

**Tricks[™]
& Traps**
eBook Series

WindowsITPro  **Books**

Essential **Security** Tips

By Randy Franklin Smith
and John Savill

Microsoft[®]



Microsoft®

Contents

Chapter 1 Passwords and Permissions	1
Tip 1: Safeguarding FTP Files	1
Tip 2: Using Passwords with Kerberos	3
Tip 3: Forcing All Users to Change Their Password At Next Logon	4
Tip 4: Understanding Event ID 560	5
Tip 5: Solving Password Problems That Involve Your PDC	6
Tip 6: Resetting the Directory Service Restore Mode Administrator Password	6
Tip 7: Disabling the Recovery Console Administrative Password	7
Tip 8: Specifying Spooler Permissions on Just One DC	7
Tip 9: Resolving a Windows XP-Related Password Error	8
Tip 10: Changing Passwords Remotely Via the Web	9
Tip 11: Configuring Pre-Staged RIS Permissions	10
Tip 12: Changing a Domain User's Password from the Command Line	11
Tip 13: Comparing Code Access Security with User Access Permissions	11
Tip 14: Enabling Users to Access Two Domain Accounts	12
Chapter 2 Event Log and Auditing	14
Tip 15: Enabling Debug Logging for IP Security	14
Tip 16: Using Log Parser to Audit Domain Logons	14
Tip 17: Monitoring For Unauthorized Scheduled Tasks	16
Tip 18: Establishing DHCP Server Log Thresholds	17
Tip 19: Operation-Based Auditing	18
Tip 20: Auditing Account Logon Events Centrally	21
Tip 21: Deciphering Security Event ID 529	22
Tip 22: Monitoring Security with Custom MMC Consoles	22
Tip 23: Audit Control List Editing Rights for a Win2K Object	25
Tip 24: Creating Multiple Event Viewer Views	26
Tip 25: Viewing Security Logs for All DCs	28
Chapter 3 Security Policy	29
Tip 26: Using Windows Update with Security Policies	29
Tip 27: Using One GPO to Control Both Windows XP and Windows 2000 Settings ..	30
Tip 28: Preventing Users from Disabling Group Policy	30
Tip 29: Connecting to a DC to Edit a GPO	31
Tip 30: Editing an IP Security Policy	32
Tip 31: Understanding Group Policy's Block Policy Inheritance and No Override Options	33

Chapter 4 IP Security	35
Tip 32: Defining IP Security	35
Tip 33: Stopping and Restarting the IP Security Policy Agent	35
Tip 34: Defining an IP Security Policy for a Group Policy Object	36
Tip 35: Changing the Authentication Method Used for IP Security	37
Tip 36: Enabling IP Security	39
Tip 37: Managing and Creating IP Security Policies	41
Tip 38: Enabling IP Security Traffic through a Firewall	42
Tip 39: Defining the IP Security/Layer Two Tunneling Protocol NAT-T Update	43
Tip 40: Disabling IP Security on a VPN Connection that Uses Layer Two Tunneling Protocol	43
Tip 41: Preventing Attackers from Bypassing IP Security Packet Filtering	44
Chapter 5 Kerberos	45
Tip 42: Defining Kerberos	45
Tip 43: Distributing a Shared Key	46
Tip 44: Distributing a Long-Term Key	47
Tip 45: Defining a Kerberos Trust	47
Tip 46: Creating a Kerberos-Based Trust Between Domains	49
Tip 47: Changing the Ticket Lifetime Used by Kerberos	49
Tip 48: Cracking Kerberos Packets	50
Tip 49: Windows NT LAN Manager Versus Kerberos Use	51
Tip 50: Exploring Kerberos Ticket Lifetime	51

Authors

John Savill (john@savilltech.com) is chief Microsoft architect for Geniant, a Dallas-based Microsoft Gold Certified Partner. He is an MCSE on Windows Server 2003 and a five-time MVP. He is the author of *The Windows XP/2000 Answer Book* (Addison-Wesley Professional).

Randy Franklin Smith (randy@winsecanswers.com) is a contributing editor for *Windows IT Pro*, an information security consultant, and CEO of Monterey Technology Group. He teaches Monterey Technology Group's Ultimate Windows Security course and is an SSCP.

Chapter 1

Passwords and Permissions

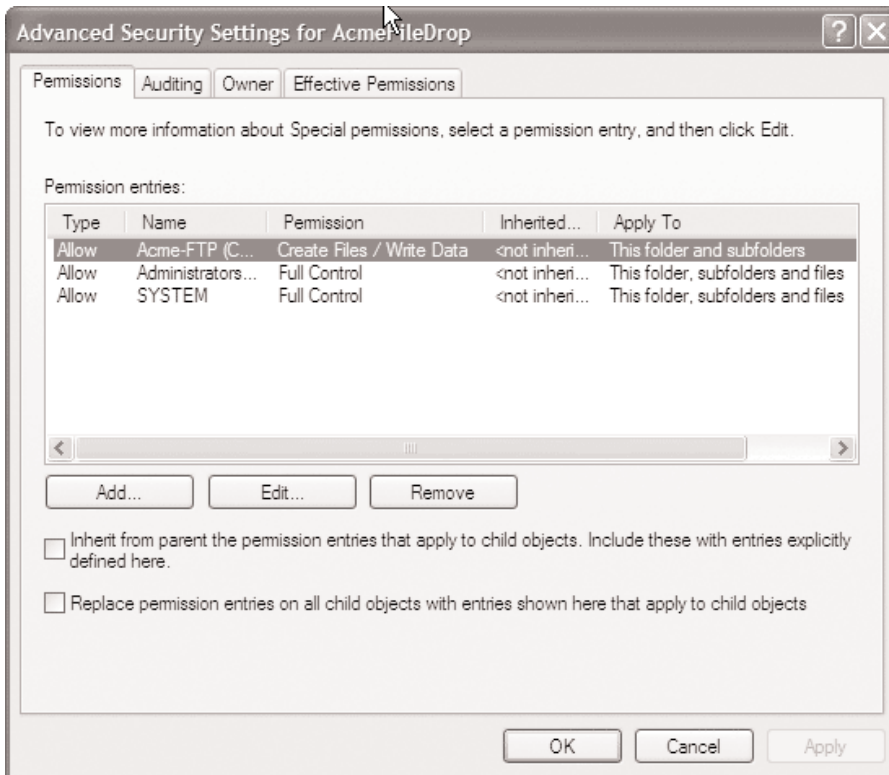
Tip 1: Safeguarding FTP Files

Q We have a business partner who regularly needs to send a file to our server over the Internet. Our only option is FTP, and we can't use VPNs, IP Security (IPSec), or FTP over Secure Sockets Layer (SSL). We've thought about encrypting the file, but we also realize FTP authentication is weak because the password is sent in clear text. We don't want an attacker who manages to capture our password to be able to wait for the transfer to complete, then log on and download or delete the file. Can we set up the transfer so that attackers between us and our partner firm can't benefit should they capture our FTP password?

A Yours is an interesting challenge, but you can meet it by using an encryption utility and implementing proper user permissions. First, create a user account for the business partner—for demonstration purposes, I'll call the account Acme-FTP. To properly limit the new account, remove Acme-FTP from the Users group to which Windows automatically adds all new accounts. Create a folder in your FTP server's root folder called AcmeFileDrop. Open the folder's Properties page, click the Security tab, then click Advanced and clear the check box that lets the folder inherit permissions from the parent folder. When Windows asks whether to copy or remove the permissions, select Remove. Then, add the Administrators and SYSTEM groups to the folder and give them Full Control. Add any other groups that need to be able to access the files that Acme delivers, and grant those groups Modify or Read access, as appropriate.

Next, add the Acme-FTP account and give it the specialized set of permissions that Figure 1-1 shows.

Figure 1-1
Securing a user account for FTP



As you can see, I grant Acme-FTP only the ability to create new files in the AcmeFileDrop folder. Finally, arrange to have Acme encrypt the file before sending it; Acme can use any of the several available encryption utilities that use shared keys or public/private key pairs.

Now, when Acme sends an encrypted file to your system, your security risk is much lower. An attacker who manages to capture Acme-FTP's clear-text username and password and use them to access your FTP server will be successful only if he or she can also spoof Acme's IP address. If the attacker manages to log on to your FTP server, he or she will be able to create new files but won't be able to download or change existing files.

The two risks that remain are that the attacker could create spurious files or fill up your system volume. Always validate new files (e.g., confirm they were encrypted with the proper key) before processing them. If you place AcmeFileDrop on a separate volume or if your FTP server is on a version of Windows that supports user disk quotas, you can prevent an attacker from filling up the system volume.

The only way an attacker can get an existing file is to reconstruct it from captured FTP packets, then work on cracking the encryption key. In case an attacker modifies the file in transit between you and your business partner, be sure you use a well-written encryption utility that uses a standard encryption protocol and detects modifications to the file after encryption. Never choose encryption software that uses proprietary encryption. One sign of a good utility is that the author explains the encryption methods used; a better sign is that he or she has made the utility's source code available for review by an encryption specialist. If the utility uses shared key encryption, make sure you select long, hard-to-guess keys, and don't exchange keys with Acme over the Internet (e.g., do it by phone). Many such utilities are available on the Internet, but I prefer WinZip 9.0 because it supports 256-bit Advanced Encryption Standard (AES) encryption and provides compression as well.

—Randy Franklin Smith

Tip 2: Using Passwords with Kerberos

Q In upgrading the last of our pre-Windows 2000 computers for security reasons, we want to address the cracking of user passwords by possible eavesdroppers on the network who could sniff and crack Windows NT LAN Manager (NTLM) authentication packets. In a pure Win2K network that uses Active Directory (AD) domain accounts, Kerberos replaces NTLM to eliminate the risks associated with NTLM authentication. However, someone claimed that Kerberos is also vulnerable to sniffing and subsequent cracking. Is that true? If it is, how can we avoid the problem?

A Any protocol can be sniffed. Kerberos's overall design and use of encryption and hashing technology makes it less vulnerable than NTLM to sniffing. However, Kerberos ultimately bases its ticket encryption on the security principal's key (i.e., the user's password), so weak passwords expose Kerberos to cracking.

Kerberos-cracking software is readily available on the Internet. Arne Vidstrom's KerbCrack, for example, uses a word list and brute force to provide sniffing and cracking functionality. KerbCrack can process a word list in a few seconds and perform a brute-force attack that uses a restricted character set in a matter of hours. As long as we use passwords, we'll need to keep them complex and avoid the use of words or other simple patterns. A fully switched network reduces the risk of someone capturing Kerberos credentials from a network drop, but switches can be tricked into rerouting traffic, and switches don't prevent network administrators from sniffing Kerberos or any other traffic from the switch itself.

If your company is willing to deploy smart cards, you can eliminate passwords from your AD domain. When a user authenticates through a smart card, Win2K automatically switches to PKINIT mode. PKINIT is a Kerberos extension that bases initial authentication on the certificate for the user whose private key is stored on the smart card. It protects all Kerberos exchanges with at least 128 bits of entropy and effectively eliminates cracking risks from today's technology.

If smart cards aren't an option, you'll need to rely on written and configured policies. Require a password at least seven characters long, and require complex passwords (e.g., passwords that include characters from at least three of four character sets—a-z, A-Z, 0-9, symbols). You can

configure both these options in the Default Domain Policy Group Policy Object's (GPO's) Computer Configuration\Windows Settings\Security Settings\Account Policies>Password Policy folder. If you use this configuration, an attacker who has a 1.5GHz Pentium processor would need as much as a year to perform a brute-force attack on every possible character set. If you bump the password to eight characters that come from the a-z, A-Z, and 0-9 character sets, an attacker with 1 processor could spend as many as 67 years cracking the password; someone with 100 processors at his or her disposal 24 x 7 could spend as long as 8 months.

Lockout policy provides no protection for offline cracking attacks, but having a good password policy and requiring password changes every few months helps you defend against cracking attempts. For example, if you require passwords to consist of at least seven characters drawn from the a-z, A-Z, and 0-9 character sets and require users to change their passwords every 60 days, the passwords would change before the attacker had worked through a quarter of the problem set. In conjunction with your domain's password policy, get management to back a written password policy that addresses the need for hard-to-guess passwords. Until you can upgrade everyone's computer to Win2K or later, you might want to implement NTLMv2—a "bandage" for NTLM that strengthens network authentication and defeats the current version of @stake's L0phtCrack.

—Randy Franklin Smith

Tip 3: Forcing All Users to Change Their Password At Next Logon

Q We recently enabled a maximum password age and want to put it into effect for all our users. The Microsoft Management Console (MMC) Active Directory Users and Computers snap-in doesn't seem to have a multiple select option for changing user accounts. What's the easiest way to select the *User must change password at next logon* check box for many users at once in Active Directory (AD)?

A The easiest way to do this task is to use the Addusers utility to produce a text file of usernames, then use the For command to execute a Net User command for each user and select the *User must change password at next logon* check box. First, run the command

```
addusers /d users.txt
```

which produces a users.txt file that contains a list of all users, global groups, and local groups in AD. Open users.txt, locate the [Global] line, and delete that line and everything after it to get rid of all the groups listed in the file. Save and close users.txt.

Next, run

```
for /f
  "skip=1 tokens=1 delims==, "
  %i in (junk.txt) do cusrmgr -u
  %i +s MustChangePassword
```

The For command skips the first line of the file (i.e., [User]), then inserts the username (i.e., the first string from each line in the file) in place of %i in the Cusrmgr command. The Cusrmgr portion of the code then equates to

```
cusrmgr -u <username> +s
MustChangePassword
```

This Cusrmgr command selects the *User must change password at next logon* check box for the username that appears in the username variable.

—Randy Franklin Smith

Tip 4: Understanding Event ID 560

Q Our Event Viewer shows occasional instances of event ID 560 (Object Open) from user Everyone on a PDC, as Figure 1-2 shows. Some of our administrators are concerned that this event comes from the Everyone group. I'd appreciate your thoughts.

