realtimepublishers.com®

The Definitive Guide™ To

Service-Oriented Systems Management

altiris®

Dan Sullivan

Copyright Statement

Chapter 2: Core Processes in Systems Management

Systems management is a multifaceted practice. The responsibilities of this domain range from ensuring servers are up and running to planning for future growth, which requires meeting the needs of business within the constraints of IT budgets and resources. This chapter examines the core processes entailed in enterprise systems management including:

- Aligning business objectives and IT
- Planning and risk management
- Business continuity and operational integrity
- Security and compliance
- Capacity planning
- Asset management
- Service delivery

These areas do, of course, overlap. For example, one cannot align IT operations with business objectives without planning for growth and potential risks. At the same time, these processes can be treated as distinct because best practices have emerged for each of these processes. In fact, much of this guide is devoted to elucidating the fundamental elements of these processes and describing the best practices that provide for effective and efficient implementation of those processes.

The best place to start a discussion about systems management is with its reason for being: leveraging IT to support business or organizational objectives.

> ✎ Throughout this guide the words “business” and “organization” are both used to describe enterprises that implement systems management practices. Even when the word business is used, the discussion can equally apply to government departments, agencies, and non-profit organizations.

Aligning Business Objective and IT Operations

IT is a means to an end for most organizations. IT is employed to increase productivity, improve communications, increase the reliability and reach of services, advance quality, and a host of other objectives. These objectives are what prompt businesses to deploy the collections of servers, desktops, mobile devices, and specialized network equipment that make up a contemporary IT infrastructure.

Ad Hoc Growth of IT Infrastructure

A common problem arises as organizations grow and shift their business focus: the IT infrastructure does not always change with the change in business objectives. Consider a simple case. A medical device manufacturer begins in business building a limited range of specialized products. The salient characteristics of the company are:

- It has a small sales force and each sales person tracks their leads with a contact management program installed locally on their laptops. Because the sales force is assigned to different markets, there is no overlap between them and no need to sales share information.
- A central office manages order fulfillment, inventory, accounts payable, and accounts receivable using a small and midsized business financials package installed on a local area (LAN) network server.
- An email server is hosted in-house on the LAN.
- The operations manager for the manufacturing process has installed and configured a database system to track production operations and track information needed to remain in compliance with government regulations.
- A Web site with basic company and product information is maintained by a local Web hosting company.

This infrastructure illustrates a typical small business IT scenario and it may work well for many businesses—at least until they start to grow. This type of organization can confront problems with:

- Management reporting—How will sales managers generate consolidated leads reports when their sales staff use standalone databases that are not integrated?
- Information sharing—Is data from the operations database re-keyed into the financials package to complete orders?
- System maintenance and trouble shooting—Who is responsible for fixing problems with the operational database?
- Systems administration—What are security policies regarding email use and antivirus scanning?
- Leveraging IT to expand business instead of simply reacting to immediate needs—How can a Web site be updated to offer online ordering and technical support for customers?

Often, hardware and applications will be procured and deployed to address a narrow problem. If the email server runs out of storage, buy more disk space. If the sales staff need the latest price sheets, export a list for the financial packages to a spreadsheet and email it to all the sales staff. These kinds of ad hoc solutions work in the short term but create an environment that is highly brittle and difficult to maintain. The problem is not that the IT staff is not providing a solution but that they are providing a solution to a series of small problems rather than providing a solution to one over-arching problem: IT operations are not aligned with business objectives.

Managing IT to the Big Picture

The quality of any decision is a product of the quality of the information used to make that decision. If an IT manager is asked to set up a Web site and is provided with the content for that site, the manager might deploy in a way that meets only those requirements. However, if the manager were aware of a competitor who has an online ordering system, the manager can plan for a Web site that includes applications as well as content. This simple example highlights this idea for a small organization—think of the complexity of this idea applied within large enterprises.

In organizations with hundreds and thousands of employees, multiple divisions and departments, a diverse range of product offerings and a geographically dispersed staff and customers, the need to align IT and business objectives is even more pressing. To effectively deploy IT in large organizations, IT managers must have a complete understanding of the current infrastructure, technical limits of existing systems, unmet business requirements, and business plans for future operations. IT cannot be a department that is kept out of the information loop when planning strategic initiatives. IT provides services to other departments and lines of business and must have adequate information to plan for change.

Planning and Risk Management in IT

Planning is a central function of IT. Many of us think of planning as it relates to new acquisitions: new servers, additional applications, extra disk storage, and so on. This is certainly a major part of the planning process, but it is not the whole picture. Planning of this nature works under the assumption that other parts of the IT operation and the business are functioning as they should, but such is not always the case. Business disruptions can result from many causes and planning for those events, known as risk management, is a core IT process.

