Windows Server® 2008 Security Resource Kit

Jesper M. Johansson and MVPs with the Microsoft Security Team

To learn more about this book, visit Microsoft Learning at
http://www.microsoft.com/MSPress/books/11841.aspx

Microsoft® Press

9780735625044

Table of Contents

What do you think of this book? We want to hear from you!

Microsoft is interested in hearing your feedback so we can continually improve our books and learning resources for you. To participate in a brief online survey, please visit:

www.microsoft.com/learning/booksurvey/

What do you think of this book? We want to hear from you!

Microsoft is interested in hearing your feedback so we can continually improve our books and learning resources for you. To participate in a brief online survey, please visit:

www.microsoft.com/learning/booksurvey/

Chapter 4

Understanding User Account Control (UAC)

— Darren Canavor

With a shift in the way people use computers, such as performing banking transactions, making online purchases, and sharing and storing personal information, a new set of security threats emerged. Windows users were largely running with administrative privileges all the time. If the user accidentally installed malicious software (malware) onto such a computer, that malware—which had administrator access—could do anything. In Windows Vista and Windows Server 2008, the new User Account Control (UAC) feature is designed to apply the principle of "least privilege": Only give enough access to perform the task with as few disruptions as possible to the user experience. That includes all interactive users, with the exception of the built-in Administrator account. This may sound simple, but the challenge required a solution encompassing extensive changes to the core operating system, changes in industry perception of the standard user desktop and broad adoption of standard user best practices by the independent software vendor (ISV) community.

> **Note** Although UAC is available in Windows Server 2008, it is primarily considered a client feature. To a systems administrator, the impact of UAC on Windows Server 2008 focuses on using Group Policy to manage Windows Vista client UAC policies.

What Is User Account Control?

UAC can help prevent unauthorized changes to a computer by allowing the user to verify actions before they happen. When a user designated with elevated privilege logs on to Windows Vista and Windows Server 2008, two access tokens are issued: a full access token and a filtered standard user access token. The filtering process removes the administrative privileges, and disables the Administrative group Security Identifiers (SIDs), resulting in a filtered standard user access token. The standard user token is then used to start the Windows desktop (explorer.exe) and all subsequent child processes. Consequently, all applications run with the standard user token by default, and only when granted permission by an administrator will the application run with a full access token. Note that because applications inherit the privilege level of the parent process, if the parent process is running with a full access token, the new child process will inherit and run without prompting the administrator for permission. For example, if you launch a command prompt as an administrator, any process you launch from within the command prompt will run as an administrator.

On the CD   Elevating Explorer

By default Explorer.exe is designed not to be elevated. Consequently, if you right-click the binary and select Run As Administrator it will launch a new window, but in the same context as the original. On the companion CD, you will find a set of elevation tools, including a tool that puts an Elevate Explorer Here item on the right-click menu of any folder. Using that tool, you can launch an elevated Windows Explorer instance anywhere you wish.

How Token Filtering Works

When a user logs on to a Windows Vista or Windows Server 2008 computer, the operating system examines the Relative IDs (RIDs) and privileges of the user. The user will receive two tokens (filtered and full) if her account possesses any of the RIDs listed in Table 4-1 or any of the privileges listed in Table 4-2.

Table 4-1   UAC List of Restricted RIDs

Restricted RIDs	Description
DOMAIN_GROUP_RID_ADMINS	Administrative domain user account
DOMAIN_GROUP_RID_CONTROLLERS	Domain Controllers group
DOMAIN_GROUP_RID_CERT_ADMINS	Certificate Publishers group
DOMAIN_GROUP_RID_SCHEMA_ADMINS	Schema administrators group
DOMAIN_GROUP_RID_ENTERPRISE_ADMINS	Enterprise Administrators group
DOMAIN_GROUP_RID_POLICY_ADMINS	Policy Administrators group
DOMAIN_ALIAS_RID_ADMINS	Administrative local user account
DOMAIN_ALIAS_RID_POWER_USERS	Power Users group
DOMAIN_ALIAS_RID_ACCOUNT_OPS	Account Operators group, Server only

Table 4-1   UAC List of Restricted RIDs

Restricted RIDs	Description
DOMAIN_ALIAS_RID_SYSTEM_OPS	System Operators group, Server only
DOMAIN_ALIAS_RID_PRINT_OPS	Print Operators group, Server only
DOMAIN_ALIAS_RID_BACKUP_OPS	Backup Operators group
DOMAIN_ALIAS_RID_RAS_SERVERS	RAS and IAS servers group
DOMAIN_ALIAS_RID_PREW2KCOMPACCESS	Pre-Windows 2000 Compatibility Access group
DOMAIN_ALIAS_RID_NETWORK_CONFIGURATION_OPS	Network Configuration Operators group
DOMAIN_ALIAS_RID_CRYPTO_OPERATORS	Cryptographic Operators group

Table 4-2   UAC List of Restricted Windows Privileges

Restricted Windows Privileges	Description
SeCreateTokenPrivilege	Required to create a primary token
SeTcbPrivilege	Identifies holder as part of the trusted computing base
SeTakeOwnershipPrivilege	Take object ownership without being granted discretionary access
SeBackupPrivilege	Required to perform system-wide backup tasks
SeRestorePrivilege	Required to perform system-wide restore tasks
SeDebugPrivilege	Can debug the memory of a process owned by another account
SeImpersonatePrivilege	Required to impersonate a client after authentication
SeRelabelPrivilege	Required to modify an object's mandatory integrity level

The filtered standard user token will have all Windows privileges removed except the list of standard Windows privileges shown in Table 4-3.

Table 4-3   UAC List of Standard Windows Privileges

Standard Windows Privileges	Description
SeChangeNotifyPrivilege	Required to receive file or folder change notifications
SeShutdownPrivilege	Required to shut down a system remotely
SeUndockPrivilege	Required to undock a laptop
SeReserveProcessorPrivilege	Required to modify user processor privilege
SeTimeZonePrivilege	Required to adjust the computer's time zone

The filtered access token has all the RIDs from Table 4-1, if present, marked as USE_FOR_DENY_ONLY. It also has the privileges listed in Table 4-2 removed. The unmodified full administrator access token is linked to the filtered access token and is used when requests are made to launch applications with a full administrator access token.

You can find more information on RIDs in Chapter 1, "Subjects, Users, and Other Actors." You can find more information on Windows privileges in Chapter 3, "Objects: The Stuff You Want."

Components of UAC

UAC is primarily perceived to be the elevation prompt. However, although that part is the most visible, it is not the most important part of UAC. UAC, in fact, consists of a number of components, all of which contribute in some way to enabling more people to run as nonadministrators, which is the ultimate goal of UAC. This section discusses the various components of UAC, starting with the various types of elevation dialogs.

UAC Elevation User Experience

The most salient impact UAC has on user experience will be seen by users who are members of the local administrator group. Standard users also have the ability to perform administrative tasks without having to log off. The prompt for standard users is identical to the administrative prompt, except it requires password entry.

The Credential Prompt

On Windows Vista and Windows Server 2008, with the exception of the built-in administrators account, all users start applications without administrator-level privilege. When a given a task requires administrator privilege the interactive standard user will be presented with an elevation credential prompt, shown in Figure 4-1, requiring the entry of a valid user name and password of a user that is a member of the Local Administrators group.
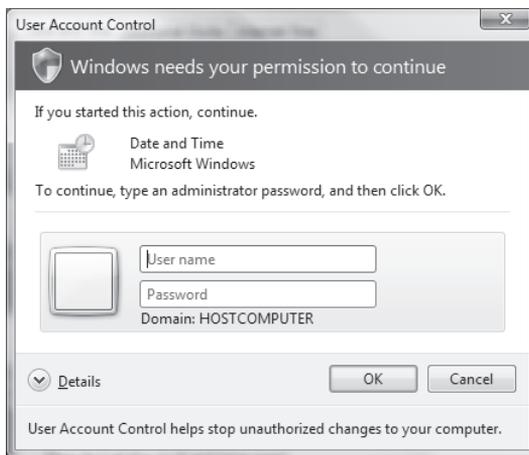
Figure 4-1   A standard user is presented with a prompt for credentials when attempting to perform an administrative action.

The Consent Prompt

By default the consent prompt, shown in Figure 4-2, is presented when a user who is a member of the local administrators group attempts to perform a task that requires administrator

privilege. This consent prompt is presented only to local administrators running in Admin Approval Mode.
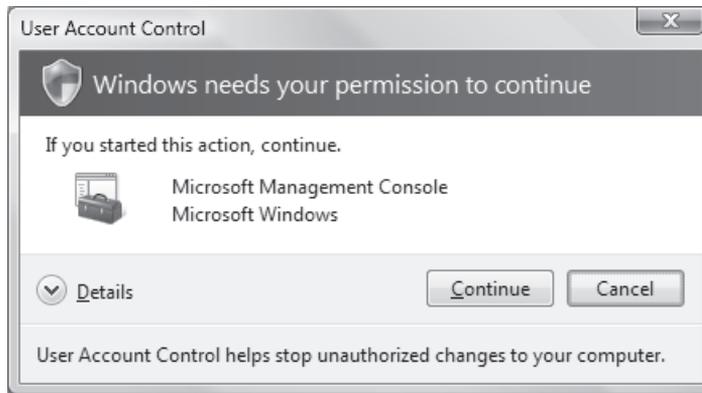
Figure 4-2   An administrator is presented with a prompt for consent when attempting to perform an administrative action.

To help users make informed decisions, the UAC elevation prompts are color-coded and use different text to indicate an application's potential security risk. For example, the color (four-color shield on blue-green bar) and text of Figure 4-2 indicate a Windows Vista or Windows 2008 application requiring administrative access, such as the Microsoft Management Console.

When an application attempts to run with an administrator's full access token, Windows Vista and Windows Server 2008 analyze the executable to determine its publisher and uses this information to determine the correct user experience.

Various alternative prompts are shown in Figures 4-3 through 4-5 and are distinguished by different colors and text. For example, in Figure 4-3 the color (yellow shield on gray bar) and text indicate that the application requiring administrative access is Authenticode signed and trusted by the local computer, such the Microsoft Firewall Client for ISA Server. In Figure 4-4 the color (yellow shield on yellow bar) and text indicate that the application requiring administrative access is unidentified and does not have a valid Authenticode signature from the publisher; therefore, take care before permitting the application to run. And in Figure 4-5 the color (red shield on red bar) and text indicate that the application requiring administrative access is from an explicitly blocked or untrusted publisher. An administrator can place the Publishers signing certificate in the local computer Untrusted certificate store to block a given publisher—this can also be set via Group Policy.

Note that UAC dialog boxes also change the displayed executable name and path details based on the trust level of the publisher's Authenticode signature. For example, in Figures 4-3 and 4-5, the user is trying to start the same application. The difference is that in Figure 4-3, the publisher is trusted, while in Figure 4-5 the publisher is explicitly blocked. When a publisher is trusted, not only does the dialog box color change, but the displayed text is also much friendlier.
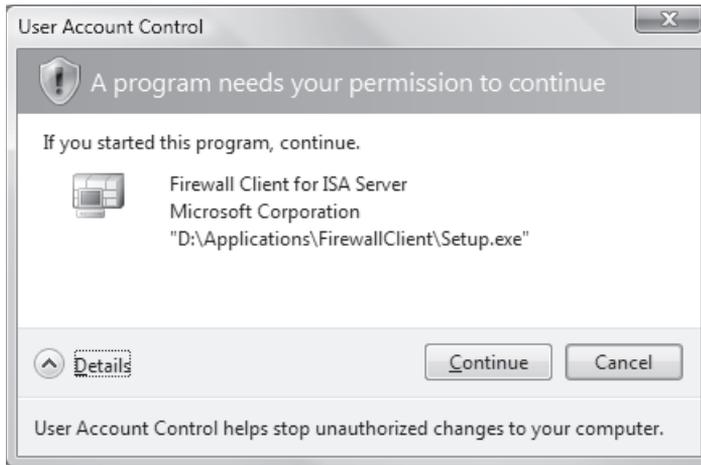
Figure 4-3   UAC prompt indicating that the application requiring administrative access is Authenticode-signed and trusted by the local computer.
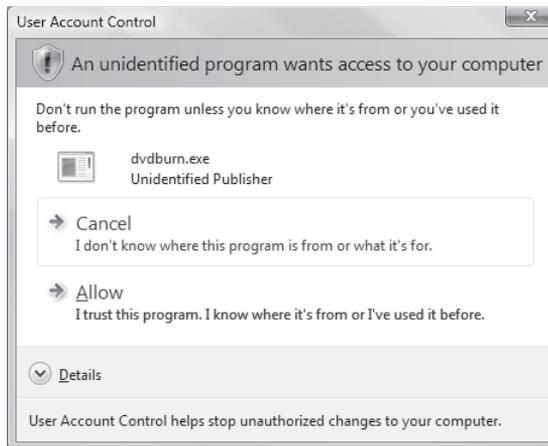
Figure 4-4   UAC prompt indicating that the application requiring administrative access is unidentified and does not have a valid Authenticode signature from the publisher.