Figure 1-2
Event ID 560

3/24/2003 (3)	6:55:58 AM 560	Security Everyone	Success Audit HOSTNAMEPDC	"Object Open:"
Object Server:	Security Account Manager			
Object Type:	SAM_USER			
Object Name:	DOMAINS\Account\Users\00001195			
New Handle ID:	73321376			
Operation ID:	{0,1371542786}			
Process ID:	2162221088			
Primary User Name:	SYSTEM			
Primary Domain:	NT AUTHORITY			
Primary Logon ID:	(0x0,0x3E7)			
Client User Name:	SYSTEM			
Client Domain:	NT AUTHORITY			
Client Logon ID:	(0x0,0x3E7)			
Accesses	ChangePassword (with knowledge of old password)			
Privileges	-			

A Windows logs event ID 560 when you enable system-level file and object auditing without enabling object-level auditing. Different versions of the OS log variations of this event, which simply indicates that a user is trying to change his or her password. Don't mistake this event for a password-reset attempt—password resets are different from password changes. Only someone who already knows the account's password can change the password. Your events might not be indicating the username because the password is expired and the user is trying to change it at logon time.

The best way to track password changes is to use account-management auditing. Make sure you enable the Audit account management security setting for success and failure on your domain controllers (DCs). Then, check your Security log for event ID 627 (Change Password Attempt), which provides better information about password changes.

—Randy Franklin Smith

Tip 5: Solving Password Problems That Involve Your PDC

Q I manage the Help desk for a company that has several remote locations. Several users at one location have received the error message *Unable to change password on this account. Code: c00000be* when their passwords have expired and they've tried to enter a new password. I'm running Windows 2000 Server and Windows NT 4.0 Service Pack 6a (SP6a) on the workstations. I suspect that the server might contribute to the problem. My current workaround is to access the BDC and use User Manager to enter the new password. But how can I address the basic problem?

A Executing a password reset from an NT workstation on a user account in an NT domain doesn't involve member servers such as the Win2K server at your remote location. Code c00000be indicates that your workstation can't reach the PDC. Check the filters on your routers between the remote location and the PDC. Try to ping your PDC from a workstation in the remote location to make sure your workstations can reach the PDC.

—Randy Franklin Smith

Tip 6: Resetting the Directory Service Restore Mode Administrator Password

Q How can I reset the Directory Service Restore Mode Administrator password?

A In Windows 2000 Server, you used to have to boot the computer whose password you wanted to change in Directory Restore mode, then use either the Microsoft Management Console (MMC) Local User and Groups snap-in or the command

```
net user administrator *
```

to change the Administrator password. Win2K Server Service Pack 2 (SP2) introduced the Setpwd utility, which lets you reset the Directory Service Restore Mode password without having to reboot the computer. (Microsoft refreshed Setpwd in SP4 to improve the utility's scripting options.)

In Windows Server 2003, you use the Ntdsutil utility to modify the Directory Service Restore Mode Administrator password. To do so, follow these steps:

1. Start Ntdsutil (click Start, Run; enter cmd.exe; then enter ntdsutil.exe).
2. Start the Directory Service Restore Mode Administrator password-reset utility by entering the argument "set dsrm password" at the ntdsutil prompt:

```
ntdsutil: set dsrm password
```

3. Run the Reset Password command, passing the name of the server on which to change the password, or use the null argument to specify the local machine. For example, to reset the password on server thanos, enter the following argument at the Reset DSRM Administrator Password prompt:

```
Reset DSRM Administrator Password: reset password on server thanos
```

- To reset the password on the local machine, specify null as the server name:

```
Reset DSRM Administrator Password: reset password on server null
```

You'll be prompted twice to enter the new password. You'll see the following messages:

```
Please type password for DS Restore Mode Administrator Account:
Please confirm new password:
Password has been set successfully.
```

- Exit the password-reset utility by typing “quit” at the following prompts:

```
Reset DSRM Administrator Password: quit
ntdsutil: quit
```

—John Savill

Tip 7: Disabling the Recovery Console Administrative Password

Q How can I configure the recovery console in Windows 2000 and later to not require me to enter the administrator password?

A To configure the recovery console to not require you to enter the administrator password, perform the following steps:

- Start a registry editor (e.g., regedit.exe).
- Navigate to the HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Setup\RecoveryConsole registry subkey.
- Double-click SecurityLevel, set its value to 1 to not require password entry (or 0 to require the user to enter the password), then click OK.
- Close the registry editor.

You can also use the Microsoft Management Console (MMC) Local Security Settings snap-in (go to Local Policies, Security Options, “Recovery console: Allow automatic administrative logon”) to configure this setting.

—John Savill

Tip 8: Specifying Spooler Permissions on Just One DC

Q I want to specify permissions for the print spooler service on one of my two domain controllers (DCs). I don't want to modify the Default Domain Controller Policy because my other DC doesn't run the print spooler service. Instead, I'd like to manually specify permissions for the print spooler service on just one DC. Is that possible?

A There are two ways to solve your problem. You can create either a Group Policy Object (GPO) or a security template and apply it to just the DC that runs the print spooler. I'll describe the latter approach first.

To start, open the Microsoft Management Console (MMC) Security Templates snap-in and create a new security template named `PrintSpooler.inf`. Select the System Services folder in your new template and find the Print Spooler service. (You must edit the template on a computer on which the service is installed.) Double-click Print Spooler and select the *Define this policy setting in the template* check box, then specify the startup mode—for Print Spooler, you'll probably want to select the Automatic startup mode so that print services are always available. Next, click Edit Security to view the permissions for the Print Spooler service and adjust them so that the appropriate users can start and stop the service and perform other necessary operations. Click OK twice to close the dialog boxes. Right-click the `PrintSpooler` template in the treeview pane of the Security Templates snap-in and select Save. Make sure the `PrintSpooler.inf` file is in a folder that's accessible from your DC.

Now, log on to your DC and open the MMC Security Configuration and Analysis snap-in. The first step in applying this template is kind of quirky—you must create a new security database, then import the template. (The database lets you import multiple templates in sequence, then apply a composite of their settings.) Right-click the Security Configuration and Analysis snap-in in the treeview pane and select Open. Enter a name for the new database, such as `PrintSpooler`, and click Open. You'll be prompted for a security template to import. Specify the `PrintSpooler.inf` file you created earlier and click Open. After the snap-in imports the template, you can right-click the Security Configuration and Analysis folder in the treeview pane and select Configure system now. Because the template defined only one setting, the snap-in will quickly process your database, and you'll be finished. Because you're using a security template to modify the system's local configuration, be aware that if for some reason the permissions on the Print Spooler service are ever changed, you'll have to reapply the template.

Alternatively, you can use Group Policy so that the permissions you specify will remain in effect and be reapplied each time the system refreshes Group Policy. Create a GPO and link it to the Domain Controllers organizational unit (OU). To apply the GPO to just one DC, open the GPO's properties and select the Security tab. Remove the Apply Group Policy permission from Authenticated Users. Then, add an entry that grants Read and Apply Group Policy permissions to the DC that runs the Print Spooler service. The settings you define in the GPO will apply to just that DC. Finally, edit the GPO and configure the `PrintSpooler` service as I described earlier.

—Randy Franklin Smith

Tip 9: Resolving a Windows XP-Related Password Error

Q Why does Windows XP prompt me to change my password, even though I haven't created one?

A If you upgraded to XP from an earlier Windows version, the OS can sometimes get confused and think you have a password. To resolve the problem, you can create a password then remove it by performing the following steps:

1. Open Control Panel, then select the User Accounts applet.
 2. Select your account, then click “Create a password.”
 3. Enter your password in both boxes, then click Create Password.
 4. Click “Remove my password,” type your password when prompted, then click “Remove Password.”
- John Savill

Tip 10: Changing Passwords Remotely Via the Web

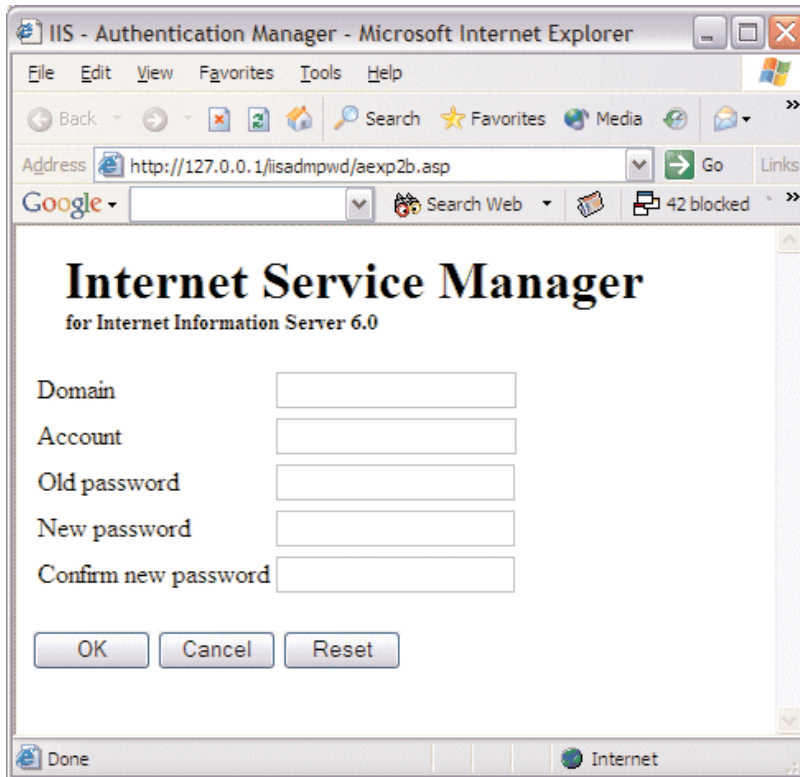
Q Does Windows Server 2003 provide a way to let users change their passwords remotely on the Web?

A The version of Internet Information Services (IIS) 6.0 that ships with Windows 2003 includes some Web-administration tools that are disabled by default. To enable the tools, perform the following steps:

1. Start the Microsoft Management Console (MMC) IIS Management snap-in by clicking Start, Programs, Administrative Tools, Internet Information Server (IIS) Management.
2. Navigate to Web Sites, Default Web Site.
3. Right-click Default Web Site. Select New, then select Virtual Directory. You’ll see the Virtual Directory Creation Wizard Welcome screen. Click Next.
4. Enter an alias of IISADMPWD, then click Next.
5. For the actual publish folder value, enter `C:\windows\system32\inetrv\iisadmpwd` (where `C:\windows` is the directory in which Windows is installed), then click Next.
6. For virtual directory permissions, select the *Read and Run scripts* check box, if it isn’t already selected. Click Next.
7. Click Finish.

You can access the new interface at `http://<server address>/iisadmpwd/aexp2.asp` to change a local account password or at `http://<server address>/iisadmpwd/aexp2b.asp` to change a domain password. Figure 1-3 shows a sample Web interface for changing a domain password.

Figure 1-3
Domain password Web interface



—John Savill

Tip 11: Configuring Pre-Staged RIS Permissions

Q What permissions does a user need at the Microsoft Remote Installation Services (RIS) client machine if the machine is pre-staged?

A If you've pre-staged a client machine for RIS, you must enter a domain account at the start of the RIS process. However, the user at the client machine won't need to have the rights to add computers to the domain because the computer account has been created in advance; instead, the user needs only the ability to read the computer account and the ability to reset the account password. To verify or add these settings, perform the following steps:

1. Start the Microsoft Management Console (MMC) Active Directory Users and Computers snap-in (go to Start, Programs, Administrative Tools, then click Active Directory Users and Computers).
2. Open the View menu and select Advanced Features to select the Advanced view.

3. Right-click the prestaged computer account, then select Properties from the context menu.
4. Select the Security tab, then click Add.
5. Select the user, or a group that the user belongs to, who will be entering his or her log on information at the start of the RIS process, then click OK.
6. Select the user or group that you added in Step 5 and verify that the user or group has read and reset password permissions; if not, select the Allow check box under the read permission and select the Allow check box under the reset password permission.
7. Click OK.

—John Savill

Tip 12: Changing a Domain User's Password from the Command Line

Q How can I change a domain user's password from the command line in Windows Server 2003?

A You can use the Dsmod command to modify directory service objects' attributes from the command line. More specific to your question, you can use "Dsmod user" to change the attributes of a user object. To modify a user's password, use the following syntax:

```
dsmod user <user's distinguished name (DN)> -pwd <user's new password>
```

For example, to change the password for user John in domain it.uk.savilltech.com, I typed

```
dsmod user CN=John,CN=Users,DC=it,DC=uk,DC=savilltech,DC=com pwd
Pa55word
```

The system returned

```
dsmon succeeded:CN=John,CN=Users,DC=it,DC=uk,DC=savilltech,DC=com
```

—John Savill

Tip 13: Comparing Code Access Security with User Access Permissions

Q How does code access security in the Windows .NET Framework affect or interoperate with user access permissions? Which one has the higher priority?

A Software written and compiled for the .NET Framework is called managed code. Classic applications, or unmanaged code, are subject only to the user's authority. Before Windows permits an application executable, such as Microsoft Word, to open a file, the Security Reference Monitor compares the file's ACL with the user's identity and with groups to which the user belongs, then grants or denies access accordingly. With unmanaged applications, which are called assemblies instead of executables, access control depends strictly on the user's authority and has nothing to do with the assembly.

The .NET Framework lets you exert very granular control over what assemblies can do based on various criteria about them. For example, you can control whether managed code can display windows; print; and access files, the network, and the registry; as well as whether the code can perform many other operations. You implement this control through code access permissions. You can grant or deny code access permissions based on a variety of criteria, including the assembly's name, the Web site from which it was downloaded, the publisher, whether the assembly originated on the local computer, and the Microsoft Internet Explorer (IE) or intranet security zone to which the code's hosting server belongs.

When an assembly tries to perform an operation or access a resource, the .NET Framework's Common Language Runtime (CLR) makes sure that the security policy allows the action. If the assembly passes muster with the CLR, execution continues as with any other Windows application. And, as with other Windows applications, the Security Reference Monitor checks whether the user's permissions, rights assignments, and group memberships permit the operation. Therefore, neither user access control nor code access control has precedence—they are equal. To complete an operation or access an object, managed code must pass code access control checks performed by the CLR, and the user must pass the same user access control checks that the Security Reference Monitor performs for unmanaged code.

