realtimepublishers.com™

The Reference Guide To™

Data Center Automation

OPSWARE INC
Automating IT™

Don Jones and Anil Desai

Copyright Statement

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

Service Level Agreements

The primary focus of IT departments should be meeting the requirements of other members of their organizations. As businesses have become increasingly reliant on their technology investments, people ranging from desktop users to executive management have specific expectations related to the levels of service they should receive. Although these expectations sometimes coincide with understandings within an IT organization, in many cases, there is a large communications gap.

Service Level Agreements (SLAs) are intended to establish, communicate, and measure the levels of service that will be provided by IT departments. They are mutually agreed-upon definitions of scope, expected turnaround times, quality, reliability, and other metrics that are important to the business as a whole.

Challenges Related to IT Services Delivery

In some areas of IT, the job can be rather thankless. In fact, it is sometimes said that no one even thinks about IT until something goes wrong. Although many organizations see investments in IT as a strategic business investment, others see it only as a cost center. The main challenge is to be able to come to an understanding that includes the capabilities of the IT department and the expectations of the “customers” it serves. That is where the idea of service levels comes in. In order to focus on these benefits, IT departments can think of themselves as outside vendors that are selling products and services to other areas of their organization. Let’s look at some details related to defining these agreements.

Defining Service Level Requirements

SLAs can be set up in a variety of ways, and there are several approaches that can be taken toward developing them. One common factor, however, is that all areas of the organization must be involved. SLAs are not something that can be developed by IT departments working in isolation. The process will require research and negotiations in order to determine an appropriate set of requirements. Figure 4.1 provides an overview of the considerations that should be taken into account. Let’s look at the process of defining SLAs.

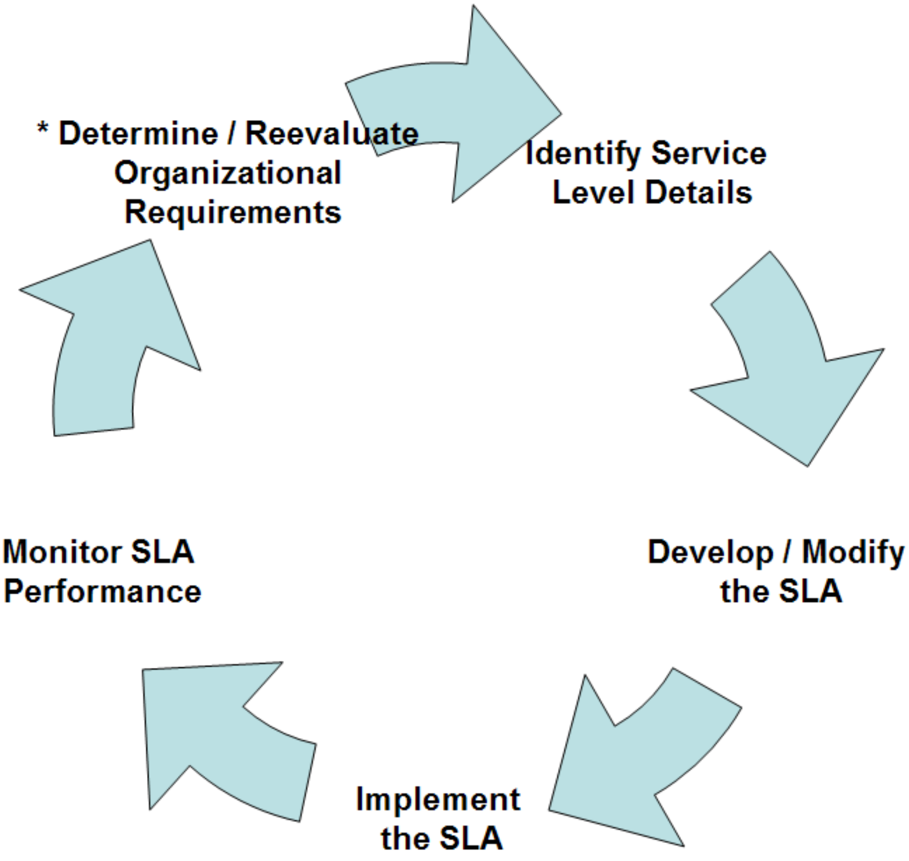
* Determine / Reevaluate Organizational Requirements
Identify Service Level Details
Develop / Modify the SLA
Implement the SLA
Monitor SLA Performance

Figure 4.1: The process of developing, implementing, and evaluating SLAs.

Determining Organizational Needs

IT departments can benefit from thinking of its services as “products” and the users and business processes it supports as “customers.” In this model, the goal of the IT department is to first determine which services the customer needs. This is perhaps the single most important part of the process: IT managers must meet with users and other managers throughout the organization to determine what exactly they need in order to best accomplish their goals. This process can be extremely valuable and enlightening by itself. It’s very important to keep the main goal in mind: To determine what organizations *truly* need, rather than what would just be nice to have.

Identify Service Level Details

The next step is to start trying to define specific details related to what service levels should be accepted. This process will ideally work as a negotiation. A manager from the Engineering department might want all new server deployments to be completed within 2 days of the request. Based on IT staff and resources, however, this might not be possible. The IT manager might present a "counter-offer" of a turnaround time of 4 days. If this isn't acceptable, the two can discuss alternatives that might allow for the goal to be more accessible. In this example, an investment in automated server deployment tools, virtualization, or additional dedicated staff might all be possible ways to meet the requirements.

