

Realtime  
publishers

"Leading the Conversation"

*The Reference Guide To<sup>tm</sup>*

Data Center  
Automation

*sponsored by*



*Don Jones and Anil Desai*

Application Infrastructure Management .....	1
Understanding Application Infrastructure .....	1
Challenges of Application Infrastructure Management .....	2
Inventorying Application Requirements .....	2
Identifying Interdependencies .....	2
Automating Application Infrastructure Management .....	3
Using Application Instrumentation .....	3
Managing Applications Instead of Devices .....	4
Business Continuity for Servers .....	4
The Value of Business Continuity .....	4
Identifying Mission-Critical Applications and Servers .....	5
Developing a Business Continuity Plan for Servers .....	5
Defining Business and Technical Requirements .....	6
Implementing and Maintaining a Backup Site .....	7
Automating Business Continuity .....	7
Using a Configuration Management Database .....	7
Change and Configuration Management .....	8
Network and Server Maintenance .....	9
Network and Server Maintenance Tasks .....	9
Configuration Management .....	9
Applying System and Security Updates .....	9
Monitoring Performance .....	10
Implementing Maintenance Processes .....	10
Delegating Responsibility .....	11
Developing Maintenance Schedules .....	11
Verifying Maintenance Operations .....	11
The Benefits of Automation .....	11
Asset Management .....	12
Benefits of Asset Management .....	12
Developing Asset Management Requirements .....	13
Identifying Asset Types .....	14
Developing Asset Tracking Processes .....	16
Automating IT Asset Management .....	16

Automated Discovery .....17

Using a Configuration Management Database .....17

Integration with Other Data Center Automation Tools .....17

Reporting.....17

Flexible/Agile Management.....18

    Challenges Related to IT Management.....18

    The Agile Management Paradigm .....18

    Key Features of an Agile IT Department.....19

    Automating IT Management.....20

## Copyright Statement

© 2006 Realtimepublishers.com, Inc. All rights reserved. This site contains materials that have been created, developed, or commissioned by, and published with the permission of, Realtimepublishers.com, Inc. (the "Materials") and this site and any such Materials are protected by international copyright and trademark laws.

THE MATERIALS ARE PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, TITLE AND NON-INFRINGEMENT. The Materials are subject to change without notice and do not represent a commitment on the part of Realtimepublishers.com, Inc or its web site sponsors. In no event shall Realtimepublishers.com, Inc. or its web site sponsors be held liable for technical or editorial errors or omissions contained in the Materials, including without limitation, for any direct, indirect, incidental, special, exemplary or consequential damages whatsoever resulting from the use of any information contained in the Materials.

The Materials (including but not limited to the text, images, audio, and/or video) may not be copied, reproduced, republished, uploaded, posted, transmitted, or distributed in any way, in whole or in part, except that one copy may be downloaded for your personal, non-commercial use on a single computer. In connection with such use, you may not modify or obscure any copyright or other proprietary notice.

The Materials may contain trademarks, services marks and logos that are the property of third parties. You are not permitted to use these trademarks, services marks or logos without prior written consent of such third parties.

Realtimepublishers.com and the Realtimepublishers logo are registered in the US Patent & Trademark Office. All other product or service names are the property of their respective owners.

If you have any questions about these terms, or if you would like information about licensing materials from Realtimepublishers.com, please contact us via e-mail at [info@realtimepublishers.com](mailto:info@realtimepublishers.com).

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


## Application Infrastructure Management

In the old days of information technology (IT), applications frequently fit on floppy disks or resided on a single mainframe computer. As long as the hardware platform met the minimum system requirements, data center administrators could be fairly sure that the application would run properly. And, ensuring uptime and reliability involved ensuring that the few computers that ran the software were running properly. Times have definitely changed. Modern applications are significantly more complicated, and can often rely on many different components of an overall IT architecture.

### Understanding Application Infrastructure

When considering hardware, software, network, and operating system (OS) requirements, the entire infrastructure that is required to support an application can include dozens of computers and devices. The actual number of independent parts adds complexity, which in turn can make it much more difficult to manage overall systems.

For example, if a user complains of slow reporting performance related to a Web application, it's not always easy to pinpoint the problem. Perhaps the database server is bogged down fulfilling other requests. Or, perhaps the problem is on the WAN link that connects the user to the Web server. Or, a combination of factors might be leading to the problem. Figure 5.1 shows a simplified path of interactions for creating a report. Each component in the figure is a potential bottleneck.



**Figure 5.1: Potential performance bottlenecks in a modern distributed environment.**

This relatively simple situation highlights the importance of understanding the entire infrastructure that is required to manage an application. For IT departments that support multiple sites throughout the world and dozens of different line-of-business applications, the overall problems can be far more complex.

## Challenges of Application Infrastructure Management

Most IT organizations attempt to manage important applications without fully understanding them. The theory is that as long as areas such as the network infrastructure are properly configured, all the applications on which it depends should also work optimally. Although such certainly can be the case in some situations, some types of issues can be far more complicated to manage.

For example, in the area of change and configuration management, making a single change might have unexpected consequences. Seemingly unrelated modifications to a firewall configuration, for example, might cause connectivity issues in another application. The main challenge for IT is to be able to identify all the inter-related components of an application and to have the ability to compare and verify suggested changes before they're made.


