Realtime publishers

"Leading the Conversation"

The Reference Guide™ *To*

Data Center Automation

sponsored by

OPSWARE INC
Automating IT™

Don Jones and Anil Desai

Copyright Statement

Server Virtualization

Virtualization refers to the abstraction between the underlying physical components of an IT architecture and how it appears to users and other devices. The term virtualization can be applied to network devices, storage environments, databases, other portions of an IT infrastructure, and servers. Simply put, server virtualization is the ability to run multiple independent operating systems (OSs) concurrently on the same hardware.

Understanding Virtualization

The concept of running multiple "virtual machines" on a single computer can be traced back to the days of mainframes. In that architecture, many individual computing environments or sessions can be created on a single large computer. Although each session runs in what seems like an isolated space, the underlying management software and hardware translates users' requests and commands so that users can access the same physical hardware. The benefits include scalability (many virtual machines can run simultaneously on the same hardware) and manageability (most administration is handled centrally and client-side hardware requirements are minimal).

Current Data Center Challenges

Before diving into the technical details of virtual machines and how they work, let's set the foundation by exploring the background for why virtualization has quickly become an important option for data center administrators. The main issue is that of server utilization—or lack thereof. The vast majority of computers in most data centers run at a fraction of their overall potential (often as little as 10 to 15 percent). The obvious solution is server consolidation: Placing multiple applications on the same hardware. However, due to the complexity of many environments, potentials for conflicts can make server consolidation difficult if not impossible. One of the many benefits of virtualization is that it allows systems administrators to easily create multiple virtual operating environments on a single server system, thereby simplifying server consolidation.

Virtualization Architecture

For modern computing environments, virtualization solutions can be quickly and easily installed on standard hardware. Figure 7.1 shows a generic example of one way in which virtualization can be implemented.

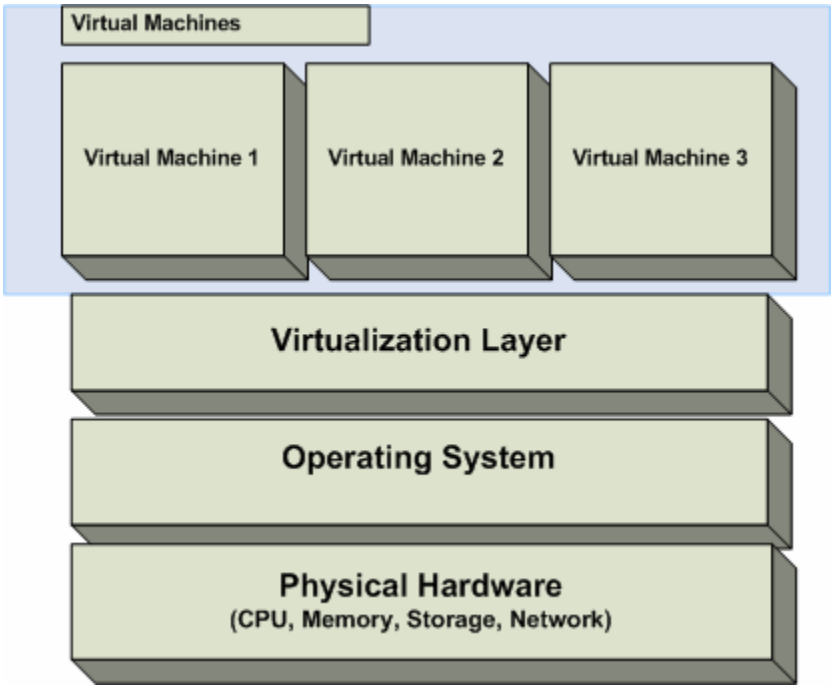

Figure 7.1: A logical overview of virtualization.

At the bottom of the figure is the actual physical hardware—the CPU, memory, hard disks, network adapters, and other components that make up the complete system. Running atop the hardware is the OS, which includes device drivers that interact with physical system components. Moving up the stack, within the OS is a virtualization management layer. This layer allows for the creation of multiple independent virtual machine environments. The virtualization layer may run as an application or as a service (depending on the product). Finally, at the top of the "stack" are the virtual machines. It is at this level that multiple OSs can run simultaneously.

The job of the virtualization layer is to translate and coordinate calls from within each virtual machine to and from the underlying hardware. For example, if the Linux-based OS within a virtual machine requests access to a file, the virtualization management application translates the request and redirects it to the actual file that represents a virtual hard drive on the host file system. Figure 7.2 shows an example of how a Microsoft Virtual Server 2005-based virtualization stack might look.