For example, say Bob has read access to FileA, read and write access to FileB, and no access to FileC. Suppose further that he can use one of two programs to open these files: Notepad, an unmanaged application; or FileEditor, an imaginary text file editor written in a Microsoft .NET language. Let's assume that FileEditor is published by a company we'll call Acme and has a corresponding Authenticode signature from Acme. The CLR on Bob's computer has a security policy that grants assemblies published by Acme no access to FileA and FileB and modify access to FileC. If Bob uses Notepad, he'll be able to open FileA or FileB, but he won't be able to modify FileA. Bob won't be able to open FileC through Notepad because his user account and the groups he belongs to have no access to FileC.

If Bob tries to use FileEditor to open FileA for read access, he'll fail. FileEditor has no access to FileA, even though Bob does. He'll be able to open FileB through FileEditor, but only for read access; although FileEditor has read and write access, Bob has only read access. And if Bob tries to use FileEditor to open FileC, he'll fail. Thus, for managed code to execute, both Windows' traditional user-level security and the Framework's CLR must allow the code to run.

—Randy Franklin Smith

Tip 14: Enabling Users to Access Two Domain Accounts

Q We want to eliminate an old domain that was added to our network after an acquisition. To migrate the accounts in the old domain, we'll first create a new account in the main domain for each user in the old domain. Then, we'll gradually have users start logging on with their new account. Finally, we'll delete the old accounts. After we create the new accounts and before we delete the old ones, users need to be able to use both their new and old accounts to log on and access their files. How can we allow users to access both accounts without a major manual effort?

A powerful, free utility called SetACL, written by Helge Klein, can help you; it's available under the GNU General Public License (GPL) at <http://www.helge.mynetcologne.de/setacl>. SetACL lets you report, modify, back up, and restore permissions, auditing, and ownership for many types of objects, including folders, files, shares, registry keys, printers, and services. SetACL fully supports Windows NT and later permissions as well as inheritance and inheritance blocking. SetACL lets you perform any of its operations on just the specified object or on the specified object and all child objects.

First, I show you a SetACL command that will search the ACL for each file and folder on a hard disk for access control entries (ACEs) for users or groups from the old domain. For each such occurrence, SetACL will duplicate the ACE but substitute the new domain name for the old domain name. The command

```
SetACL.exe -on \\server1\share1
  -ot file -actn domain
  -rec cont_obj
  -dom "n1:OldDomain;
      n2:MainDomain;da:cpydom;
      w:dacL"
```

processes each file and subfolder in \\server1\share1. (Although this command appears on several lines here, you must enter it on one line in the command-shell window, making sure there are no spaces after the semicolons.) When SetACL encounters an ACE for a user or group in OldDomain, it looks up the same username or group name in MainDomain and finds the SID for that user or group. Then, SetACL uses the SID to create a new ACE with the same permissions in MainDomain. After you execute this command, user Fred will have the same access whether he logs on as olddomain\fred or maindomain\fred. After you complete the migration, you can use the command

```
SetACL.exe -on \\server1\share1
  -ot file -actn domain
  -rec cont_obj -dom
  "n1:OldDomain;da:remdom;w:dacL"
```

to delete all the old ACEs.

—Randy Franklin Smith

Chapter 2

Event Log and Auditing

Tip 15: Enabling Debug Logging for IP Security

Q How can I enable debug logging for IP Security (IPSec)?

A To enable logging for IPSec, which will result in logs being written to the %systemroot%\debug\oakley.log, perform the following registry change:

1. Start the registry editor (e.g., regedit.exe).
2. Navigate to the HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\PolicyAgent registry subkey.
3. From the Edit menu, select New, Key.
4. Enter the name Oakley, then click OK.
5. Select the Oakley key, then from the Edit menu, select New, DWORD Value.
6. Enter the name EnableLogging.
7. Double-click the new value and set it to 1.
8. Close the registry editor.
9. Restart the policy agent by typing

```
C:\> net stop policyagent  
C:\> net start policyagent
```

—John Savill

Tip 16: Using Log Parser to Audit Domain Logons

Q For auditing purposes, I'd like to keep track of logons to my domain whether they are successful (event ID 528—Successful Logon) or not (event ID 529—Logon Failure). I don't need to track all event ID 540 (Successful Network Logon) events. I've tried configuring filters, but I could set up only one filter at a time. Can you help me?

First, let me say that using Audit logon events, the audit category that generates event IDs 528, 529, and 540, isn't an accurate way to track logons to your domain. The Audit logon events category captures logons to the actual physical domain controller (DC) but doesn't capture logons by workstation users who use domain accounts to log on to the domain. To capture those events, you need to enable auditing on the local workstation. With Windows 2000 and later, you should use the Audit account logon events audit category, which logs each authentication that the DC performs regardless of whether the logon is local or originates from a workstation or server on the network.

To monitor successful domain logons, check all your DCs for event IDs 672 (Authentication Ticket Granted) and 680 (Account Used for Logon by %1). To monitor for logons that failed because of a bad password, look for event ID 675 (Pre-Authentication Failed) with failure code 0x18 and event ID 681 (The Logon to Account: %2 by: %1 from Workstation: %3 Failed. The Error Code was: %4) with error code 3221225578.

Second, you're right that Event Viewer lets you configure only one filter at a time. However, you can create a custom console that includes multiple instances of Event Viewer. (For details about how to create a custom console that contains multiple Security log views, see "Monitoring Security with Custom MMC Consoles," Security Administrator, March 2004, InstantDoc ID 41574, <http://www.windowsitpro.com/Article/ArticleID/41574/41574.htm>.)

The best way to filter your Security logs for logon-related events, however, is to use the Log Parser tool. Log Parser lets you use SQL-like queries to extract data from log files. You can download the tool from <http://www.microsoft.com/windows2000/downloads/tools/logparser/default.asp>.

Listing 2-1 shows a LogParser command that queries the local computer's Security log for event ID 675 with failure code 0x18 and event ID 681 with error code 0xC000006A, then sorts the results by event ID, date, and time.

LISTING 2-1: Sample Log Parser Command

```
logparser "select EventID, TimeWritten, Message from security where
(EventID = 675 and EXTRACT_TOKEN(Strings,4,'|')='0x18') or
(EventID = 681 and EXTRACT_TOKEN(Strings,3,'|')='0xC000006A')
order by EventID, TimeWritten"
```

(If you want to run this command on Windows Server 2003, you must first change EventID = 681 to EventID = 680 because of changes in the Windows 2003 Security log.) Figure 2-1 shows the resulting text file.

Figure 2-1
Resulting text file

EventID	TimeWritten	Message
675	1/16/2004 18:42:49	Pre-authentication failed: User Name: Administrator User ID: %S-1-5-21-2121316058-685099279-904526279-500} Service Name: krbtgt/ACME.LOCAL Pre-Authentication Type: 0x2 Failure Code: 0x18 Client Address: 127.0.0.1
681	1/16/2004 18:42:49	Logon attempt by: MICROSOFT_AUTHENTICATION_PACKAGE_V1_0 Logon account: Administrator Source Workstation: W03-IMAGE Error Code: 0xC000006A
681	1/16/2004 18:43:22	Logon attempt by: MICROSOFT_AUTHENTICATION_PACKAGE_V1_0 Logon account: administrator Source Workstation: CPQ Error Code: 0xC000006A

You can view the file directly or import it into Microsoft Access and format it into a report.

—Randy Franklin Smith

Tip 17: Monitoring For Unauthorized Scheduled Tasks

Q How can I monitor scheduled tasks on my server? I need to determine whether unauthorized jobs have been added. Can I glean this type of activity from the Security log?

A If you use Windows Server 2003, you can obtain information about unauthorized jobs. Microsoft added a new event ID to Windows 2003 for tracking newly created scheduled tasks. First, you need to enable Audit process tracking. Then, you'll see event ID 602 in the server's Security log, which Figure 2-2 shows, whenever someone adds a scheduled task.

Figure 2-2
Event ID 602

```

Event Type:      Success Audit
Event Source:    Security
Event Category:  Detailed Tracking
Event ID:        602
Date:           10/16/2003
Time:           11:08:05 AM
User:           W3\Administrator
Computer:       W3
Description:
Scheduled Task created:
  File Name:     C:\WINDOWS\Tasks\Calculator.job
  Command:      C:\WINDOWS\system32\calc.exe
  Triggers:     At 11:07 AM every day, starting 10/16/2003.
  Time:         10/17/2003 11:07:00 AM
  Flags:        0x18000C0
  Target User:  W3\Administrator
By:
  User:         Administrator
  Domain:       W3
  Logon ID:     (0x0,0xAA4D)

```

As you can see, Windows logs the user profile that added the task, the task's filename, the command the task will run, and the task's triggers.

—Randy Franklin Smith

Tip 18: Establishing DHCP Server Log Thresholds

Q While tracking down some suspicious activity in our logs, we noticed that our daily DHCP server logs have holes, apparently because the server stopped logging events. In each case, the server resumed logging at midnight with the next day's log. The only pattern we observed for the holes is that they occurred on busy days when DHCP activity was highest. Is the logging process hanging for some reason on busier days, then restarting at midnight?

A Windows Server 2003's and Windows 2000 Server's DHCP servers have built-in logic for controlling how much disk space the daily DHCP server logs consume. If a given day's activity exceeds the configured threshold, Windows stops logging DHCP events until either more disk space is available or the next day starts.

Windows lets you configure two thresholds for controlling event-log size. First, you can configure a maximum number of megabytes for all DHCP server audit logs combined. This threshold defaults to 7MB, and Windows restricts each day's log to one-seventh of the maximum space allowed for DHCP server audit logs. Thus, by default, each day's log can grow to a maximum of 1MB. Windows automatically overwrites week-old audit logs, so you retain only 1 week's activity.

Second, Windows lets you configure a minimum amount of space that must be preserved on the disk on which you store your audit logs. The default minimum is 20MB. If free space on the disk drops below the minimum threshold, Windows stops logging. DHCP starts logging again when disk conditions permit.

You can reconfigure the DHCP event log thresholds by using REG_DWORD values in the DHCP server's registry under the HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\DhcpServer\Parameters registry subkey. The DhcpLogMinSpaceOnDisk subkey specifies the minimum amount of space that must be preserved on the disk that contains your audit logs. Specify this value in megabytes—for example, if you want to stop DHCP logging when disk free space falls below 80MB, simply set the value to 80. The registry value that's more important to resolve your problem is DhcpLogFilesMaxSize. If this value is absent from the registry, Windows defaults to a value of 7, which gives you only 1MB for each day's log. Configure this value large enough to accommodate a full week's activity. If you want to keep more than 1 week's activity on hand, you'll need to copy each day's audit log file before Windows overwrites it the next week.

—Randy Franklin Smith

Tip 19: Operation-Based Auditing

Q We need to closely and accurately track changes to some important files. Windows 2000 Server and Windows NT Server let us use file auditing only when a user opens a file for a given type of access (e.g., write access). But a file that's been opened for write access isn't always changed—the user might have simply closed the file. I've heard that Windows Server 2003 lets you distinguish when a user has closed a file without using the access level and when a user has actually used the access he or she requested. Does Windows 2003 provide that ability? If so, how does it work?

A Windows 2003 supports operation-based auditing, which Microsoft originally introduced in Windows XP. Earlier versions of Windows log just two events whenever you work with a file. When you open a file, Windows logs event ID 560. Event ID 560 identifies you, the file you opened, the program with which you opened the file, and the types of access the program requested (e.g., read, write). When you close the file, Windows logs event ID 562. You can use the handle ID that appears in the details of both events to link events. However, as you've observed, earlier versions of Windows let you know only the types of access a user had for a file when the file was opened—the OS can't tell you whether the user actually performed an operation on the file. Operation-based auditing introduces event ID 567 (Object Access Attempt). This event shows you the specific access types that were used on the object.

Say that event ID 560 tells you that Fred used WordPad to open the C:\audittest\test.txt file for all types of read and write access. The Object Open section of this event lists handle ID 96 for Fred's file-open session, as Figure 2-3 shows.

Figure 2-3
Event ID 560

Event Type:	Success Audit
Event Source:	Security
Event Category:	Object Access
Event ID:	560
Date:	9/15/2003
Time:	3:34:52 PM
User:	W03RBC\Fred
Computer:	W03RBC
Description:	
Object Open:	
Object Server:	Security
Object Type:	File
Object Name:	C:\audittest\test.txt
Handle ID:	96
Operation ID:	{0,11218399}
Process ID:	3992
Image File Name:	C:\Program Files\Windows NT\Accessories\wordpad.exe
Primary User Name:	Administrator
Primary Domain:	W03RBC
Primary Logon ID:	(0x0,0x7DA0)
Client User Name:	-
Client Domain:	-
Client Logon ID:	-
Accesses:	READ_CONTROL SYNCHRONIZE ReadData (or ListDirectory) WriteData (or AddFile) AppendData (or AddSubdirectory or CreatePipeInstance) ReadEA WriteEA ReadAttributes WriteAttributes
Privileges:	-
Restricted Sid Count:	0
Access Mask:	0x12019F

But did Fred actually use his write access to change the file? Looking at subsequent events in the Security log, you find the event ID 567 that Figure 2-4 shows; handle ID 96 tells you that this event corresponds to Fred's file-open session.

Figure 2-4*Event ID 567*

Event Type:	Success Audit
Event Source:	Security
Event Category:	Object Access
Event ID:	567
Date:	9/15/2003
Time:	3:36:55 PM
User:	W03RBC\Fred
Computer:	W03RBC
Description:	
Object Access Attempt:	
Object Server:	Security
Handle ID:	96
Object Type:	File
Process ID:	3992
Image File Name:	C:\Program Files\Windows NT\Accessories\wordpad.exe
Accesses:	WriteData (or AddFile) AppendData (or AddSubdirectory or CreatePipeInstance)
Access Mask:	0x6

Looking at the Accesses section of that event, you see that Fred used WriteData and AppendData on the test.txt file at 3:36 p.m. In other words, Fred changed the file, then saved it in WordPad. Then, you find the event ID 562 that Figure 2-5 shows, and it tells you that Fred closed the file at 3:37 p.m.