## Inventorying Application Requirements

To get a handle on the complete requirements for complex applications, IT departments should start by taking an inventory of important applications. For example, a Web-based CRM tool that is hosted by an external provider might have relatively simple requirements: As long as users' workstations can access the Internet, they will be able to use the application (although even a Web application might impose other requirements, such as specific browser features). The infrastructure requirements might include network connectivity to the desktop and the firewall and access through edge routers.

Data center applications that require multiple servers can be significantly more complex. Often, multi-tier applications consist of components that include routers, switches, firewalls, and multiple servers. From a logical organization standpoint, the requirements for the application should include all these devices.

## Identifying Interdependencies

Infrastructure components that are shared by multiple applications can be identified after taking an inventory of the application requirements. Often, the results can help provide greater insight into operations. Figure 5.2 provides an example of a shared component that might be used by multiple applications.



**Figure 5.2: Shared application infrastructure requirements for a modern, distributed application.**

As an example, a single low-end switch might be identified as a single point-of-failure for multiple important applications. In this case, an investment in upgrading the hardware or implementing redundancy might help protect overall resources. Also, whenever changes are being planned, test and verification processes should include examining all of the applications that use the same components.

### **Automating Application Infrastructure Management**

It's probably evident that even in relatively simple IT environments, identifying and managing application infrastructure components can be a complicated task. Fortunately, through the use of data center automation solutions, much of the work can be managed automatically.

By storing application infrastructure information centrally in a Configuration Management Database (CMDB), IT staff can quickly find details about all the devices that are required to support an application or service. High-end solutions provide the ability to be able to visualize the interdependencies between hardware, software, and network resources in ways that are easy to understand. Change and configuration management features can also help keep track of the effects of modifications and can help avoid potential problems before they occur.

### **Using Application Instrumentation**

Many third-party applications provide built-in methods for monitoring performance and application activity. Collectively known as “instrumentation,” these features may take the form of a custom Application Programming Interface (API), OS performance monitor counters, or log files. IT departments should look for data center automation solutions that can collect and report on this data.

## Managing Applications Instead of Devices

Although it's easy to get bogged down in the heads-down technical details of maintaining an IT environment, the overall success of operations is not based on routers, servers, and workstations. Instead, the real goal is to manage the applications and services upon which users depend. Well-designed data center automation tools can help IT staff visualize complex inter-dependencies even in widely distributed environments. By focusing on the management of entire applications, IT departments can significantly improve performance, reliability, and availability.

## Business Continuity for Servers

It's no secret that the success of enterprise environments is at least partially based on reliable and available computing resources. In modern business environments, even minor disruptions in service can result in large financial losses. Normally, outages can be caused by power failures, hardware failures, or even the unavailability of an entire data center. As most organizations have become increasingly reliant on IT, technical managers have been tasked with ensuring that services can continue, even in the case of major disasters.

### *The Value of Business Continuity*

Generally, most IT and business leaders have a good idea about the value of business continuity planning. Simply put, the goal is to avoid downtime and to minimize potential data loss. Although it might be tempting to imagine a large meteor heading towards your data center (perhaps, targeting your mission-critical systems), there are many other reasons to protect against disaster. Security breaches or malicious intruders to your system could cause a tremendous amount of damage to systems. In addition, most organizations must rely on infrastructure that is out of their immediate control, such as electric grids and Internet infrastructure. Finally, good old-fashioned user or administrator error can lead to downtime. When all of this is taken into account, the reasons for implementing business continuity are compelling.

Unfortunately, maintaining complete redundancy can be an expensive proposition. Therefore, the organization as a whole should work together to determine the high-level reasons for undertaking a business continuity initiative. In some cases, the main drivers will be related to contractual obligations or complying with regulatory requirements. In other cases, the financial impact of downtime or data loss might create the business case. The important point is for the entire business to realize the value of disaster planning.

Inevitably, organizations will need to determine what needs to be protected and how much is appropriate to spend. The main point is that a successful business continuity approach will include far more than just the IT department—the organization's entire management team must be involved in order for it to be successful.



## Identifying Mission-Critical Applications and Servers


Given infinite resources, implementing business continuity would be simple: multiple redundant environments could be created, and the infrastructure to support real-time synchronization of data would be readily available. In the real world, financial and technical constraints make the process much more difficult. Therefore, before looking at the technical aspects of implementing disaster recovery measures, IT management should meet with business leaders to identify the critical portions of the infrastructure that must be protected.

Assuming that not all resources can be completely protected, it's important to determine the value of each important asset. The first step in prioritization is to take an inventory of the most important high-level functions of the IT department. For example, an online financial services firm might rely heavily upon stock trading software. Next, the technical details of supporting the application should be identified. Modern applications will have many different requirements, including network connections and devices, authentication and security services, and many physical computer systems. In order to provide continuity for the entire end-user service, it's important that none of these components is overlooked. Ideally, IT management will be able to provide an estimate of the cost required to protect each system. In most environments, this process can be challenging, but it's absolutely critical to ensuring a reliable business continuity plan.

### *Developing a Business Continuity Plan for Servers*

When developing a plan for managing servers during disaster situations, it's important to keep in mind the overall goal—to allow business to continue. Often, systems and network administrators will focus on the lower-level technical details of high availability. For example, redundant power supplies and RAID disk configurations can help reduce the likelihood of downtime and data loss. However, the overall approach to high availability should include details related to all areas of operations. For example, even if data and hardware is protected, how will an actual failover occur? Will users be required to implement any changes? What is the process for the IT team? Immediately after a failure occurs is probably the worst time to “rehearse” this process.

