

realtimepublishers.comtm

*The Administrator
Shortcut Guidetm To*



**Active Directory
Security**

SCRIPTLOGIC

*Dave Kearns
and Beth Sheresh*

Introduction

By Sean Daily, Series Editor

Welcome to *The Administrator Shortcut Guide to Active Directory Security*!

The book you are about to read represents an entirely new modality of book publishing and a major first in the publishing industry. The founding concept behind Realtimepublishers.com is the idea of providing readers with high-quality books about today's most critical IT topics—at no cost to the reader. Although this may sound like a somewhat impossible feat to achieve, it is made possible through the vision and generosity of corporate sponsors such as ScriptLogic, who agree to bear the book's production expenses and host the book on its Web site for the benefit of its Web site visitors.

It should be pointed out that the free nature of these books does not in any way diminish their quality. Without reservation, I can tell you that this book is the equivalent of any similar printed book you might find at your local bookstore (with the notable exception that it won't cost you \$30 to \$80). In addition to the free nature of the books, this publishing model provides other significant benefits. For example, the electronic nature of this eBook makes events such as chapter updates and additions, or the release of a new edition of the book possible to achieve in a far shorter timeframe than is possible with printed books. Because we publish our titles in “real-time”—that is, as chapters are written or revised by the author—you benefit from receiving the information immediately rather than having to wait months or years to receive a complete product.

Finally, I'd like to note that although it is true that the sponsor's Web site is the exclusive online location of the book, this book is by no means a paid advertisement. Realtimepublishers is an independent publishing company and maintains, by written agreement with the sponsor, 100% editorial control over the content of our titles. However, by hosting this information, ScriptLogic has set itself apart from its competitors by providing real value to its customers and transforming its site into a true technical resource library—not just a place to learn about its company and products. It is my opinion that this system of content delivery is not only of immeasurable value to readers, but represents the future of book publishing.

As series editor, it is my *raison d'être* to locate and work only with the industry's leading authors and editors, and publish books that help IT personnel, IT managers, and users to do their everyday jobs. To that end, I encourage and welcome your feedback on this or any other book in the Realtimepublishers.com series. If you would like to submit a comment, question, or suggestion, please do so by sending an email to feedback@realtimepublishers.com, leaving feedback on our Web site at www.realtimepublishers.com, or calling us at (707) 539-5280.

Thanks for reading, and enjoy!

Sean Daily

Series Editor

Introduction.....	i
Chapter 1: Directory Security	1
Using Directories to Manage Network Access.....	1
Directory Security Protects Information and Service Assets.....	2
Why Directory Security Is Essential.....	3
Basic Security Mechanisms	4
Authentication.....	4
Authorization	5
Auditing	6
How AD Provides Security.....	7
Policy-Based Security.....	7
Threats, Vulnerabilities, and Attacks.....	8
Threat	8
Vulnerability	9
Attack.....	10
User-Based Attacks.....	11
Software-Based Attacks.....	12
Environment-Based Attacks	13
Threat Analysis	13
Spoofing.....	13
Tampering.....	15
Repudiation.....	15
Information Disclosure	15
DoS	16
Elevation of Privilege	16
Managing the Directory Service	16
Best Practices for Service Administrator Account Management	17
Best Practices for Managing Directory Information.....	19
User Management in AD	20
Creating a User Object.....	21
Creating a Group.....	22
Summary	22

Copyright Statement

© 2004 Realtimedpublishers.com, Inc. All rights reserved. This site contains materials that have been created, developed, or commissioned by, and published with the permission of, Realtimedpublishers.com, Inc. (the "Materials") and this site and any such Materials are protected by international copyright and trademark laws.

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

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

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

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

If you have any questions about these terms, or if you would like information about licensing materials from Realtimedpublishers.com, please contact us via e-mail at info@realtimedpublishers.com.

Chapter 1: Directory Security

For all networks, systems administrators must keep track of who is accessing the network as well as control each user's access to the various network resources. In most networks, information about users and their access rights are stored in a directory that provides user authentication and access control services.

A directory service typically contains sensitive information about the user and service accounts that have access to the enterprise network and information regarding directory-enabled applications and services as well as other network resources. This information is sensitive in that the unregulated disclosure and/or disruption in the provision of this information and related services can interfere with business operations.

Directory security is fundamentally focused on protecting information, service, and resource assets accessible through the enterprise network. In addition to protecting information stored within the directory, the authorization and access control mechanisms provided by the directory service protect the services and information stored within your network.

Implementing security for the information contained in and the resources protected by Microsoft's directory service implementation—Active Directory (AD)—is not a simple task. Although AD provides powerful management capabilities, these features introduce complexities. You must understand AD, the network, the corporate environment, and the potential threats and vulnerabilities before you can effectively implement security.

In this chapter, we'll explore directory security at a high level before moving onto an exploration of the possible threats and approaches to managing the directory service and information from a security perspective. In later chapters, we'll delve into how the design of the directory impacts security and administration, then we'll take an in-depth look into Group Policies and delegation of directory administration.

Using Directories to Manage Network Access

In a network environment, you need to be able to control which people are able to access information and resources. To accomplish this control, authentication mechanisms must be used to ensure that only the designated people can access your directory, network, or other enterprise resources. To protect the resources on the network, access control mechanisms must be implemented to prevent unauthorized access to network resources and services as well as to prevent the illegitimate modification or deletion of information.

For most companies, a directory service integrated with their network operating system (OS) is the fundamental point of entry and access to all of the resources available on the enterprise network. The directory service handles authentication of users attempting to access the network and the authorization needed to access network resources by managing the users' rights and permissions.

Because the directory service plays a central role in controlling access to enterprise resources throughout the network, carefully securing the directory is essential to maintaining control of access to your information and operations. Directory-based security enables a high degree of granularity in managing user access to information and protects against the disclosure of confidential information to unauthorized users. Unlike some authentication systems, a directory service provides a hierarchical structure in which access permissions that are applied at a higher level can be inherited by directory objects, such as user accounts, that exist in organizational units (OUs) lower in the directory tree.

AD is a Lightweight Directory Access Protocol (LDAP)-compliant directory integrated with the Windows 2000 (Win2K) and Windows Server 2003 (WS2K3) OSs. AD leverages the Windows server security subsystem that provides authentication by validating the logons of users (and other security principals), and AD protects network resources by enforcing strict adherence to access control permissions assigned to resources. When access to a directory or network resource is requested, the user's security information is parsed against the security descriptors assigned to the resource; if a match is found between the user's security identifiers (SIDs) and the security descriptors, the user is provided with the level of access that has been granted.