Figure 2-5*Event ID 562*

Event Type:	Success Audit
Event Source:	Security
Event Category:	Object Access
Event ID:	562
Date:	9/15/2003
Time:	3:37:13 PM
User:	W03RBC\Fred
Computer:	W03RBC
Description:	
Handle Closed:	
Object Server:	Security
Handle ID:	96
Process ID:	3992
Image File Name:	C:\Program Files\Windows NT\Accessories\wordpad.exe

As you can see, event ID 567 provides the handle ID that the program obtained for the file when Fred opened it—the event doesn't provide the filename. To identify the filename, you must find an event ID 560 that has the same handle ID.

Operation-based auditing generates event ID 567 the first time a user exercises a specific permission on a particular file-open session. In other words, if Fred made further edits to the file, saved the file again, then closed it, Windows wouldn't log another event ID 567. However, if Fred reopened the file, changed it, then saved and closed it, Windows would generate another sequence of event IDs 560, 567, and 562.

No additional setup is required for operation-based auditing beyond what you must usually do to set up auditing on a file. To configure one computer, click Start, click Control Panel, double-click Administrative Tools, then double-click Local Security Policy. Navigate to Security Settings\Local Policies\Audit Policy, then click Audit Policy. In the Details pane, double-click Audit object access, then configure the Success and Failure auditing you want to use. To configure multiple computers, use the Microsoft Management Console (MMC) Group Policy Editor snap-in to enable the audit category in an appropriate Group Policy Object (GPO) under Computer Configuration\Windows Settings\Security Settings\Local Policies\Audit Policy.

The only difficult thing about using operation-based auditing is linking related event IDs 560 and 567. If you find an event ID 560 for an important file and you want to know which access types were exercised while the file was open, you must note the handle ID in event ID 560, then look for subsequent event ID 567s that have the same handle ID. When you find an event ID 567 and want to identify the file that it refers to, you must look for an earlier event ID 560 that has the same handle ID.

—Randy Franklin Smith

Tip 20: Auditing Account Logon Events Centrally

Q Two security events new to Windows 2000—event ID 680 and event ID 681—seem to mirror event ID 528 (user logon) and event ID 529 (failed logon: bad username or password), which have existed since Windows NT. How do these two sets of events differ?

A The difference lies in distinguishing local logon activity from domain controller (DC) authentication. One problem with NT has been that the OS records logon activity on the local computer only—not centrally at the DC. Whenever you log on at your NT workstation or connect to a server—or, more important, fail to log on or fail to connect—NT logs the event on the workstation and server, respectively. Thus, although NT captures logon events, it scatters them across several computers in the domain.

In Win2K, Microsoft added a new audit category, Audit account logon events, which lets you track logon activity centrally. If you enable Audit account logon events on a computer, whenever you attempt to log on with a domain account, Win2K logs the event at the local computer. If you enable Audit account logon events enabled on the DC, Win2K logs an event there, too.

Audit account logon events logs one set of event IDs when you use Kerberos authentication and a different set of event IDs when you use NT LAN Manager (NTLM). Event ID 680 (successful authentication) and event ID 681 (failed authentication) are NTLM authentication events. When you see event IDs 680 and 528 in proximity for the same user, a user has logged on to the DC itself

rather than on to a workstation or server. When you see event IDs 681 and 529 in proximity for the same user, it means that someone tried to log on to the DC itself with a bad password.

—Randy Franklin Smith

Tip 21: Deciphering Security Event ID 529

Q Why do I receive event ID 529 in my Security event log?

A Windows will generate event ID 529 if the machine environment meets the following criteria:

- The machine is running Windows XP.
- The machine is a member of a domain.
- The machine is using a machine local account.
- You've enabled logon failure auditing.

When the user logs off, Windows will write event ID 529 to the log file because the OS incorrectly tries to contact the domain controller (DC), despite the fact that the machine is using a local account. You can safely ignore this event ID.

—John Savill

Tip 22: Monitoring Security with Custom MMC Consoles

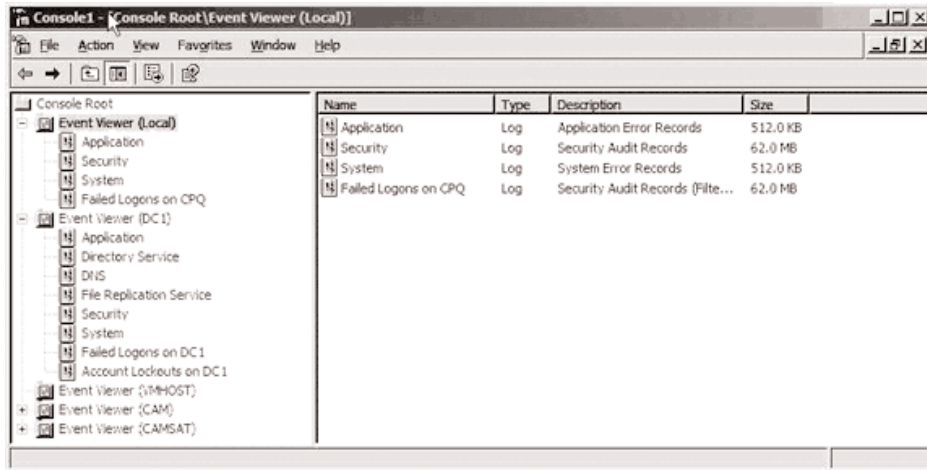
Q Every day, I check all servers I manage for security events. When I make my rounds, I have to connect to each computer individually and redefine the Event Viewer filters. Do you have a better approach?

A The solution is to build a custom Microsoft Management Console (MMC) console and use the MMC Event Viewer snap-in's New log view feature. First, create a new console by clicking Start, Run, then typing

```
mmc
```

and clicking OK. Select Add/Remove Snap-in from the File menu and click Add in the Add/Remove Snap-in window. In the Add Standalone Snap-in window, select Event Viewer from the list of available snap-ins and click Add. In the Select Computer window, select the Another computer check box and enter the name of one of your servers, then click Finish. Click Close, then click OK. Repeat this procedure to load a copy of the Event Viewer snap-in for each server you manage, as Figure 2-6 shows.

Figure 2-6
Creating multiple Event Viewer snap-ins

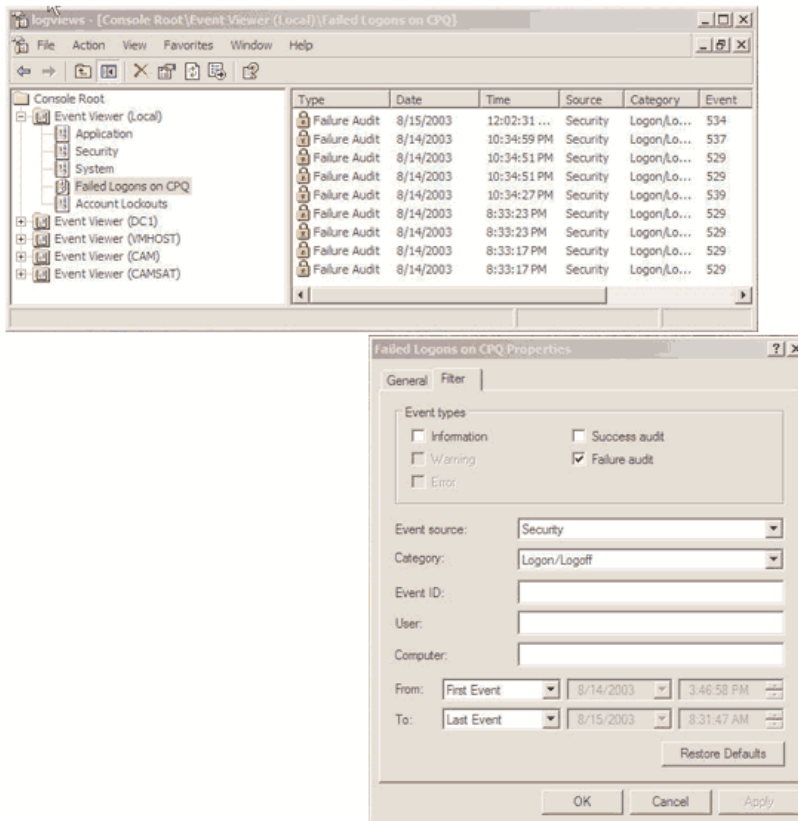


Next, expand Event Viewer for a server in the treeview pane, right-click the Security log, and select New Log View. MMC creates a new view of the Security log called Security (2). Right-click Security (2) and rename it to reflect the first type of event you regularly check for. For example, you might name the view Account Lockouts. Right-click the view again and select View, Filter. On the Filter tab, define the filter according to your needs. For example, to filter for account lockouts, enter 644 in the Event ID field, then click OK. Add another view to the instance of Event Viewer for each type of security event you want to monitor for on that server.

Select the other Event Viewers one by one and create Security log views for the types of security events you want to monitor on that server. After you create all your views, select Save As from the File menu and save the console in a place that you can find it later. Finally, create a shortcut to the new console on the Start menu. When you open the console, you'll be able to access the logs and customized views you created.

One of my favorite views on all types of computers shows failed logon attempts. To configure this view, select the Failure audit check box on the Filter tab, then choose Security from the Event source drop-down list and Logon/Logoff from the Category drop-down list, as Figure 2-7 shows.

Figure 2-7
Configuring the view for failed logon attempts



This view shows all failed attempts to log on to the computer either interactively or over the network. For domain controllers (DCs), I also like to create a view that shows failed authentication attempts, which filters for Failure audits for the Account Logon category. This view shows all failed attempts to log on through a domain account in the entire domain, whether the attempted logon used Kerberos or Windows NT LAN Manager (NTLM). You need to create and monitor this view on each DC. You can keep up with new users being added to groups in your domain by creating three views—New Members in Global Groups, New Members in Local Groups, and New Members in Universal Groups—and filtering for event IDs 632 (global group member added), 636 (local group member added), and 660 (security enabled universal group member added), respectively.

—Randy Franklin Smith

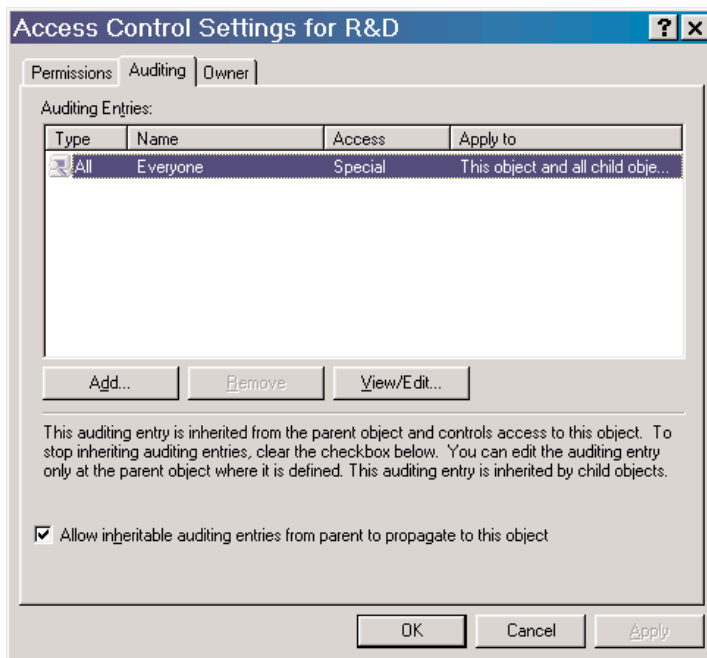
Tip 23: Audit Control List Editing Rights for a Win2K Object

Q Is the owner of a Windows 2000 object's audit control list the only one who can edit it? If not, is there a permission that lets me delegate this ability?

A You can't delegate the ability to edit the audit settings on an object as you can delegate other types of access (e.g., Read, Modify, Full Control). Win2K controls who can change audit settings on objects through the *Manage auditing and security log* right. To view or edit the audit control list for an object such as a file in Windows Explorer or an organizational unit (OU) in the Microsoft Management Console (MMC) Active Directory Users and Computers snap-in, right-click the object, select Properties, click the Security tab, then click Advanced. If you have the Manage auditing and security log right on the computer on which the object resides, you'll find the Auditing tab that Figure 2-8 shows.

Figure 2-8

Selecting the Access Control Settings Auditing tab



The Manage auditing and security log right not only governs who can configure auditing for all the objects on the system but also lets you view and clear the Security log.

—Randy Franklin Smith

Tip 24: Creating Multiple Event Viewer Views

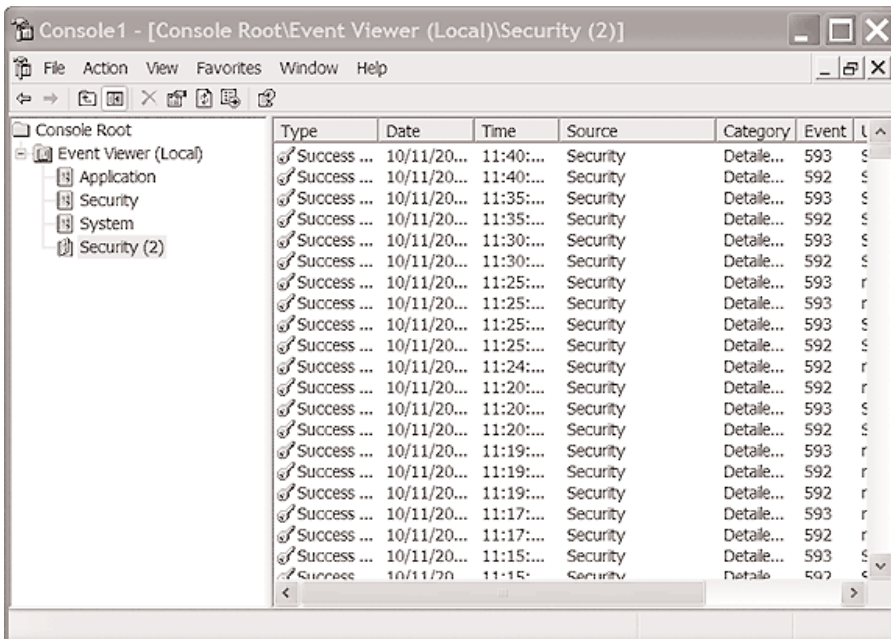
Q I'd like to check my domain controllers (DCs) every day for failed attempts to log on to the DC (failed events in the Logon/Logoff category), failed authentication attempts throughout the domain (event IDs 675 and 681), and account lockouts (event ID 644). The filtering feature of the Event Viewer (eventvwr.exe) lets me display only one event type at a time, and changing the view three times a day becomes onerous. Do you know of a better way?