Business continuity planning generally involves several major steps (see Figure 5.3). The process begins with identifying which systems must be protected. Then specific business and technical requirements should be defined. Finally, based on this information, the organization will be ready to look at implementing the business continuity plan.



**Figure 5.3: Steps to include in a server continuity plan.**

## Defining Business and Technical Requirements

A general best practice related to performing backups is to base the actual processes that are performed on recovery requirements. When developing business continuity implementations, there are several important factors to take into account:

- **Acceptable data loss**—Although most business managers would rather not think about it, the potential for data loss during a disaster is difficult to avoid. Businesses should come up with a realistic idea of how much data loss is acceptable. An important consideration will be approximate costs. Is it worth a \$1.2M investment to ensure that no more than 2 minutes of transactions will ever be lost? Or is it acceptable to lose up to an hour's worth of transactions to lower the implementation cost? Other considerations include the impact to actual production systems. For example, two-phase commit (2PC) replication for database servers can add single points of failure and can decrease overall production performance.
- **Automated failover**—A disaster or system failure can occur at any time. One requirement to ensure the highest levels of availability is that of automatic failover. Like other factors, however, this comes at a significant cost. For a seamless failover to occur, many aspects of the infrastructure must be prepared. Starting from the server side, machines must be able to coordinate the removal of one server from active service and the promotion of another one to take its place. This process usually requires a third “witness” server. Additionally, the network infrastructure and configuration must be adapted. Finally, changes might be required on the client-side. Although Web applications can often failover without users noticing, full client-side applications might require users to change connection settings or to log out and log back into the system. Clearly, there is a lot of work to be done to ensure automatic failover, but in some business cases, this work is unavoidable.
- **Time for failover**—When primary production servers become unavailable, it will generally take some period of time for the backup site to take its place. There are many challenges related to minimizing this time. For example, how long should systems wait before determining that a failover should take place? And, how is a failure defined? Business should decide on acceptable failover times, taking into account the cost and feasibility of supporting those levels of availability. Furthermore, the entire process should be tested to ensure that there are no unexpected surprises. Even multi-million dollar disaster recovery plans can fail due to seemingly minor configuration discrepancies.

Now that we have a good idea of some of the business and technical considerations, let's look at how you can use this information to create a plan.

### ***Implementing and Maintaining a Backup Site***

The most important aspect of implementing a business continuity plan involves the creation of a secondary site that can be used in the event of a failure. A backup site will generally contain enough hardware and infrastructure services to support critical backup operations from a remote location. Setting up this new site generally involves purchasing new hardware and duplicating the configuration of current production equipment. Although systems administrators are generally aware of the steps required to perform these processes, it can be difficult to replicate configurations exactly.

Once a backup site has been implemented, it's time to look at details related to maintaining it. In some cases, business requirements might allow for periodic backups and restores to be performed. In those cases, some data loss is acceptable. In other situations, however, the backup site must be kept up to date in real-time and must be ready for a loss-less failover in a matter of seconds. For servers, various solutions such as clustering, replication, backup and recovery, and related methods can be used. Regardless of the technical approach, however, a lot of time and effort is usually required to implement and monitor synchronization for a disaster recovery site.

### ***Automating Business Continuity***

Implementing business continuity is generally no small undertaking. IT staff must have a complete understanding of the resources that are to be protected, and all technical information must be kept up to date. It's simply unacceptable for changes to be made in the production environment without an accompanying change within the disaster recovery site. Fortunately, data center automation tools can greatly help reduce the amount of time and effort that is required to maintain a disaster recovery site.

