Realtime publishers
"Leading the Conversation"

July 2008
Volume 1, Number 6

Windows Administration
in Realtime

SmartShedding™ Technology

Allows the master outlet to sense when your computer has either been turned off or has gone into sleep mode, so it can shut off power to peripherals plugged into the controlled outlets–saving you power and money.

Your data should last forever.

And so should our planet.

Save $40* a year on your electric bill with the most efficient battery backup yet.

Let's protect what's important

What's in your computer? Photos, music, personal files, financial data, broadband access, videos, and more. Your computer has never been more important, and yet it has never been at higher risk for damaging power surges and other disturbances.

So like most people, you need to protect your assets. But like most people, you'd also like to protect the environment. With our new energy conscious products, you can do both. Energy efficient by design, our new smart products protect the power going in your computer, at a cost that is quickly offset by big energy savings. How? Not only do the new Back-UPS ES® and SurgeArrest® use power very wisely, they also boast a master/controlled outlets feature, which automatically powers down idle devices to conserve energy.

APC power protection products are available at:

BEST BUY · CDW · circuit city · STAPLES (that was easy.) · Office DEPOT

"The price tag on the new UPS is $99.99. While I'm not in the habit of endorsing products in this blog, if you're in the market for a workstation-class UPS, why not opt for the greener option?"

- Heather Clancy
ZDNet.com

In fact, while protecting your power supply, we're up to 5 times more energy efficient than any other solution. By saving you $40 a year in energy costs, our Back-UPS ES pays for itself in 2 short years. The high frequency, low copper design has a smaller transformer and environmental footprint. Even the packaging has been carefully selected and manufactured to maximize use of recycled materials and minimize waste.

In this world, every decision you make counts. So protect your power with a battery backup that works to protect the environment. It conserves power, it pays for itself, and it's backed by APC's 20-plus years of legendary reliability. For more information on this or our other great products, or for information about environmentally responsible disposal of your old battery, visit ***www.apc.com***

2008 Design & Engineering Showcase Honors Innovations International CES

Energy Efficient Solutions for Every Level of Protection:

Save $25 per year* on your electric bill!

Surge Protection

Starting at $34

Guaranteed protection from surges, spikes, and lightning.

7 outlets, Phone/Fax/Modem Protection, Master/Controlled Outlets

SurgeArrest® P7GT

Save $40 per year* on your electric bill!

Battery Back-UPS®

Starting at $99

Our most energy efficient backup for home computers.

10 outlets, DSL and Coax protection, Master/Controlled Outlets, High Frequency Design, 70 minutes of runtime†

Back-UPS® ES 750G

APC can help with your other power protection needs. Visit apc.com to see our complete line of innovative products.

Enter to Win a Back-UPS® ES 750G! (a $99.99 Value)

Also, enter keycode to view other special offers and discounts.

Visit www.apc.com/promo Key Code c439w or Call 888.289.APCC x9544 or Fax 401.788.2797

APC® Legendary Reliability®

e-mail: esupport@apc.com • 132 Fairgrounds Road, West Kingston, RI 02892 USA • 998-0968 †Runtimes may vary depending on load.
*Average savings are based on comparable competitive models, and are comprised of two energy saving features: An ultra efficient electrical design, and the master/controlled outlets feature.

Letter from the Editor

The State of IT Conferences

by Greg Shields

By the time you get this month's eJournal, we should be coming to the close of this year's spring conference season. The April/May/early-June timeframe always seems to bring with it a slate of IT conferences on topics ranging from Windows to security to developer issues. IT conferences, and the company-paid privilege of going once or twice per year, have long been considered part of the IT worker's job. In no other place do IT people come together for the purpose of learning new tips, tricks, and technologies in a condensed and high-energy format.

Well, at least for most of the time you're there. Many IT conferences are well-known for their blow-out afterschool "specials" as well…

But recent changes in the economic landscape as well as an explosion in the sheer number of conferences available for people to choose from have made attendance a much more difficult activity. High travel costs plus high conference fees plus more days missed from work and lower training budgets can combine to make it difficult to get the approval to attend.

And yet there's a problem intrinsic to the not going as well. If IT personnel are not given the privilege to attend these events, what other avenues do they have for bringing back tips and tricks to more effectively manage and secure the business network? Books and online articles are great, but when you're running with a hose between fires all day, getting the 30-minute breather to sit down and read a little is a difficult proposition.

I'm a big fan of IT conferences for just this reason. For the harried IT worker, they're the only opportunity to get away from the office and spend a week or so doing nothing but learning. No firefighting. No resolving of work orders. No fixing the boss' Blackberry problems. Merely a much-needed 15%-vacation-and-85%-new-ideas break from the day-to-day grind.

So my questions for you this month are actually many in number:

- What drives you to IT conferences? The location? The price? The content? The speakers?
- What are the inhibiting factors to getting you there? The cost? The time away? The justification to the powers that be?
- What is one critical component that absolutely must be present at a conference to get you there?

Let us know your thoughts at feedback@realtimepublishers.com. We'll post some of the best and most enlightening responses in a future issue. In the meantime, this month, take a gander at a great article by Darren Mar-Elia on Group Policy Preferences. I present a narrative on the tools available for rapidly deploying Windows Server 2008. Peter talks on Exchange high availability, and Jeff gives us his view on antivirus reporting. ◆

Answers from the Experts

Whitelisting Applications for Windows

by Don Jones

Q: This month's question comes from my blog at www.ConcentratedTech.com, where we've been having an ongoing discussion about application whitelisting. The question is: What third-party application do I have to buy in order to implement whitelisting in an all-Windows XP/Windows Server 2003 environment?

A: The answer? While companies such as Bit9, Symantec, and Lumesion all provide high-end whitelisting solutions, a basic whitelisting capability is built right into Windows, beginning with Windows XP. It's called Software Restriction Policies (SRP), and it's implemented in Windows Group Policy.

Obviously, the whole point of whitelisting is to first develop a list of "allowed" applications. I'm not going to understate the difficulty of doing so, especially for environments that have zillions of permitted applications. But, once that inventory process is complete, it'll be easy to maintain over the long term. Once you know your list of allowed applications, you need to determine how to best identify them.

Start by creating or modifying a Group Policy Object (GPO). Expand the Windows Settings section, then Security Settings, and locate the Software Restriction Policies section—it appears under both Computer and User configuration containers. Right-click Software Restriction Policies to create a new policy. The first setting you'll modify is Enforcement, shown in Figure 1.

Normally, you'll enforce SRP on everything except DLLs, which means your inventory has to focus on only primary executables—much easier. The Designated File Types setting allows you to configure everything that is considered an executable, which can include an Access database or whatever you like, in addition to EXE, DLL, VBS, and other built-in executable types. Under Security Levels, you'll define which restriction level is the default for all programs not specifically called out. Set Disallowed to the default, and you're in application whitelisting mode. A nice trick in Vista and Windows 2008 is that you can allow, by default, any software that doesn't need to run as an admin, and then whitelist software that does need Admin privileges—a great way to cut back on malware.

Finally, you identify applications that you want to specifically allow by adding Additional Rules. There are a few ways to identify apps:

- Path—Pretty insecure—malware could simply run itself from an allowed path
- Hash—Pretty secure—even altered executables will be detected
- Certificate—Extremely secure but requires executables to be digitally signed by their author or publisher

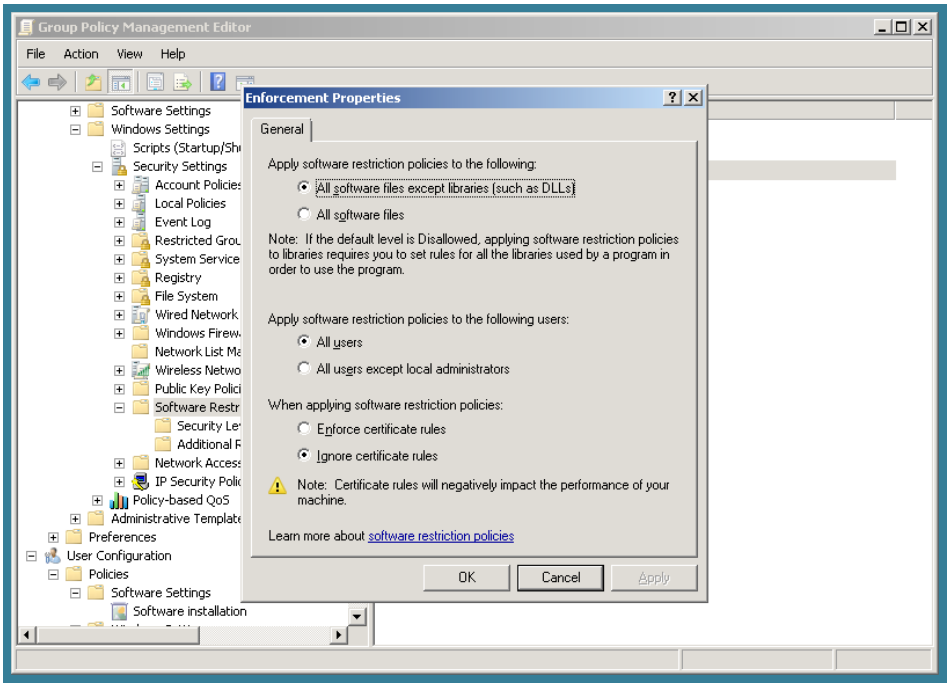

Figure 1

Two default rules ensure that Windows itself is permitted to execute as well as anything installed in the Program Files folder.

Do you have an IT question you'd like answered? Send your question to answers@realtimepublishers.com. We cannot provide personal replies but will select questions with broadest appeal for inclusion in a future issue.

Don Jones is a Series Editor for Realtime Publishers, and the Director of Training and Publishing for SAPIEN Technologies. Visit him online at www.ScriptingAnswers.com.

QUEST SOFTWARE®

Read an excerpt from the new Quest Software white paper *"Be the Master of Your Domain – Understanding Windows Server 2008 Active Directory Domain Services."*

—by Tony Murray.
Directory Services MVP

Microsoft recently announced the release of Windows Server 2008 RTM. Codenamed "Longhorn," this latest version of the server operating system from Microsoft marks a significant departure from its predecessors. This white paper introduces the changes made to Active Directory in Windows Server 2008 and addresses the impact for organizations with Active Directory already installed.

New Forest and Domain Functional Level

In Windows Server 2003, Microsoft allowed administrators to set the functional level of the domain or forest to a specific value, assuming certain conditions were met. The available AD functionality was determined by the functional level. For example, a Domain Functional Level of 2 (Windows Server 2003) permitted domain controller renames, updated and replicated last logon time stamp attribute and certain other features not available with other levels. Similarly, a Forest Functional Level of 2 allowed for cross-forest trusts, domain renames, and so on. Windows Server 2008 provides a new level: 3 (also known as Windows Server 2008).