In Windows Vista and Windows Server 2008 the shield icon shown in Figure 4-6 denotes that when a user clicks a shielded control or program, UAC will prompt for authorization before continuing.

Some Control Panel components, such as the Date and Time, contain both administrator and standard user operations. For example, standard users can view the clock and change the time zone, but a full administrator access token is required to change the local system time, as shown in Figure 4-7. One reason for this is that a user who changes the system time can reorder events in the event log or impact the ability for a computer to authenticate to a Windows domain.
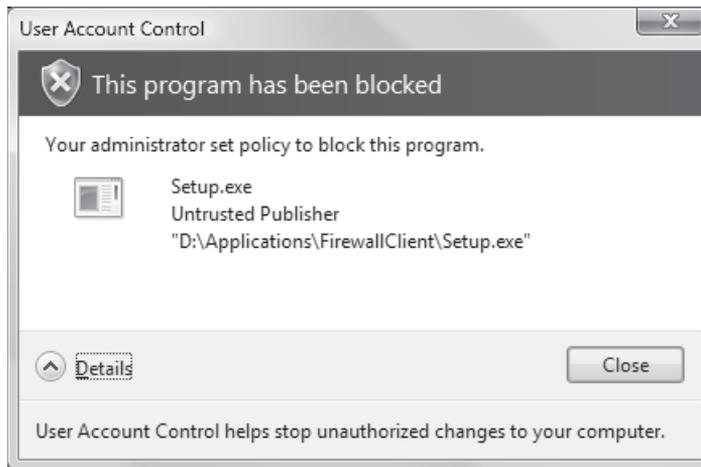
Figure 4-5   UAC prompt indicating that the application requiring administrative access is from an explicitly blocked or untrusted publisher.

Figure 4-6   The shield icon denotes an administrative action in Windows Vista and Windows 2008.
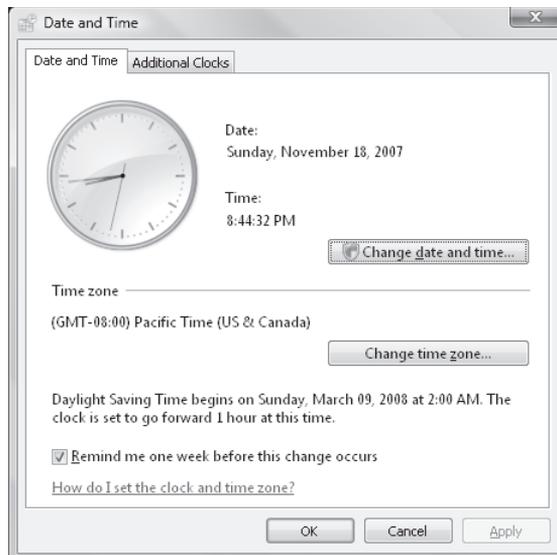
Figure 4-7   The Date And Time Control Panel utility is used to configure local computer time and time zone.

Application Information Service

The Application Information Service (AIS) is a new system service in Windows Vista and Windows Server 2008 that controls the launching of programs that require one or more elevated privileges, restricted rights, or privileged integrity levels to run. AIS is the component that actually launches these processes and attaches the right token to them. You could say that AIS is the heart of UAC. Note that AIS is disabled in Safe Mode; therefore, users who are members of the local administrator's group log on with their full administrative tokens. Windows took this approach because of the recovery and maintenance nature of Safe Mode scenarios.

File and Registry Virtualization

Windows Vista and Windows Server 2008 include file and registry virtualization, which is a new application compatibility technology to address issues encountered by applications that historically required an administrator's access token to run. Virtualization helps mitigate these applications without burdening the ISV to make changes. A large number of legacy applications that previously failed to run without the administrators access token now work on Windows Vista and Windows Server 2008, thanks to virtualization.

When a legacy application running with a filtered standard user access token attempts to write to a protected directory, such as Program Files, the application is given a virtualized view of the resource it is attempting to change. The virtualized copy is maintained under the user's profile (or registry). Each user has a completely separate copy of the virtualized file. This means that two users playing the same game on the same computer may not see the same list of high scores, because each user could have his or her own virtualized vision of the game's % PROGRAMFILES %\Game\highscores.txt file. Therefore, IT administrators must understand file and registry virtualization and may potentially need to implement custom virtualization settings within the enterprise to overcome application compatibility issues. The following section examines file and registry virtualization.

File Virtualization

File virtualization addresses the situation in which an application relies on creating or modifying files, such as a configuration file, in a protected location (%PROGRAMFILES%, %PROGRAMDATA%, or %SYSTEMROOT%) writeable only by administrators. Running such a program with a filtered standard user token may result in unexpected failures, or in some cases might be entirely blocked from running because of insufficient file or registry access.

When a program writes to a protected system location, the file virtualization filter driver (%SYSTEMROOT %\System32\Drivers\Luafv.sys) "traps" the operation and redirects it to a per-user location under the Virtual Store directory, located at %LOCALAPPDATA%\ VirtualStore. When the program later reads the file, Luafv.sys traps the operation and again redirects it to the user's Virtual Store. If the file is not found in the Virtual Store, Luafv.sys will query the nonvirtualized location. Because file virtualization happens automatically, the

program believes it was successful in writing to %PROGRAMFILES%*appName*. For security reasons file virtualization by default will not allow the redirection of known executable file types such as .exe, .dll, .sys, .bat, and .cmd. If, because of application compatibility constraints, the program needs to virtualize a .bat file, you can reconfigure the file virtualization filter to support this. The following examples demonstrate how to configure file virtualization.

> **Configuring file virtualization to improve application compatibility**   The FileList registry is not present by default and must be manually created to configure file virtualization.

Scenario: An enterprise relies on a legacy accounting application that writes a log file back to the application's restricted program folder. To enable virtualization on the accounting program's folder C:*appNameX*, create a new DWORD named **Exclude** with a value of 0 under the following registry key:

```
[HKLM\SYSTEM\CurrentControlSet\Services\luafv\Parameters\FileList\Device\
HarddiskVolume1 \appNameX]
```

Scenario: An enterprise forces all users to save their data to a specific location by locking down all user-writeable locations except the designated backup location. With virtualization enabled, a user can potentially store data in any virtualization-enabled location. To disable virtualization on a specific folder C:\Program Files*appNameY*, create a new DWORD named **Exclude** with a value of 1 under the following registry key:

```
[HKLM\SYSTEM\CurrentControlSet\Services\luafv\Parameters\FileList\Device\HarddiskVolume1
\Program Files\appNameY]
```

Scenario: An enterprise relies on a legacy accounting application that happens to write a .bat file back to the application's restricted program folder. To enable virtualization of .bat file extension types, create a new REG_MULTI_SZ named **ExcludedExtensionsRemove** with a value of bat under the following registry key:

```
[HKLM\SYSTEM\CurrentControlSet\Services\luafv\Parameters]
```

> **Note**   To expose virtual files and folders, browse to the virtualized file location using Windows Explorer and click Compatibility Files on the Explorer toolbar.

Registry Virtualization

Registry virtualization is similar to file virtualization but applies to registry keys under HKLM\SOFTWARE. This feature permits applications that rely on the ability to store configuration information in HKLM\SOFTWARE to continue to operate when running without administrative privilege. The keys and data are redirected to HKEY_CLASSES _ROOT\

VirtualStore\SOFTWARE. Note that the VirtualStore location is created on demand by the first application utilizing virtualization. As with file virtualization, each user has a virtualized copy of values that an application has stored in HKLM. If, because of application compatibility constraints, a program needs to configure registry virtualization, this is supported. The following examples demonstrate how to configure registry virtualization.

Configuring Registry Virtualization to Improve Application Compatibility

Scenario: An enterprise wants to prevent the virtualization of registry values under the key *DontVirtMe*. To do so, run the following command from an elevated command prompt:

```
Reg.exe flags HKLM\Software\DontVirtMe SET DON'T_VIRTUALIZE
```

Scenario: An enterprise wants to prevent the virtualization of all registry values and subkey values under the parent registry key *DontVirtMe*. To do so, run the following command from an elevated command prompt:

```
Reg.exe flags HKLM\Software\appName RECURSE_FLAG DONT_VIRTUALIZE
```

Although virtualization allows the overwhelming majority of pre-Windows Vista applications to run, it is a short-term fix rather than a long-term solution. In addition, some applications cannot be fixed, including applications that contain specific checks for user privileges. For example, many process-control applications check whether the user is an administrator, and exit if the user is not. You can get those applications to run on Windows Vista by attaching an application manifest that states the application needs to be run with administrative privileges and redeploy. Developers should modify all applications to comply with the Windows Vista and Windows Server 2008 Logo Program rather than relying on file and registry virtualization.

Manifests and Requested Execution Levels

Applications running on Windows Vista and Windows Server 2008 can use application manifests to describe or declare requirements to the operating system at run time. Administrative applications can declare their privilege requirements in the application manifest and the system will prompt the user for permission accordingly. Most pre-Windows Vista administrative applications, however, can run smoothly without modification even though they lack an entry in the application manifest. This is due to the vast array of Windows Vista and Windows Server 2008 application compatibility fixes, most of which depend on UAC being enabled. Application compatibility fixes enable applications to run that would normally fail if they ran without administrative access. For example, imagine a game that checks during start-up to see whether the user is a member of the local administrators group. Running with a filtered standard user access token, this check will fail—causing the application to fail. Using the application compatibility database, the operating system can discover that the application must run with a full token and prompt

the user accordingly or discover that the application runs fine without a full token and makes the application perceive it was started with a full token. These types of application compatibility fixes are called *shims*.

All Windows Vista and Windows Server 2008 logo-compliant applications must have a valid manifest with a defined requested execution level. The application uses the *requestedExecutionLevel* attribute to declare its access requirements. If the application requires administrative access, the application manifest specifies a requested execution level of *requireAdministrator*. This will ensure that the system identifies this program as an administrative application and provide the necessary elevation experience. Note that an application can also have mixed functionality—administrative and standard user—depending on the user. For example, the Microsoft Management Console (MMC) is marked *highestAvailable*. If a standard user runs the MMC, it will start with standard user privilege and will not prompt. If the user has a filtered access token, such as a local administrator or network operator, the operating system will prompt the user to launch MMC with the user's highest available privilege, allowing the administrator to have a different level of access than the network operator and the standard user.

Installer Detection Technology

Installation programs are applications designed to deploy software, and most write to system directories and machine registry keys. These protected system locations typically require administrator-level privilege, which means that standard users do not have sufficient access to install most programs. Windows Vista and Windows Server 2008 heuristically detect installation programs, updaters, and uninstall programs that require administrator access to run. Installer detection is a key component of the UAC design. It facilitates the correct elevation experience and prevents installations from being executed without the user's knowledge.

Installer detection only applies to the following:

- 32-bit executables
- Applications without a *requestedExecutionLevel*
- Interactive processes running as a standard user with UAC enabled

The operating system will heuristically determine whether an application is an installer. Heuristics are based on the following attributes:

- Keywords included in the filename, such as *install*, *setup*, *update*, and other language equivalents
- Keywords in the following Versioning Resource fields of the executable: Vendor, Company Name, Product Name, File Description, Original Filename, Internal Name, and Export Name

- Keywords in the side-by-side manifest that are embedded in the executable

- Keywords in specific *StringTable* entries that are linked in the executable

- Key attributes in the RC data that are linked in the executable

For example, if you have an application called setup.exe or install.exe, it will be detected as an installer and will automatically get a prompt. You can find general information and an overview of the Microsoft Windows Installer at MSDN: *http://go.microsoft.com/fwlink/ ?LinkId=30197.*

User Interface Privilege Isolation

User Interface Privilege Isolation (UIPI) is a new technology in Windows Vista and Windows Server 2008 to help isolate administrator-level processes from processes running with lower privileges on the same interactive desktop. UIPI prevents a lower-privilege application from using Windows messages to send input to a higher-privilege process. Sending input from one process to another allows a process to "inject" input into another process without the user providing consent.

UIPI defines a set of permitted Windows messaging interactions controlled by the highest of the different process levels. Higher privilege levels can send Windows messages to applications running at lower levels, but lower levels cannot send certain Windows messages to application windows running at higher levels. UIPI does not interfere or change the behavior of window messaging between applications at the same privilege level. UIPI comes into play for a user who is a member of the administrators group and chooses to run both administrator and standard user privileged applications on the same interactive desktop.