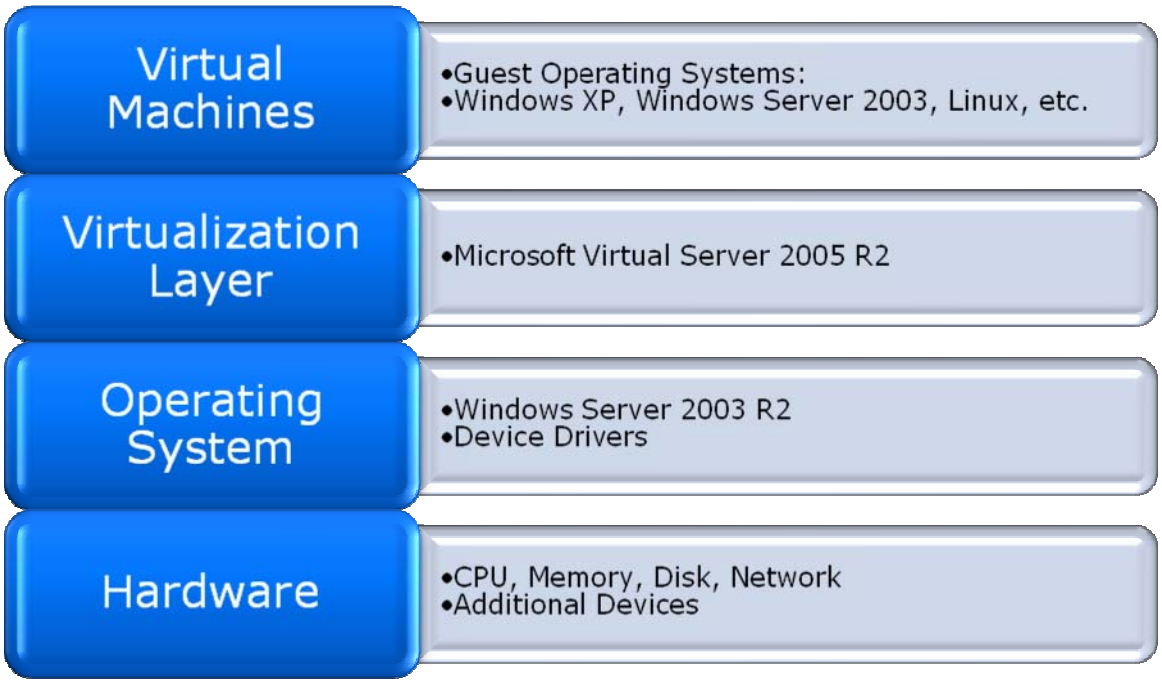

Figure 7.2: An example of a virtualization configuration using Microsoft Virtual Server 2005 R2.

Virtualization Terminology

Virtualization provides new ways in which to refer to standard computer resources, so it's important to keep in mind some basic terminology. The physical computer on which the virtualization platform is running is known as the host computer and the primary OS is referred to as the host OS. The OSs that run on top of the virtualization platform are known as guest OSs.

An additional concept to keep in mind is the virtual hard disk. From the perspective of the guest OS, these files appear to be actual physical hard disks. However, physically, they're stored as files within the host OS file system.

Finally, another major advantage of virtual machines is that they can be "rolled back" to a previous state. This is done by keeping track of all write operations and storing them in a file that is separate from the primary virtual hard disk.

Other Virtualization Approaches

It's important to note that, in addition to the OS-based virtualization layer shown in Figure 7.1, there are other virtualization approaches. In one such approach, the virtualization layer can run directly on the hardware itself. This model (also referred to as a "Hypervisor") offers the advantage of avoiding the overhead related to running a primary host OS. The drawbacks, however, include more specific requirements for device drivers and the potential lack of management software.

Another virtualization approach is "application-level virtualization." In this configuration, application environments are virtualized—in contrast with running entire OSs. The main benefit is that scalability can be dramatically improved—often hundreds of applications can run simultaneously on a single physical server. There are drawbacks, however; some complex applications might not be supported or might require modifications. In addition, OS versions, device drivers, updates, and settings will affect all virtual environments because they're defined at the machine level.

The following sections focus on the type of virtualization described in Figure 7.1.

Benefits of Virtualization

The list of benefits related to working with virtual machines is a long one. Let's take a brief look at some of the most relevant advantages from the standpoint of data center management:

- Increased hardware utilization—By allowing multiple virtual machines to run concurrently on a single server, overall resource utilization can be dramatically improved. This benefit can lead to dramatic cost reductions in data center environments, without significant costs for upgrading current hardware.
- Hardware independence—One of the major challenges related to managing data center environments is dealing with heterogeneous hardware configurations. Although it's easy to physically relocate an array of hard disks to another machine, chances are good that OS and device driver differences will prevent it from working smoothly (if at all). On a given virtualization platform, however, virtual machines will use a standardized virtual environment that will stay constant regardless of the physical hardware configuration.
- Load-balancing and portability—Guest OSs are designed for compatibility with the virtualization platform (and not the underlying hardware), so they can easily be moved between host computers. This process can allow users and systems administrators to easily make copies of entire virtual machines or to rebalance them based on overall server load. Figure 7.3 provides an illustration. This method allows systems administrators to optimize performance as business and performance needs change over time. In addition, it's far easier than manually moving applications or reallocating physical servers.

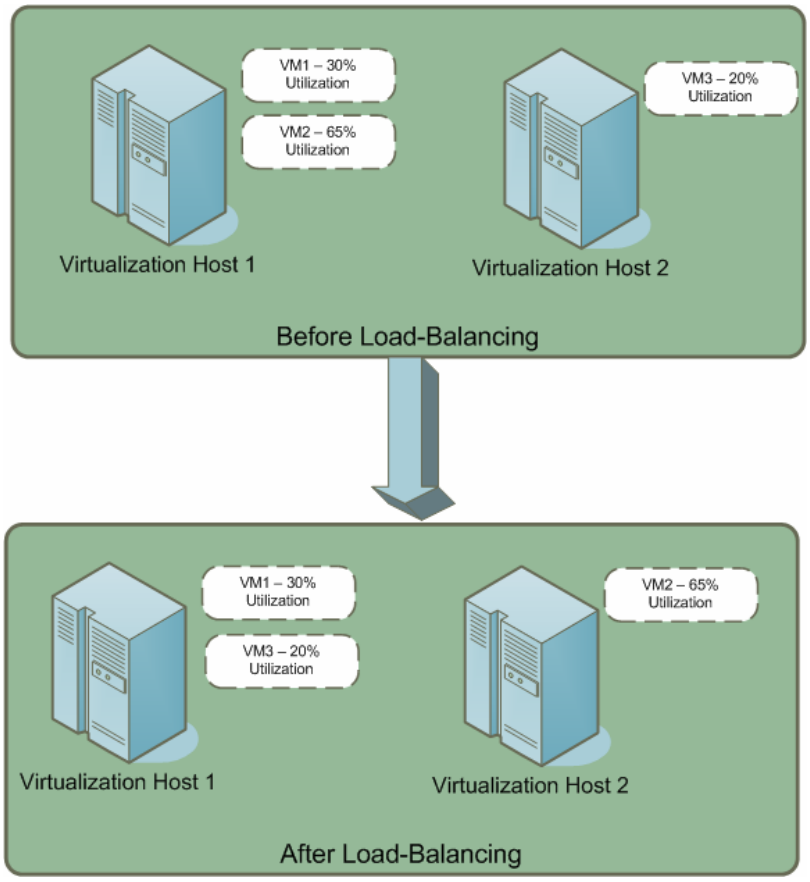

Figure 7.3: Load-balancing of virtual machines based on utilization.

