Realtime publishers

The Definitive Guide™ To

Virtual Platform Management

sponsored by

ca technologies

Anil Desai

Copyright Statement

Chapter 4: Virtualization Management Challenges

Although virtualization technology can provide many significant advantages to IT environments, the use of virtual machines is not without potential drawbacks. Many IT organizations have come to the realization that initial purchase costs of hardware and software represent only a relatively small portion of overall expenditures. Additional factors that contribute to the bottom line include expenses related to administration, maintenance, and monitoring.

The focus in this chapter is on highlighting the types of issues that environments are likely to face as they manage infrastructures that contain a mix of physical and virtual computing technologies. The focus is on identifying the problems and how they can have a huge effect on overall IT environments. Organizations that are just getting started with virtualization should keep in mind what's coming up. Those that have already invested significantly in the technology have probably already started to run into some of these problems.

The overall organization of this chapter is based on the process of managing the complete life cycle of physical and virtual machines. It will start with deployment, then look at administration, configuration management, and monitoring. Although the focus on identifying problems related to virtualization might seem somewhat negative, later chapters will build upon this information as the guide evaluates various solutions that can address these challenges.

Deploying Virtual Machines

One of the primary advantages of working with virtual machines rather than their physical counterparts is the ease with which they can be created and deployed. Most data center administrators are well aware of the process of purchasing and receiving new physical computers. The process also involves placing machines in racks and ensuring that they're connected to the proper power and network connections.

Next steps involve the configuration of special hardware (such as PXE boot ROMs and RAID disk configurations). Then it's time to install and configure the operating system (OS). The last steps of the ritual generally involve the installation of any required software as well as all applicable security updates. This process is generally reserved for IT staff, because end users lack the expertise and resources to perform all these steps by themselves.

Benefits of Virtualization

As an exercise in contrast, let's look at the process for deploying a new virtual machine. If a systems administrator is starting from scratch, the administrator can create a new virtual machine and a couple of virtual hard disks in a matter of minutes. The new "system" will be ready for use as-is. More commonly, an administrator will want to connect the machine to one or more virtual networks (which might, in turn, have access to physical ones via the host's network adapters). Once this process is complete, it's time to start the OS installation process. The main benefits are obvious. All the steps can be performed remotely, and the process is quick and easy.

Creating a new virtual machine is even easier if an administrator does not need to start from nothing. That is, an existing virtual machine can be copied and deployed to a host. For example, if a new virtual machine running Windows Server 2003 is required, a systems administrator can simply copy the virtual machine and make any required changes in a matter of a few minutes. Some standard modifications will likely be required. For example, changes to security identifiers, the virtual machine's host name, and network settings might be necessary. The overall process is similar to that of deploying a new physical machine but can be completed much more quickly. Similarly, if a virtual machine needs to be relocated to another host computer (for capacity or other reasons), the machine can be quickly shut down so that its files can be copied or moved. It's easy to see the many benefits of this process. But there are drawbacks as well.

Determining Optimal Placement

Generally, IT environments will evaluate the needs of a particular service and application before they choose to purchase and deploy the necessary hardware. Typical considerations include CPU, memory, disk space, and network configuration. Additionally, the base OS and associated features should be taken into account. Often, IT staff need only consider one or a few different types of workloads when determining the requirements of the computer. The primary drawback is that it's likely that some resources will be left underutilized. However, from a planning standpoint, ensuring adequate capacity is simple.

In the world of virtualization, however, many different workloads are often combined on the same physical server. Although each workload runs in an independent environment, each virtual machine is ultimately competing for access to the same physical resources on the host computer. For example, if several virtual machines require heavy disk I/O, they could run into significant performance constraints when running on the same host server. Other resources, such as network bandwidth, CPU utilization, and total physical memory, can create scalability issues.

Resource-Based Placement

One of the primary drivers for virtualization adoption is that of increasing hardware utilization. Of course, it goes without saying that applications and services should still continue to function at acceptable levels (as defined by users). Therefore, the particular management challenge of determining the optimal host server on which to place a virtual machine becomes very important. IT organizations can begin the process by characterizing the hardware and software requirements of the workloads that they plan to virtualize. Table 4.1 provides an example.

Workload	CPU Utilization	Memory Utilization	Disk Utilization	Network Utilization
Public Web server	Low	Low	Low	High
Intranet application server	Medium	Medium	Low	Medium
CRM application server	Medium	High	Low	Low
CRM database server	High	High	High	Medium

Table 4.1: Estimating workload resource utilization to determine ideal virtual machine placement.

In Table 4.1, each workload is evaluated based on its expected resource usage. For simplicity, this table uses a fairly subjective relative weighting. A more scientific approach would include running performance monitoring tools to get hard numbers such as percentage of CPU utilization, amount of network bandwidth used (in Mb/sec), amount of memory used (in MB), and so on. Regardless of the approach, however, IT staff can evaluate the "compatibility" of various types of workloads. For example, a disk-intensive workload might be a good candidate to run alongside a virtual machine that performs few disk reads and writes but that needs a lot of CPU time.

Other Workload Characteristics

In addition to looking at straight hardware-based resource utilization, IT staff should take into account other workload characteristics. One of these is to look at the time-based usage of resources. Some applications and services might place the heaviest load on a server during the morning hours. For example, many network environments will see a large amount of authentication-related traffic during the morning. Other types of virtual machines might exhibit peak load during what are traditionally "off hours." For example, a database extract, transform, and load (ETL) process might run in the middle of the night. Figure 4.1 illustrates two examples. By taking these resource requirements into account, IT staff can make better decisions about which hosts can support which virtual machines.