Basics of IT Planning

Once an IT department understands the objectives of the enterprise and has aligned the strategic plan of IT with those of lines of business, the planning phase can begin. The planning process entails a number of areas, including:

- Technical architecture
- Organizational structure
- Budget and staff management
- Communications

Of these, the technical architecture is the one most often addressed in IT planning.

Planning Technical Architecture

The technical architecture of an IT organization encompasses the hardware, software, data models, and network platforms that comprise the infrastructure. To successfully develop a technical architecture, one should begin with an enterprise data model.

Enterprise Data Models

Enterprise data models are logical descriptions of the data that is used to conduct the business of an organization. This model can include customer information, order data, employee records, sales information, performance measures, and other pieces of information that describe the state of operations. An enterprise data model should include entities and descriptions of processes that correlate with entities and processes in strategic plans. For example, if an objective of the strategic plan is to increase market share in a particular region, the enterprise model should include products, sales volumes, time periods, channel partners, distribution channels, and other entities involved in meeting the objective.

The data model will also have to include descriptions of how information flows throughout the organization. Continuing with the same example, the data model would need to describe where data originates (for example, in an order entry system), when and where it flows to other systems (for example, to an operational reporting system or data warehouse), how it is backed up and archived, along with any other changes to it or uses for it during the data’s life cycle.

> 🕮 Perhaps one of the most complex data reference models is the U.S. Federal Enterprise Architecture Data Reference Model designed for cross-agency information sharing and analysis. For more details, see http://xml.coverpages.org/ni2005-12-28-a.html.

The other part of planning technical architecture focuses on the systems that manipulate enterprise data.

Hardware, Software, and Network Components

Today, most IT departments deploy distributed systems based on open standards. This is a great advantage for IT managers as well as designers and developers, as they are no longer restricted to a single vendor’s hardware platform, operating system (OS), or application offerings. It is common to mix mainframe hardware from IBM with Sun Microsystems and HP UNIX servers as well as Dell servers running a variety of Windows and Linux OSs. Applications may be custom built or purchased from ERP vendors, such as SAP and Oracle, or a plethora of other software vendors that provide specialized applications.

Key enablers of this flexibility in mixing-and-matching components include:

- The widespread adoption of Internet protocols, such as IP, TCP, UDP, HTTP, and others
- The general use of just two types of OSs: the UNIX/Linux family and the Windows family
- The emergence of distributed applications based on Java 2 Platform, Enterprise Edition (J2EE) and .NET frameworks
- The use of standard communication protocols and data exchange formats based on XML

📖 The OASIS organization coordinates a large number of XML standards in a wide range of areas, from e-government and financial services to printing and plumbing. For more information see http://www.oasis-open.org/home/index.php.

Planning enterprise architecture is no longer a matter of committing to a particular vendor's product line, it is more a process of adopting one or more frameworks for organizing a collection of hardware and software components.

Organizational Structure

Planning around organizational structure is about answering questions related to who is responsible for parts of IT infrastructure and services. Common assignments include:

- Help desk support
- Network management
- Server and storage management
- Training
- Security and compliance
- Application and database administration
- Auditing

One goal of organizational structure planning is to ensure that all critical functions are identified and clearly assigned to a business unit. This does not necessarily mean there is a department or group within IT dedicated solely to a single task, but that all tasks are covered. For example, auditing may be assigned to the same group as security and compliance, while Help desk support and training are managed by the same staff.

Another goal of organizational planning is operational efficiency. For example, it is far more cost effective if a single group evaluates anti-malware systems and selects applications best suited to the organization than if every department purchases their own antivirus software. Similarly, allowing disparate lines of business to install different database systems will increase development and support costs as well as introduce application dependencies that can drive up cost long after the initial purchase.

Clearly demarcating lines of authority and responsibility is essential for efficient and effective IT resource management. With an overall organizational structure in place, the next step is to address budgeting and staffing.

Budget and Staff Management

Budgeting and staff management is one of the most difficult areas of system management to discuss in general terms. Obviously, resources should be allocated according to availability and priority of objectives served, but finding the ideal balance is rarely easy. Consider the following issues in the planning process.

First, budgets will be allocated to staff and to tools (or labor and capital in economist's parlance). From a cost management perspective, IT managers do not want to have a 1:1 increase in staff as the size of their infrastructure grows. Enabling staff to manage larger infrastructures requires appropriate tools to:

- Monitor hardware and software systems
- Resolve issues remotely
- Apply patches and upgrades systematically from a central source
- Receive and manage support requests effectively
- Collect data on current uses and capacities to aid in other planning operations

A second consideration is cross training and rotation of duties. This serves two purposes: it enables backup staff to take over in the event a primary support person is unable to meet demands, and it helps reduce the likely success of an internal security breach, such as fraud.