- Rapid provisioning—New virtual machines can be set up in a matter of minutes, and hardware changes (such as the addition of a virtual hard disk or network interface) can be performed in a matter of seconds. When compared with the process of procuring new hardware, rack-mounting the devices, and performing the entire installation process, provisioning and deploying virtual machines usually takes just a small fraction of the time of deploying new hardware.
- Backup and disaster recovery—The process of creating a complete backup of a virtual machine can be quicker and easier than backing up a physical machine. This process also lends itself well to the creation and maintenance of a disaster recovery site.

Virtualization Scenarios

Earlier, we mentioned how virtualization can help data center administrators in the area of server consolidation. This, however, is only one of the many ways in which this technology can be used. Others include:

- Agile management—As virtual machines can be created, reconfigured, copied, and moved far more easily than can physical servers, virtualization technology can help IT departments remain flexible enough to accommodate rapid changes.
- Support for legacy applications—IT departments are commonly stuck with supporting older servers because applications require OSs that can't run on newer hardware. The result is higher support costs and decreased reliability. By placing these application within a virtual machine, the application can be moved to newer hardware while still running on an older OS.
- Software development and testing—Developers and testers often require the ability to test their software in many configurations. Virtual machines can easily be created for this purpose. It's easy to copy virtual machines to make, for example, changes to the service pack level. Additionally, whenever a test is complete, the virtual machine can be reverted to its original state to start the process again.
- Training—Dozens of virtual machines can be hosted on just a few physical servers, and trainers can easily roll back changes before or after classes. Students can access their virtual machines using low-end client terminals or even over the Internet. Usually, it's far easier to maintain a few host servers than it is to maintain dozens of client workstations.

Limitations of Virtualization

Despite the many benefits and applications of virtualization technology, there are scenarios in which this approach might not be the perfect solution. The first and foremost concern for most systems administrators is that of performance. All virtualization solutions will include some level of overhead due to the translation of hardware calls between each virtual machine and physical hardware device. Furthermore, virtual machines are unaware of each other, so competition for resources such as CPU, memory, disk, and network devices can become quite high. Overall, for many types of applications and services, organizations will likely find that the many benefits of virtualization will outweigh the performance hit. The key point is that IT departments should do as much performance testing as possible before rolling out virtualized applications.

There are additional considerations to keep in mind. For example, for physical servers that are currently running at or near capacity, it might make more sense to leave those systems as they are. The same goes for complex multi-tier applications that may be optimized for a very specific hardware configuration. Additionally, for applications require custom hardware that is not supported by the virtualization platform (for example, 3-D video acceleration), running within a virtual machine will not be an option. Over time, virtualization solutions will include increasing levels of hardware support, but in the mean time, it's important to test and verify your requirements before going live with virtualization.

Automating Virtual Machine Management

In many ways, IT environments should treat virtual machines just like physical ones. Virtual machines should be regularly patched, monitored, and backed up and should adhere to standard IT best practices. This leads to the issue of automating the management of virtualization solutions. IT departments should look for tools that are virtualization-aware. Specifically, these solutions should be able to discern which virtual machines are running on which hosts systems. Ideally, virtualization management tools should be integrated with other data center automation features such as change and configuration management and performance monitoring and should coordinate with IT policies and processes.

Developers can also automate virtual machine management. Most virtualization solutions provide an Application Programming Interface (API) that allows for basic automation of virtual machines. You can generally write simple scripts that enable tasks such as creating new virtual machines, starting and stopping virtual machines, and moving virtual machines to other computers. More complex programs can also be created.

Overall, through the use of virtualization technology, IT departments can realize numerous benefits such as increased hardware utilization and improved management of computer resources. And, through the use of automation, they can ensure that virtual machines are managed as well as physical ones.

Remote/Branch Office Management

In an ideal world, all of an organization's technical and human resources would be located within a single building or location. Everything would be within arm's reach, and systems administrators would be able to easily access all their resources from a single data center. The reality for all but the smallest of organizations, however, is that it's vital to be able to support a distributed environment. The specifics can range from regional offices to home offices to traveling "road warriors." In all cases, it's important to ensure that users can get the information they need and that all IT assets are properly managed.

Challenges of Remote Management

Before delving into the details of automating remote management, it will be helpful to discuss the major challenges related to performing these tasks. The overall goal is for IT departments to ensure consistency in managing resources that reside in the corporate data center as well as resources that might be located in a small office on the other side of the planet. Let's look at some details.

Technical Issues

In some ways, technology has come to the rescue: network bandwidth is more readily available (and at a lower cost) than it has been in the past, and establishing physical network connectivity is usually fairly simple. In other ways, improvements in technology have come with a slew of new problems. Storage requirements often grow at a pace that far exceeds the capacity of devices. In addition, almost all employees of modern organizations have grown accustomed to high-bandwidth, low-latency network connections regardless of their locations. IT departments must meet these demands while working within budget and resource constraints.

Perhaps one of the most pertinent issues related to remote office management is that of network bandwidth. Usually the total amount of bandwidth is constrained, and factors such as latency must be taken into account. This process has often lead to remote office systems being less frequently updated. Servers sitting in a wiring closet of a branch office are often neglected and don't get the attention they deserve. The result is systems that are likely out of compliance with IT policies and standards.