AD further facilitates network security management by providing a pass-through authentication mechanism in which authentication via the directory service enables access to other enterprise resources. For example, logon authentication to an AD domain can provide authenticated access to Microsoft Exchange Server email and to all of the Microsoft SQL Server databases on the network.

AD also provides a policy-based implementation of the security constraints applied to computers or users. This implementation enables streamlined control of forest, domain, or OU-wide capabilities or restrictions.

Directory Security Protects Information and Service Assets

Securing your directory is critical in that the directory plays a central role in providing network security not only in the authentication of users but also in the operations of other network services and applications. Fundamentally, directory security is designed to protect against:

- Unauthorized access to the directory or network
- Disclosure of information
- Unauthorized modification of data
- Disruption of service

A directory service, such as AD, controls access to its objects and attributes by assigning security descriptors to each object or attribute, enabling administrators to provide differential access to each bit of information stored in the directory. The directory service also plays a central role in the management of identity information within your enterprise. The directory service is commonly responsible for storing identity information for enterprise network users as well as protecting this information by limiting access to authorized users and preventing unnecessary information disclosure.

AD is a distributed directory service that supplies centralized access to all of your information resources on the network. By using AD, you can search for and locate users, network devices, and shared information repositories.

Most directory services operate within a heterogeneous environment of platforms, services, and enterprise applications, so the need to support multiple security standards is common and some degree of pass-through authentication may be supported (particularly within a given platform). Increasingly, vendors of enterprise applications are writing their software to leverage the security mechanisms built-in to the directory service, thus increasing the ability to implement security for information and services from a centralized directory.

Why Directory Security Is Essential

Securing your network resources is important for a variety of reasons, not the least of which is the possible adverse impact on enterprise assets. The impact of a security breach could include disruption of service, destruction of enterprise information, and disclosure of sensitive information in ways that could damage investor or customer confidence—even a public perception of a lack of adequate security could negatively impact the confidence of your customers, investors, or other business partners. The protection of identity and directory information has also come under government regulation in recent years, which could lead to fines and even jail time for those responsible for any security breach. Therefore, maintaining effective security and promptly correcting any real—or perceived—vulnerability should be a constant priority.

Directory security mechanisms protect against unauthorized disclosure of information as well as modification or destruction of information. Directory security also protects the network services critical to your network operations against unauthorized tampering or disruption.

Securing AD is critical to your overall network and enterprise security, so you should give considerable thought to general AD security as well as the specifics of implementing AD security in your network environment. Not only do you have to determine how to initially implement effective security for your AD installation but you must also define procedures for maintaining that security on an ongoing basis. Additionally, you will need to plan for failure—when systems or services fail (and they will) or security breaches occur (and they may), which procedures do you have in place to contain the breach/failure and to recover from it?

Consider the administration of AD:

- Which services and/or data administrative tasks are assigned to whom and why?
- How are these administrative privileges and tasks delegated?
- How is the directory service monitored?
- How are security breaches detected and responded to?
- How are policy settings determined, deployed, and monitored?

These are the questions this guide will help you to answer.

Basic Security Mechanisms

Every directory uses the same basic security mechanisms to establish and maintain security—they all provide some means to authenticate validated users; control access to file shares, file systems, databases, network services, and other resources; and monitor the activities of user access, manipulation, and modification of these resources. Although the implementation might differ substantially between different directory services, the fundamental mechanisms include authentication, authorization, and auditing.

Authentication

Authentication is the process of determining that a user, computer, service, or application is actually the user, computer, service, or application that they claim to be. Authentication matches the submitted account name with the corresponding account name in the directory, and compares the logon credentials—account name and password, smart card, and so on—with the credentials stored in the directory for that account. If the submitted credentials match the corresponding credentials in the directory, the logon is authenticated and regulated access to the directory and network is granted. Authentication is not the same as validation, however. It is assumed that the identity is validated by external proofs when the account is created. Authentication then matches an asserted identity with one stored in the directory.

In AD, this authentication occurs during the logon process (see Figure 1.1). At a Windows workstation, the logon and authentication process begins when the user presses Ctrl+Alt+Delete, which invokes the logon screen (the graphic identification and authentication module). The user then inputs a username and password and selects a domain—these credentials are passed to the Winlogon service, which hands them off to the Local Security Authority. The LSA hashes the user's password and invokes the designated security support provider (by default, in Win2K and WS2K3, this provider is the Kerberos security support provider), which authenticates the submitted credentials with a domain controller. If the user credentials are authenticated with the domain controller, the user is allowed to access the directory and network.