Communications

Communications across lines of business and operational units is sometimes difficult. Each part of the organization has its own priorities and they are not always in sync. What is important to one department is a marginal issue to another. At the same time, vertical communications up and down the organizational structure is an important aspect of keeping IT operations aligned with business objectives. By formally planning and implementing a communication plan, IT systems managers can keep executives informed of the status of operations and projects and keep lines of businesses appraised of service changes, development backlogs, and dependencies on systems that can impact their performance.

Communications across the organization must include more than technical details, project plans, and delivery schedules. Understanding and planning for risks is major factor in IT planning.

Risk Management in IT

Risk management is the process of identifying and assessing potential loss to an organization. This process includes three main steps:

- Prioritizing business objectives
- Assessing risks and impact
- Mitigating risks

Together, these provide the means to identify risks as well as options for dealing with them.

Prioritizing Business Objectives

The first step in risk management is prioritizing business objectives. This has nothing to do with hardware, software, or network infrastructure. The driving question here is, "What are the core operations of the business, and what is the order of priority for protecting those?" An example prioritized list might include:

- Ability of customers to place orders
- Ability of customers to check order status
- Order fulfillment
- Customer support
- Operational reporting
- Management reporting

A prioritized list such as this focuses attention on the most important operations. The next question, what threatens these operations?

Assessing Risks and Impacts

Risks to IT operations come from a number of sources and include damage to physical infrastructure, operator error, hardware and software failure, and security breaches by hackers or malware. Physical infrastructure can be damaged by fire, flood, earthquakes, hurricanes, and other natural disasters. They can also be damaged by failures of other systems; for example, a spike in the electric grid that overloads power conditioners and surge suppressers can do serious harm. The impact of physical risk can range from moderate to total loss.

Operator errors are less likely to cause the loss of physical assets but more likely to result in the loss of information. For example, an operator might accidentally overwrite a backup tape that contains necessary data, or a data entry clerk might accidentally delete records from a transaction processing system. In the first case, the information may be permanently lost or recoverable from other backup tapes. In the case of a data entry error, the lost data might be recovered from database redo logs if caught before changes are committed or from backups in other cases. The recovery methods range from quick and inexpensive to slow and costly procedures.

Hardware and software failures as well as security breaches can range from annoyances to significant disruptions. When assessing the impact of these types of failures, one should address both the direct consequences—for example, an order entry system is down—as well as dependencies, such as the data warehouse cannot be updated and management reports cannot be generated because of the delay in getting operational data. Clearly, the range of impacts is broad; mitigation strategies should be selected based on that range.

Mitigating Risks

Risk mitigation is a balancing act. Formally speaking, risk mitigation strategies should not cost more than the value of the lost resource multiplied by the probability that loss will occur. Unfortunately, quantifiable measures are only available for a small set of risks. For example, hardware manufactures can cite mean time between failure statistics about a device, but there are not good statistics on the mean time between significant bugs in an ERP system, or the likelihood of Denial of Service (DoS) attack, or the chances an operator will accidentally corrupt a backup script that then fails to execute the backups properly. Often, risk mitigation strategies are based on best guesses and past experience.

Risk mitigation strategies, therefore, tend to fall into general approaches that address a number of different risks. Typical examples include:

- Multiple, overlapping backups of critical data
- Failover servers in the event of a hardware failure
- Off-site storage of backups and alternative servers in case of physical damage
- Preventive measures, such as firewalls, intrusion prevention systems (IPSs), and content-filtering applications to prevent breaches and the introduction of malware
- Application user interfaces (UIs) designed to prevent accidental destruction of data
- Database integrity constraints to prevent accidental loss of information—for example, deleting a customer record when the customer has open orders in the database

Understanding the types of risks that confront IT operations is a fundamental part of the planning process. It is also closely related to another core IT process: business continuity planning.

Business Continuity

The goal of business continuity planning and management is to minimize the chance of a business disruption. The risks outlined can lead to an outage of business service. Although the risk management planning process tries to minimize the chance of these risks actually disrupting operations, business continuity addresses what to do when those risks are realized.

Business continuity planning creates policies and procedures that dictate what to do in the event of a business disruption. These plans leverage the resources put in place as part of the risk mitigation strategy. For example, offsite backups can be restored to a backup server at a remote site in the event the primary site is destroyed by fire. To be effective, these plans must be:

- Detailed, application managers and network administrators should not have to think of undocumented but necessary steps to restore operation (for example, updating a DNS record to point to a backup instead of a primary server)
- Tested to ensure the procedures accomplish the proscribed goals
- Rehearsed so that staff are not executing these procedures for the first time during a disruptive event

Business continuity is not an isolated set of tasks that are done at one time, documented, and put on the shelf until the next audit. They are tightly linked to the risk management aspects of IT planning as well as to the security operations of an IT organization.