Personnel Issues

Ideally, organizations would be able to place senior-level systems and network administrators at each remote office. Unfortunately, cost considerations almost always prohibit this. Therefore, certain tasks must be performed manually (and often by less-trained individuals). Common tasks include the installation of security updates or the management of backup media. Dedicated technical staff is not available, so it's common for these important operations to be overlooked or to be performed improperly. Even when using remote management tools, some tasks cannot easily be accomplished from a remote location.

Business Issues

Functions served by remote offices can be mission critical for many of an organization's operations. From a business standpoint, new initiatives and changes in standard operating procedures must apply through the entire organization. The concept of "out of sight, out of mind" simply is not acceptable for remote locations. All of the hardware, software, and network devices that are under IT's supervision must be maintained to ensure overall reliability and security.

Automating Remote Office Management

Clearly, the task of managing remote locations and resources can be a difficult one. There is some good news, however: data center automation solutions can make the entire process significantly easier and much more efficient. IT departments that need to support remote offices should look for several features and capabilities in the solutions that they select:

- Change and configuration management—Keeping track of the purpose, function, and configuration of remote resources is extremely important in distributed environments. Often, physically walking up to a specific computer just isn't an option, so the data must be accurate and up to date. Whenever changes are required, an automated solution can efficiently distribute them to all the IT department's resources. In addition, they can keep a record of which changes were made and who made them. Doing so helps ensure that no devices are overlooked and can help avoid many common problems.
- Use of a configuration management database (CMDB)—Collecting and maintaining information across WAN links in distributed environments can require a lot of bandwidth. When IT managers need to generate reports, it's often unacceptable to wait to query all the devices individually. A CMDB can centrally store all the important technical details of the entire distributed environment and can facilitate quick access to the details.
- Notifications—In fully staffed data centers, trained support staff is usually available to resolve issues around-the-clock. For remote offices, however, an automated solution must be able to notify the appropriate personnel about any problems that might have occurred. In addition to IT staff, those alerted might include the branch manager or other contacts at the remote site.
- Monitoring—The server and network resources that reside in remote offices are often critical to the users in those offices. If domain controllers, database servers, routers, or firewalls become unavailable, dozens or hundreds of users might be unable to complete their job functions. Furthermore, staff at these locations might be unqualified to accurately diagnose a problem and determine its root cause. Therefore, it's important for computing devices and resources to be closely monitored at all times.
- Scheduling—When supporting remote sites that are located in distant locations, factors such as time zones and normal work hours must be taken into account. When performing tasks such as applying updates, it's important to have the ability to specify when the changes should be committed. The main benefit is the ability to minimize disruptions to normal activity without placing an unnecessary burden on IT staff.
- Support for low-bandwidth and unreliable connections—Remote sites will have varying levels of network capacity and reliability. The automation solution must be able to adapt to and accommodate situations such as the failure of a connection during an important update or the application of security changes as soon as network connections become available again. Also, client agents should be able to automatically detect low-bandwidth states and reduce the number and length of messages that are sent accordingly.

In addition, most of the best practices covered through this guide also apply to remote sites. By incorporating all these features in an IT automation solution, organizations can be assured that their remote resources will enjoy the same level of care and management as resources in corporate data centers.

Patch Management

One of the least glamorous but still important tasks faced by systems and network administrators is that of keeping their hardware and software up to date. The benefits of applying patches for all devices within an environment can range from reducing security vulnerabilities to ensuring reliability and uptime. More importantly, the cost of *not* diligently testing and applying updates can be extremely high.

The Importance of Patch Management

Although many of the reasons to keep systems updated might be obvious, let's take a quick look at the importance of having a patch management process. First and foremost, security is an important concern for all the components of an IT infrastructure. Ranging from physical hardware to operating systems (OSs) to applications, it's important for known vulnerabilities and issues to be addressed as quickly as possible. Due to the nature of security-related updates, it's difficult to predict which systems will be affected and when updates will be made available. Thus, organizations must be ready to deploy these as soon as possible to prevent exposure to security attacks.

Other reasons for managing patches are just as relevant. By using the latest software, IT departments can avoid problems that might lead to downtime or data corruption. Some patches might increase performance or improve usability. In all cases, there are many advantages to deploying patches using an organized process.

Challenges of Manual Patch Management

Although some environments might handle patches on an ad-hoc "as-needed" basis, this approach clearly leaves a lot to be desired. Even in relatively small IT environments, there are numerous problems related to performing patch management through manual processes. Due to the many demands on IT staff's time, it's often easy to overlook a specific patch or a specific target device when updates are handled manually. The time and costs related to deploying updates can also present a barrier to reacting as quickly as possible.

In larger IT environments, coordinating downtime schedules and allocating resources for keeping hundreds or thousands of devices up to date can be difficult (if not impossible). Often, entire remote sites or branch offices might be out of compliance with standard IT best practices and policies. These seemingly small challenges often result in problems that are very difficult to troubleshoot or that can allow network-wide security breaches. With all these factors in mind, it's easy to see how manual patch management is not ideal.

Developing a Patch Management Process

An important step in improving patch management is to develop a well-defined process. Figure 7.4 provides an example of the high-level steps that should be included in the process.

