

# Windows Administration

## *in Realtime*

### **2 Letter from the Editor**

*Virtualization, Meet the Small Business*

### **3 Answers from the Experts**

*Do you need digital certificates in your organization?*

### **5 Product Review**

*Shavlik HFNProtect*

### **7 DFS-R**

*Replication for the Rest of Us*

*By: Jaime Halscott - DFS-R is an unsung hero in the battle for true data replication.*

### **11 The Deep Dive**

*Learning Through Listening*

*By: Greg Shields - A tag-up on Realtime Windows Server podcasts.*

### **15 Practical PowerShell**

*Managing WMI Events with Windows PowerShell, Part 1*

*By: Jeffery Hicks - Use PowerShell and WMI to create an event monitoring solution that costs only your time.*

### **24 Exclusively Exchange**

*Working with Transport Rules*

*By: J. Peter Bruzzese - Transport rules afford you a greater level of control and, by extension, security for your organization.*

CONCENTRATED  
TECHNOLOGY

MAXIMUM KNOWLEDGE  
MINIMUM TIME

# MONSTERS of TECH

T E C H N O L O G Y   T R A I N I N G   T O U R

## GREG SHIELDS

author, "Windows Server 2008: What's New | What's Changed"  
*TechNet Magazine* columnist  
Top national conference speaker

## DON JONES

co-author, "Windows PowerShell: TFM"  
*TechNet Magazine* Contributing Editor and columnist  
World-renowned PowerShell speaker & trainer

Quickly learn today's most important and leading-edge technologies for Windows Administrators! Jump-start your existing Windows administration skills to **Windows Server 2008** with Greg Shields' world-famous "What's New, What's Changed" class. Start automating Windows and Active Directory administration today with Don Jones' unique, Practical Method training in **Windows PowerShell**. Go further by ramping your skills up to **Windows Server 2008 R2** and **Windows PowerShell v2**. All of this in just five days of the most unique training you'll ever enjoy!

We laugh at classes where the instructor relies on slide decks and reading from a prepared script. In our class, we don't even use a projector! Instead, you view the slides on your computer, you watch demo videos on your computer, and you work hands-on almost continuously on your computer - meaning you spend more time with your hands *on the technology* and less time craning your neck to see the screen. You get lab guides, slide decks, demo videos, and virtual machines on a Passport USB hard drive - which is *yours to keep*, meaning you'll be able to take a self-paced "refresher course" anytime you want!

We're offering this training to a *maximum of 60 people* in 2009. Be one of the elite - one of the few who is truly prepared for today's latest production-class technologies. Be a MONSTER of TECH!

More info: <http://ConcentratedTech.com/class>

Dates announced for April • More locations & dates coming soon • Early Bird Pricing Now Available

# Letter from the Editor

## *Virtualization, Meet the Small Business*

---

*by Greg Shields*

At a recent IT conference I found myself again presenting a few sessions on virtualization. Although these sessions weren't necessarily new, I found something very unique about their presentation at this conference: The attendance.

Having presented these sessions before, I found myself patently amazed at the number of people showing up to hear this information. The room filled itself to capacity, new chairs were brought in, and eventually the last stragglers found themselves with standing room only. This in comparison with the same presentation only 6 months prior where an equal-sized room was filled but nowhere near full.

I asked those in attendance about their reasoning for being there, "Is virtualization the hot topic in your organization today? More importantly, is it one you don't have enough information about?" The response to both questions was a resounding "Yes!"

In digging a little further, I discovered that virtualization appears to be finally filtering out of the enterprise data center and into the small business. With new products in the market, the early virtualization vendors' high cost of entry is being replaced with affordable solutions that are friendly to the IT jack-of-all-trades. More importantly, network-based storage solutions affordable only by the big guys just a few years ago are similarly coming down in cost. The result is that small businesses who before couldn't fathom a move to virtualization are now putting it on their project lists.

By this point, we all know the players. Each comes with their list of benefits and detractors. Yet some now arrive with an immediate ROI that's compelling to businesses that must minimize the "I" if they're ever to see any "R."

*If you're a small business in this situation, what has been your organization's tipping point?* The price of quality storage such as iSCSI is growing cheap enough that we may soon see penetration into the consumer markets. Hypervisors themselves are growing ubiquitous. Extending the value of backups and enabling cost-effective disaster recovery are all now obtainable goals on the smallest of budgets.

Most importantly, we in IT are settling in with this newfangled idea that the hypervisor is becoming the computer. Virtualization, I'd like to introduce you to the small business. I think you're quickly developing a new friend. ♦

## Answers from the Experts

# Do you need digital certificates in your organization?

by Don Jones

**Q: Do I need digital certificates in my organization? If so, why and how do I get them?**

A: The short answer: Yes, you do. Certificates are the basis for a new wave of security that's been creeping into your enterprise since roughly 2002, and it's high time you started taking advantage of them.

There are several basic classes of certificate: A class I certificate is intended for individuals, is linked to an email address (usually), and is typically

intended for encrypting and signing email. A class II certificate identifies an organization, while a class III certificate identifies servers and software authors—these are the certificates issued for Web sites using SSL, for example, and for software developers and publishers who digitally sign their code. Although class I certificates typically only require you to prove that you “control” an email address (for example, you have access to the inbox), a class III certificate requires independent verification—such as

providing a business license or other legal document.

It's a common misperception that a Web server SSL certificate is intended to provide encryption; in fact, it's possible to have an SSL session without encryption! The point of SSL is more to prove that the Web server is operated by a particular company, so that you know who you're sending your data to; the encryption is great, but it's almost a bonus. If the identity factor weren't there, you could be sending personal information to a

## CONCENTRATED TECHNOLOGY

MAXIMUM KNOWLEDGE • MINIMUM TIME

Join columnists Don Jones and Greg Shields for informative articles on Windows PowerShell and Windows Server, freebies, techno-geek arguments, off-topic amusements, and even some free tools and resources. Get smarter, faster, and smile while you're doing it.

<http://concentratedtech.com>

bogus Web site, in which case it wouldn't matter if it was encrypted.

All certificates are really about proving identity in the otherwise anonymous world of the Intertubes. Encryption—that is, ensuring only the intended recipient can see data—and signing—proving that the data hasn't been tampered with—are practical uses that stem from proof of identity. Signing, for example, not only verifies the identity of the sender or software publisher but also tells you that the data or software is exactly as it was when it was signed; signed code, for that reason, is unlikely to be malicious because the signer's identity could be easily determined.

Where do you get certificates? Commercial certification authorities (CAs) such as Thawte, VeriSign, GoDaddy, and EnTrust will be happy to sell them to you—higher-class ones typically require more strenuous identity checks and cost more per year. Windows Server also includes

a Certificate Services role that can be used to deploy your own CA. A certain amount of planning is required before deploying a CA: You need to ensure that the first CA—referred to as the root—is protected and can be recovered in the event of a disaster; you also need to deploy half of its certificate—its public key—to your client computers so that they trust the CA, any subordinate CAs it authorizes, and any certificates it and its subordinates issue. You also need to plan for certificate revocation, a way of publishing compromised certificates so that clients won't trust them any more.

For more information—especially on Windows' Certificate Services, as well as on how and why certificates work—check out the forthcoming Tips and Tricks Guide to Windows Certificate Services, available in December 2008 at [www.RealtimeNexus.com](http://www.RealtimeNexus.com).