Maintaining Security and Ensuring Compliance

The fundamental objective of information security is to ensure confidentiality and integrity of information while maintaining the availability of the systems and applications that manage and process that information. Complying with government regulations is largely a matter of meeting the first two objectives—confidentiality and integrity of information.

These objectives are defined as follows:

- Integrity is a property of information that ensures that information is not changed, either intentionally or unintentionally by unauthorized users, or unintentionally by authorized users. In addition, different copies of the same information are consistent.
- Confidentiality is a property of information that ensures that information is not disclosed to unauthorized users.
- Availability is the property of systems that ensures that they are functioning and accessible to users according to Quality of Service (QoS) requirements established for the system.

As the goals of the two are so close, with some variation in specific requirements, it is often helpful to consider security and compliance together. Consider some of the regulations targeted to maintaining individuals' privacy and others designed to ensure integrity in public reporting.

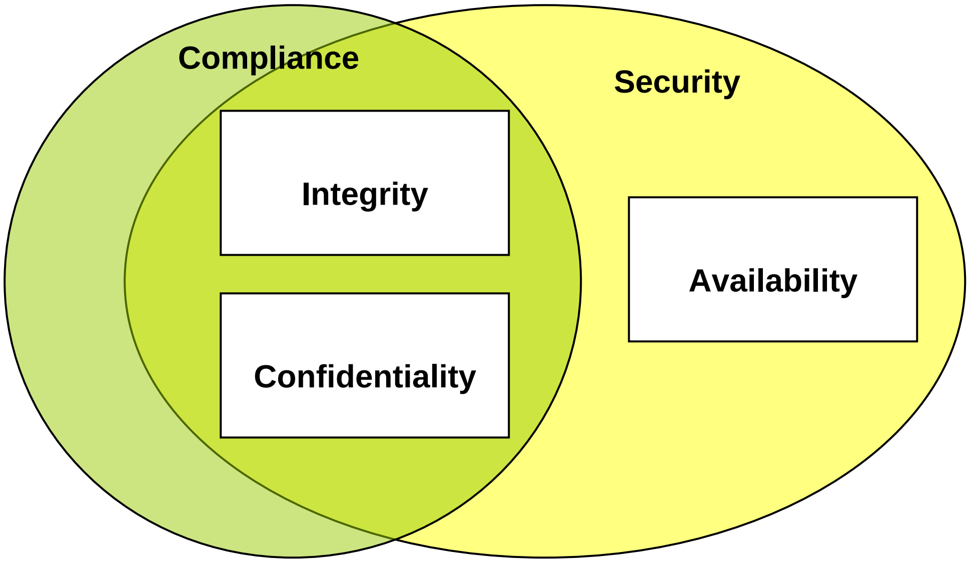

Figure 2.1: Information security and compliance share the common goals of information integrity and confidentiality.

Regulations and Compliance

Regulations related to privacy and information integrity have been established by a wide range of governing bodies. Many of these are formal regulations with the force of law; others are established frameworks within an industry or discipline that are largely accepted and expected to be followed by members of the industry.

Regardless of the source of the regulation or framework, they all require a governance process within IT to ensure full compliance. This is a significant challenge, but fortunately a widely recognized set of best practices, known as Control Objectives for Information and Related Technology (COBIT), provides a set of controls, activities, and measures for governing IT operations. When supported with comprehensive information about the operational state of an IT infrastructure, such as provided by a configuration management database, COBIT is an effective means for maintaining compliance.

📖 COBIT is discussed in detail in Chapter 3.

Privacy and Confidentiality

In the case of government regulations, there are many regulations addressing privacy and confidentiality of information:

- Health Insurance Portability and Accountability Act (HIPAA), which addresses protected healthcare information
- California S.B. 1386, which defines procedures that must be followed if personally identifying information about California residents is compromised
- European Privacy Directives (95/46/EC and 2002/58/EC), which define minimum standards for protecting private information about EU citizens
- Canada’s Personal Information Protection and Electronic Documents Act (PIPEDA) and the Australian Federal Privacy Act, which protects citizens of those countries

It is worth noting that the United States has adopted a decentralized approach to privacy, protecting, for example, healthcare information at the federal level while leaving general privacy regulations to states. Unlike the U.S., many other countries, including the European Union members, Australia, and Canada have adopted comprehensive privacy legislation at the national and transnational levels.

Information Integrity

Maintaining the integrity of business and government information is essential to maintaining the trust of markets, constituents, and others outside those organizations. This reality became abundantly clear with the fiscal reporting scandals that occurred at Enron, WorldCom, Tyco, and other large businesses just a few short years ago.

In response to the growing awareness of the importance of maintaining the integrity of publicly reported information, governments passed a number of regulations to minimize the chance of any further corporate accounting debacles. The most well-known legislation is probably the Sarbanes-Oxley Act (SOX), which defines responsibilities for maintaining and reporting accurate information on publicly traded companies in the United States.