Secure Desktop Elevation Prompts

Credential and consent prompts are displayed on the secure desktop by default in Windows Vista and Windows Server 2008. Every application must run on a desktop, and each interactive user receives a desktop upon logon where all her applications run. The Secure Desktop is used by the operating system for services and sensitive user interfaces such as the log-on interface.

By presenting the elevation prompt on the secure desktop, the operating system guarantees that the information being presented cannot be tampered with. When an executable requests elevation, the user is switched from the user's interactive desktop to the secure desktop. The secure desktop renders a dimmed background of the user desktop and displays a highlighted elevation prompt. When the user clicks Continue or Cancel, the desktop automatically switches back to the user's interactive desktop. While malware can paint over the interactive desktop and present an imitation of the secure desktop (spoofing), authorizing consent does not allow the malware elevation. If UAC is configured to prompt for credentials, malware imitating the credential prompt may gather the user's credentials; however, the malware will be unable to use those credentials remotely to obtain administrator privilege. Somewhat

bizarrely, the malware will gain absolutely nothing from spoofing the admin approval mode dialog box. Malware cannot enter the user name and password into a valid UAC dialog box presented on the Secure Desktop, nor can it use runas.exe to invoke a process with elevated privilege or automate a legitimate UAC dialog box.

Using Remote Assistance

In Windows Vista and Windows Server 2008, a domain user can run as a standard user and have a centralized IT group provide all administration tasks. Microsoft provides both Remote Desktop (RD) and Remote Assistance (RA) access to computers for different administration purposes. RD sessions are useful when an administrator does not require end-user interaction but does require full control of the remote computer. RA is useful for diagnosing and trouble-shooting problems when the end user needs to demonstrate the problem to an IT expert. RA has been impacted by UAC; it is important that you understand how.

IT experts will experience two typical problems using RA. The first is that by default, the UAC prompts use the secure desktop and consequently are not available to the remote user. The second is if the UAC enterprise policy Behavior Of The Elevation Prompt For Standard Users is configured to Automatically Deny Elevation Requests, elevation is blocked entirely.

Windows Vista SP1 has a new UAC policy to address the challenge of the secure desktop prompting: User Account Control: UIAccess Applications To Prompt For Elevation Without Using The Secure Desktop. With this policy configured, AIS dynamically disables secure desktop prompting for UIAcess accessibility applications such as Remote Assistance and re-enables it once the program exits. For more details, see "What Is New in UAC in Windows 2008 and Windows Vista SP1" later in the chapter.

If the policy Behavior Of The Elevation Prompt For Standard Users is set to Automatically Deny Elevation Requests, the IT expert who connects using RA will be unable to launch an application with administrative privilege. To work around this issue, the IT expert can use runas.exeto launch a Command Prompt window using her own user name and password and then start a process that requires elevation. UAC will use the IT expert's UAC prompt policy.

The following procedure could be used by an IT expert for running the Registry Editor with administrator privilege:

1. Open a command prompt and type **runas /user:domain\ITExpert cmd.exe.**
2. In the new Command Prompt window that opens up, type **regedit.exe.**
3. Respond to the UAC elevation prompt.

UAC Remote Administrative Restrictions

When an administrator logs on to a Windows Vista or Windows Server 2008 computer remotely, using normal Windows networking, he logs on in Admin Approval mode, just as if he were logging on locally. To augment this behavior, UAC restricts remote administration to

prevent admin loopback attacks and help protect against local malicious software running remotely with administrative privilege. For example, admin loopback would occur when a user logs on with a filtered access token and then malware simply performs a **net use \\127.0.0.1\c$** to obtain administrative access to the file system. When UAC remote restrictions are enabled, the loopback would also obtain a filtered access token and not full administrative access. This behavior works differently for different types of user accounts, as described in the following sections.

Local User Accounts

Imagine that a user who is local to the server and a member of the local Administrators group on the server establishes a remote connection by **net use * ***server******share**. In this scenario, the token used for that user on the server will not be a full administrative token as in previous versions of Windows. The user has no elevation potential on the remote computer and cannot perform administrative tasks. If the user wants to administer the workstation with a local account, she must interactively log on to the remote computer by Remote Assistance or Remote Desktop if available.

Domain User Accounts

When a user with a domain user account logs on to a computer remotely, and he is a member of the local Administrators group, the domain user will run with a full administrator access token on the remote computer and UAC will not be in effect.

Managing UAC Remote Restrictions

To disable UAC remote restrictions for local accounts and obtain Windows XP and Windows 2003 parity, create a DWORD named **LocalAccountTokenFilterPolicy** with a value of 1 under the following registry key:

```
HKLM\SOFTWARE\Microsoft\Windows\CurrentVersion\Policies\system
```

Mapping Network Drives When Running in Admin Approval Mode

When an administrator in Admin Approval mode maps a network share, that share is only associated to the current log-on session for the current process access token. This means that if a user running a command prompt (cmd.exe) with a filtered access token explicitly maps a network share, that network share would not be exposed to any elevated cmd.exe instances running with a full administrator access token. Note that only in the case of UNC paths will the sessions be automatically linked by the system.

You can configure a registry value to share network connections between processes started with the filtered access token and full access token for a member of the Administrators group only. When you enable this registry setting, if a network resource is mapped to an access token, the LSA checks whether another access token is associated with the current user

session. If the LSA determines that there is a linked access token, it adds the network share to the linked location.

To enable a linked network drive, create a DWORD named **EnableLinkedConnections** with a value of 1 under the following registry key:

```
HKLM\SOFTWARE\Microsoft\Windows\CurrentVersion\Policies\System
```

Direct from the Field: Which Accounts Are Accepted for Elevation

Recently I was asked to troubleshoot some elevation problems for a friend of mine. She was unable to elevate to change some networking parameters on her laptop. The laptop was domain-joined, but the DC was unavailable at the time. After a few minutes of troubleshooting, I wrote up the following scenario, which I think helps highlight how UAC is not always intuitive, as well as how it interoperates with other features of Windows:

- The computer is called Denise-PC.

- The computer is joined to example.com.

- The DC for example.com is offline—in other words, Denise-PC is roaming.

- She has only previously logged on to Denise-PC using EXAMPLE\Denise.

- EXAMPLE\Denise is a member of BUILTIN\Users.

- BUILTIN\Administrators on Denise-PC contains BUILTIN\Administrator and DENISE-PC\Denise.

- When she attempts an administrative action she gets an elevation prompt asking for an admin account.

We have several options for how to elevate:

1. Attempt to elevate to BUILTIN\Administrator.

2. Attempt to elevate to EXAMPLE\Denise.

3. Attempt to elevate to EXAMPLE\Administrator.

4. Attempt to elevate to EXAMPLE\Foo, where Foo is a member of EXAMPLE\Domain Admins.

5. Attempt to elevate to DENISE-PC\Denise.

Option 1 will fail because BUILTIN\Administrator is disabled by default in Windows Vista as long as there is another local admin account. Because DENISE-PC\Denise is a local admin, and it is enabled, BUILTIN\Administrator is not available for use.

Option 2 will fail as well. EXAMPLE\Denise is only a member of users. It is not an admin and therefore you cannot elevate to it.

Option 3 will fail because although EXAMPLE\Administrator is a member (indirectly) of BUILTIN\Administrators, it has never logged on to Denise-PC. Because the computer is offline, authentication of domain accounts has to happen against the password verifier. (See Chapter 2, "Authenticators and Authentication Protocols," for information on cached credentials.) Cached credentials exist only for accounts that have previously logged on interactively; therefore, we have nothing to verify EXAMPLE\Administrator against. It should also be pointed out that elevating to a domain administrator on a member workstation would be an extraordinarily bad idea. For more information on why, see Chapter 13, "Securing the Network."

Option 4 fails for the same reason as Option 3.

Option 5 will succeed. DENISE-PC\Denise is a local account. Therefore, no cached credentials are necessary. It is a member of BUILTIN\Administrators, so it is legal to elevate to this account. And it is not disabled, so it can be actively used to log on with.

I have found this write-up very helpful in explaining to people which accounts can be used for UAC elevation, as well as the relationship between domain accounts, password verifiers, and UAC.

Jesper M. Johansson
Windows Security MVP

Application Elevations Blocked at Logon

Windows Vista and Windows Server 2008 block administrative applications that try to start in the user's log-on path. Many ISVs place programs in the user's log-on launch path to ensure that they run each time the user logs on. While this solution may be convenient, it often results in application compatibility problems when the user logging on is not an administrator but the application requires him to be. This behavior is also convenient for malware, which can simply place itself in the user's log-on launch location. From that point forward, every time the user logs on the malware runs silently with administrator-level access and without the user's consent. To block this behavior, Windows Vista and Windows Server 2008 create a workflow to help the user manage the blocked list of programs. An elevation balloon notifies the user, as shown in Figure 4-8, and the tray icon allows the user to run the blocked program or enter the management UI, as shown in Figure 4-9.

With UAC, applications that require administrator-level privileges to run are blocked when launched from the following locations:

- Per-User Startup Folder   %USERPROFILE%\Start Menu\Programs\Startup.

- Per-Machine Startup Folder   %ALLUSERSPROFILE%\Start Menu\Programs\Startup.
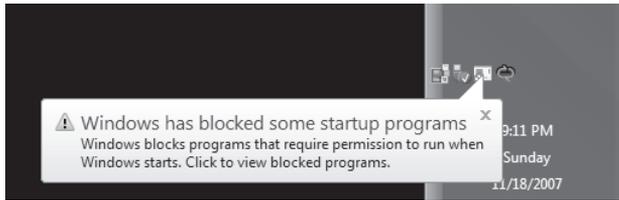
Figure 4-8   Blocked application balloon: Windows has blocked some start-up programs.
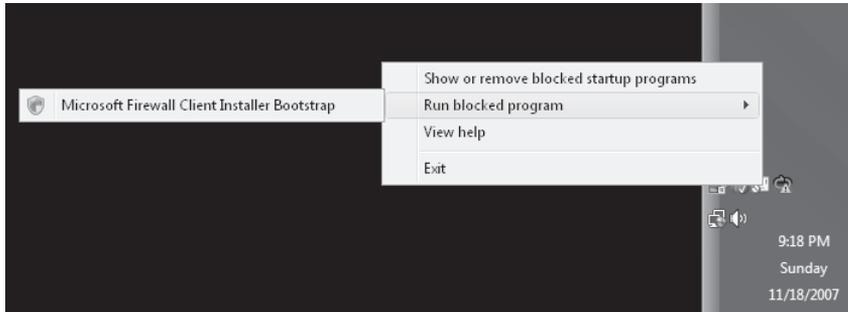
Figure 4-9   Blocked application tray icon: Run Blocked Program/Show Or Remove Blocked Start-up Programs.

- Per-User RUN Key   HKEY_USERS*\Software\Microsoft\Windows \CurrentVersion\Run

- Per-Machine RUN Key   HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft \Windows\CurrentVersion\Run.

Note   In the preceding section, the asterisk ("*") denotes all user security identifiers (SIDs) including the .Default SID.

It is important to note that enterprise Group Policy supports a user log-on script that will use the currently logged-on user's highest available access token; if the user is a member of the local administrators group, the script will elevate without prompting the Administrator in Admin Approval mode.

Configuring Pre-Windows Vista Applications for Compatibility with UAC

The final and most important step in configuring UAC is ensuring that your software is either designed to be UAC compliant by following logo requirements or has been configured to run with Windows Vista or Windows Server 2008.

For new applications that are Windows Vista and Windows Server 2008 logo-compliant, the application must either run with a standard user privilege or, in the case of an administrative application, be marked with an application manifest entry. For more information, visit the Microsoft Windows logo home page at *http://www.microsoft.com/ whdc/winlogo/hwrequirements.mspx.*

During the deployment of Windows Vista and Windows Server 2008, IT departments may discover some existing line-of-business (LOB) applications that will not function properly. In most cases, the problem is due to application incompatibility with the enhancements incorporated in the new operating systems. Microsoft provides an Application Compatibility Toolkit that assists in identifying the compatibility problems and aids in the creation of application compatibility fixes or shims. Some programs may need to perform administrative operations. For this to work correctly on Windows Vista and Windows Server 2008 under UAC, the program needs declare this to the operating system so that users will be prompted for approval before the application can run with a full administrator access token. The Application Compatibility Toolkit 5.0 with the Standard User Analyzer provides the means to test, build, and install the application compatibility database entries, which facilitate the requested execution level marking mechanism.

For information about application compatibility and the Application Compatibility Toolkit 5.0 featuring the Standard User Analyzer, visit TechNet at *http://go.microsoft.com/fwlink /?LinkId=23302.*

UAC Group Policy Settings

The following section explores each of the eleven UAC group policies supported on Windows Vista and Windows Server 2008. These settings can be applied locally using the Local Security Policy editor or across an enterprise by using Group Policy.