Domain Functional Level 3

Domain Functional Level 3 provides the following features:

- All features from the Windows Server 2003 domain functional level
- Distributed File System Replication support for SYSVOL, which provides more robust and detailed replication of SYSVOL contents
- Advanced Encryption Services (AES 128 and 256) support for the Kerberos protocol
- Last Interactive Logon Information, which displays the time of the last successful interactive logon for a user, the workstation used, and the number of failed logon attempts since the last logon
- Fine-grained password policies, which make it possible for password and account lockout policies to be specified for users and global security groups in a domain

Forest Functional Level 3

Forest Functional Level 3 provides all the features available at the Windows Server 2003 forest functional level, but no additional features. The sole purpose of raising the forest functional level to 3 is to prevent any new downstream domains or domain controllers from being joined to the forest.

One other item to note is that Windows Server 2008 domain controllers can be added to domains at Functional Level 2 (i.e., Windows Server 2003). If your organization has an urgent need for read-only domain controllers (RODCs), you can deploy them into your existing Windows Server 2003 forest without having to first upgrade your existing domain controllers.

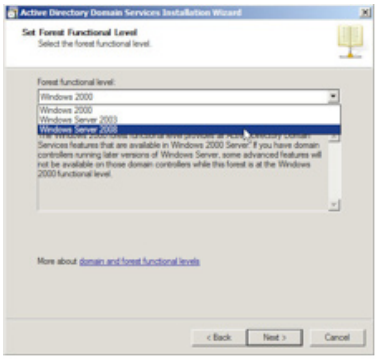

Read-only Domain Controller (RODC)

The introduction of the RODC in Windows Server 2008 may represent the biggest change to AD since its Windows 2000 inception. The RODC is intended to reduce the risk of security compromises in locations where the threat is highest (such as a perimeter network) or where the physical security of the domain controller is not optimal (for example, a branch office).

Unlike standard domain controllers (called writable domain controllers), the RODC does not replicate any changes to other domain controllers. This means that an attacker cannot use a compromised RODC to gain control of the forest by replicating permissions or schema changes. An attacker would be limited to using a compromised RODC to gain access to data held within the local credential cache. To reduce this risk, administrators have the ability to configure the RODC to cache only the password hashes of the accounts that will actually use the RODC for authentication.

The risk of compromises can be reduced even further by installing the RODC in combination with the Server Core version of Windows Server 2008. This effectively lowers the surface area for attack and reduces the patching requirements.

The RODC also supports the filtered attribute set, a new feature that enables administrators to define a set of attributes with values that do not replicate to RODCs in the forest. An example would be an application that uses certain attributes to store credential information for authentication to the application. If these attributes are added to the filtered attribute set, their values are replicated between writable domain controllers as normal, but they are not replicated to your RODCs.

About Quest Software, Inc.

Quest Software, Inc. delivers innovative products that help organizations get more performance and productivity from their applications, databases, Windows infrastructure and virtual environments. Through a deep expertise in IT operations and a continued focus on what works best, Quest helps more than 90,000 customers worldwide meet higher expectations for enterprise IT. Quest Software can be found in offices around the globe and at **www.quest.com**.

Want more? Read the entire white paper at **www.quest.com/GetActive**

Quest-MasterofYrDomain-13336-Q22008

Product Review

Lieberman Random Password Manager

by Don Jones

A major bane in the existence of every Windows administrator is password management. Not necessarily domain passwords, which are easily and centrally changed in Active Directory (AD), but rather all those local account passwords, especially the sensitive local Administrator accounts located on every Windows client and member server. Lieberman's Random Password Manager, now in version 4.0, aims to make life easier by automating the management of these passwords.

Installation and initial configuration of Random Password Manager is straightforward. Helping to simplify deployment, it does not require you to deploy agents to targeted systems. You will need to provide it with the credentials necessary to initially log on to each system and change passwords for the Administrator account. The Windows-based administration console, shown in Figure A, is easy to use and allows you to define the systems that will be managed by the tool. Interestingly, the tool by default randomizes Administrator and Guest accounts, and can locate those accounts by SID rather than name, so if you've renamed those accounts, the tool will still work properly.

Essentially, the tool goes out and changes your existing local Administrator account passwords to new passwords consisting of random characters. When an authorized individual needs to use one of these sensitive accounts, he or she logs onto the

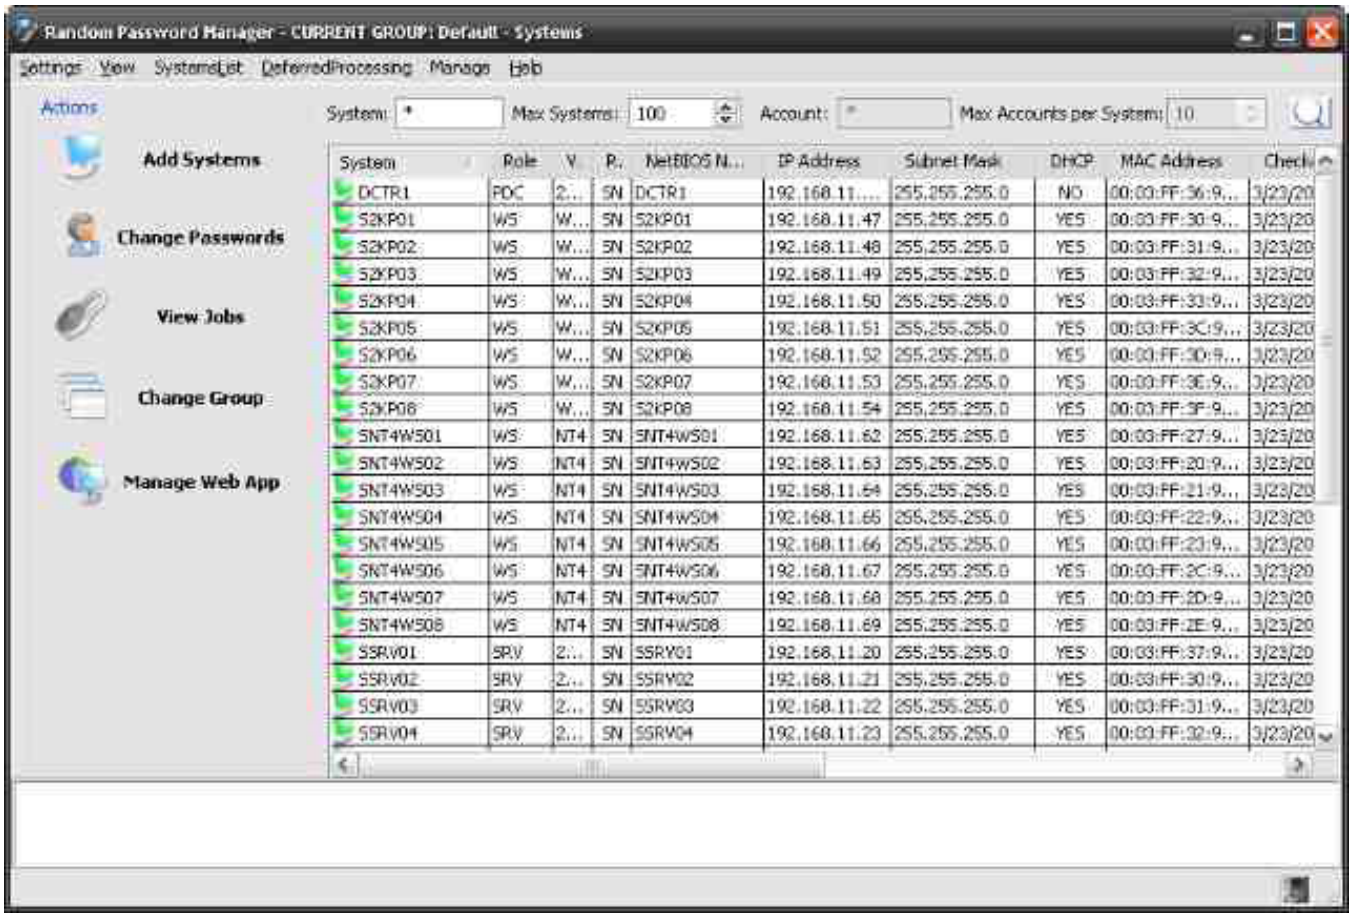

Figure I

Random Password Manager application and retrieves the current password. The tool can be set to automatically change the password after it's been retrieved, which provides an important control point for access to these important accounts.

The tool isn't limited to just Administrator passwords: You can have it manage any accounts. It's also not limited to local Windows accounts, and can manage passwords for Linux and UNIX computers, Microsoft SQL Server, MySQL, Oracle, OS/390, AS400, Juniper, and Cisco. You can also store passwords for other devices or applications, although the tool will not automatically manage these for you.

Passwords can be retrieved via a simple Web-based console, making passwords easy to access from anywhere in the organization. You can require a comment for password recovery, allowing you to track the purpose for each recovery. Email alerts can also be sent when a password is recovered, and you can configure the amount of time that a password can be "checked out;" at the expiration of that time, the password is automatically re-randomized.

The current version also supports hardware-based encryption (using PKCS #11-interface hardware), achieves FIPS 140-2 levels 2 and 3 encryption, and utilizes two-factor authentication for password retrieval. In fact, if your organization uses SecurID hardware tokens, you can have the tool require their use in order to look up passwords.

Interestingly, you can configure Random Password Manager to enable a "Personal Vault," a means by which normal users can utilize the tool's Web interface to store their own Web site, application, and other passwords. This encourages users to use different passwords for different services, and ensures they can always retrieve their passwords. Of course, the tool does not randomize these passwords; it simply makes them easier to remember. You can also enable "Self Recovery," a feature that requires you to map users to the computers they use, and which allows users to retrieve the passwords for their own computers. This would allow, for example, a software developer to retrieve her local computer Administrator password—and would allow the tool to automatically track such retrievals and to re-randomize the password afterward. The Web portal has a specific view optimized for use with Windows Mobile devices.