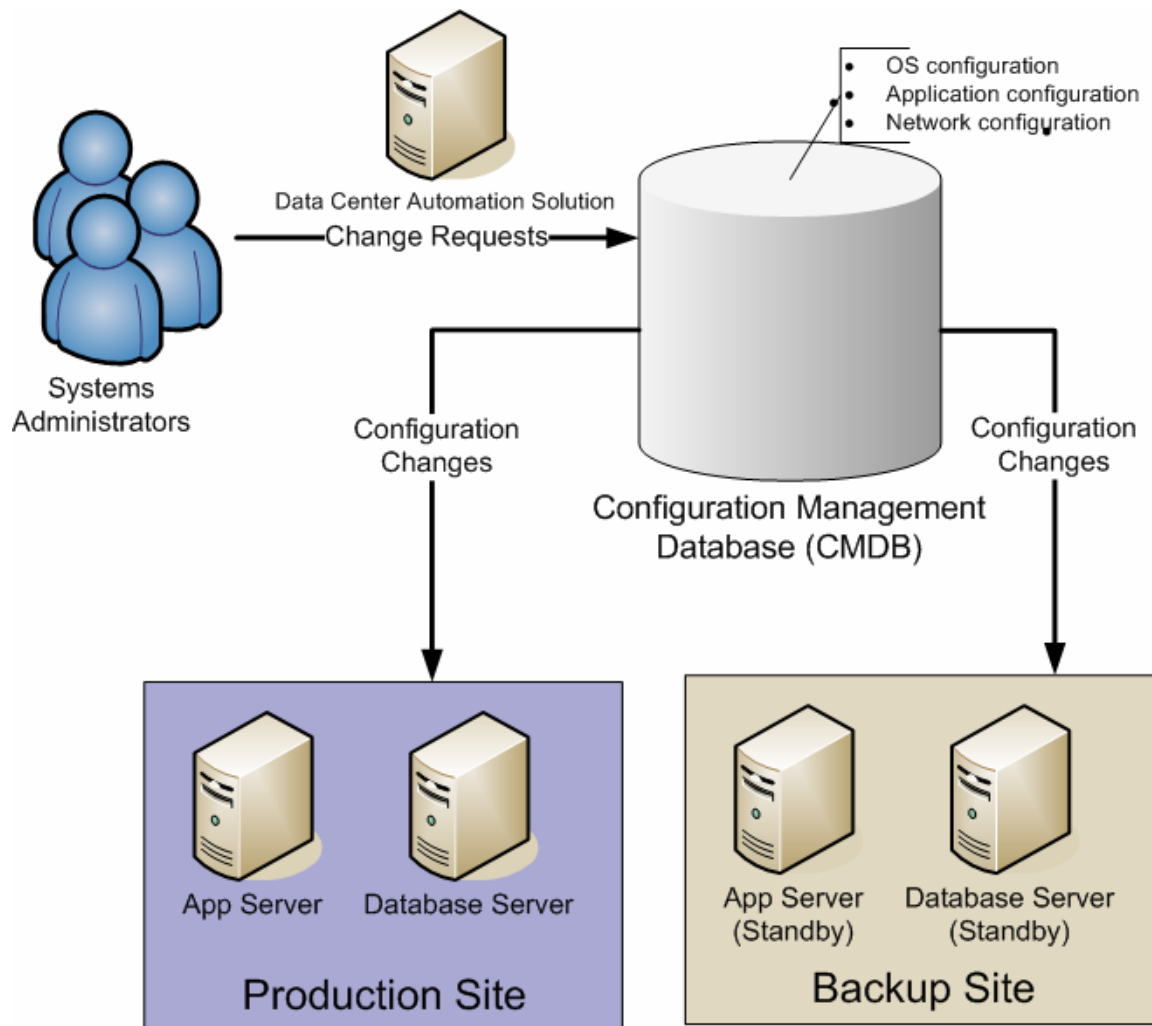
### ***Using a Configuration Management Database***

The purpose of a Configuration Management Database (CMDB) is to centrally store information related to the entire infrastructure that is supported by an IT department. Specifically, related to servers, the CMDB can store configuration details about the operating system (OS), security patches, installed applications, and network configuration.

Using this information, systems administrators can quickly view and compare configuration details for the disaster recovery site. One of the potential issues with maintaining redundant sites is ensuring that a site that is effectively "offline" is ready for a failover. Therefore, reports can be centrally run in order ensure that there are no undetected problems with the backup site.

## Change and Configuration Management

The operations related to keeping a backup site up to date leaves a lot of room for error. If done manually, the process involves a doubling of effort whenever configuration changes are made. Data center automation tools that provide for server change and configuration management can automatically commit the same change to multiple servers (see Figure 5.4). This is ideal for situations in which a backup site must remain synchronized with the production site, and it dramatically reduces the chances of human error.



**Figure 5.4:** Automating configuration management using data center automation tools.

Overall, the process of developing and implementing a business continuity plan for servers will be a major undertaking for IT staff and management. However, through the use of data center automation tools, the process can be significantly simplified, and administration overhead can be minimized. The end result is increased protection of critical data and services at a reasonable overall cost.

## Network and Server Maintenance

Although it might not be the most glamorous aspect of IT, maintaining network and server devices is a critical factor in managing overall IT services. Most systems administrators are well-versed in performing standard maintenance tasks manually, but there are many advantages to automating routine operations.

Most IT organizations have established at least some basic procedures and processes related to the maintenance of devices in the data center. For example, patches might be installed on servers, as needed, and device configurations might be routinely backed up.

### *Network and Server Maintenance Tasks*

Perhaps one quick way of building a list of maintenance tasks is to ask IT administrators what they least enjoy doing. Although a complete list of routine IT tasks could fill many books, this section will focus on common maintenance areas. Each section will explore how data center automation tools can provide significant benefits over manual processes.