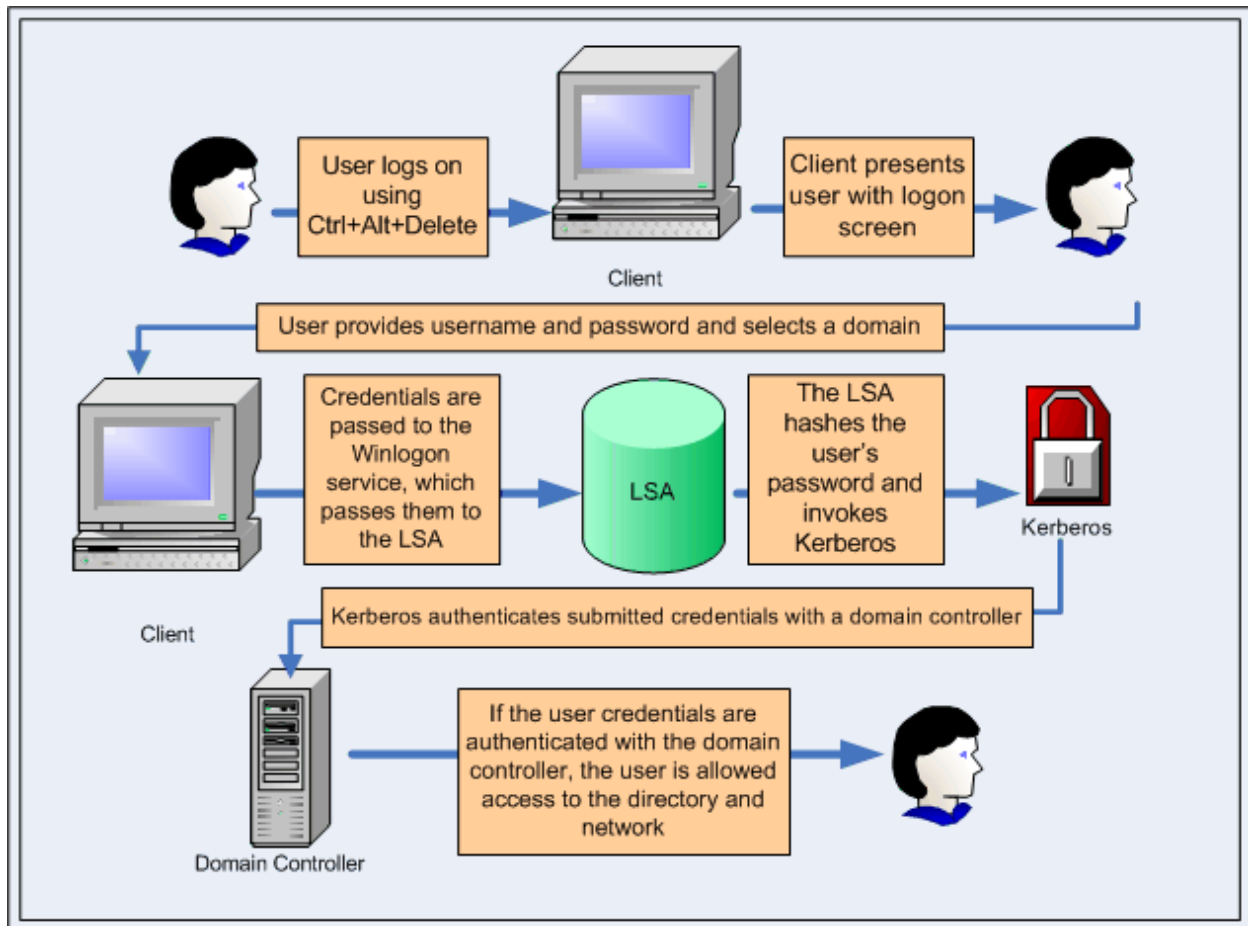


Figure 1.1: An overview of the AD authentication process.

Authorization

Authorization, also referred to as access control, is a bit more involved than authentication; authorization controls access to directory and network resources on a case-by-case, resource-specific basis. Upon logon of a security principal, the security subsystem works with AD to create an access token containing the SIDs of the user account and all security groups to which the user belongs. When a user attempts to access a protected resource, this access token is used to compare the information stored in the account's SIDs with the information stored in the security descriptor access control lists (ACLs) for each directory object or network resource the account attempts to access. Based on this comparison, the security subsystem allows or denies the specified types of access (read, modify, create, delete, take ownership, and so on) to the resource.


Auditing

Auditing is a mechanism that can be implemented by administrators to track user account logons, system events, and changes to directory objects and policy. Auditing is essential for maintaining security of your AD implementation and the security of access to your network.

In AD, auditing is implemented by first enabling the auditing capability within a Group Policy, then selecting the events to be audited for specific objects within the directory. To track changes to directory service access and objects, you must enable auditing within the scope of interest (domain, OU, and so on). To enable auditing on AD objects within a domain, you must first enable auditing within the default Domain Controller Policy Group Policy object (GPO—located in the Domain Controller OU). Then, to turn on auditing for any specific object within the directory, select the object within the directory, then access the object properties. Within the object's properties, select the Security tab, then select the Auditing tab, then specify the specific events—either success or failure—that you want to audit for that object. For the purpose of most auditing of directory-related events, you will usually want to select the Everyone group so that you can track changes made by any user to the objects in question.

Auditing in AD allows you to track both successful and failed attempts at access or modification of directory objects. Tracking all successful changes to directory objects will provide you with an audit trail for all actual alterations to the directory; whereas, tracking all failed attempts at changes can provide you with information indicating unauthorized or invalid attempts at access or directory modification. The following list highlights examples of events to audit:

- System events on domain controllers—particularly unexpected or unexplained reboots of the domain controller, which should always be investigated as they might indicate hacking attempts.
- Account management—auditing account management can provide you with an audit trail of changes to user or group accounts (when auditing success events) and can supply feedback indicating illicit access attempts (when auditing failure events).
- Excessive activity on directory partitions could indicate a Denial of Service (DoS) attack being launched by adding large amounts of data to the directory database, which could fill up the directory partition and render AD inoperative.

 However, there is a caveat to auditing within AD: When you enable auditing for a given domain, every time the audited event occurs, an entry is made in the Security log (which you can see with Event Viewer). If events occur frequently, the Security log can rapidly fill up. Depending upon your audit policy settings for how to respond to excessive security log entries, this situation could result in either the overwriting—thus loss—of auditing information or the shutdown of the domain controller. When you use auditing in AD, you will need to set the Security log policies to increase the size of the log so that log entries are not overwritten if the log is full. In addition, monitor the Security event log for excessive or inappropriate entries.

How AD Provides Security

AD implements access control to network resources by managing which security principals have access to each specific resource. In AD, security principals can be users, computers, groups, or services (via service accounts) and are validated by way of the authentication process (for users, authentication occurs at logon, for computers, it occurs at startup).

A SID is assigned to security principals at the point of creation—when the object is created in the directory. Each SID is comprised of a domain identifier (common to all security principals within the domain) and a unique relative identifier (RID). When a user logs on to the network, an access token is created that contains the user's SID, the SIDs for each group to which the user is a member, and the assigned user rights or privileges.

All resources within AD (objects and their properties), network folder and printer shares, and folders and files within the NTFS file system are protected by the assignment of security descriptors—access control entries (ACEs) contained within access control lists (ACLs)—that are associated with each object or resource. A security descriptor is comprised of two distinct ACLs assigned to each object or resource: the discretionary access control list (DACL) and the systems access control list (SACL). In brief, the DACL contains a list of the SIDs of all security principals that are either granted or denied access and the degree of access that is allowed (read, modify, full, and so on). The SACL contains a list of all the SIDs of the security principals whose access or manipulation of the object or resource needs to be audited, and the type of auditing that needs to be performed.

When a user attempts to access a directory object or network resource, the security subsystem checks to see whether the SIDs for the user (or security groups to which the user is a member) match the security descriptors assigned to the resource. If they match, the user is granted the degree of access to the resource that is specified in the ACL. Most commonly, users are assigned to security groups within AD, and the security groups are granted varying degrees of access to the network resources or AD objects. By assigning users to groups and applying security descriptors to objects and resources, groups of users can be granted or denied access to or control over entire classes of objects and sets of resources.