In addition to SOX, some less well-known integrity measures include:

- Computer Fraud and Abuse Act
- Electronic Signatures in Global and National Commerce Act
- Gramm-Leach-Bliley Act

Like privacy protections, the movement to preserve accurate business reporting is a transnational undertaking. For example, the Bank for International Settlements established the Basel II standards to ensure that banks accurately report risks associated with their investments.

Information integrity regulations have also targeted other industries. The U.S. Food and Drug Administration (FDA), for example, has established policies governing the recording, reporting, and storing of information related to the production of pharmaceutical products in the 21 CFR Part 11 regulations.

📖 For more information about compliance from an IT perspective, see the IT Compliance Institute at http://www.itcinstitute.com/.

With so many regulations, it is easy to become overwhelmed. Fortunately, many of these regulations are seeking the same objective: preserve the integrity of information that others depend on and protect personal privacy. The steps that are needed to do so are encompassed by practices in place to ensure a secure information infrastructure. Thus, if one has a well-managed IT environment, it is probably secure and close to, if not already, in compliance with a number of regulations.

Information Security

Of all the areas that comprise systems management, information security is the largest and most difficult. It is the most difficult because there are adversaries who are trying to compromise security measures. It is the largest because there are so many areas that have to be addressed; virtually every aspect of IT is touched by security issues or play a role in security maintenance.

The areas of information security most closely associated with systems management include:

- Threat assessment
- Vulnerability management
- Managing countermeasures
- Auditing
- Incident response
- Change control
- Information security management

These domains within information management require individual planning and management yet depend on each other to be effective.

Threat Assessment

Threats to IT seem ubiquitous since the widespread adoption of the Internet. A threat is a person, program, or process that can compromise the confidentiality, integrity, or availability of information or systems.

> ✎ Threats should not be confused with vulnerabilities, which are weaknesses, deficiencies, or errors in applications, OSs, network devices, or procedures that can be exploited by a threat. Vulnerabilities are addressed in the next section.

Threat assessment is the practice of determining who and what can damage an IT system. Of course, a human is ultimately responsible for all threats, but direct actions carried out by a hacker trying to break into a system require different responses than a malware writer who unleashes a virus to delete randomly selected files from victim's hard drives. For this reason, it is useful to think in terms of categories of threats, such as:

- Information theft, a threat to confidentiality
- Information tampering, a threat to integrity
- DoS attacks, a threat to availability
- Viruses, worms, and other malware, potential threats to confidentiality, integrity, and availability
- Spam, a threat to availability
- Phishing attacks, a threat to confidentiality
- Spyware and other potentially unwanted programs (PUPs), a threat to confidentiality and availability

With an understanding of the broad category of threats, the next step is to understand how these threats are executed. For example, information theft can occur when a hacker compromises a database server and steals credit card information; it can also occur when a disgruntled employee uses legitimate access rights to collect data for unauthorized purposes. In the case of malware, virus can be downloaded along with email through an organization's email server; it can also occur when a laptop user browses a compromised Web site from a poorly secured network at home.

Threat assessment is the practice of discovering potential threats and understanding the motives for those threats. In general, you cannot prevent threats—they exist outside of your control. You can, however, minimize the chances that a threat can successfully compromise your infrastructure. This is the role of vulnerability management.

Vulnerability Management

Vulnerability management is the practice of identifying and compensating for weaknesses in systems, applications, and procedures that can be exploited by threats to breach a system. Like threats, there are a variety of types of vulnerabilities, including:

- Misconfigured network software that allows hackers to use those programs to gain access to protected resources
- Errors in OS software that allows malware writers to gain elevated privileges and execute destructive programs on a compromised host
- Poorly designed programs that do not check for proper parameters and result in a commonly exploited condition known as a buffer overflow
- Organizational policies and procedures that do not account for the potential for attacks or thefts from internal personnel—for example, not rotating duties of employees in critical functions

There are several ways to combat vulnerabilities. First, keep OSs, applications, and network software up to date with security patches. Some systems, such as Microsoft Windows, make it relatively easy for single users or small organizations by offering tools such as Windows Update. As the size of an enterprise increases, more sophisticated tools are required that include centralized management and rollback capabilities. Of course, not all applications have tools for automatically downloading patches from a vendor site. For example, updating a database typically requires a manual download of a patch, which is then applied by a database administrator.

Applying patches to production systems can introduce as well as resolve problems. See the section on change management for more details and caveats.

Second, employ code reviews and software analysis tools to check for common vulnerabilities in custom developed software. This is more within the realm of software engineering than systems management, but systems administrators should be confident that reasonable and prudent measures have been taken to ensure the quality and safety of any application before they deploy it on their networks.