Pricing is based on the number of systems managed, with a minimum of 100; at $29 per system; that's a minimum price of $2900. Discounted pricing is available for larger volume purchases, and the initial purchase includes 1 year of product upgrades. The license permits the installation of as many administrative consoles as needed. Random Password Manager requires Microsoft SQL Server; a customized edition of Microsoft's SQL Server Express is available for an additional fee, which is less than a full SQL Server license and provides full reporting capability. The Web portal utilizes IIS, rather than having its own embedded Web server, and the admin console can automatically install the Web portal to a designated IIS Web site. ◆

Leveraging the New Group Policy Preferences in Windows Server 2008

by Darren Mar-Elia

With the release of Windows Server 2008, Microsoft has upped the ante in terms of what you can configure on a Windows system using Group Policy. They have done so by releasing the technology formerly called DesktopStandard PolicyMaker (Microsoft acquired DesktopStandard in 2006) as a new set of Group Policy extensions known as **Group Policy Preferences**. In this feature, we will look at how you can leverage Group Policy Preferences to greatly enhance what you can bring under control using Group Policy. We'll show you the support platforms for this new feature, define what's required to leverage it, and run through some of the highlights of the technology. By the end of this article, you should have all the tools required to put this powerful technology to work in your own environment.

What's Required to Run Group Policy Preferences?

Microsoft released the Group Policy Preferences feature as a built-in capability within Windows Server 2008. That is, when you install Windows Server 2008, you will already have everything you need to both edit and process Group Policy Preferences features on that Server 2008 box. But what if you want those Group Policy Preferences features to apply to your Vista, Server 2003, or XP systems? No problem! Group Policy Preferences support all those platforms for the processing of these new settings. What you will need to do, for those other platforms, is download an update from Microsoft's Web site that needs to be installed on every system on which you want to receive Group Policy Preferences settings. This update provides the "Client Side Extensions" to Group Policy that allows these systems to read and process the new Group Policy Preferences settings that you define. Table 1 provides the download information for each of the platforms supported.

When you install Service Pack 1 (SP1) on Windows Vista, if you have not already installed the Group Policy Preferences Client Side Extensions, you won't need to download a separate package to enable this feature. SP1 already includes the Client Side Extensions required for Group Policy Preferences.

Platform	Download Location
Windows XP, Service Pack 2 and later	http://www.microsoft.com/downloads/details.aspx?FamilyID=E60B5C8F-D7DC-4B27-A261-247CE3F6C4F8&displaylang=en *
Windows Server 2003	http://www.microsoft.com/downloads/details.aspx?FamilyID=BFE775F9-5C34-44D0-8A94-44E47DB35ADD&displaylang=en *
Windows Vista	http://www.microsoft.com/downloads/details.aspx?FamilyID=AB60DC87-884C-46D5-82CD-F3C299DAC7CC&displaylang=en

Table 1: Download locations for Group Policy Preferences Client Side Extensions.

***NOTE:** For Windows XP and Server 2003, you will also need to install the "XML Lite" installation package found at http://www.microsoft.com/downloads/details.aspx?familyid=D7B5DC81-AD14-4DE2-8AD5-8C4A9AAB5992&displaylang=en for each client that needs to process Group Policy Preferences.

Editing Group Policy Preferences

The previous information relates specifically to enabling the ability for client systems (servers and desktops) to process the new Group Policy Preferences. In order to create those new settings, there are a few restrictions to be aware of. First, and most important, you can only see and edit Group Policy Preferences settings from systems running Windows Server 2008 and Windows Vista SP1. Vista without SP1 won't work. Neither will Windows XP or Server 2003. You can only create, edit, and report on Group Policy Preferences from Server 2008 or Vista SP1. On Windows Server 2008, you will need to install the Group Policy Management Console (GPMC) feature in order to get the tools necessary to edit Group Policy

Preference settings. On Vista SP1, you will need to install the separate Remote Server Administration Tools (RSAT) download (available at http://www.microsoft.com/downloads/details.aspx?FamilyID=9ff6e897-23ce-4a36-b7fc-d52065de9960&DisplayLang=en). Once RSAT is installed on your Vista SP1 system, you will need to explicitly install the GPMC feature in order to gain access to the Group Policy Preferences features.

To perform that install, start the "Programs and Features" applet from the Vista Control Panel. Select the "Turn Windows Features on or off" link from the left side of the applet. Once the Features dialog box appears, scroll down to "Remote Server Administration Tools" and expand the "Feature Administration Tools" folder below it. Select the "Group Policy Management Tools" check box, and click OK to install the GPMC and related tools, as shown in Figure 1.

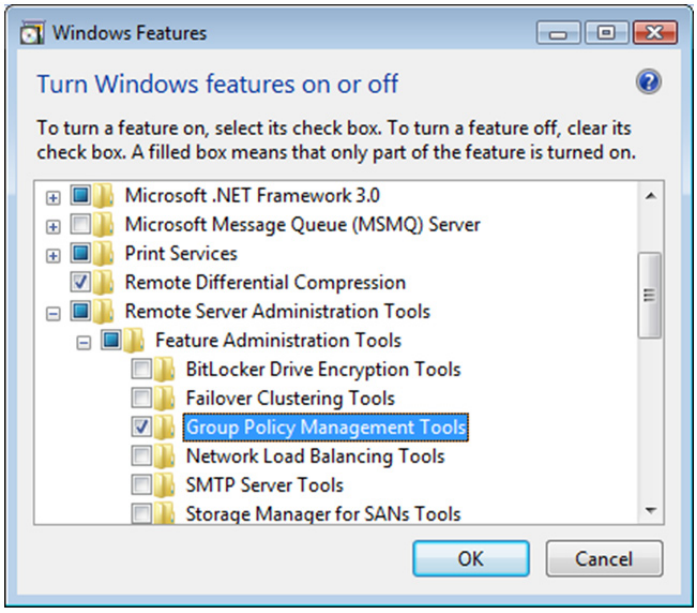

Figure 1: Installing the GPMC feature on Vista SP1.

Once GPMC is installed on Vista SP1, you can use the GPMC to view and edit Group Policy Preferences settings just like any other Group Policy settings. Let's take a look now at how you can leverage this new Group Policy technology.

Creating Policy Using Group Policy Preferences

Once you've installed the requisite bits on the administrative side to edit Group Policy Preferences from Windows Server 2008 or Vista SP1, it's a short step to begin using the technology. Simply create a new GPO or select an existing one from the GPMC. Once you've got the GPO you want to edit, right-click it, and choose Edit. You will immediately notice a difference in what you see in the Group Policy Management Editor compared with the pre-Group Policy Preferences days. You will still see the top-level Computer Configuration and User Configuration nodes, but you will notice new node names directly under those two top-level nodes. Namely, you will see two containers named Policies and Preferences, as shown in Figure 2.

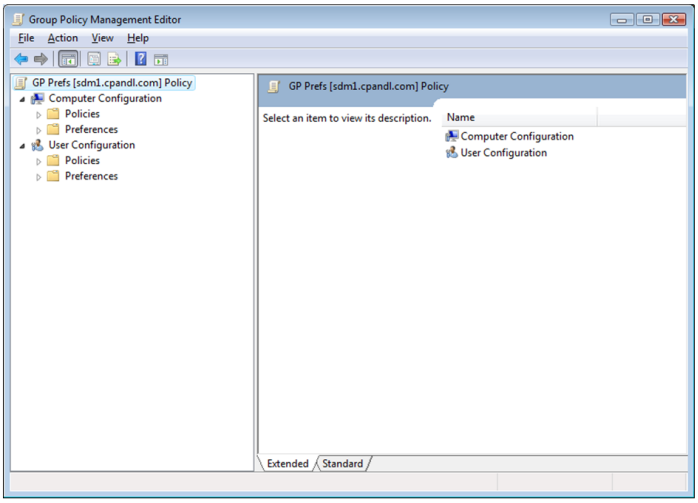

Figure 2: Viewing the new Group Policy Management Editor with Group Policy Preferences.

As you can imagine, "original" policies that you are used to seeing can be found under the Policies container, while the new Group Policy Preferences features are located under the Preferences container. If you expand the Preferences node, you will see two sub-folders—Windows Settings and Control Panel Settings. These two folders help organize the two main areas where Group Policy Preferences add functionality, namely general Windows configuration settings such as environment variables and shortcuts, and Control Panel elements such as power configuration, scheduled tasks, and local user or group creation.

Now that we know how to find the new features, let's walk through using one of them to see how it works. There are a lot of new configuration capabilities that Group Policy Preferences brings. In fact, there are too many to detail here, but in general, they all work the same basic way, albeit each with different configuration options. The first consideration about Group Policy Preferences is that they are just that—preferences. What does that mean in the context of Group Policy? It means that unlike regular policies that you are used to, Group Policy Preferences settings are processed by clients but the user is not prevented from modifying the setting after Group Policy has delivered it.

Ultimately, what this means to the user is that, unlike true Policies, Preferences are not "greyed out" in the associated UI to prevent the user from changing the policy. Now, there are ways to mitigate this fact within the Group Policy Preferences dialog boxes, but this fact is important to keep in mind as you are deploying Group Policy Preferences. That being said, let's walk through implementing some Preference settings to see how it works.

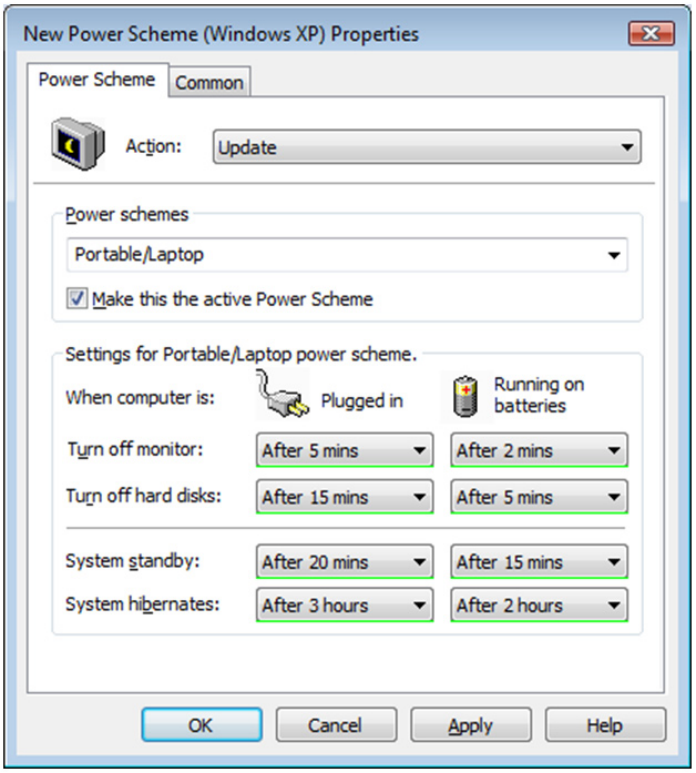

Figure 3: Choosing power scheme options.

