

# Windows Administration

## in Realtime

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

*I'm Noticing the Excitement Surrounding Server 2008 R2*

### 4 **Break the Habit**

*Haphazard Access Management*

### 6 **Product Review**

*UltraEdit*

### 8 **SQL Server High Availability**

*Is Your Corporate Data Always Available?*

By: Richard Siddaway - Explore techniques for providing SQL Server high availability to ensure your vital corporate data is continuously available.

### 18 **The Deep Dive**

*A Business Perspective on SharePoint*

By: Robert Emmert - How to get your SharePoint project approved.

### 20 **Practical PowerShell**

*Excel-lent PowerShell*

By: Jeffery Hicks - Discover a unique way to integrate PowerShell and Excel.

### 25 **Exclusively Exchange**

*Exchange Server 2007 SP2 Coming 3Q*

By: J. Peter Bruzzese - What will this new service pack bring our way?

# Letter from the Editor

## *I'm Noticing the Excitement Surrounding Server 2008 R2*

---

*by Greg Shields*

Think back a few years: When Windows Server 2003 R2 released, did you actually upgrade? My suspicion for most is “probably not.” Microsoft’s first attempt at an R2 release wasn’t filled with all that many new goodies that got us salivating for its update bits.

Yes, if you needed its built-in identity management features or the improved interoperability with UNIX workloads, this upgrade might have been something on your horizon. However, for many of us, the capabilities added into the first R2 weren’t much to write home about. Heck, that first R2 attempt didn’t even feel like a whole new operating system (OS), with its “install-this-DVD-atop-an-existing-2003-server” upgrade process.

However, the story with Windows Server 2008 R2 appears dramatically different. Administrators everywhere are talking—a lot—about the new capabilities they’re looking to immediately implement once it hits the streets. Those capabilities go far beyond the much-ballyhooed Hyper-V improvements. They ring home towards solving problems that environments of all sizes actually have and need to resolve. What are those new capabilities? In a nutshell:

- ▶ Terminal Services becomes Remote Desktop Services, advancing Microsoft’s presentation virtualization play and adding some very interesting VDI support options.
- ▶ DirectAccess enables even small environments to create always-on, secured, Internet-based access to their internal LANs, reducing the pain of VPNs.
- ▶ BranchCache reduces the impact of transferring files over the network by caching those files in the branch office itself.
- ▶ AppLocker finally provides a manageable framework for whitelisting applications, dramatically reducing the capability of malware to infect your environment.
- ▶ BitLocker-to-Go takes Microsoft’s BitLocker full-drive encryption solution to the streets, protecting the data on the very external drives that keep security folks up at night.
- ▶ VPN Reconnect makes the process of using VPNs even easier, allowing VPN connections to effectively “roam” as necessary.
- ▶ File Classification Infrastructure creates what might be the first actually-useable native solution for enabling special protection on sensitive data.

All of these are features over and above the Hyper-V version 2 press machine that’s been getting more than its fair share of pundit processor cycles. But where this editor finds even more excitement in this release is in the target audience for each of these features: the small- and medium-sized environment. It is my opinion that one of Windows Server 2003 R2’s shortcomings related to whom it was targeted for: enterprises. Although those in the SMB markets might have interest in identity management, a richer Web platform, and UNIX interoperability, these topics are bandied about far more in enterprise circles. The problem is that many of these features were also already available by third parties, often with better support or more features; third parties that enterprises tend to already have relationships with and funding to purchase.

With this second go at an R2 release, Microsoft re-targets its new features directly towards those environments that need the features and are most likely to use them. Plus, you can argue that there are far more SMB environments than enterprise ones consuming Windows licenses today. As such, the excitement surrounding R2 version 2 feels palpable. Here at Realtime, we'll keep that excitement to a peak as we continue our team coverage on its topics. Keep it tuned to your regular Realtime Publishers station for continued reporting.

Let me know how it goes. The best story gets written up in a future edition of this eJournal. Drop me a line at [gshields@realtimepublishers.net](mailto:gshields@realtimepublishers.net). ♦

---

## Exciting New Training from Greg Shields and Don Jones

### **Citrix Presentation Server 5**

After watching Greg Shield's Citrix 5 Training Videos from CBT Nuggets, you'll know how to build remote application delivery infrastructure from start to finish. And you'll be fully prepared for the Citrix CCA 1Y0-A05 certification exam. [Click here](#) to watch a free video from Greg's series.