## Configuration Management

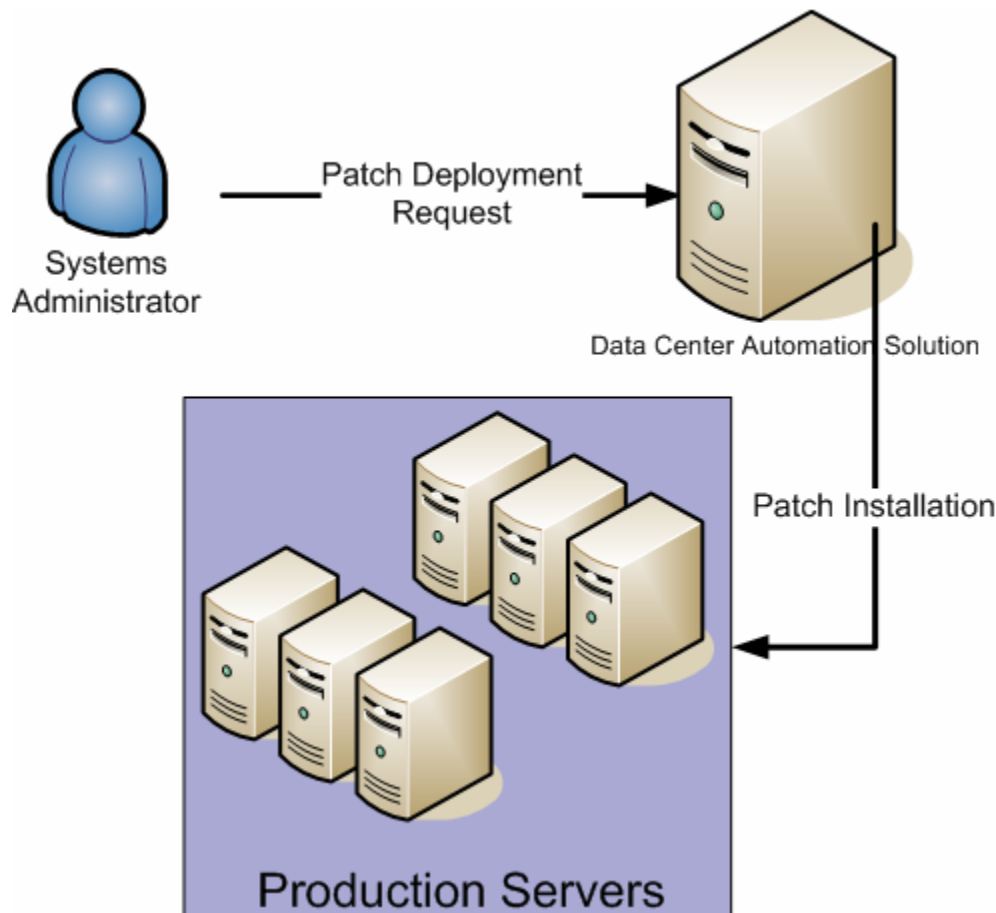
Over time, servers and network equipment will likely need to be updated to meet changing business needs. For example, when a router is reconfigured, network address information may need to change on multiple servers. Alternatively, the organization might implement stricter security policies that must then be applied to hundreds of devices. In very small and simple network environments, it might be reasonable to perform these changes manually. In most IT environments, the process of manually making changes is one that is tedious and leaves a lot of room for error.

Data center automation solutions can ease the process of making changes on even hundreds of devices. The process generally involves a member of the IT staff specifying the change that should be made. Assuming that the staffer has the appropriate permissions, the actual modifications can be scheduled or applied immediately. Often, the task can be completed in a matter of minutes, and all that is left for the administrator to do is verify the change log.

## Applying System and Security Updates

Computers and network devices often benefit from periodic updates. For example, operating system (OS) vendors often release updates that can fix potential functional problems or add functionality. And, security updates are critical to ensuring that important systems and data remain protected. An automated patch management solution can quickly deploy a single update to even thousands of devices with minimal effort.

Figure 5.5 illustrates an automated patch deployment process. In this example, a systems administrator has tested an OS patch and has verified that it is ready to deploy to a set of production servers. Instead of connecting to the servers individually, the change request is sent to a data center automation solution. This server identifies which machines require the update and then automatically manages the patch deployment process. While the updates are being performed, the administrator can view the progress by using the central configuration console.



**Figure 5.5:** Applying updates using an automated system.

## Monitoring Performance

All modern OSs require some standard maintenance operations in order to perform at their best. Actions such as deleting unnecessary files and performing defragmentation can help keep systems performing optimally. For certain types of applications, such as databases, other tasks such as index defragmentation or consistency checks might be required. By implementing automated monitoring solutions, administrators can often be alerted to potential problems before users experience them. And, many types of problems can be fixed automatically, requiring no manual intervention at all.

## Implementing Maintenance Processes

In addition to the various categories of tasks we've covered thus far, there are several considerations that IT departments should keep in mind when performing maintenance operations.

## Delegating Responsibility

An important best practice to keep in mind is that of delegating responsibility. Without coordination between members of the IT team, statements like, “I thought you were going to take care of that last week,” can lead to significant problems. Data center automation solutions can allow IT managers to create and configure schedules for their staff members, and can assign specific duties. This can make it much easier to handle vacation schedules and to ensure that no area of the environment is left completely uncovered at any time.

## Developing Maintenance Schedules

Systems and network administrators are likely all too familiar with spending cold nights and evenings in the server room in order to work during “downtime windows.” Although the goal is a good one—to minimize disruption to production services—it’s still one that is dreaded by IT staff. Through the use of data center automation solutions, downtime windows can be scheduled for any time, and changes can be applied automatically. Administrators can review and verify the changes the next day to ensure that everything worked as planned.

## Verifying Maintenance Operations

The very nature of performing maintenance often relegates important tasks to “back burner” status. When performed manually, it’s far too easy for a network or systems administrator to forget to update one or a few devices, or to be called off on other tasks. In addition to automatically making changes, data center automation solutions can store expected configuration information within a Configuration Management Database (CMDB). IT managers and staff can then compare the actual configuration of devices to their expected configuration to find any discrepancies. This process can quickly and easily ensure that maintenance operations are not overlooked, and that all systems are up to specifications.

