Microsoft®

MOF Service Management Function Security Administration

patterns & practices

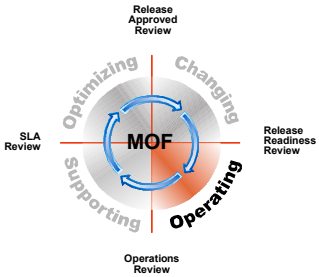

Microsoft®

Solutions for Management

Contents

Document Purpose

This guide provides detailed information about the security administration service management function (SMF) for organizations that have deployed, or are considering deploying, Microsoft technologies in a data center or other type of enterprise computing environment. This is one of the more than 20 SMFs defined and described in Microsoft® Operations Framework (MOF). The guide assumes that the reader is familiar with the intent, background, and fundamental concepts of MOF as well as the Microsoft technologies discussed.

An overview of MOF and its companion, Microsoft Solutions Framework (MSF), is available in the *Introduction to Service Management Functions* guide. This overview guide also provides abstracts of each of the service management functions defined within MOF. Detailed information about the concepts and principles of each of the frameworks is also available in technical papers available at www.microsoft.com/solutions/msm/.

Executive Summary

Security administration is the process of maintaining a safe computing environment. Security is an important part of enterprise infrastructure: An information system with a weak security foundation will eventually experience a security breach.

Security can be divided into six basic requirements, or tenets, that help ensure data confidentiality, integrity, and availability. The six security tenets are:

- *Identification.* This deals with user names and how users identify themselves to the system.
- *Authentication.* This deals with passwords, smart cards, biometrics, and so on. Authentication is how users prove to the system that they are who they claim to be.
- *Access control (also called authorization).* This deals with access and the privileges granted to users so that they may perform certain functions on the system.
- *Confidentiality.* This deals with encryption. Confidentiality mechanisms ensure that only authorized people can see data stored on or traveling across the network.
- *Integrity.* This deals with checksums and digital signatures. Integrity mechanisms ensure that data is not garbled, lost, or changed when traveling across the network.
- *Nonrepudiation.* This is a means of providing proof of data transmission or receipt so that the occurrence of a transaction cannot later be denied.

Process and Activities

Security Administration Overview

Security is an important part of system infrastructure. An information system with a weak security foundation will eventually experience a security breach. Examples of security breaches include data loss, data disclosure, loss of system availability, corruption of data, and so forth. Depending on the information system and the severity of the breach, the results could vary from embarrassment, to loss of revenue, to loss of life.

Security can be broken up into six requirements, or tenets. All of the tenets are equally important for ensuring the confidentiality, integrity, and availability of data. The tenets are listed as follows:

- *Identification.* Identification is concerned with user names and how users identify themselves to a computer system.
- *Authentication.* Authentication is concerned with passwords, smart cards, biometrics, and so forth. Authentication is how users prove to the system that they are who they claim to be.
- *Access control (also called authorization).* Access control is concerned with access and privileges granted to users so that they may perform certain functions on a computer system.
- *Confidentiality.* Confidentiality is concerned with encryption. Confidentiality mechanisms ensure that only authorized people can see data stored on or traveling across the network.
- *Integrity.* Integrity is concerned with checksums and digital signatures. Integrity mechanisms ensure that data is not garbled, lost, or changed when traveling across the network.
- *Nonrepudiation.* Nonrepudiation is a means of providing proof of data transmission or receipt so that the occurrence of a transaction cannot later be denied.

Another very important aspect of security is auditing. Audit logs may give the only indication that a security breach has occurred. Or, if the breach is discovered some other way, proper audit settings generate an audit log that can help administrators pinpoint the location and the perpetrator of the breach.

Goals and Objectives

The primary goals of security administration are to ensure:

- *Data confidentiality.* No one should be able to view an organization's data without authorization.
- *Data integrity.* All authorized users should feel confident that the data presented to them is accurate and not improperly modified.
- *Data availability.* Authorized users should be able to access the data they need, when they need it.

Scope

Security administration is concerned with all aspects of security necessary for maintaining a safe computing environment:

- *Personnel security.* Ensuring employees are properly cleared to handle the data that they access and that adequate checks have been completed before employees are granted access to a system.
- *Application security.* Ensuring that business-critical applications are secure from unauthorized access. This includes a means of identifying and authorizing users of the system.
- *Middleware security.* Middleware includes messages that pass between parts of a service and data that is stored in databases. These must be secured to ensure that data is not viewed, garbled, or modified in any way.
- *Operating system security.* The operating system controls access to hardware and provides access to higher-level services such as databases. If the operating system is not secure, then all the systems and services dependent on the operating system can be compromised.
- *Hardware security.* Security of the computing hardware, storage media, and print output must be ensured. More than ever, hardware such as portable computers (for example, laptops or notebooks), backup tapes, and smart cards contain or provide access to business systems. These assets must be protected both within and without the corporate campus.

- *Network security.* The network carries system data in electronic form. A proper security system protects that data from unauthorized viewing and tampering.
- *Facility security.* Ensuring that physical locks and alarms are in place to keep the computing system safe and that access to the facility is limited to properly identified and authorized personnel. For example, it is useless to secure data electronically if an intruder can simply open an unlocked door and steal the computer.
- *Egress security.* Anything that comes into or out of the facility needs to be secured. This includes but is not limited to mail, electricity, and trash. The loss or compromise of these systems should be assessed to determine the impact on critical business systems.

Key Definitions

Access control. Access and privileges granted to users so that they can perform certain authorized functions on a system.

Authentication. The method by which users prove to the system that they are who they claim to be. Authentication is used in passwords, smart cards, biometrics, and so forth.

Authorization. A process that verifies that the user has the correct rights or permissions to access a resource in a domain.

Confidentiality. A component of encryption. Confidentiality mechanisms ensure that only authorized people can see data stored on or traveling across the network.

Digital certificate. A digital certificate is a data structure that contains the public key of a public/private key pair and identification information and is signed by the private key of the issuing certificate authority (CA). The certificate binds the public key to the security principal (that is, users and computers). The information included includes the name of the owner of the certificate, the uses of the certificate (authentication, data encryption, smart card logon, and so on), and the origin of the certificate (which CA or CA hierarchy created it). The certificate is digitally signed by the CA's private key. To check the authenticity of the certificate, the public key of the CA can be used.

Identification. Any mechanism used to uniquely identify a user or a set of privileges on a system. Identification can be likened to a key. Access control can be likened to a lock. Both the key and lock must match, or "fit," in order to gain access.

Integrity. Data integrity mechanisms ensure that data is not garbled, modified, or lost during transmission across a network. Data integrity mechanisms also ensure that the data is from the intended sender, and not from an impostor. Data integrity mechanisms include checksums and digital signatures.

Nonrepudiation. Nonrepudiation is the security concept that applies to proving the transmission of a particular message. If a system does technical nonrepudiation, then the sender of a message cannot later deny having sent the message, and the receiver of a message cannot later deny having received the message. Furthermore, if the message contains contractual information, the presence of a digital signature with the message can verify that the contractual information was not improperly altered.

Public key infrastructure (PKI). The term generally used to describe the laws, policies, standards, and software that regulate or manipulate certificates and public and private keys. In practice, it is a system of digital certificates, certification authorities, and other registration authorities that verify and authenticate the validity of each party involved in an electronic transaction. Standards for PKI are still evolving, even though they are being widely implemented as a necessary element of electronic commerce.

Virtual private network (VPN). The extension of a private network that encompasses encapsulated, encrypted, and authenticated links across shared or public networks. VPN connections can provide remote access and routed connections to private networks over the Internet.

Major Processes

Security management comprises two main processes and a number of subprocesses as follows:

- Technical security fundamentals
  - Identification
  - Authentication
    - Biometric authentication systems
    - Smart card authentication systems
    - Password authentication systems
    - Web access authentication
  - Access control
    - Authorized usage warning
    - Accountability and shared user IDs
    - Account lockout
    - Settings to limit unauthorized session use and systems access
    - Setting privileges and permission on objects
    - Role-based access control and delegation of authority
- Confidentiality
  - Private key encryption
    - Kerberos authentication
      - The ticket-granting ticket exchange
      - The client service exchange

- Public key encryption and PKI
  - Digital certificates
    - Certificate policies and practice statement
  - Virtual private networks
    - Firewalls, VPNs, and packet filters: protecting the Internet interface of the VPN server
    - Layer 2 tunneling protocol
      - Carrier/ISP model
      - Client model
      - L2TP security
      - Internet protocol security
        - Security protocols
          Authentication header
          Encapsulating security payload
          Security associations
    - File system confidentiality
  - Integrity
  - Nonrepudiation
  - Auditing
    - Planning auditing
    - Implementing auditing
    - Testing auditing

Technical Security Fundamentals

This section describes the six security tenets: identification, authentication, access control, confidentiality, integrity, and nonrepudiation.