Do you have an IT question you'd like Don to answer? Send it to [answers@realtimepublishers.com](mailto:answers@realtimepublishers.com) for consideration! ♦

*Don Jones is a co-founder of Concentrated Technology ([www.concentratedtech.com](http://www.concentratedtech.com)), helping to deliver IT knowledge in less time using innovative content techniques. He also serves as CTO and Series Editor for Realtime Publishers. Don is the author of more than 30 IT books, including Windows PowerShell: TFM; VBScript, WMI, and ADSI Unleashed; Managing Windows with VBScript and WMI; and many more. He is a multiple-year recipient of Microsoft's "Most Valuable Professional" (MVP) Award with a specialization in Windows PowerShell.*

# Product Review

## Shavlik HFNetProtect

---

*by Eric Schmidt*

Every month, IT departments are being tasked with the installation of security patches and fixes, be it the monthly patches that are released by Microsoft on the second Tuesday of each month or those that are released from other vendors on less predictable schedules. Patching strategies are very different between servers and workstations. With workstations, things are very unpredictable due to the fact that they are often turned off or rebooted; thus, in order to patch effectively, a tool is needed that will interact with the patch server on a regular basis (Windows Software Update Services, Systems Center Configuration Manager). Servers, however, are much more predictable. Generally speaking, they are up all the time and are only rebooted during scheduled downtime or to resolve issues. Often there are Service Level Agreements (SLAs) that need to be maintained or uptime requirements during critical business hours, so patch deployment must be performed in a very controlled manner during very specific maintenance windows. Shavlik's HFNetProtect is a product that enables the deployment of patches in this way.

### **Installation**

The product installation and configuration is very straight forward. One excellent feature of the installation is that it first detects and installs all the prerequisites that the product requires (MDAC, MSXML, and so on).

By default, it uses the Microsoft SQL Desktop Engine, but it can also leverage an existing Microsoft SQL Server. SQL Server is a good option if the number of servers that are going to be patched is more than 25 or so. There were noticeable performance improvements when patching 100+ servers and leveraging a SQL Server system.

### **Scan What/Scan Group**

Configuring patch deployments starts with creating groups of servers. This can be done a number of ways, including by domain, OUs, subnets, and manual list or by linking the group to a text file list. I recommend linking text files because there are occasions in which one or two of the servers in a list will need to be rescanned, and the process goes much quicker by removing the servers from the list that were already patched and don't need to be scanned again.

### **Scan Template**

The next step is to create a scan template. A scan template contains all the settings and products that will be scanned. This is an effective way to control exactly what products will be scanned. For example, you might have several terminal servers with desktop products installed such as Office, Acrobat Reader, and so on. All the other servers in your environment probably don't have desktop products installed and therefore don't need to

be scanned for them. By creating scan groups, in conjunction with server lists, you can target scans based on roles, which will make the process run quicker and more efficiently. Another feature of scan templates is the ability to email the results, which is useful if the scans are scheduled to run non-interactively.

### **Deployment Template**

The next step in the process is to create a deployment template. These templates contain all the settings that will be applied when patches are deployed. This includes how reboots will be handled. With HFNetProtect, reboots can be performed before and after patch installation, only after patch installation, or not at all. In cases in which a server has been up for a long time, it might be beneficial to reboot before the installation of patches and then again when the installation has completed. Deployment templates also offer the ability to email the results, which again, is very helpful if the deployments are scheduled and not run interactively.

### **Performing a Scan**

Scans are initiated by selecting the server list and identifying what scan group to use. Generally speaking, the scans are performed very quickly depending on how many products are being scanned and the number of servers that are in the list. Once the



scan has completed, there is a summary that indicates patch compliance and a detail report that lists each patch. Clicking on each patch will show details about it and the systems that need to have it installed. Deploying patches starts by selecting one or more patches and downloading them. An important point to make here is that the binaries come directly from the manufacturer of the product (that is, Microsoft, Adobe, and so on) and not from Shavlik. This is worth mentioning because if the Shavlik server is behind a proxy server with a white list, you will need to make sure the server has access to all the manufacturers' Web sites. Once they have been downloaded, one or more can be selected and then deployed to the list of servers that were scanned with the ability to select a path group to control the deployment process. Another important item to note is that the application is profile specific. What that means is that the application configuration is specific to each profile, so if there are multiple people that perform patching, it is best to create a service account that will be shared. This way, all the patch and scan groups have to be created only for the one service account profile and all the users will have the same configuration.

## Status

Once deployment has started, a job monitor will display the status of each server so that the deployment can be monitored. I have found this to be rather inconsistent at times and have not relied on the output. To ensure patch deployment, the most effective method is to scan, patch, and scan again. Then redeploy any patches that didn't get installed the first time. The scan, patch, scan again process is repeated until all systems are fully patched.

## Spyware Scanning

A new addition to HFNetProtect is spyware scanning. The configuration and execution of spyware scans involves the same steps as patch scanning. This is a very nice addition to the product, especially if it is being used for servers because administrators may not want to install normal client-based spyware products that are better suited for desktops. Although you might not want a full spyware client running on servers, there are plenty of reasons to have the ability to scan them, and HFNetProtect provides this functionality in a more passive way. Another reason this is a great addition to the product is that it provides a method to validate the effectiveness of client-based spyware products on workstations. By scanning a select list of workstations from time to time, you can compare the results in HFNetProtect with those of the full spyware client.

## Deploying in Isolated Environments

One situation that arises occasionally is the need to scan and patch systems that aren't connected to the Internet—test labs are a good example of this. HFNetProtect offers a means to accomplish this, but setting it up is a little cumbersome. Licensing the product requires entering a key and then activating it. When connected to the Internet, this process is simple and automatic. For environments that are not connected, there is an extra step where a file must be carried out and emailed to Shavlik. Shavlik then sends a file back that has to be carried back in to the isolated environment. Once activation has completed, an Internet-connected Shavlik installation must be maintained, which is used to scan and download the content. The content is then copied to removable media and

carried into the isolated environment. One important note about this process is that all platforms and applications that are going to be patched in the isolated environment must be scanned for content downloaded in the Internet-connected environment. Once the content is moved into the isolated environment, the patching process works very well.

## Summary

Shavlik's HFNetProtect is a great product for patching both large and small numbers of servers during maintenance windows because it provides very tight control over deployment and reboots and the ability to scan for spyware is a welcome addition. I wouldn't recommend it for workstation patching as there are other products that are better equipped to deal with the variability in uptime that workstations have. I've found that some servers have to be rescanned and repatched occasionally, but that is usually limited to 2 out of 50 servers and the benefits that HFNetProtect provides far outweighs the occasional need to rescan and repatch. I have also found that Shavlik's product support is very good. The staff is very knowledgeable and consistently provides timely answers to questions. ♦

*Eric Schmidt works as Enterprise Microsoft Security Technologist, with Honors, for Raytheon Company and has worked in Information Technology for 13 years. Eric has a Masters degree in Computer Information Technology and has developed extensive experience in systems administration, engineering, and architecture specializing in Microsoft Active Directory and Systems Management. Eric has been well recognized throughout his career for his contributions to designing and implementing enterprise-wide solutions using Microsoft Windows-based technologies.*

## Replication for the Rest of Us

---

*by Jaime Halscott*

File replication isn't a new concept. Microsoft has had some form or another built-in to Windows since Windows 2000 when the File Replication Service (FRS) was introduced. Initially, FRS was meant to provide a mechanism for replicating files between the SYSVOL directories on domain controllers. Although it worked well for the most part, it was known to have a few problems, mainly in the areas of reliability and scalability.