When discussing goals, it's important for business leaders to avoid diving too far into technical details. For example, rather than requesting a "clustered database solution for the CRM application," it is better for a Marketing manager to state the high-level business requirement, "We need to ensure that, even in a worst-case scenario, our people can access the CRM application." In this particular case, it might well be that the best technical solution doesn't involve clustering at all. The bottom line is that it's the job of IT to figure out *how* to meet the requirements.

A major benefit of this negotiation process is that it forces both sides to communicate details of their operations, and it allows each side to compromise to find a solution that works within given constraints. Occasionally, it might seem impossible for an IT department to meet the needs of a particular business area. In this case, either expectations have to be adjusted or budgetary and staffing resources might be required. In any case, communicating these issues makes the topics open and available for discussion. Once acceptable terms have been reached, it's time to determine what to include in the SLA.

Developing SLAs

There are several important points to include in a complete SLA. Of course, it begins with a description of what level of service will be provided. At this point, the more detailed the information, the better it will be for both sides. Details should include processes that will be used to manage and maintain SLAs. For example, if a certain level is not being met, points of contact should be established on the IT and business sides.

In many cases, IT departments might find that many different service level requirements overlap. For example, several departments might require high availability of Virtual Private Network (VPN) services in order to support traveling users and remote branch offices. This can help IT managers prioritize initiatives to best meet their overall goals. In this example, by adding better monitoring and redundancy features into the VPN, all areas of the organization can benefit.

Delivering Service Levels

IT managers might have some level of fear when committing to specific service levels. Due to the nature of technology, it's quite possible that situations could arise in which SLAs cannot be met (at least not for all areas of the organization). An extreme example might be the "perfect storm" of industry-wide hardware shortages combined with a lack of staff. In such a case, circumstances beyond the control of an organization can cause failures to meet the predefined goals.

Overall, IT departments and business leaders should treat SLAs like they would any other target (such as sales-related goals or Engineering milestones). Ideally, the levels will always be met. But, when they're not, everyone involved should look into the issues that caused the problem and look at how it can be resolved and avoided in the future. Even in the worst case, having some well-defined expectations can help avoid miscommunications between IT and its customers.

The Benefits of Well-Defined SLAs

When implemented properly, SLAs can help make the cost and challenges related to IT operations a part of the entire organization. By providing some level of visibility into IT operations and costs, other departments can get an idea of the amount of work involved. This can help manage expectations. For example, once the Accounting department understands the true cost of ensuring automated failover, perhaps it might decide that some unplanned downtime is acceptable.

IT management can benefit greatly from the use of SLAs. They can use these agreements to justify expenditures and additional staff if appropriate resources are not available to meet the required levels. By communicating these issues up front, either their service levels must be lowered or necessary resources must be made available. Either way, the decision is one that the organization can make as a whole.

Another major benefit of using SLAs is that investments in technologies such as data center automation products can become much more evident. When relatively small investments can quickly return increases in service levels, this is a clear win for both the IT department and the users it supports.

Enforcing SLAs

When dealing with outside parties, an agreement is often only as strong as the terms of any guarantee or related penalties. Because most IT departments tend to be located in-house, it's generally not appropriate to add financial penalties. Thus, the enforceability of SLAs will be up to the professionalism of the management team. When goals are not being met, reasons should be sought out and the team should work together to find a solution. SLAs should be seen as flexible definitions, and business leaders should expect to adjust them regularly. As with other performance metrics, organizations might choose to attach salary and performance bonuses based on SLAs.

Perhaps the biggest challenge is that of prioritization. Given a lack of labor resources, what is more important: uptime for the CRM application or the deployment of new Engineering servers? To help in these areas, IT managers might want to schedule regular meetings, both inside and outside of the IT department, to be sure that everyone in the organization understands the challenges.

Examples of SLAs

The actual details of SLAs for organizations will differ based on specific business needs. However, there are some general categories that should be considered. One category is that of application, hardware, and service uptime. Based on the importance of particular portions of the IT infrastructure, availability and uptime goals can be developed. Other types of SLAs might focus on deployment times or issue resolution times.

Table 4.1 provides some high-level examples of the types of SLAs that might be developed by an organization. The examples focus on numerical metrics, but it's also important to keep in mind that "soft metrics" (such as overall satisfaction with the Service Desk) might also be included.

SLA Area	Metrics	Goal	Notes/Terms
CRM Application Uptime	Percent availability	99.9% availability	Excludes planned downtime for maintenance operations and downtime due to unrelated network issues; major application updates might require additional planned downtime
Service Desk: Level 1 Issue Resolution	Issue Resolution Time	4 business hours	Include definition of "Level 1 Issues"
Service Desk: Level 2 Issue Resolution	Issue Resolution Time	8 business hours	Time is measured from original submission of issue to the Service Desk; include definition of "Level 2 Issues"
Engineering: New Server Deployments (Physical machine)	Time to deployment	3 days	Time is measured from when formal change request has been approved; SLA applies only to servers that will be hosted within the data center
Engineering: New Server Deployments (Virtual machine)	Time to deployment	2 hours	Virtual machines must use one of the three standard configuration profiles; time is measured from when formal change request has been approved.

Table 4.1: Examples of SLAs.

Now that we've looked at some examples, let's see how IT organizations can keep track of SLAs.

Monitoring and Automating SLAs

Once SLAs have been put into place, it's up to the IT department to meet the goals that have been agreed upon. Although some environments might attempt to handle issues only when they arise, the ideal situation is one in which IT managers regularly produce reports showing SLA-related performance. This can be done manually, but in many cases, the management and process overhead related to tracking issue resolution times and uptime can be significant.