Policy-Based Security

One of the strengths of AD is its support for policy-based networking. Through the use of the Group Policy feature, security and usage policies can be established for both computer accounts and user accounts separately. These policies can be applied at multiple levels: a policy can be applied that affects the computers or users in a specific AD site, an entire AD domain, or only the users or computers residing in a specific OU.

Although this Group Policy capability provides a substantial degree of control over the network environment through the use of hundreds of different policy settings for computers or users, it can be a bit complicated to assess which specific cumulative set of policies are controlling the environment for a specific user or computer. In an improvement over Win2K, WS2K3 provides the ability to track and report the Resultant Set of Policy (RSOP), which is essentially the net effect of each of the overlapping policies on a specific user or computer within the domain.

Even more challenging is trying to monitor and track changes to the multiple and overlapping Group Policies implemented throughout the forest and domains. When you are managing AD in a distributed enterprise in which you have multiple administrators with the authority to implement and alter Group Policies, changes to Group Policies might occur without all administrators being aware of what has changed, when it changed, and the implications of the change for directory and network operations. For this reason, and others we'll discuss later, it is a good idea to limit the number of people who manage Group Policy.

Threats, Vulnerabilities, and Attacks

Protecting against attacks on your enterprise information or operations requires you to understand the nature of the types of vulnerabilities, threats, and attacks that might and to implement appropriate prevention, detection, and recovery strategies. In general, the degree of protection implemented should be related to the degree of value of the enterprise information or operations. For example, in most networks you probably wouldn't need to or want to implement fingerprint and retinal scanning to control access to the average user's workstation. You might, however, want to implement the use of smart cards to control access to critical domain controllers. In this section, we'll first explore what constitutes a threat, vulnerability, and attack, then examine some of the most common forms of attacks conducted against enterprise assets such as networks, directories, and the associated information.

Threat

In its most generic sense, a threat is someone or something that has the capability or potential to compromise the security of your directory, network, or information. In general, three factors are commonly required in order for a person to be a threat to the security of your directory: motive, method, and opportunity.



There are threats that do not have motive such as fire or flood; however, for threats that involve a person, motive is an applicable factor.

Another way to define the concept of a threat to your enterprise IT systems or information is as any action by a user, condition, or process that has the potential to disclose, damage, or disrupt your operations or information (see Figure 1.2). A user attempting unauthorized entry into your network, a fire that breaks out in the building that houses your network servers, and a virus that attempts to corrupt or delete needed information are all examples of viable threats to the security of your directory and your network.

Although threats to network security are commonly thought of as arising from external attackers exploiting some kind of vulnerability in your network or application software, it is not uncommon for both deliberate and inadvertent threats to the integrity of your network resources and operations to occur from people internal to your organization. According to some industry sources, internal threats are more prevalent than external ones.

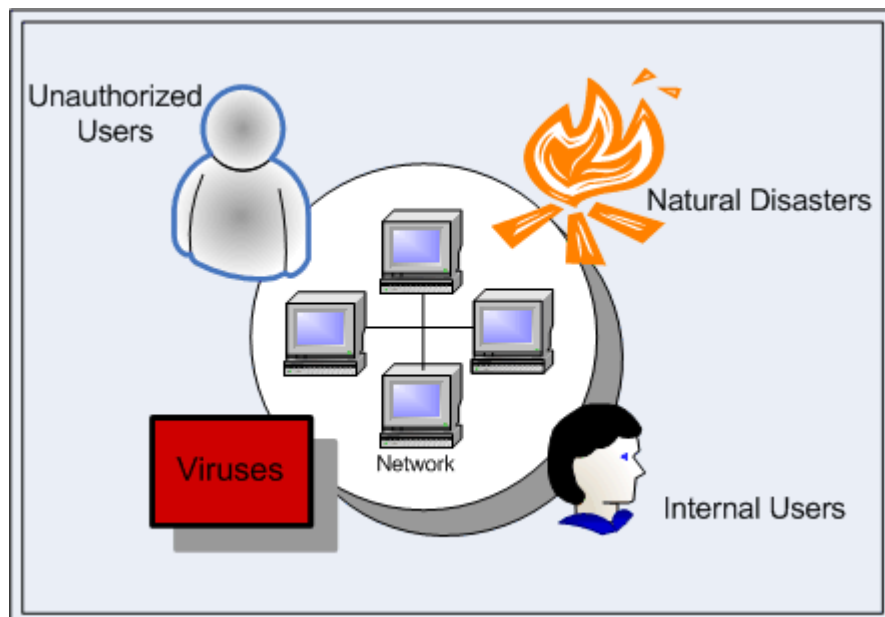


Figure 1.2: Threats to the security of your network.

Vulnerability

It seems that the IT industry magazines and even the mainstream press are constantly talking about new vulnerabilities discovered in software that is commonly used on workstations or services within your enterprise. Nevertheless, far fewer incidents are discussed by the press than exist, and it's rare for the general nature of such vulnerabilities to be discussed.

A vulnerability can be defined as any weakness in your security that provides an opportunity for an attack and that, by its utilization, can allow an attack to succeed. Vulnerabilities can occur in many different aspects of your network—software, hardware, social or physical environment—and you need to protect against all of them all of the time in order to ensure security. This requirement requires constant vigilance on many fronts—it sometimes seems as though there is a new weak spot revealed every day. It's not an easy task, but it is a critical one.

One of the most obvious areas of vulnerability is software, starting with the OS. If you are running Windows on your servers, you must ensure that each system has the latest service pack and patches, which requires you to monitor Microsoft's Web site for updates. To make this task a bit easier, you can subscribe to Microsoft's security updates newsletter and security update notification service (see the following resources for Web site information).

📄 Useful Microsoft Security URLs

Microsoft Security Update at <http://www.microsoft.com/technet/security/signup/default.mspx>

Microsoft Security Bulletins at http://www.microsoft.com/security/security_bulletins/

Microsoft Security Guidance Center at <http://www.microsoft.com/security/guidance/default.mspx>

Microsoft Security Anti-Virus Information at <http://www.microsoft.com/security/antivirus/>

Microsoft Security Newsgroups at
<http://www.microsoft.com/technet/community/newsgroups/security/default.mspx>

Patch Management, Security Updates, and Downloads at
<http://www.microsoft.com/technet/security/topics/patch/default.mspx>

In addition to the Microsoft resources, be sure to check with vendors of other software that is deployed on your network and independent security organizations such as the System Administration, Networking, and Security Institute (SANS) at <http://www.sans.org> and the Computer Emergency Response Team (CERT) at <http://www.cert.org>, both of which issue bulletins about security problems arising from many vendors' products.