### **SQL Server 2008**

Don Jones's new CBT Nuggets training covers the SQL Server 2008 basics and loads of advanced concepts. Plus, it prepares you for Microsoft's 70-451 exam — your final step in SQL Server 2008 MCITP Certification. [Click here](#) to watch a free video from Don's series.

### **Nugget Streaming Subscription**

[Click here](#) to learn how you can watch all of Greg and Don's videos — and thousands of IT training videos by other great trainers — at one low annual price.



[www.cbtnuggets.com](http://www.cbtnuggets.com)  
888-507-6283 toll free  
541-284-5522 international

# Break the Habit

## *Haphazard Access Management*

---

by Don Jones

“We have a new user starting Monday—John Doe,” the boss says.

“Okay,” you reply. “What access does he need?”

“Oh, the same as Joe.”

“What access does Joe have, exactly?”

“Oh, just give him everything.”

Riiiiight. The fact is that this is too common a scenario, and managers—especially those outside IT—often don’t understand the importance of not giving everyone access to everything. They’re not thinking about it, and they just want the question—and the problem—to go away. Too

many times, however, we have to just let them get what they want—meaning we *know* we’re doing a poor job of managing access to resources, but what else can we do?

A first approach is to seek out some kind of tool that will help us report on what permissions someone actually has. These tools are complicated because Windows doesn’t centralize permissions. Instead, you basically have to query every single file server, every registry, as well as Exchange, Active Directory (AD), and maybe even databases stored in something like SQL Server. It’s a mess, but tools exist to do it: Dumpsec and Showacl

from the Windows Resource Kit provide a start, and commercial tools—like those from companies such as Quest and NetIQ—provide a more mature way of inventorying permissions. Tools like this help answer the question “What access does Joe have,” making it easier to “clone” that access.

A better approach is to use role-based management of permissions: *Never assign permissions to a person.* Instead, break down the job roles—titles like “Sales Rep,” for example—into roles, and then assign permissions to those roles. Again, commercial tools can help: Quest offers ActiveRoles

---



## 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>

Server for this purpose, and other companies have similar tools. However, in small- to medium-sized environments, you might be able to roll your own approach (although you'll obviously enjoy less functionality than a commercial toolkit supplies). Just use Windows' global domain groups to represent job roles within your organization and assign permissions to them. Yes, this seems simplistic—even the oldest Windows certification exams espoused this approach—but the trick is being consistent and thorough. For example, if one user, let's call him Dave, really does need a special and unique set of permissions, then create a global group to receive those permissions. It becomes easier to give someone else the exact same access that way.

If your organization is dealing with compliance—meaning you have to prove who has what permissions, and likely even track how permissions change over time—then trust me when I tell you that buying a commercial solution designed for this purpose is the only way to go. I know it's annoying to have to buy something to manage something as seemingly straightforward as permissions. However, having a database to track current and historical permissions is exactly what auditors need, and there's no way to provide that information (the historical bit, mainly) without a tool designed especially to do so.

But the end goal is to never assign permissions to a user account, no matter how tempting it is to do so just for “that one special user” or to get the job done quickly when you're busy with other tasks. ♦

*Don Jones is a co-founder of Concentrated Technology. Join him and cohort Greg Shields for intense Win2008 and Windows PowerShell training—visit [ConcentratedTech.com/class](http://ConcentratedTech.com/class) for more details. Ask Don a question by visiting [ConcentratedTech.com](http://ConcentratedTech.com) and using the “Contact” page.*

# Product Review

## UltraEdit

by Eric Schmidt

There are a number of multi-function text editors on the market today. Some people choose to use those that are offered for free, such as Notepad and WordPad, but there are many circumstances when one needs more functionality that the free products don't provide. UltraEdit by [IDM Computer Solutions](#), Inc is now in its fifteenth version and offers features that many of the free editors lack. Because UltraEdit is such a comprehensive tool, it would be difficult to examine all the features in this article. As such, the scope here will be limited to some of the features that are designed for systems administrators.

UltraEdit is organized into what they refer to as "Environments." An environment is the layout of the interface, including menus and toolbars. The benefit of this setup is that different environments can be created to support the current task. By creating layouts based on the work that needs to be performed, UltraEdit is able to be adapted to the task at hand instead of requiring the user to adapt to a static environment, thus improving efficiency.

The free tools that UltraEdit can replace are Notepad and WordPad. There are many circumstances when systems administrators need to review logs as part of troubleshooting an application. With some applications, these files can be so large that standard tools such as Notepad are unable to open them. The next step when Notepad fails is to move to WordPad, which is only marginally better at handling large files. This is where UltraEdit can really shine. It is capable of efficiently opening files of 4GB or more in size. Searching log files is also made easier because of the ability to perform multiple line searches. For example, this feature can be used to search a log where a date/timestamp appears over and over but only certain text below the timestamp needs to be found. With this feature, a log file search can be limited to specific events that occur on multiple lines.