UAC Policy Settings Found Under Security Options

You can find the following nine UAC settings in the Group Policy Editor or the Local Security Policy editor under: Local Computer Policy\Computer Configuration\Windows Settings \Security Settings\Local Policies\Security Options.

User Account Control: Admin Approval Mode for the Built-in Administrator Account

This security setting controls the behavior of the Built-in Administrator (BA) account. If you use the BA account for daily administrative tasks, you may consider disabling this setting. However, if you do so, you also lose Internet Explorer Protected Mode. All applications will run as a full administrator. By default this policy is enabled.

User Account Control: Behavior of the Elevation Prompt for Administrators in Admin Approval Mode

If an operation requires administrator privilege to start, this policy will control the UAC prompt experience for administrators in Admin Approval Mode. Although the default consent configuration is convenient, enforcing credentials may be desirable. For example, if a parent and child share the same user account, the child will be unable to perform elevated tasks without knowledge of the password. Also, in some cases administrators may want to disable elevation prompting without disabling UAC and therefore can set the elevation prompt to Silent. This retains Internet Explorer Protected Mode, but removes the elevation prompts. By default this policy is Prompt For Consent.

User Account Control: Behavior of the Elevation Prompt for Standard Users

UAC provides an in-context elevation prompt experience, and if the user can provide a valid administrator user name and password, the elevated operation will succeed. For enterprises that do not want their users to have the opportunity to elevate, you can set this policy to automatically deny all elevation requests. By default this setting is Prompt For Credentials.

User Account Control: Detect Application Installations and Prompt for Elevation

This setting enables or disables application installer detection. It is best to leave this setting enabled, which is the default.

User Account Control: Only Elevate Executables That Are Signed and Validated

This setting will enforce Authenticode signature validation on any interactive applications requesting elevation. If an enterprise runs only Authenticode-signed programs, this setting can increase security by controlling which application publishers are allowed to run with elevated privileges. However, most users would experience significant application compatibility problems if they tried to use this setting, which is why it is disabled by default.

User Account Control: Only Elevate UIAccess Applications That Are Installed in Secure Locations

UIAccess applications are most often accessibility programs that need to interact directly with the Windows UAC elevation dialogs. Windows Vista and Windows Server 2008 UAC elevation dialogs are protected with a high integrity level. For UIAccess applications to interact they must declare this requirement in the application manifest. When the program

starts, it receives a special integrity level permitting interaction. Because UIAccess applications are powerful, this setting enforces that such programs be started only from a secure directory file path. UIAccess applications must also have a valid and trusted Authenticode signature. By default this setting is enabled.

User Account Control: Run All Users, Including Administrators, as Standard Users

This is the UAC on/off switch. Don't disable UAC! If UAC is disabled, all related features also become disabled. File and registry virtualization no longer function and all virtualized data appear lost to the user. Users who were running in Admin Approval Mode now log on with full administrative rights, and all applications run with administrator privilege, silently! Application compatibility shims designed to increase compatibility with pre-Windows Vista applications are also disabled. Internet Explorer's Protected Mode is disabled, forcing Internet Explorer to run with administrative privilege. Have we convinced you to leave UAC on? By default this setting is enabled.

User Account Control: Switch to the Secure Desktop When Prompting for Elevation

This setting determines whether elevation requests are presented on the interactive user's desktop or on the secure desktop. The secure desktop prevents output spoofing, which means that whatever is presented on the secure desktop cannot be tampered with. UAC dialog boxes on the interactive user's desktop can be spoofed and therefore are less secure than those presented on the secure desktop. By default this setting is enabled.

User Account Control: Virtualize File and Registry Write Failures to Per-User Locations

This setting enables or disables the redirection of write failures for the file system and registry. Disable this feature if you use only Windows Vista or Windows Server 2008 logo-compliant software. If you require custom virtualization settings, see "Configuring Registry Virtualization to Improve Application Compatibility" earlier in the chapter. By default this setting is enabled.

Related UAC policies

Windows Vista and Windows 2008 also have two complementary policy settings: the Require Trusted Path For Credential Entry and Enumerate Local Administrator Accounts On Elevation settings. You can find both settings in the Group Policy Editor under: Local Computer Policy\Computer Configuration\Administrative Templates\Windows Components\ Credential User Interface.

Require Trusted Path for Credential Entry

This setting controls whether the user must enter Windows credentials using a trusted path. The trusted path is a secure key sequence–sometimes referred to as a Secure Attention Sequence (SAS)–which prevents malware from stealing your Windows credentials. When a standard user tries to perform a task requiring administrator privilege, the system forces the user to enter Ctrl+Alt+Delete before being redirected to the secure desktop to enter a valid administrator user name and password to complete the operation. This trusted path credential workflow prevents input spoofing and output spoofing, making this the most secure Windows credential input configuration. By default this setting is disabled.

Enumerate Administrator Accounts on Elevation

This setting enables the automatic enumeration of local administrator accounts in the UAC credential UI, as shown in Figure 4-10. Note that in some domain-joined environments that encounter networking connectivity issues, this setting can cause unexpected delays when enumerating the local administrator accounts. By default this setting is disabled.
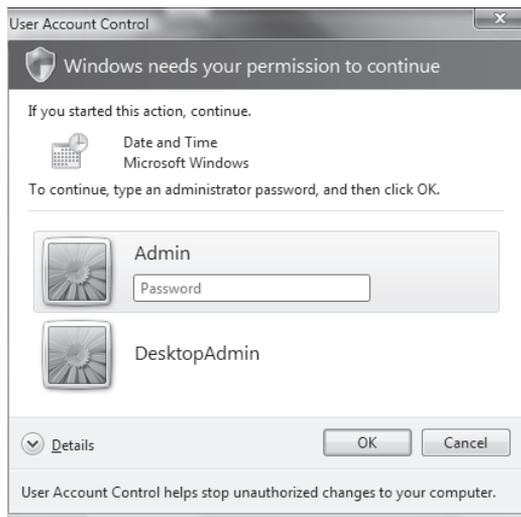
Figure 4-10   Automatic local administrator account enumeration in the UAC credential dialog box.

What's New in UAC in Windows Server 2008 and Windows Vista SP1

UAC underwent only a few small changes in Windows Server 2008 and Windows Vista SP1 when compared to the original release in Windows Vista. The following sections summarize the UAC changes.

New Group Policy Setting: UIAccess Applications to Prompt for Elevation without Using the Secure Desktop

This setting enables UIAccess applications such as Remote Assistance to request the disabling of secure desktop prompting. When the UIAccess application is complete, the secure desktop prompting is automatically enabled, thus removing the necessity for the end user to allow the desktop admin elevation access. (See Figure 4-11.) As discussed earlier in this chapter, this is a convenient setting for those enterprises that rely on Remote Assistance to provide end user desktop help desk support. By default this setting is disabled.
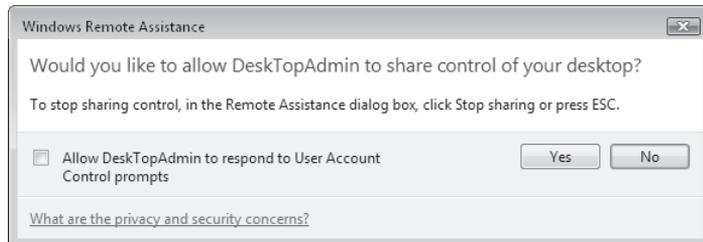
Figure 4-11   Windows Remote Assistance: Allow Helpdesk To Respond To User Account Control Prompts.

UAC Prompt Reduction When Performing File Operations in Windows Explorer

When a user creates a new folder in a protected location, the user will be prompted only once to create and name the folder. This was a two-prompt scenario in Windows Vista RTM.

More Than 40 Additional UAC-Related Application Compatibility Shims

The UAC team in conjunction with the Application compatibility team produced over 40 new application shims to help increase Windows Vista and Windows 2008 compatibility.

UAC Best Practices

Managing UAC is not as hard as it seems. How you deploy in an organization depends largely on your organization's security needs and tolerance to implement the required policies to meet those needs. The following solutions are presented in reverse order of preference (good, better, best) with respect to security value.

Good Practice

Run users in Admin Approval Mode. If an administrative user requires elevated privileges, the enterprise UAC policy should enforce that the user enters a valid administrator user name and

password instead of simply clicking the Consent dialog box. This configuration prevents unauthorized elevations on the off-chance that a user leaves his workstation unattended. To improve security you could also require the Ctrl+Alt+Delete key sequence for any elevation to complete. This makes entering administrative credentials far more secure.

Better Practice

Enforce that all users who require administrator privilege have two accounts: one standard user account for day-to-day activities such as reading e-mail and one for the occasional administrative operation. The standard user can log on and when needed can elevate using a UAC credential prompt. This is not the best solution because now the user is running both standard user and administrator-privileged applications in the same interactive session. To increase security an enterprise can enforce that the user must use Fast User Switching (FUS) anytime she needs to perform an elevated operation. Although FUS is more secure, it does have user experience drawbacks. To improve security you could also require the Ctrl+Alt+Delete key sequence as with the previous option.

Best Practice

Run all users as standard users. The IT department must then assume that standard users will generally not be able to install applications and therefore must deploy software on their behalf. Windows provides an installation service to do this called the Microsoft Software Installer (MSI) Service. In addition, the Group Policy Software Installation (GPSI) extension allows applications to be distributed to a user's computer without any user interaction required. See the Group Policy Software Installation Extension documentation at *http://go.microsoft.com/fwlink/?LinkId=71356* for more information.

Summary

UAC is probably the most talked-about feature in Windows Vista. It is even the subject of advertisements from rival software vendors. It is hard to say whether to be flattered or annoyed that Microsoft's competitors are now advertising their products as more desirable because Windows is too secure. Regardless, UAC is a critical step for Windows. The status quo, where users run as administrators to get normal tasks done, is unacceptable and has led to a malware pandemic. Only by helping users run as nonadministrators can we ever hope to stem the flood of malware and reduce desktop total cost of ownership (TCO).

The future is one where users only use administrative privileges where necessary. UAC is a step in that direction, but it will only work if people use it, and if they demand that their ISVs get software that works as a standard user. You can do your part in protecting the IT ecosystem by using UAC, and by buying software that works with it and rejecting software that does not.

Additional Resources

- Microsoft Corporation (2006). "The Windows Vista and Windows Server 2008 Developer Story: Windows Vista Development Requirements for User Account Control (UAC)," at *http://msdn2.microsoft.com/en-us/library/aa905330.aspx.*

- Mark Russinovich (2007). "Inside Windows Vista User Account Control," at *http://www.microsoft.com/technet/technetmag/issues/2007/06/UAC/.*

- Raymond Chen (2006). "An Administrator Is Not the Administrator," at *http://www.microsoft.com/technet/technetmag/issues/2006/03/WindowsConfidential/ ?related= /technet/technetmag/issues/2006/03/WindowsConfidential.*

- Wole Moses (2007). "Services Hardening in Windows Vista," at *http:// www.microsoft.com/technet/technetmag/issues/2007/01/SecurityWatch/?related= /technet/technetmag/issues/2007/01/SecurityWatch.*

Chapter 14

Securing the Branch Office

— Byron Hynes

In this chapter, we will examine the challenges of a branch office, and some solutions, techniques, and tools to address them. We will start by discussing the characteristics of a branch office, and why they present unique challenges, and then we will look at some of the common infrastructure optimization scenarios that can apply to branch offices. We will go through the features of Windows Server 2008 that support branch offices, and identify those that can particularly help secure branches, focusing on the role of Read Only Domain Controllers (RODCs) and BitLocker Drive Encryption.

An Introduction to Branch Office Issues

IT security often invokes the image of a well-maintained data center, where IT professionals are constantly striving to improve policies, procedures, settings, and software to protect the physical and knowledge assets in their control. Only a few people can enter the inner sanctum–by "carding in"–to find orderly arrangements of neatly placed computers, efficiently humming along while being securely watched.

Many of our readers–you–strive to build this sort of environment. But, did you realize that about one-third of IT budgets are spent on branch offices, or remote locations? Between one-quarter and one-third of servers running Windows Server software are installed in what we call a branch office.

Branch offices are found in many diverse environments. I worked with a customer in the fast-food industry who had a Windows server in every one of thousands of restaurants, nationwide. A colleague of mine did an early Active Directory design and deployment project with domain controllers in each one of 800 branches of a major European bank. Branch offices can include retail stores, medical facilities, hotels–perhaps even the check-in counter

where you drop off your child for daycare, or the elementary school office where the secretary enters daily attendance information.

Why Do Branch Offices Matter?

Branch offices are not going away anytime soon. Quite the opposite, actually. Branches are where the vast majority of business is won or lost for a corporation; branches are where services are delivered for nonprofits and government agencies. In addition to online services, branches are where source data is collected and interpreted, and where more savvy consumers are demanding that decisions can be made by local representatives.