## The Benefits of Automation

Overall, without data center automation solutions, the process of maintaining server and network equipment can take a significant amount of time and effort. And, it’s rarely any IT staffer’s favorite job. Through the use of automation, however, tasks that used to take days, weeks, or months can be implemented in a matter of minutes. And, the process can be significantly more reliable, leading to improved uptime, quicker changes, and a better experience for IT departments and end users.



## Asset Management

The goal of asset management is to track the fixed assets that an organization owns and controls. From a general standpoint, asset management can include everything ranging from racks to servers and from buildings to storage devices. IT departments are often tasked with managing budgets and keeping track of inventory, even in highly distributed environments. Without proper technology and processes in place, it can quickly become difficult to find and manage all of these assets. The following sections focus on what to track and how to develop a process that will allow for easily maintaining information about hardware, software, and other important aspects of a distributed IT environment.

### **Benefits of Asset Management**

Although many organizations perform some level of asset tracking using a variety of methods, usually these processes leave much to be desired. For example, IT managers might be tasked with performing a complete audit of software used by the PCs in a particular department, based on the request of the CFO. The process of collecting, analyzing, and verifying the data can be extremely painful. Often, systems administrators must manually log into many different devices to get the information they need. It can take weeks or even months to complete, and still the accuracy of the information might be difficult to verify.

By implementing best practices related to asset management, IT departments and the organizations they support can quickly realize many important benefits:

- Lowering costs—Many IT departments are in a position to negotiate deals and discounts with hardware and software vendors. Often, however, IT departments can leave “money on the table” by not leveraging their bargaining power. In some cases, too much equipment may be purchased, leading to unused systems. In other cases, the IT departments are unaware of exactly how much they’re spending with a vendor, making it difficult to use this information during pricing negotiations. Asset management practices can shed light on the overall resource usage of the entire organization and can lead to better decision making.
- Security—“Out of sight, out of mind,” can apply to many of the assets that are managed by an IT department. During audits or when troubleshooting problems, IT staff might find network devices that have not been patched, or servers for which there is no known purpose. This ambiguity can lead to security problems. Through the use of asset management tools, IT departments can be sure that the purpose and function of each device is known, and they can help ensure that no system is overlooked when performing critical system maintenance.
- Improved service levels—IT departments that are unaware of the location and purpose of devices that they support are generally unable to provide high levels of service and responsiveness whenever problems arise. When asset management can be used to provide the entire IT staff with visibility of the entire environment, monitoring and troubleshooting systems can become significantly easier and more efficient. The end result is quicker and more thorough issue resolution.

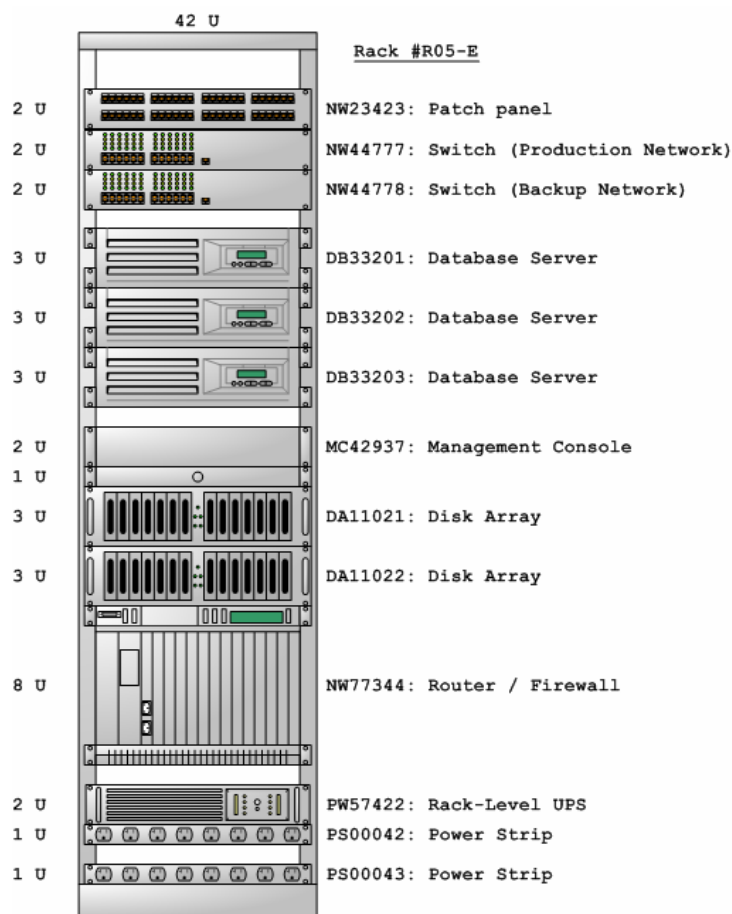


- **Regulatory compliance**—The proper management of fixed computing assets is an important part of many regulatory requirements. It is also an important financial practice. IT managers must be able to identify and locate various capital purchases during an audit, and must be able to provide details related to the purpose and history of the device.
- **Software licensing**—In most IT environments, a significant portion of overall capital expenditures is related to software licensing. Operating systems (OSs), office productivity applications, and specialized tools all incur costs related to purchasing, installation, configuration, and maintenance. It's not uncommon for an IT department to support dozens or even hundreds of different applications. Without an asset management solution, it can be difficult to produce an up-to-date picture of which software is installed (and what is actually being used). However, with this information, IT departments can quickly identify how many licenses are actually needed and whether licenses can be reallocated. The information might indicate that reduced investments in some software and upgrades of other products might be in order.
- **Budgeting**—Providing service while staying within budgetary constraints is one of the most challenging tasks for IT departments. Often, purchasing is handled in a case-by-case, ad-hoc manner. Whenever new assets are needed, IT managers must justify the related expenditures to upper management. And, there are often surprises related to unexpected expenses. By efficiently tracking current assets (and their levels of usage), IT management can provide significantly more accurate predictions about ongoing and future capital asset requirements to the rest of the business.