Another benefit for systems administrators is the ability to create and edit script files of various types. UltraEdit leverages extensible wordfiles that contain language-specific keywords in order to provide syntax highlighting. By leveraging editable wordfiles that are user-customizable, UltraEdit has the ability to adapt to the environment in which it is being used. A prime example where this flexibility can be leveraged is in cases where non-standard, legacy, or in-house tools and languages are employed. In these situations, wordfiles can be created specifically to color-code those languages, which makes them much easier to maintain.

UltraEdit itself has an internal scripting engine. The scripting engine is referred to on the product page as "macros on steroids" because it enables the automated execution of tasks like macros do in Microsoft Word or Excel. In addition, this feature leverages the power of JavaScript, which allows programming logic to be included. This functionality could prove very useful when parsing/analyzing log files.

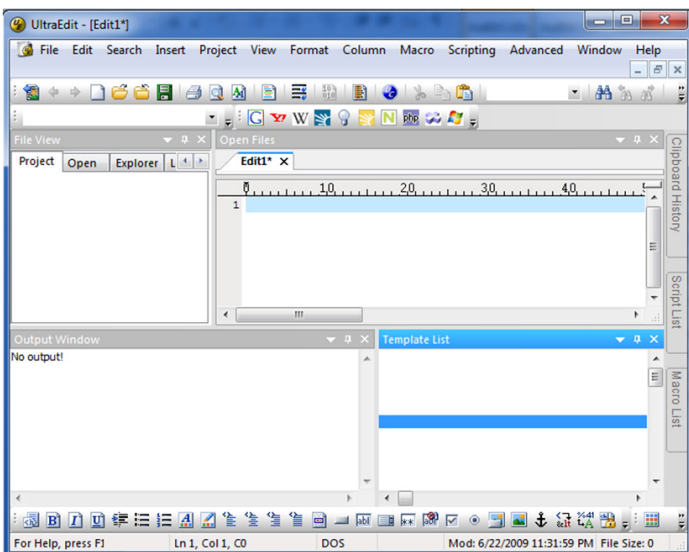



Figure 1: The UltraEdit default interface.

UltraEdit also contains an integrated ftp, telnet, and ssh console, which easily enables users to connect to the files they need to edit and thus eliminates the need to switch between tools when retrieving, editing, and re-posting files.

Overall, the capability that UltraEdit provides to systems administrators is a robust and fully customizable tool that is an ideal solution for many of the tasks that systems administrators perform on a daily basis. ♦

*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.*



Authored by **Eric Schmidt**

sponsored by



## New Essentials Series!

### Increasing Performance in Enterprise Anti-Malware Software

Featuring the Following Articles:

*Why is Traditional Anti-Malware So Slow?*

*Considerations for Evaluating Performance in Anti-Malware Products*

*Best Practices in Deploying Anti-Malware for Best Performance*

⌘ **Download it today** ⌘

# SQL Server High Availability

by Richard Siddaway

SQL Server is one of the major pieces of our technology jigsaw. In some respects, the whole of our enterprise computing effort can be said to centre on SQL Server. Why? Because SQL Server holds that most important of commodities—our corporate data. Whatever your organization does, it will need to store data that is important, if not vital, to its continued well being.

In a previous article, we looked at how we can protect that data through ensuring that it is backed up correctly. However, backup is not enough. Many organizations require their data to be continuously available. In some cases, the business will suffer financially if the database is down. That is known in admin speak as “a bad thing”—especially in the current economic climate. Losing a database can become an invitation to revise your resume. What we are looking for is a method of providing high availability to our SQL Server databases.

There are a number of techniques for providing high availability that SQL Server can utilize. This article will examine those techniques and provide information on the positive and negative aspects of the various techniques. Implementation details are available in the SQL Server documentation and will not be repeated here. The aim of this article is to provide sufficient information to choose the high-availability method that matches your requirements.

Before we get into the technical details, we need to think about what we mean by high availability.