The Distributed File System (DFS) actually predates FRS but was meant to address a completely different need. As Windows file servers multiplied, keeping track of exactly where data was stored on a growing number of file servers became a challenge for systems administrators. DFS was meant to address this challenge by creating a logical namespace system where resources on distributed servers could be accessed through one logical path.

The first appearance of DFS in the Windows world was with Windows NT 4.0. Although it was not a part of the core operating system (OS), it could be downloaded freely from Microsoft. It only supported standalone DFS roots and as such wasn't particularly scalable.

Windows Server 2003 saw the merging of FRS and DFS with the introduction of FRS as a replication mechanism for DFS. This solution proved to be problematic for large organizations that needed enterprise-level support and features from a DFS replication scheme. FRS simply wasn't up to the task.

DFS received its first major overhaul in Windows Server 2003 R2 with the addition of Distributed File System Replication. DFS-R ripped out FRS completely and replaced it with an integrated replication solution. This tight integration and revamped replication engine gave DFS-R the functionality it needed to truly be successful in larger environments.

### Setup

Setting up DFS-R isn't nearly as complicated as one might think. The DFS Management MMC snap-in even gives you step-by-step guides to publishing content and data collection, explaining along the way what each option really means and the implications it will have for how you use DFS.

There are two parts to DFS setup in Windows Server 2003 R2. The first is DFS Namespaces. In DFS Namespaces, you create a logical means of storing and locating data. Although this step is not necessary for setting up replication, it does make the process of managing your replicated folders a bit easier. This is also the core functionality that comes to mind when one thinks of DFS. To set up DFS Namespaces, you simply need to take a few actions through a wizard, which guides you quite verbosely through the proper steps.

Initially, you will need to determine what the namespace will be called and where it will be located. A server in your organization that has the DFS component or role installed will host the namespace. This is the first item you must enter. Once you have completed this part, you must determine exactly what the namespace will be called. It could be called anything that you want and will follow the name of the server or domain depending on the type of DFS namespace you choose in the next step.

If you create a namespace called FileShare, it will be called `\\Domain\FileShare` or `\\Server\FileShare` depending on the type of namespace you want to create. During this step, you will see an Edit settings button. This button is used to tell DFS where the namespace share will be. This can be quite confusing if you haven't used DFS before. This is simply a logical share name and will not be hosting any data, but it must exist in order for the DFS namespace to work. This defaults to `C:\DFSRoots\` with the name of the DFS namespace following. In this example, the folder name would be `C:\DFSRoots\FileShare`. Unless you have a pressing reason to change this location, it should be left alone. Again, it doesn't hold any data itself, but serves as a gateway for the browsing of the folders you will be adding into the namespace.

You must now select the permissions for the shared folder. These are share-level permissions, not NTFS. The wizard defaults to the same permissions as any new share created on Windows Server 2003 R2, Everyone with Read access. The wizard calls this "All users have read-only permissions" and it can be changed to one of four permissions templates or you can define custom permissions. It is easiest to set the custom permissions to Everyone with Full control and then rely on NTFS access controls for granularity. This makes figuring out access denied errors a bit easier if you accidentally locked down something too far.



Once you have this step complete, you will have to decide whether this will be a Domain-based namespace or a Stand-alone namespace. Each has advantages and the DFS help onscreen goes a long way toward assisting with this decision. In most cases, electing Domain-based namespace is the best option. All that is left to do is review the settings and create the namespace.

A namespace isn't much use if it doesn't point to any data. Once the namespace wizard is complete, you will need to add folders to the namespace. The great thing about DFS namespaces is that these folders can be anywhere, again assuming that the server is running Windows Server 2003 R2 or Windows Server 2008 with DFS installed.

To add a new shared folder to the DFS namespace, you simply right-click the namespace and select New Folder. Here, you call the logical shared folder whatever you want. You will see a preview of the namespace for the new folder and then you will need to specify at least one folder target. Folder target is a fancy way of referring to the shared folder you want the new folder you have just created to point to. You can browse for an existing shared folder or even create a shared folder on a remote server. As you had in the previous namespace wizard, you will be able to use permissions templates or set custom permissions. The real fun begins if you choose two folder targets. DFS now tightly integrating replication will ask whether you want to create a replication group.

This brings us to the second part of DFS and the main focus of discussion, DFS-R. This is the area in which the greatest changes have come to DFS. Setting up a DFS replication group will enable seamless replication of data between the folder targets in a namespace. As noted earlier, DFS-R can be set up without the need to configure namespaces and folders, but it is not as common or easy to manage.

Creating a DFS-R group requires a few details about how you want to have replication flow. Depending on whether you select Multipurpose replication group or Replication group for data collection, you will be presented with different options in the wizard. In either scenario, you will need to name the replication group, select the members who will participate, choose what data will replicate, when the data will replicate, and how much bandwidth it will be allowed to use.

The most common option is to create a Multipurpose replication group. This choice is best for controlling how the data flows between servers by allowing topology selection and more granular control.

DFS also has several enterprise-level features such as delegation of management permissions. You don't need to over-privilege a user just so that they can manage DFS.

Not all data replicates exactly as one would expect, but DFS-R does a great job of compensating for different uses of

the technology. Each replicated folder will contain a super-hidden folder for DFS-R-related items. The folder is not visible even when you choose the option to show hidden folders and files within Windows Explorer. The folder is accessible under the root of the replicated folder target and is called DFSRPrivate. DFSRPrivate contains the information about the replicated files and the status of each file. Depending on the status of the file, it is placed in one of the following folders:

- ▶ ConflictAndDeleted
- ▶ Deleted
- ▶ Installing
- ▶ Pre-existing
- ▶ Staging

At the root of the DFSRPrivate folder is a hidden file called ConflictAndDeletedManifest.xml. It contains information about the files that have been deleted as a result of the conflict resolution algorithm. DFS-R was designed with many possible scenarios in mind. These folders each attempt to address issues such as changes to two files at once or files that are present when replication first begins.

## Advanced Configuration

One of the most overlooked features in DFS-R is the Staging Quota. The Staging Quota sets aside a predetermined amount of space to be used for staging replicated data. By default, this is 4096MB for all replicated folders. This works great for small amounts of data, but large file shares quickly consume this staging space and DFS-R performance is adversely affected. When the staging data passes 90% of the staging quota, it is above the high watermark and this event will be logged in the event viewer. Some time after this, the older files will be purged to drop this value to 50% of the staging quota. The increased disk I/O and calculations required by this process will hurt DFS-R performance. There is no exact formula for setting the staging quota, but it can be modified for individual folder targets as needed.

## Tips and Tricks

Perhaps the most prevalent configuration for DFS-R is the branch office/central office scenario. In this scenario, a branch office either has a need for local copies of files that exist on a central office server or the central office has a need for copies of files that exist on a branch office server. Sometimes in this scenario, both offices will have a need for files from the other office. There can also exist a situation in which there are several branch offices with the same or similar needs.