Identification

Identification is the mechanism by which the system asks the user, "Who are you?" Users identify themselves to the system by means of a user ID (also referred to as a user, or logon, name). User IDs must be unique (that is, no two users in a system can have the same user ID). To ensure that user IDs are unique, it is important to develop a logon-naming standard that clearly addresses all name characteristics. This is especially important if the system limits user ID length to eight characters, although this consideration is not an issue with Microsoft® Windows® operating systems. Difficulties in naming conventions arise when people have:

- Hyphenated last names.
- Last names that contain *de*, *de la*, *van*, *van der*, and so forth.
- A name identical to another user who is already on the system.

A well-defined naming convention has the following characteristics:

- User IDs are easy for users to remember.
- User IDs are easy for administrators to create.
- Administrators can easily determine the owner of any user ID.

Following are some suggestions for creating standard user IDs:

- Use first initial plus last name. For example, Lori Kane would have user ID lkane.
- For users who have elements in their last names, such as *de*, *de la*, *van*, *van de* and so on, retain the full last name, but remove the spaces.
- If user IDs are limited to eight characters, truncate the last name at eight characters. For example, Andreas Berglund would now have user ID aberglun).
- If two or more users have identical user IDs, replace the last character with a number. For example, if Andreas Berglund joins the company first, his eight-character user ID would be aberglun. If a user with a similar first and last name joins the company, the user ID would be aberglu1. If another user with similar names joined the company, that individual's user ID would be aberglu2, and so forth.

Many organizations use a four-digit identifier at the end of logon names because this method often eliminates the possibility of duplicate IDs. The following number sequences might be used:

- Last four digits of a social security number.
- Last four digits of an employee number.
- Birth date in the form mmddyy.
- Office phone extension. These four digits should be a number that is easy for the user to remember.

The following practices are ***not*** recommended:

- Using a full government ID (social security number) as part of the user name. This could violate privacy laws.
- Creating user IDs that are entirely numerical. User names should be easy to associate with the users to which they belong. It is much easier to recognize a name than a number.
- Allowing users to use nicknames as part of their user ID. This goes back to easily associating user names with users. For example, if a user uses a school nickname as part of a user ID, many people may not recognize to whom the user name is referring.
- Using the birth year as part of the user name. This raises age discrimination issues.

An additional point to consider is that the user ID is half of the information needed to get into the system if traditional user name/password authentication is used. It is therefore not advisable to allow the system to pre-fill the user ID into the logon window. If outsiders sit down at a computer to break into the system, they have to defeat the authentication mechanism. If they get a blank logon window, they have to guess both a user name and the associated valid password. If the user name is already entered in the logon window, half of their work is done for them. Therefore, it is strongly recommended that the logon window always come up blank for each new user, even if a particular computer is normally only used by one person.

Authentication

Authentication is the mechanism by which the system asks the user, "Is that really you?" If a system has a good logon naming standard, but no authentication, then anyone could log on to the system by using someone else's account, since user IDs would be easy to guess. Therefore, to make sure that only the true owner of an account can get into the account, the system must enforce some sort of authentication mechanism. This can include the use of passwords, personal identification numbers (PINs), biometrics, smart cards (sometimes referred to as tokens), and so forth.

Biometric Authentication Systems

The best authentication mechanisms are biometrics (bio = life, metric = to measure). Examples of biometric measurements include retinal scans, facial feature scans, palm prints, fingerprints, and voice recognition. With biometrics, users may or may not need to enter user IDs, but they are clearly authenticated with features that only they possess.

Smart Card Authentication Systems

A smart card is actually a physical card, about the size of a credit card. It is used to store information that a user needs to access the network. It typically contains such information as the user name, private key, and other credentials.

Smart cards are physical entities and are small enough to fit into a wallet. They require a PIN number to provide access to the network. Many security experts consider them an added level of security because they are a form of two-factor authentication. Two-factor authentication consists of something one possesses, and something one knows. A user ID/password combination is simply something one knows. Additionally, while users may be unaware that their user name and password have been compromised, users discover the loss of a smart card the next time they try to log on to the system.

Some organizations use smart cards for authentication for temporary help, including consultants. If a temporary worker leaves and does not return the card, the account information is deleted, and the card is useless.

Whether users are contractors or permanent employees, they sometimes lose or forget their cards. Part of the security plan for smart card implementation should include a strategy for dealing with lost cards.

Password Authentication Systems

The most frequently used authentication mechanism is the use of a PIN or password. This mechanism is ubiquitous because it bears no additional cost (biometric readers and smart card readers can be expensive), and it is familiar to the user (many users are wary of biometric scans because of health worries or concerns about government using the information inappropriately). However, the security of the computing environment can be weakened if users are allowed to select character combinations for their passwords or PINs that are easy to guess or if they share their passwords with other people.

It is very important for an organization that utilizes passwords as its only authentication mechanism to establish good password guidelines and to ensure that every user in the organization follows the guidelines. A good password has the following characteristics:

- It is at least eight characters long.
- It is alphanumeric (consisting of both letters and numbers).
- There are at least two letters, one number, and one special character (such as a punctuation mark or asterisk).
- No proper names, pet names, or dictionary words are used.
- The number is not at the beginning or end of the password string (unless there are multiple numbers used).
- Lowercase and uppercase letters are used if possible (some systems are case insensitive).
- It appears random.
- It is changed at least every 60 to 90 days.
- It is not reused for at least six months.
- It is significantly different from previous passwords created by the same user.

There are many software tools available that can "crack" passwords by checking them against dictionary words and combinations of dictionary words with numbers, and so forth. These tools are becoming increasingly sophisticated (for example, including foreign words in the dictionary attacks and employing faster processing times) and increasingly inexpensive.

Often, users ask, "How do I create a password with all the characteristics listed above, and still remember it?" Here is an example of a password generation scheme that is easy to remember:

- Come up with a short sentence, with proper capitalization and punctuation. For example, "Windows 2000 is the product for me!"
- Apply the set of substitution rules that follow. The sentence above, after implementing the substitution rules in Table 1, becomes the password: W2K=tp4m! Notice that this password meets all of the characteristics of a good password (that is, mixed case, at least eight characters, at least one number, at least one special character, seemingly random). However, the sentence is easy to remember, so be sure to choose something more difficult.

Because all sentences begin with a capital letter and end with a period, passwords created in this way will likely contain a capital letter and a special character without any additional effort. The following replacement rule suggestions ensure the usage of numbers and special characters:

Table 1. Substitution Rule

Replace this item:	With:
The letter “s”	$
The letter “t”	+
Is	=
For	4
To	2
You	U
And	&
At	@
Not	- or ~
The letter “v”	^
The letter “l”	!
The letter “o”	The number “0”
The letter “c”	(

Note It is insufficient to simply take a proper name or dictionary word and make these replacements. For example, M!cr0$0f+ is not a good password, because the basis of the password is Microsoft, a proper name. All good password decryption tools are programmed with the ability to make the types of replacements discussed above for all words in a dictionary, and then some.

Also note All of the example passwords given in this section are now bad passwords, because anyone can access them. Any password, no matter how “good” it is according to the guidelines, becomes a bad password if it is disclosed. Disclosure of passwords often comes in two forms: either a user tells the password to other users, thereby allowing sharing of his or her account, or a user writes the password down and then places the note on his or her monitor, under the keyboard or mouse pad, or some other place. Proper training of users on the creation and use of passwords minimizes or eliminates these practices. In any case, the security policy should mandate that passwords may not be written down or stored in a text file on a computer. The security policy might make an exception for administrators, who might have multiple, complex passwords. However, the policy should state that if an administrator needs to write down passwords, those passwords must be stored in a sealed envelope and placed in a locked drawer.