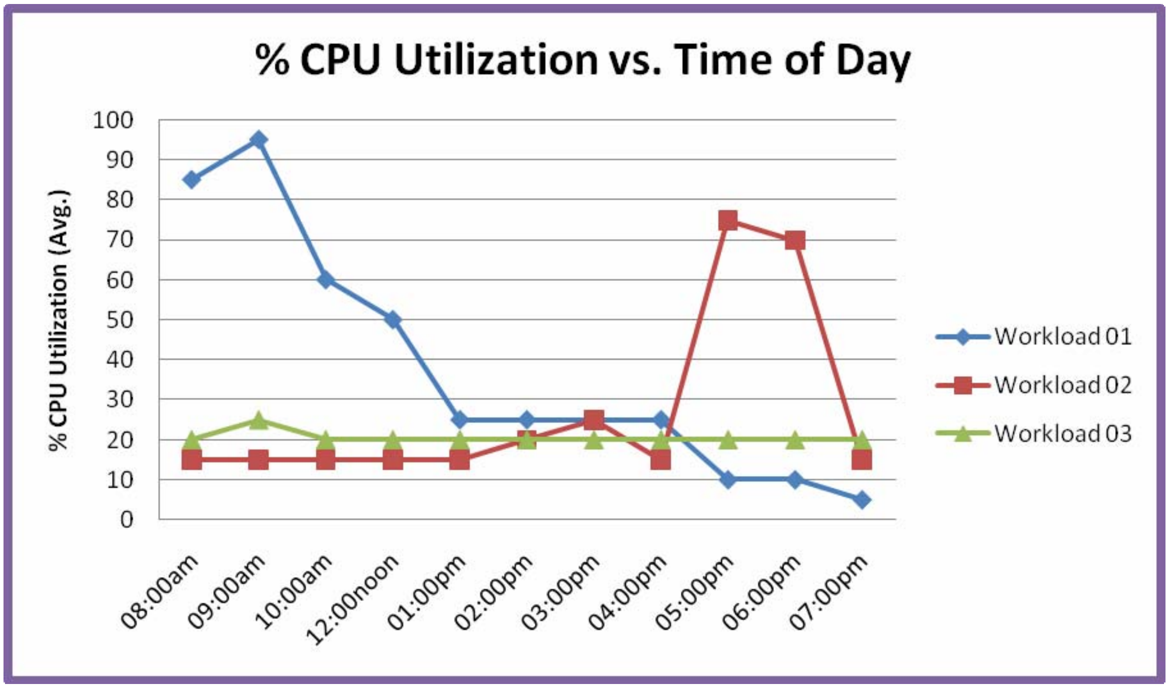

Figure 4.1: CPU utilization over time for several workloads.

When it comes to running business OSs, applications, and services, not all workloads have an equal importance. A production Customer Relationship Management (CRM) application might be considered mission critical. If this application is slow or unavailable, dozens of users might be unable to complete their jobs. A virtual machine running this application should clearly have a high importance. Other virtual machines might be significantly less important, and reduced performance will affect only a few users. For example, a development virtual machine that is used primarily by a small group of software testers should obviously have a lower priority. Figure 4.2 provides examples of typical types of workloads, arranged by importance.

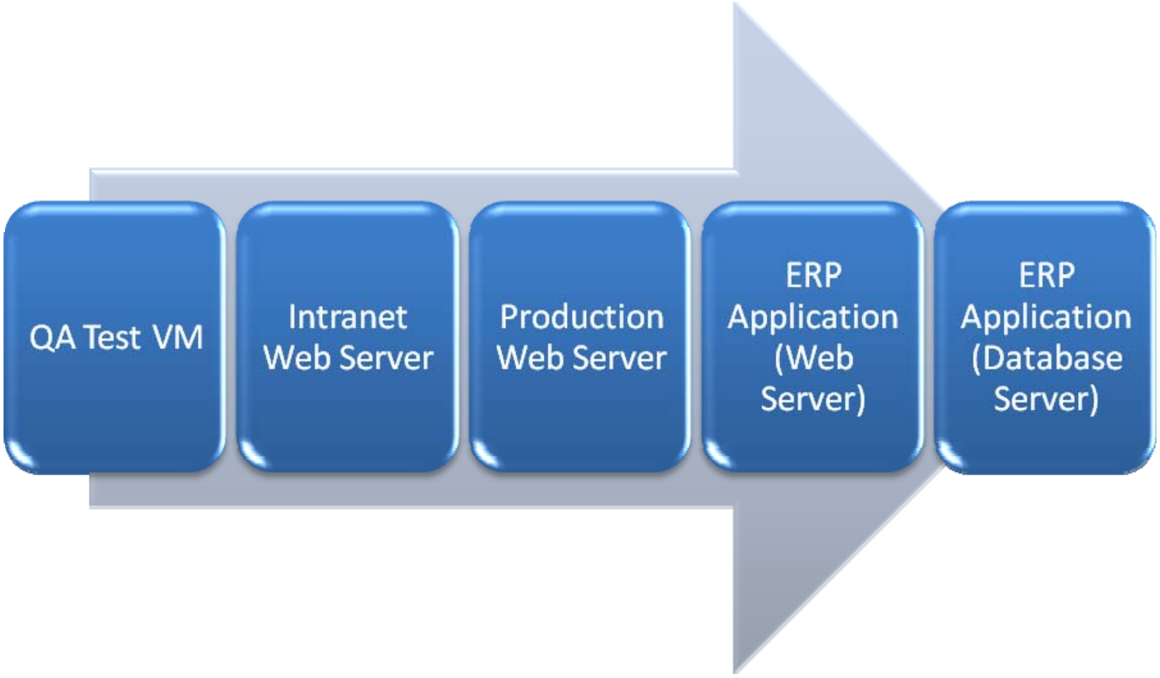
QA Test VM
Intranet Web Server
Production Web Server
ERP Application (Web Server)
ERP Application (Database Server)

Figure 4.2: Comparing the importance of virtual machines (increasing importance, from left to right).

In summary, IT organizations must take into account as much information as possible when determining the optimal placement for virtual machines. Although the technical process of moving or copying a virtual machine can be fairly simple, the process for making the best capacity-focused decisions requires a significant amount of time and attention.

Evaluating Virtualization Candidates

Once organizations have decided to incorporate virtualization solutions into their environments, the next important options are related to how to best implement it. Although virtual machine technology can provide dramatic benefits in a variety of areas, this is clearly not the ideal solution for all types of workloads. In some cases, hardware and software requirements might preclude the use of running within a virtual machine. So how should systems administrators decide which workloads are good candidates for virtualization?

Although there is no single technical answer to this question, one general rule is to use virtualization wherever you *can* and use dedicated physical servers wherever you *must*. This simple statement clearly places the preference on running in a virtual environment. For a wide array of different reasons discussed in earlier chapters, virtual machines can provide many deployment-related advantages. This changes the question to listing characteristics that might prevent running with a virtual machine. Some examples include:

- Resource requirements—Some applications scale almost linearly based on the capacity of the hardware configuration on which they run (see Figure 4.3). These workloads are able to take full advantage of the underlying hardware, so they are also most likely to see the effects of the performance penalty incurred by the addition of a virtualization layer. Also, many virtual machine layers are limited in the amount of physical memory they can expose to a virtual machine. Wherever performance is absolutely critical, IT departments might find that it is most appropriate to run them on dedicated computers.