Unfortunately, a number of characteristics can be limiting or even negative for the branch office.

What Is Different in a Branch Office?

Branch offices come in many different configurations, but most share the following traits:

- **Remotely connected over a WAN**   While network speeds continue to increase, most branch offices depend on an externally controlled network, such as a leased line, frame relay connection, or shared fiber circuit. Some branch offices are connected only over the public Internet, using a virtual private network (VPN). Remember that when you depend on a third party, such as a telephone or network company, you have no control over what happens to your data on that network.

- **Reduced physical security**   Servers in branch offices are rarely accompanied by the same level of physical security found in a corporate data center. I hope that your server is not stuck under a receptionist's desk in a public area, or sitting on a shelf in a restroom, but I have seen both.

- **Lack of local skilled IT personnel**   In a branch office, particularly a small one, you often find that an employee with little or no formal IT skill "inherits" the care and feeding of your server in the branch. This might be the business manager (or highest-ranking employee on site), the local tech enthusiast, or even just the staff member who is "always there" and never in a meeting.

- **Distance**   As well as not being staffed with IT professionals, a branch office is more often than not geographically far from IT support. It may be impossible to get one of your own analysts or technicians on-site in less than a day or two.

- **Reduced infrastructure**   Unlike the data center, with its precise climate control, redundant power supplies, and multiple communication paths, branch offices may not be engineered to the same level of fault tolerance—both in terms of IT systems and in terms of civil infrastructure, such as power and communications.

Branch offices face growing pressures, including cost and performance pressures, pressure to remain secure and in compliance with growing laws and regulations, and the pressure to be agile and responsive to business needs. For example, branch offices face the high costs of

maintaining complex, disparate, and often nonintegrated hardware sets, coupled with the high cost of bringing in IT support when needed, either from the data center or perhaps by subcontracting expensive local support agencies. Of course, branch office managers—and certainly the home office management team—are aware of the threats posed by data breaches, attacks against the system, and the difficulty of recovering from a disaster and protecting that locally collected data. (After all, why go after Jesper's well-hardened servers—you know, the ones he keeps in his data center and has guarded by men with guns, 24 hours a day—if a much less secure server is sitting in the local branch.)

But even with all of that, branch offices are expected to have and use more and more information as a "point of service," where data is collected, analyzed, and used to deliver immediate results, all while being integrated across the company. I remember opening a bank account as a teenager where the ledger was still kept by hand—and of course, that account could only be accessed at that particular branch. I do not expect that many banks would be able to compete using that kind of operation today! (And, just so you do not think I am too old, I should add that I walked down the street to a competitor with computers and multi-branch services.)

On the whole, a branch office represents a simple challenge for the consummate IT professional. All you have to do is secure remote data and access; protect systems; provide secure connectivity with reduced on-site support, simple backups, and efficient use of hardware and bandwidth—while also improving your infrastructure, delivering new capabilities quickly, and establishing a platform for the future and maintaining compliance. (Whew!)

IDC describes it this way:

> *Branch locations often are looking for high levels of flexibility and autonomy in implementing IT solutions. However, expanding corporate requirements call for centralized management, enhanced security, and regulatory compliance.*

Source: **IDC White Paper, Addressing Operational Efficiencies in Branch Office, May 2006,** located at *http://download.microsoft.com/download/6/d/0/6d032455-bf07-42ac-b006-ee0e4c8ab606/idc_branch_whitepaper.pdf*

Building Branch Offices

A number of design scenarios can address the branch office situation. Basically, they range on a continuum from one or more gargantuan servers in the branch office to "no servers in the branch office, do it all remotely." As you can imagine, each of these possible solutions has strengths and weaknesses.

To leave myself some room in this chapter to talk about securing Windows Server 2008 in the branch office, I am only going to present a brief overview of design approaches. Some of these approaches are very common; others are only occasionally used. Some are becoming more popular as lower costs and new technologies become more commonplace. And many branch offices have a combination of elements from more than one of the following options.

■ **Virtualized server in the branch office**   A server capable of running a virtual server image for each workload provides maximum flexibility and configurability—and agility—with a solution that allows all workloads and services to be run off of a single, physical server computer. This option is best where there are numerous stand-alone servers that could be consolidated into a single unit. Services are run in-house in order to allow 100 percent uptime and functionality even when the wide area network (WAN) or VPN is offline. By managing only one server, you can gain some centralized management and data backup.

■ **Single server**   Using one server offers the benefits of a server in the branch office, but without the cost and complexity of supporting virtualization and server consolidation. In this case, services that must be constantly available or that require the lowest possible latency can be run from within the remote office, while others are centralized to the data center. You must strike a balance between reliance on WAN uptime, remote management, centralized backup, and local responsiveness.

■ **Branch office appliances**   Branch office appliances strive to provide a complete solution that is easy to deploy and configure out-of-the-box. These appliances typically provide WAN acceleration and HTTP compression, making it easier to centralize workloads such as storage and databases. The branch office appliance may come preloaded with some line-of-business (LOB) applications or other services to allow the branch to continue to operate during a WAN outage. Appliances can be designed to connect to a WAN or to use the Internet as VPN appliances.

■ **Secure centralization**   In this sort of design, the server is removed from the branch location and all workloads are centralized. Application virtualization (such as the sort provided by Microsoft SoftGrid Application Virtualization) is used to allow LOB applications to be run from the data center but executed on the local desktop, reducing the burden of deploying and managing updates to those applications that are virtualized. A security appliance, intelligent firewall, or gateway enables secure connectivity (including packet filtering and VPN access) to the data center.

■ **No infrastructure**   In this case, the branch office simply has no servers or significant networking hardware. Everyone relies on network connectivity to the data center to enable all remote office functionality. Essentially, everyone is a network client. Remote office employees either exist on the same corporate network as the head office using a WAN, or they connect via VPN to the data center to access applications, or all branch office users are treated as Internet users and applications are delivered via secured Internet protocols. This reduces remote infrastructure deployment to a bare minimum and eliminates remote server management.

Microsoft offers a number of products and solutions for the branch office, as well as extensive infrastructure optimization (IO) guidance to help you select and build an effective and secure remote or branch office infrastructure. For the rest of this chapter, we will focus on branch office features and enhancements in Windows Server 2008, primarily those that are directly

related to securing the branch office. For more information about branch offices in general (or using Windows Server products in branch offices), visit *http://technet.microsoft.com/ en-us/branchoffice/default.aspx.*

Windows Server 2008 in the Branch Office

Improving the situation in branch offices was one of the main goals during the development of Windows Server 2008. Features in Windows Server 2008 address the concerns of a branch office in three areas: cost control, security, and business agility. Table 14-1 summarizes the features designed in Windows Server 2008 to overcome branch office challenges.

Table 14-1   Windows Server 2008 Branch Office Challenges and Solutions

Challenge	Requirement	Windows Server 2008 Features
Cost control	Reduce the cost of managing and supporting remote offices (including making most efficient use of network links).	Server Core installation option
		Read-Only Domain Controller (RODC)
		Server Message Block (SMB) 2.0
		Data Protection (Backup)
		Windows Server Virtualization
Security	Improve security of data and access.	BitLocker Drive Encryption
		Read-Only Domain Controller (RODC)
		Server Message Block (SMB) 2.0
Agility	Provide a flexible infrastructure that maximizes IT investment.	Server Message Block (SMB) 2.0
		TCP/IP redesign

Nonsecurity Features

I want to spend most of this discussion covering security features, but other features deserve a mention too, because it is often difficult—and usually counterproductive—to look at security in complete isolation. By understanding how all of the pieces fit together, you can be more effective—and that leads to being more secure.

In a very pragmatic sense, it is also a really good idea to strive to keep your network solutions cost-effective, agile, reliable, and available. When your solutions fail to meet end users' needs (or perceived needs), your end users (and especially management) will seek workarounds or shortcuts. These shortcuts are *rarely* secure. This is not to say that these users are being malicious. They will, however, go around you or your security when what you have put in place does not do what they need, or what they think they need.

Server Core Installation Option

Server Core is a minimal server installation option for Windows Server 2008 that contains a subset of executable files and makes available only a few server roles. Server Core runs on either 32-bit or 64-bit architectures.

I consider Server Core to be a security feature because it reduces the attack surface in the server. Only the binaries need to support the role and the base operating systems. This means that generally fewer processes are running. But Server Core is also a cost-control feature because it reduces maintenance and management requirements.

The Server Core installation option installs only a subset of the executable files and supporting dynamic-link libraries (DLLs). Specifically, only those features that are required by nine specific server roles are installed. Components not installed include Graphical User Interface (GUI), the .NET Framework Common Language Runtime (CLR), the Windows Explorer user interface, and Internet Explorer. Because there is no GUI interface, administrators must access all of these services through a local command shell, either remotely from a management server or workstation (with a remote command shell or with a remote tool, such as an MMC snap-in), or by using the new remote management feature called WS-Management, or by using the Windows Management Interface (WMI).

One of the biggest single benefits of a Server Core installation is a much reduced need for updating. While updates are the bane of an IT professional's existence in the data center, updating remote and branch offices has its own special kind of pain. Based on comparisons with Windows Server 2003, Microsoft estimates that less than 60 percent of updates would have applied to a Server Core installation than would be needed for a full installation of Windows Server, if that kind of installation existed for Windows Server 2003. Server Core is described in more detail in Chapter 12, "Securing Server Roles."

Hyper-V (Windows Server Virtualization)

Windows Server 2008 Hyper-V, formerly called Windows Server Virtualization, is a virtualization platform designed to provide the flexibility through a dynamic, reliable, and scalable platform capabilities. You can also use a single set of integrated management tools to manage both physical and virtual server resources. Hyper-V is included in most Windows Server 2008 licenses, but you may need to download it from Microsoft separately.

A full discussion of virtualization is outside the scope of this chapter. But because many branch office situations will benefit from Server Virtualization, you should read more about it on the Web at *http://www.microsoft.com/windowsserver2008/virtualization/default.mspx.*

Complete Redesign of TCP/IP Stack

Sometimes called the Next Generation TCP/IP stack, the improved networking stack in Windows Server 2008 and Windows Vista has some direct benefits to the branch office. By improving network communication between servers in the data center and the branch office, and between remote clients and centralized services, you can maximize the use of WAN links.

From a security point of view, you should be aware of the following aspects of the new TCP/IP functionality:

■ Expanded IPsec integration and easier IPsec management tools.

■ TCP/IP includes an integrated filtering platform. Windows uses this platform for Windows Firewall and IPsec, but other applications and vendors can also make use of the Windows Filtering Platform (WFP).

Both of these issues are discussed in Chapter 5, "Windows Firewall(s)," and Chapter 13, "Managing Security Dependencies to Secure Your Network." As well, the redesign of TCP/IP is discussed extensively in TechNet online (*http://technet.microsoft.com/en-us/network/ bb545475.aspx*) and *Windows Server 2008 TCP/IP Protocols and Services* by Joseph Davies (Microsoft Press, 2008).

Server Message Block

Continuing with the theme of connectivity and protocols, Windows Server 2008 includes a new remote file sharing protocol, Server Message Block 2.0 (SMB 2.0). SMB is the protocol generally used between Windows computers, such as Windows XP, Windows Vista, Windows Server 2003, and Windows Server 2008.

SMB 2.0 greatly improves the scalability of SMB 1.0, increasing the number of open files and allowed shares. The protocols have been enhanced to reduce the "chattiness" that makes file sharing on a WAN sometimes painful. Real-world testing has shown improvements with SMB 2.0 copying large files up to 35 times faster than with SMB 1.0. However, SMB 2.0 benefits only really appear when communicating between Windows Server 2008 computers or between Windows Server 2008 and Windows Vista clients. SMB 2.0 is not available separately for older Windows operating systems.

To learn more about the WAN performance improvements in SMB 2.0, watch the TechNet webcast "Wide Area Network Performance Improvements in Windows Server 2008" at *http://msevents.microsoft.com/CUI/WebCastEventDetails.aspx?culture=en-US& EventID=1032348651&CountryCode=US*.

Data Protection—Windows Server 2008 Server Backup

Remember that security includes keeping data available and being able to recover from disasters. Windows Server Backup uses the Volume Shadow Copy Service (VSS) and block-level backup technology to back up your servers, locally or remotely, with less downtime. Support has also been added for backing up to optical media (DVD) and to another file server.

Security Features for the Branch Office

Two more major components of Windows Server 2008 have a direct impact on securing your branch offices: Read-Only Domain Controller (RODC) and BitLocker Drive Encryption.

RODC

RODC is a new feature in Active Directory Domain Services (AD DS) in Windows Server 2008. Please refer to Chapter 9, "Designing Active Directory Domain Services for Security," for an overview of AD DS features and information about using and installing RODCs. In this chapter, I want to focus on the security implications of an RODC in a branch office.