Web Access Authentication

When a user accesses a Web server, the usual form of authentication is anonymous (that is, the Web site information is publicly available, and no user identification or password is required). However, when access needs to be restricted, or the identity of the server or client needs to be validated, authentication algorithms are implemented. A Web server may be set to allow multiple types of authentication and thus contain both public and private data. In this case, when an attempt to access data made unavailable to the anonymous account is attempted, the user is required to enter an ID and password.

Access Control

Once users have properly identified and authenticated themselves to the system, they can access the system. Most users, however, have only limited access to the system. Authorization is the mechanism by which user access is determined. User access must always follow the least privilege principle, which states that users may only have the access required to perform their job functions, and no more. Some administrators, on the other hand, should have full access to the system.

The most common way of implementing authorization is with role-based access control (RBAC). RBAC means that people with the same job function should have the same access to the system. This is important for auditing purposes because, if a relatively large number of people are grouped into a relatively small set of access roles, it becomes easier for an auditor to see if someone has access that he or she should not have.

Conversely, there should be enough access roles so that each type of user has only the access needed to do his or her job (in keeping with the least privilege principle). When setting up access roles, there is a very fine line between maintaining least privilege and maintaining RBAC: True least privilege requires many roles with fine distinctions between them, whereas true RBAC requires a minimum number of roles for a maximum number of users. The balance is somewhere in the middle and varies from organization to organization.

Other access control mechanisms include:

- Physical access control
- Authorized usage warnings
- Security settings

Authorized Usage Warning

It is recommended that the system display a warning message to users before allowing them to log on to the system. (The company's legal department might want to check the message's wording.) The message should inform users that the system is for authorized use only and that they could be prosecuted if they misuse the system.

The following is an example of a usage warning scenario: A company decides that its computer systems may only be used for official company business and that unauthorized use or intentional misuse of this system could result in criminal prosecution. The company monitors the computer systems to ensure proper use and to ensure that security mechanisms are not circumvented.

Although the company displays a warning message to this effect before allowing users to log on to the system, some people do not read this warning. However, it is the first place in the system where the policy on proper use can be presented. If some individuals get into trouble for misusing the system, they cannot argue that they were not warned.

Accountability and Shared User IDs

The purpose of auditing is to enforce accountability. Whenever a user creates a security-relevant event on the system, whether it is authorized or not, that event is recorded in the audit log, along with the user ID that caused the event to occur. Therefore, for accountability, it is important that each user have an individual user account and that no one shares user accounts. Most companies now have the policy that if something illegal or improper occurs under a user ID, the user is responsible, even if he or she did not do it. This encourages people to protect their passwords and not share their accounts with others.

Account Lockout

Another way that hackers try to gain access to a system is by what is called a brute-force attack. This is when they keep trying guessed passwords until they guess correctly. A good way to reduce this risk is by configuring the system to lock out a user's account after a specified number of failed logon attempts due to incorrect password entry. It is possible to configure systems so that a locked out account can only be unlocked by an administrator. Thus, users who get locked out have to identify themselves with an administrator and verify that they made a mistake or forgot their password. If users find their accounts locked out and they were not at fault, this may be an indication someone is trying to break into the system.

The account lockout policy controls what happens when incorrect passwords are entered. The policy should be established to allow users who mistype passwords to have the ability to give it another try or two. However, a limit should be established that discourages intruders from continuously trying until they guess a user's password.

Settings to Limit Unauthorized Session Use and Systems Access

One common way for malicious users to access the system is to wait for someone to leave a workstation unattended and gain access through that individual's account. There are several ways to minimize the possibility of this happening. The first is procedural: the company should provide regular security awareness training to all users and educate them regarding the following:

- The importance of locking their screens or logging out of their sessions prior to leaving their desks, even if they think they will only be gone for a few minutes.
- How to lock their screens or log out of their sessions. Proper training is key to any effective security policy.
- The implications of doing so. When a system is locked, all current session activity is still active, but unsaved information is not automatically saved. When the user unlocks the system, the user can continue from where he or she left off. However, if an administrator logs on and unlocks the system (but not the user session) all session activity is closed. Any unsaved data will be lost.

The other mechanism is a simple technical implementation. Configure systems to automatically lock out a user's session after no more than 15 minutes of inactivity. When the system (or the user) locks out the session, the screen appears covered by a solid color or a screen saver so that an observer cannot see what the user was working on, and the user must re-authenticate before returning to the session.

Setting Privileges and Permission on Objects

Privileges and permissions can be set to control access to specific objects (files, folders, printers, and so on). The granularity to which this can be accomplished is dependent on the structure of the operating system and the nature of the file system.

Role-based Access Control and Delegation of Authority

The most effective way to implement proper use of the system is to establish effective access roles within the system. As discussed previously, this is the premise behind role-based access control (RBAC). Access roles in the system are based on user job functions. For example, suppose 10 users all perform the same job function and their job title is "Data Entry Clerk." If they all do the same job, they should all have the same system access, allowing them to perform their assigned job function, but nothing else. To enable this, the administrator would create a Data Entry Clerk access role.

Important advantages to creating access roles include:

- If a new data entry clerk joins the company, it is very easy to assign that person the proper systems access.
- When auditing user access on the system, it is very easy to determine if certain users have somehow exceeded their access and to rapidly fix it.
- If a data entry clerk leaves the company or transfers to another role within the company, it is easy to remove the clerk's system access by removing him or her from the group.

A record of the various access roles in the organization may be entirely documented within the access settings of the operating system or within a spreadsheet or small database containing more complex access roles across the entire enterprise system. For example, data entry clerks might have a particular set of access privileges to the operating system, but they may also have additional access to the company's database and other applications or legacy systems. Both sets of privileges should be documented in one central location, making it easy for administrators to:

- Assign the complete access role to new employees or employees in new positions.
- Thoroughly delete or disable the complete access role for employees who have left the company or changed positions.
- Properly add some access to, or remove some access from, the access role if the job functions for a certain type of employee changes.

In general, an organization's users can be divided into three broad categories:

- *Functional users.* These are the client-facing users that update system data as sales are made or work orders are completed.
- *System administrators.* These are the networking infrastructure support people. If these users also require functional user privileges, they should have two separate user accounts. At no time should a system administrator log on with a privileged account to perform ordinary functional tasks—there is too much risk of making a costly mistake (such as accidentally deleting important files or data that can compromise the system). Also, it is recommended that system administrators' e-mail accounts be attached to their functional user accounts, and not their system administrator accounts.
- *Managers.* These may require a lot of view-only access to generate reports on company status and progress, but they do not need to perform any functional user or system administrator tasks. Managers never update any data in the system.

Functional users are the people with the highest concentration of update access to data, but their access role limits their access to just the data they need to do their jobs. When assigning update access to functional users, system administrators should take care to assign access on a need-to-know basis. For example, a user who has been granted rights to perform administrative duties in the finance department probably does not need extended rights to perform duties in the engineering or human resources departments.

As functional users' areas of responsibility expand (determined by job description and need-to-know access), their access role must also expand to allow them to perform more update functions and become a "senior" functional user. In turn, a senior functional user who has a great deal of expertise and a significant amount of update access to the system would be an ideal application support person.

Users that only have an application support role, however, should have comprehensive view access but should not be able to update anything. There has been a great deal of controversy among industry security experts about how to handle access for application support users. These users are generally considered to be power users. However, general industry consensus is that application support users should be limited to view-only access, unless they also have a functional user role.

Those who are performing limited administrative tasks do not need full administrative rights. It is recommended that systems administrators have segregation of duties just as functional users do. For example, some administrators might handle user accounts, while others handle such server maintenance activities as backup and restore, while still others may handle performance issues. Delegation of authority is appropriate with these types of users. More experienced administrators with higher levels of administrative privileges may choose to delegate lesser administrative functions, such as user administration, to junior administrators.

However, users who have a need to load software on individual workstations may have a legitimate claim to be granted additional access rights to those individual workstations. On the other hand, it is not recommended that any update authority be delegated to any member of management who otherwise has view-only access in the system.

All managers requesting update access for their users should be aware of the above limitations, but it is up to the security manager to ensure that the administrative hierarchy is appropriate.