We typically find that these branch offices are connected by VPN tunnels to the central office and possibly each other. However, there isn't always enough bandwidth to make the initial replication of data happen in a realistic time period. Most companies can't afford to wait several weeks or months for replication to occur. Although DFS-R supports Remote Differential Compression (RDC) to compress data and send only the changes within files, the initial data still needs to get from the source server to the destination server the first time.

To get around the issue of time needed to initially replicate, the concept of DFS Seeding or DFS Pre-Staging came about. For the purposes of this article, we will use the term DFS Seeding. DFS Seeding can occur in a few ways, but care must be taken to ensure that it doesn't adversely affect performance.

If the servers being used at the branch offices are newly provisioned or dedicated to this purpose or otherwise able to be taken offsite, they may be physically moved to the central office so that the initial replication can happen. Replicating the data over gigabit Ethernet can happen at up to 4000 times the speed of a 256Kb broadband upload. Even T1 speeds can seem terribly slow compared with a fast local area network (LAN) connection. Scenarios involving 100 gigabytes or more can greatly benefit from this type of physically close DFS Seeding.

Some scenarios do not or will not lend themselves to physically moving servers. When the first option is not available, the next best option is to use removable storage such as an external hard drive. Typically, using Robocopy from the Windows Server 2003 Resource Kit Tools will prove very effective. However, using Robocopy for this purpose requires a few counter-intuitive options. DFS-R will treat a file as changed unless it is exactly the same on the source and target. As such, one would think that you would have to use the /COPYALL switch in Robocopy to get the files, with all permissions and attributes intact, to copy over to the target disk. This is not the case. DFS-R will see the files as changed if the /COPYALL switch is used and try to replicate them again.

Microsoft simply states that DFS Seeding can occur by copying the files, restoring from backup, or copying from tape, DVD, or removable hard disk. Using a removable hard drive is usually the easiest and most cost-effective mechanism.

## Troubleshooting

As reliable as DFS-R is, no system is perfect. Fortunately, Microsoft has provided a considerable amount of diagnostic functionality with DFS-R. DFS-R actually gets an event log just for itself. Here you can get detailed information about everything from staging quota usage to when initial replication has completed, to problems with contacting domain controllers for replication topology information.

DFS-R also features diagnostic reporting. For any replication group, you simply highlight the group in which you are interested and select Create Diagnostic Report from Actions. This will create an HTML report with the options for which members to include and whether backlogged files are to be part of the report. The reports are stored by default in C:\DFSReports and are tagged with the replication group name and date/time stamp for later reference. The reports scour the event log entries for relevant issues, suggest fixes, and even show RDC bandwidth savings that have been calculated.

For advanced troubleshooting, there is DFSRDIAG.exe. This command-line tool is not only useful for diagnostics but also triggering and stopping replication and viewing low-level configuration such as Active Directory (AD) settings for a particular member.

## Changes to DFS-R in Windows Server 2008

Windows Server 2008 brings with it a large number of evolutionary changes to DFS-R. The first of these involves using DFS-R to replace the legacy FRS for AD-related functions. In particular, DFS-R replaces FRS for SYSVOL replication. Using DFS-R instead of FRS for SYSVOL replication requires that the domain be at Windows Server 2008 Domain Functional Level. In order to migrate the SYSVOL from FRS to DFS-R, there is a new command-line tool called DFSRMIG.exe shipped with Windows Server 2008.

The Windows Remote Procedure Call (RPC) subsystem was overhauled in Windows Server 2008. DFS-R benefits from this by implementing RPC Asynchronous Pipes. RPC Asynchronous Pipes give three primary benefits:

- ▶ Multiple outstanding calls from a replication partner
- ▶ Slow or delayed partners
- ▶ Replication of large amounts of data

Windows Server 2008 boosts DFS-R performance through other OS-related tuning. In particular, Windows Server 2008 provides Unbuffered I/O that increases throughput by decreasing the number of data copy operations that would normally occur during replication.

Windows Server 2008 also introduces Asynchronous Low Priority I/Os, which take the familiar action of running a process in a lower priority and apply this concept specifically to hard disk access. As long as an application or service is low-priority I/O aware, as DFS-R is in Windows Server 2008, it can function without negatively impacting overall server responsiveness.

The number of concurrent files that can be replicated has been increased four-fold in Windows Server 2008 DFS-R from 4 to 16. This is due to the Asynchronous Low Priority I/O changes. This combined with enhancements to the replication algorithm ensure that replication from multiple branch sites occurs optimally even over slow WAN links.

DFS-R benefits from an improved ability to recover from “Dirty Shutdowns” (that is, power loss, dismounted volume, or other I/O-related problem). These types of events would cause inconsistencies between the NTFS Update Sequence Number (USN) journal entries and the values for files in the DFS-R database. A “Dirty Shutdown” in Windows Server 2003 R2 would force a complete rebuild of the DFS-R database, which is quite resource and time intensive. A new validation algorithm has been built-in to Windows Server 2008 to keep database rebuilds to a minimum.

Finally, scalability has been addressed by replacing the performance counters library used by DFS-R. Windows Server 2003 R2 was limited by a formula for the number of objects that could reliably be replicated. The formula, the result of which should be kept to 1024 or fewer on each server, was published in the Microsoft TechNet blogs and is as follows:

*(number of replicated folders in replication group<sub>x</sub> \* number of simultaneously replicating connections in replication group<sub>x</sub>) + (number of replicated folders in replication group<sub>y</sub> \* number of simultaneously replicating connections in replication group<sub>y</sub>) + (number of replicated folders in replication group<sub>n</sub> \* number of simultaneously replicating connections in replication group<sub>n</sub>)*

Windows Server 2008 has no effective limit to the number of performance counter objects that can be created and as such does not need to follow this formula.

Of note is the fact that Windows Server 2003 R2 and Windows Server 2008 can successfully coexist in the same DFS-R environment, but all the advanced features of DFS-R incorporated into Windows Server 2008 require all servers to be running Windows Server 2008.

## Summary

DFS-R is an unsung hero in the battle for true data replication. Microsoft introduced a mature system for data access, availability, and replication on the first try. The scalability and reliability of DFS-R make it suitable for organizations of all sizes. For uses ranging from branch office file sharing to warm-spare data storage for high availability without the complexity or costs of clustering, the enhancements to DFS alone make the upgrade to Windows Server 2003 R2 or Windows Server 2008 worth it. ♦

*Jaime T. Halscott is a founder of UnbreakableIT, a global leader in managed services, managed security, and outsourced IT solutions. He has served in Director positions at several software companies including ScriptLogic. Jaime is a graduate of the University of Central Florida where he studied Computer Science. In addition to his undergraduate work, Jaime holds an extensive list of industry certifications from Microsoft, Cisco and CompTIA. In addition, Jaime is a feature contributor to the Windows Administration in Realtime eJournal from Realtime Publishers.*

# The Deep Dive

## Learning Through Listening

---

*by Greg Shields*

In our line of work, we spend far too much time sitting in front of computer screens. We're used to reading to glean information out of magazines, knowledgebase articles, and the occasional IT blog. But all that reading is only one medium that's available to us to learn about the goings-on in the world of IT.