RODC Read-Only Database, Replication, and Group Memberships   As the name suggests, RODCs are read-only copies of the AD DS database. This means that they require only unidirectional replication for Active Directory, as well for the File Replication Service (FRS) and Distributed File System Replication (DFSR).

One-way replication imparts a security benefit. Any compromise or other issue that introduces poisoned data into the RODC's local copy of the AD DS database cannot be replicated back to the rest of the organization from the affected RODC. This is certainly a mitigation that can help stop a local problem from becoming a global problem! On the other hand, be aware that if malware, attacks, or other issues cause you to have a corrupt, or "owned" local domain controller—even a read-only one—you still have a big problem to deal with!

Additional steps have also been taken to minimize the problems that could potentially be caused by a rogue RODC. Each RODC has its own, unique KrbTGT account for the Key Distribution Center (KDC), allowing the system to know when a ticket has been issued by an RODC (and which one) and when it has been issued by a normal, writable DC. This is a form of *cryptographic isolation*, a phrase you will sometimes hear used in conjunction with RODC discussions.

One-way replication brings benefits in terms of designing your replication topology and controlling replication traffic, as well. Bridgeheads and hubs do not have to poll the RODC for changes. The RODC performs normal inbound replication for AD DS and FRS and any DFSR changes.

Because the RODC is a member of the domain, sometimes it has a legitimate need to write to Active Directory. However, it does not write to the local database copy, but will instead connect to a writable DC, just like a workstation. The RODC computer account is a workstation account, so it has very limited rights to write to AD DS—again to minimize any damage to the enterprise AD DS if the RODC is compromised. Because they are "workstations" in this sense, RODC computer accounts are not members of the Enterprise Domain Controllers (EDC) or Domain Domain Controllers groups.

Database Contents and Credential Caching   With the exception of account passwords and attributes specifically added to the filtered attribute set, an RODC holds all the AD DS objects and attributes that a writable domain controller holds.

Local applications that request read access to the domain directory information can obtain access directly from the RODC, while Lightweight Directory Access Protocol (LDAP) applications that request write access receive an LDAP referral response. This referral response directs them to a writable domain controller, normally in a hub site.

In the domain database, each security principal—user, computer, or iNetOrgPerson—has a set of up to 10 passwords or secrets, collectively called credentials. An RODC does not store these credentials, except for its own computer account and the special krbtgt account for each RODC. The RODC is advertised as the Key Distribution Center (KDC) for its site, which should normally be the branch office in which it is located. When the RODC issues and signs a ticket-granting ticket (TGT), it uses its own krbtgt account and keys, which are different from the krbtgt account shared by the KDCs on all writable domain controllers in the domain.

The first time a user attempts to authenticate to an RODC, the RODC sends the request to a writable domain controller at the hub site. If the authentication is successful, the RODC also requests a copy of the appropriate credentials. The writable domain controller recognizes that the request is coming from an RODC and consults the Password Replication Policy that is in effect for that RODC.

The Password Replication Policy determines whether the credentials are allowed to be replicated and stored on the RODC. If so, a writable domain controller sends the credentials to the RODC, and the RODC caches them for future use. After the credentials are cached on the RODC, the next time that user attempts to log on, the request can be directly serviced by the RODC until the credentials change.

Obviously, for this to work the RODC has to know whether it is likely to have a copy of the credentials. To avoid needless lookups, the RODC checks the signature on various requests. If a TGT has been signed with its own krbtgt account, the RODC immediately recognizes that it has a cached copy of the credentials. Conversely, if another DC has signed the TGT, the RODC will always forward requests.

The key to limiting risk by installing RODCs in branch offices is correctly using the Password Replication Policy to control which credentials are cached. By default, no user credentials will be cached on an RODC, but that is an inefficient design. If you think about most domains with a branch office, you realize that even in a large branch (say 100 users), those users are a tiny percentage of users in the domain. If you have 20,000 users in the domain, that is one half of one percent.

Because only a few domain users need to have credentials cached on any given RODC, use the Password Replication Policy to specify which groups of users can even be considered for caching. By limiting RODC caching to only users who are frequently at that branch office, you can reduce the potential exposure. You can—and should—also explicitly block the caching of credentials for high-value accounts. By denying the caching of accounts for Domain Admins, for example, you guarantee that Domain Admin credentials will never be stored on the RODC in that branch office. This is where the real security benefit comes in over Windows Server 2003,

where you had no option but to have a local copy of every account in every branch office. This is a huge security improvement.

You access the Password Replication Policy by selecting the Properties of the RODC computer object. In Figure 14-1, you can see that only accounts in a group representing that Branch Office are allowed to be cached, and high-value credentials are explicitly denied.
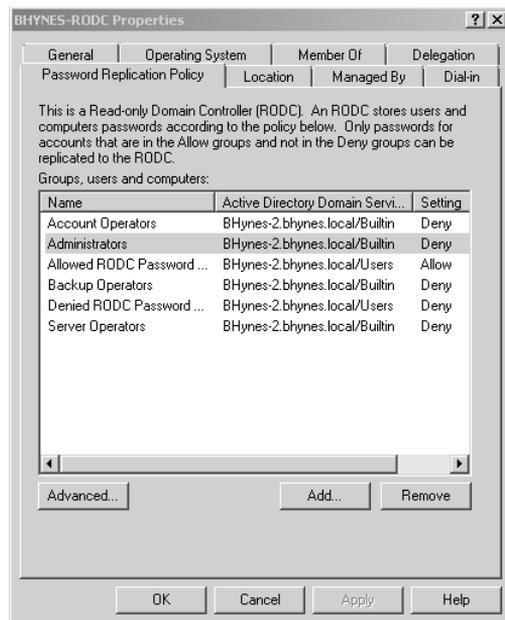
Figure 14-1   An RODC's property sheet showing its Password Replication Policy.

Let us consider for a moment that the unthinkable has happened and your RODC is stolen. The evildoer mounts the AD DS database and attempts to retrieve passwords using a commonly available, umm, "security research" tool. However, instead of seeing more and more passwords appear as the password cracking progresses, the tool only shows empty or blank passwords for almost all accounts.

In the past, if your DC was stolen, you would really have no choice but to reset all of the passwords in the domain—a formidable task if you have thousands of users. However, in this case, you only need to deal with the specific accounts that have already been cached on that RODC. To be clear: With an RODC, you know exactly which passwords have been cached. You can check which ones have been replicated to the RODC by looking at the Password Replication Policy dialog box. Initially, only the computer itself and the special KRBTG account have passwords stored locally, as you can see in Figure 14-2. After clients begin to authenticate to the RODC, additional passwords may also be stored, as shown in Figure 14-3.

You can also determine which users have attempted to authenticate to the RODC, but whose passwords have not been allowed to be cached. For example, notice in Figures 14-2 and 14-3

that no password has been cached for the Administrator account. However, as shown in Figure 14-4, the administrator has logged on in the site served by the RODC.
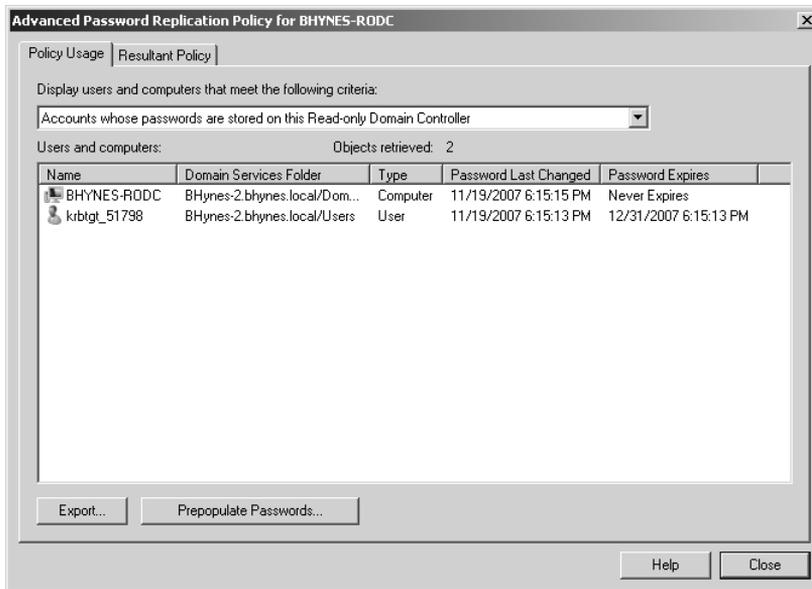
Figure 14-2   Initially, only two accounts have passwords stored on the RODC.
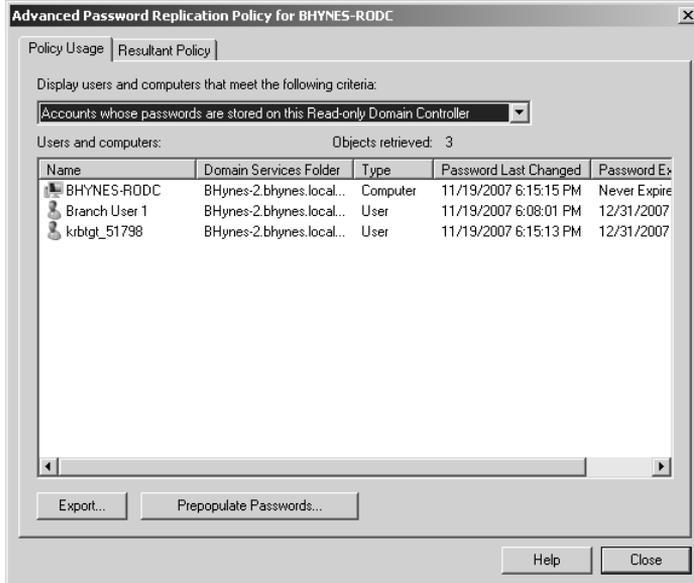
Figure 14-3   You can store additional passwords on the RODC.

Finally, as shown in Figure 14-5, you can model which accounts will be allowed or not allowed to have passwords cached on the RODC.
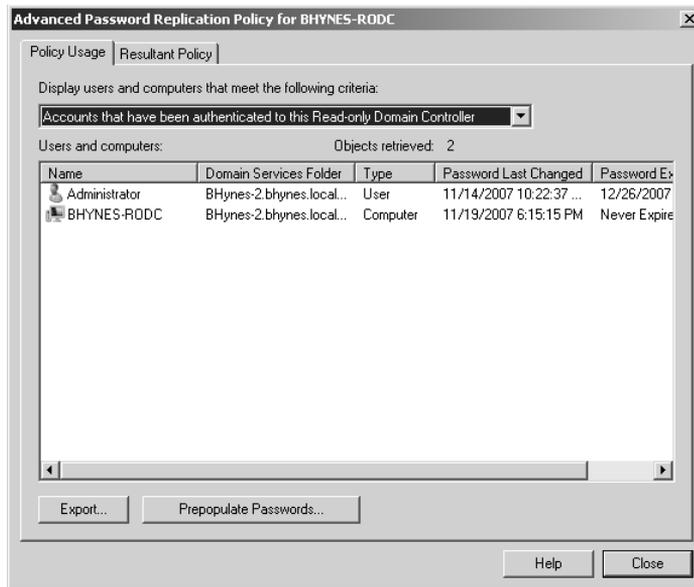
Figure 14-4 Users who have authenticated to the RODC, but have not had passwords cached.
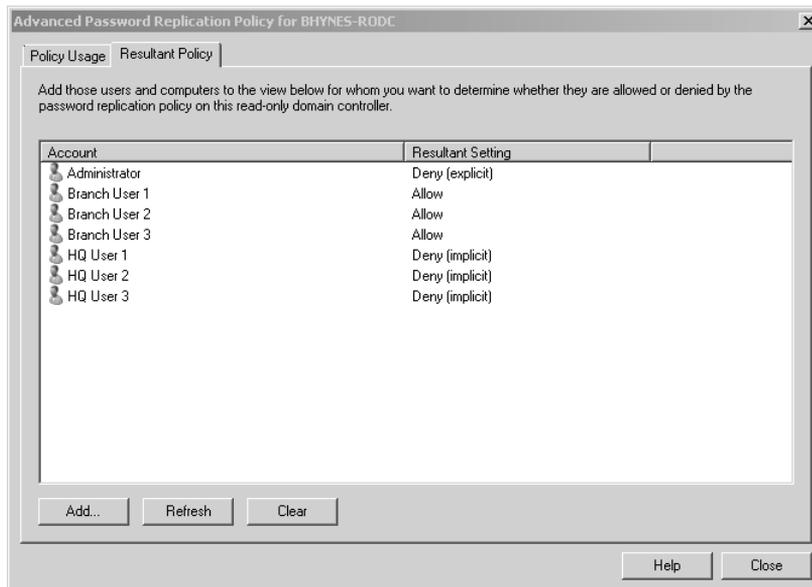
Figure 14-5 Using the Resultant Policy page to model password caching.