Third, implement organizational policies and procedures that minimize the chance of a breach or theft by an internal staff member. Unfortunately, these crimes are more common and serious than you might expect. For example, a Florida man who was the controlling owner of a Internet advertising company was recently convicted and sentenced to 8 years in federal prison for stealing information about more than 1 billion records containing personal information, such as names, physical addresses, and email addresses from Axciom Corporation, a personal information repository and distributor (details at http://www.cybercrime.gov/levineSent.htm).

📖 For more examples of internal-based breaches, see the U.S. Department of Justice Cybercrime site at http://www.cybercrime.gov/cccases.html.

Finally, understand that vendors are not always the first to detect a vulnerability in their software. Researchers, developers, systems managers, and others may discover and report vulnerabilities to the public through one of the large, public repositories of system vulnerabilities.

Tracking Vulnerabilities

A number of public databases and related tools are available to systems managers in addition to information provided by vendors. These include:

- The National Vulnerability Database (http://nvd.nist.gov/) is a government-sponsored database of all publicly known vulnerabilities. It contains tens of thousands of vulnerabilities as well as a number of cybersecurity alerts cross referenced from the U.S. Computer Emergency Response Team (CERT) from http://www.us-cert.gov/cas/techalerts/.
- The Open Source Vulnerability Database (OSVDB—http://www.osvdb.org/) project also maintains a database of known vulnerabilities. The OSVDB includes support for exporting entries to XML files for importing into other databases.
- The Common Vulnerability and Exposure dictionary (http://cve.mitre.org/) is a standard naming convention for identifying vulnerabilities. It is not a separate database of vulnerabilities but a tool for sharing information across vulnerability databases and making it easier for systems administrators, developers, and other users to query those databases.

Once a vulnerability is found, it should be addressed by either patching the vulnerable code or deploying a workaround. This part of vulnerability management overlaps with some of the tasks associated with change control.

Change Control

Change control in IT is like maintaining a plane while it is in flight. Too often, systems administrators do not have the luxury of shutting down systems and keeping them offline to update software and hardware, test it thoroughly, and bring back users in a controlled manner. Instead, software patches, upgrades, and software installations have to be done with minimal disruption to operational systems.

Information for Change Control

To effectively manage changes to software, hardware, and configurations, systems administrators should have:

- Detailed configurations of deployed systems, including desktop, servers, and network systems
- Information about dependencies between applications, OSs, and hardware components
- An understanding of operational patterns—for example, the frequency and duration of large data loads, batch reports, peak usage of servers, expected growth patterns in storage space, and so on.
- Troubleshooting history for applications, OSs, and hardware. (It is one thing to know how a system is supposed to work; it is another to know how it actually is working.)

The diversity of this kind of information makes it difficult to track in an ad hoc manner. Ideally, systems manager would have a centralized configuration management database that maintains relevant details about the IT infrastructures as a basis for managing change.

Accounting for Dependencies

Although it is easy to track an inventory of devices and applications, it is much more challenging to track dependencies between these components. Some dependencies are clear, such as when an application states that it must run on Windows Server 2003 (WS2K3) or Red Hat Linux AS. Other dependencies are more difficult to discern. Consider a financial software management package that runs on an Oracle database on a UNIX server. The database requires a particular version of a C library. This may or may not be documented, and the database administrator may discover requirements the hard way during installation. Dependencies like this are notorious for throwing off the best laid plans and schedules.

Other dependencies are even more difficult to work around. For example, a server may run two different applications that both use the same database management system. However, a patch to the database may correct a problem encountered by one of the applications but introduce a bug that breaks the other applications. There are no good options in this case. One could continue without the patch and tolerate the vulnerability, bug, or missing feature that the patch corrected, or the systems administrator could run two instances of the database, one patched and one unpatched. Although the latter option solves an immediate problem, it introduces another application instance that must be maintained and managed.

Change management is an especially vexing challenge in systems management and will be addressed in depth throughout this guide.

> 🕮 Configuration management databases are essential to efficient change management. See Chapter 4 for more information about this topic.

Although change management tools help to plan for infrastructure-level changes, auditing helps understand what is happening within those systems now.

Auditing for Security and Systems Management

Auditing is the process of reviewing significant events at various levels, including:

- System
- Application
- User levels

The goal of auditing is to ensure that systems perform as expected, policies are enforced, and unusual and potentially disruptive activities are detected.

System Events

System events occur within OSs. Three of the most important types of events are access control events, configuration change events, and performance measurements. Access control events include:

- Successful and failed logins
- User lockouts due to multiple failed attempts
- Failed file access due to access control violations

When these events occur, the identity of the user as well as the time and device (for example, IP address) should be tracked.

Configuration change events occur when, in the case of Windows, a registry setting is changed or, in UNIX OSs, when configuration files are changed. The identity of the user making the change, the old and new values of the change, and the time and the device from which the change is made are some of the characteristics that may be tracked.