Once you're sold on the benefits of asset management, it's time to look at how you can implement this technology.

### ***Developing Asset Management Requirements***

Before implementing an asset management solution, organizations should look at what information they need to track. Although the basic needs of asset management are well-defined, additional data can often help improve operations. At a minimum, IT departments should be able to enumerate all the devices that they manage. They should be able to uniquely identify each asset and find the current physical location of the device. In the case of a data center, this might involve the row, rack, and position numbers. Figure 5.6 shows an example of a simple rack diagram.



**Figure 5.6: Developing a rack diagram for asset management.**

In addition to basic information, IT departments should consider capturing details related to the initial and ongoing costs for the device, details about its purpose, and any configuration information that can help in troubleshooting and management.

## Identifying Asset Types

Loosely defined, IT assets can include many different items. The granularity of what is tracked could extend to physical buildings, office spaces, and even network cables. So that raises the consideration of what an IT department should practically track. The main rule is usually based on asset value. It might be determined that only devices that cost more than \$250 should be tracked. Table 5.1 provides a list of the types of assets that should generally be included by an asset tracking solution.

Category	Examples	Information to Collect
Software	Operating systems Office productivity applications Line-of-business applications Standard utilities (firewall software, anti-virus, anti-spyware)	The purpose and cost of each supported application Where software applications are installed Actual application usage Unauthorized software installations
Workstations	End-user desktop computers Training and test lab computers	Computer name and model Hardware and network configuration details Asset cost and related information
Servers	Intranet servers Application servers Database servers	Purpose of the server Computer name and model Hardware and network configuration details Asset cost and related information Support contract details
Mobile devices	Laptop computers PDAs Other “smart” portable devices	Current location of the asset Current “owner” of the device Purpose of the device Information about the capabilities of the device Security information
Networking devices	Routers Switches Firewalls Content caches Intrusion detection/prevention systems	Device manufacturer and model Purpose of each device Physical location

**Table 5.1: An example list of asset types.**

It’s likely that IT departments will need to take into account other types of devices, as well. For example, if a business uses specialized hardware in a testing lab, that hardware should be included. Additionally, IT departments should take into account assets that are committed to remote sites.

## Developing Asset Tracking Processes

As with many other IT initiatives, developing solid asset-tracking processes is critical to ensuring that information remains consistent and relevant. Although software solutions can meet some of these needs, defined and enforced processes are crucial to the overall process. To facilitate the accurate tracking of asset data, organizations should physically place asset tags on devices. Doing so helps uniquely identify a device and requires no technical expertise to match up an asset with its information.

All IT staff must be responsible for keeping asset management up to date through the use of the asset management system. For example, if a router is removed from service, this information should be captured by the asset management tool. A best practice is to include asset-tracking details in any change and configuration management process. Figure 5.7 shows some possible steps to the process.





Figure 5.7: Steps in an asset management process.

 For organizations that have implemented the IT Infrastructure Library (ITIL) best practices, the Software Asset Management topic can be particularly useful. For more information, see the ITIL Web site at <http://www.itil.co.uk/>.

## Automating IT Asset Management

It's likely obvious at this point that the process of asset management can be greatly simplified through the use of automation tools. The tasks of collecting, storing, and maintaining up-to-date data are often well-suited for computer systems. When examining asset management solutions, IT departments should look for features that fit into their overall automation tools frameworks.

For example, a Web-based user interface (UI) can make accessing asset-related data easy for non-IT users. In addition, support for regular updates can help maintain the accuracy of information. Many IT industry hardware and software vendors have included asset tracking features in their solutions. Asset management products that can utilize this type of information should be preferred. The following sections look at other useful features that should be considered when evaluating asset management solutions.

## Automated Discovery

One of the largest barriers related to implementing asset management is the difficulty associated with collecting data about all the devices that must be supported in an IT environment. In some cases, this task might be done manually by physically or remotely connecting to each device and recording details. Of course, apart from the tedium of the process, it's easy for certain devices to be overlooked altogether.

Many asset management solutions can leverage an automated discovery feature to programmatically scan the network and find devices and nodes that are currently active. The process can often be performed very quickly and can include details about devices located throughout a distributed environment. Furthermore, routine audits can be performed to ensure that devices are still available and to track any changes that might have occurred.

## Using a Configuration Management Database

Asset-related data is most useful when it can be combined with other details from throughout an IT environment. For this reason, using a Configuration Management Database (CMDB) is beneficial. The CMDB can centrally store details related to assets and their configuration. The CMDB should also store change-related data in order to ensure that data is always up to date.