Example: Using Group Policy Preferences Power Options

In this example, I'm going to create per-computer power settings for my Windows XP clients. Group Policy Preferences added support for power management into Group Policy for Windows XP, which is a great thing in these times of energy consciousness. Windows Vista systems already include power management in "regular" Group Policy, so I'm going to create this policy for managing power on my Windows XP systems only. The first thing I need to do is locate the policy. In this case, it's under **Computer Configuration \ Preferences \ Control Panel Settings \ Power Options** within the Group Policy editor. To create a power policy, right-click the Power Options node, and select New, Power Options (Windows XP) or Power Scheme (Windows XP). The Power Options section lets you set general power options while the Power Scheme section lets you choose which scheme is active and what its behavior looks like. Let's choose the Power Scheme option to set up a default power scheme for the laptop computers in my environment. Once I choose that option, I'm presented with a dialog box similar to that shown in Figure 3.

From this dialog box, I can choose the scheme I want to make active on the target computers (in this example, it's the Portable/Laptop scheme), and then choose the behavior settings for this scheme. Once I've got the scheme configured the way I want it, I can adjust the behavior of this policy by selecting the "Common" tab, as shown in Figure 4.

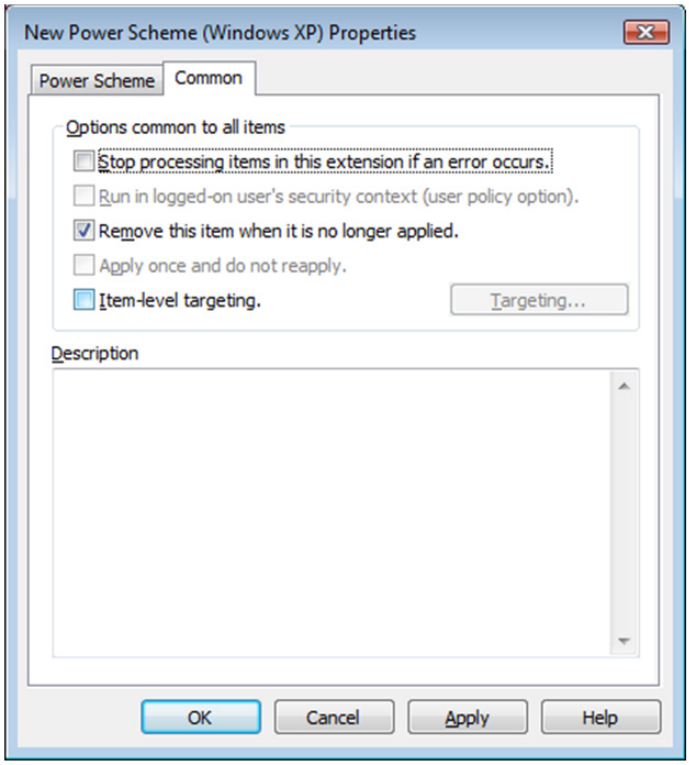

Figure 4: Viewing the Common tab in a Group Policy Preferences setting.

The Common tab is unique to Group Policy Preferences, and lets us change the default behavior of the setting being managed. For example, in Figure 4, I've selected the check box to remove the setting when the GPO no longer applies to the target computers. This is the default behavior for normal Group Policy Policies but not for Preferences. Thus, if we want to get that "non-tattooing" behavior for Group Policy Preferences, we need to select this check box. Another feature of Group Policy Preferences is **item-level targeting**. In normal policy, we have only a couple of ways to control which computer or user gets a policy. We can link a GPO, security filter it, or apply a WMI filter to it. However, in Group Policy Preferences, each **setting**

within a policy can have its own item-level filter. You can imagine that this capability can quickly get out of hand, but there are certain circumstances where you may want this additional level of filtering control. For example, in my Power Options policy, wouldn't it be nice to filter this power scheme such that only laptop computers received it? To do so using WMI filter may be difficult, but by selecting item-level targeting, I can very easily select a criteria (for example, selecting laptop computers only) that pretty much guarantees that only laptop computers will process this policy, as shown in Figure 5.

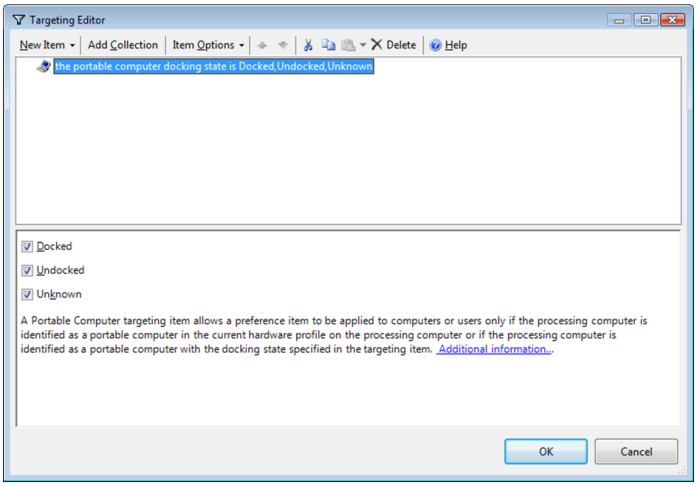

Figure 5: Targeting laptop computers using item-level targeting.

Group Policy Preferences provide a dizzying array of item-level targeting options, and you can group multiple targeting criteria using AND, OR, IS, or ISNOT operators from the Item Options menu. This provides a large degree of flexibility when targeting individual settings, but again, you should proceed cautiously with the use of item-level targeting because each of these filters must be evaluated by the client receiving the policy during policy processing. This could lead to performance issues if you have too many item-level targets applied to a given computer or user.

Summary

Group Policy Preferences provides additional configuration capabilities for Group Policy. It greatly expands what you can configure via policy to include features such as power management, USB device restrictions, drive and printer mapping, shortcut distribution, and much more. As long as you have a Windows Server 2008 or Vista SP1 system to manage this new feature, your existing Windows XP, Windows Server 2003, and Vista clients can take advantage of it to give you unprecedented control over your Windows desktop and server environments.

Darren Mar-Elia is CTO & Founder of SDM Software, Inc., a Windows Group Policy management solutions company. He maintains the popular Group Policy resource web site at www.gpoguy.com and has been a contributing editor for Windows IT Pro Magazine *since 1997. He has written and contributed to twelve books on Windows including, most recently, the* Windows Server 2008 Security Resource Kit Guide, the Windows Group Policy Guide, *published by Microsoft Press in 2005, and* The Definitive Guide to Windows 2000 Administration, The Definitive Guide to Windows 2000 Group Policy, *and* The Tips & Tricks Guide to Group Policy, *all published online by Realtimepublishers.com.*

Don't Leave Your Group Policy Health to Chance.

ELIMINATING RISK IS AS SIMPLE AS SPOTTING RED FROM GREEN.

Your Group Policy health is not something to gamble with. You use it to configure critical security and lockdown settings on your Windows systems. Why leave its performance to chance?

The GPExpert™ Troubleshooting Pak gives you four products that work together to ensure that all the moving parts in your complex Group Policy infrastructure are working properly. Collect information from a variety of sources with the GPExpert Health Reporter and get a quick visual "red or green" health status so you know when something is dicey. When trouble strikes, use the GPExpert Log Analyzer to pinpoint the source of your Group Policy problem. Or, use the GPExpert Group Policy Spy to watch registry policy activity in real time to find conflicts. Finally, give your help desk staff the GPExpert Status Monitor, running on your users' desktops, to easily see when Group Policy is functioning correctly on their machine.

With the GPExpert Troubleshooting Pak you never have to leave Group Policy health to chance. Visit **http://www.sdmsoftware.com/products** today to get your free trial.

sdmsoftware

The Group Policy Experts

The Deep Dive

Easing the Jump to 2008

by Greg Shields

In this journal as well as IT publications everywhere, you're probably hearing a lot of talk about the features and functionality that compel an upgrade to Windows Server 2008. There's plenty to talk about. Upon last check, there are already more than 30 books on the topic available on Amazon. Through the sheer mass of fallen trees alone, one can gather the conclusion that there's a lot of excitement about Server 2008 in our industry.

Obviously, before your environment makes the decision to start an upgrade, there's likely to be a vetting period. Testing, evaluation, learning, and finding the value in a server OS upgrade can be a long and heavily political process. Your organization has to recognize value in the updates and changes, and make the financial and personnel decisions necessary to start the project. The bad news is that is the hard part.

The good news is that once the decision is made, getting the server OS installed to available hardware is relatively easy. With Windows Server 2008, Microsoft brings to the table three major ways in which to install the OS: the traditional manual method, scripted, and image-based installations.

Old School Manual

If you sit back and think on the tens or hundreds of Windows servers you've installed over your IT career, have you ever wondered how many hours, or days, or years you've spent sitting in front of its installation screens. The traditional mechanisms for installing a Windows server involved a substantial amount of time to go from zero to installed. Answer a few questions, wait, answer a few more questions, wait a little more. Lather, rinse, repeat.

With Windows Server 2008, Microsoft has taken a long and hard look at its manual installations and their unfriendly process steps. Initiating an installation of Server 2008 can happen with as few as seven mouse clicks, all of which occur at the beginning of the installation. Thus, to install a Server 2008 instance, you need only insert the disk, boot to the Windows Pre-installation Environment (WinPE), answer a few quick questions, and your involvement with the installation is complete.

All this streamlining of the actual installation does come at a cost. The questions that you used to answer in previous versions still require answering. With Server 2008, they have been relocated to after the installation completes. After the post-installation reboot, you are asked to login to the server and change your initial administrator password. You are then presented with the Initial Configuration Tasks wizard, seen in Figure 1. This wizard brings into one place the rest of the initial configuration questions that you used to answer during the installation.