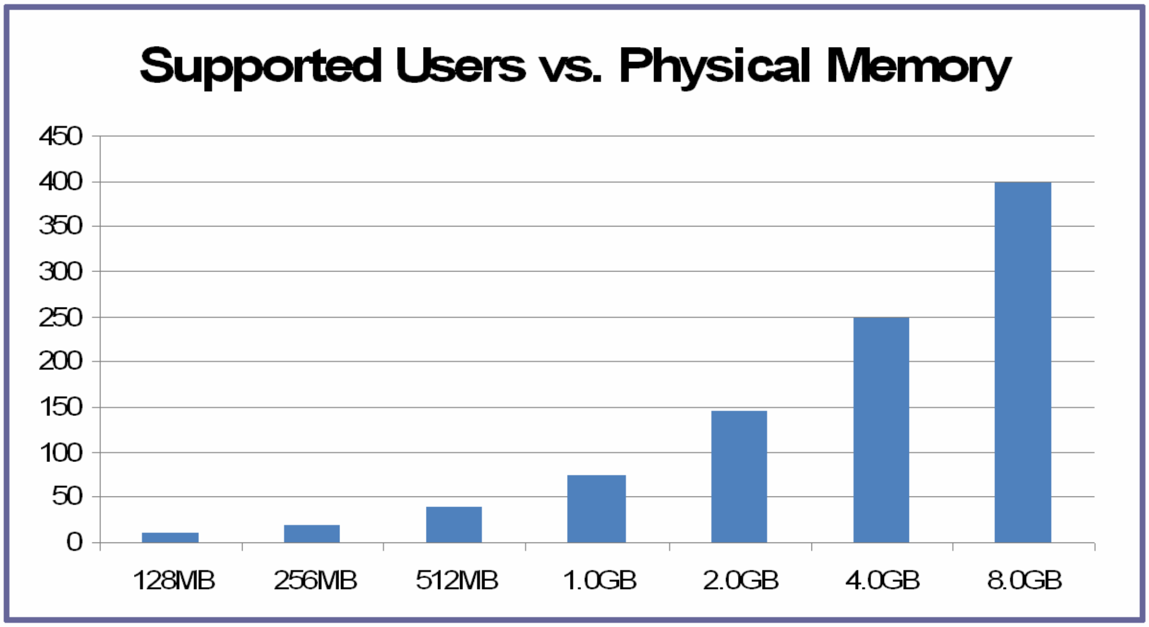

Figure 4.3: The number of supported users vs. the amount of physical memory for an example workload.

- Special hardware requirements—Most virtualization solutions provide an emulation layer that is compatible with a broad range of OSs and applications. Although the configurations may not be anywhere near state of the art, they generally provide a solid foundation upon which to install a guest OS. Some types of applications and technologies might have specialized requirements that don't fit neatly within these restrictions. For example, if an application requires 3-D acceleration, direct access to custom types of hardware, or specific Fibre Channel adapter cards, there is a good chance that these features are not available from within the virtual machine.

- Software compatibility—In some cases, the desired OS might not be supported on the virtualization platform. Or, even if the OS does appear to be working correctly in a virtual machine, the virtualization platform vendor may not support it. Although this might be less of a concern for testing, software development, and related scenarios, it should be an important concern when running mission-critical enterprise applications in a virtual machine. Specific details such as a particular OS version (along with service releases and other updates) should be evaluated against the officially supported configurations of the virtualization platform.
- Licensing and support issues—Although the entire IT industry is showing signs of moving towards embracing virtualization technology, some vendors may withhold technical support for applications that are running with a virtual machine. In some cases, there may be known incompatibilities with the emulated hardware platform. In other cases, the vendor may lack the time or expertise to adequately test it. Also, some licensing agreements can make it worthwhile to run directly on hardware. For example, if pricing is based on the number of physical CPU cores on the host computer, IT organizations might get more value by running directly on hardware.

Managing Physical and Virtual Assets

An old business aphorism states that, “If you can’t measure it, you can’t manage it.” In general, the idea is that you must have a good idea of statistics and details in order to be able to ensure the best business operations. This is certainly true in IT. IT departments and staff members are often under a significant amount of scrutiny, both from within and outside of their organizations. Internal business units want to ensure that their information is kept secure and that they’re reasonably well protected from malware and other threats. External entities such as regulatory auditors place additional requirements.

Regardless of the reasons, it’s important that IT departments are aware of all the systems in their environment. A potentially embarrassing situation occurs when a CTO claims compliance for all computers in the environment and subsequently finds several dozen computers that she was unaware of altogether. This section will look at challenges related to corralling and inventorying physical and virtual assets.

Server Sprawl

In many organizations, IT departments were historically seen as more of a cost center than as a strategic business partner. In these situations, the utility aspect of the IT team was stressed. For example, if a Human Resources manager thought he needed a new server to host an application, he would simply request a new physical server. The data center team would often deploy the new computer without even verifying the requirements. Often, end users and non-technical business managers would misinterpret the specifications they required. In some cases, systems were not powerful enough to support these workloads. In the vast majority of cases, however, the use of a new physical machine was overkill.

Over time, hardware platforms have grown dramatically with respect to shear computing capacity. For the vast majority of servers in a typical data center environment, it's difficult to find enough work to keep them busy. In fact, numerous industry surveys have shown that the majority of IT assets are significantly underutilized (often averaging 15% or less).

Regardless of this, systems and network administrators must take into account every machine within their environment when planning for monitoring and maintenance. For example, every physical computer is a potential vulnerability if it is not patched. It's often difficult enough to ensure that known machines are kept up to date with the latest standards. But it's impossible to verify the same for unknown assets. Clearly, it's important to be able to take an inventory of all the systems in the environment.

Virtual Machine Sprawl

It's somewhat ironic that one very popular solution to server sprawl—server consolidation through virtualization—has led to a related problem: virtual machine sprawl. There are numerous ways to detect physical computers. Although it may not be an elegant or glamorous process, one method is to simply perform a visual count of the number and types of machines that are present in each server rack.

In the virtual world, things are significantly more challenging. Virtual machines can be quickly and easily created (and moved) between physical server hosts, so it can be all but impossible to manually maintain a list of where these systems are located. Because each virtual machine is running its own independent OS, IT staff must keep track of these systems to ensure that they remain properly patched. In many cases, the virtual machines will participate on production networks and are potentially vulnerable to security threats, malware, and other issues.

Discovering Physical Servers and Virtual Machines

Going back to the original challenge—taking an inventory of current physical and virtual systems—it is clear that an automated solution is required. Many enterprise management solutions have the ability to perform an automatic "discovery" operation. Generally, this involves watching the network for traffic from various computers or initiating a scan. In the world of physical servers, the process is fairly reliable. It can identify servers, network devices, and even applications.

But what about virtual machines? IT organizations should specifically look for the ability to enumerate virtual machines on their host computers. Virtual machines can be deployed in a matter of minutes, so changes might occur every day. Additionally, some virtual machines might not have access to the network at all. In those cases, the only way of detecting them would be to query the virtualization services on the host computer.

Perhaps the most important point related to the discovery of virtual and physical assets is that it is not a one-time task. Systems and network administrators must be able to routinely run scans to detect changes and the presence of new assets.

Monitoring Virtual Machines

In the early days of virtualization technology, the most common configurations involved workstation "power users" running one or a few virtual machines for testing or application compatibility purposes. Although these virtual machines tended to be important to their users, downtime or data loss would rarely have a significant impact on the entire organization. Today, it's much more common to run mission-critical applications within a virtual environment. Hundreds or even thousands of users might rely on these services to complete their tasks. All this makes the task of ensuring reliability a high priority for IT departments. Let's look at some details.