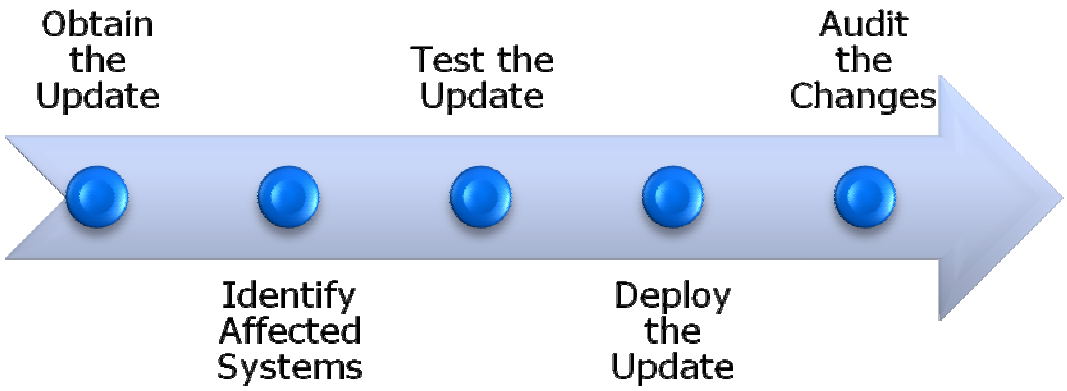

Figure 7.4: Steps in a typical patch management process.

Obtaining Updates

It's important for IT staff to be aware of new updates and patches as soon as possible after they're released. Although many vendors provide newsletters and bulletins related to updates, most IT environments must continuously monitor many sources for this information. This requirement makes it very likely that some updates will be overlooked.

Identifying Affected Systems

Once a potential patch has been made available, systems administrators must determine whether the issue applies to their environment. In some cases, the details of the update might not necessitate a deployment to the entire environment. In other cases, however, dozens or hundreds of systems might be affected. If the patch is relevant, the process should continue.

Testing Updates

A sad-but-true fact about working in IT is that sometimes the "cure" can be worse than the disease. Software and hardware vendors are usually under a tremendous amount of pressure to react to vulnerabilities once they're discovered, and it's possible that these updates will introduce new bugs or may be incompatible with certain system configurations. This reality highlights the need for testing an update. Developers and systems administrators should establish test environments that can be used to help ensure that a patch does not have any unintended effects.

Deploying Updates

Assuming that a patch has passed the testing process, it's time to roll out the update to systems throughout the environment. Ideally, it will be possible to deploy all the changes simultaneously. More likely, however, the need for system reboots or downtime will force IT departments to work within regularly scheduled downtime windows.

Auditing Changes

Once patches have been deployed, it's important to verify that all systems have been updated. Due to technical problems or human error, it's possible that some systems were not correctly patched. When done manually, this portion of the process often requires the tedious step of logging into each server or manually running a network scanning tool.

Automating Patch Management

Clearly, the process of implementing patch management is not an easy one. After multiplying the effort required to perform the outlined steps by the frequency of updates from various vendors, performing the process manually might simply be impossible. Fortunately, data center automation tools can help to dramatically reduce the amount of time and error related to distributing updates. Figure 7.5 provides an example of how an automated patch management solution might work.

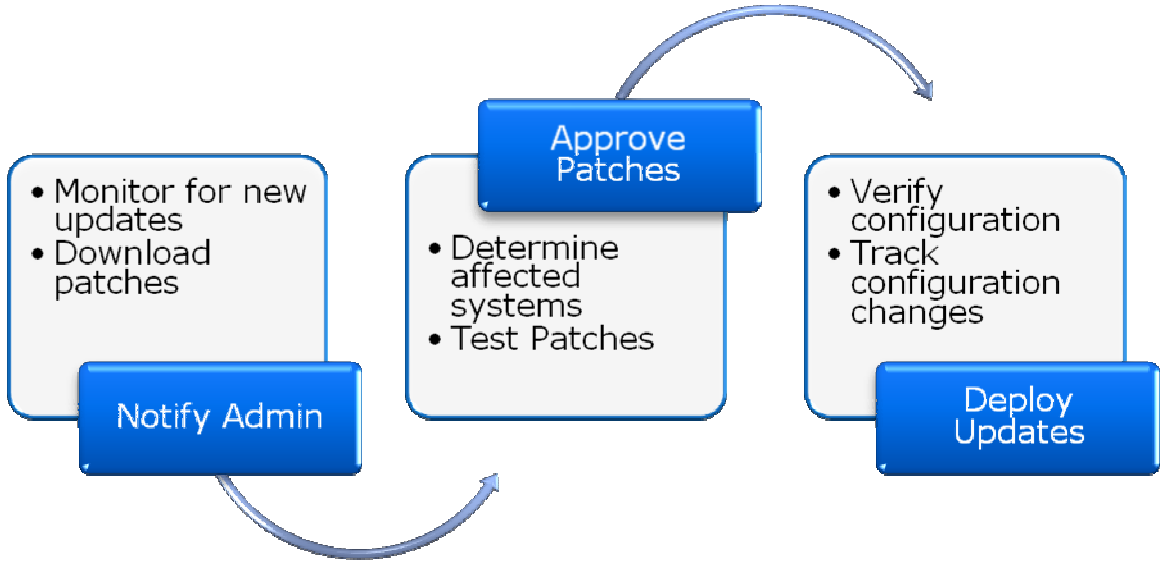

Figure 7.5: An overview of an automated patch management process.

The process begins with the detection of new patches. Ideally, the system will automatically download the appropriate files. If systems administrators determine that the update is relevant and that it should be tested, they can instruct the solution to deploy the update to a test set of servers. They can then perform any required testing. If the update passes the tests, they can instruct the automated patch management system to update the relevant devices. Patches are then applied and verified based on the organization's rules. The entire process is often reduced to a small fraction of the total time of performing these steps manually.

Benefits of Automated Patch Management

The main purpose of an automated patch management solution is to help carry out all the steps mentioned earlier. This includes obtaining updates, testing them, deploying the changes, and auditing systems. In addition to automating these tasks, other benefits include:

- Obtaining updates—The process of discovering and downloading updates can be automated through various tools. This is often done through a database that is managed by the solution vendor. Broad support for many different device types, OSs, and applications is a definite plus. IT staff can quickly view a "dashboard" that highlights which new patches need to be deployed.
- Identifying patch targets—It's often difficult to determine exactly which systems might need to be patched. Automated tools can determine the configuration of IT components and allow administrators to easily determine which systems might be affected.
- Auditing—Expected system configurations can be automatically compared with current configuration details to help prove compliance with IT standards.
- Simplified deployment—Patches can be deployed automatically to hundreds or even thousands of devices. When necessary, the deployment can be coordinated with downtime windows.