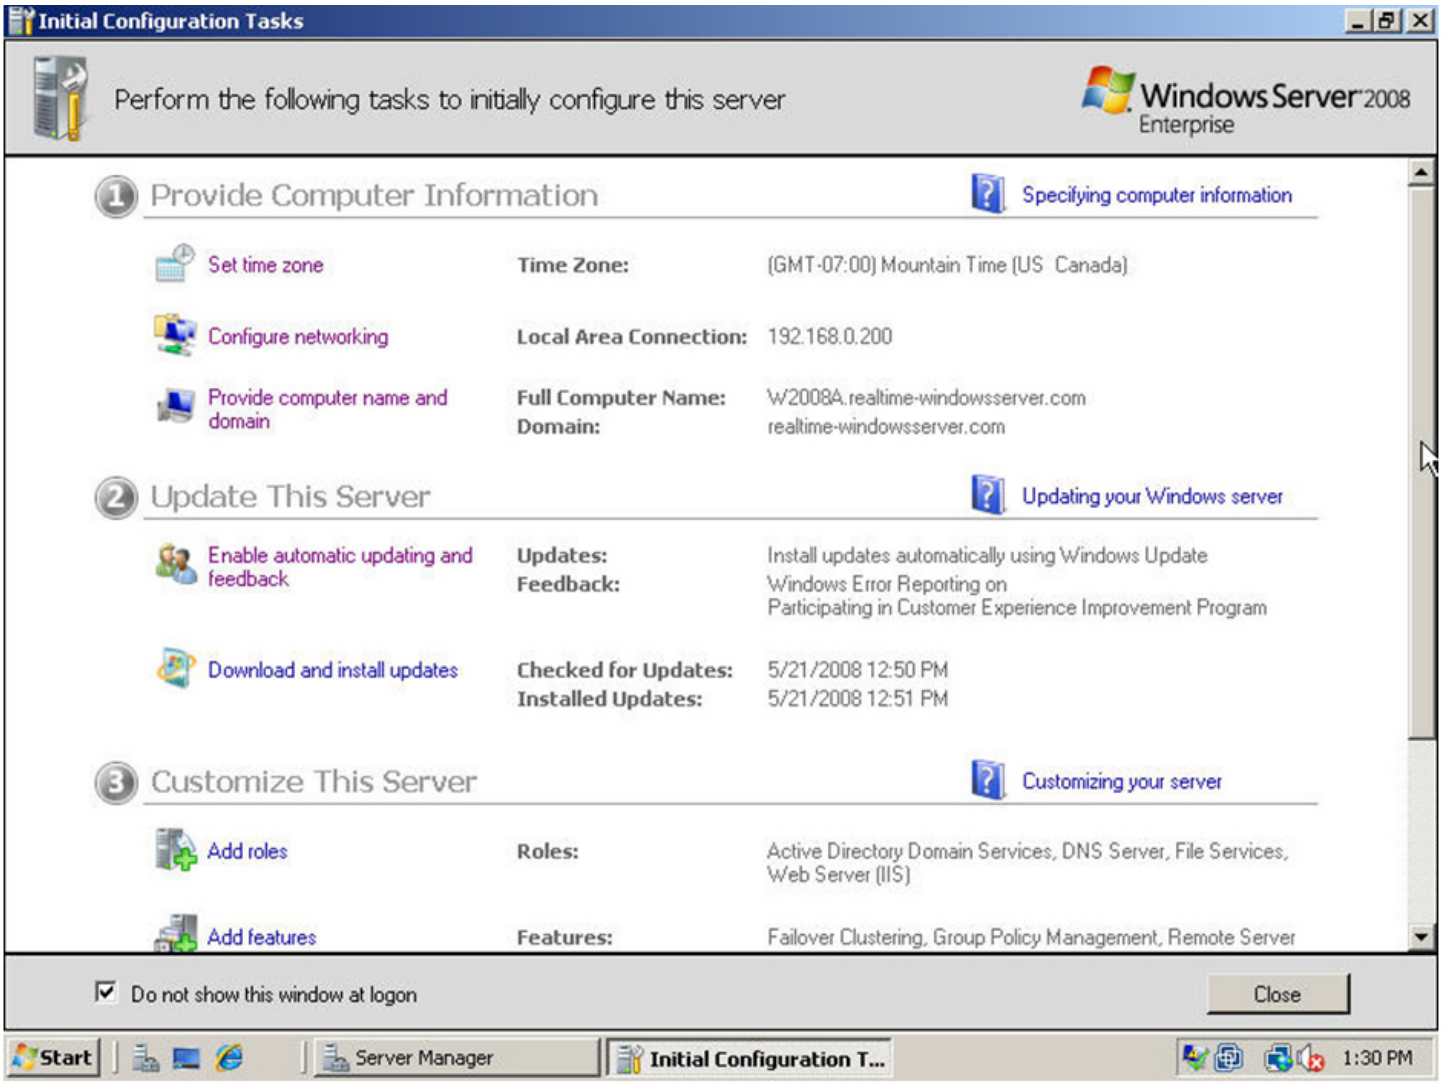

Figure 1: Initial Configuration Tasks replaces the mid-installation questions of previous OS versions.

As you can see in Figure 1, Initial Configuration Tasks provides one location for configuring time zone, computer name and domain, automatic updates, roles and services, remote desktop, and the Windows firewall. Once this screen is closed, the post-installation assignment of server roles is then done through Server Manager.

> If you accidentally close the Initial Configuration Tasks wizard and need to bring it back, run the command oobe from the command prompt.

A Line on Scripted Installations

Even with the better-tuned manual installation method available with Windows Server 2008, there comes the time that you want even more automated control over installation settings. In these cases, Microsoft also makes available the Windows Automated Installation Kit (WAIK). This large download from the Microsoft Web site includes numerous WinPE boot images that can be used to boot servers of various processor classes. It also includes a graphical tool that eases the process of creating scripted installations. This tool, called the Windows System Image Manager and shown in Figure 2, has the power to interrogate Windows installation media to determine the possible settings that can be configured during installation. With this information in hand, the tool then gives you the ability to set your desired customization for settings of interest.

Once you've determined the settings you want to include as part of your scripted installation, WSIM produces a file named *autounattend.xml* that contains your customizations. Simply copy this file to removable media and insert it into the server with the Windows media as you start the installation. The server will recognize the file and use its information to configure the server.

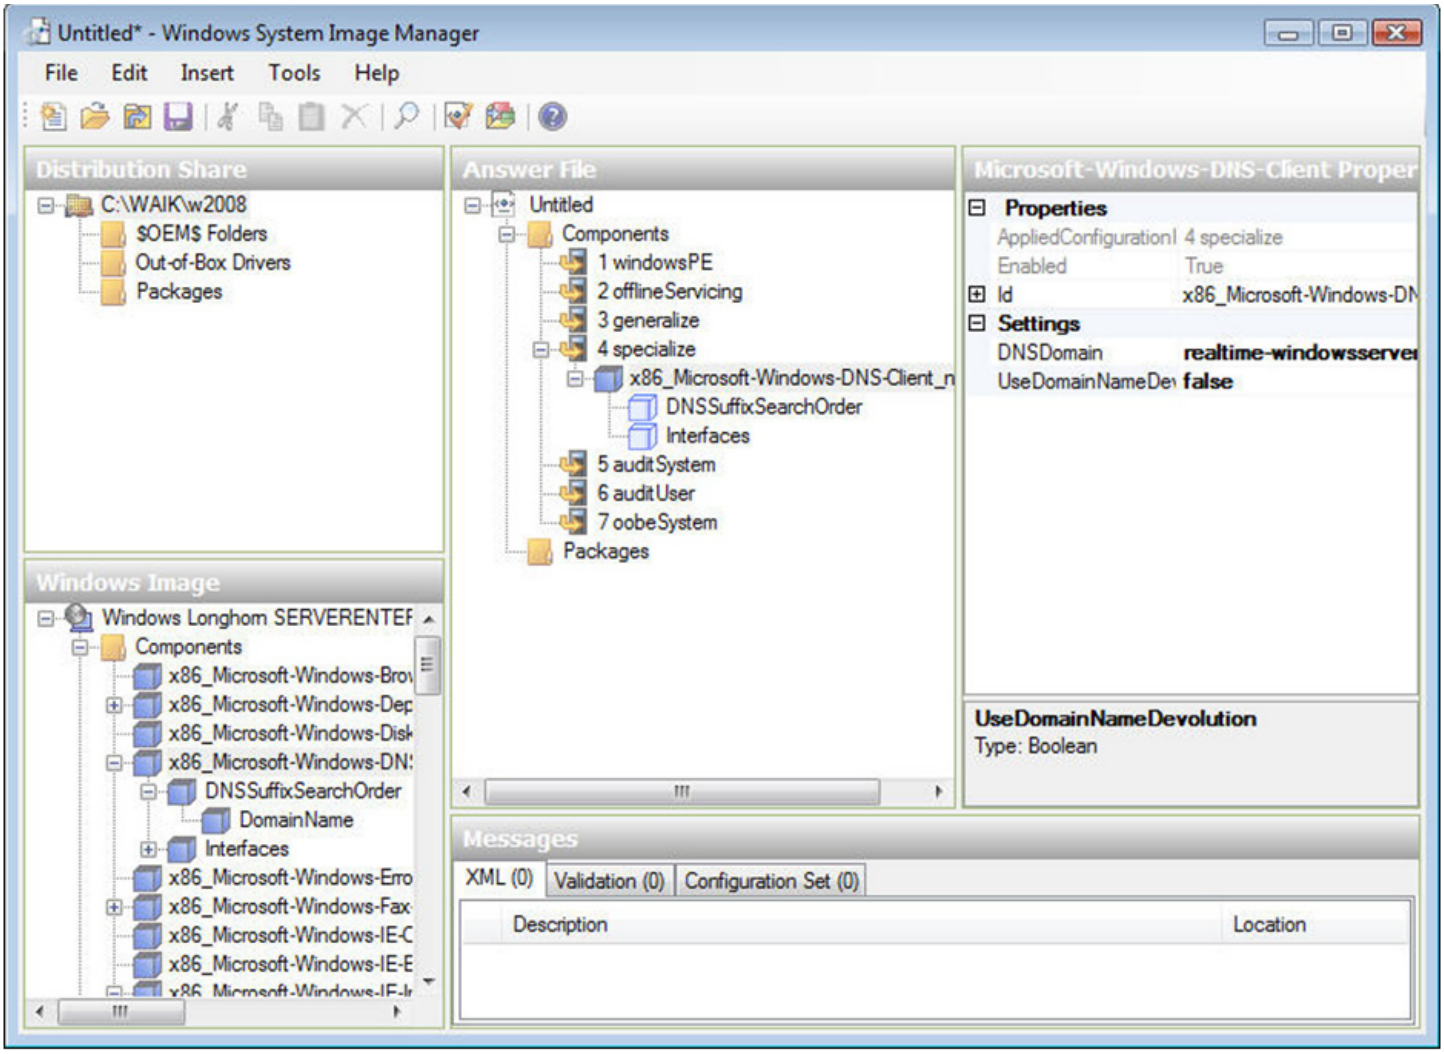

Figure 2: WSIM provides a central clearinghouse for all the possible settings that can be configured at the time of installation.