Performance measurements indicate levels of system activity. There are a wide variety of performance measurements that may be collected, including:

- Disk I/O rates
- Page fault rates
- Percent of CPU time in different modes
- Number of files open
- Number of network connections established and connected
- Network segments received per second

These measures are specific to OS and network performance; individual applications may be monitored as well.

Application-Level Auditing

The type and volume of audit information tracked by applications varies widely. Some applications will log details of startup and shutdown processes, error events, database accesses, files opened, and other details of normal operations.

In addition, some applications support a detailed, debugging level of auditing that provide much more detail than is normally recorded in audit logs. Debugging detail is designed to log information about the execution path of a program, indicating which modules are executed, conditions of key variables at the time of execution, and other details that help programmers and support personnel identify problems. This level of detail is not normally needed for application monitoring, only for problem resolution.

User Auditing

In some especially secure environments, it is important to have a record of user activity. This record can include login attempts, use of various resources (including files and applications), and programs executed. It may also record details of commands issued. For example, if someone attempts to copy a file from a secure server to another server using ftp, the file name, the target ftp site, and the date, time, and user identity should be recorded.

Auditing information is useful for systems management as well as for security purposes. It can be especially useful for incident response.

Incident Response

The purpose of incident response is to limit the damage caused by a security breach. Ideally, organizations will have incident response plans in place that dictate how IT staff and management should respond to a security incident. Depending on the type of incident (for example, a virus infection, a database break-in, or a DoS attack), the incident response plan should describe the steps to mitigate the risks of damage. These steps can include:

- Removing a compromised server from the network
- Blocking traffic at a firewall
- Monitoring user activity if an unauthorized action is underway
- Notifying management
- Securing audit logs for forensic analysis

Like so many other security and systems management activities, incident response is most effective when a comprehensive set of information is available about servers, applications, and other devices within the IT infrastructure. An accurate and up-to-date centralized configuration database is as important to enterprise security management as it is to operational systems management.

Capacity Planning and Asset Management

Capacity management deals with the problem of having enough resources to accomplish a given task in the required amount of time. Asset management is closely related but deals more with the details of particulars of acquisition, configuration management, and asset life cycles.

Capacity Planning

Capacity planning is one of the better examples of a systems management domain that leverages the information and practices of other domains. To accurately gauge how much storage space, how many CPUs, or how much bandwidth will be required to support operations at some point in the future requires information about:

- Current loads on servers and the network, which is gathered during performance monitoring
- Growth in application loads, which in part, is determined when aligning IT operations with business strategy
- Dependencies between existing systems and proposed additions to infrastructure, which uses data from change management practices
- Trends in security issues, such as the rate of growth in spam and malware targeted to the enterprise network

Capacity planning requires a combination of looking backward for data and looking forward to anticipated changes. It also requires a firm understanding of existing resource and their levels of use. This is one of the elements of asset management.

Asset Management

Assets are hardware and software components that provide for particular services within the IT infrastructure. Servers, desktops, routers, firewalls, databases, ERP applications, LDAP directory servers, and a range of other devices and applications fall into this category. The scope of asset management, at a minimum, includes:

- Acquiring assets
- Deploying assets
- Configuring assets
- Maintaining assets
- Retiring assets

The specific details of each of these will vary with the type of asset but some general principals hold for all.

Acquiring Assets

The acquisition of assets is closely tied to capacity planning. During capacity planning, when a finding is made that additional resources are required, the acquisition process is initiated. Requirements are defined, designs are formulated, configurations are determined, and the necessary assets are purchased. Also during this phase, dependencies are analyzed to determine how the introduction of the new asset will impact other parts of the infrastructure.

This process is especially important with assets that serve multiple business services. For example, firewalls provide a core network service and could potentially affect every other service and device on the network. A single-user desktop application, however, would have limited impact on others in the organization and could be introduced with less thorough planning.

Deploying and Configuring Assets

Once an asset is acquired, putting it into place may sound relatively straightforward, but that is not necessarily the case—especially when the asset is deployed to a production environment. It may be useful to note at this point that many organizations use development, testing, and production environments for their software development efforts. The development environment is used to create or configure new systems and, after passing basic unit and integration testing, they move to the test environment for user acceptance testing. At that point, representative end users work with the system to determine whether it works as expected and will adequately meet their needs. Only after passing that acceptance testing is the application moved to production.

Moving an asset into production can be a challenge, especially when the time windows for such operations are minimal or when rolling back in the case of problems is difficult. For example, if a new version of a database is to be deployed to production, database administrators and systems managers may have to:

- Replicate the operational database to a secondary, mirror database that will continue to operate during the deployment steps
- Take the operational database offline and create a backup
- Install the new version of the database
- Validate the installation
- Restore the backup to the new version of the database
- Replicate changes made to the secondary database while the primary was down
- Place the primary database back online

Testing and releasing an asset into production should be a highly structured process. The release management processes described in the IT Infrastructure Library (ITIL) defines a method for controlling this process. As with the COBIT framework for governance, the successful implementation of ITIL processes is dependent on adequate operational information such as that found in a configuration management database.