We have made it much easier to delete missing domain controllers, including RODCs, with a simple interface that can remove the departed RODC from the domain (no more metadata cleanup needed!) and also automatically reset the potentially compromised passwords. When you start to delete the computer object representing an RODC, you see the dialog pictured in Figure 14-6.
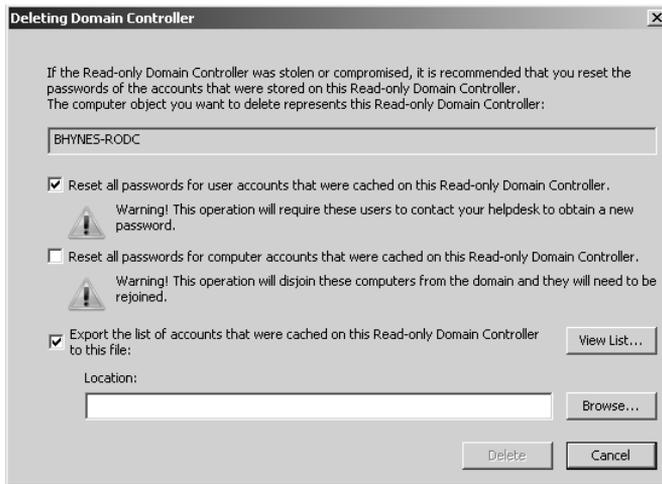
Figure 14-6   Removing an RODC that cannot be brought online from AD DS.

You should also read ahead to the section on BitLocker Drive Encryption. By using BitLocker on the Windows operating system volume of the RODC, you can prevent the thief from even being able to access the local copy of the AD DS database—which, at a minimum, buys you time to reset passwords in an orderly fashion.

Administrative Role Separation   With Role Separation you can delegate the local administrator role of an RODC computer to any domain user without granting that user any rights to the domain itself or to other domain controllers. In Windows Server 2003, DCs didn't have a local administrator; if you could administer a DC, you could administer the whole domain.

Administrative Role Separation can allow a local branch user to log on to an RODC and perform maintenance work on the server, such as upgrading a driver, without allowing that user to log on to any other domain controller or manage the domain.

RODC Benefits   RODCs provide a way to deploy domain controllers more securely in a branch office location because they are designed to be placed in locations that require rapid, reliable, and robust authentication services but that might also have a security limitation that limits or prevents deployment of a writable domain controller. With an RODC, organizations can mitigate risks with deploying a domain controller in locations where physical security cannot be guaranteed.

The RODC feature is clearly designed for branch offices, but it is an integral part of AD DS as well. Please review Chapter 9 for information about installing an RODC, using the RODC filtered attribute set, and configuring read-only DNS.

BitLocker Drive Encryption

BitLocker Drive Encryption is a component of the operating system in Windows Vista Enterprise Edition and Windows Vista Ultimate Edition, and it is also included as a feature in all

editions of Windows Server 2008. You may already be familiar with the need to protect data on client hard drives, but that need can be just as important on servers, particularly in a branch office.

According to Forrester, data breaches are expensive: "The cost of a single, significant security breach may run into millions or even billions of dollars." ("Calculating the Cost of a Security Breach," Khalid Kark, April 10, 2007: *http://www.forrester.com/Research/Document/Excerpt/ 0,7211,42082,00.html*) and, unfortunately, they are also common: "34% of respondents experienced at least one personal data breach." ("Aligning Data Protection Priorities with Risk," Jonathan Penn, April 13, 2006: *http://www.forrester.com/Research/Document/Excerpt/ 0,7211,39257,00.html*)

Most IT professionals are well aware of the risk of a user's laptop being left in a taxi or stolen. Hopefully, you are not likely to leave a production server in a taxi, but have you considered what would happen if your servers were stolen? I worked with a military customer once who told me that his organization had equipped domain controllers with C4 explosive charges because they were deployed "in theater" and would be destroyed rather than risk authentication data falling into enemy hands. Now, your risk is probably not the same as those on the battlefield, but there is always some risk, and it is not just hypothetical. A colleague of mine (okay, it was Jesper), was involved with a customer's organization that did have a thief back a truck into and through the wall of their building to steal their servers.

BitLocker gives you another tool to help mitigate the risk at the level you need. I hope, however, that you are thinking beyond the data center. Here are some contrasting examples from my work experience:

■   I visited one particular customer whose data center happened to be in a 100-year-old former sanatorium. The servers were all behind two-foot-thick stone walls, and staffed 24 hours a day. It would be hard to steal those servers.

■   I volunteered at a large youth event, where the "data center" was in a tent. Servers strung around several hundred acres lost connectivity whenever someone plugged a kettle into the same electrical outlet as the fiber-optic repeater.

Lest anyone dismiss my second example, let me point out that both servers held confidential personal information including health information about clients or participants—and both were where they *needed to be to meet the particular business demand*.

So what is BitLocker Drive Encryption and how can it help secure the data on your servers? BitLocker performs two functions:

■   BitLocker encrypts all data stored on the Windows operating system volume (and configured data volumes). This includes the Windows operating system, hibernation and paging files, applications, and data used by applications.

■   BitLocker is configured by default to use a Trusted Platform Module (TPM) microchip to help ensure the integrity of early start-up components (components used in the earlier stages of the start-up process), and locks any BitLocker-protected volumes so that they remain protected even if the computer is tampered with when the operating system is not running.

BitLocker Configuration in Windows Server 2008   In Windows Server 2008, BitLocker is an optional component that must be installed before it can be used. (In Server Manager, optional components are called *features.*) To install BitLocker, select it in Server Manager or type the following at a command prompt: **ServerManagerCmd -install BitLocker−restart**.

Installing BitLocker on a production server means giving consideration to a few points. BitLocker has two optional components (OCs) included with Windows Server 2008. The BitLocker OC is used to protect volumes on that particular server. There is also a remote administration OC called RSAT-BitLocker that installs the binary files and command-line scripts required to manage BitLocker on remote servers. (RSAT stands for Remote Server Administration Tool.) If you install or remove the BitLocker feature using Server Manager, both OCs are installed or removed. If you only want to install the remote tool, you must use either the **pkgmgr** command or the **ocsetup** command. Because Server Manager is not supported with the Server Core installation option, you must use **pkgmgr** or **ocsetup** to install BitLocker on Server Core.

Because BitLocker uses a filter driver for encryption and decryption, installing or removing BitLocker requires a restart of the computer. Installing the RSAT-BitLocker component alone does not. After the component is installed, you can use the command-line tool, manage-bde.wsf, or Control Panel to enable BitLocker on particular volumes.

BitLocker is different from existing technologies like EFS because once you turn BitLocker on, it is automatic, transparent, and includes the entire volume. (Note that BitLocker is fully compatible with EFS, and you can use them both together to mitigate different risks. Active Directory Rights Management Service, or AD RMS, adds the third part of Microsoft's data encryption strategy. EFS and RMS are outside the scope of this chapter.)

BitLocker does require that a computer be configured with multiple volumes (partitions). A computer will start from the active partition, sometimes called the system partition or system volume. The operating system is installed on a second volume, which Windows Server 2008 documentation usually calls the Windows OS volume, but older documentation refers to as the boot partition. (Yes, for those keeping score: You never did boot from the boot partition and no Windows system files are on the system partition.)

The active partition is not encrypted, but the Windows OS volume is. In other words, BitLocker encrypts everything written to a protected volume: code, data, and temporary files.

BitLocker's full-volume encryption protects against offline attacks—the kind of attacks that are mounted by trying to bypass the operating system. It is absolutely essential to understand that BitLocker does not protect a *running* operating system. Once the Windows OS Volume has been unlocked at start-up, BitLocker continues to encrypt and decrypt sectors on the fly—whenever an application or the operating system itself reads or writes data on a protected volume—but it has no more protection function to perform. Instead, BitLocker helps protect you against offline attacks such as having the disk (or entire server) stolen.

Without BitLocker, it is a trivial exercise to take the disk out of a server and put it in another computer to bypass NTFS permissions and similar defenses. With BitLocker, though, nothing on that disk can be read in another computer, unless you have the required recovery password to unlock the volume. It is also true that files protected by RMS and EFS will also not be easily read, but BitLocker extends protection to the entire volume, including files that cannot be protected by RMS or EFS (such as the Active Directory database or the registry).

Direct from the Source: Automatically Unlock Data Volumes

In Windows Server 2008, Microsoft added support for fixed data volumes to BitLocker. Along with that comes the ability to automatically unlock (auto-unlock) any encrypted data volumes. To do this, BitLocker stores keys for encrypted data volumes (the "auto-unlock keys") in the system registry, where they are retrieved by BitLocker to unlock the volume.

It is very important to protect the auto-unlock keys, since they can be used to decrypt the volume. Therefore, these keys are always themselves encrypted before being stored in the registry. They are encrypted by using a new key called the Auto-Unlock Master Key (AMK). The AMK is very similar to the FVEK. The FVEK is used to encrypt the data on the volume, while the AMK is used to encrypt the data in the registry. Like the FVEK, the AMK is encrypted by the OS Volume's VMK, which explains the necessity of enabling BitLocker first on the OS Volume before it can be enabled on any data volume.

With this design, you can safely enable BitLocker on a data volume and set it to auto-unlock (which is the by default) without waiting for the OS Volume to be fully encrypted.

Also, since decrypting the OS Volume causes the AMK to be deleted and, therefore, the auto-unlock keys in the registry to become unusable, if you decrypt the OS volume, we also force decryption of all mounted data volumes that are set to auto-unlock.

Narendra S. Acharya
Software Design Engineer in Test (SDET)
System Integrity, Core OS Division

BitLocker Encryption Algorithms and Keys   BitLocker uses the Advanced Encryption Standard (AES) algorithm with a 256-bit key space. However, with the default setting only 128 bits of the key are used. For stronger encryption, the key length can be increased to 256-bit keys using Group Policy or the BitLocker Windows Management Instrumentation (WMI) provider. As with all encryption, stronger encryption and longer keys result in more processing power being used. (To take effect, any changes to the encryption settings need to be made before BitLocker is turned on for a volume.)

Before data is passed to the AES algorithm, it is run through two separate diffuser steps. A *diffuser* is a complex cryptographic algorithm with a simple purpose. The diffusers ensure that any change (even only one bit) made to the plaintext will result in the entire block being

changed in the encrypted ciphertext. This makes it harder, if not impossible, to guess keys by making changes in the plaintext to see predictable output in the ciphertext.

With BitLocker, each sector in the volume is encrypted individually. Part of the encryption key is calculated from the sector number itself. This means that two sectors with the same unencrypted data will also result in two totally different encrypted blocks being written to the disk. This also makes it harder to discover the keys by encrypting known pieces of information.

One of the developers of the algorithm, and a respected cryptographer who worked on BitLocker, Niels Ferguson, presented and published a paper on the details of the BitLocker encryption algorithm. It is available on the book's companion CD-ROM and online at *http://www.microsoft.com/downloads/details.aspx?familyid=131dae03-39ae-48be-a8d6-8b0034c92555&displaylang=en*: "*AES-CBC + Elephant Diffuser: A Disk Encryption Algorithm for Windows Vista.*" A simplified overview is shown in Figure 14-7.
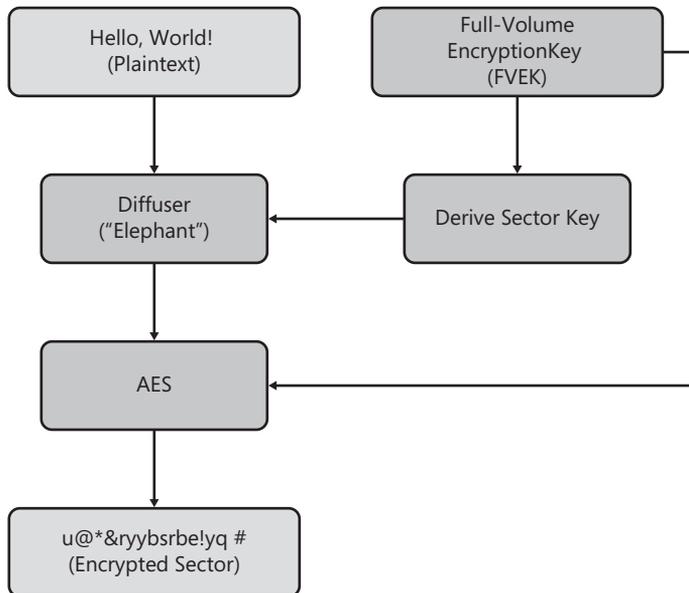
Figure 14-7   Simplified overview of BitLocker encryption algorithm.