The downside of WSIM is that using it is a complex, multi-step process. In Chapter 1 of *The Definitive Guide to Building a Windows Server 2008 Infrastructure* (which you can download from free from http://nexus.realtimepublishers.com/DGBWS2K8I.htm), I explain this process in more detail. To quote directly from that chapter, the following steps show you how to use the WSIM to create a scripted installation:

1. *Download and install the WAIK.* You'll want to install this onto a workstation and not the server you intend to build using the unattended installation file. The workstation you use can be your desktop.

2. *Locate and copy install.wim.* WIM, or Windows Imaging Format files, is a file-based disk image format. These files are used for booting to WinPE and ultimately installing OSs. The *install.wim* file needed by WSIM is found on the Server 2008 media in the *\sources* folder and contains the metadata information necessary to do an installation of Server 2008. Copy this file to your desktop.

3. *Launch WSIM.* Upon launching WSIM, you'll see a relatively busy screen similar to what is shown in Figure 2, though initially your screen will start with much less information. Each of the windows within the WSIM link to each other. As an example, a distribution share (in the upper-left) can have multiple Windows Images (in the lower-left). You'll interrogate the possible settings of those Windows images and configure the ones of interest using the middle and right panes. The answer file in tree format will show the configured settings. The right-most pane will include the options both available and set for a desired configuration.

4. *Create a new distribution share.* The first step with WSIM is to create a share that will contain your working folders as you create your unattended installation file. Do this by right-clicking *Select a Distribution Share*, and clicking *Create Distribution Share*. Select an appropriate folder and click *Open*.

5. *Select a Windows image or catalog file.* Next, you'll need to locate and interrogate the install.wim file you copied from the Server 2008 media. Do this by right-clicking *Select a Windows image or catalog file* and choosing *Select Windows image*. Locate your install.wim file and click Open. WSIM will search the file for available installations and provide you with a list of those available. In our case, we'll install *Windows Longhorn SERVERENTERPRISE*. Yours may be slightly different. An error will appear noting that no catalog file exists and asking if you want to create a new one. Select *Yes* to continue. The process to create the catalog file can take an extended period of time. WSIM at this point is discovering and cataloging all of the potential configuration options available.

6. *Create a new answer file.* In the center pane, right-click *Create or open an answer file* and select *New answer file* to create a new answer file to be stored in your distribution share. The terms "answer file" and "unattended installation file" here are used interchangeably. You'll notice that the answer file has seven components. Each of these components, called a "pass" is one phase of the installation process. Depending on the configuration you intend to include, that configuration may be set in one of the seven possible passes.

7. *Configure the settings for your unattended installation file.* The lower-left pane will now include the list of options available for configuration in your answer file. Here's where the hard part begins. Discovering which of the options are necessary and interesting is a time-consuming process. You'll want to dig through the components to find configurations of interest. For example, if you want to set the DNS suffix search order, you can right-click *x86_Microsoft-Windows-DNS-Client_6.0.6001.16659_neutral* and select *Add setting to Pass 4 specialize*. Then highlight the new setting under *4 specialize* and in the upper-right screen look for *DNSDomain*. There, you'll want to enter the correct value for suffix search order. Continue with setting your desired configurations as necessary.

8. *Validate the answer file.* Due to linkages between settings, some settings require others to also be set. Once you have completed setting those of interest to you, you'll want to validate the answer file to ensure that it includes the necessary components. Do this by clicking *Tools* and then *Validate Answer File*. Saving the answer file also completes a validation prior to the save.

9. *Save the answer file.* From the *File* menu, select *Save Answer File*. Save the file as *autounattend.xml* to use it for a server installation.

10. Copy the file to removable media and boot from the Server 2008 media. Once you've completed your autounattend.xml file, you can copy it to the root of your removable media and use it to boot a candidate system along with the Server 2008 DVD. If you've done everything properly, the installation should complete and a fully-configured server will be the result.

WSIM is rich in the level of configuration control it provides. But that richness also means that there are a lot of settings to wade through to find just what you want. Finding the bare minimum of settings needed to successfully create a scripted installation can also be a challenge, so thorough testing will be required for you to find what works for your environment.

The third mechanism for getting Windows installations from media to hardware is using image-based installations, not unlike what you may have used before with tools such as Symantec Ghost. Microsoft provides both command-line and GUI tools to assist with the creation, manipulation, and deployment of images. Microsoft's GUI-based tool for doing this arrives as a Windows Server 2008 Role, namely the Windows Deployment Services (WDS) Role.

At its simplest, WDS includes the native ability to read the *install.wim* file from the Server 2008 media and deploy that installation over the network. After installing the WDS Role, right-click the server name within Server Manager, and select Configure Server. WDS' initial configuration allows you to choose a location for storing images as well as the initial PXE server response settings. These settings allow you to choose whether the WDS server will automatically respond to PXE clients that attempt to connect. Also part of the initial configuration is the ability to add initial images to the server. Point the server to the */sources* folder of your Server 2008 media and it will discover and ingest the images into WDS for later distribution. A picture of how this looks within Server Manager is shown as Figure 3.

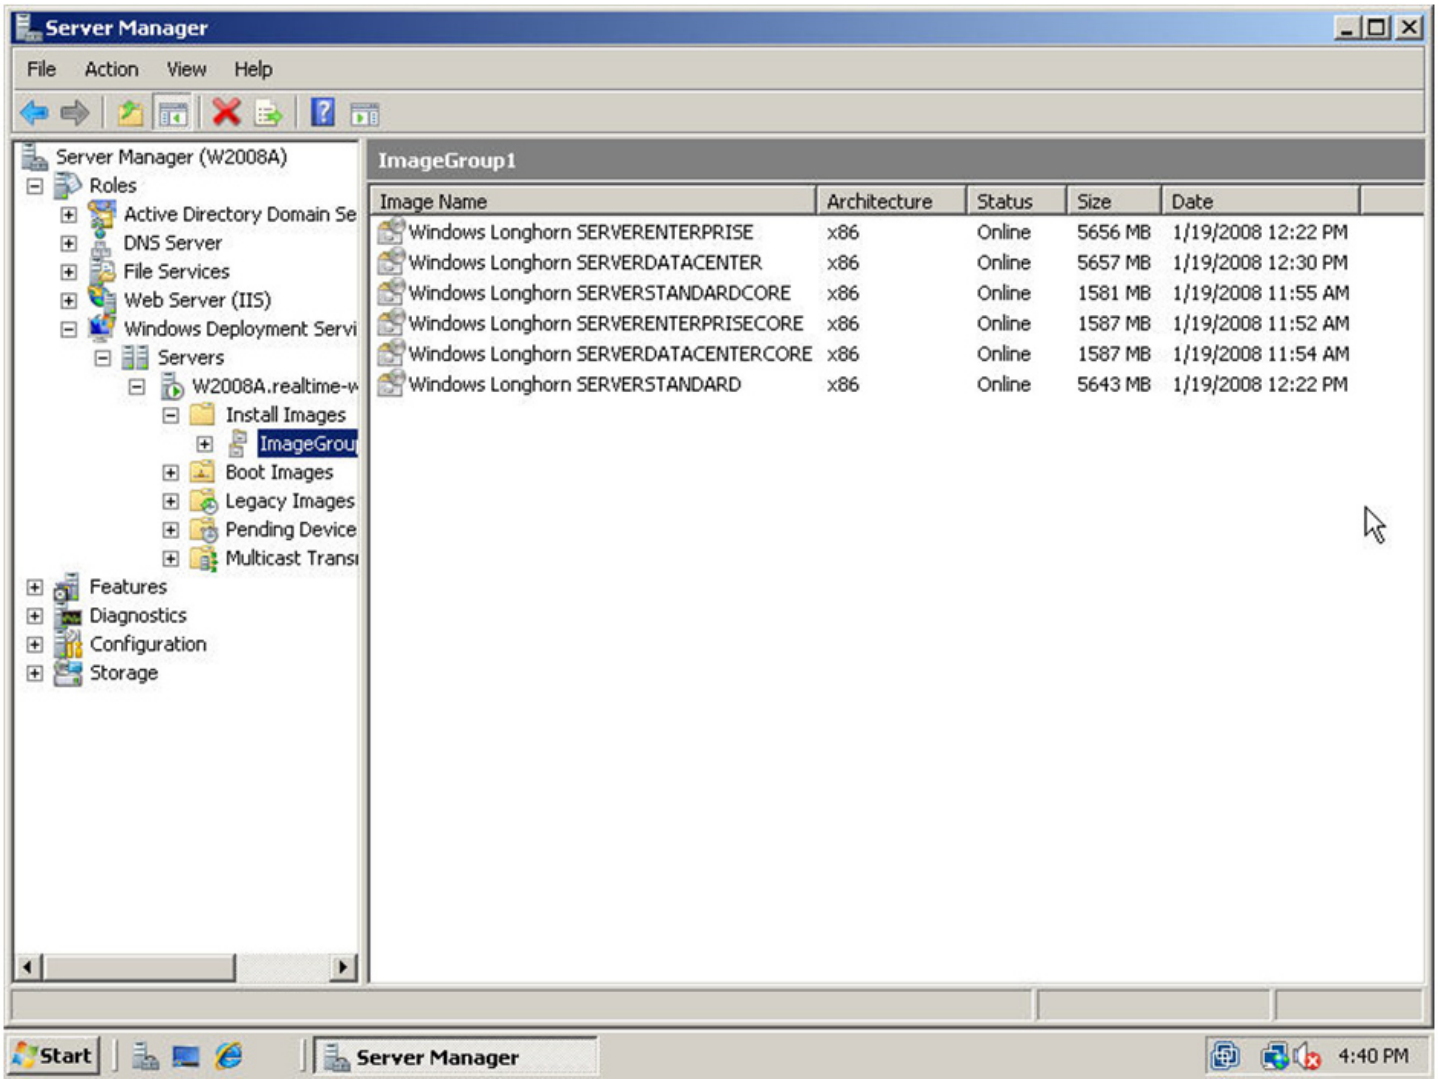

Figure 3: WDS has a simple-to-use interface for deploying OS images to multiple machines simultaneously.

Among other improvements over previous versions, WDS includes one major new capability that significantly improves its use in all environments. Multicast support adds the ability to image more than one client simultaneously. Thus, if you have 10 servers to install, you can use their PXE boot functionality to connect them to a multicast transmission on your WDS server.

To start a multicast transmission of the default Server 2008 installation, simply right-click any of the images you just added, and select Create Multicast Transmission. In the resulting wizard, you'll be asked for a friendly name for the transmission and information about the type of transmission. Transmissions can occur either immediately or based on a number of connected clients or countdown timer.