Confidentiality

Confidentiality mechanisms ensure that only authorized persons can see a particular set of data. Confidentiality is achieved by means of encryption. When data is encrypted, it is scrambled using an encryption key. Encryption keys can be of varying lengths, but various cryptographic algorithms have different key size limits. The key size may be restricted by law. Typical lengths for shared, or secret, key algorithms are 40, 56, 64, or 128 bits in length. The encrypted data is then unscrambled using a copy of the key or another member of a key pair. Encryption keys are created with a random number generator and are manipulated using a mathematical formula. The strength of the encryption key does not lie in the secrecy of the formula used, but rather in the secrecy of the key value and the size of the numbers used.

Two common divisions of encryption are: private (also known as secret, shared, or session) key encryption and public key encryption.

Private Key Encryption

Private key encryption is the older method and is currently more widely used than public key encryption. Private key encryption is used for secure data storage and the protection of transmitted data. Kerberos version 5 protocol uses private key encryption to protect data transferred and stored as part of its authentication algorithms. Private key encryption of transmitted data works as follows:

1. The sender of data generates a secret key.
2. The receiver of data must generate, or be provided with, an identical key.
3. The sender of data encrypts the data with his or her copy of the key.
4. The sender sends the data to the receiver.
5. The receiver decrypts the data with his or her copy of the key.

Note A similar mechanism could be used for secure data storage.

In private key encryption, the encryption keys must remain secret, otherwise the mechanism is useless. Thus, the biggest concern with this mechanism is how to securely exchange and distribute secret keys.

Kerberos Authentication

A client and a server that are each trying to determine the legitimacy of the other might use Kerberos authentication. Kerberos protocol works with private key cryptography. That is, both the Kerberos-enabled client and server share a secret key. Kerberos authentication is a complex process consisting of multiple message exchanges. Two types of secret keys are used. First, each client and server in the domain has its own secret key or password. Second, session keys are generated by the server and distributed by encrypting them with the secret keys of the clients and resources. Because the secret key of the client and the secret key of the server are both required to decrypt and obtain the session key, communications between the two can be kept confidential.

The Kerberos authentication process is described as follows:

1. The client sends several pieces of information about itself (including the current time in clear text) in a request for authentication to a Kerberos Distribution Center (KDC).
2. The KDC encrypts the client's current time with its copy of the client's hashed password in its database. If there is a match, then other checks are made to assure the request is not *spoofed* (the practice of tricking users into providing passwords and other information to allow unauthorized access into a system). For example, the difference in time between the clocks on the client and authenticator must not be more than a set amount, called the Kerberos policy skew time, or the request is rejected.
3. The KDC replies by returning a ticket-granting ticket (TGT) and a server authenticator and other information in a reply to the client. The TGT includes a generated session key and other information all encrypted with the secret key of the server. In the client information, a copy of the session key is encrypted with the hashed password (secret key) of the client. The server's authenticator is a copy of the client's system time encrypted with the session key.
4. The client authenticates the server by decrypting the information, retrieving the session key, and decrypting the authenticator. It stores the TGT.

The Ticket-Granting Ticket Exchange

The ticket-granting ticket (TGT) exchange is the process by which the client receives credentials that it can use to authenticate particular resources. Before the client can access resources, even the desktop environment of its own workstation, this exchange must produce a session ticket for that resource. The possession of a session ticket does not guarantee access to files, folders, and other resources. At that point, normal access control procedures take place. The TGT exchange process is described as follows:

1. The client sends the TGT and other information to the KDC along with a request to access a resource. The client includes a fresh authenticator. A TGT may be reused, but an authenticator may not. This helps to protect the system from replay attacks.
2. Since the TGT is encrypted in the key of the server, the server can use it to decrypt the ticket and retrieve the session key embedded in the TGT. Several checks are used to verify the validity of the TGT.
3. If all is correct, the KDC creates a session ticket for the resource the client has requested. The session ticket includes a session key for use by the resource and the client. The ticket is encrypted in the secret key of the resource computer. A copy of the session key is encrypted in the session key of the client. The session ticket is returned to the client.

The Client Service Exchange

The Client Service Exchange is the process that allows the client to begin negotiating access to the resource. Its sequence is as follows:

1. The client sends the session ticket to the resource along with its request.
2. The resource server decrypts the session ticket and retrieves the session key.
3. Various checks are used to prove the validity of the session ticket.
4. The process of negotiation of the resource access can be followed.

Public Key Encryption and PKI

Public key encryption is a newer encryption method than private key encryption and is becoming increasingly popular with organizations involved in e-commerce. Public key infrastructure (PKI) is the sum of the tools, components, and processes involved in the use of public key encryption. PKI is more expensive to implement, is slower than private key encryption, and a bit less user-friendly. However, this mechanism does not rely on the exchange of secret keys. In this system, two keys, one public and one private, are generated for each user. Private keys are never exchanged or distributed.

Public key encryption is often used in key exchange algorithms to securely transmit secret keys. The secret keys are used to encrypt the bulk of messages.

PKI works on the basis of key pairs. The key pair is mathematically related in such a way that data encrypted with one key can be decrypted by the other. Each user has a public key and a private key. The public key is contained in a digital or key certificate so that it can be identified and made publicly available. A user applies for a key certificate from a certificate authority (CA). The private key is not part of the certificate but is stored securely on the client or in a hardware device that is accessible to the client.

Note A company can be its own CA for internal needs. That is, the company can issue and manage its own digital certificates for internal needs. Certificate generation and management are executed using a feature such as Windows Certificate Services. For external digital certificate needs, other providers, such as VeriSign, can be used as a CA.

The key certificate contains the user's public key. As the name implies, a public key is an encryption key that anyone can access. The user's private key is generated on the user's local computer and is never sent anywhere else. The user must ensure the secrecy of the private key.

PKI encryption can be used to encrypt data for transmission across a network or for data storage. Typically, because of the slow speed of public key encryption, a combination for public key and private key encryption may be used. Public key encryption works in the manner described as follows:

1. The data sender obtains the data receiver's public key.
2. The data sender encrypts the data with the data receiver's public key and sends the data.
3. The data receiver receives the data and decrypts it with its own private key.

PKI is a better mechanism for ensuring data confidentiality and integrity. Digital signatures, which are an integral feature of PKI, can also be used when nonrepudiation is required.

Digital Certificates

A digital certificate is a data structure that contains the public key of a public/private key pair and identification information and is signed by the private key of the issuing CA. The certificate is used to bind the public key to the security principal (that is, users and computers). The information included includes the name of the owner of the certificate, the uses of the certificate (authentication, data encryption, smart card logon, and so on), and the origin of the certificate (which CA or CA hierarchy created it). The certificate is digitally signed by the CA's private key. To check the authenticity of the certificate, the public key of the CA can be used. The certificate may be in a directory, file, or browser, or it can be attached to an electronic message. The certificate is used for security purposes. For example, a digital certificate enables the receiver of a message to verify the identity of the sender.

Certificates are used by user accounts, by computers, in IPSec communications, and by special services. The CA itself has a certificate.

Certificate Policies and Practice Statement

While certificate policies define what a certificate should be used for, the Certificate Policies and Practice Statement spells out the management practices that control the CA and the certificates it issues. It identifies how certificate policies are implemented. It also explains the operating policies, system architecture, physical security, and computing environment.

- Certificate policies might include:
- Uses of available certificate type.
- Authentication of users to the CA.
- How private keys are managed (such as requirements for storing them on smart cards) and whether they can be exported.
- User policies, such as what users should do if their private key is compromised.
- Cryptographic algorithms to be used.
- Length of public and private keys.
- Certificate lifetime.
- A Certificate Policy and Practice Statement should include:
- Identification of the CA.
- Certificate policies.
- Types of certificate issued by the CA.
- Policies and procedures.
- Cryptographic algorithms and key length for the CA.
- CA certificate lifetime.
- Physical network of the CA.
- Security of the CA.
- Policies for renewal of the CA certificate.
- Policies for Certificate Revocation Lists (CRLs).

Virtual Private Networks

A virtual private network (VPN) uses encryption to provide both confidentiality and data integrity. VPN technology is new and preferred by organizations over more traditional ways of securing communication lines because of its versatility and cost-effectiveness. VPN technology creates a virtual tunnel between a remote client and a central server over the existing WAN infrastructure. This can eliminate the need for leased lines (by using a public network such as the Internet), yet it has the same privacy benefits. If all data packets sent through the VPN tunnel are encapsulated and encrypted, they are protected from being viewed, lost, modified, or garbled along the way.