One important way in which SLAs can be better monitored and managed is through the use of data center automation tools. Integrated platforms include features for monitoring uptime, automating deployment, and tracking changes. They can also provide IT managers with the ability to define service levels and measure their actual performance against them. Reports can be generated comparing actual performance with expected performance. Without these reports, people might have had to guess whether SLAs were being met. And the inevitable perception issues can negate many of the advantages of having created the SLAs in the first place.

Overall, through the establishment of SLAs, IT departments can verify that they are meeting their customers' requirements and ensure that the organization is receiving the expected value from their IT investments.

Network Business Continuity

Network business continuity focuses on ensuring that network operations will continue to function as quickly as possible after a major outage or disaster. The goal is to limit the disruption to service caused by the failure of a device, a network, or even an entire data center. Most implementations will involve a backup site and a process for failing over to that site, when needed. There are many factors that IT managers should keep in mind when developing network business continuity plans.

The Benefits of Continuity Planning

Business continuity, in general, has become increasingly important for many types of organizations. Customers and business partners have become increasingly reliant on applications and services, and even minor downtime can cause significant financial losses. For example, the loss of connectivity lasting a few minutes for a financial institution can result in lost revenues and reduced customer confidence, both of which would be difficult to regain. The list of things that can go wrong is a long one, ranging from issues with electricity to widespread natural disasters. Business continuity planning attempts to mitigate these risks by planning for processes that will resume normal operations, even in a worst-case scenario.

Developing a Network Business Continuity Plan

The success of any continuity process hinges on its accuracy and alignment with business needs. This section will look at the many considerations that should be taken into account when developing a network business continuity plan. Figure 4.2 shows a high-level view of the processes that should be included.

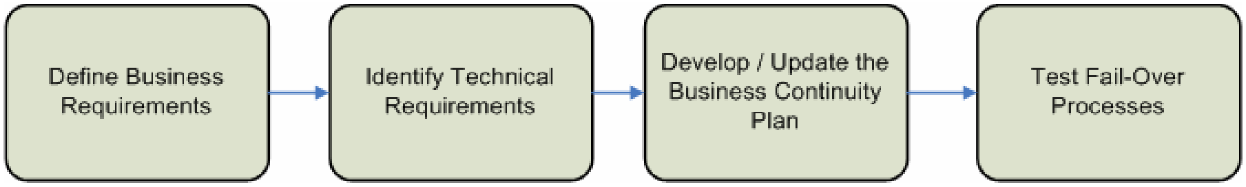
Define Business Requirements
Identify Technical Requirements
Develop / Update the Business Continuity Plan
Test Fail-Over Processes

Figure 4.2: Example steps of a network business continuity plan.

Defining Business Requirements

The first step in developing a network business continuity plan is to determine the organization's requirements. Although all systems are important, certain areas of the network might be more important than others. The most important aspect of determining requirements is to involve an entire organization. The IT department shouldn't rely on its own knowledge to make important decisions related to the most important areas of the computing infrastructure. Given infinite resources, multiple duplicate network environments might be possible. In the real world, it's much more likely that budget and labor constraints will restrict the reasonable level of protection against failures and disasters.

A realistic plan should include discussions of the costs of downtime, the effects of data loss, and the importance of various areas of the network. Ideally, a list of critical systems will be developed based on input from the organization's entire management team.

Identifying Technical Requirements

Modern IT networks tend to be complicated. There are many interdependencies between devices such as switches, routers, firewalls, and network caching devices. And this list doesn't even include details related to which devices are relying on that infrastructure. When planning for business continuity, IT staff should first develop a high-level overview of the network topology and should outline critical systems. The goal is to ensure that the base levels of the infrastructure (which will be required by all other systems) are identified.

The next step is to enumerate which devices will be required in the event of a failover process. Core routers, switches, and firewalls will probably be the first items on the list. Next would be devices required to support the most important applications and services on the network. Considerations should include how the network can run with reduced capacity (particularly if the budget doesn't allow for full redundancy).

Preparing for Network Failover

In the event of a network outage, failover processes must be performed. But before these steps can be taken, IT departments must ensure that they have the tools and information required. This section will take a look at some of the most important considerations.

Configuration Management

Keeping track of network configuration files is an important first step to enabling the failover process. In the event of a failover, restoring this information will help bring a network back to a usable state. Whenever configuration changes are made, network administrators must be sure that the change is recorded and replicated to any backup or standby devices.

Managing Network Redundancy

The implementation of redundancy is a major component of most business continuity plans. When planning for redundancy, it's important to start with defining acceptable downtime limits and appropriate failover times. Most enterprise-level solutions offer options for enabling automatic failover of routers, switches, firewalls, content caches, and other network devices. It is important to keep in mind that, in the case of most failovers, the process might be noticeable to users (although the impact will hopefully be limited to a few connections that need to be reestablished).

Simulating Disaster Recovery Operations

An important—but often overlooked—aspect of any recovery process is to rehearse the failover and business continuity plan. There are many benefits to walking through this process. First, through a trial run, it's likely that business and technical staff will find areas for improvement in the plan. Even the best planning can overlook some of the details that are revealed when performing the "real thing." In the worst case, perhaps a critical system was completely overlooked. Or, there may be various time-saving changes that can be made to improve the process.

Another major benefit of simulating disaster recovery is that practice builds expertise. IT staff should be well-versed in what is required to perform failover processes. There is one iron-clad rule related to testing recovery processes: Immediately after the failure of a critical system is not the time to start learning how to recover it.