Another that's gaining momentum is a throwback to the old school radio interview, or at least our modern-day representation of that classic art form. Over at the [Realtime Windows Server Community](#), you might have missed some of the great podcasts queued up for listening if you simply click the Podcast link. For this month's feature, let's take a look through some of the very best podcasts you can listen to right now by navigating over to the community. Here are but a few.

### [Getting the Network Administrator Home at Night with PacketTrap](#)

The job of the network administrator appears never-ending. Constantly searching through monitoring data to isolate and solve the problem of the day means long hours and few nights home with family. The problem often isn't that there is not enough data to identify network problems and performance issues, but usually that there's too much.

This week, I sat down with Matt Bolton of PacketTrap, a brand new network monitoring solutions provider in the market. This company sells a network monitoring product called Perspective that is designed to assist the network administrator with just these sorts of problems. A product not unlike Orion from Solarwinds or Openview from HP, Perspective gives a new...perspective...on the network by dialing down the amount of data shown to the administrator. Matt tells me that his main priority is in "getting network administrators back home for dinner." If you're an overworked network administrator, you'll want to pull up a chair and check out this great podcast with PacketTrap where we talk about what you need to do just that.

### [Hosted Exchange: A Compelling Idea for SMBs](#)

If you're a small business (or even a "medium" one), you might suffer from the "accidental IT person" syndrome. That is, the person with a knack for computers somehow accidentally becomes the IT person for the business. That's perfectly fine when it comes to keeping people's desktops running or setting up the file server. But email is growing to become a mission-critical need for all businesses. When email goes down, the business goes down.

Adding to the problem is that the skills required to administer email correctly are hard to find and require a lot of experience to develop. Unlike a lot of Microsoft servers, you can't just drop an Exchange Server in your network and expect it to work perfectly for long—unless you build its environment correctly.

Alternatively, you can outsource this critical service to an outside vendor with layers upon layers of redundancy. In this podcast, I talk with Apptix, one of those very vendors that prides themselves on being a perfect fit for the SMB. In fact, they're the largest email outsourcer in the world. Learn from Apptix Director of Engineering James Bond about what's cool and exciting in the world of Exchange outsourcing, and what you need to know to do outsourcing correctly.

Oh, and yes, that is his name, and we even talk about a few funny stories about his time at the Pentagon, seriously. James is also an entertaining speaker, which makes this interview a whole lot of fun. It's worth a few minutes of your time.

### **Virtualization, Backups, Disaster Recovery, and Vizioncore**

Adding virtualization to an environment solves a lot of problems: consolidation, power use, rapid deployment, and server management, to name a few. But the move to virtualization also makes some otherwise trivial issues fairly problematic if you don't plan for them properly. Backups and disaster recovery are two of these situations.

With virtualization, both backups and disaster recovery have the potential to grow more complicated simply because you have more options in which to best deploy them. With all the options available, it's easy to go down the wrong path and implement an architecture that won't ultimately bring you the restore-ability you need.

Vizioncore recognizes that both of these topics are challenging when considered within the frame of a virtual environment. So they attempt to solve some of these complexities through a suite of software add-on products that make easy the process of backups, restore, and disaster recovery (as well as some other needs). In this highly technical discussion with Jason Mattox, CTO of Vizioncore, we talk about the hard concepts associated with virtual backups and the best ways to resolve some of those problems.

If you're a user or a potential user of VMware's products for virtualization, you definitely want to take a listen. Jason's advice could save you and your environment plenty of time and headache.

### **The State of Computer-Based Training**

Instructor-led training can be expensive, and it's tough to learn through reading multiple-thousand page books. Computer-based training with all its multimedia capabilities lies somewhere in the middle—reasonably priced, entertaining enough to keep your attention, and a visual source of learning while doing.

In this podcast, I interview Roy Furr with CBT Nuggets, makers of all kinds of computer-based training for IT technologies as well as preparation for IT certifications. In this podcast, we talk about the state of CBT today and how we think the job market is dealing with both training and certification needs. We'll talk about employers and their desire for certified employees, and how businesses can get exceptional training at exceptionally inexpensive prices through the right training channels.

### **A Smarter User Account Control with BeyondTrust Privilege Manager**

User Account Control is for many the bane of Windows Vista. Its encompassing ideas of reducing the spread and impact of administrator rights are great in theory but difficult in practice. At the same time, getting rid of administrator rights in our networks is a key need that we all wish we could solve.

In this podcast, I interview John Moyer, CEO of BeyondTrust. We talk about the regulatory, compliance, configuration control, and malware problems associated with the widespread distribution of administrator rights to those people who shouldn't have them. We discuss existing tools that are improvements on UAC and that eliminate this need while still providing laser-focused elevation for needed applications.

I'll admit that I wasn't aware of BeyondTrust's Privilege Manager product, but after talking with John, I'm impressed with how this product can secure a network, eliminate unnecessary admin rights, and do so in a manner that's easy to use and fully supported by Microsoft. This one's worth a listen.



## Another exclTing Interview: Travis Morrison of New Belgium Brewery

In this exclTing interview with Travis Morrison of the New Belgium Brewery, we talk about what its like to be a Sr. Systems Administrator at a place that sells some of the best small-batch craft beers in the nation. We talk about New Belgium's recent migration from a Novell-based directory to Microsoft Active Directory, their success with Microsoft's Unified Communications platform, and just how much fun it can be to work in the IT Department of a major brewery. Now I'm off to kick back a few...

## Automated Documentation and Compliance - An Interview with Bryan Cote of Ecora Software

Understanding your network is one thing, but proving that you do is fully another. With compliance regulations requiring a database of records showing what on your network changed and when, it is growing more and more impossible every day to fulfill the needs of documenting and auditing your environment.

That's why tools such as Ecora Auditor Pro exist. Inventorying and storing thousands of elements about every machine on your network, this toolset means you can pass compliance audits with ease. All the while, you get a better-running network because you have more information about that network.

In this podcast, I interview Brian Cote, senior product manager for Ecora Software. Here we talk about the challenges of network documentation and compliance and how Ecora Auditor Pro resolves some of the critical missing auditing pieces in our native operating systems (OSs).



WHAT'S NEW  
WHAT'S CHANGED



Microsoft has released its next server operating system – Windows Server 2008 – and you need to know more about it. But you don't need the basics. You already know Windows 2003. You just need to know what's new and what's changed in Windows Server 2008. Read-Only Domain Controllers, the Group Policy Central Store, Terminal Server RemoteApps, Fine-Grained Password Policies. This quick and entertaining guide, written by Windows insider Greg Shields does just that. Focusing on the new technologies for installing, managing, and securing Windows Server 2008, you'll quickly ramp up your skills. Save yourself some time and money by skipping the basics and using your existing skills to master Microsoft's new server O/S.

Automate server installations \* More effectively manage servers through Server Manager \* Gain insight with Reliability and Performance Monitor \* Implement powerful new Group Policy \* Reduce your attack surface with Server Core \* Complete better Active Directory backups \* Deploy apps using Terminal Services \* Secure your servers with the new Windows Firewall

TABLE OF CONTENTS	
<b>Chapter 1:</b>	Introduction to Windows Server 2008
<b>Chapter 2:</b>	Installing Windows 2008
<b>Chapter 3:</b>	Server Management
<b>Chapter 4:</b>	Group Policy
<b>Chapter 5:</b>	Server Core
<b>Chapter 6:</b>	Windows Server Virtualization
<b>Chapter 7:</b>	Active Directory
<b>Chapter 8:</b>	Terminal Services
<b>Chapter 9:</b>	Security & the Windows Firewall with Advanced Security
<b>Chapter 10:</b>	IIS 7.0
<b>Chapter 11:</b>	Other New & Compelling Features

[http://www.sapienpress.com/Windows\\_Server\\_08.asp](http://www.sapienpress.com/Windows_Server_08.asp)

Greg Shields



## There Are Plenty More

This is only a sampling of the very best—and most interesting—the community has to offer. If any of these tickle your fancy or your iPod, drop on by <http://www.realtime-windowsserver.com/podcast> for a listen.