A Yes. The Eventvwr tool lets you create multiple views of each log you open. You can customize the filter criteria for each view, then save your settings in a Microsoft Management Console (MMC) console, which will recreate all the views and filter criteria each time you reopen it.

First, create a new MMC console. Select Console, Add/Remove Snap-in, then click Add to open the Add Standalone Snap-in window. Select Event Viewer from the Available Standalone Snap-ins list, and click Add to open the Select Computer window. Select Another Computer and enter the name of your first DC. Click Finish.

Repeat the process to add additional instances of the Event Viewer snap-in to the console for each of your DCs. Click Close in the Add Standalone Snap-in window, and click OK in the Add/Remove Snap-in window. In the tree pane of the main console window, double-click the first instance of Event Viewer, right-click Security, and select New Log View to create another view of the Security log called Security (2), as Figure 2-9 shows.

Figure 2-9
Creating a view of the Security Log

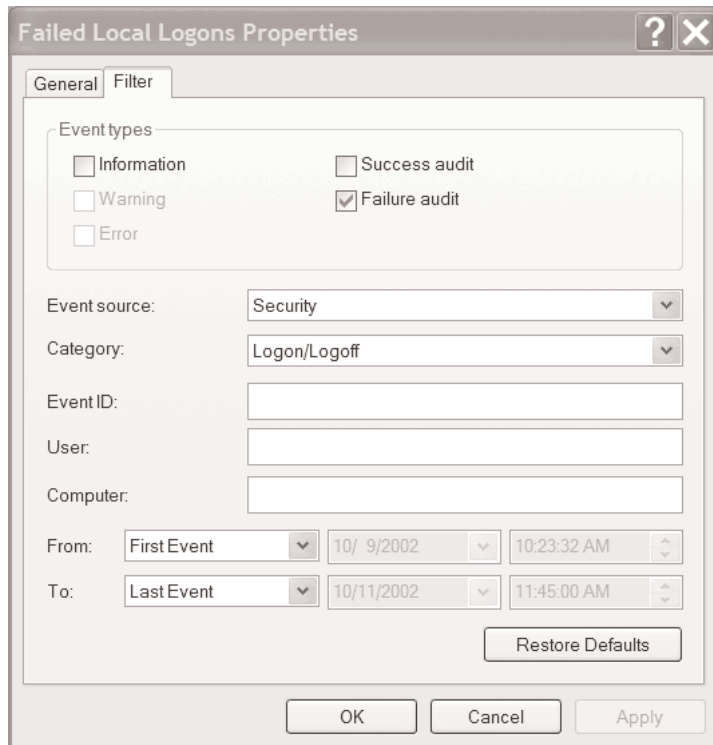


Right-click Security (2), select Rename, and enter Failed Kerberos Pre-Authentication. Right-click Failed Kerberos Pre-Authentication, select View\Filter, enter 675 in the Event ID field, and click OK.

Next, create another Security log view called Failed NTLM Authentication and filter it for event ID 681. This time, configure the filter to limit event types to Failure audit, and set Event source to Security and Category to Logon/Logoff, as Figure 2-10 shows.

Figure 2-10

Configuring a view of failed authentication



Create one more view called Account Lockouts and filter it for event ID 644. Repeat the process for the other instances of Event Viewer so that you have the same views for each DC. Finally, click File, Save As and name your new console. By default, MMC saves new consoles to Administrative Tools on the Start menu. Now, you can simply open your console each morning and check the logs without having to change filters.

—Randy Franklin Smith

Tip 25: Viewing Security Logs for All DCs

Q I want to let a group of users who aren't members of Enterprise Admins or Domain Admins view the Security logs of all domain controllers (DCs) in a forest. What user rights do I need to apply?

A To view the Security log, users must have the *Manage auditing and security log* user right on each computer whose log they want to view. To grant this right on all your DCs, open the Microsoft Management Console (MMC) Active Directory Users and Computers snap-in. Create a global group, give it an appropriate name, such as DCSecurityLogViewers, and populate the group with the users who need access to the Security log. Then, simply edit the Default Domain Controllers Policy Group Policy Object (GPO), which is linked to the Domain Controllers organizational unit (OU). To edit this GPO, go to Computer Configuration\Windows Settings\Security Settings\Local Policies\User rights assignment and assign the Manage auditing and security log right to the DCSecurityLogViewers group. Everyone in the DCSecurityLogViewers group will be able to view the Security log of each DC in the current domain. To grant DCSecurityLogViewers the same access on the DCs in other domains, make the same change to the Default Domain Controllers Policy GPO in each domain.

Be aware that this user right also lets users clear the Security log. You can't give someone the ability to view the Security log without giving them the authority to clear it as well. And you can't restrict administrators from viewing or clearing the Security log—although you can deny administrators the Manage auditing and security log user right, they can grant the right to themselves. If you don't want your users to be able to clear the Security log or to be able to grant the Manage auditing and security log user right to themselves, you must set up a process that frequently exports each computer's Security log to a central database, to which you can then control access as you see fit.

—Randy Franklin Smith

Chapter 3

Security Policy

Tip 26: Using Windows Update with Security Policies

Q After reading an article about implementing IP Security (IPSec) packet filters to protect Web servers, I blocked all traffic at my test server, then created exception rules to allow incoming packets to TCP ports 80 (HTTP), 20 and 21 (FTP), and 3389 (Terminal Services) and to let the server send packets back to clients. No other ports are open, and I feel much more secure. However, Windows Update no longer works. When I try to browse to <http://windowsupdate.microsoft.com>, Microsoft Internet Explorer (IE) fails to connect. How can I keep my system locked down but still let it download Microsoft updates?

A Setting up an IPSec policy like the one described above prohibits not only incoming packets but also outgoing packets unless they're explicitly allowed. Although static filters (such as those in IP security policies) are great for controlling connections to inbound ports, they don't work well for outgoing connections because client-side port numbers are numerous and unpredictable. This drawback is why stateful inspection firewalls are so important. As powerful as they are, IPSec policies fall short in this area because they aren't stateful.

You have at least one alternative: You can create an exception rule in your IPSec policy that lets you communicate through port 80 to <http://windowsupdate.microsoft.com>. Or, if you can upgrade to Windows Server 2003, you can use its built-in firewall to solve your problem.

The first alternative involves creating a filter that looks for inbound packets that have the windowsupdate.microsoft.com source address and source port TCP 80. Then, create a rule that allows traffic through that filter. Configure the rule to include mirror image packets (i.e., outgoing packets to <http://windowsupdate.microsoft.com> and destination port TCP 80), and you'll be able to download updates from that Web site. However, be aware that this approach opens you up to incoming connections through any port on your server if the attacker can spoof packets to look like they come from <http://windowsupdate.microsoft.com>.

If you run Windows 2003, you can solve your problem simply by enabling Internet Connection Firewall (ICF) or RRAS's Basic Firewall, depending on which edition of Windows 2003 you have. Evidently, you use IPSec policies only for their packet-filtering ability, not for IPSec communications. ICF and Basic Firewall let you easily implement that filtering functionality while letting your computer make outgoing Web requests. First, enable ICF or Basic Firewall. Configure the firewall to publish your public TCP ports 80, 20, 21, and 3389, then disable your current IPSec policy. Your private ports will still be blocked, but because ICF and Basic Firewall are both stateful inspection firewalls, your server can make outgoing Web requests.

—Randy Franklin Smith

Tip 27: Using One GPO to Control Both Windows XP and Windows 2000 Settings

Q Our domain organizational unit (OU) structure is set up in such a way that I have a mix of Windows XP and Windows 2000 computers in the same OU. By looking at the local Group Policy Object (GPO) on an XP computer, I can tell that XP's Group Policy settings are different from Win2K's settings. How can I use GPOs that are stored in Active Directory (AD) to manage these Group Policy settings? Also, can I use one GPO linked to my OU to manage both kinds of computers?

A You can manage both XP's new Group Policy settings and Win2K's settings from the same GPO. But first, you must update the GPO to include XP's new settings. To do so, copy the `system.adm`, `wmplayer.adm`, `conf.adm`, and `inetres.adm` files from `%systemroot%\inf%` on an XP computer to a folder somewhere on your domain controller (DC). Then, from a Win2K computer, edit the GPO you want to update. Right-click `Computer Configuration\Administrative Templates` and select `Add/Remove Templates`. In the `Add/Remove Templates` dialog box, remove the existing copies of the above files and add the templates you copied from an XP computer.

After you load the new templates, the GPO won't look different when you edit it on a Win2K computer—you'll still see all the Win2K policies you're accustomed to. However, when you edit the same GPO from an XP computer, you'll see all of XP's new settings.

The reason you see different settings for the different OSs is that the same GPO holds both versions of policies, and each computer that applies the GPO looks for and applies the appropriate policies for that computer's version of Windows. However, to edit both versions of the policies, you must use an XP computer to edit the XP settings and a Win2K computer to edit the Win2K settings.

—Randy Franklin Smith

Tip 28: Preventing Users from Disabling Group Policy

Q I've heard about a registry setting that local users can configure to prevent their computers from applying security settings and other restrictions that we've defined in Group Policy. Does such a setting exist? Could users use other techniques to disable Group Policy?

A You're referring to a registry value that existed in a beta version of Windows 2000. The setting was a `REG_DWORD` value under the `HKEY_LOCAL_MACHINE\SOFTWARE\Policies\Microsoft\Windows\System` registry subkey; the value was called `DisableGPO` and set to 1. However, this tweak has no effect on the final release of Windows XP or Win2K. The computer always applies Group Policy whether or not this value exists.

With regard to other techniques by which users could circumvent Group Policy, no methods have been published. I've tested some likely methods, but I've never succeeded in disabling Group Policy even when using the local administrator account. For example, the most obvious possible method for disabling Group Policy is deleting the program that applies it—`gpupdate.exe` on XP or `secedit.exe` on Win2K. But deleting these files simply causes Windows File Protection (WFP) to replace them.

I've also tried to deny everyone access to these files. Doing so prevents me from refreshing Group Policy manually but doesn't stop the system from applying Group Policy, including changes made to Group Policy after I denied access to `secedit.exe` or `gpupdate.exe`. Group Policy is deeply embedded in the OS and can't easily be circumvented by users or local administrators.

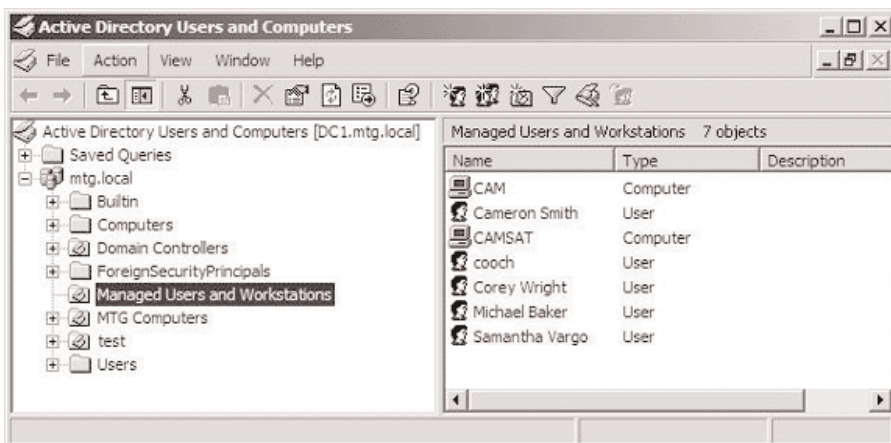
—Randy Franklin Smith

Tip 29: Connecting to a DC to Edit a GPO

Q I've been monitoring changes made to Group Policy Objects (GPOs), and I've noticed that, apparently, only one of my two domain controllers (DCs) will write GPO change events to the Security event log. Regardless of where I make the change (i.e., on DC1, on DC2, or on a workstation), the log always shows that DC1 wrote the change, and the event shows up only on DC1. DC1 has several Flexible Single-Master Operation (FSMO) roles; is that why it's the only computer that tracks GPO change events?

A FSMO roles don't come into play when you edit GPOs. You can edit a GPO's copy on any DC within the domain. The DC behavior that you describe is an indicator of which DC the Microsoft Management Console (MMC) Active Directory Users and Computers snap-in happens to connect to. When you open the Active Directory Users and Computers snap-in on a DC, the snap-in doesn't necessarily connect to the local DC—it might connect to another DC on the network. To determine which DC you're connected to, check the computer name displayed in brackets beside the root in the treeview pane. For example, Figure 3-1 shows a connection to DC1.mtg.local.

Figure 3-1
Determining which DC you're connected to



You can connect to a different DC in the same domain by right-clicking the root of the tree and selecting *Connect to Domain Controller*.

—Randy Franklin Smith

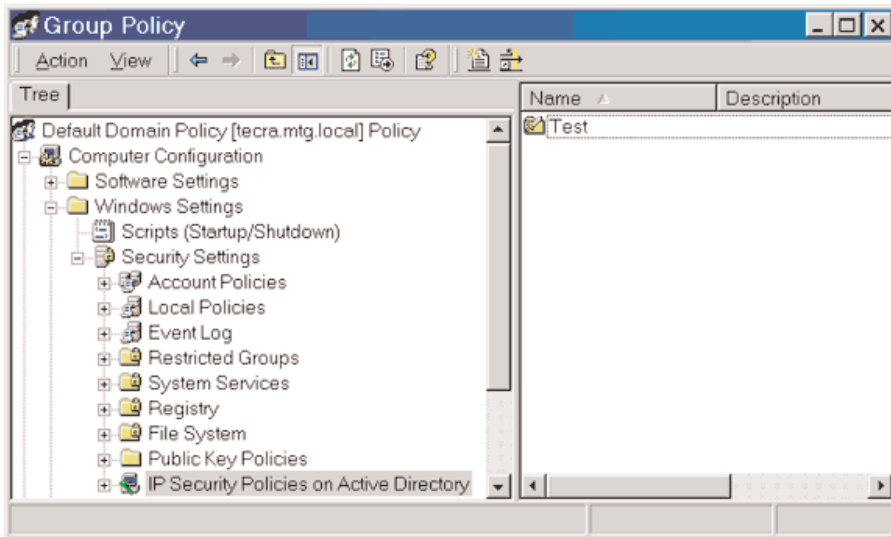
Tip 30: Editing an IP Security Policy

Q When I edit an IP Security (IPSec) policy in a Group Policy Object (GPO), the changes don't take effect unless I first delete and recreate the IPSec policy object. How can I avoid this unnecessary work?