Automating Network Business Continuity

There are many aspects of an organizations' network that must be considered when developing and preparing a business continuity plan. For most organizations, the tasks involved will require a lot of work. Fortunately, automated data center management tools can help make the process easier. For example, through the use of automated network discovery, network administrators can easily look at the overall network and discover interdependencies. And, through the use of configuration management (ideally with a configuration management database—CMDB), accurate network device configuration details can be collected. The process of keeping routers, switches, and firewalls up to date at a backup site can also be performed automatically. Figure 4.3 provides an example of how this process might work.

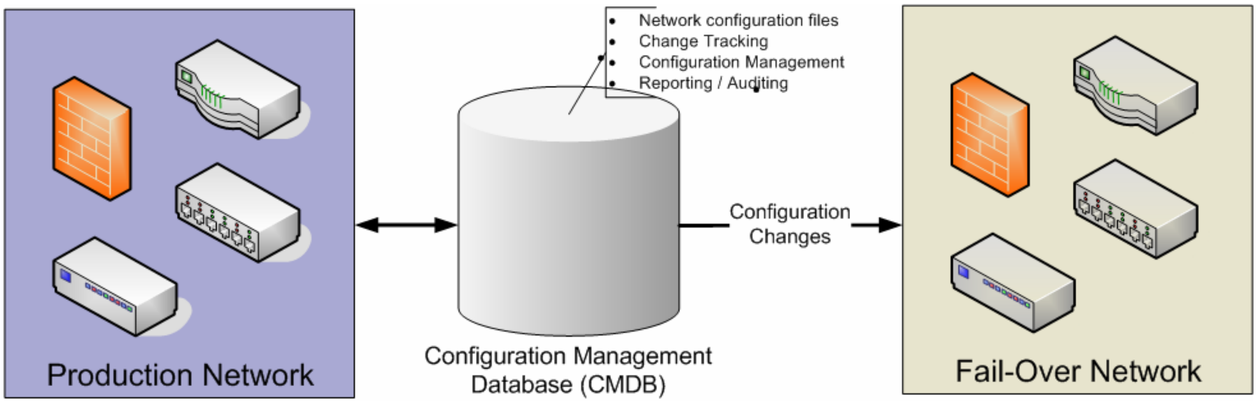

Figure 4.3: Maintaining a failover configuration using data center automation tools.

Developing a network business continuity plan is no small task for most IT departments. Through the use of data center automation solutions, however, this critical task can be made much more manageable.

Remote Administration

In modern IT environments, systems and network administrators are often tasked with managing increasing numbers of devices without additional time and resources. In addition, the systems might be spread out over numerous sites. Centralized management can help meet these needs by increasing overall efficiency. IT staff should be able to manage devices that are located across the world just as easily as they can manage the computing devices on their desks. Remote administration can be used to improve systems and network administration in an IT environment.

The Benefits of Remote Administration

When thinking of desktop administration, the term “SneakerNet” (referring to the fact that systems administrators often spend much of their time and effort walking between systems) might come to mind. For this reason, remote administration is a concept that is usually an easy “sell” to IT departments. When you factor in the labor costs and time associated with physically traveling to remote offices and departments, it’s difficult to find a method that is *less* efficient.

Before looking at specific requirements related to remote administration, let’s quickly cover some of the potential benefits of remote administration. First and foremost, by centrally managing the configuration of hardware, software, and network devices, systems administrators can work from the comfort and convenience of their own workstations. Although having to deal with left-handed mice and custom keyboards might be a fun challenge, it’s clearly not efficient. Time saved by avoiding walking around is also another obvious benefit. For managing data center operations, you can increase security by limiting physical access to servers. From an end-user standpoint, having problems solved quickly and with minimal disruption to work are important goals. By now, the benefits are probably pretty obvious. Let’s delve into what you should look for in a remote management solution.

Remote Administration Scenarios

From a technical standpoint, remote administration can take many forms. Perhaps the most familiar to systems administrators is that of managing servers located in the data center or troubleshooting end users’ desktop machines. Network administrators can also perform remote administration tasks to configure routers, switches, firewalls, and other devices. In distributed environments, the remotely managed device might be located a few feet away or half-way across the world.

Some terms to be familiar with include the remote management host (the computer or device to which you are connecting), and the remote management client (which is usually implemented as software that is run on users’ workstations). Additionally, an organization might have specific tools for monitoring and managing machines remotely.

Remote Management Features

There are several important features to consider when evaluating and selecting a remote management solution:

- Broad support—The ideal remote management solution will be able to support a variety of device types, platforms, and versions. For example, in the area of desktop administration, the remote administration client should be able to connect to all of the operating systems (OSs) and versions that an organization regularly supports. Support for future OSs and products should also be taken into account. All of these platforms should be managed in a consistent manner.
- Reliability—As organizations depend on remote administration features for both routine and emergency operations, reliability is a major concern. The client- and server-sides of the remote management solution should be robust and dependable. Features that allow for remotely restarting a non-responsive host device can be helpful in a pinch. In addition, the ability to perform "out-of-band" management (that is, connections to a system by using non-standard connection methods) can help ensure that services are available when you need them most.
- Efficient bandwidth utilization—Remote management features should efficiently use network bandwidth. In some cases, remote administration connections may be made over high-bandwidth connections, so this won't be an issue. However, when managing remote data centers, small branch offices, and international locations, using an efficient protocol can really help. Potential issues include low throughput rates and high latency on networks (both of which can make a remote connection practically unusable). Specific features to look for include the ability to provide for data compression, low average data rates, and ways to minimize latency given a variety of different network scenarios. In the area of desktop administration, for example, reducing the color depth, hiding desktop backgrounds, and changing screen resolution can help decrease requirements (see Figure 4.4).