And, don't forget that as the host of the Realtime Windows Server Podcast Series, I'm always out to find IT professionals who work in interesting and exciting industries. If you've got an IT job in an interesting business, drop me a line at [gshields@realtimepublishers.net](mailto:gshields@realtimepublishers.net). We'll have a little interview and give you the platform to tell others about what's neat and exciting in the sometime droll world of IT. ♦

*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](http://www.realtime-windowsserver.com).*

# Practical PowerShell

## Managing WMI Events with Windows PowerShell, Part I

by Jeffery Hicks

You can download a zip file with all these scripts from [http://www.realtime-windowsserver.com/code/v1n11/Practical\\_PowerShell.zip](http://www.realtime-windowsserver.com/code/v1n11/Practical_PowerShell.zip).

You've probably seen or even used monitoring solutions to manage your servers and networks. These tools watch for certain things to happen, such as a service stopping, and then do something, such as notify you. Most of these tools are likely using Windows Management Instrumentation (WMI). With a little effort on your part, along with PowerShell, you can create your own event monitoring solution and it won't cost you a dime, just your time.

There are a few approaches you can take with WMI and Windows PowerShell v1.0. Let me show you how to use a notification query to alert you when a Windows event fires. Where a traditional WMI query would attempt to retrieve objects from a Win32 class, a *notification query* typically returns objects that match one of several system classes. You'll most likely want to watch for when things are created, modified, or deleted. You'll use these corresponding classes: \_\_InstanceCreationEvent, \_\_InstanceModification, and \_\_InstanceDeletion. Here's how our query might start:

```
Select * from __InstanceCreationEvent
```

The system class is essentially generic and would return a result whenever anything was created. What we need to do is narrow the query by specifying what class of created object we are looking for. This is referred to as the TargetInstance.

Let's suppose you want to watch for when a new file is created in a particular folder. The object class is CIM\_DATAFILE. We'll use the ISA operator in our query like this:

```
Select * from __InstanceCreationEvent WHERE TargetInstance ISA 'CIM_DATAFILE'
```

Remember, this is a notification query, which means WMI must constantly check to see whether any objects match the query. From a practical as well as performance view, this isn't very desirable. Instead, we'll use WMI polling and instruct WMI to only check x number of seconds using the WITHIN operator:

```
Select * from __InstanceCreationEvent WITHIN 30 WHERE TargetInstance ISA 'CIM_DATAFILE'
```

This query will check every 30 seconds for matching objects. There's no right or wrong polling interval. It depends on what you are monitoring and how critical it is that you are notified as soon as the event fires. What I do suggest is that you make your query as specific as possible, especially when using a class like CIM\_DATAFILE:

```
Select * from __InstanceCreationEvent WITHIN 30 WHERE TargetInstance ISA 'CIM_DATAFILE'
AND TargetInstance.drive='C:' AND TargetInstance.Path='\\files'
```

Now the query will check every 30 seconds to see whether a new file has been created in C:\Files. Remember, in WMI you need to escape the \ character, thus \files becomes \\files. But how do we use it?

The **Get-WMIObject** cmdlet in v1.0 isn't designed for notification queries, so we'll use the underlying .NET management classes. The query we just created will be defined as a System.Management.WQLEventQuery object. To execute the query, we'll need a special management class, the System.Management.ManagementEventWatcher. We'll bind everything to this object and wait:

```
$watcher = New-Object System.Management.ManagementEventWatcher $scope,$EventQuery
$watcher.start()
$watcher.WaitForNextEvent()
```

However, there is one more critical step. What if the script ends before a matching event fires? Well, you'll never know. Your script needs to stay "alive" waiting for an event to fire. A common approach is to use a simple loop.

```
While ($true) {
    $watcher = New-Object System.Management.ManagementEventWatcher $scope,$EventQuery
    $watcher.start()
    $evt=$watcher.WaitForNextEvent()
}
```

You would need to use CTRL-C to break the loop and exit your script. In a moment, I'll give you a more elegant approach.

When the event fires, you'll have a new TargetInstance object of the type you were monitoring, like CIM\_DATAFILE in the query we've been using.

```
$evt.TargetInstance | select @{Name="Server";Expression={$_.CSName}},`
Name,FileType,FileSize,@{Name="Created";Expression={
[System.Management.ManagementDateTimeConverter]::ToDateTime($_.CreationDate)}}
```

The result is that when a new file is created, the server name, filename, file type, file size, and its creation date are captured by PowerShell. Let me give you a complete script that you can use to monitor for new file creation called Get-NewFileEvent.ps1.

```
#Get-NewFileEvent.ps1

# $path is the full path to the folder you
# want to monitor like C:\files\PowerShell. It is relative
# to the computer you are monitoring
```

```

# $poll is the polling interval in seconds.
$ESCkey = 27

$namespace="\\"$computername\root\cimv2"
$drive=Split-Path $path -Qualifier
$Folder=(Split-Path $path -NoQualifier).Replace("\","\\")+"\\"

$query="Select * from __InstanceCreationEvent Within `
$poll where TargetInstance ISA 'CIM_datafile' AND `
TargetInstance.drive='$drive' AND TargetInstance.Path='$folder'"

$EventQuery = New-Object System.Management.WQLEventQuery $query
$scope       = New-Object System.Management.ManagementScope $namespace

if ($Credential) {
    #use alternate credentials if passed
    $scope.options.Username = $credential.GetNetworkCredential().Username
    $scope.options.Password = $credential.GetNetworkCredential().Password
    $scope.options
}

$watcher      = New-Object System.Management.ManagementEventWatcher $scope,$EventQuery
$options       = New-Object System.Management.EventWatcherOptions
$options.Timeout = [timespan]"0.0:0:1"
$watcher.Options = $options

cls
Write-Host "Waiting for:" $EventQuery.querystring " on $computername. Press ESC to quit."
-back cyan -fore black
#start waiting for events
$watcher.Start()

#keep looping and waiting
while ($true) {

#trap any errors and keep going
trap [System.Management.ManagementException] {continue}

$evt=$watcher.WaitForNextEvent()
#if an event has fired get the target instance and select a subset of properties
if ($evt) {
    $evt.TargetInstance | select @{Name="Server";Expression={$_.CSName}},`

```

```

Name,FileType,FileSize,@{Name="Created";Expression={
[System.Management.ManagementDateTimeConverter]::ToDateTime($_.CreationDate)}}
#clear the evt object and wait for the next event
Clear-Variable evt
}

#watch for ESC key
if ($host.ui.RawUi.KeyAvailable)
{
    $key = $host.ui.RawUI.ReadKey("NoEcho,IncludeKeyUp")
    if ($key.VirtualKeyCode -eq $ESCkey)
    {
        $watcher.Stop()
        break
    }
}

} #end while Loop