Extending the customization capabilities of WDS are the ability to capture images from existing hardware as well as the ability to integrate *autounattend.xml* files for server customization as a component of the imaging process. Leaving the image capture process for another column, you can import an *autounattend.xml* file by viewing the properties of any image and selecting the *Allow image to install in unattended mode* check box. You'll need to choose the unattended installation file to associate with the image by clicking Select File.

You can see that you have options when it comes to installing your new server OS. With all the easy options available to you, the hard part now is getting the justification to start your upgrade project. If you've got questions or comments about this process or have a story about your own upgrade to share, drop me a line at gshields@realtimepublishers.com.

Greg Shields, MCSE: Security, CCEA, is an independent author, speaker, and consultant, based in Denver, Colorado. With more than 10 years of experience in information technology, Greg has developed extensive experience in systems administration, engineering, and architecture. Greg is a contributing editor for both Redmond magazine and MCPmag.com, authoring two regular columns along with numerous feature articles, webcasts, and white papers. He is also the resident editor for Realtime Publishers' Windows Server Community at www.realtime-windowsserver.com. Greg is currently finishing his new book Windows 2008: What's New, What's Changed through SAPIEN Press.

Practical PowerShell

Antivirus Management the PowerShell Way

by Jeffery Hicks

If you work for a midsize to large company or have a healthy IT budget, you most likely have an enterprise-wide antivirus solution. Typically, these solutions include a reporting feature so that you can tell which clients are up to date with definition files and the like. But if you don't use such a product, you probably thought you were out of luck. Well, not really, especially because we can leverage the combination of Windows Management Instrumentation (WMI) and PowerShell to gain access to this information.

Most modern antivirus products register themselves with WMI when installed. Within the root\SecurityCenter WMI namespace, you will find an AntiVirusProduct class. Armed with this information, you can easily use the Get-WMIObject cmdlet in PowerShell to query the local computer:

```
PS C:\> get-wmiobject -namespace root\securitycenter `
>> -class antivirusproduct
>>


__GENUS                          : 2
__CLASS                          : AntiVirusProduct
__SUPERCLASS                     :
__DYNASTY                        : AntiVirusProduct
__RELPATH                        : AntiVirusProduct.instanceGuid="{4…
__PROPERTY_COUNT                 : 10
__DERIVATION                     : {}
__SERVER                         : PUCK
__NAMESPACE                      : ROOT\securitycenter
__PATH                           : \\PUCK\ROOT\securitycenter:Anti...
companyName                      : Grisoft
displayName                      : AVG 7.5.524
instanceGuid                     : {41564737-3200-1071-989B-0000E87B…
onAccessScanningEnabled          : True
pathToSignedProductExe           : C:\Program Files\Grisoft\AVG7\avg…
productHasNotifiedUser           :
productState                     :
productUptoDate                  : True
productWantsWscNotifications     :
versionNumber                    : 7.5.524
```

I think you'll recognize some key properties. Let's refine our initial PowerShell expression:

```
PS C:\ > get-wmiobject -namespace root\securitycenter `
>> -class antivirusproduct | Select CompanyName,Displayname,`
>> VersionNumber,ProductUpToDate
>>

CompanyName      Displayname     VersionNumber   ProductUpToDate
-----------      -----------     -------------   ---------------
Grisoft          AVG 7.5.524     7.5.524         True
```

All I've done is select the pertinent properties. Of course, I can easily check a remote machine, even using alternate credentials like this:

```
PS C:\> $cred=get-credential "jdhitsolutions\da_jhicks"
PS C:\> get-wmiobject -namespace root\securitycenter `
>> -class antivirusproduct -computername XPDESK02 `
>> -credential $cred |
>> Select CompanyName,Displayname,VersionNumber,ProductUpToDate
>>

CompanyName      Displayname     VersionNumber   ProductUpToDate
-----------      -----------     -------------   ---------------
Grisoft          AVG 7.5.524     7.5.524         True
```

How easy was that? The Get-Credential cmdlet securely stores a PSCredential object for jdhitsolutions\da_hicks. I use this credential to connect to XPDESK02 and execute my query. For the rest of my examples, I'll use my current credentials, but if you need to specify alternate credentials, this is how it is done.

> During my testing, I discovered some odd behavior with alternate credentials and the root\SecurityCenter namespace. If you are working in a domain, alternate credentials should work just fine. Where you might run into an issue is when attempting to connect to a remote computer that isn't trusted, such as in a workgroup. Even though alternate credentials in these situations work just fine with the root\cimv2 namespace and the Win32 classes, querying the root\securitycenter namespace in PowerShell would give me an Access Denied error. Oddly enough, I could use alternate credentials with VBScript and WBEMTest. Only PowerShell was problematic, which leads me to believe there is some odd .Net problem. Although I was generally successful when using passthrough authentication—that is, the account and password I was logged in as also existed on the remote computer. In this situation, I didn't need to use alternate credentials. So where does this leave you? If you are in a domain environment and need alternate credentials, you shouldn't have a problem. For everyone else, using passthrough authentication is your best bet.

Of course, whatever I can do for one computer in PowerShell I can do for 10 or 100. First, let's examine a quick-and-dirty approach. First, we'll take the PowerShell expression we know works and save it as a script block:

```
$sb={get-wmiobject -namespace root\securitycenter `
-class antivirusproduct -computername $_  |
select __Server,CompanyName,Displayname,VersionNumber,`
ProductUpToDate}
```

Notice I've use $_ for the computername and I've also added __Server to the Select expression. The former is so that I can pipe names into the script block and the latter so that I can tell what settings belong to a particular computer. I'll test it first by piping in a single computer name:

```
PS C:\> "puck" | ForEach { &$sb $_}


__SERVER        : PUCK
CompanyName     : Grisoft
Displayname     : AVG 7.5.524
VersionNumber   : 7.5.524
ProductUpToDate : True
```

The string "puck" is piped to ForEach, which executes the script block, passing the piped name as $_. Once I know it works, scaling it out is as simple as using this expression:

```
PS C:\> get-content computers.txt | foreach {&$sb $_}
```

If all I'm interested in are computers that aren't up to date, I can create a report by filtering using Where-Object and sending the results to a file.

```
PS C:\> get-content computers.txt | foreach {&$sb $_} |
>> Where {!$_.ProductUpToDate} | Select __Server |
>> Tee-Object C:\AVReport.txt
```

Because the ProductUpToDate property is Boolean, I can check for NOT true using the ! character in my Where expression. I select the server name for computers that meet this expression, and pipe the result to Tee-Object. This lets me see the output and save it to a text file.

If your network is small and you can be assured that all the computers in your list are up and accessible, what we've looked at so far may be sufficient. However, it doesn't handle connection errors gracefully, limits what information is returned, and could be revised to better use the PowerShell pipeline.

> Listing 1 offers a more robust solution which you can download the code sample from: http://www.realtime-windowsserver.com/code/v1n6_Practical_PowerShell.zip.

```
Function Get-Antivirus {
    Param([System.Management.Automation.PSCredential]$credential)

BEGIN {
        #set to Continue if you want to see Warning messages
        $WarningPreference="SilentlyContinue"

        #Set to Continue if you want to see Debug messages
        $DebugPreference="SilentlyContinue"
```

```
        #Set to Continue if you want to disable the
        #function's error handling
        $errorActionPreference="SilentlyContinue"

      Write-Debug "Starting Get-Antivirus function"

        If ($credential) {
            Write-Debug ("Using alternate credentials:"+ `
($credential.username))
         }

    }

PROCESS {
        Trap {
           if ($_.Exception -match "RPC server is unavailable") {
         Write-Warning "$computername is not available via RPC."
             continue
           }
         elseif ($_.Exception -match "invalid namespace") {
                Write-Warning "root\securitycenter namespace not found on $computername."
           continue
         }
         elseif ($_.Exception -match "access is denied") {
              Write-Warning "Access denied to $computername."
              continue
           }
          else {
            Write-Warning "There was an error"
            Write-Warning $_
              continue
          }
     }


        if ($av) {
             Write-Debug "Removing leftover `$av variable"
             Remove-Variable av
    }
```

```
  if ($_) {

    $Computername=$_
}
else {
    $Computername=$env:computername
}

Write-Debug “Scanning $computername”

if ($credential) {
    #use alternate credentials if supplied
    $av=Get-WmiObject -namespace root\securitycenter `
    -class antivirusproduct -computername $Computername `
    -Credential $credential -ErrorAction Stop
}
else {
    $av=Get-WmiObject -namespace root\securitycenter `
    -class antivirusproduct -computername $Computername `
    -ErrorAction Stop
    }

if ($av.CompanyName) {
    Write-Debug ($av | Out-String)
    # add computer name and ScanDate properties
    Write-Debug “Adding custom properties”
    $av | Add-Member -MemberType “Noteproperty” `
      -Name “ProductFound” -value $true
    $av | Add-Member -MemberType “NoteProperty” `
      -Name “Computername” -Value $Computername.toUpper()
    $av | Add-Member -MemberType “NoteProperty” `
      -Name “ScanDate” -Value (Get-Date)
}
else {
    Write-Debug “No Antivirus product found on $computername.”
    Write-Warning “No Antivirus product found on $computername.”

    $av | Add-Member -MemberType “Noteproperty” `
      -Name “ProductFound” -value $False
    $av | Add-Member -MemberType “NoteProperty” `
      -Name “Computername” -Value $Computername.ToUpper()
    $av | Add-Member -MemberType “NoteProperty” `
```

```
                -Name "ScanDate" -Value (Get-Date)
        }

            #write result to pipeline
            $av

        } #end Process script block

        END {
        Write-Debug "Ending Get-Antivirus function"
    }

}
```

Listing 1: The Get-Antivirus function.

Load this function into your PowerShell session. This function is written as a filtering function, which means you can pipe objects directly to it. The Begin script block will execute once before any objects are processed and the End script block runs once after every object has been processed. The Process script block executes once for every pipelined object. The $_ represents the current object in the pipeline. If you want to get a feel for how it works, try this:

```
PS C:\ "MyComputer" | Get-Antivirus
```

You should get output similar to what I showed earlier.

Let me quickly go over some function details. In the Begin scriptblock, I've configured variables that control different PowerShell pipelines. I've added code to display Warning and Debug messages. You can set these to Continue if you want to see these messages.

The function can take a PSCredential as an alternate credential, but it must be the same for all computers.

```
PS C:\ $cred=get-credential "Mycompany\Admin"
PS C:\ get-content "Desktops.txt" | get-antivirus $cred |
>> Select Computername,Displayname,VersionNumber,ProductUpToDate
```

The function will check every computer in Desktops.txt using the mycompany\admin credential.