When it comes to network services, the fewer the better—install the minimum number of network services required. You should install the services you need, of course, but make sure you disable any unnecessary services and, if at all possible, avoid installing unnecessary services on AD domain controllers. Every service has its own soft spots (vulnerabilities)—the fewer weaknesses you need to keep track of, the easier your job will be.

Don't ignore physical security; all the software-based security in the world won't help you if someone can walk into the server room and lay hands on the machine. Physical access to a server means that it is open to a variety of forms of attack such as:

- Rebooting the server, possibly with another OS on a floppy disk or CD-ROM
- Attaching devices that allow capturing keystrokes or copying data from the server
- Adding or removing system components such as hard drives or network devices
- Copying data to removable media
- Picking up the entire computer and walking out with it (don't laugh, this has happened)

Attack

An attack is any action by a user or software process that, if successful, results in the disruption, disclosure, or damage to enterprise information, services, or operations. Attacks, like threats, share the characteristics of motive, method, and opportunity, which assume the intent on the part of the attacker to deliberately be attempting to damage or steal information or disrupt operations. In a directory context, an attack is an action that uses or exploits the directory to gain access to or deny service from the directory or network resource.

There are many forms of attacks that can be carried out against your network, directory, and information; there are also many sources of attacks—both intentional and unintentional.


User-Based Attacks

The most common source of attacks are those initiated by people—whether by anonymous users attempting external penetration of the enterprise network or by an authenticated user working from inside the network. User-based attacks can either be physical attacks on the equipment supporting your directory or network or based on using the network or directory environment. Physical attacks can be as simple as stealing the physical computers (workstations, servers, and domain controllers), damaging the physical computers, and/or damaging the physical network infrastructure.

Network and directory-based attacks can come from anonymous users, authenticated users, or even administrators. Each of these sources has its own approaches to an attack and associated potential risks, which we'll briefly explore.

Anonymous Users

Anonymous user attacks commonly attempt to use vulnerabilities in the network, service, or application software. An attacking user might gain access via scanning tools or by exploiting a well-known but not patched error condition in operating software. When a known vulnerability is patched, the software update is generally accompanied by a description of the weakness, often providing all the information needed to hack an unpatched system.

 It is critical to stay on top of released patches and security updates.

In an AD environment, an anonymous user might be able to use LDAP to flood domain controllers with lookup queries, read domain information, identify user account security policies, find account names and SIDs, and identify shares on domain computers. Although some of these anonymous attacks can be mitigated by tightening security settings, thwarting anonymous DoS attacks requires monitoring of the domain controllers for unreasonably high levels of LDAP queries.

Some anonymous attacks are amazingly easy to carry out. Breaking into the typical user account requires only two pieces of information—username and password. Every Windows installation has a default account named Administrator, providing half the information needed to gain enormous power over the system. Similarly, many Windows computers have well-known hidden file shares (C\$, D\$, and so on) for administrative purposes. Disabling these file shares and renaming the Administrator account are a couple of easy steps to take that will help protect you against attackers guessing or hacking passwords as a means of gaining access to your network.

Authenticated Users

Authenticated user-based attacks use an authenticated account as the starting point in the attack. These attacks might be from spoofed-account access (via hacking/cracking tools), the illicit use of a valid account (obtained through some social engineering scheme), or a valid user who has decided to attack information, services, or operations for some personal or professional reason.

One of the problems with attacks by authenticated users is that the accounts have legitimate access to a range of resources and information on the enterprise network; thus, it is more difficult to detect when such attacks are taking place. Authenticated users, for example, can validly start processes that will have the effect of creating DoS conditions by consuming inordinate amounts of service resources (for example, a flood of LDAP queries or connections) or disk space (for example, storing many extremely large objects in the directory). Security for attacks by authenticated users requires a significant degree of monitoring, analysis, and responsiveness to anomalies occurring in the directory.

Authenticated users can also identify members of sensitive security groups, for example, determine sensitive account information (names, addresses, phone numbers, password, delegation status, and so on), discover linkage of Group Policies, identify sites, identify the OSs of domain controllers, and discover and disclose much additional information stored in the directory. The ability to read most objects in the directory is also contained in permissions assigned to all authenticated users by default; thus, the possibility of information disclosure by authenticated users is high.

Administrators

Illicit administrator attacks include conditions in which an Administrator account has been spoofed, the account has invalidly elevated privileges, or a trusted administrator has decided to attack the directory or network. Attacks using an account with administrative capability present some of the most serious threats to the directory, the network, and to the enterprise information accessible via the network.

Although they do not offer the range of capability of service administrators, accounts with limited delegated administrative rights can modify permissions on objects within their scope, enable accounts to be trusted for delegation, change passwords on other user accounts to be used for further (spoofing and repudiation) attacks, and change security settings causing DoS conditions.

Software-Based Attacks

The directory information structure is defined in the directory schema, which specifies the objects, attributes, and syntax permissible within the directory. Because the AD forest and domain directory structure is based on a correctly specified schema, any software application that corrupts the schema could render the entire directory—and your enterprise network—inoperative.

Likewise, automated attacks via viruses or worms that are not necessarily directed against your company can nevertheless have a damaging or disruptive effect. Email attachments present a huge risk and user education doesn't seem to stop people from opening every attachment that shows up in their inboxes. If such is true in your company, consider having your messaging system block, or at least scan, all attachments. Additional measures, such as turning off preview panes that automatically display messages, converting HTML mail to plain text, and blocking email clients from accessing the Internet can save you many headaches.

Environment-Based Attacks

In the physical environment that houses the domain controllers, any condition that has the effect of damaging or destroying the server hardware (fire, flood, tornado, hurricane, lightning, and so on) could also render the AD environment inoperative. These types of threats to IT operations are consistent across platforms and are usually well addressed by IT management in planning and implementing strict backup and restoration procedures.

Make sure that your disaster preparedness and recovery plans include provisions for offsite data backups, then make sure that the backups are actually taken offsite, and consider a secondary physical site that is ready to go in case the worst happens and your primary site is disrupted for an extended period of time.