Monitoring Uptime and Reliability

When managing physical servers, uptime and reliability can often be measured using percentage-based statistics. Table 4.2 provides an example of actual downtime for each level of "nines." Although most workloads would benefit from 100% uptime, real-world constraints often make this impossible. Tasks such as installing software and OS updates often require, at the very least, the reboot of the computer.

% Uptime	Downtime (minutes/month)	Downtime (minutes/year)
99%	438.00	5,280 (~88 hours)
99.9%	43.80	526.00 (~9 hours)
99.99%	4.38	52.60
99.999%	.44	5.26

Table 4.2: Comparing availability levels and acceptable amount of downtime.

There are several additional challenges that organizations face when measuring and ensuring availability of virtual machines. At a basic level, the same standards can be used. Tools that regularly poll for service availability or perform sample test transactions can help ensure that the components are working as expected. However, virtual machines can easily be moved to other host computers. A monitoring tool must take this into account. Additionally, when many virtual machines are running on a single physical server, the importance of protecting against downtime due to hardware issues is dramatically increased.

Finally, processes related to installing updates and maintaining virtual machines must be scheduled with respect to other virtual environments on the same server. If, for example, one virtual machine is rebooted during a particularly busy time on the host server, it could adversely impact performance on the other virtual machines. In this case, a monitoring tool that is unaware of the physical-to-virtual (P2V) relationship might point systems administrators to troubleshooting a virtual machine. The root cause of the performance issue, however, could be better correlated with other operations on the same host server.

Monitoring Performance

Well-managed IT environments try to perform as much proactive performance monitoring as possible. They try to detect potential issues as early as possible and are often aware of problems before users start to notice them. Although this can take significant time and effort, it's often a worthwhile endeavor in the long run as it leads to improved user satisfaction. Organizations can use a variety of tools to monitor performance characteristics of both physical and virtual machines.

Most OSs provide basic monitoring tools that can be used to collect and analyze system performance. In some cases, they can be helpful for troubleshooting isolated issues on a single or a few computers. In other cases, they can be used to track performance details over longer periods of time.

When monitoring performance in larger-scale environments, enterprise-ready software solutions can provide numerous advantages. For example, they provide methods for centrally monitoring dozens or hundreds of servers. In addition, they allow for the creation of alerts based on thresholds, along with ways to manage how alerts are communicated to IT staff.

> 📖 Later chapters will provide more detail about how organizations can use automated performance monitoring and optimization tools to better manage their environments.

Monitoring Host Servers

When adding virtual machines to the mix, there are some important performance relationships that must be kept in mind. There are two approaches that can be taken. The first is to start with monitoring physical computers. When monitoring a virtualization host, for example, a systems administrator might find that CPU utilization is often sustained at or near 100%.

In this case, it's likely that the system is either overworked or a CPU upgrade might improve the end users' experience. However, it isn't clear which workload specifically is causing the problem. The problem might be related to a process that is running on the host itself (such as a server agent or monitoring application). Or, the primary user of CPU time might be a specific virtual machine.

Monitoring Virtual Machines

Another approach to monitoring performance is to collect data from within virtual machines and then use this to determine resource requirements. Perhaps two virtual machines are trying to monopolize the host computer's physical adapter. Or, several workloads have the same periods of peak usage.

The virtualization management challenge is to be able to analyze and correlate performance information from both virtual and physical machines. Administrators should have an easy way of determining which virtual machines are trying to use the most resources and to compare this information with resource usage on the host computer. When working with a relatively small number of virtualization host computers, features such as CPU resource reservations and scheduling can be used.

Performance monitoring and management can become significantly more complicated when working in larger environments that contain dozens or hundreds of virtual machines and host computers. Manual monitoring and analysis are simply not efficient enough to quickly identify problems. In these situations, IT departments must look for tools and solutions that have the ability to automatically characterize virtual workloads and compare this information with the physical performance characteristics of the host computer.

Overall, the use of virtual and physical machines can make the process of monitoring workload performance significantly more complex. IT staff should take into account the P2V relationships when planning to administer virtual environments.

Change and Configuration Management

An important IT best practice is that of implementing and enforcing policies related to change and configuration management. Configuration management refers to the process of ensuring that systems are configured as they should be, based on technical and business requirements. This can go a long way toward ensuring that systems have not been inadvertently or maliciously modified. Change management involves ensuring that any modifications to an IT environment are properly tracked and that the necessary approvals have been obtained. Let’s look at each of these important topics, with a focus on how virtualization can affect these best practices.

Implementing Configuration Management

In the world of physical servers, this process can be performed by the use of standard network scanning tools and techniques. The best way to track all the information is within a centralized data store known as a Configuration Management Database (CMDB). Figure 4.4 provides an example of a CMDB that can store information about a variety of resources.

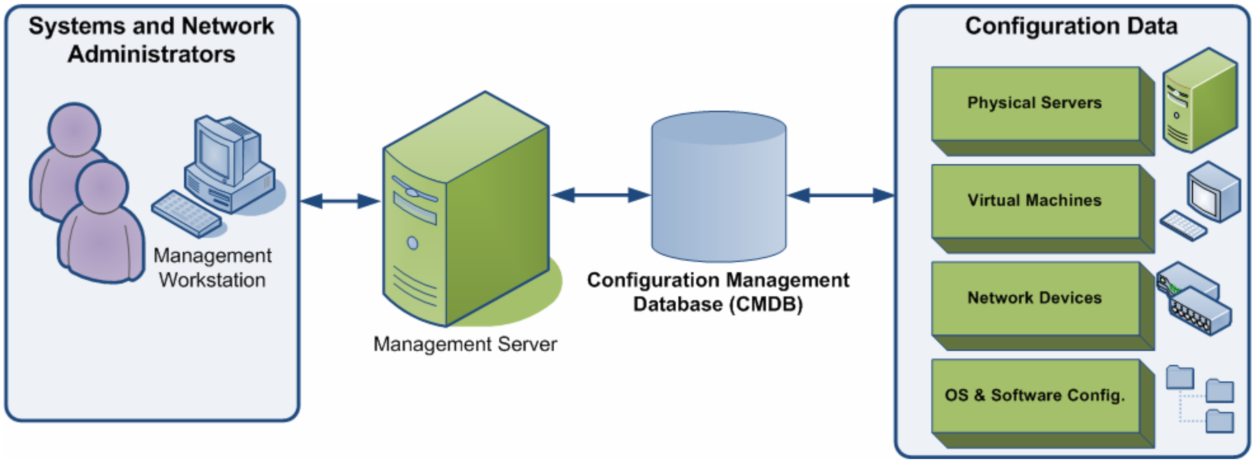

Figure 4.4: Implementing a CMDB.