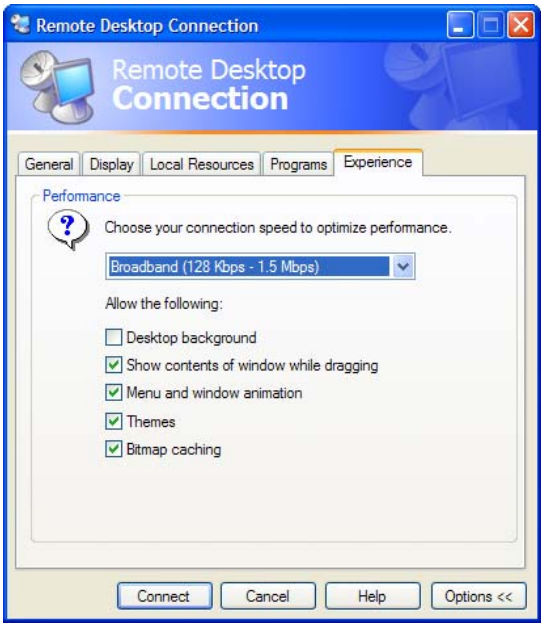

Performance

Choose your connection speed to optimize performance.

Broadband (128 Kbps - 1.5 Mbps)

Allow the following:

□ Desktop background
☑ Show contents of window while dragging
☑ Menu and window animation
☑ Themes
☑ Bitmap caching

Figure 4.4: Configuring video settings in the Windows XP Remote Desktop client.

- File transfers—In addition to controlling remote computers, Help desk staff and systems administrations might need a quick and easy way to transfer files. In some cases, transfers can be handled outside of the remote administration solution by using standard network file transfer methods. In other cases, such as when a connection is made to a remote office or across multiple firewalls, a built-in solution that uses the same protocol and connection as the remote connection can be helpful.
- Shadowing support—For training and troubleshooting purposes, the ability to "shadow" a connection can be helpful. In this method, the remote user might have view-only privileges on the remote device. Or, a trainer might be able to demonstrate an operation on a remote computer without worrying about interruptions from a user.

In addition to these basic features, let's look closer at details related to security.

Securing Remote Management

A critical concern related to remote management features is security. After all, if you're adding a new way in which users can access your users' computers (and the data they contain), what is to keep unauthorized users from doing the same? Fortunately, most modern remote management tools offer many capabilities to help address these concerns.

First, authentication security—controlling who can remotely access a machine—must be implemented. Through the simplest method, authentication security can take the form of a simple "shared secret" username and password combination. But this approach leaves much to be desired—by creating new login information, many potential security problems can be introduced. In addition, it's difficult to manually manage these settings. For example, what happens when systems administrators enter and leave the company? For this purpose, reliance on directory services (such as Windows Active Directory—AD) can help greatly. By centrally managing security settings and permissions, systems administrators can keep track of which users have access to remotely manage which resources.

The next type of security to consider is encryption. Most remote management tools will transfer sensitive information in some form. Even keystrokes and converted video displays can be misused if they're intercepted. The security solution should provide for verification of the identity of local and remote computers (through the use of certificates or machine-level authentication) and should implement encryption of the packets that are being sent between the client and server.

Finally, a remote management solution should provide administrators with the ability to configure, review, and manage permissions related to remote management. In some cases, being able to remotely manage a computer or other device will be an all-or-nothing proposition—either the user will be able to fully control the device or they won't. In other cases, such as in the case of remote desktop management, you might choose to restrict the operations that some users can perform. For example, a Level-1 Service Desk staff member might be allowed to only view a remote desktop machine while the user is accessing it. This can help in the area of troubleshooting, while maintaining adequate security and avoiding potential problems that might be caused by accidental changes.

Choosing a Remote Management Solution

Most systems and network administrators already commonly use remote management features. For example, on Windows desktop and server computers, the Remote Desktop feature is easily accessible. And, for network devices, it's a simple and straightforward process to connect over the network rather than to a physical serial port or dedicated management port on the device itself. Although these features might meet the basic needs of systems management, they do leave a lot to be desired. Managing permissions, keeping track of logins, and controlling connection details can make the process cumbersome and error-prone.

An ideal remote management solution will integrate with other IT data center operations tools, utilities, and processes. For example, in the area of security, existing directory services will be used for authentication and the management of permissions. Security can also be improved by maintaining an audit log of which staff members connected to which devices (and when).

The remote management features may also integrate with change and configuration management tools to keep track of any modifications that have been made. This functionality can greatly help in isolating problems and ensuring compliance with IT standards. Additionally, processes should be put in place to ensure that remote management features are used only when necessary. For example, if automated tools can be used to change network address information on a server, systems administrators should only connect to those machines if they need to perform a more complicated task.

Overall, there are many potential benefits of working with remote management tools in environments of any size. When managed and implemented correctly, remote administration can save significant time and effort while improving IT operations and the end-user experience.

Server Configuration Management

The servers that an IT department manages for other members of the organization are one of the most visible and critical portions of the infrastructure. From hosting file shares, databases, and other critical applications services, servers must be available and properly configured at all times. The challenge for IT staff is ensuring that these computers are properly configured and problems don't crop up over time. This section will talk about details related to server configuration management, including important things to keep in mind when documenting and configuring servers. Based on that, it will then look at details related to simplifying and improving the process through automation.