## Integration with Other Data Center Automation Tools

Ideally, an asset management solution will integrate with other automation tools used by an IT department. For example, service desk application users should be able to quickly and easily access details about workstation, server, and network devices. This ability can help them more quickly isolate and troubleshoot problems. In addition, systems administrators should be able to update configuration details about a server and have the information automatically update asset-related details such as physical location, network details, and purpose of the computer. Many integrated data center automation solutions provide the ability to make assets easier to track and maintain with minimal effort from systems administrators and IT managers.

## Reporting

The key goal of asset management is to facilitate reporting. IT managers should be able to generate on-demand information about hardware, software, and network devices, as needed. Many asset management solutions will provide the ability to create real-time reports. Products often allow for Web-based report design and customization. By making asset-related information available to managers throughout the organization, IT departments can better ensure that they are meeting overall business needs. Overall, by developing an asset management approach and selecting an appropriate data center automation tool, IT organizations can realize the benefits of tracking the devices they support with minimal time and effort.

## Flexible/Agile Management

In just about any IT environment, changes are frequent and inevitable. Successful businesses must often make significant modifications to business and technical processes to keep pace with customer demands and increasing competition. In business and IT terms, agility refers to the ability to quickly and efficiently adapt to changes. The faster an IT organization can react to changes, the better aligned it will be with business units—and that will translate to overall success for the entire enterprise.

### ***Challenges Related to IT Management***

In some cases, the problems related to agile management might seem paradoxical. On one hand, IT managers work hard to define and enforce processes to ensure that changes are performed consistently and in a reproducible manner. This often requires additional steps to record and track changes, and processes to support them. On the other hand, IT departments must remain flexible enough to support changes that might come at a moment's notice. This raises the question: How can an IT department plan for the future when anything could change at a moment's notice?

It's important not to confuse agility with a lack of processes. As is the case with all areas of the business, chaos resulting from ad-hoc changes is rarely productive and can lead to far more complicated problems in the future. The main point for IT managers to remember is that they must preserve standard operating best practices, even when making large changes in a small period of time.

### ***The Agile Management Paradigm***

The term *agile management* is often heard in reference to managing software development projects. The central theme is to ensure that designers and programmers are ready to accommodate change at a moment's notice. For many environments, the standard year-long cycles of designing, prototyping, implementing, and testing are no longer adequate. Business leaders want to see changes occur with little delay and are unwilling to accept the time and cost related to entire application rewrites. Therefore, the teams must work in much smaller cycles, and each portion of the development process should result in usable code.

Many of the same goals and concepts also translate into the area of managing data center environments. Rather than setting up servers and network infrastructure and considering the job “done,” systems and network administrators must be ready to make major changes when they're required.

## **Key Features of an Agile IT Department**

Although there are many aspects of IT management that can affect the overall quality of operations, there are common areas that should be kept in mind. The key features of an agile IT department include the following:

- **Coordination with business objectives**—Agile IT departments recognize that their main function is to support business initiatives. IT managers and systems administrators must have a high level of awareness of the systems they support, and the reasons that they exist. This awareness can help IT immediately identify which areas might change due to shifts in business strategy rather than waiting until it becomes completely obvious that the systems no longer fit requirements. To keep on top of changes that might be coming, IT representatives should be included in business strategy meetings.
- **Consistent and repeatable processes**—A well-managed IT environment will adhere to best practices and processes such as those presented by the Information Technology Infrastructure Library (ITIL). Although it might seem that processes could get in the way of quick reactions, well-designed processes can usually be adapted to meet new requirements. Specifically, change and configuration management practices can help IT departments quickly react to new needs.
- **Communications**—Too often, IT departments tend to work in a way that is isolated from other areas of the organization. In some cases, IT doesn't find out about the needs of its users until just before changes are required. This situation should be avoided by proactively talking with users and business managers to help prioritize any changes that might be required. In many cases, simple solutions can be developed that minimize disruptive impact while meeting business goals.
- **Efficient administration**—Most IT departments lack the resources to spend days, weeks, or months manually making system configuration changes. Rather, the changes must be made as quickly as possible, but still in a reliable way. Tasks such as the deployment of new software, upgrades to existing equipment, and the deployment of new computing hardware can take significant amounts of time and effort when performed manually. Through the use of dedicated tools for managing the IT infrastructure, even organization-wide changes can be implemented quickly and reliably.

Many other features can also help make IT departments more agile. However, the general rule is that greater agility comes from efficient and coordinated IT departments.

### ***Automating IT Management***

Obviously, all these requirements related to automating IT management can necessitate a significant amount of expertise, time, and effort. As with many other areas of improving IT efficiency, data center automation tools can significantly help IT departments increase their flexibility. Especially when budgets and personnel resources are limited, investments in automation can decrease the overhead related to changes.

Specific areas from which organizations can benefit include change and configuration management, server and network provisioning deployment, automatic updates, asset management, and reporting. For example, there are significant benefits to storing all IT-related information in a centralized Configuration Management Database (CMDB). The combined data can help IT and business leaders quickly identify which systems might need to be updated to accommodate business changes.

Overall, the process of making an IT department more flexible and agile can provide tremendous advantages throughout an entire organization. By quickly adapting to changing needs, the role of IT can transform from a rate-of-change limitation to a strategic advantage. And, through the use of data center automation technology and best practices, IT organizations can quickly work towards the features that can help make them agile.

### **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>.