Firewalls, VPNs, and Packet Filters: Protecting the Internet Interface of the VPN Server

The usual configuration of a server used as a VPN end-point requires at least two network interfaces. One interface connects the server with the internal network, while the other connects it to the Internet. To protect the VPN server, one should do the following:

- Routing on the interface should be static and point to the internal network by using static routes instead of using a routing protocol.
- A routing protocol should reside on the private network interface.
- Configure packet filtering in the remote access policy profile for user groups, permitting or denying specific types of Internet protocol (IP) traffic.

To accomplish these goals, one approach includes the use of a firewall. The firewall may be placed between the VPN server and the Internet, or between the VPN server and the intranet.

If the VPN server is between the Internet and the firewall, packet filters can be configured on the Internet interface of the VPN server. Packet filters for Point-to-Point Tunneling Protocol (PPTP) should be:

Drop all packets except those that meet the following criteria for input filters:

- Destination IP address of VPN server Internet interface, subnet mask 255.255.255.255. and TCP destination port 1723 (PPTP tunnel maintenance traffic from PPTP client to PPTP server).
- Destination IP address of VPN server's Internet interface, subnet mask 255.255.255.255 and IP protocol ID 47 (Generic Route Encapsulation routing protocol used by PPTP – PPTP tunneled data from client to server).
- Destination IP address of VPN server Internet interface, subnet mask of 255.255.255.255 and TCP (established) source port of 1723 (only if VPN server also acts as VPN client in a router-to-router VPN. The TCP traffic can only be accepted if the VPN server initiated the connection.

Drop all packets that do not meet the following criteria for output filters:

- Source IP of VPN server's Internet interface, subnet mask 255.255.255.255 and TCP source port of 1723 (PPTP tunnel maintenance form VPN server to client).
- Source IP or VPN server Internet interface, subnet mask of 255.255.255 and IP protocol ID of 47 (tunneled data from VPN server to VPN client).
- Source IP address of VPN server Internet interface, subnet mask of 255.255.255.255. TCP established port of 1723 only if VPN server is acting as VPN client.

Packet filters for L2TP over IPSec

Create input filters to drop all packets that do not contain one of the following:

- Destination IP address of the VPN server's Internet interface, subnet mask of 255.255.255.255, and UDP destination port of 500
  (Internet Key Exchange (IKE) traffic to the VPN server).
- Destination IP address of the VPN server's Internet interface, subnet mask of 255.255.255.255, and UDP destination port of 1701
  (L2TP traffic from the VPN client to the VPN server).
- Configure output filters that drop all but the following:
- Source IP address of the VPN server's Internet interface, subnet mask of 255.255.255.255, and UDP source port of 500
  (IKE traffic from the VPN server).
- Source IP address of the VPN server's Internet interface, subnet mask of 255.255.255.255, and UDP source port of 1701
  (L2TP traffic from the VPN server to the VPN client).

If the firewall is between the VPN server and the Internet, the firewall should be configured to allow these packets, with the exception of 1701. At the firewall, all L2TP packets are encrypted with IPSec and, therefore, pass through the firewall. Other filters may be configured to allow other packets for other purposes (such as those to Web servers).

The firewall knows nothing of the certificates required for IPSec connections. Therefore, all tunnel traffic is allowed to pass. Since all packets are authenticated, this policy should not be a security risk.

Layer 2 Tunneling Protocol

Layer 2 Tunneling Protocol (L2TP) is a protocol that encapsulates, or tunnels, Point-to-Point Protocol (PPP) traffic between different networks and multiple communication links. The key benefits of L2TP are:

- It is an industry standard: it encompasses the attributes of a variety of other proprietary protocols developed by, among others, Microsoft Corporation and Cisco Systems.
- It supports multiple protocols—not just IP, although IP is the primary protocol used.
- It supports multiple, simultaneous tunnels for a single client.
- It supports dynamic addressing.
- It permits centralization of logon and authentication by co-locating a network access server (a point of access to the private network from the public network) with an L2TP access concentrator.

Layer 2 Tunneling Protocol does not provide security. Instead, security features, such as authentication and link-level encryption, are provided by Point-to-Point Protocol (PPP) and the Internet Protocol Security (IPSec). IPSec uses the Internet key exchange (IKE). Certificates are used to authenticate the computers. User authentication can be by smart card, certificate, or token card by using the Extensible Authentication Protocol (EAP). Computers on either end of the tunnel establish trust by using computer certificates. Trust must also be negotiated between the two computers using these certificates. This allows the securing of transactions from other computers on the public network that might try to spoof an authorized session. The connection process is described as follows:

1. The L2TP client uses IPSec to negotiate a security association (SA—a set of parameters that defines services necessary for IPSec secure communications).
2. The remote access policy of the client is checked. A remote access policy is defined by properties that enable or disable remote access for a user account in the domain.
3. The connection is made.

There are two L2TP models. In the carrier/ISP model, either the carrier or the Internet service provider's access device needs to support L2TP. In the client model, the user devices that connect to the carrier or ISPs network themselves need to support L2TP. Both models are described in the following sections.

Carrier/ISP Model

In the carrier/ISP model, the L2TP access concentrator (LAC—one endpoint of an L2TP tunnel) is generally the network access server (NAS) to which remote users dial in for their VPN service. The L2TP network server (LNS—the other endpoint of an L2TP tunnel) is usually a perimeter device (for example, a firewall or router) on the enterprise network.

The end user usually connects to the LAC without encryption. The LAC then identifies the client via Client Identification (CLID), Data Network Identification Code (DNIC), or Password Authentication Protocol (PAP), retrieves an appropriate LNS address from the VPN database, and creates a tunnel to the destination LNS. From then on, the tunnel carries the PPP datagrams between the LAC and the LNS.

Many sessions can be multiplexed over a single tunnel. A control connection manages the initiation, termination, and maintenance of sessions and of the tunnel itself. Additional security is provided by the use of IPSec features by the LAC and the LNS data traveling inside the tunnel.

Client Model

In the client model, the LAC is generally a small office, home office, or branch office router, or an L2TP client running on a remote computer. The LNS is usually a perimeter device (for example, a firewall or router) on the corporate network.

The client, typically a remote device or computer, dials into the nearest NAS or to an Internet Service Provider (ISP) Point of Presence (POP) server. The NAS authenticates the client and allocates a global address to it. A tunnel connection is established with the remote LNS, which then carries a new PPP session during which the LNS at the central site again performs the authentication and assignment of a private (or global) address via a second exchange. Key exchange, compression, and encryption are also negotiated.

L2TP Security

As mentioned earlier (in the "Layer 2 Tunneling Protocol" section), L2TP relies on other protocols for its security. L2TP authentication is best for the exchange of packets between the LAC and the LNS. Therefore, it is advisable to make IPSec-based authentication a part of L2TP.

If confidentiality is desired with L2TP, the client is generally responsible for encryption.

For L2TP, authentication is typically LCP-negotiated by Password Authentication Protocol (PAP) or Challenge Handshake Authentication Protocol (CHAP) or is negotiated via the PPP Extensible Authentication Protocol (EAP). However, PPP has no per-packet authentication mechanism. As L2TP is now thought to need better security for the exchange of IP/L2TP packets between the LAC and the LNS. IPSec-based authentication is recommended.

Internet Protocol Security

Internet Protocol Security (IPSec) is a new standard developed by the Internet Engineering Task Force (IETF) to provide secure tunnels across untrustworthy IP networks, including the Internet. Unlike L2TP, which functions at the media access control (MAC) layer (layer 2), IPSec operates at the IP layer (layer 3). The purpose of IPSec is to provide interoperable, high-quality, cryptographic-based security. IPSec is being touted as being the next preferred tunneling protocol.

The key benefit of IPSec is that it provides authentication, confidentiality, and integrity, but only the two endpoints of the communication need to be IPSec-enabled. Regardless of what components (routers, switches, servers, and so forth) the data touches while en route, none of those components need to know anything about IPSec for the data to safely and intactly pass.

The two endpoint computers create a virtual IPSec tunnel by enacting a security policy. The security policy includes the addresses of the endpoints, the encapsulation method, the encryption algorithms, parameters for things such as key length and lifetime, and so forth. The IPSec tunnel is the embodiment of all of the procedures and protocols that enable data to travel safely between the endpoints. Put together, this set of procedures and protocols is called a security association (SA).