I've added error handling for typical problems you might encounter, such as an unreachable computer or a system running an operating system (OS) that doesn't support the AntiVirusProduct class such as Windows 2003. In all instances, I'm using the same Get-WMIObject code from earlier.

```
            $av=Get-WmiObject -namespace root\securitycenter `
            -class antivirusproduct -computername $Computername `
            -ErrorAction Stop
            }
```

I'm assuming you'll be sending multiple computer names into the function, and some computers might be problematic. So, in addition to the WMI properties, I added custom properties to the WMI object such as computer name, whether a product was found, and a scan date.

```
            $av | Add-Member -MemberType "Noteproperty" `
              -Name "ProductFound" -value $true
            $av | Add-Member -MemberType "NoteProperty" `
              -Name "Computername" -Value $Computername.toUpper()
            $av | Add-Member -MemberType "NoteProperty" `
              -Name "ScanDate" -Value (Get-Date)
```

Now you can be a little more selective. For example, perhaps you want to find all computers on which an antivirus product isn't found:

```
PS C:\ get-content computers.txt | get-antivirus |
>> where {!$_.productFound} | select Computername
```

Or perhaps you'd like a report of computers on which the antivirus product is not up to date:

```
PS C:\> "computerA","computer","ComputerC" | get-antivirus |
>> Where {!$_ProductUpToDate} |
>> Select Computername,CompanyName,Displayname,VersionNumber
```

Or maybe you want to save all this information to a CSV file so that you can slice and dice it in Microsoft Excel

```
PS C:\> cat computers.txt | get-antivirus |
>> Export-Csv "c:\AVReport.csv"
```

Hopefully, you've discovered how much information you can get from WMI and PowerShell without leaving your desk. And if you are ready to take this even further, you can use this function as a template for the AntiSpywareProduct and FirewallProduct classes in the same namespace. But I'll leave that PowerShell pleasure to you. ◆

Jeffery Hicks, MCSE, MCSA, MCT, and Microsoft PowerShell MVP, is a Scripting Guru for SAPIEN Technologies. Jeff is a 16-year IT veteran. He has co-authored and authored several books, courseware, and training videos on administrative scripting and automation. His latest book is WSH and VBScript Core: TFM (SAPIEN Press 2007). You can contact him at jhicks@sapien.com.

High-Availability Solutions with Exchange 2007 SP1

by J. Peter Bruzzese

With the release of Exchange 2007, we were introduced to new forms of high availability under the continuous replication heading. Exchange 2007 pre-Service Pack 1 (SP1) offered us Local Continuous Replication (LCR) and Cluster Continuous Replication (CCR). Another form of high availability (more of an enhanced update from Exchange 2003) was offered as a Single Copy Cluster (SCC). With the release of SP1 for Exchange 2007, we have been given yet another type of high availability called Standby Continuous Replication (SCR).

That is certainly a lot to take in all at once. So, the purpose of this column is to break down the concept of high availability as well as explore how continuous replication and each of the four flavors of high-availability solutions work for Exchange 2007.

How High Availability Works

The requirement for 24/7 data availability is not new. The tools and methods we use to accomplish this goal are ever-changing, however. When we speak of high availability, we are not simply discussing server uptime, but more accurately, server uptime that allows for accessibility to server resources, in our case, messaging services.

To provide for higher levels of availability, there are both hardware and software solutions that must be utilized. For example, having redundancy of disks and servers is a key element to high availability. In the event one server goes down or a disk fails or the data on that disk becomes corrupt, there must be a secondary means of providing access to the data on that server or disk that doesn't include literally restoring from a backup. The time wasted getting the backup restored is availability lost to your users—and, by extension, money lost. So, RAID solutions, Volume Shadow Services, multiple servers, even expensive mirror servers across geographically dispersed regions to prepare for the ultimate in disasters can all be considered for a heightening of your availability expectations.

From a software perspective, there are third-party vendors such as NeverFail and Double-Take that offer solutions for your consideration. However, you might want to consider what those companies offer when compared with the free solutions directly within Exchange. Three of the four offerings in Exchange 2007 SP1 utilize continuous replication; let's explore how that works.

The Lowdown on Continuous Replication

To truly grasp what is happening behind the scenes in any of the three continuous replication options, you need to know a bit about Exchange messaging architecture. Literally, what happens when a message comes over the wire and into your email server.

The message, when it hits your server, is redundantly located in two places. One is the Exchange database file itself. Exchange can determine the mailbox for which the message is destined, and from that, it will locate the mailbox database and storage group that you have on that server. It will then place the message in the database, but it will also break that message into 1MB chunks and place these pieces into transaction logs.

> In previous versions of Exchange, the transaction logs were 5MB in size; thanks to the new high-availability solutions, it was determined that smaller logs would be better.

The recommended scenario for proper storage group/database management is for one mailbox database to be placed within one storage group. This setup prevents multiple database log files from building up within the same location. In the Standard Edition of Exchange 2007, you can have five storage groups and five databases. In the Enterprise Edition, you can have 50 storage groups and 50 databases. The proper structure of your database and transaction logs is to, first and foremost, remove them from the system drive (the C drive) and place the database on a Raid-5 disk set and the transaction logs on a mirrored drive set. Doing so will enhance performance and increase fault tolerance and recoverability for your messaging server.

With continuous replication, in each of its flavors, you tell Exchange to create a secondary copy of the database. That second copy will either go on another disk (in the case of LCR) with the same system or on another server. Once the initial database copy is created, the transaction logs will be replicated over to the 'passive' disks once they are closed, and they will be replayed into the secondary database. In this manner, the passive copy is always ready to step in as the active copy with (hopefully) minimal data-loss.

The Four High-Availability Options

When disaster strikes, before you go for your backup and begin the restore process, having one of these four free solutions in place may be just what you need to keep messaging available for your users. However, they don't all provide the same level of availability or cost.

SCC

Similar in design to your Exchange 2003 high-availability solution, SCC allows you to cluster two servers under the cluster services (so you will not only need two servers but you will also need to be running the Enterprise Edition of these servers in order to have Cluster Services function), with both working from a single storage location. The positive side here is that you have automatic failover for server failure. The negative is that the data has no high-availability solution in place for your SAN (although, most likely your SAN is already prepared with some form of RAID solution in place).

LCR

Called the poor man's cluster, LCR allows you to simply place another disk in a server and have the data 'mirror' over using continuous replication (with the transaction log shipping and replay technology we discussed earlier). The positive side is that this is the cheapest solution you can implement, requiring only an additional drive (or drives), and you can perform volume shadow copies off the passive side of the data as you like. The negative side is that you aren't running Cluster Services (which may seem like a positive), which means

you have to manually switch from one disk to the other if a failure occurs. The time it takes for you to switch is considered unavailable time for your users. This is high-er availability…but not the ultimate solution.

CCR

CCR utilizes Cluster Services from your servers but also uses continuous replication. Thus, you have both the benefit of an automatic failover solution combined with both a server and a disk redundancy for your messaging solution (as shown in Figure 1). In addition, there are features in place to ensure that even data that might not be synchronized between the active and passive sides to the cluster can be retrieved from the Hub Transport servers 'transport dumpster,' which retains email passing through the server for a period of time. The negative side is that this requires a bit more knowledge and work, especially in relation to clusters, quorums, heartbeats, and all that fun tech-talk.

SCR

The latest flavor of high availability thanks to SP1 is SCR. The concept here is that you have the same conceptual idea as LCR, but you can use the technology now to go from one server to another (rather than just from one disk to another). This allows you to use SCR to provide both server and disk redundancy. But thinking a bit beyond the obvious, you can also use SCR to replicate a storage group from a CCR or SCC cluster over to a remote location. Another positive feature to SCR is that it includes a built-in delay for replay activity (which is excellent if you are looking to prepare for database corruption scenarios where the delay could save that corruption from making its way over to your SCR copy). On the downside (for some), you can only manage SCR from the Exchange Management Shell (EMS). And, being that the activation process doesn't have clustering services running behind the scenes, you have to perform a manual intervention to activate the SCR database in case of a failure.

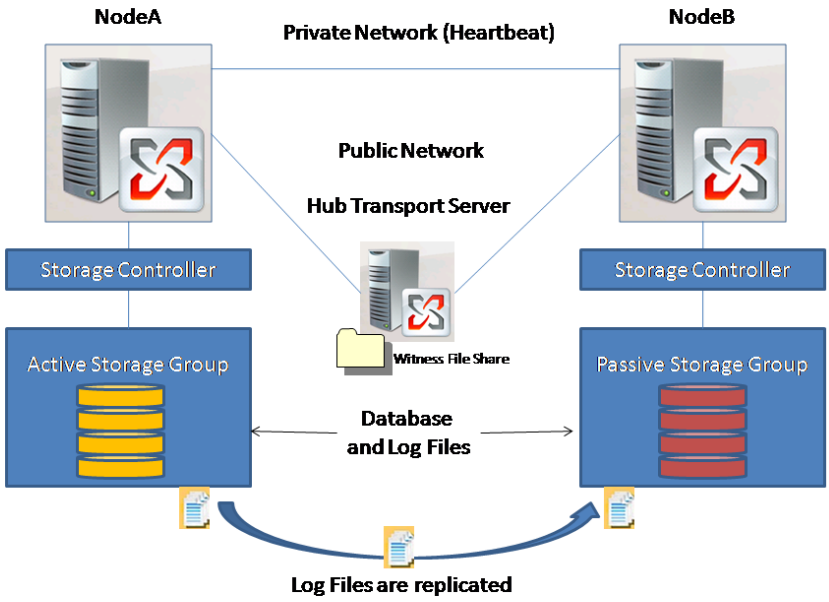

Figure 1

You can learn a lot more about high-availability solutions by going directly to the Microsoft TechNet site (http://technet.microsoft.com/en-us/library/bb124558%28EXCHG.80%29.aspx) and by checking out the Exchange Team blog site (http://msexchangeteam.com/).

High, Higher, or Highest Availability

The level of availability to choose to implement will often come down to the amount of money you are willing to spend. The measurements are usually in percentages of 99.9, 99.99, or 99.999. Do you want 100%? Better yet…do you really require 100%? These are decisions you have to consider, and, as we've explored, the options are varied.

J. Peter Bruzzese is an MCSE (NT,2K,2K3)/MCT, and MCITP: Enterprise Messaging Administrator. His expertise is in messaging through Exchange and Outlook. J.P.B. is the Series Instructor for Exchange 2007 for CBT Nuggets. His latest book is "Tricks of the Vista Masters". He is co-founder of ClipTraining.com, a provider of short, educational screencasts on Exchange, Windows Server, Vista and Office 2007. You can reach Peter at jpb@cliptraining.com.

Copyright Statement