Server Configuration Management Challenges

When working in production data center environments, there are many challenges that can make managing server configurations more difficult. They can broadly be categorized as technical challenges and process-related challenges.

Technical Challenges

Regardless of the operating system (OS) platform or the applications that are supported, all servers must be kept up to date by systems administrators. Common tasks that must be performed include installing security patches, managing changes to system and network configurations, and taking an inventory of installed services. These operations are fairly simple to perform on one or a few servers, but in most data center environments, IT staff members must manage dozens or hundreds of machines.

Technical challenges include the actual deployment of updates and configuration changes. Performing this task manually is time-consuming and tedious, even when using remote administration features. Also, it's far too easy for systems administrators to accidentally overlook one or a few machines. In the case of implementing security patches, the result could be serious security vulnerabilities.

Other challenges are related to actually performing configuration changes. IT departments should ensure that changes are made consistently, that they adhere to best practices, and that any modifications are tracked and documented. It's also important to ensure that only authorized administrators are making changes and to track who made modifications. Although most systems administrators would agree to this process, in the real world, it can be difficult to spend the time and attention required to follow these steps every time.

Process-Related Challenges

It's important for IT departments to implement and enforce processes related to change and configuration management. The goal is to ensure that all changes are valid and authorized and to avoid problems that might appear due to inappropriate modifications to server configurations. Unfortunately, ensuring communications between IT staff, management, and the users they support can be difficult. The result is that some changes can cause unexpected problems due to a lack of coordination.

IT management should also consider "quality assurance" processes and auditing of server configurations. Ideally, management would be able to quickly and easily view up-to-date details related to the configuration of all servers in the environment, regardless of location. This can help identify machines whose configurations are outdated or not in compliance with IT policies.

Automating Server Configuration Management

Server configuration management is an excellent candidate for automation in most data center environments. Many of the tasks that must be routinely performed can occur within minutes rather than days, weeks, or months. Let's take a look at the many features and benefits of automating server configuration management.

Automated Server Discovery

An important first step in managing an entire IT environment is to discover what is out there. Instead of manually connecting to individual servers and collecting configuration details, automated server discovery features can scan the network and discover all the servers that are present on the network. Often, this will include computers that systems administrators weren't aware of, and machines whose purpose is unknown. The computers may be located in the organization's data centers, or within remote branch offices.

Applying Configuration Changes

Once an IT department has decided to make a change on all of its servers, it must begin the tedious and time-consuming process of performing the changes. By using an automated solution, however, a single change can be propagated throughout an entire network environment in a matter of minutes or hours. The changes can be scheduled to occur during periods of low activity and results can be automatically collected. An automated process enforces consistency and helps ensure that some systems are not accidentally overlooked during the update process.

Configuration Management and Change Tracking

A basic fact of working in an IT environment is that server configurations will change over time. In most cases, changes are based on authorized modifications due to business and technical initiatives. A server configuration management tool can collect network configuration information and OS details and conduct application inventories. All these details are obtained automatically, either through the use of agent software or standard OS methods. This reduces the chance for human error and allows for frequent validation of changes.

Additionally, all of the configuration-related data can be stored centrally in a single configuration management database (CMDB). The data can then be correlated with other information about the environment to ensure that configurations are consistent.

Monitoring and Auditing Server Configurations

The process of auditing server configurations ensures that all servers are compliant with server configuration policies. By using these solutions, IT managers can confidently state that all their assets are being properly managed. When configuration details are properly tracked, systems administrators can easily identify which servers might need to be updated. The process of monitoring ensures that only authorized changes have been made and helps avoid unexpected problems. In addition, it applies to the entire environment—not just one or a few servers that an administrator might work with at a particular point in time.

Enforcing Policies and Processes

The importance of strong and consistent policies and processes cannot be overstated. IT departments should develop and enforce methods for making changes to servers. Although many IT managers might have developed approvals processes on paper, in reality, many ad-hoc changes often occur. An automated server configuration management solution can greatly help enforce processes by restricting changes to only authorized users and validating that the proper approvals have been obtained.

From a technical standpoint, security permissions can be greatly restricted. For example, only the automation solution might have permissions to perform actual changes, and systems administrators must make all their modifications through the tool (see Figure 4.5). This serves the dual purpose of increasing accountability and ensuring that only authorized users are accessing server assets.

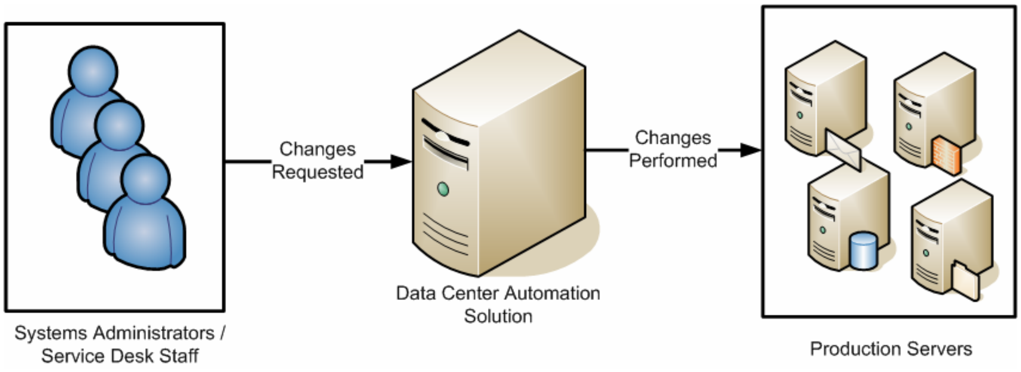
Changes Requested
Changes Performed
Data Center Automation Solution
Systems Administrators / Service Desk Staff
Production Servers

Figure 4.5: Making configuration changes using data center automation tools.