```

The script takes several parameters, such as `–computername` and `–path`. Remember that the path is relative to the remote computer, so `C:\Files` would be the folder on the remote computer:

```
PS C:\Scripts\> .\Get-NewFileEvent -path C:\files -computername SERVER02
```

You should recognize parts of the script's beginning. To keep the script alive, it keeps looping until the ESC key is pressed. This is accomplished by watching for a key press using the `$host.ui.RawUI` object:

```

if ($host.ui.RawUi.KeyAvailable)
{
    $key = $host.ui.RawUI.ReadKey("NoEcho,IncludeKeyUp")
    if ($key.VirtualKeyCode -eq $ESCkey)
    {
        $watcher.Stop()
        break
    }
}

```

If the Escape key is detected, the event watcher is stopped and the script exits. Until that time, the script keeps looping. When a new file is detected, the target instance is populated and the script returns a custom object with a subset of file properties:

```

Server    : CHAOS
Name      : c:\files\file0232.dat
FileType  : dat File
FileSize  : 3267
Created   : 10/14/2008 1:29:07 PM

```

You, of course, could insert any PowerShell code you'd like, such as an email notification, copy or moving the file, or anything else you can think of. The script will continue waiting and display information as other files are created.

Although the script can be used in a production environment, it is really intended as an example of how to monitor an \_\_InstanceCreation event. But you might need to watch for other events, such as when a service object is modified (for example, a service has stopped) or a process is terminated, which would fire an \_\_InstanceDeletion event.

I created a second script, Get-WMIInstanceEvent.ps1, which you can use either as is or as a prototyping tool for building your own script.

```
#Get-WMIInstanceEvent.ps1

# $event must be either Modification, Creation, or Deletion
# $class is the name of a WMI class like Win32_Service. This is the TargetInstanceClass
# $poll is the polling interval in seconds. Default is 10
# $filter is some filtering criteria like Name='Spooler'
# $credential is a saved PScredential

#example
# .\get-WMIInstanceEvent.ps1 -class win32_service -event modification -filter "targetinstance.name='spooler'"
# .\get-WMIInstanceEvent.ps1 -class win32_NTLogevent -event Creation -filter "targetinstance.logfile='system'"
```

# World's hottest IT topics

Windows PowerShell™: TFM® 2nd Edition  
Windows PowerShell™: TFM® 3rd Edition  
(covers Windows PowerShell v2.0)  
ADSI Scripting: TFM®  
WSH and VBScript Core: TFM®  
PrimalScript 2007: TFM®  
Windows Server 2008: What's New/What's Changed  
Exchange Management Shell TFM®  
Managing Active Directory Windows PowerShell TFM®



For more information:  
[www.sapienpress.com](http://www.sapienpress.com)





```

Param([string]$computername=$env:computername,
      [string]$event,
      [string]$class=$(Throw "You must specify a class like win32_service"),
      [int32]$poll=10,
      [string]$filter,
      [System.Management.Automation.PSCredential]$credential)

#define the Escape key value
$ESCkey = 27

#validate the event type and return a warning message
#if an invalid option is used
Switch ($event) {
    "modification" {$Instance="__InstanceModificationEvent" }
    "creation" {$Instance="__InstanceCreationEvent"}
    "deletion" {$Instance="__InstanceDeletionEvent"}
    Default {
        $msg="You must enter a valid event: 'modification','creation' or 'deletion'."
        Write-Warning $msg
        Return
    }
}

}

#define the basic query
$query="Select * from $Instance within $poll where TargetInstance ISA '$class'"

if ($filter) {
    #if a filter has been specified, append it to the basic query
    $query=$query + " AND " + $filter
}

$namespace="\\"$computername\root\cimv2"

$EventQuery = New-Object System.Management.WQLEventQuery $query
$scope      = New-Object System.Management.ManagementScope $namespace

```

```

if ($Credential) {
    #use alternate credentials if passed
    $scope.options.Username = $credential.GetNetworkCredential().Username
    $scope.options.Password = $credential.GetNetworkCredential().Password
    $scope.options
}

$watcher      = New-Object System.Management.ManagementEventWatcher $scope,$EventQuery
$options      = New-Object System.Management.EventWatcherOptions
$options.Timeout = [timespan]"0.0:0:1"
$watcher.Options = $options

cls

Write-Host "Waiting for:" $EventQuery.querystring "on $computername. Press ESC to quit."
-back cyan -fore black

$watcher.Start()

while ($true) {
    trap [System.Management.ManagementException] {continue}

    $evt=$watcher.WaitForNextEvent()

    if ($evt) {
        $evt.TargetInstance
        if ($event -match "modification") {

            #the __InstanceModification class will also return an PreviousInstance
            #object that shows the object's state before the modification
            write "PREVIOUS INSTANCE"
            $evt.previousInstance
        }

        Clear-Variable evt
    }

    #watch for ESC key

```

```

if ($host.ui.RawUi.KeyAvailable)
{
    $key = $host.ui.RawUI.ReadKey("NoEcho,IncludeKeyUp")
    if ($key.VirtualKeyCode -eq $ESCkey)
    {
        $watcher.Stop()
        break
    }
}

} #end while Loop

```

The script takes several parameters and executes a WMI notification query. The script defaults to the local computer, but you can specify a remote computer with the `–computername` parameter. If you want to use alternate credentials, you can pass a stored PSCredential with `–credential`. You use the `–event` parameter to define the instance event you want to monitor. Acceptable choices are Creation, Modification, and Deletion. You must specify a WMI class for the target instance with the `–class` parameter. This might be something like `CIM_DATAFILE` or `Win32_Service`. The script will poll every 10 seconds by default, but you can modify that with the `–poll` parameter. Finally, the `–filter` parameter can be used to fine-tune your query. You’ll need to pass a string that is in the proper WMI format, for example, “`targetinstance.name='spooler'`”.

The script builds a notification query based on the parameters you specified and executes it. I use the same looping technique to keep the script alive. Pressing the ESC key will end the script. If a matching event fires, the complete instance object is written to the pipeline. If you are running an `__InstanceModification` query, the script will also show the previous instance object.

```

if ($evt) {
    $evt.TargetInstance
    if ($event -match "modification") {
        #the __InstanceModification class will also return an PreviousInstance
        #object that shows the object's state before the modification
        write "PREVIOUS INSTANCE"
        $evt.previousInstance
    }
}

```

Here are some examples of how you might use this script.

```

PS C:\> .\get-WMIInstanceEvent.ps1 -computer Print02 -class win32_service -event
modification -filter " targetinstance.name='spooler'"

```

This example will monitor the Spooler service on PRINT02. If the service is modified in any way, for example if it stops, an event fires and writes the new server instance object and its previous instance to the pipeline.

```

PS C:\> .\get-WMIInstanceEvent.ps1 -computer Server01 -class win32_NTLogevent -event
Creation -filter "targetinstance.logfile='system'"

```

Or you can use notification queries to watch for new event log entries. This example will notify you when a new event is written to the System event log on SERVER01.

Although you can use the script as is, I expect you will use it as a tool for developing other scripts. By monitoring specific types of events and specific classes, you can discover what information is available from the different instance objects. Once you know what properties are available, you can write additional code to do something with it, such as sending an email, writing an event to a log file, restarting a service, or terminating a process.