Threat Analysis

To prevent attacks, you must first determine the nature and purpose of the attacks you need to protect against. Threats to the security of a directory service and the information it contains are varied, yet they can be usefully subdivided into several common categories. There are types of attacks that rely upon false authentication and subsequent access to the directory information (including spoofing, repudiation, and information disclosure). Some attacks are focused on preventing normative access to the directory information or service (for example, DoS attacks). Other types of attacks involve deleting or corrupting information in the directory, network, databases, or other information repositories. Still other attacks are based on changing access rights to allow an unauthorized user to gain access to or control over directory information (elevation of privilege).

To discuss common threats to information systems such as directory services, databases, and networks, the acronym STRIDE is used to summarize the Spoofing, Tampering, Repudiation, Information disclosure, DoS, and Elevation of privilege types of attacks.


Spoofing

Spoofing commonly refers to the type of attack in which the attacker is pretending to be someone or some process that otherwise has legitimate access to the directory. In a spoofing attack, the attacker obtains and uses the account and password of another user or service that has sufficient permission to access the directory information. The attack might impersonate an actual user or software process, leveraging the account information and security credentials to conduct unauthorized or malicious actions.

To defend against such spoofing attacks, implement policies to protect the username and password information for all user and service accounts that have access to the directory. Most network administrators have seen passwords on sticky notes prominently stuck to the side of a monitor, so don't forget that users must be educated about your security policies as well as the consequences of ignoring them.

Stringent authentication measures can help provide some degree of protection against spoofing attacks. Use of biometrics, for example, in addition to the use of usernames and passwords can help insure that the user is who they claim to be.

In AD, the spoofing of Administrator accounts is the most serious risk. If the spoofed user account is a service administrator (member of the Enterprise Admins, Domain Admins, or Schema Admins security groups), the attacker could damage or disrupt domain-wide, or possibly forest-wide, directory operations. Because these service administrators can modify the configuration of the entire AD environment and especially domain controllers, the compromising of these accounts to an attacker is particularly problematic. Directory schema management, replication, DNS service configuration, and domain addition and deletion are directly under the control of the service administrators. Service administrators also manage the installation and configuration of all software (including the OS), patches, and updates, and configure the settings on all network servers. As a result, if an attacker gains entry to your directory via a service administrator account, they can wreak unparalleled havoc throughout your AD environment—and mostly likely much of the rest of your enterprise network.

 We'll explore this topic in a bit more detail later in this chapter.

Although service administrator accounts are particularly sensitive, the compromising of data administrator accounts by spoofing can be just as disastrous. The data administrator accounts don't have the ability to change the directory configuration or operations; they are used to administer and modify user and group data contained within a portion of the directory and control the configuration of network file and printer shares. Spoofing of accounts that have been delegated administrative authority can allow the attacker to add or remove users and to modify user information. The last of these would allow the attacker to change a user's password and carry out additional attacks impersonating that user.

Even the spoofing of a domain user with nominal privileges can allow the attacker to access information stored on your enterprise network, potentially stealing, disclosing, or damaging important information. There are many horror stories of the chaos created by a single determined user; the same sorts of risks apply to a user account that is compromised by an external person.

Tampering

A tampering attack occurs when the information contained in the directory is changed, deleted, or corrupted by an unauthorized user in order to accomplish subterfuge, disrupt operations, or damage the directory information. A tampering attack might be conducted directly by an unauthorized user or indirectly by software constructed specifically to modify or damage the directory information (such as using a script that exploits a security flaw).

Keeping your security patches up to date in order to block inadvertent security holes in your applications, services, and OSs is a good starting point for protecting against such tampering attacks. In addition, lock down the directory service by using permissions that allow only necessary and authorized users to change directory information and access. Doing so will limit the window of vulnerability to unauthorized attacks.

Repudiation

Repudiation refers to the type of attacks that are designed to perform unauthorized operations wherein administrators of the attacked system are unable to prove who performed the attack. If changes to a directory or database are not being audited, for example, or the Security log is modified or deleted, any unauthorized change to the information the log contains wouldn't be traceable back to the source of the change.

To defend against a repudiation-based attack, both stringent authentication and auditing of the directory needs to be performed. Detailed event logging auditing access and changes to directory information can provide you with essential data to help track and stop such attacks. Consistent real-time off-server backups of the Security log need to be made in order to effectively track attempted repudiation-based attacks.

Information Disclosure

An information disclosure attack is designed to cause protected information to be exposed to one or more people who are not authorized to have access to that information. Information disclosure can take many forms; inappropriate access to documents in the file system; unauthorized access to databases containing sensitive user, financial, or medical information; and access to user accounts and other information stored in the directory. Information disclosure attacks can also occur when the information is on-the-wire—being transmitted across network connections. In such attacks, a network sniffer (or custom application with similar packet monitoring capabilities) is used to capture the information.

Use of network protocols that encrypt packets prior to transmission can protect against the latter type of attack, yet there are many ways that attackers can bypass standard security measures to inappropriately access information. Using permissions to control access to sensitive directory objects, file systems, and databases is a baseline necessity to defend against information disclosure attacks. Nevertheless, social engineering attacks—convincing an authorized user to unwittingly provide an attacker access—are a type of attack that technological mechanisms will not prevent. In addition to technical security mechanisms, security training needs to be provided to all people that have access to the directory and other sensitive data stores in order to prevent such information disclosure attacks from succeeding.

DoS

DoS attacks take many forms: as simple as remotely shutting down a server or as complex as an attack that hijacks many (tens, hundreds, thousands) of client systems and overloads a network service with bogus requests so that the network service cannot provide services to authorized users. In all of its variations, the purpose of a DoS attack is to render the network service unavailable to the users or systems that depend upon the service.

In the recent past, several high-profile DoS attacks have targeted the Web servers of well-known companies—attacks that have effectively prevented the normal operation and usage of their public Web sites. Within a company's internal network, DoS attacks can be substantially more problematic, potentially bringing all IT-dependent activities to a halt until the attack is neutralized. Real-time performance monitoring and automated alerts are a necessary starting point for defending against DoS attacks.

Elevation of Privilege

An elevation of privilege attack is one in which a user—authorized or not—has changed his or her access permissions to allow enhanced, or complete, control over directory, network, or file system settings. In this situation, an attacker can alter the permissions assigned to users, services, files, and directory objects. This attack could include turning off auditing, monitoring, or other tracking mechanisms, which would effectively allow subsequent attacks to be untraceable.