The key weaknesses of IPSec are twofold. First, IPSec tunnels do not support all the protocols currently supported by PPP and L2TP. Second, since IPSec tunnel peers have global addresses, they could be open to denial of service attacks. However, there are simple solutions to overcome this second issue.

IPSec operates in two modes: tunnel mode and transport mode. IPSec tunnel mode is designed to secure one or more tunnels between a pair of hosts, or between a pair of security gateways, or between a security gateway and a host. Tunnel mode is the mode that applies to VPNs using IPSec technology. IPSec transport mode applies only to communications between two hosts.

Security Protocols

IPSec utilizes two protocols to support secure operations: the Authentication Header (AH) and the Encapsulating Security Payload (ESP). AH and ESP can be used individually or together during an IPSec session. Both protocols are described in greater detail in the following sections.

Authentication Header

The Authentication Header (AH) provides authentication support for the IP packet. In transport mode, the AH is inserted in the datagram between the IP header and the Transmission Control Protocol (TCP) header. The AH requires the use of keys to verify the integrity of the information in the Authentication Header itself. However, the AH does not encrypt data, so it does not provide any kind of data confidentiality.

The AH does not mandate a specific authentication protocol; rather, the two endpoint systems must agree on a common authentication protocol through a security negotiation process.

Encapsulating Security Payload

The Encapsulating Security Payload (ESP) protocol provides authentication, integrity, and confidentiality. However, although ESP does provide authentication, the original IP header prior to the addition of the ESP header is not authenticated. If that original IP header needs authenticating, then AH must be used as well.

In transport mode, the datagram has the following components: the original IP header, the ESP header, the TCP header, the data, the ESP trailer, and ESP authentication. In tunnel mode, the datagram has slightly different components: an additional IP header is inserted, designating the destination tunnel endpoint. The ESP header follows, then the original IP header (which is encapsulated by the other IP header), and the TCP header. The last three components are the same as in transport mode.

Security Associations

As mentioned above, two systems communicating with IPSec require a security association (SA) to define the security policies by which they communicate. The benefit to this method is its versatility. VPNs can involve communications among a variety of users (for example, internal users, customers, suppliers, subcontractors, and mobile users). These different users likely have different security needs and probably use different security protocols. The security association enables users to negotiate a common set of security attributes in an authenticated and protected manner.

The SA defines how the communicating systems use security services, including information about the traffic security protocol, the authentication algorithm, and the encryption algorithm that is used. SAs also contain information on dataflow, lifetime, and life data as well as sequence numbering for anti-replay. It should be noted that if a server has multiple sessions with multiple clients, then the server has a different SA for each of the clients.

An SA is unidirectional. That is, for each pair of communicating systems there are at least two security connections—one from A to B and one from B to A. Also, a given SA can only use one ESP or AH, but not both. If a connection needs both AH and ESP, two SAs must be established for each direction (giving a total of four for a two-way connection).

There are three key parameters connected with a security association:

- An encryption algorithm—typically Data Encryption Standard (DES) or 3DES.
- A session key—via IKE.
- An authentication algorithm—typically Secure Hash Algorithm (SHA-1) or Message Digest 5 (MD5).

Another component of importance is the Security Parameters Index (SPI). The SPI is derived from a randomly generated number and the destination host's IP address. The purpose of the SPI is to identify each SA as a unique entity, separate from any other IPSec connections that may be in existence concurrently.

Two databases are required for security associations. The first is a security policy database (SPD). The SPD specifies all the security services and policies that apply to a given communications session. The second is a security association database (SAD). Each entry in the SAD defines the parameters associated with one SA.

Once the security attributes have been established, initial identity authenticated, and required keys generated, the SA can be used for subsequent communications by the invoking system. There are no specific requirements for encryption algorithm, key generation technique, or security mechanism, provided that strong authentication mechanisms are used (specifically, digital signatures and digital certificates created by a trusted third party or certificate authority).

As with AH and ESP, security associations can be used in either transport or tunnel mode. The major advantage of tunnel mode is that the end systems do not need to be modified to obtain the benefit of IP security, only the gateways. Tunnel mode also protects against traffic analysis, in that an attacker can only determine the tunnel endpoints, but not the true endpoints of the packets, even if they are the same as the tunnel endpoints. Transport mode is less secure than tunnel mode because it does not conceal the IP control information

File System Confidentiality

Data stored on disks or other storage media needs to be kept secure. The most obvious method to secure the data is to prevent unauthorized access to the disk or directory. If the physical disk is stolen, however, special tools can be used to access the data it contains by bypassing the security mechanisms provided by the operating system. Portable computers (laptops, notebooks, and so on), backup tapes, and other assets that are easily stolen are especially at risk. Therefore, it is recommended that sensitive data be encrypted, in addition to being protected via access controls. That way, even if the disk falls into the wrong hands, no sensitive information can be retrieved from it.

Integrity

Data integrity mechanisms ensure that data is not garbled, modified, or lost during transmission across a network. It also ensures that data comes from the intended sender, and not from an impostor. Data integrity mechanisms include checksums and digital signatures.

For the most part, data integrity is automatically maintained if one also has data confidentiality. That is, if only the right people have access to the data, then it should be safe from improper modification. This theory breaks down, however, when talking about viruses. A key goal of the computer virus is to delete or otherwise mutilate system data. It is very important that enterprise systems have virus-checking software installed to guard against the possibility of viruses entering the system and destroying data. However, as with any security feature, the virus scanner is only as good as its implementation. The scanning tool must be properly installed (with vendor-provided updates promptly loaded) and the tool run as often as recommended by the vendor to ensure that viruses do not sneak into the system unnoticed.

Nonrepudiation

Digital signatures are very useful for nonrepudiation, as well as for data integrity. Nonrepudiation is the security concept that applies to proving the transmission of a particular message. If a system does technical nonrepudiation, then the sender of a message cannot later deny having sent the message, and the receiver of a message cannot later deny having received the message. Furthermore, if the message contains contractual information, the presence of a digital signature with the message can ensure that the contractual information was not improperly altered at a later time.

E-mail can use public key encryption to digitally sign messages.

Digital signatures are created by using the private key of the account that is used to encrypt a message digest. This signature is sent with the message. The message may or may not be encrypted. Upon receipt, the public key of the sender is used to decrypt the signature. This message digest is compared to a message digest created by the same algorithm as the original, encrypted digest. If the two messages (the one sent and the one decrypted from the signature) match, the message is considered to have been sent by the signer. Since only the private key that is paired with the public key could have been used to encrypt the digest, the message is also considered to be free of tampering. If an attacker changes the message, the new digest created by the receiver does not match the decrypted one. If the attacker attempts to change the signature by using another key, the receiver will not be able to decrypt it with the public key of the supposed sender.

Auditing

An audit log records an entry whenever users perform certain specified actions. For example, the modification of a file or a policy can trigger an audit entry. The audit entry shows the action performed, the associated user account, and the date and time of the action. One can audit both successful and failed attempts at actions.

The state of the operating system and applications on a computer is dynamic. For example, security levels may be required to change temporarily to enable immediate resolution of an administration or network issue; this change can often go un-reversed. This means that a computer might no longer meet the requirements for enterprise security.

Regular analysis enables an administrator to track and ensure an adequate level of security on each computer as part of an enterprise risk management program. Analysis is highly specified information about all system aspects related to security. This enables an administrator to tune the security levels and, most importantly, detect any security flaws that may occur in the system over time.

Auditing is not considered one of the six security tenets because it is not a security-specific function. However, security auditing is extremely important for any enterprise system, as audit logs may give the only indication that a security breach has happened. If the breach is discovered some other way, proper audit settings generate an audit log that contains important information about the breach.

For proper security audit settings, the following should be recorded:

- Log on and log off activity, including network and remote connections.
- File and object access, including access to files and directories, and sending print jobs.
- File and object creation or deletion.
- Access to user privileges, except those related to logon and logoff activities.
- User and group management, including creation, deletion, renaming, and other changes to user accounts and passwords.
- Changes to security policy.
- System administrator functions such as system restarts, shutdowns, and security functions of the system.
- Process tracking for processes running in the system, including duplicate processes and program exits.
- Both the success and failure of the items listed above should be audited. Often, failure logs are much more informative than success logs, since failure more often indicates an error. For example, if a user successfully logs on to the system, this would be considered normal. However, if a user unsuccessfully tries to log on to the system multiple times, this may indicate someone is trying to break into the system using someone else's user ID.