Reporting

One of the most visible benefits of automating the server configuration management process is the ability to generate on-demand reports. The information provided can range from software installation details to security configurations to server uptime and availability reports. All configuration and change data is stored in a central CMDB, so systems administrators and IT managers can quickly obtain the information they need to make better decisions.

Reporting might also be required in order to demonstrate compliance with various regulatory requirements. A process that was formerly time-consuming and inaccurate can be reduced to a few simple steps. Better yet, individuals from areas outside of the IT department can view details that are relevant to performing their jobs.

Evaluating Automated Solutions

In addition to looking for the already mentioned features, there are several factors IT decision makers should keep in mind when evaluating automated server configuration management solutions. They should be sure that most of the platforms they support are manageable using the product. Considerations include hardware platforms, OS versions, and various system updates. Ideally, the technology will be regularly updated to keep pace with new systems. Additionally, the tool should enforce policies and processes to ensure that all changes are authorized and coordinated. Finally, all details should be tracked centrally, and the ability to perform audits and regular reporting can greatly help IT better manage its server investments.

Overall, through the implementation of an automated server configuration management solution, IT departments can perform the vital task of keeping servers updated while avoiding much of the manual work involved. The benefits are that servers are configured consistently and accurately and IT staff is free to perform other important tasks.

IT Processes

Processes define a consistent set of steps that should be followed in order to complete a particular task. From an IT standpoint, processes can range from details about Service Desk escalations to communicating with end users. The goal of IT processes is to improve overall operations within the IT department and the organization as a whole.

It’s often a fact that the implementation of processes requires additional effort and may add steps to some jobs. The steps can be time-consuming and may result in resistance or non-compliance. That raises the challenge: Processes must be worth far more than the “trouble” they cause in order to be considered worthwhile. This section will look at details related to what makes a good process, how you can enforce processes, and the benefits of automating process management.

The Benefits of Processes

Let's first talk about the upside of designing and implementing processes. The major goals and benefits include:

- Consistency—Tasks should be performed in the same way, regardless of who is performing them. In fact, in many cases, it can be argued that having something done consistently in a sub-optimal way is far better than having tasks sometimes completed well and sometimes completed poorly. Ad-hoc changes are difficult to manage and can lead to complex problems.
- Repeatability—It's often easy for IT staff to make the same mistakes over and over or to "reinvent the wheel." The goal of defining processes is to ensure that the same task can be completed multiple times in the same way. Simply allowing everyone to complete goals in their own way might be good for tasks that involve creativity, but they often don't work well for operations that require a lot of coordination and many steps.
- Effectiveness—The process should indicate the best way to do things with respect to the entire organization and all that are involved. The steps involved in the process should enforce best practices.

These benefits might make the decision to implement processes an easy one, but the real challenge is not related to "why" but rather "how" to implement processes.

Challenges Related to Process

For some IT staff members, the very thought of processes might conjure up images of the Pointy-Haired Boss from the Dilbert comic strips. And there are some very good reasons for this: Mainly, many processes are poorly implemented and can actually make work more difficult with little or no benefit to anyone. Some of the problems with poorly implemented processes are based on a lack of knowledge of the details of a particular task.

When out-of-touch management tries to single-handedly implement steps in an operation that it does not understand, the result can be disastrous. Consequently, many IT staffers tend to resist processes. They tend to circumvent them and do the bare minimum in order to meet managements' requirements. Worse, they don't see that there are benefits at all. OK, so that's the bad news. Let's look into what makes a good process (and one that people will like and follow).

Characteristics of Effective Processes

There are several aspects of processes that should be taken into account when implementing new methods of doing things. First, the purpose of the process should be clearly defined before going into the details themselves. Usually, the purpose is to define how a particular set of actions should be performed. Change management processes are a typical example. Organizations might implement formal change request documents and a Change Advisory Board (CAB) to keep track of modifications.

Effective processes should be well aligned with the business and technical goals they're trying to accomplish. "Process for the sake of process" is often counter-productive. Some questions to ask might include "Is this is the best and most efficient way to accomplish a particular goal?" and "Is the extra effort required by the process really worth it?" In some cases, if reporting and documentation of actions aren't useful, perhaps they can be removed to make the process simpler.

The reasoning behind processes should be well-understood. IT staff will be much more likely to adhere to processes that they understand and agree with. Managers should avoid implementing unnecessarily rigid rules: Processes should not attempt to describe every possible action an employee must take. Instead, implementers should be given some leeway in determining the best method by which to complete smaller portions of the tasks. Presenting processes as flexible and evolving guidelines can go a long way toward ensuring compliance.

Designing and Implementing Processes

When you choose to design and implement a new process, it's important to solicit input from all the individuals and business units that might be involved. Ideally, the process will have collective ownership. Although you might be able to coerce employees to follow specific sequences of steps, you might reduce overall productivity by hurting morale and overlooking better ways to do things. The best processes will solicit and incorporate input from all of those involved. Although it might be painful, sometimes one of the best things that IT managers can do is get out of the way.