Although auditing can alert you to changes in privilege elevation, there is commonly significant delay between the action and the awareness by IT management that such a change has occurred. Even after a change has been discovered, staff must determine whether each specific privilege elevation was authorized or not. Use of intrusion detection systems (IDSs) and intrusion prevention systems (IPSs) can greatly enhance the responsiveness of the network security team in identifying and preventing such attacks.


Managing the Directory Service

AD administration can be divided into two significant areas:

- Service administration
- Data administration

Both areas are critical and entail security risks and vulnerabilities; however, service administration requires a higher degree of access to the directory and you must take great care in determining who will do these tasks and how they will carry out the tasks.

Service administration involves managing AD operations, including replication, schema changes, domain creation and removal, and the delegation of tasks to data administrators. Because part of the job of a service administrator involves the installation and configuration of software—including service packs and other software updates on domain controllers—service administrators will need physical access to domain controllers.

 Although Microsoft often describes a domain as a security boundary, this definition is only partially accurate. The domain boundary *does* act as a block to the inheritance of security policy data (such as password and account policies), but it does *not* protect against attacks by service administrators. Thus, you *must* trust your service administrators across the entire forest, even if you do not intend them to administer outside of a designated domain.

You *must* trust the people who will be service administrators absolutely—you will still want to implement auditing and other security controls, of course, but a rogue service administrator can wreak havoc on your directory, your network, and therefore your business. Only employees who have demonstrated that they are responsible and understand both AD and your business operations should be entrusted with these tasks. This job should not be outsourced, given to temporary staff, or, in most cases, delegated to a brand-new hire. For similar reasons, don't add user accounts from another forest to service accounts in your domain; security lapses in the other forest—which you cannot control—can easily compromise the security of your forest.

AD uses several built-in service administrator accounts, such as Enterprise Admins, Domain Admins, Schema Admins, and so on. Some groups, such as Schema Admins, should not have permanent members but should instead have membership granted only when those tasks need to be performed. The built-in Admin groups allow more access than you will necessarily want to provide, even to the people who need to do those tasks. In this case, create a custom group that provides only the specific access that is needed rather than using a default group.


Make sure that you define clear administrative policies and ethical standards, establish consequences for breaches, then educate administrators about these policies. Obtaining signed copies of critical policies from each administrator provides a paper trail so that people can't say, "I didn't know..." Senior administrators should set an example by adhering to both the spirit and the letter of these policies—it doesn't make much sense to expect compliance with policies that managers are ignoring.

Best Practices for Service Administrator Account Management

There are several best practices to ensure that service administrator accounts are not misused in any way. The following highlights these considerations:

- Limit the use of service administrator accounts to actual service administration tasks. Although it might be easier to simply log on with an account that provides all the privileges you might need, such an account introduces security risks. If a service administrator is also a data administrator, ensure that there is one account for each function, plus a standard user account for normal work logon.
- Restricting the computers that the service administrator accounts can log on to will provide further security by insuring that these accounts are only used on a limited number of workstations (perhaps in a secure environment) and not from just any computer. These workstations should have all the usual protective software such as antivirus, anti-spyware, and so on because a virus running on such a machine will be running with a high level of privileges and could cause much damage.


- Consider requiring strong authentication of some sort using a token or biometric identifier or even split credentials for service account logon. The use of split credentials means that two people are required to log on to a single service account. This process might involve giving the physical token to one person and the password to another or creating a complex password and telling each person half of it. Thus, access to the account requires that two people be present, further reducing the possibility of attacks because it is unlikely that someone will attempt an attack with someone literally looking over their shoulder.
- Create only the service administrator accounts that are actually needed. Every person with service administrator access is a potential security risk—even if the person associated with an account is 100 percent trustworthy, each account that exists is another account that could be hacked.
- Keep service administrator accounts restricted from unnecessary access and possible exposure by not providing Internet access, email accounts, and so on. Doing so can help limit the possible exposure of these accounts to hacking attempts and serves as a constant reminder to the administrator that he or she should be using a regular user account for non-administrative work.
- The built-in Administrator account is an obvious target, so make sure that you rename it. When you rename the account, make sure that all the information in the account is altered to resemble a standard user account—it won't do you much good to rename the account if the text in the description still identifies it as the Administrator account. You should also create an account that looks like the standard Administrator account but has no privileges to act as a decoy.
- Create an AD subtree that will contain all service administrator user and group accounts and the workstations from which they can perform service administration tasks. You should not allow a data administrator—or even numerous service administrators—access to service account management; only allow trusted service administrators to manage these accounts. Doing so will help protect against data administrators elevating their privileges. This subtree should be fully audited and audit logs should be checked regularly.
- Don't forget DNS! AD relies on DNS and you should monitor its correct operation as part of good directory management. When you design your internal AD namespace, use a namespace that is different from any of your public DNS names. AD-integrated DNS stores the data in the directory, which is more secure than standard zone files and supports secure DNS updates. Every authenticated user has the ability to create DNS resource records. WS2K3 supports quotas to ensure that a rogue user can't flood the DNS service with spurious or malicious DNS records as part of a DoS attack.

 Microsoft has several whitepapers available that describe specific steps to take when securing AD (<http://www.microsoft.com>). Consult these resources when you are designing your directory and administrative practices.

Best Practices for Managing Directory Information

Data administrators manage the contents of the directory, user accounts, groups, network resources such as computers, and so on. Thus, you need to take the precautions when determining data management practices. Depending on the information that is being stored in AD, something as simple as the disclosure of user data could present serious risks. If, for example, the Human Resources department has decided to store confidential information such as pay grades and social security numbers in the directory, that information must be protected for reasons that range from internal policy (many companies forbid employees disclosing their compensation) to legal (if your staff's personal and financial data ends up on the Internet, you're in huge trouble).

Most of the people who will be administering your directory will be data administrators, and most of them will have access only to a limited subset of directory data. Some things are obvious—data administrators do not control delivery of the directory *service*, they don't need access to domain controllers, and, in general, they can do their work from any available workstation. In most organizations, the privileges afforded to a data administrator are restricted to a portion of the directory—a single OU representing a workgroup, for example, or only a printer and related services.

 One of the major advantages to AD when compared with pre-Win2K versions of Windows Server products is this ability to delegate control of portions of the directory. It's a very simple process in practice; you create a group, add members, and provide the required privileges. This functionality is almost deceptively easy as it hides the risks that are inherent in delegating management of a portion of the directory.

