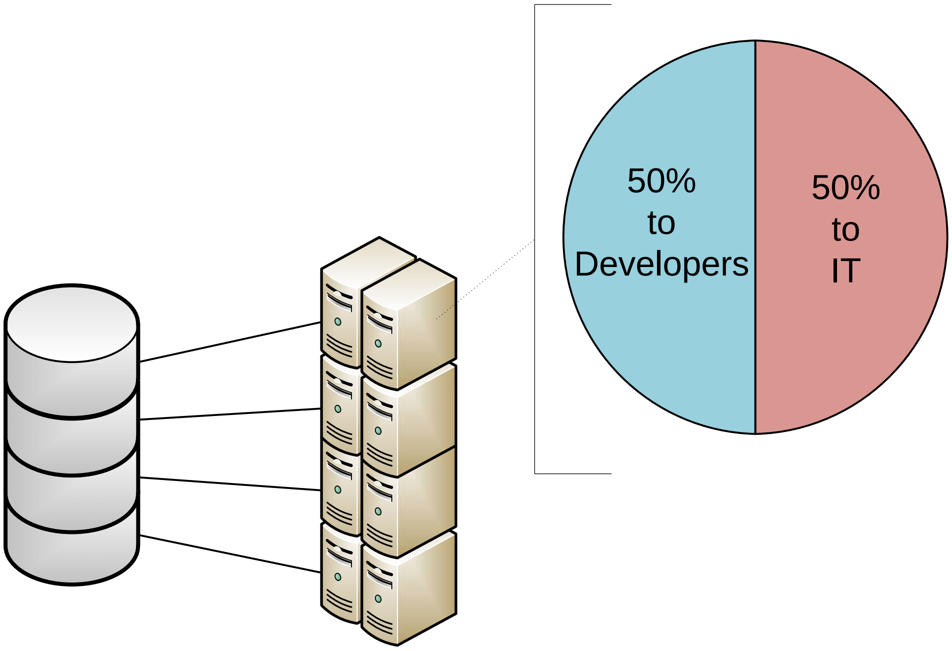

Figure 3.2: With resource delegation, resources can be assigned to different groups. Those groups can be given the rights to use their resources in whatever way they see fit.

Self-Management

It is further possible to couple this delegation ability with granular rights management. Continuing our earlier example, the virtualization environment can make use of server templates along with an automated management interface. Adding in permissions and delegation controls, it is now operationally possible for the systems administrator to delegate many server creation, startup and shutdown, and even packaged application management duties to the developer group. The administrator is no longer the bottleneck to necessary data center operations. New servers and services can be brought online through the virtualization solution’s management interface.

All this enables a sense of self-management for non-administrative users. The systems administrators can retain the ability to manage the environment as a whole. But their involvement is no longer necessary for the daily mundane tasking of basic server administration.

> ✎ Virtualization enables the offloading of much of this work to other individuals. Thus, highly talented and highly paid administrators can focus on strategic initiatives rather than basic tasks.

Usage Accounting & Chargebacks

Lastly, taking this entire example a final step is the capability of adding usage-based accounting to managed resources in the virtualization environment. We've already shown how it is possible to abstract server resources into buckets of resources. Those resource buckets can be then assigned to individuals or groups based on their needs. Through granular permissioning, non-administrators can assign specific rights and privileges to manipulate virtual machines within their assigned quantity of resources. This concept of workload management allows non-administrators to easily provision new resources on-the-fly with a much-reduced need for administrator involvement.

Once computing resources are abstracted away from individual machines, it is possible to assign a dollar value to those resources. That value can be used on a per-use model to ensure that the people using resources within the virtualization environment are properly paying for their use. Chargebacks are a mechanism to "charge back" to the user a fee associated with the use of the resource. This is commonly done in hosting environments in which desktops or servers are hosted by one entity for another. But it can also be done within corporate environments in which different organizations have different budgets and needs for computer resources.

Different virtualization solutions have different ways of enabling usage-based accounting and chargebacks. Depending on your need for enabling this functionality, the mechanism in which the interface enables these features may be a compelling reason to drive towards a particular virtualization solution.

Incorporating a Best Practices Approach to Virtualization Implementation Ensures Maximum RoV

As we've seen in this chapter, there are several common best practices associated with the move to virtualization. Specific to the implementation, different solutions can provide different benefits to the business. Finding which solution works best for the type of implementation planned by your business is critical to choosing the right virtualization solution.

In the next chapter, our last, we'll conclude our discussion on best practices. There, we'll focus on the management of virtual servers and their physical hosts. You'll see that there are a number of ways that virtualization management can further assist with ensuring a maximum return on your virtualization investment.

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.