Systems and network administrators can then use a management server (either through a Web or GUI-based interface) to query and analyze the information they need. The CMDB is responsible for storing a wide variety of details about the entire IT environment, including:

- Server hardware configurations
- Network device configuration files
- IP address range assignments on a router
- OS versions
- Installed software versions
- Services that are running on a particular server

Traditionally, IT departments have collected these types of data separately. For example, backups of firewall configuration files might have been stored on a secure file server. And server configuration information might have been recorded in a spreadsheet. Although the information can certainly be tracked this way, it makes the process of reporting much more difficult. Through the use of a CMDB and a suitable front-end application, the relationships between devices, applications, and services can be determined. All the data is stored in one place and in a consistent format, so IT managers can proactively identify potential problems.

There are numerous advantages of using a CMDB-based solution when managing virtualization. The most important requirement is for the solution to be able to adequately detect and store the relationship between a guest virtual machine and the host in which it is running. In many ways, virtual machines share characteristics of physical computers. They have a “hardware configuration” (though the hardware is virtually emulated), they run OSs, and they support applications and services. Details such as network addresses and host names are similar between physical machines and virtual machines.

Where virtualized workloads differ is in their relationship to host servers. This is where a CMDB solution must have special support for virtual machines. Virtual machines can be quickly moved between hosts, so a CMDB must be able to uniquely identify each virtual machine. When comparing resources, the CMDB should be able to make a distinction between virtual hard disks, virtual memory, virtual CPUs, and other devices. They should then be able to compare this data against their physical counterparts on the host system. Fortunately, all these features are certainly possible as long as the solution has been designed to support virtualization.

Implementing Change Management

As many IT professionals can report from experience, a large number of IT-related problems can be caused by improper configuration changes. In some cases, the changes are complete mistakes. For example, a junior-level database administrator might make a modification to memory parameters on a production database server without understanding all the potential effects the change may have. In other cases, the changes are "good," but they're not communicated to other members of the team. Modifications to the OS configuration of a new server, for example, can have a significant impact on server monitoring, maintenance, and troubleshooting processes.

From an IT process standpoint, technical staff may have agreed to implement some measures to prevent ad-hoc changes. For example, network administrators might be required to notify all potentially affected server administrators prior to changing rules on a router or firewall. In addition, a series of approvals might be required for major changes to the environment. When handled manually, however, it's difficult to ensure that these policies are being followed. In many cases, numerous administrators might have access to a particular device or application, and it can be difficult to track down who made which changes and when.

Automated change management tools can provide numerous advantages in IT environments. The ideal method of operation is through the use of a change management solution such as the one illustrated in Figure 4.5.

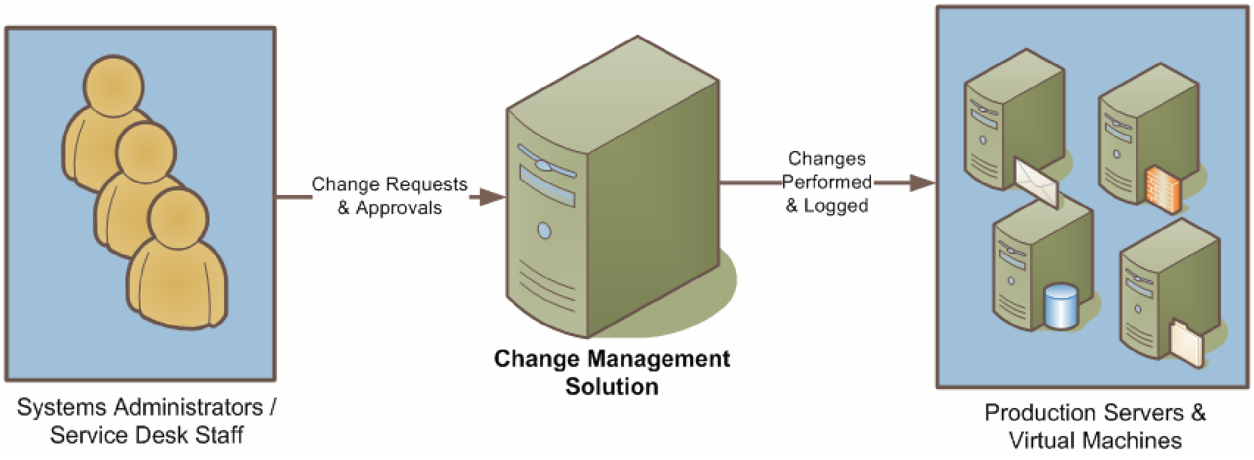

Figure 4.5: Making changes using a change management solution.

In this case, the change management application is directly responsible for carrying out authorized modifications to the environment. Rather than allowing network and systems administrators to directly commit modifications, the solution can ensure that the proper steps are first carried out. For example, several approvals might be required before a server can be removed from the environment. IT managers can define the overall workflow and requirements within the system. Change implementers can place requests within the system and be assured that they will be presented to the appropriate staff members for review.

When it's time to apply the change, the configuration management product is able to perform it directly. This approach provides numerous advantages. First, systems and network administrators no longer require direct access to devices such as servers, routers, and firewalls. This reduces potential security problems by minimizing the numbers of logins and passwords for a system.

Second, the changes themselves can be scheduled to occur during a period of low activity. The change management solution can verify that the modifications were successful and report on the results using a variety of notification mechanisms. Finally (and perhaps most important), all changes are automatically tracked in a database so that any authorized member of the team can access the details while performing troubleshooting.

In general, these same principles apply to managing virtual machines. For example, when changes are required to virtual OSs, applications, and services, a configuration management solution should be able to perform them on behalf of administrators. It's important that the solution is virtualization-aware so that it is able to make a distinction between physical and virtual systems. In some cases, IT departments might choose different changes for virtual machines, and the results might need to be carried out differently. Some changes might require host modifications and guest OS modifications to be coordinated.

Moving From Physical to Virtual (and Back)

For a variety of reasons, systems administrators might want to move workloads between virtual and physical environments. Usually, a particular OS, application, or service will run faster when running directly on the host OS rather than within a virtual machine. Additionally, as only one or a few different workloads are competing for resources, performance will be more predictable.

Although virtualization technology is constantly improving, there will always be some level of performance overhead that must be taken into account. This section will look at some of these conversion options, along with reasons why they might be required.

P2V Conversions

Earlier in this chapter, I mentioned factors to consider when selecting virtualization candidates. Once an IT department has realized the value of moving various workloads to virtual machines, it's time to go about carrying out that process. In some cases, it might make sense to deploy only new OSs and applications within virtual machines. In other cases, data center constraints and the benefits of server consolidation will place a priority on moving existing applications and services.

One approach to accomplishing this goal is to manually build the new virtual machine. This task often involves the creation of a new virtual machine, installation of a new OS, installation of the appropriate applications, and finally, configuration of the system. If all went perfectly, the virtual machine will be a nearly identical copy of the physical one. However, it's likely that some settings were overlooked. Or, in some cases, the resources and expertise to reinstall and reconfigure the application may be unavailable. In those cases, an automated P2V conversion tool can be very helpful.