A A Windows 2000 bug causes this problem; fortunately, the bug has a simple workaround. An IPSec policy is an Active Directory (AD) object that's separate from the GPO in which you define it. To demonstrate how the bug works, create but don't assign an IPSec policy called Test in your Default Domain Policy GPO. (When you create an IPSec policy, it doesn't take effect until you right-click it and select Assign.) Open the Microsoft Management Console (MMC) Active Directory Users and Computers snap-in, right-click the root domain, then select Properties. Click the Group Policy tab, highlight Default Domain Policy, then click Edit. In the Group Policy window that Figure 3-2 shows, drill down to \Computer Configuration\Windows Settings\Security Settings\IP Security Policies on Active Directory.

Figure 3-2

Accessing the IPSec policy called Test



Right-click in the right pane, then select Create IP Security Policy. Edit your Default Domain Controllers Policy GPO. You'll see the new IPSec policy Test there, too. At boot up and about every 90 minutes thereafter, a Win2K computer checks AD to see whether any of the GPOs that apply to the computer have been changed—either because you've linked new GPOs to the computer's site, domain, or OU or because you've deleted relevant GPOs. If the check reveals changes, the Win2K computer reapplies group policy to apply the changes.

To determine whether a GPO has changed, Win2K checks the GPO's version number. By reappling Group Policy only when changes have occurred, Win2K saves network and system

resources. However, this approach causes a problem when it comes to IPSec policies. Because an IPSec policy object is separate from the GPO, when you edit an IPSec policy, Win2K doesn't update the GPO's version number. Thus, the computers to which that GPO applies miss the updated IPSec policy. To work around this problem, remember to unassign and assign the IPSec policy after you edit it. Reassigning the IPSec policy increments the GPO's version number. Also, remember to reassign the IPSec policy in any other GPOs in which you use it. Otherwise, you might think you've deployed an important security change related to IPSec in your network that never takes effect.

—Randy Franklin Smith

Tip 31: Understanding Group Policy's Block Policy Inheritance and No Override Options

Q I want to make sure that I'm applying the security settings in my Group Policy Objects (GPOs) correctly. In Group Policy, what's the relationship between the Block Policy inheritance and No Override options, and how can I best use them?

A In short, No Override takes precedence over Block Policy inheritance, but read on. Remember that Windows 2000 applies GPOs in a specific sequence. Win2K first applies a local computer's GPO, then (in order) any site-linked GPOs, domain-linked GPOs, and organizational unit (OU)—linked GPOs. When two or more GPOs define a value for the same policy (with very few exceptions, such as logon scripts), the last policy wins. For example, if you define the Audit account management category with Success, Failure at the domain level but specify Failure for the same policy in a GPO linked to a lower-level OU (i.e., OUs beneath the domain), computers in that lower-level OU will end up with the Audit account management category set to Failure.

You can specify the Block Policy inheritance setting on domains and OUs. To do so, open the Microsoft Management Console (MMC) Active Directory Users and Computers snap-in, double-click a domain or OU, and click the Group Policy tab. If you select the Block Policy inheritance option at the domain level, when computers in this domain apply Group Policy, they won't apply any site-linked GPOs. If you select the Block Policy inheritance option on an OU, computers in this OU won't apply site-linked GPOs, domain-linked GPOs, or GPOs linked to higher-level OUs. Note that Win2K always applies a computer's local GPO regardless of Block Policy inheritance settings, but because the local GPO is the first one applied, any conflicting policies in subsequent GPOs override the local GPO. You can use the Block Policy inheritance option when you have a subset of computers or users that you want to insulate from policies you set at the domain or higher level. Put those users or computers in an OU and select the Block Policy inheritance check box. Now, you can manage those computers exclusively through GPOs linked to that OU.

What I've described is default behavior, but consider what happens when you use the No Override option. You select the No Override option by clicking that column in the list of GPOs. No Override is a GPO link-level setting instead of a domain- or OU-level setting. Therefore, if you link the same GPO to more than one site, domain, or OU, the No Override setting won't follow the GPO. You can control No Override at each point at which a GPO is linked. If you specify No Override on a GPO link, the policies you've defined in that GPO override any conflicting policies in GPOs processed later in the Group Policy application sequence. Policies that you define in No Override

GPO links defeat conflicting policies even in GPOs that specify the Block Policy inheritance setting or other subsequently applied GPOs that specify the No Override setting.

You can use the No Override setting to configure mandatory policies. For example, you might have certain default domain-level policies (i.e., you can override them at lower OUs to manage legitimate exceptions). You can configure these policies in the Default Domain Policy GPO. You might also have policies that you want to apply without exception to all computers or users in the domain. If so, define these mandatory policies in a new GPO that you create called Mandatory Domain Policies, link the Mandatory Domain Policies GPO to the domain, and configure the new GPO link with the No Override setting. Rest assured that policies that you define in Mandatory Domain Policies will override any policy conflicts that OU-linked GPOs inadvertently create at lower levels in the domain.

—Randy Franklin Smith

Chapter 4

IP Security

Tip 32: Defining IP Security

Q What Is IP Security (IPSec)?

Although TCP/IP is widely used in most networks and forms a compulsory part of any Windows network, it does present a few challenges. Because TCP/IP doesn't encrypt data sent across the network, that data is vulnerable to a number of attacks, including eavesdropping. Being able to view data sent over the network can allow an attacker to view data such as passwords when connecting to services such as FTP, which doesn't encrypt passwords sent over the network.

To help solve this security risk, IPSec was developed as an industry standard based on end-to-end security—only the transmitting and receiving computers know about any encryption that takes place. In cooperation with Cisco Systems, Microsoft developed an implementation of IPSec in Windows 2000 that along with Win2K's Group Policy settings lets you define how your environment handles IP security. One of the great things about IPSec in Windows is that it operates at layer 3 of the OSI model, so any application of IP and upper-layer protocols such as TCP and UDP gain the advantage of IPSec without requiring any modifications to the applications.

—John Savill

Tip 33: Stopping and Restarting the IP Security Policy Agent

Q How can I stop and restart the IP Security (IPSec) policy agent on a machine?

Windows 2000's policy agent is responsible for handling IPSec negotiations between machines. If you experience problems and want to restart the agent, you can stop and restart its service by typing the following commands at the command prompt, respectively:

```
C:\> net stop policyagent  
C:\> net start policyagent
```

—John Savill

Tip 34: Defining an IP Security Policy for a Group Policy Object

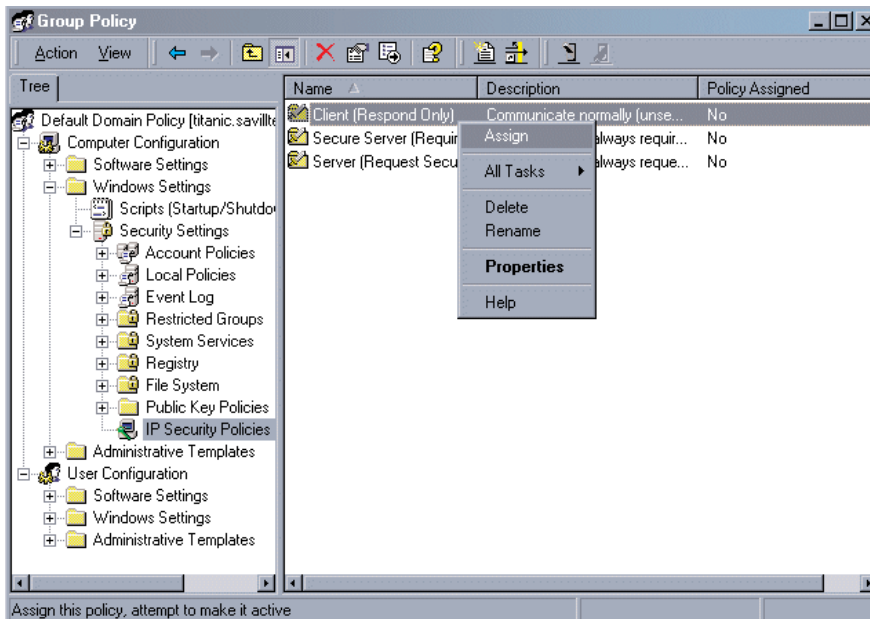
Q How can I define an IP Security (IPSec) policy for a Group Policy Object (GPO)?

A In a typical networking environment, IPSec is defined using a GPO configured on a domain or organizational unit (OU). To define an IPSec policy, perform the following steps:

1. Start the Microsoft Management Console (MMC) Active Directory Users and Computers snap-in (go to Start, Programs, Administrative Tools, and click Active Directory Users and Computers).
2. Right-click the container that has the GPO, then select Properties (e.g., the domain).
3. Select the Group Policy tab.
4. Select the Group Policy Object, then select Edit.
5. Expand the Computer Configuration root.
6. Expand Windows Settings, Security Settings, IP Security Policies.
7. Right-click the policy you want to assign, then select Assign from its context menu, as Figure 4-1 shows.

Figure 4-1

Assigning an IPSec policy to a GPO



8. You can only assign one IPSec policy to the GPO—if you assign more than one, the previously assigned one will be unassigned.
9. Close Group Policy Editor (GPE).

If you want to remove a policy, you need to right-click the assigned policy and select Un-assign from the context menu. Unlike other Group Policy settings, IPSec policies will remain even if you delete the GPO, so make sure you un-assign the policy before deleting the GPO.

To force a GPO update, from the command prompt type

```
C:\>secedit /refreshpolicy machine_policy /enforce
```

You'll see the following message:

```
Group Policy propagation from the domain has been initiated for this computer. It may
take a few minutes for the propagation to complete and the new policy to take effect.
Please check Application Log for errors, if any.
```

—John Savill

Tip 35: Changing the Authentication Method Used for IP Security

Q How can I change the authentication method used for IP Security (IPSec) by a policy?

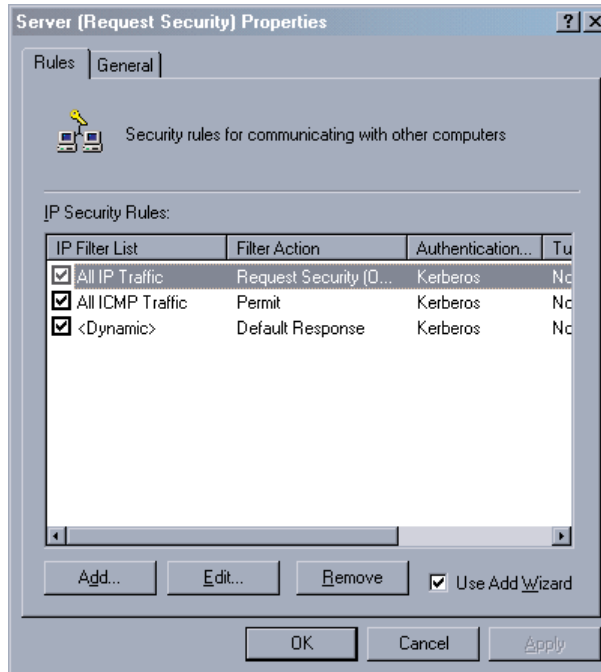
A By default, IPSec uses the Kerberos V5 protocol for its authentication method. However, other options include

- using a certificate from a selected certificate authority
- using a predefined string (a preshared key)

To modify an existing IPSec policy, start the Microsoft Management Console (MMC) IP Security Policy snap-in and perform the following steps:

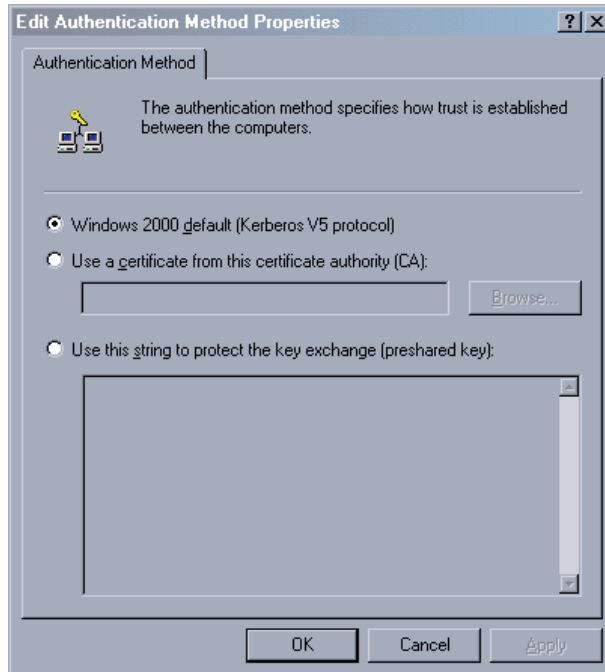
1. Right-click the policy, then select Properties from the context menu.
2. From the list, as Figure 4-2 shows, select one of the security rules you want to change the authentication method for, then click Edit.

Figure 4-2
Selecting IP Security Rules



3. Select the Authentication Methods tab. The current authentication method will be shown (e.g., Kerberos by default).
4. Click Edit.
5. Select the preferred authentication method from the dialog box that Figure 4-3 shows.