Some types of data administration are fairly straightforward and easy to delegate securely. Printer management, for example, requires a limited set of rights to directory information and isn't likely to provide many security holes. GPO management, however, provides many opportunities for security breaches, either inadvertent or deliberate; even though GPO management might be considered data management, you might want to restrict this task to a few trusted service administrators. There are default groups available for a number of administrative tasks each with predefined sets of privileges. You can also create custom groups to accommodate the specific needs of your organization. When assigning data administration tasks, you will have to find a balance between enough access to do the job and too much access, which puts directory information or network security at risk. The following list highlights best practices for data administration:

- Limit data administrator's scope of access to the minimum required to accomplish the assigned tasks. AD allows highly granular assignment of security, so use it. It's better to create a group that has too few rights and go back and add an additional privilege later than to grant too many privileges initially and only discover after the fact that security has been compromised.
- A single person should be delegated management for each group, if possible, to avoid group membership conflicts that could occur if several people are changing members of the same group.
- The application of Group Policy has wide-ranging implications as it controls the application of security for the network. Accordingly, GPO management should be restricted to service administrators (not data administrators) in most cases.

- Watch for abuse of Creator/Owner status. When a directory object is created, the user creating that object owns that object as well as any objects created underneath it. Thus, a data administrator who has the right to create OUs can create a subtree and then block access to it. Although this problem is not permanent—members of the Administrators group can take back ownership—it is something to keep an eye on if for no other reason than it might indicate a person whose activities you should closely watch.
- Monitor user reports (and Security logs in Event Viewer) of odd behavior. An occasional occurrence might be nothing to worry about, it could be an indication of internal tampering—such as indicating that an administrator is resetting passwords to allow them to log on as another user. By doing so, the administrator can impersonate that user and carry out activities such as reading or even sending email from the user’s account.

User Management in AD

A great deal of the administration of AD is focused around user and group management, generally referred to collectively as user management. User accounts provide the means of identifying and authenticating individuals, while the rights that are granted to each user are normally controlled by group membership. Even these simple user management tasks, however, have the potential to compromise the security of your directory if performed incorrectly.

If you’ve been managing Windows networks since before Win2K, the changes to the User account with AD will seem significant. User accounts in Windows NT had less than a dozen settings to configure; however, an AD user account has more than 250 possible attributes to deal with! Everything from extensive work-related information, such as a user’s manager, to personal data, such as home address and phone number, can be securely stored in the directory.

All this configurability comes with a price: the complexity of user management has grown exponentially. Luckily, only about half a dozen attributes are mandatory and many of them are not even displayed in the standard AD management tool. In fact, quite a few are entirely hidden from the usual management tools and require special utilities to view and configure (although they can generally be included as part of a search in the standard UI).


To simplify the process of securing this information, there are default security settings applied to sets of user properties, which can be either good or bad. Using predefined property sets streamlines setting security; however, allowing access to the home phone number of a user, for example, also exposes the user’s home address. You should study the details of user properties and consider the related security implications carefully rather than assuming that you understand how it all works.

Creating a User Object

There are three types of objects employed for user accounts in AD:

- **User**—This account is the standard user account and will be used most in corporate networks. A User object is a security principal and a member of the domain in which it is created.
- **InetOrgPerson**—The InetOrgPerson object was added in WS2K3 to facilitate the use of AD as an LDAP directory that complies with Request for Comments (RFC) 2789. The InetOrgPerson is a security principal and a domain member and functions in the same manner as a User.
- **Contact**—A contact object is used to represent a person who needs an email account but no other access to the network. It is not a security principal, so it cannot be used for a person who needs access to network resources. You might use a contact to allow the inclusion of a person in the directory for searching purposes—as a phonebook entry of sorts—or to include members of other forests in the global catalog. Exchange Server 2000 and later use contact objects for custom recipients, which are external email addresses that are included in the local address book.

To create a User, select the container in which you want the object to be created, then choose the option to create a new User, InetOrgPerson, or Contact object. Configure the various forms of name for the new user, making sure it is not a duplicate of another name in the domain and set other mandatory and desired properties such as password.

 A clear naming standard is useful to prevent the existence of duplicate names.

Set the values for other properties that need to be configured. Doing so might require navigating through several property sheets to find all of the needed properties as well as going into the special properties that are available from the Advanced button on the Security tab. Next, add the new user account to the appropriate groups.

You can perform this add-user process manually or by using a bulk import method such as Ldifde or Csvde. Creating user accounts manually is quick and easy if you only have one or a few to add. If you are populating a new directory, however, or have many new users to add every week, automating the process is useful as it can lighten your workload and avoid many opportunities for error. Automated creation methods can also allow you to set hidden properties without the need for an additional tool.


Creating a Group

Before creating a group, you must make two decisions to determine which type of group you need. The two areas that you must consider are the type of group and its scope. The type of group determines what you can do with the group. There are two types of groups used in AD:

- Security groups provide a means to apply permissions and are the primary type of group used in AD.
- Distribution groups are used only by messaging systems that are integrated with AD, such as Exchange Server 2000. These groups provide no means of applying security permissions, but rather function solely as a collection of user accounts for email and other messaging products.

The group scope determines where it can be used and which object types can be a group member:

- Domain local—Applies to a single domain, can contain users as well as global and universal groups
- Global—Used within an entire AD forest, can contain users from the same domain and global groups
- Universal—Can be used throughout an entire forest and can contain users, global groups, and other universal groups; for compatibility reasons, Universal groups are only enabled in native-mode domains

 We'll explore Universal groups in more detail in later chapters.

To create a group, select the container in which you want to create the group, then choose the option to create a new group object. Next, name the group and select the type and scope of the new group. Finally, add user and group accounts to the new group.

Summary

In this chapter, we introduced you to how directory security works by looking at the big picture. After exploring some general security concepts, we took a look at how AD manages those aspects of security. Possible threats were discussed, and we talked a bit about best practices for managing both the directory service and directory data.

With this background on directory security and some ideas about how to approach secure management, we'll move on to the details of doing so. In the next chapter, we'll look at the big picture of directory administration starting with how the design of your AD tree impacts your management processes (hint: it's more than you might think), then explore the details of delegating administration. After that, we'll move on to an in-depth look at Group Policies and finish up with a chapter on delegating administration.