P2V tools generally work by running either directly on a physical computer or on another machine in the environment. Some of these utilities require a physical computer to be taken offline and rebooted several times. Others can perform the operation "hot" (that is, while the server continues to run). In either case, the result will be a virtual hard disk that can be attached to a new virtual machine. Examples of P2V conversion tools include VMware's VMware Convertor (a free utility available from the company's Web site) and Platespin's PowerConvert product. Figure 4.6 shows an example of migrating several physical machines to a single virtualization host.

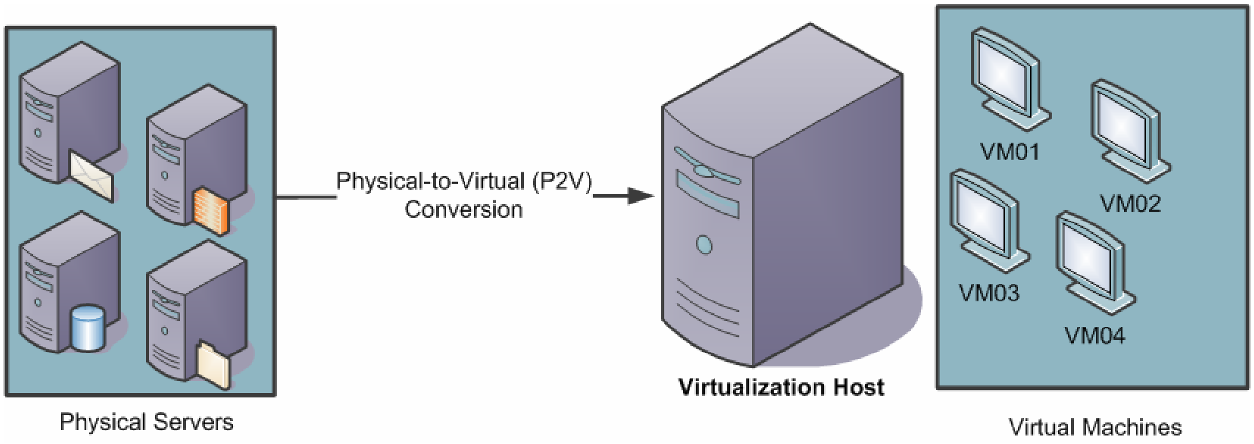

Figure 4.6: Performing a P2V conversion for server consolidation.

P2V tools must to be guest OS-aware in order to abstract the expected hardware configuration of the base OS. This process will differ significantly for a Windows Server 2003 virtual machine versus one running RedHat Enterprise Linux. Although no technology solution is perfect, it's especially important to keep in mind that not all conversions will run smoothly. IT departments should reserve adequate time and resources for testing the resulting virtual machine before deploying it into production.

V2P Conversions

Although it's generally not as common as moving a workload to a virtual machine, there are several reasons why an IT department might want to move a virtual machine to a physical server. The most obvious reason is related to performance and scalability. Perhaps an application or service has outgrown the constraints of running within a virtual machine and can truly take advantage of a dedicated physical server.

As with P2V conversions, the V2P process could be performed manually by reinstalling the OS and related applications. Third-party tools and utilities are available for performing the conversion, as well. The same P2V caveats related to testing also apply here.

V2V Migrations

With most technology solutions, IT departments have several options related to vendors and platforms. Virtualization technology is no exception. A virtual-to-virtual (V2V) migration refers to the process of moving a virtual machine from one virtualization platform to another. In some cases, this process might be required during an upgrade of the underlying virtualization services.

A new version might necessitate changes to the emulated hardware platform and/or the virtual hard disk platform. In other cases, the IT department may have standardized on a particular platform and might want to move existing virtual machines to it. Regardless of the reason, automated virtual machine conversion software might be available. At a basic level, the conversion requires that the virtual hard disk format be converted. Other details require the removal of any hardware-specific device drivers or configuration settings that might prevent the migration from being successful. Figure 4.7 provides an overview of the technical steps that must be carried out.

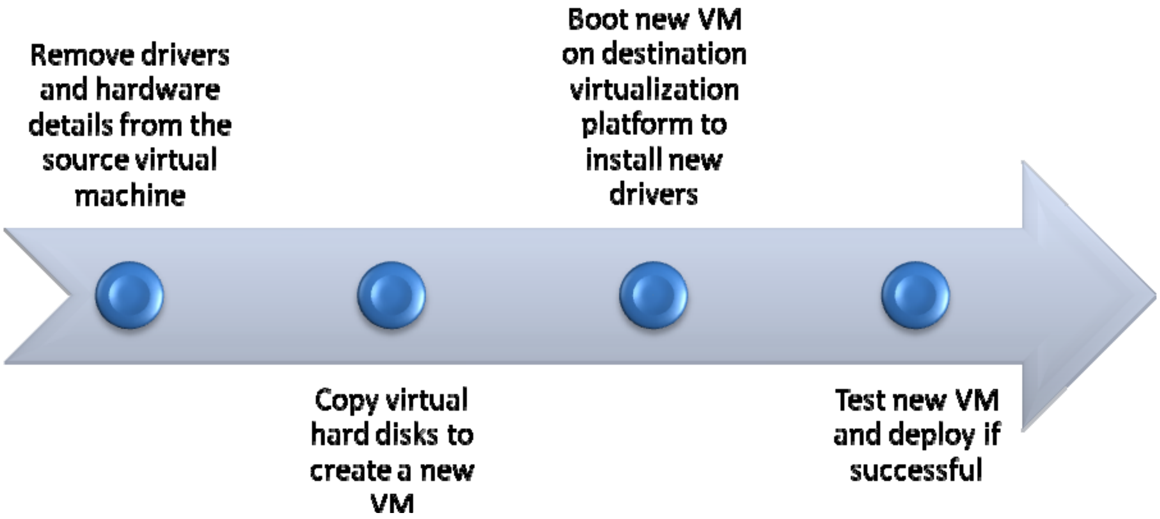
Remove drivers and hardware details from the source virtual machine
Copy virtual hard disks to create a new VM
Boot new VM on destination virtualization platform to install new drivers
Test new VM and deploy if successful

Figure 4.7: The process of performing a V2V migration.

Overall, the ability to move between physical and virtual machines can provide greatly improved flexibility and freedom in designing IT architectures. Of course, increasingly dynamic environments make it even more important for virtualization management solutions to be able to tell the difference between various platforms.

Additional Virtualization Management Challenges

In addition to the many system life cycle issues explored thus far, there are potential problems that should be considered by IT departments that are investing in virtualization. In some cases, the benefits of virtual machine technology can exacerbate current management issues. In other cases, the issues are completely new ones.

One important goal to keep in mind is that IT departments should remain involved in the creation and deployment of new virtual machines. Although users may have the technical expertise to build new virtual machines, they must understand the implications of doing so in a production environment. This section will look at additional concerns that systems administrators should consider.