Figure 4-3
Selecting an authentication method



6. Click Apply, then click Close.
7. Close all dialog boxes.

If the change was made on a domain Group Policy Object (GPO), you can force the change to take effect by going to the command prompt and typing

```
C:\> secedit /refreshpolicy machine_policy
```

—John Savill

Tip 36: Enabling IP Security

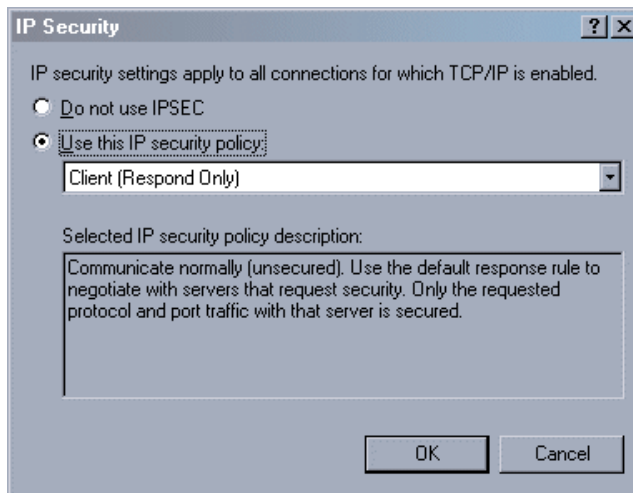
Q How can I enable IP Security (IPSec) on a machine?

A You typically use group policies to assign IPSec in a Windows 2000 domain; however, you can configure IPSec on a per-computer basis by performing the following steps:

1. Right-click My Network Places, then select Properties.
2. Right-click Local Area Connection, then select Properties.

3. Select Internet Protocol (TCP/IP), then click Properties.
4. Click the Advanced button.
5. Select the Options tab.
6. Select *IP security*, then click Properties.
7. Select *Use this IP security policy*, as Figure 4-4 shows, then select a policy to use:
 - Client (Respond Only) - It will only use IPsec if asked to by the other end of a session
 - Secure Server (Require Security) - All IP traffic requires security using Kerberos trust
 - Server (Require Security) - Use IP security if possible


Figure 4-4
Selecting an IPsec policy



8. Click OK to close all dialog boxes.

To set IPsec on a machine, you must be a member of the local Administrators group. Also, if a domain IPsec policy has been defined, you can't override it with a local policy—the options will be grayed out, as Figure 4-5 shows.

Figure 4-5
Confirming that a domain IPsec policy is in effect



—John Savill

Tip 37: Managing and Creating IP Security Policies

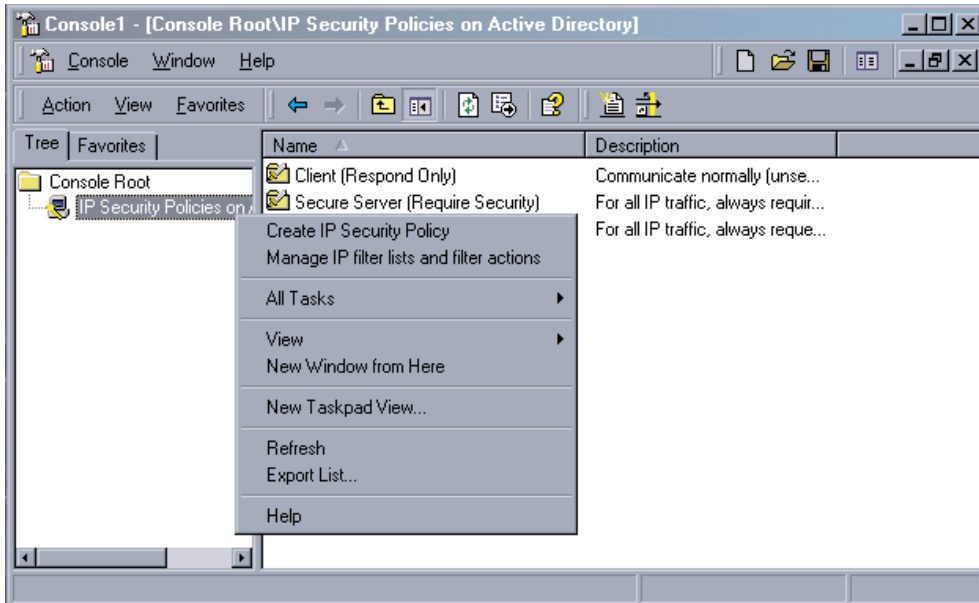
Q How can I manage or create IP Security (IPsec) policies?

A In Windows 2000, you can use the Microsoft Management Console (MMC) IP Security Policies snap-in to modify and create IPsec policies, which you can then assign to computers and Group Policy Objects (GPOs). To open the IP Security Policies snap-in, perform the following steps:

1. Start the MMC (go to Start, Run, and type `mmc.exe`).
2. From the console menu, select Add/Remove Snap-in (or press `Ctrl+M`).
3. From the Standalone tab, click Add.
4. Select IP Security Policy Management snap-in, then click Add.
5. Select either *Local computer* or the domain policy, then click Finish. If it's for a domain, select *Manage domain policy for this computer's domain*, then click Finish.
6. Click Close to close the dialog box, then click OK.
7. Double-click the root to display the three built-in options:
 - Client (Respond Only)
 - Secure Server (Require Security)
 - Server (Request Security)

If you right-click the root, you can create a new policy by selecting Create IP Security Policy, as Figure 4-6 shows.

Figure 4-6
Creating an IPSec policy



If you right-click an existing policy and select Properties, you can modify its settings.

—John Savill

Tip 38: Enabling IP Security Traffic through a Firewall

Q How can I enable IP Security (IPSec) traffic through a firewall?

A IPSec is generally invisible to routers because it operates at layer 3 of the OSI model and Windows encrypts all IP and upper-layer protocols. However, firewalls and gateways in the data path must be able to forward the following IP protocols for IPSec to correctly work:

- IP Protocol ID 50—This protocol is used for both inbound and outbound filters and is needed for forwarding Encapsulating Security Payload (ESP) traffic.
- IP Protocol ID 51—This protocol is used for both inbound and outbound filters and is needed for forwarding Authentication Header (AH) traffic.
- UDP Port 500—This protocol is used for both inbound and outbound filters and is needed for forwarding Internet Security Association and Key Management Protocol (ISAKMP) traffic.

Layer Two Tunneling Protocol (L2TP)/IPSec traffic looks the same as IPSec traffic on the wire and you need to open IP Protocol ID 50 and UDP Port 500.

—John Savill

Tip 39: Defining the IP Security/Layer Two Tunneling Protocol NAT-T Update

QWhat's the IP Security (IPSec)/Layer Two Tunneling Protocol (L2TP) NAT-T update for Windows XP and Windows 2000?

AThe IPSec/L2TP NAT-T update is a Microsoft update for L2TP and IPSec for XP and Win2K that lets you operate VPN clients behind Network Address Translation (NAT) software or hardware. The update is available from the Windows Update Web site and requires XP Service Pack 1 (SP1) or later or Win2K SP3 or later.

After you install the update, clients behind the NAT device will be able to create IPSec connections and monitor those connections through the updated monitoring tool that installs as part of the update. For more information, see the Microsoft article “L2TP/IPSec NAT-T Update for Windows XP and Windows 2000” at <http://support.microsoft.com/?kbid=818043>.

—John Savill

Tip 40: Disabling IP Security on a VPN Connection that Uses Layer Two Tunneling Protocol

QHow can I disable IP Security (IPSec) on a VPN connection that uses the Layer Two Tunneling Protocol (L2TP)?

AWindows automatically creates an IPSec policy for L2TP connections because L2TP doesn't encrypt data. However, you might want to test a VPN L2TP connection without the security of IPSec (e.g., when troubleshooting). Although you must disable IPSec on both the client and server in this situation, make sure you re-enable the security policy after you resolve any problems; otherwise, your systems are vulnerable to attack. To disable IPSec, perform the following steps on both ends of the connection (client and server):

1. Open a registry editor (e.g., regedit.exe).
2. Navigate to the HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\RasMan\Parameters subkey.
3. From the Edit menu, select New, DWORD Value.
4. Enter a name of ProhibitIpSec and press Enter.
5. Double-click the new value, set it to 1, then click OK.
6. Restart the machine.

For more information, see the Microsoft article “How to Configure a L2TP/IPSec Connection Using Pre-shared Key Authentication” at <http://support.microsoft.com/default.aspx?scid=kb;EN-US;q240262>.

—John Savill

Tip 41: Preventing Attackers from Bypassing IP Security Packet Filtering

QI've been using the Windows 2000 IP Security (IPSec) feature's packet filtering to block access to vulnerable ports on my Web server as a fail-safe measure in case my firewall is ever compromised or misconfigured. However, I understand that a sophisticated attacker can use specially formed packets to bypass Win2K's IPSec packet filtering. Is this true? If it is, how can I eliminate that vulnerability?

A By default, Win2K doesn't block packets from source TCP port 88 (Kerberos) or UDP port 500 (Internet Key Exchange—IKE) regardless of the block filters you configure in your IPSec policy. Consequently, an attacker who's a bit more sophisticated than the average script kiddy can send packets to any destination port on your computer by spoofing the source port to make the packet look like a legitimate Kerberos or IKE packet. In Win2K Service Pack 1 (SP1), Microsoft responded to this problem with the new NoDefaultExempt REG_DWORD registry value under the HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\IPSEC registry subkey. If you set NoDefaultExempt to 1, Win2K will no longer exempt Kerberos packets from your IPSec packet filters. For more information about NoDefaultExempt, see the Microsoft article “IPSec Does Not Secure Kerberos Traffic Between Domain Controllers” (<http://support.microsoft.com/?kbid=254728>).

NoDefaultExempt doesn't, however, block IKE packets on UDP port 500. If you want to have full control over packet filtering, you might consider using Jean-Baptiste Marchand's PktFilter freeware utility, which is available at <http://www.mirrors.wiretapped.net/security/firewalls/pktfilter>. PktFilter runs as a Win2K service. By editing PktFilter's rules.txt file, you can block or allow packets based on the NIC, direction, protocol (e.g., TCP, UDP), source and destination addresses and ports, and other criteria.

—Randy Franklin Smith

Chapter 5

Kerberos

Tip 42: Defining Kerberos

Q What's Kerberos?

Named after a three-headed dog that guards the entrance to Hades, Kerberos replaces the Windows NT LAN Manager (NTLM) native communication standard for Windows 2000 computers; however, Win2K still supports NTLM for compatibility with older Windows NT 4.0 and Windows 9x clients but doesn't support NTLM version 2.


If two people share a secret, they can communicate by encrypting a message with the secret and the first person can know that the second person is who he says he is and vice versa. But how do you securely share the secret? The problem is that you can't just send the secret as plain text over the network because anyone with a network sniffer might discover the secret.

Kerberos solves this problem by using secret key cryptography. Rather than sharing a password, communication partners share a cryptographic key that's symmetric in nature, meaning the single key can both encrypt and decrypt.

To communicate, the first person (e.g., John) sends the second person (e.g., Bob) an encrypted message containing the first person's name and local time; the second person's computer then uses the symmetric key to decrypt the packet. The second computer also factors in time as part of the encryption/decryption process. (The fact that time is part of the encryption technology is why Win2K machines need to be time synchronized with a Simple Network Time Protocol—SNTP—service.) If the time in the message is close to the second computer's time, then the match is OK. Figure 5-1 shows this time calculation, where K_{JB} is the symmetric key shared by John and Bob.

Figure 5-1

Using the Kerberos symmetric key to decipher a message



—John Savill

Tip 43: Distributing a Shared Key


Q How's the shared key used by Kerberos distributed?

A The previous tip explained that two machines can use a shared symmetric key to communicate encrypted data. However, you might be wondering how the shared key gets distributed to these two machines (a client and a server) in the first place.

Kerberos uses a Key Distribution Center (KDC), which consists of a service that runs on all Windows 2000 domain controllers (DCs), as Figure 5-2 shows.

Figure 5-2

Viewing the KDC running as a service



The KDC generates the shared key for the client and server and sends it in an encrypted form to the client. The KDC responds to the client's request to talk to a server by sending both copies of the session key to the client (one for the client and one for the server—it's the same key, just packaged differently). The client's copy of the session key is encrypted with the key that the KDC shares with the client. The server's copy of the session key is embedded, along with authorization data for the client, in a data structure called a session ticket. The entire structure is then encrypted with the key that the KDC shares with the server. The session ticket—with the server's copy of the session key safely inside—becomes the client's responsibility to manage until it contacts the server.

The client uses its key shared with the KDC to extract the session key (the client can't, however, decrypt the session key because it doesn't know the key shared between the server and the KDC). The session key information is stored in a secure cache on the client in memory (never written to disk). When the client wants to communicate with the server, it sends its name and time encrypted in the shared key (which it extracted) to the server along with the server's session ticket. The server then decrypts the session ticket using the key shared with the KDC, extracts the session key, decrypts the client authenticator, and replies back to the client with the client's time encrypted with the session key.

All these steps mean that the server doesn't have to store session keys for clients; it's the client's responsibility to send the server's session ticket to the server as part of the communication. Also, session tickets are good for a defined period of time based on the Kerberos policy. This time period is typically 8 hours (a normal logon time) so the KDC isn't contacted every time the client wants to talk to a server; instead, the client has a cached copy of the session ticket that's good for the day.

—John Savill

Tip 44: Distributing a Long-Term Key

Q How's a long-term key between a client and the Key Distribution Center (KDC) distributed?

A The final piece of the Kerberos key exchange is the key used between a client (or server) and the KDC for distributing the session key used to encrypt information between a client and a server. For this step, the client and the KDC share a long-term key, and this long-term key is just a hashed version of the client user's password. The password is never sent over the network.

When John logs on, the Kerberos client on his workstation runs his password through a one-way hash algorithm to produce a cryptographic key. The client machine then sends this key to the KDC where the KDC extracts this hash from its record and checks to see whether the hash matches what the KDC thinks it ought to be. This only happens only on initial logon to the network.

At initial logon, the Kerberos client asks the KDC for a session key and ticket that it can use for further KDC communication during the logon session; this step avoids the need for the client to constantly generate and the KDC to constantly check the hashed password. The KDC replies with a copy of the session key it's generated for communication between the client and itself encrypted with the clients long-term key (hashed password). The client can decrypt using its cached long-term key.