📖 For more information about ITIL, see Chapter 3.

If testing was thorough, the configuration of the new asset should function in the production environment, but there is always the potential for overlooking a configuration parameter or missing a dependency, and configuration changes may be needed after an asset is deployed in production. These steps begin to boarder on maintenance.

Maintaining and Retiring Assets

Although an asset may not change functions during the course of its use, other assets that either depend on that device or the device depends on may change. Maintenance is a routine part of systems management. Sometimes maintenance is driven by outside factors, such as the release of a security patch, or by internal factors, such as the need for additional capacity or a change in network architecture.

Even if a device does not change and its related assets do not force changes, there are still tasks that must be attended to. License management is a prime example. Regardless of what the device is doing, if the software or hardware is licensed from a third party, organizations must ensure that they are in compliance with license renewals and other terms of their agreements.

A short list of aspects of asset management that systems managers must attend to includes:

- Application life cycles
- Dependency management
- User management and security
- Asset acquisition
- Asset deployment and management
- Asset decommissioning
- License management
- Leases
- Warranties

Again, the details necessary to effectively manage this diverse array of tasks are best managed in a centralized configuration management database. The functions of systems management clearly require a broad range of overlapping and multi-functional information. That same information is useful for one other area of systems management: service delivery.

Service Delivery

Service delivery is the process of ensuring that functions and resource needed by the organization are provided in a reliable and cost-effective manner. As is common in systems management, there is some overlap with other core processes. The main components of service delivery are:

- Service level management
- Financial management for IT services
- Capacity management
- Availability management
- IT service continuity management

Service Level Management

Service level management ensures that business units have the amount of resources needed. For example, the e-commerce group in a retail company will need peak network bandwidth and server response time during the holiday shopping season. The financial management group might need significant network bandwidth late in the night to move large data files to the data warehouse, and the customer call center will need the lowest query response time from the database during normal business hours. Balancing the needs of different groups and ensuring each receives the levels and QoS required to do their jobs is the responsibility of service level management.

Financial Management of IT Services

IT is a business within a business. It has both significant labor and capital costs that must be managed. As IT has both operational and project-oriented work, the financial management of the group must support both. Some of the key elements of financial management in IT include:

- Labor, recruiting, and retention
- Support service contracts
- Procurement management
- Consulting and staff augmentation
- Project management
- Capital investment
- Financial risk management

Financial management does not occur in isolation. Many of these tasks, such as project management and risk management are closely tied to other areas of systems management.

Capacity and Availability Management

Capacity is the amount of a resource for performing and operation and meeting a need; availability is having that resource operational when it is needed. A core operation of IT systems management is planning for the future by understanding current usage, trends in growth, and significant changes in business operations that will impact the need for IT resources.

Availability management ensures that resources are functional and providing the level of service required. Clearly, this is closely related to service level management; the distinction lies in the focus. In the case of availability management, the focus is on the minute-to-minute continuity of basic service. For example, if Internet access is functional and the network is operational, networking services are available; if the network latency is too long or there is not enough bandwidth to meet the peak demands, that is an issue addressed as a service delivery problem.

Closely related to availability management is continuity management.

IT Service Continuity Management

IT service continuity management deals with potential business disruptions. There are many causes of business disruptions, from power failures to natural disasters. Businesses plan for such disruptions, and one of the major elements of that planning is how to continue IT operations. The solution entails a number of factors, including:

- Ensuring off-site facilities are available outside the geographic area affected by the disruption
- Brining backup servers and other hardware online
- Restoring backups or ensuring replicated data is up to date at the off-site facility
- Getting key staff to the off-site facility
- Switching services, such as e-commerce servers, to the backup facility

Continuity management must also assess the disruption and determine when it is feasible and prudent to return to the main site as well as plan for the transition back from the backup facility.

Summary

The core operations of systems management are designed to support the strategic objectives of an organization. That is the starting point for the core services of systems management. With a clear and well-defined alignment of business objectives, IT professionals can plan for the capacity needs of the organization, weigh potential risks and mitigate appropriately, and insure the continuity and operational integrity of IT operations. Systems management professional have always had significant responsibility in the area of systems security, and those responsibilities have expanded to support organizational efforts to remain in compliance with a host of government regulations. Other areas of systems management attend to the needs for capacity planning, asset management, and service delivery. As this chapter has demonstrated, the range of systems management is broad and extends beyond the boundaries of the traditional IT department into the business units which they serve.

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.