Planning Auditing

Planning is an important step in the auditing process. One should be selective in determining the objects to audit. Auditing creates system overhead, and auditing too many objects causes security logs to become large and difficult to manage. Be selective, record the selections, create a plan, and test it.

An auditing policy must also be established. The policy defines the types of events to be audited for a specific user or group of users.

Implementing Auditing

Implementing auditing includes the following steps:

1. Enable auditing.
2. Select the object to audit, and set the system access control list (SACL) for the object.
3. Configure the event log.
4. Test the audit configuration.

Testing Auditing

When the configuration is complete, log on multiple times and accesses the objects that have been configured. Log on using a variety of test user accounts that have been configured with differing rights. Then log on using an administrator account and check the logs.

One can do this by using the Event Viewer console. Browse through the security log and view the events, which can be displayed according to the event category to which they apply. These entries are the audit trail of those who have gained (or attempted to gain) access to the audited objects. The information in the trail is dependent upon the defined configuration.

Roles and Responsibilities

Principal roles and their associated responsibilities for security administration have been defined according to industry best practice. Organizations might need to combine some roles, depending on organizational size, organizational structure, and the underlying service level agreements existing between the IT department and the business it serves.

In a small organization, it may be sufficient to have a security manager who oversees all of the security aspects of the organization. However, for most large organizations, it may be desirable to have a security manager and a dedicated security team, which should include one or more security administration experts.

In addition, it is advisable to form a security council whose members include representatives from the different business and technical support organizations in the enterprise. Security councils typically have regular meetings to develop and update policy, review and respond to recent security incidents, and plan the future of security in the organization.

The following sections describe the roles and responsibilities in a security organization. The roles are part of the security role cluster defined in the MOF team model. A single member of the organization may assume responsibility for more than one role.

Security Manager

The security manager is the process owner for the security administration process. The responsibilities of the security manager include:

- Detecting intrusions and protecting against viruses.
- Defining policies for data retention and secure data disposal.
- Performing audit tracking and reporting.
- Providing effective network domain security design and management.
- Testing and implementing strategic security technology.
- Monitoring and assessing network vulnerability.
- Monitoring and assessing third-party vulnerabilities.
- Providing fast, real-time network intrusion response.
- Managing authentication and access methods requirements.
- Managing user policy usage and requirements (such as a password policy).
- Managing external and physical security requirements (such as access to computer rooms).
- Managing secure messaging requirements.
- Providing ongoing technical support and subject matter expertise for security initiatives within the company.

Personnel Security Technician

The personnel security technician works with the security manager to:

- Ensure that only authorized personnel are granted access to critical systems and facilities.
- Conduct background checks on employees to verify identity.

Antivirus Technician

The antivirus technician works with the security manager to:

- Understand the systems present and the type of vulnerabilities to which they are susceptible.
- Ensure that antivirus systems are in place and operating correctly.

Application Security Technician

The application security technician works with the security manager to:

- Ensure that only authorized users gain access to critical business applications.

Database Security Technician

The database security technician works with the security manager to:

- Ensure that data is confidential.
- Ensure that data is available only to authorized personnel.
- Ensure that database auditing and journaling are in place where appropriate.

Messaging Security Technician

The messaging security technician works with the security manager to:

- Ensure that messages are confidential and free from tampering and repudiation.

Operating System Security Technician

The operating system security technician works with the security manager to:

- Ensure that strong security measures are in place, including but not limited to:
  - Strong passwords.
  - Encrypted file systems.
  - Biometric authentication systems.
- Ensure that all users and/or processes default to the least privilege required.

Hardware Security Technician

The hardware security technician works with the security manager to:

- Ensure that hardware computing resources are secure from pilfering and sabotage.

Network Security Technician

The network security technician works with the security manager to:

- Ensure that network communications are secure and free from tampering and/or eavesdropping.

Facilities Security Technician

The facilities security technician works with the security manager to:

- Ensure that only authorized personnel gain physical access to the building and/or computing assets.
- Create emergency response plans so that personnel and assets are safe in the event of a mishap.

Egress Security Technician

The egress security technician works with the security manager to:

- Ensure that critical utilities are unadulterated and free from tampering.
- Ensure that proprietary information is disposed of in a secure way and rendered inaccessible.
- Ensure that corporate refuse is disposed of in accordance with environmental regulations.

Outsourcing Manager

The outsourcing manager works with the security manager to:

- Assessing and minimizing the security risk that a supplier possesses.

Security Compliance Auditor

The security compliance auditor works with the security manager to:

- Audit the efforts of all members of the security team (personnel security technician, anti-virus technician, application security technician, database security technician, messaging security technician, operating system security technician, hardware security technician, network security technician, facilities security technician, egress security technician, and outsourcing manager) for compliance with the standards set by the security manager.
- Evaluates risks to the enterprise as a result of the security audit.

Relationship to Other Processes

Security administration can impact an entire information system and many other processes as well. In fact, Security administration acts as an "umbrella" process in the operating quadrant of the MOF process model. Every other process must conform to the guidelines set forth in this document. The graphic below depicts the relationship between security administration and the other MOF SMFs.

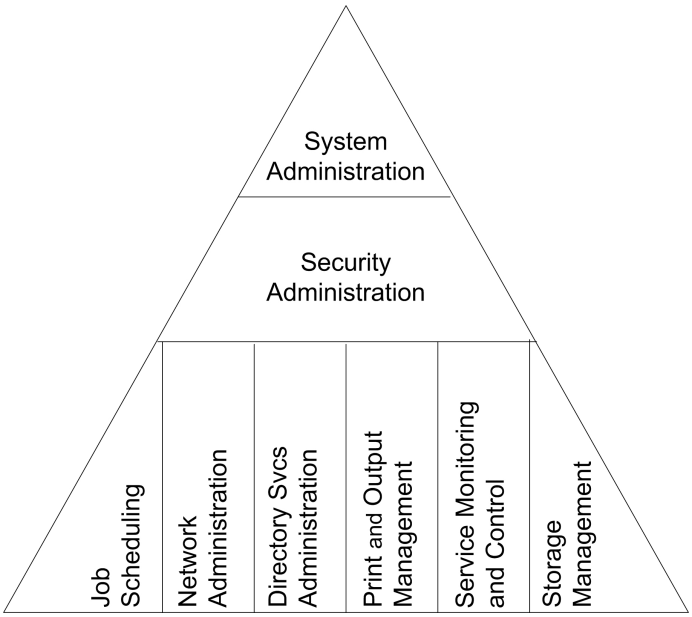

Figure 1
Relationship to other SMFs in the Operating quadrant

System Administration

System administration deals with the administration model used by an organization. Some organizations prefer a model where all IT functions are performed at a single site by a team of IT professionals collocated at that site. Other organizations prefer a distributed branch-office model where both technologies and support staff are geographically distributed. System administration examines the trade-offs of each model. Each type of system administration model has unique security requirements. Security administration ensures that system security is not compromised regardless of the model chosen.

Service Monitoring and Control

Service monitoring and control oversees the various aspects of system performance to ensure that service level agreements are being met. One security component that frequently affects performance is auditing. Administrators in charge of monitoring systems performance must assist the security manager in generating a security audit log. They must also be familiar with the performance cost associated with generating the security audit log so that ongoing planning properly accounts for this feature.

Job Scheduling

Job scheduling deals with performing batch processing tasks during times that maximize system resources but do not compromise business and system functions. It is important that any security-related jobs be appropriately scheduled and executed and that problems with and results of security jobs are promptly communicated to the security team. Examples of security batch jobs might include a script that validates that active users on a system are still valid employees in the personnel database, or a script that compares audit logs between the operating system and the database to ensure that things like logon and logoff times correspond.

Depending on the environment and the schedule, security tasks may need to be scheduled and consistently executed at certain times of the day, week, or month to ensure that only authorized users are active in a system. Such tasks often provide users with warnings prior to the end of their authorization period, and then finally log them off a system if they remain active. These jobs may be critical to ensuring the integrity of the data in some environments where, for example, the balancing of financial figures is essential at the end of the business day. The execution of such jobs must be coordinated carefully within the IT environment, as well as with business management, and then be monitored for successful execution prior to starting the next phase of system operation.