Rapidly Changing Environments

It's no secret that virtual machines can be quickly and easily deployed. They can also be simply moved and copied across host servers. Additionally, users can download and install free desktop virtualization products such as Microsoft Virtual PC and VMware Server. Although these are primarily advantages of virtualization technology, they can make it very difficult for IT departments to keep track of the systems they manage. When compared with similar issues for physical servers, virtual machines are much more fluid in their design.

As virtual machines migrate between physical host computers, for example, associated network issues will need to be addressed. One such case is when a virtual machine is transitioned from a development server to a production one. It's likely that the IT department has specified different security and configuration standards for production computers versus those that are run in a development environment. One method of managing this is to ensure that production host computers are locked down so that only authorized administrators are able to deploy new virtual machines. All new virtual machines that are scheduled for deployment should undergo an IT process that ensures that they adhere to requirements.

Standardization

Most IT departments have developed standards to help reduce the complexity inherent in managing heterogeneous environments. The list of approved technologies might include a list of supported software, OSs, services, and hardware platforms. Generally, these decisions are factored into purchasing and deployment plans to ensure that only appropriate configurations are put into production. In the virtual world, however, users and administrators can run an almost limitless range of OSs, versions, and software. As long as their choices are supported by the emulated hardware platform, they can then deploy these systems into production or other environments.

To help combat this problem, IT departments must specify standards for production deployments. Although the list of supported software and OSs might be larger for virtual environments, configurations should always be tested. Users and systems administrators should verify that they are running a supported platform. Automated change and configuration management tools can help audit current virtual machines and highlight any systems that do not adhere to IT standards and specifications.

Centralized Management

The goal of centralized management is to ensure that an organization's data center devices can be managed remotely, ideally from within a single system. Generally, administration for a data center should involve minimal direct access to a server or network device console. In many data centers, the only reason to physically visit a machine is to manage the hardware configuration or to change physical network connections. This approach applies equally to supporting distributed environments such as branch offices, mobile users, and employees that work from home. In all of these cases, remote management tools can dramatically reduce administrative costs.

The same centralized management approach can greatly simplify access to virtual machines. After the discovery of host servers and virtual machines has been performed, administrators can connect to them over the network to carry out administrative tasks. A centralized management system that is virtualization-aware will allow administrators to clearly identify and locate which virtual machines are running on which physical servers. In some cases, virtual machines might not have direct access to a network environment at all. In those cases, it's usually necessary to connect first to the host computer and then to use either a GUI- or Web-based management interface to perform tasks.

Managing Security

For most practical purposes, virtual machine security should be on par with physical machine security. Common IT best practices such as relying upon directory services for authentication and centralized security are important concerns. Especially for production servers, systems administrators should ensure that each virtual machine is configured with a minimal set of permissions. Virtual machines should run only services and features that are required to support their particular workload. Management utilities should be able to identify physical and virtual systems that are out of compliance with corporate standards and allow a way to easily resolve the issue if possible.

Keeping Virtual Machines Updated

If you were to ask a random sample of systems administrators to rank tasks that they prefer to perform, it's likely that performing updates and managing patches will appear close to the bottom. When done manually, the process is tedious, time-consuming, and risky. Often, administrators need to determine which patches should be applied. They then must test them for compatibility. When it comes time to deploy the patches, they must minimize downtime and disruption to service.

Modern OSs and production environments are potentially vulnerable to a wide array of threats and attacks. Some potential issues are related to known vulnerabilities in commonly used software. Unauthorized users from both within and outside an environment can attempt to perform actions such as deleting or modifying critical system files, obtaining sensitive information, or launching Denial of Service (DoS) attacks. Additionally, issues related to malware can cause system stability and performance issues. Clearly, there is a need to keep all systems—whether physical or virtual—up to date. The potential risks and liabilities can be significant, and this is a critical operation for all IT organizations.

Virtual machines present some additional challenges. One issue is that virtual machines may not always be present and responding on the network. For example, development teams might use a set of virtualization servers only during the later stages of testing. Often, these virtual machines are neglected because they don't appear to be on the network. IT departments should keep this situation in mind when trying to keep virtual systems updated. Another virtual machine-related issue is that often the number of running OSs can be extremely large (as mentioned earlier, related to "virtual machine sprawl"). That makes manual management all but impossible.

Implementing Virtual Machine Backups

Data protection concerns increase as critical resources and services are moved to virtual environments. Like their physical counterparts, virtual machines must be backed up. There are two primary approaches to performing backups of a virtual machine. The first is to treat the virtual machine like a physical server. Usually, this involves the installation of a backup agent that is able to communicate with a central management station. Systems administrators can easily specify which data must be backed up, and where it will be stored. Storage destinations might include the host file system or a central network-based storage server or device. Regardless of the approach, the IT department must have a good understanding of which information must be protected, based on business requirements related to data loss.

The other approach for performing backups is from within the host file system. Administrators can identify all the files that are related to important virtual machines and schedule them to be copied. Generally, the list includes the following types of files:

- Host server configuration data
- Virtual hard disks
- Virtual machine configuration files
- Virtual network configuration files
- Saved-state files

The primary challenge is that several of these files are often locked as in-use. Of particular importance are virtual hard disks, which must be taken offline to get a "clean" backup. Figure 4.8 provides an overview of the steps that are generally required to complete this "warm backup" process. Though there are several steps, they can be automated, resulting in a minimal period of downtime.

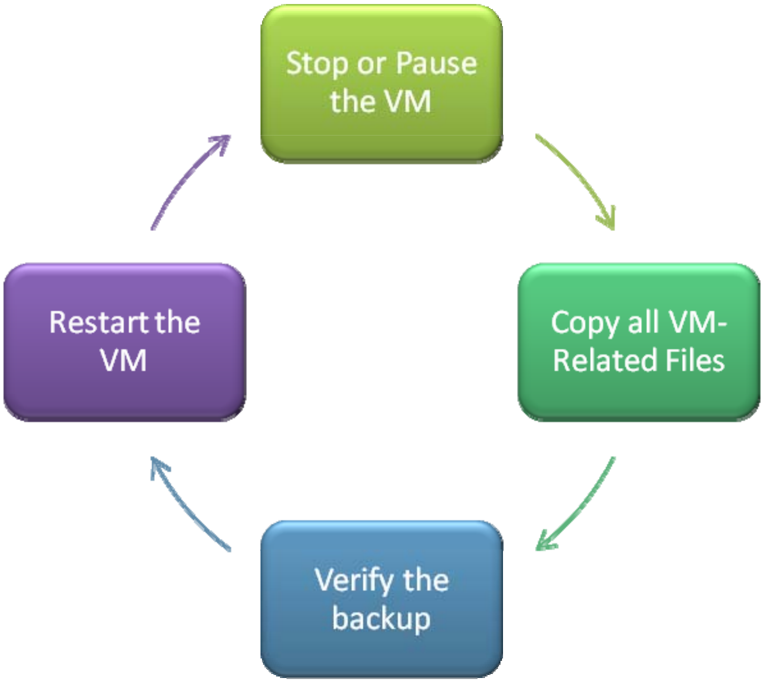

Figure 4.8: Implementing "warm backups" of virtual machines.