The KDC also sends a copy of the session key encrypted with the KDC's own long-term key, which is known as the ticket-granting ticket (TGT) and the client sends this key in future communication with the KDC in the same way the client-server session ticket works. Only the KDC can decrypt the TGT because only the KDC knows the long-term key that was used to encrypt the session key. The TGT is valid only for the logon and discarded on logoff; hence, it's called a "logon session key."

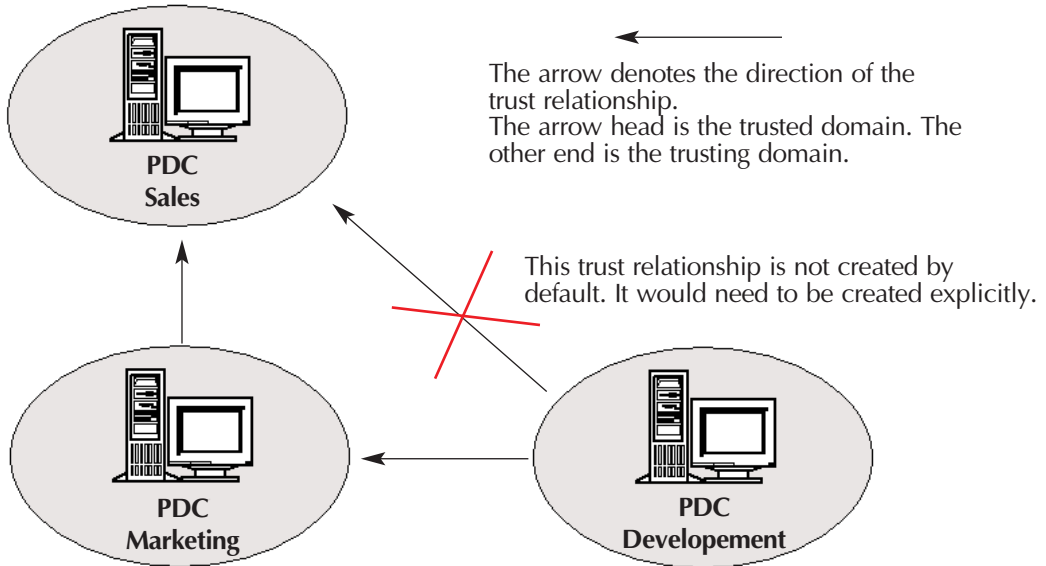
In future KDC communication, the client would send the normal requests encrypted with its session key shared with the KDC and the KDC's TGT from which the KDC can extract the session key needed to talk to the client. As a result, the KDC doesn't have to maintain a list of session keys for each client.

—John Savill

Tip 45: Defining a Kerberos Trust

Q What's a Kerberos trust?

A In Windows NT 4.0, trust relationships weren't transitive. So, for example, if domain2 (e.g., Marketing, in Figure 5-3) trusted domain1 (Sales), and domain3 (Development) trusted domain2 (Marketing), domain3 (Development) didn't trust domain1 (Sales).

Figure 5-3*Understanding non-transitive trusts in Windows NT 4.0*

In Windows 2000, the trust relationships that connect members of a tree or forest are two-way, transitive Kerberos trusts. Thus, all domains in a tree implicitly trust all other domains in the tree or forest. Because trusts occur automatically when a domain joins a tree, time-consuming trust administration is unnecessary.

Kerberos is Win2K's primary security protocol. Kerberos verifies a user's identity and a session's data integrity. Each domain controller (DC) has Kerberos services on it, and every Win2K workstation and server has a Kerberos client. A user's initial Kerberos authentication gives the user one logon session to enterprise resources. Kerberos isn't a Microsoft protocol but is based on MIT's Kerberos 5.0. For more information about Kerberos, see the Internet Engineering Task Force (IETF) Requests For Comments (RFC) 1510, The Kerberos Version 5 GSS-API Mechanism, at <http://www.ietf.org/rfc/rfc1510.txt?number=1510>>The%20Kerberos%20Network%20Authentication%20Service%20(V5), %20and%20RFC%201964,%20<A%20HREF=.

<http://www.winnetmag.com/Article/ArticleID/13379/13379.html>

—John Savill

Tip 46: Creating a Kerberos-Based Trust Between Domains

Q Why can't I create a Kerberos-based trust between two domains in different forests?

A When you manually create trusts, you can select one of two authentication protocols.

Kerberos—The Kerberos V5 authentication protocol is the default authentication service for Windows 2000. You use it to verify that a user/host is who it says it is. This protocol is used for trusts between domains in a tree and between the root domains in a forest.

Windows NT LAN Manager (NTLM)—The NTLM authentication protocol is the default for network authentication in Windows NT 4.0 and earlier, but Win2K still supports it (although not as the default). NTLM is a challenge/response authentication protocol.

A transitive Kerberos-based trust links domains *within* a forest. Thus, when you create a trust between two domains in different forests, you can select only NTLM because Kerberos isn't available for cross-forest trust relationships. This limitation isn't a Kerberos one, but a limitation of the Microsoft implementation. If you use a third-party Kerberos implementation (e.g., MIT), you can use Kerberos for cross-forest trusts.

—John Savill

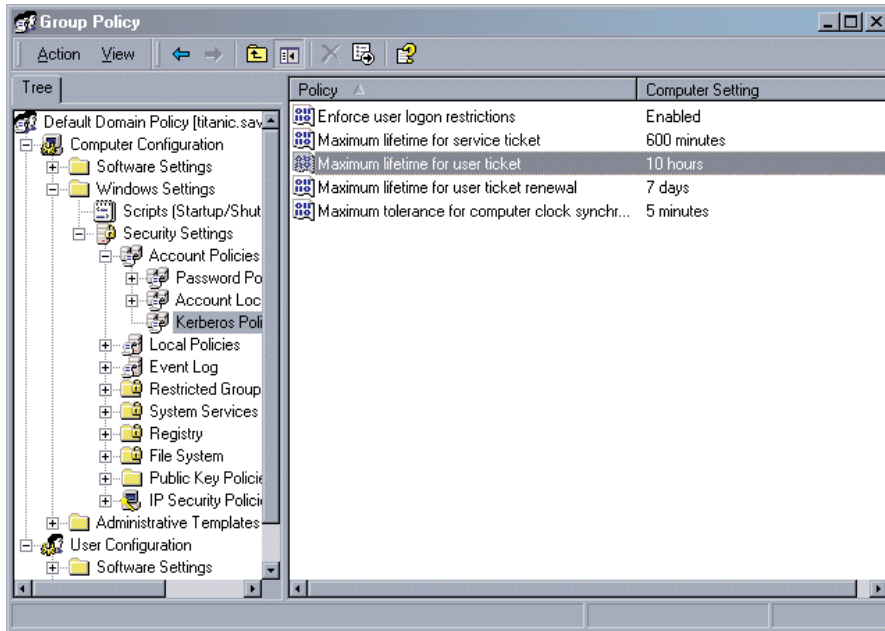
Tip 47: Changing the Ticket Lifetime Used by Kerberos

Q How can I change the ticket lifetime used by Kerberos?

A The default lifetime for a Kerberos ticket is defined by the Group Policy for the domain, which is 10 hours by default. You can change the default lifetime but 10 hours will typically suffice (unless people work very long days). To change the default lifetime, perform the following steps:

1. Start the Microsoft Management Console (MMC) Active Directory Users and Computers snap-in (go to Start, Programs, Administrative Tools, and click *Active Directory Users and Computers*).
2. Right-click the domain, then select Properties from the context menu.
3. Select the Group Policy tab.
4. Select the domain Group Policy Object (GPO), then click Edit.
5. Expand the Computer Configuration root then Windows Settings, Security Settings, Kerberos Policy, as Figure 5-4 shows.

Figure 5-4
Changing the Kerberos ticket lifetime



6. Double-click the time you wish to change, modify the time, then click OK.
7. Close Group Policy Editor (GPE).

To force the GPO change to take effect, from the command prompt type

```
secdit /refreshpolicy machine_policy /enforce
```

—John Savill

Tip 48: Cracking Kerberos Packets

Q Can you sniff Kerberos packets and crack them to obtain the user's password?

A Although stronger than Windows NT LAN Manager (NTLM), Kerberos is still based on user passwords. A weak user password remains vulnerable even if your Windows XP or Windows 2000 workstation uses Kerberos to authenticate to the domain controller (DC). Arne Vidstrom wrote a Kerberos sniffer and cracker, KerbCrack (<http://ntsecurity.nu/toolbox/kerbcrack>), that demonstrates this vulnerability.

You have a few options for protecting yourself from attackers who might sniff and crack NTLM or Kerberos authentication traffic either on your intranet or on the Internet. One option is to try to convince your users to select strong, hard-to-guess passwords, enforce minimum password lengths and password complexity, then back up those measures by periodically using a password cracker such as @stake's LC5 to audit password strength. However, this method is a lot of work and usually isn't successful because users resist selecting strong passwords. Some organizations try to secure their internal networks against password sniffing by implementing a fully switched network so that each computer receives only the packets destined for it. However, attackers can use Address Resolution Protocol (ARP) redirects to sniff across switches or can hack switches.

The best solution available is to eliminate NTLM by upgrading all computers to XP or Win2K, then eliminate Kerberos-associated password risks by implementing smart cards for interactive logon. Win2K has good Plug and Play (PnP) support for smart card readers and uses the PKINIT Kerberos extension, which replaces passwords with public/private keys for the initial ticket-granting ticket (TGT) that the workstation obtains when the user initially logs on.

—Randy Franklin Smith

Tip 49: Windows NT LAN Manager Versus Kerberos Use

Q I've read that Kerberos replaces the much weaker Windows NT LAN Manager (NTLM) authentication in Windows 2000 and later. Are there any circumstances under which Win2K still uses NTLM?

A Yes, Win2K still uses NTLM in certain situations. You should know the circumstances under which this occurs because NTLM is much more vulnerable to eavesdropping and subsequent cracking. For Win2K to use Kerberos when a user logs on, all computers involved—workstations, domain controllers (DCs), and servers—must be Win2K or later and members of the same domain or at least the same forest. In addition, the user account that's logging on must be an Active Directory (AD) user account, not an account in a computer's local SAM or an account from an NT domain.

When a user with an AD domain account logs on at an NT or Windows 9x workstation, NTLM will authenticate the logon rather than Kerberos because pre-Win2K versions of Windows don't support Kerberos. For the same reason, even when a user logs on with an AD domain account to a Win2K workstation but maps a drive to an NT server, NTLM will authenticate the logon rather than Kerberos. Also, when a user maps a drive to a Win2K server but uses a local account in that server's SAM, Win2K uses NTLM instead of Kerberos—even if the workstation and server are part of an AD domain.

—John Savill

Tip 50: Exploring Kerberos Ticket Lifetime

Q My domain controllers' (DCs') Security logs are recording frequent occurrences of event ID 677 (i.e., Service Ticket Request Failed) with failure code 32. The failure code 32s seem directly related to the User Ticket Lifetime and the Service Ticket Lifetime parameters. Increasing these parameters' times results in far fewer event ID 677s, but the events still

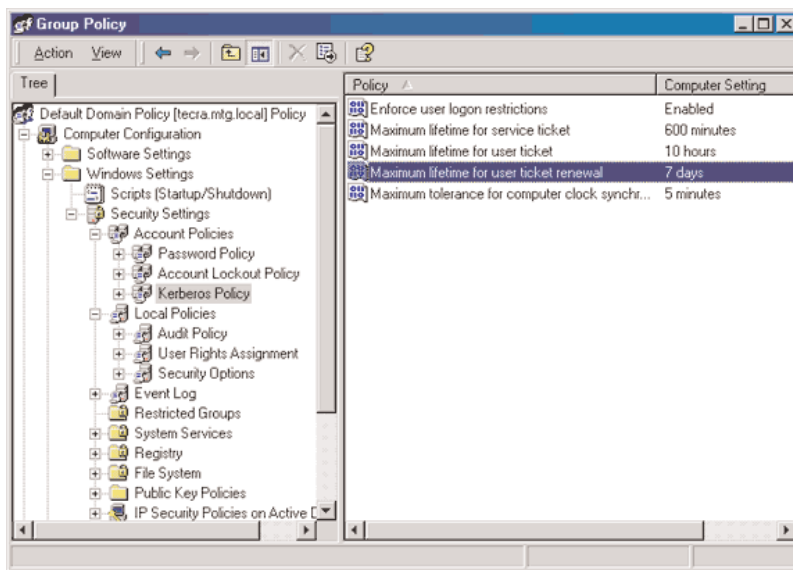
occur. The renewal process seems to cause an error even though renewal request succeeds. What's happening?

A You're correct about ticket lifetime. Whenever a user or client application needs to access a domain service (e.g., file sharing, Active Directory—AD), Windows 2000 obtains a Kerberos ticket that lets the client access the service. Kerberos tickets have a maximum lifetime specified in hours and a maximum renewal limit specified in days. The default maximum lifetime for user tickets and service tickets is 10 hours. The default maximum renewal for user tickets is 7 days, and service tickets have no maximum renewal policy.

Win2K uses user tickets when users access resources on the network. Win2K uses service tickets when a computer's service needs to access another computer's service. For example, local workstations must regularly check the DC for Group Policy changes. The maximum ticket-renewal limit is 7 days. Win2K automatically renews tickets when they expire, and eventually Win2K tries to renew a ticket beyond the ticket's renewal limit. The renewal then fails and generates event ID 677 with failure code 32. Failure code 32 is benign: Kerberos reissues a ticket despite the warning.

The failure codes you find in event ID 677 come directly from the Kerberos error codes in the Internet Engineering Task Force (IETF) Request for Comments (RFC) 1510. You can edit ticket lifetime and other Kerberos policies in Group Policy Objects (GPOs) under \computer configuration \windows settings\security settings\account policies\kerberos policy, as Figure 5-5 shows.

Figure 5-5
Editing the Kerberos policy settings



Don't extend ticket lifetime too much, however, because doing so gives an attacker more time to attempt to break the ticket—an unlikely, but not impossible, event.

—Randy Franklin Smith