With all these benefits in mind, let's look at some additional features that can help IT departments manage updates.

What to Look for in Patch Management Solutions

IT organizations should look for patch management solutions that integrate with other data center automation tools. Through the use of a configuration management database (CMDB), all details related to servers, network devices, workstations, and software can be collected centrally. The CMDB facilitates on-demand reporting, which can help organizations demonstrate compliance with regulatory requirements as well as internal patch policies. Other features include automated notifications, support for remote offices, easy deployment, and support for as many systems and devices as possible.

Overall, the important task of keeping servers and network devices up to date can be greatly simplified through the use of data center automation tools. This approach provides the best of both worlds: ensuring that systems are running in their ideal configuration while freeing up IT time and resources for other tasks.

Network Provisioning

Perhaps the most critical portion of modern IT environments is the underlying network infrastructure. Almost all applications, workstations, and servers depend on connectivity in order to get their jobs done. In the "old days" of computing, networks were able to remain largely static. Although new switches may be added occasionally to support additional devices, the scope of the changes was limited. In current environments, the need to react to rapidly changing business and technical needs has made the process of network provisioning increasingly important.

Defining Provisioning Needs

From a high-level view of the network, it's important to keep in mind several main goals for managing the configuration of the infrastructure. The main objective should be to allow systems and network administrators to efficiently design, test, and deploy network changes. The quicker the IT team can react to changing requirements, the better will be its coordination with the rest of the organization. The list of types of devices that are supported by network teams is a long one, and usually includes many of the items shown in Figure 7.6.

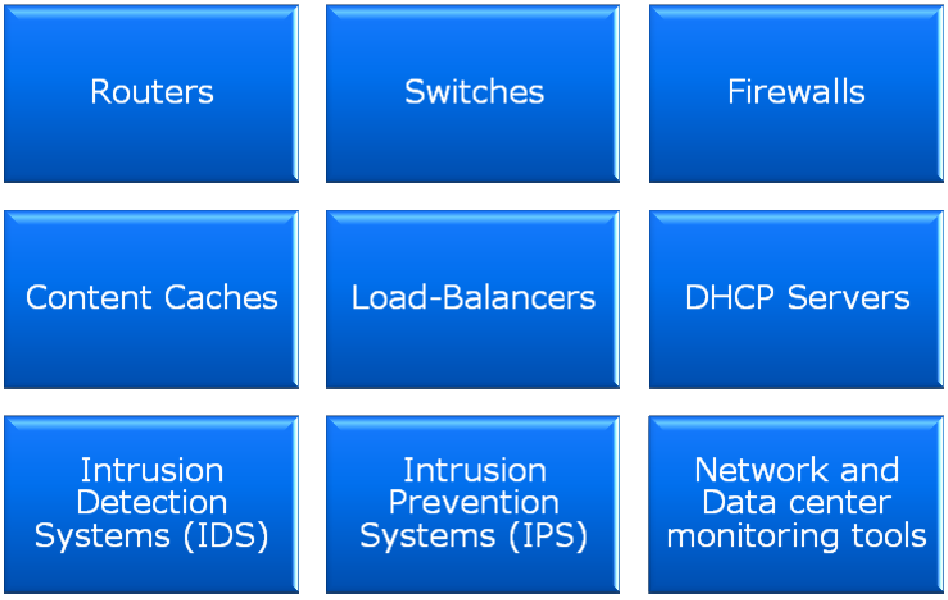

Figure 7.6: Examples of commonly supported network device types.

Common operations include the deployment of new devices and making network-wide changes. Additional tasks include making sure that devices are configured as expected and that they meet the organization's business and technical requirements. Figure 7.7 provides an overview of the types of tasks that are required to perform network provisioning. Let's take a look at some of these requirements in more detail, and how using an automated network provisioning solution can help.

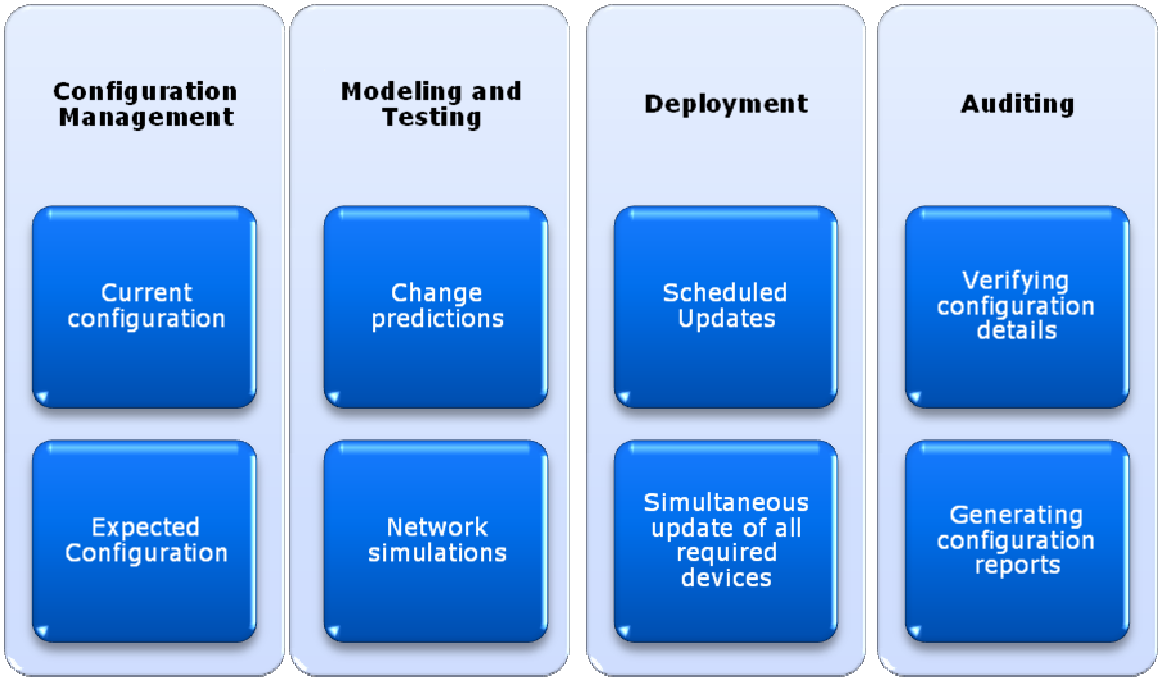

Figure 7.7: An overview of network provisioning goals.