Finally, third-party utilities and features are available for performing what is often referred to as a "hot backup." These backups can be completed while a virtual machine is running. Solutions can use a snapshot-based method to access the files while they are in use, thereby eliminating downtime. IT departments should keep in mind that the added disk I/O during the backup process might decrease overall virtual machine performance. Figure 4.9 summarizes these four main approaches.

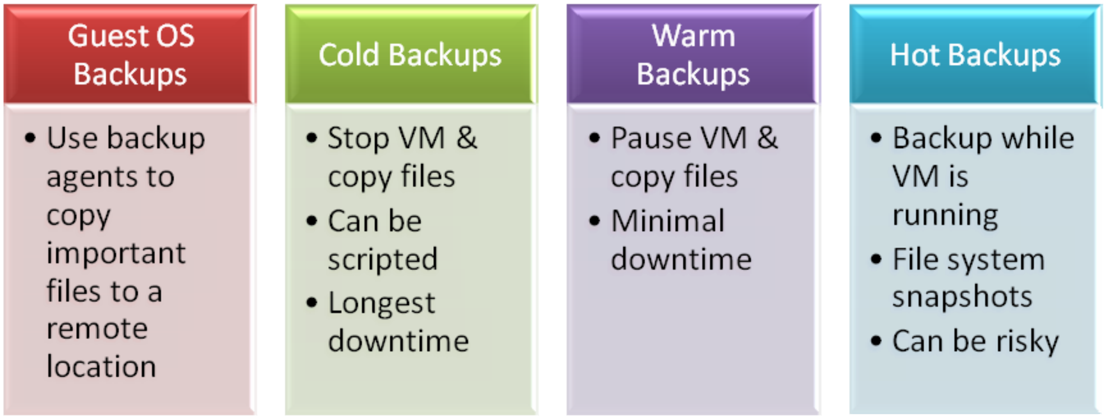

Guest OS Backups	Cold Backups	Warm Backups	Hot Backups
• Use backup agents to copy important files to a remote location	• Stop VM & copy files • Can be scripted • Longest downtime	• Pause VM & copy files • Minimal downtime	• Backup while VM is running • File system snapshots • Can be risky

Figure 4.9: Comparing virtualization backup approaches.

Managing Licensing

Software and OS licensing represents a major cost for most IT departments. To remain in compliance with purchased licenses, systems administrators must be able to perform a complete inventory of physical and virtual system configurations. Tracking these details on physical machines is fairly easy because configuration management utilities can often report on the details.

Virtualization presents several additional challenges, however. Often, end users will make a copy of their virtual machines for a variety of different reasons. Depending on details of vendors' licensing policies, the organization may be required to purchase additional licenses to support this. This puts the burden on IT departments to educate users about the total costs of deploying new virtual machines. Recently, many vendors have provided special provisions for virtualized configurations. If such information is not available, IT departments will be safest by licensing virtualized software and OSs as if they were running on physical computers.

An Ideal Virtualized Environment

Throughout, this chapter has focused on identifying particular management challenges that tend to arise when IT organizations move toward virtualization. Although the list of potential issues is a long one, there are numerous ways in which enterprise management tools can help address them. Let's wrap up the discussion by envisioning a few of the characteristics of an ideal virtualized environment.

Technology Independence

The overall goal for IT managers should be to create an environment in which they can seamlessly manage both virtual and physical assets. It's generally a given that more variations in technology can lead to more difficult management challenges. To counter this, administrators should be able to use the same policies, tools, and technologies to enhance the operations of workloads that run on physical servers and those that run within virtual machines. In the end, the focus should be on the user experience. IT manages a wide array of physical servers and virtual platforms from which to select. If the technology team is doing its job properly, internal and external "customers" will not care how the systems are implemented. The organization as a whole will benefit from increased efficiency and better utilization of existing physical assets.

Simplified Movement of Workloads

Earlier, this chapter mentioned methods by which systems administrators and IT staff can move workloads between physical and virtual environments. This flexibility is a crucial part of creating a fluid data center that includes a pool of readily available resources. Whether the underlying system is a high-end multi-core system, a mid-range rack-mounted server, or a blade configuration, the required applications, OSs, and services should be portable.

Automated Administration

From the standpoint of keeping systems up to date, the same administration practices should apply to virtual and physical systems. After all, there are many similarities in the OS and application configurations. The sheer volume of virtual machines can present additional issues, so organizations that haven't done so already will need to evaluate and deploy enterprise management systems that are virtualization-aware.

Summary

This chapter covers a lot of ground related to real-world issues raised by adopting virtualization technology. It began by detailing the process of virtual machine deployment, including ways in which organizations can determine the optimal placement for a particular workload. Considerations include hardware resource requirements, usage patterns, and business requirements. Based on these details, IT teams can determine whether a particular workload is suitable for virtualization.

Next, the chapter looked at the process of managing physical servers and virtual machines, after they've been deployed. A familiar issue in many IT departments is that of "server sprawl"—the phenomenon that results in a large number of underutilized computers that must be managed. Although virtualization can help in the area of sever consolidation, a related problem—"virtual machine sprawl"—has become a reality in many environments. By implementing auto-discovery, IT departments can get a better grasp on what is running throughout their environments.

Availability and reliability have always been important concerns for IT departments. The use of virtual machines often increases the number and types of systems that must be supported. A standard IT best practice is to implement change and configuration management policies and processes to ensure that physical and virtual assets are within compliance with IT standards. I presented ways in which automated solutions could help.

The ideal application of virtualization technology will vary based on the type of workload and the particular applications and services that must be delivered to users. I covered details of moving from physical servers to virtual machines (and vice versa) and the process of converting between different virtualization platforms. Overall, this flexibility can provide IT departments with the best mix of various technology options.

The list of virtualization management challenges is rounded out with some additional concerns. Keeping up with rapidly changing environments and implementing standardization can be daunting tasks. This is especially true for environments that rely on manual administration. Other issues include managing security, keeping physical and virtual systems updated, and providing for virtual machine backups.

Finally, I covered some characteristics of a hypothetical "ideal" virtualized environment. In some cases, the management issues can be summed up as "more of the same." Standard IT practices apply equally to physical and virtual assets. In other cases, the use of virtual machines can amplify certain issues that departments were already struggling with.

The good news is that there are several ways in which these problems can be addressed. I'll dive into those in upcoming chapters. The primary goal should be to take advantage of the many benefits of virtualization technology and to minimize the numbers and types of management headaches.

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.