Network Administration

Network administration deals with the maintenance of the physical components that make up the organization's network, such as servers, routers, switches, firewalls, and so forth. An improperly set up or inadequately protected network can be a tremendous security risk. Network administrators must ensure proper physical security of network components to prevent unauthorized access. Network administrators must also be familiar with proper firewall configuration and maintenance.

Directory Services Administration

Directory services contain all user and system profiles. Directory services administration deals with properly configuring and modifying object profiles to optimize functionality and security in a system. It is extremely important that directory services administrators be familiar with the organization's security requirements, as they are the ones who manage the area where security information is stored and maintained.

Print and Output Management

Print and output management deals with all data that is printed or compiled into reports that are distributed to various members of the organization. The print and output management team must ensure that any sensitive printed material is properly secured (for example, quickly removed from the printers, not left out where anyone can view the material, and so forth) and that reports in electronic format are protected either with encryption or some access control mechanism so that only authorized users can view that information.

Storage Management

Storage management deals with on-site and off-site data storage for the purposes of data restoration and historical archiving. Storage management has a strong security component in that it is often a requirement that sensitive data—for example, cost or revenue information—be kept in storage the longest time for auditing purposes. The storage management team must ensure that the physical security of backups and archives is sound. As encryption technologies improve and become more ubiquitous, physical security is required to ensure that backup media are not damaged. However, as most backups are not encrypted, physical security of the backup media must also ensure that there is no unauthorized access to those archives. The storage management team must also ensure proper purging of backup media once the data is no longer needed. The backup media must either be overwritten methodically enough times to ensure that no one can obtain any of the original data (in which case the medium can be reused) or, in the case of sensitive data, the backup medium should be destroyed (for example, by degaussing or dipping in acid) so that it cannot be used or read.

Configuration Management

There is a strong security component to configuration management. Configuration management deals with keeping track of the hardware owned by an organization and the versions of internal software that are used. Administrators should be aware of and in full control of the versions of the operating system, database management system, and all applications running on network computers. Poor configuration management could facilitate the introduction of malicious code into an operating system(s) or into an application.

Availability Management

Availability management deals with overall system availability versus downtime. Since most organizations today are virtually paralyzed when a system is down, it is extremely important that administrators properly configure and monitor a system to maximize uptime and mean time between critical failures. Availability management also has a strong security component in that an increasingly popular outsider security attack is the denial of service attack. A denial of service attack attempts to bring system availability down to zero for as long as possible.

Capacity Management

Capacity management deals with planning for additional resources as current system resource use increases and begins to near the point of full capacity. Capacity management therefore ties directly to monitoring and measurement, as well as to availability management. Security affects capacity management in the same way, by affecting the amount of system resources required to generate the security audit log. The audit log grows larger as additional servers, databases, applications, and users are added to an enterprise system. If the enterprise begins to store and use data that is more sensitive than data previously stored and used in a system, it may have to turn on even more auditing features to make sure that data is adequately protected. This, in turn, requires greater capacity for the security audit logs. It should be noted that systems dealing with highly sensitive information are often programmed to shut down if the audit log gets filled to capacity. The capacity management team needs to be aware of this condition and help ensure that it does not occur.

Release Management

Release management deals with all aspects of the decision to upgrade some or all system software, including operating system, database management system, and all applications. Release management is closely associated with configuration management. Security is an important concern when doing any type of upgrade. Release management ensures that the existing security infrastructure is maintained. Improvement is ideal, but no upgrade should ever be implemented if it degrades system security. Be sure to understand the implications that the proposed upgrades have on security up front. Also, be sure to test upgrades in a development environment.

Change Management

Change management deals with the coordination of any change that occurs within the organization, including software upgrades, entire system overhauls, organizational or personnel changes, business changes, and so forth. It is the responsibility of the change management team to ensure that all affected parties are involved in the change process. It is extremely important that security experts are aware of changes within the organization.

Problem Management

Problem management deals with any network problem that affects a number of system users. Often, the help desk staff discovers problems when a number of people call in with the same complaint. Problem management can have a big security component. If, for example, the problem happens to be that the network is down or unusually slow, a denial of service attack should be considered as a possible cause. In such a case, the security team should be involved in fixing the problem. Security personnel should be involved if there is any question that a problem is being caused by a security feature (or inadequate security) or if the fix to the problem might affect security.

Service Desk

When users have trouble with the system, their first (and sometimes only) line of support is from the service desk. One of the most common problems users have is that of locking themselves out of their accounts or forgetting their passwords. Service desk personnel must be very clear on what they can and cannot do to help users with password problems. Service desk personnel must also be very familiar with the guidelines for good passwords, as well as other user security policies.

Service Continuity Management

Service continuity management deals with automatically changing to an alternate server when a server goes down temporarily, and then transferring back to the main server when it becomes available again. The key security component in this process is that all security information (such as user access rights and audit log configurations) and processes (such as the generation of an audit log) are maintained on the alternate system in the same condition that they are normally maintained on the original system. Also, security audit logs should be protected from deletion when a server suddenly goes down.

Workforce Management

Workforce management deals with all aspects of employee management, from determining what kinds of skills are needed to perform certain tasks, to determining how many people are required for a particular role, to hiring and managing people to do the job. Workforce management ties into security in two ways: performing background checks when employees are hired and performing on-going performance checks. The workforce management team must ensure that proper background checks are performed before allowing the employee entrance into company facilities or access to the organization's computer systems. In addition, disgruntled employees are more likely to attempt to do damage to the organization's resources or data, so it is imperative that employee access be terminated as soon as possible after termination.

Employee turnover must also be carefully managed. As employees move from one role to the next, their security access may need to be adjusted to properly reflect their new areas of responsibility. Employees who leave the organization should have their access disabled after their last day. For these reasons, the security team needs to work closely with organizational management at all levels to coordinate their activity.

Turnover in the IT team itself might pose some additional risk. If individuals leaving a company have access to multiple systems, numerous passwords must be reset, and personal logon information must be disabled.

Financial Management

Financial management deals with the analysis and management of the cost of running a network, data center, or system. This includes the cost of hardware, training, administrator time, CPU time (if it is outsourced and billed), new software upgrades, and so forth. Frequently, security is either left entirely out of the cost picture, or an inadequate assessment of the cost of security is made, and thus the allotted security budget is insufficient. It is extremely important that a thorough risk analysis be conducted before annual budgets are determined so that it is clearly understood how much budget is required to properly implement security.

Contributors

Many of the practices that this document describes are based on years of IT implementation experience by Accenture, Avanade, Microsoft Consulting Services, Fox IT, Hewlett-Packard Company, Lucent Technologies/NetworkCare Professional Services, and Unisys Corporation.

Microsoft gratefully acknowledges the generous assistance of these organizations in providing material for this document.

Program Management Team

Jeff Yuhas, Microsoft Corporation

William Bagley, Microsoft Corporation

Lead Writer

Ioana Bazavan, Accenture

Contributing Writers

William Bagley, Microsoft Corporation

Roberta Bragg, Have Computer Will Travel, Inc.

Vicky Howells, Fox IT

Jeff Yuhas, Microsoft Corporation

Editors

Steve Morgan, Fox IT

Patricia L. Rytkonen, Volt Technical Services

Sybil Wood, Volt Technical Services

Other Contributors

Tom Baker, Microsoft Corporation

Vladimir Bakhmetiev, Microsoft Corporation

Alicia Berend, Microsoft Corporation

Deniz Chan, Microsoft Corporation

Andrew Cheeseman, Microsoft Corporation

Patrick Conlan, Microsoft Corporation

Sandy Coyne, Microsoft Corporation

Curtis Cummins, Microsoft Corporation

Brian Davies, Microsoft Corporation

Walter Dickson, Microsoft Corporation

William Dixon, Microsoft Corporation

Praerit Garg, Microsoft Corporation

Dave Gasiewicz, Microsoft Corporation

Clive Jacobs, Microsoft Corporation

Chunyang Jia, Microsoft Corporation

Lura Johnson, Microsoft Corporation

John Norby, Microsoft Corporation

Greg Parsons, Microsoft Corporation

Laurence Reffold, Microsoft Corporation

Susan Saranovich, Microsoft Corporation

Don Schmidt, Microsoft Corporation

Kirk Soluk, Microsoft Corporation

Diana Spickerman, Microsoft Corporation

Shane Van Jaarsveldt, Microsoft Corporation

Marcus Vilcinskas, Microsoft Corporation

Richard J. Wood, Microsoft Corporation