Modeling and Testing Changes

Simple types of network changes might require only minor modifications to one or a few devices. For example, if a new port or protocol should be allowed to cross a single firewall or router, the change can safely be performed manually by a knowledgeable network administrator. The modification is also likely to be fairly safe.

Other types of network changes can require the coordination of changes between dozens or even hundreds of network devices. Often, a relatively simple error such as a typo in a network configuration file or overlooking a single device can lead to downtime for entire segments of the network. Furthermore, applying changes to numerous devices at the same time can be a tedious and error-prone process.

This additional complexity can best be managed through the use of an automated system. By allowing network administrators to design their expected changes in an offline simulation or test environment, they can predict the effects of their changes. This can help catch any configuration problems before they are actually committed in a production environment.

Managing Device Configurations

Once an IT organization has decided which changes need to be made, an automated solution can apply those changes. The process generally involves defining which modifications are to be made to which devices. Data center automation tools can verify that the proper approvals have been obtained and that standard change and configuration management processes have been followed. The actual modifications can be deployed simultaneously to many different devices, or they can be scheduled to occur in sequence. From a network standpoint, the coordination of changes is extremely important in order to avoid configuration conflicts or unnecessary downtime.

Automated network provisioning systems also provide additional useful features. Common operations might include copying the relevant portions of the configuration of an existing device (for de-provisioning or re-provisioning), or defining templates for how network devices should be configured. For environments that often need to scale quickly, the ability to define standard configuration templates for devices such as routers, switches, firewalls, load balancers, and content caches can dramatically reduce deployment times and configuration errors.

Auditing Device Configurations

Even in well-managed IT environments, it's possible for the configuration of a device to deviate from its expected settings. This might happen due to simple human error or as a result of an intrusion or unauthorized modification. Automated network provisioning solutions should be able to regularly scan the configuration of all the devices on the network and report on any unexpected values that are encountered. These reports can be used to demonstrate compliance with regulatory requirements and IT policies and standards.

Using a Configuration Management Database

An excellent method for managing the complexity of network environments is through the use of a centralized configuration management database (CMDB). This central repository can store details related to all the devices in the environment, including networking hardware, servers, workstations, and applications. The information can be combined to provide reports such as insight into overall network utilization or finding the root causes of any performance problems or failures that might have occurred.

Additional Benefits of Automation

By automating network provisioning, IT departments can also realize numerous additional benefits. For example, automatic notifications can be sent whenever problems occur on the network. Also, overall security is often greatly increased because network administrators will no longer need to share passwords, and IT managers can ensure that only authorized personnel are able to make changes. Overall, data center automation tools can greatly simplify the process of network provisioning and can increase the responsiveness of an IT department.

Network Security and Authentication

It is commonly accepted that network security is one of the most important aspects of IT management, but the methods by which users and computers are granted access to communicate within an organization can vary greatly between environments. The goal of most security measures is to ensure that only authorized users can access resources while still allowing all users to do their jobs with a minimal amount of hassle.

Understanding Security Layers

If you were to imagine a house with a concrete front door that includes numerous locks and that has flimsy single-pane windows, it's unlikely that you would consider the house to be secure. The same applies to networks—security must be implemented and managed throughout the organization and at all entry points to the network. The best-implemented security plan will include multiple layers of security. Figure 7.8 provides an overview of some of these layers.