Another important consideration to keep in mind is that processes should never be considered "final." Instead, they should evolve when business and technical needs change. If you hear systems administrators explaining that processes reflect the way things "used to be done," it's probably time to update the process. In order to ensure that the proper steps are being followed, however, IT staff should be encouraged to propose changes to the process. In fact (and at the risk of sounding like a management fad), a process to control process changes might be in order.

Often, processes can require many steps, and it can be very difficult for all of those that are involved to understand them. One useful method for communication processes is that of flowcharts. Figure 4.6 provides an example of a server deployment process. Note that decisions and responsibilities are clearly identified, and roles for each step have been defined.

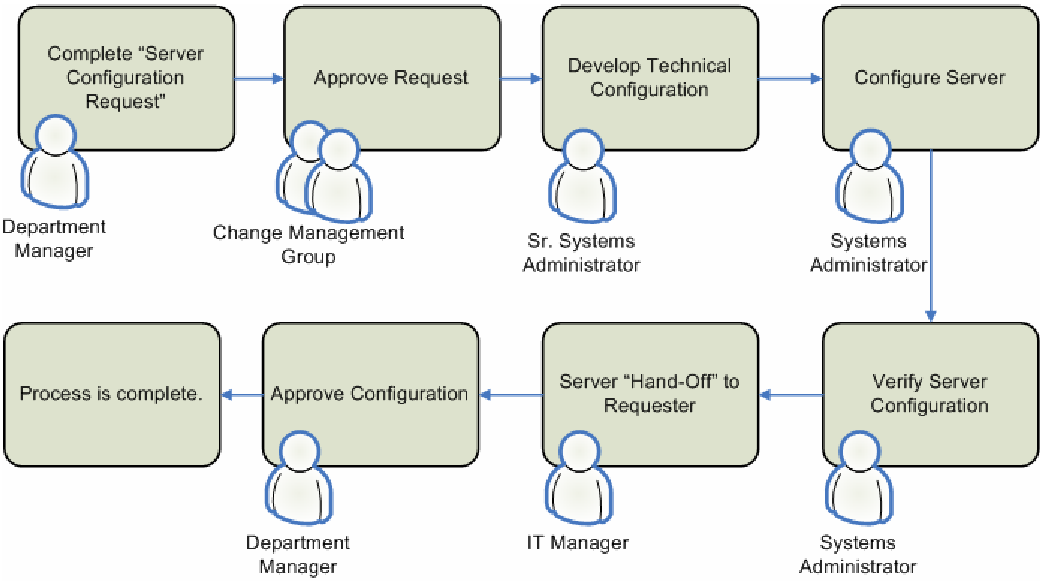

Figure 4.6: An example of a server deployment workflow process.

Overall, the key goals are that those who follow processes should clearly understand the benefits. Without buy-in, the process will be seen as a chore that is forced by management.

Managing Exceptions

An unfortunate fact related to working in the real world is that most rules will have at least occasional exceptions. For example, in an emergency downtime situation, you might not have enough time to walk through all the steps in a change and configuration management process. Clearly, in that case, resolving the problem as quickly as possible is the most important factor. However, the goal should be for exceptions to be relatively rare. If exceptions do occur frequently, it's probably worth considering adding them to the current process or developing a new process.

Delegation and Accountability

One crucial aspect related to developing and managing processes is the people involved. Although it might be easy to define a process and just expect everyone to follow it, there will be many cases in which this simply will not happen. Rapidly approaching deadlines and juggling multiple responsibilities and handling related concerns can often cause diligence related to processes to slip.

One way to ensure that processes are consistently enforced is to ensure that specific individuals are tasked with reviewing steps and ensuring that they're followed. Management can add accountability and metrics to the individuals based on how closely processes are followed and how many exceptions are made.

Examples of IT Processes

By now, it's likely that you're either considering updating existing procedures or putting new processes in place. That raises the question of which operations can benefit most from well-defined processes. In general, it's best to focus on tasks that require multiple steps and multiple operations in order to be completed. The tasks should happen frequently enough so that the process will be used regularly. Other characteristics include business goals that are often not met due to miscommunications or inconsistent ways of handling the tasks that are involved.

Some specific examples of IT processes that organizations might either have in place or might be considering are shown in Table 4.2.

Business Process	Possible Steps	Notes
Change and Configuration Management	Formal documentation of change requests and approval by a CAB	Standard forms for communicating changes can be helpful
IT Purchasing	Requests for multiple quotes (if possible), cost justification, ROI/TCO analysis, and approvals from senior management	Different processes or approval levels might apply based on the cost and business area related to the purchase
Server Deployments	Server configuration review, security configuration checklist, and management acceptance of new configuration	The server should be based on one of the predefined supported configurations
Service Desk	Documentation of new requests, prioritization based on relevant Service Level Agreements (SLAs), and escalation of process details	At any given point in time, the issue must be "owned" by a specific individual

Table 4.2: Examples of IT processes.

Automating Process Management

One important way in which IT managers can better implement, enforce, and manage processes is through the use of data center automation tools and utilities. Ideally, these tools will provide the ability to quickly and easily define processes and workflow. The steps might involve branching logic, approvals, and metrics that must be met along the way. By storing this information consistently and in an accessible way, all people involved should be able to quickly and easily view details about the steps required to complete a particular task.

If the solution can lead the individual or user through the steps required to complete a process correctly, compliance will increase significantly. Additionally, automated process management tools should provide the ability to audit and report on whether particular processes were closely followed.

Overall, when implemented and managed properly, IT processes are a significant characteristic of a well-managed IT environment. Processes can help ensure that tasks are performed consistently, efficiently, and in accordance with business requirements.

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.