Next month, I'll be back with a different approach for managing WMI events with PowerShell. In the meantime, if you have any questions about these scripts or PowerShell in general, I hope you'll use the forums at [ScriptingAnswers.com](http://ScriptingAnswers.com). ♦

*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](mailto:jhicks@sapien.com).*

# Exclusively Exchange

## Working with Transport Rules

---

by J. Peter Bruzzese

One of the designed functions of the Hub Transport role is that it is responsible for all mail flowing in and out of your organization to go through one of your organization's Hub Transport servers. What this means is that you can actually apply a set of rules to your incoming and outgoing messages, while that mail is "IN TRANSIT". In addition, you can create rules that are located on your Edge Transport servers so that mail coming and going in the perimeter might be affected prior to that mail coming into your internal network or going out on the Internet.

Currently, a user can employ Outlook (or another mail application) and establish Rules that are applied to messages when they arrive in the user's Inbox. But, with transport rules, you, the administrator, have the first line of defense and control.

### *The Structure of a Transport Rule*

Rules created on the Hub Transport server have a Transport Rules Agent that ensures rules are applied. On the Edge Transport server, it is called an Edge Rules Agent. The types of rules you can create on the Hub Transport or Edge Transport roles varies slightly due to the nature of each server role; however, the most important thing to remember is that rules created at the Hub Transport level are applied across all Hub Transport servers. It's an organization-wide rule. That is not the way it works for the Edge Transport role. Rules are created per server; thus, if you have more than one Edge Transport server, you are going to have to make sure that rules are set on all of them, if that is your goal.

There are three parts to the rule: conditions, actions, and exceptions. The condition determines which messages need to be affected by the rule. The action determines what will happen to those messages. And the exception allows for circumstances when a message might not have the rule applied, even if the condition were met.

To create a Hub Transport rule, open the Exchange Management Console (EMC) and expand out the Organization Configuration work center, click Hub Transport, and then select the Transport Rules tab. Next, select the New Transport Rule link from the actions pane:

Note: On the Edge Transport server, you would expand the Edge Transport work center and select the Transport Rules tab, then the New Transport Rule link from the actions pane.

- ▶ On the Introduction screen, you are asked for a name for the rule and a comment where you can explain the purpose of the rule you are creating. You can also select/deselect a check box for the rule to be automatically enabled or disabled after it is created. Enter the information, and click Next.
- ▶ The Conditions screen, see Figure 1, enables you to choose one or more conditions you want met before action is taken on a message. In Step 1, you select the condition, and in Step 2, you alter the condition by selecting an underlined value and indicating your personal need. If you want to have emails from a group moved into a certain folder, you can select the option in the first part and indicate the exact group in the second. After you have your condition(s) established, click Next.
- ▶ The Actions screen enables you to select an action or actions to perform in the event the condition is met. After the action(s) is determined, click Next.
- ▶ The Exceptions screen is not mandatory. Here you can select a variety of exceptions to this transport rule if you need them. Configure your exceptions, and choose Next.
- ▶ After the three aspects of the rule are in place, the Create Rule screen allows you to review the Configuration Summary for the rule. If something is amiss, click Back. If everything is fine, click New.
- ▶ After the rule has been created, click Finish.

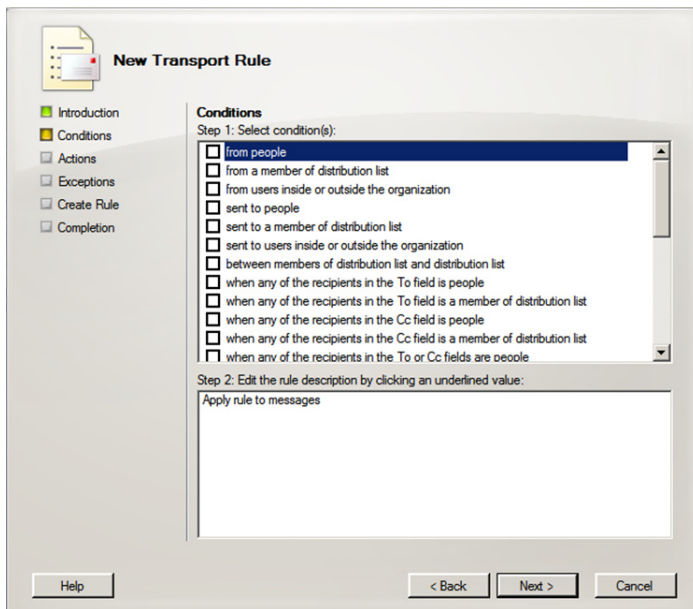


Figure 1: The Conditions tab of the Transport Rule wizard.

## Common Transport Rule Ideas

You might wonder, on an enterprise level, what kinds of rules you might need to implement that would affect a wide range of users. Here are a few possibilities:

- ▶ You have a legal disclaimer that must be attached to all mail leaving your company. At the same time, you do not need that disclaimer attached to mail that is simply going from one person in the company to another only to outsiders.
- ▶ You may have a legal requirement that persons in one department are not permitted to communicate openly to members of another department. You can control that through a transport rule.
- ▶ You may have a virus with specific file type extension that your virus software hasn't been updated to remove just yet. You can create a transport rule to filter email with that known file type.



## VISTA / OFFICE 2007 ROLLOUT

"The key to a smooth **Vista / Office 2007 ROLLOUT** is **ClipTraining**."

- Chris Nichols - Director of IT, Tax Education Support of Iowa

When you give your team the latest software; give them the latest training. ClipTraining supports your team and creates a confidence unattainable with traditional classroom and video training.

LEARN WHAT YOU NEED...  
**WHEN YOU  
NEED IT.**



**www.ClipTraining.com**

Email: **info@ClipTraining.com**

Phone: **1-888-611-CLIP (2547)**



As you can see, you can use transport rules to filter confidential information, to prevent known virus extensions from running rampant through your organization, to track or archive messages that are sent to or received from specific individuals, to redirect inbound and outbound messages for inspection before delivery, and much more. There are hundreds of possibilities.

### ***How to Rule with Transport Rules***

Obviously you don't want to overdo it with too many rules. Every new rule causes your Hub Transport servers to work that much harder scrutinizing email for the sake of the conditions, actions and exceptions. But you should know what your ability is thanks to the Hub Transport features. You are afforded a greater level of control and, by extension, security for your organization. Use this power for good and not for evil.♦

*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. In harmony with the joy of writing Exclusively Exchange for Realtime Publishers, he has created a free Exchange training site at [www.exclusivelyexchange.com](http://www.exclusivelyexchange.com). He is co-founder of ClipTraining.com, a provider of short, educational screencasts on Exchange, Windows Server, Vista, Office 2007 and more. You can reach Peter at [jpb@cliptraining.com](mailto:jpb@cliptraining.com).*

## Copyright Statement

© 2008 Realtime Publishers, all rights reserved. This eJournal contains materials that have been created, developed, or commissioned by, and published with the permission of, Realtime Publishers (the “Materials”) and this work and any such Materials are protected by international copyright and trademark laws.

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

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

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

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

If you have any questions about these terms, or if you would like information about licensing materials from Realtime Publishers, please contact us via e-mail at [info@realtimepublishers.com](mailto:info@realtimepublishers.com). ♦