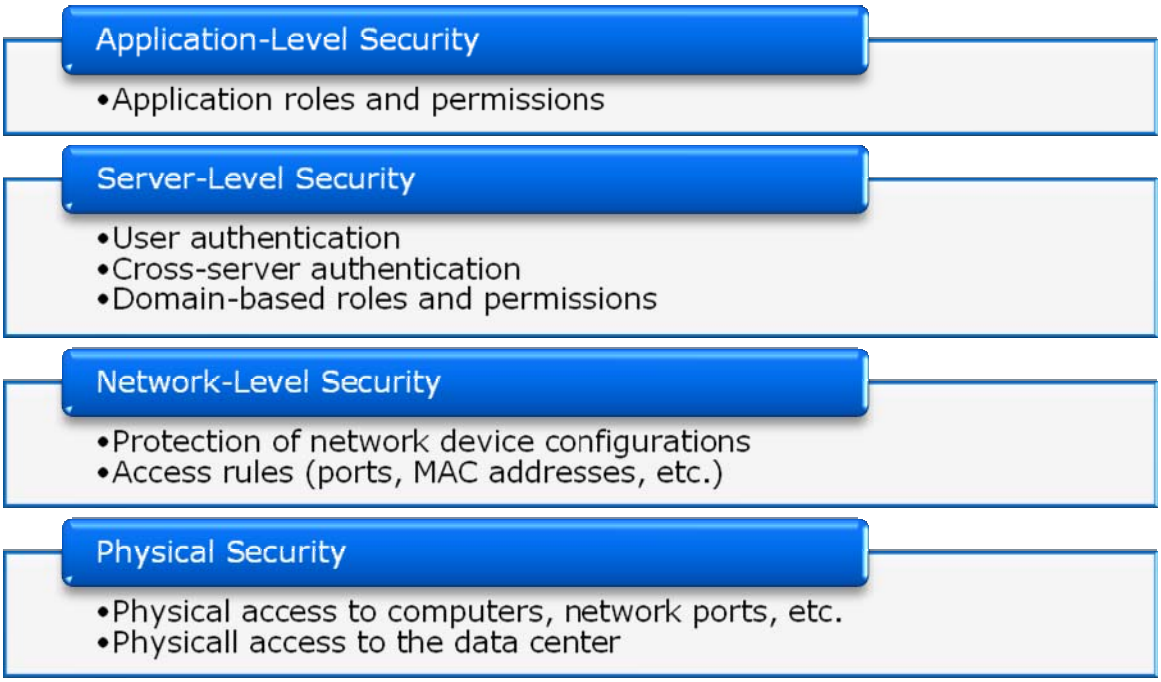

Figure 7.8: An overview of various IT security layers.

All these layers work together to form the links in an organization's armor. For example, before an employee or consultant can access a specific database application, the employee will first have to have access to a physical network port. He or she will then be verified at the network and server levels, and finally at the application level. The user must meet all these challenges in order to be able to access the application.

Choosing a Network Authentication Method

When working in all but the smallest of IT environments, it's important to use a centralized authentication mechanism. One of the most commonly used systems is Microsoft's Active Directory. AD domains provide an organization-wide security database that can be used to control permissions for users, groups, and computers throughout the environment. All administration is managed centrally without requiring security to be configured on individual computers. As long as a user has the appropriate credentials, he or she will be able to access the appropriate devices or services.

Security Protocols

For managing authentication in a distributed network environment, one of the most common protocols is Kerberos. This protocol allows computer systems to be able to positively identify a user in a secure way. It can help avoid security problems such as the interception of security credentials through the use of encryption. Generally, Kerberos is implemented at the server or the application level. However, network devices and other components can also take advantage of it.

There are also several other authentication methods that can be used. Older versions of the Microsoft Windows platform use the NTLM authentication protocol and method. Although this method is less secure than Kerberos, NTLM is a widely supported standard that might still be required to support down-level clients and servers. Also, numerous Lightweight Directory Access Protocol (LDAP)-compliant solutions can integrate with or replace AD. Remote Authentication Dial-In User Service (RADIUS), which was originally developed for the purpose of authenticating remote users, can help centralize security for mobile users and remote locations.

Authentication Mechanisms

The goal of authentication is to ensure that a specific user is who he or she claims to be. By far, the most common authentication mechanism is through the use of a login and password combination. Although this method meets basic security requirements, it has numerous drawbacks. First, users are forced to memorize these pieces of information, and handling lost passwords is a tedious and time-consuming process. Additionally, passwords can be shared or stolen, making it possible that a person is not actually being positively identified. So much is dependent on having the right credentials that this method leaves much room for improvement.

Newer authentication mechanisms include biometrics and the use of specialized security devices. Biometric devices are most commonly based on the use of fingerprints or voice identification to identify individuals. Other methods such as retinal scans are available (though they're most commonly seen in spy movies). Security devices such as an encryption card or "fob" can also be used to verify individuals' identities, especially for remote access. All of these methods involve a certain level of management overhead, and IT departments must be able to keep track of security principals, regardless of the method used.

Authorization

Figuring out how administrators can control access to a system is only part of the security puzzle. Just as important is defining what exactly these users can do. Restrictions can range from determining which files and folders can be accessed to limiting the time of day during which a user can log on. Authorization is the process of granting permission to security principals (such as users or computers) in order to granularly manage what tasks they can perform.

Automating Security Management

With the many methods of managing and determining network permissions, IT departments are faced with a difficult challenge. On one hand, administrators must make systems as usable and accessible to authorized users as is practical. On the other hand, the IT team must ensure that all the different levels and layers of security include consistent information to prevent unauthorized access. Even a single device or database that is out of compliance with policies can create a major security hole in the overall infrastructure.

So how can security be managed across all these disparate systems? A commonly used method is through the use of a centralized security management solution. Figure 7.9 shows an example of how this might work from a conceptual standpoint. The goal of the solution is to coordinate details between multiple security providers. It can do so through the use of a centralized security database that might contain either a master set of credentials or mappings between different types of security systems. The actual implementation details will vary based on the overall needs of the environment. From the user's standpoint, this can help achieve the benefit of single sign on (SSO).

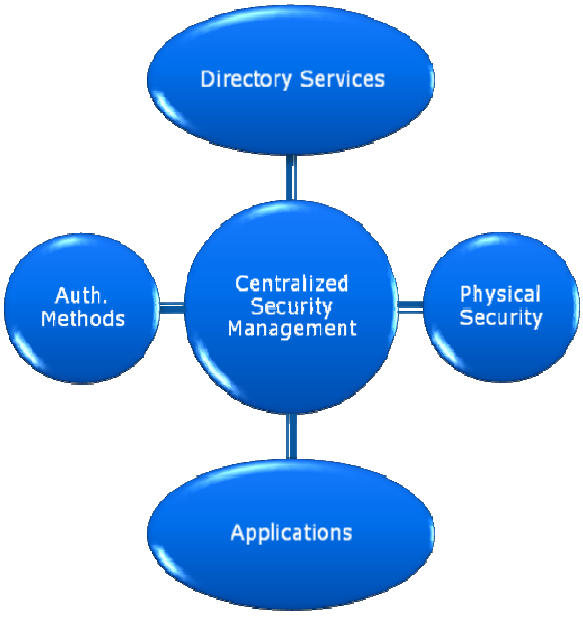

Figure 7.9: Coordinating security between multiple systems.

Overall, by integrating the management of overall security, IT departments and organizations can be sure that all their systems remain coordinated and that only authorized users can access the network.

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.