As I mentioned, each sector is encrypted uniquely, with part of the key being derived from the sector number. The symmetric encryption key that is used for the bulk of the encryption is called the Full-Volume Encryption Key (FVEK). The FVEK is never revealed to a user or to an administrator. If the FVEK were compromised, security best practices would dictate that everything it encrypted would need to be decrypted and re-encrypted with a new key. Because we could be dealing with very large quantities of data, that would be a time-consuming process (during which noticeable performance degradation could occur). Instead, the FVEK is kept as a very closely guarded secret, and the system works with another key called the Volume Master Key (VMK). The VMK is used to encrypt the FVEK. The encrypted copy of the FVEK is written to the volume metadata. (It can also be backed up as part of a binary key package to AD DS.) This is not like leaving the key under the mat—it is more like leaving the key in a safe bolted to

two tons of concrete, and then hidden under the mat. In addition, if the VMK is somehow compromised, you can re-encrypt the FVEK with a new VMK very quickly.

To decrypt sectors, the system needs to get the FVEK. To get the FVEK, the system needs to get the VMK. The VMK is also stored in the volume metadata, and it is also never written to disk without being strongly encrypted itself.

The VMK can be encrypted by any number of key protectors or authenticators. The default key protector is the computer's Trusted Platform Module (TPM), as discussed in the following sections. You can combine the TPM with a numeric PIN or a partial key stored on a USB key or both. By adding a PIN or USB requirement, you are introducing a form of two-factor authentication.

If the computer does not have a compatible TPM, you can configure BitLocker to store a key protector completely on a USB Flash Drive (although this does not provide any boot integrity checking, as discussed later). The VMK is also encrypted with a 48-digit recovery password that can be backed up to AD DS, stored in a database, or printed to be used for emergency recovery.

Finally, in a case where you want to disable BitLocker but not fully decrypt the drive, BitLocker uses a *clear key* key protector. The clear key is very much like putting the key to the door under the mat. The VMK is encrypted with a new key, but that key is stored in plaintext on the volume. This is effective for operations that require BitLocker to be deactivated, such as some kinds of upgrades, but avoids the time required to decrypt and re-encrypt the volume. However, because the VMK could be somewhat easily discovered in this state, you should not leave BitLocker in this disabled state for any longer than necessary.

BitLocker Start-up   At start-up, all BitLocker-protected volumes are locked and cannot be used, even to start the operating system. Early start-up code in the Boot Manager and Windows Loader looks for an appropriate key protector by querying the TPM, checking the USB ports, or, if necessary, prompting the user for a PIN or a recovery password. Finding a key protector lets Windows decrypt the VMK, which decrypts the FVEK. At this point, the disk is unlocked and the BitLocker filter driver decrypts the data stored on disk as each sector is read.

BitLocker Integrity Checking   When a server starts, the components used early in the process need to be unencrypted so that they can be used in start-up. If an attacker could change the code in the earliest start-up components, compromising a system is much, much easier. This is the basis of many rootkits. Once a system is compromised this way, secrets such as passwords can be much more easily obtained.

In the case of a system protected by BitLocker, a compromise of the earliest components could lead to the disclosure of the VMK or FVEK and render the protection of BitLocker useless, even though data on the disk is encrypted.

If the computer has a compatible TPM, BitLocker can use integrity checking functionality to help prevent this sort of attack. With a TPM, each time the computer starts, each of the early start-up components—such as the BIOS, the master boot record (MBR), the boot sector, and

the boot manager code—examines the code about to be run, calculates a hash value, and stores the value in specific registers called platform configuration registers (PCRs) in the TPM. Once a value is stored in a PCR, the value cannot be replaced or erased unless the system is restarted.

A TPM can encrypt data (in this case, the VMK) and tie that encryption to specific PCR values, which is called *sealing* a key to the TPM. Once sealed, only that specific TPM can decrypt that key, and the TPM will decrypt the key only if the current PCR values match the values specified when the key was sealed.

By default, BitLocker seals keys to the PCRs that record the Core Root of Trust Measurement (CRTM), the BIOS and any platform extensions, option ROM code, MBR code, the NTFS boot sector, and the boot manager. If any of these items change unexpectedly, the TPM will not decrypt its key protector, and BitLocker will not be able to unlock the volume. This prevents the volume from being accessed or decrypted.

By default, BitLocker is configured to look for and use a TPM. You can use Group Policy or a local policy setting to allow BitLocker to work without a TPM and store keys on an external USB Flash Drive, but without a TPM, BitLocker cannot verify system integrity.

BitLocker in the Branch Office   "That's all well and good," you might be thinking, "and sure, it helps with laptops. But what does it have to do with my server in a branch office?" I'm glad you asked.

In several specific scenarios BitLocker can help immensely with servers, particularly those not housed in the central data center. These include:

- **Reducing exposure to malware, especially rootkits.** Most rootkits and other malware that can do serious damage on your servers require a restart to install. By using a TPM module and BitLocker's integrity checking, the server will not start if the early start-up components have been manipulated. (Windows Code Integrity features take over protecting the operating system against unauthorized changes after BitLocker unlocks the volume.) It is true that this could result in a denial of service because the server will not restart, but it is generally better to be stopped than to be running with active malware on a production server.

- **Theft of a server.** Thieves target laptops and mobile devices because they are often convenient targets. While some thieves steal them because they may contain some interesting bits of information, most laptop theft is for the hardware, not the infinitely more valuable data that is on them—information. But a server or its disks can the mother lode to a serious thief. A domain controller could contain secrets for thousands of users (see the section on RODCs!); a file server could be filled with your company's intellectual property; a database server could have reams and reams of personally identifiable information—handy for identity theft and fraud, to name but two. Note that to effectively mitigate against the risk of the entire server being stolen, using the default of TPM protection only will not be enough.

- **Theft of a disk**. It is often much easier to remove a single disk than to steal the entire server. Most server disks are swappable, so it only takes a moment for someone to remove one. Like a server, a disk contains valuable information and should be protected.

- **Shipping a server (or disk) to a remote location.** In many cases, servers are first provisioned centrally or large amounts of data need to be shipped to a remote server. In either case, you can enable BitLocker on the volume at the data center, and then ship the volume to the remote location. In the case of an entire server, it can be configured to require a start-up PIN. For a single volume, all key protectors except the recovery password can be removed. Then, after authorized personnel have possession of the server or disk at the destination, the PIN or recovery password can be provided by telephone, fax, or e-mail. Once the server is in production, if desired, you can remove the PIN requirement or even decrypt the volume. If the disk or server is lost or stolen in transit, the data cannot be accessed.

- **Secure decommissioning.** At some point, a server or disk needs to be removed from service. It is absolutely essential that no server disk with company data leave your possession with that data still readable. Most processes that remove confidential data from disk drives are time consuming, costly, or result in the permanent destruction of the hardware. Instead of removing the data after the fact, BitLocker helps ensure that confidential data is not stored on disk in a risky way in the first place. Because everything written to the disk is encrypted, you can permanently render the data completely inaccessible by destroying all copies of the encryption keys. The hard disk itself is completely unharmed and can be reused. Removal and destruction of the keys contained in the volume metadata is instantaneous and can be performed locally or remotely by an administrator. The format utility in Windows Vista and Windows Server 2008 deletes the volume metadata and overwrites those sectors to securely delete any BitLocker keys. (Note that you cannot use the format utility from Windows versions earlier than Windows Vista to achieve the same result.)

BitLocker Caveats and Trade-offs   As you consider using BitLocker on your servers, bear a few things in mind.

Remember that BitLocker protects against certain specific threats—from offline attacks. It is imperative that other security controls, such as strong passwords, remain in place. Likewise, BitLocker does not prevent the destruction of data or the denial of service. Backups are still essential. In fact, if you do not take care to ensure the accessibility of the recovery passwords, BitLocker could prevent you from accessing your data! We recommend using BitLocker's built-in ability to back up recovery information to AD DS. For more information about backing up recovery information, see "Configuring Active Directory to Back Up Windows BitLocker Drive Encryption and Trusted Platform Module Recovery Information" at *http://technet2.microsoft.com/WindowsVista/en/library/3dbad515-5a32-4330-ad6f-d1fb6dfcdd411033.mspx?mfr=true*.

In Windows Server 2008 and Windows Vista SP1, BitLocker uses a cryptographic signature to check for certain types of deliberate attacks. If one of these attacks occurs, the cryptographic

signature is invalid, and BitLocker will lock out the recovery password as well as routine key protectors. This scenario will only occur if the attacker has physical possession of the volume or computer. If you regain possession of the volume, you will need to use the BitLocker Repair Tool and a backup copy of the FVEK, as well as the recovery password. Therefore, you should ensure that these items are backed up using the AD DS system or the WMI provider.

Also keep in mind that BitLocker is an optional feature on Windows Server 2008. Because the encryption and decryption operations are implemented in a filter driver, you cannot use BitLocker until you install the feature, as described earlier in this chapter. Note, however, that installing or removing BitLocker requires restarting the server.

BitLocker does not support clusters. If you configure servers in failover clusters, you cannot use BitLocker to protect shared volumes.

BitLocker—like all encryption products and technologies—incurs a performance hit. In most user scenarios, this performance is not noticeable, and can be measured at well less than 5 percent. However, if your server is performing intensive disk operations or is very heavily loaded, you may need to take this into consideration in capacity planning.

Security is often a matter of trade-offs. One significant trade-off to make with regards to BitLocker is whether to require user intervention at start-up. By default, BitLocker operates using the TPM as a key protector and, providing that the integrity validation completes successfully, BitLocker will unlock the Windows operating system volume without user intervention. However, this means that a thief who takes the entire server can also turn it on without any interference from BitLocker.

To mitigate this risk, use the TPM with a PIN, with a USB Flash Drive (called a *startup key*), or both. This increases security tremendously, but it also means that a server cannot automatically restart. (Leaving a start-up key in the server all the time would be a foolish idea.) Because servers are normally running this may not be a problem—until a power failure results in the branch being down because the one manager who had the PIN is on vacation.

My recommendation, in most cases, is to use a TPM plus PIN, but to ensure that there is an operational procedure that describes who knows the PIN, how to use it, and how to obtain the PIN from central support if needed. This also requires that any planned restarts (such as for an operating system upgrade or update) are scheduled when someone with the PIN is on-site. Only your organization can make the call about which is the more important risk: potential additional downtime or potential theft of the server.

Other Security Steps

I believe you need to invest in two other areas to protect your branch office servers: physical security and user education.

While features such as RODC and BitLocker help mitigate risk, a bad guy having physical possession of your server is still a big deal. Even if he cannot gain access to the data, he can

still prevent your users from gaining access to it. Reasonable investments in physical security also increase the availability of services.

I continue to hear stories of cleaning staff unplugging servers, unnecessary water damage, and accidental shutdowns. Use locked closets or cabinets. Plan for UPS protection. Ensure proper ventilation. For large enterprises, this may be obvious—but then again, in a big firm, do you know the condition of every server in the truck shop or out on the plant floor?

Also invest in user education. Explain the importance of security and how to treat both the physical and data assets in their care. This may be more important than ever in small branch offices with no dedicated IT staff, let alone compliance and security officers.

Summary

Windows Server 2008 has been designed from the ground up to support branch offices, and introduces tools to make securing them easier than ever. The RODC and BitLocker features are key weapons in your arsenal of protection.

Additional Resources

- Windows BitLocker Drive Encryption Frequently Asked Questions
- BitLocker Drive Encryption Technical Overview
- *Data Encryption Toolkit for Mobile PCs*
- BitLocker Drive Encryption Glossary
- BitLocker Drive Encryption (from Changes in Functionality from Windows Server 2003 with SP1 to Windows Server 2008)
- BitLocker Drive Encryption Step-by-Step Guide for Windows Server 2008
- Windows BitLocker Drive Encryption Step-by-Step Guide (for Windows Vista)
- Windows BitLocker Drive Encryption Design and Deployment Guides
- *Configuring Active Directory to Back Up Windows BitLocker Drive Encryption and Trusted Platform Module Recovery Information*
- BitLocker Drive Encryption Events and Errors
- AD DS: Read-Only Domain Controllers (*http://technet2.microsoft.com/ windowsserver2008/en/library/ce82863f-9303-444f-9bb3- ecaf649bd3dd1033.mspx?mfr=true*)
- Step-by-Step Guide for Read-Only Domain Controllers (*http://technet2.microsoft.com/ windowsserver2008/en/library/ea8d253e-0646-490c-93d3- b78c5e1d9db71033.mspx?mfr=true*)
- Branch Office Infrastructure Solution for Microsoft Windows Server 2003 Release 2 (*http://www.microsoft.com/technet/solutionaccelerators/branch/default.mspx*)