## How High?

One measure of availability often used is the percentage availability.

| \$ Time Available | Downtime Allowed in 1 year | Cost       |
|-------------------|----------------------------|------------|
| 90                | 36.5 days                  | \$         |
| 95                | 18.25 days                 | \$+        |
| 99                | 3.65 days                  | \$\$       |
| 99.9              | 8.76 hours                 | \$\$\$     |
| 99.99             | 52.6 minutes               | \$\$\$\$   |
| 99.999            | 52.6 minutes               | \$\$\$\$\$ |

One question that must be asked immediately is “Does this include planned downtime?” Planned downtime is the time needed to perform maintenance such as patching or upgrades. If you have a system that is supposed to be available for 99.99% of the time, that means you have less than 1 hour per year to take it offline to perform any required maintenance. Better work quickly. In a 99.999% available scenario, your available downtime could disappear with a couple of failovers between your cluster nodes!



Before committing to a high-availability service level agreement (SLA) some questions need to be asked:

- ▶ Do we include planned down time?
- ▶ What time is available for a maintenance window?
- ▶ Is the system required 24x7 or just during business hours? Creating, and managing, a 99.99% available system to cover just business hours is easier than 24x7.
- ▶ How is availability to be measured?
  - Just because the server is up doesn't mean the data is available. If the users can't connect and work with the data, they won't regard it as available.
  - Likewise performance has to be at an acceptable level.

We also need to think about what we are trying to make resilient. This can be thought of in two ways:

- ▶ We need our servers to be resilient
- ▶ We need our data to be resilient

If we build a "traditional" active/passive two-node cluster that is attached to a SAN, we have made our servers resilient. If the active node fails, the SQL Server instance will failover to the passive node and the users can continue working. However, a single point of failure exists. There is a single copy of the data on the SAN and a problem with the SAN can take out your database, application, and business process even though the servers are up and running.

There are storage devices that will mirror your data automatically and, in the event of a failure on the primary storage, the secondary will take over the service. This protects the data but doesn't protect the service at the server level.



**Microsoft**  
GOLD CERTIFIED  
Partner

2009  
Preferred Partner  
GOLD  
hp

## THE QUICK WIN WIN:

Release payback on ROI IN 8 MONTHS  
*around simplification and operating cost reduction*

### SERVER CONSOLIDATION

### Get in Touch

For an Optimisation Assessment that will establish the benefits of consolidation and proven ROI in the shortest possible time, email [consolidate@scc.com](mailto:consolidate@scc.com) or visit [www.scc.com](http://www.scc.com)

“ We help customers **optimise** their infrastructure to drive out costs, release immediate cash savings and achieve long term operational efficiencies. ”



These two examples cover high availability within the data center. What happens if you lose access to the data center, or as a worse case, lose the whole data center? Do you need to build any disaster recovery capability into your high-availability solution?

When thinking about high availability, one aspect that is bound to quickly come to the foreground is cost.

### *How Much?*

As a general statement, high availability costs, and the more availability you need, the more it will cost. On a number of occasions, I have designed a high-availability solution having been told that budget is available only to be greeted by “How Much?” when the first cost estimates were produced. I would advise discovering the approximate budget before you start rather than later determining part way through that you don’t have the money available.

The preceding table gives a rough indication of the relative costs of various levels of availability. Most modern servers should provide 95 to 99% availability as a (rather sweeping) generality. After that, it starts to cost.

### *High Availability in the Organization*

High availability is not just a technical problem. There are three requirements for a system to remain in a state of high availability.



The right technologies have to be utilised. This includes the hardware and software that is involved in the system. Ideally, the system should be considered end-to-end (that is, from the users’ workstations to the server and storage). We can design, and build, a server system that ticks all of the high-availability boxes, but if the network goes down, the users will say the system isn’t available.

We need to ensure that the people running our systems are correctly trained and have the right mindset. High availability is no place for the “cowboy” who won’t follow procedures.

The third and probably most important aspect covers the procedures we adopt:

- ▶ Is our system documented?
- ▶ Is every procedure documented and followed?
- ▶ Is change control used correctly?


With this foundation established, we can look the various ways we can supply high availability for our SQL Server instances.

### **Backup and Restore**

We can implement every high-availability option we have but that does not take away the need to back up the data. The heart of every piece of availability planning is a restoration plan. The important aspect of this is the ability to restore—the backup is just a step on the path to restoration.

There is always the chance that everything is lost. Your only hope at that point is the offsite backup tape containing a recent version of the data.

Ensure backups are taken and that they are usable. The middle of a database failure incident is not the time to discover that none of the backups you have taken in the past 6 months are usable. A planned test restore is a necessary action.



Authored by **Dan Sullivan**

sponsored by



## **New eBook!**

### **Productivity and Cost Savings through Collaboration and Information Exchange**

Featuring the Following Chapters:

*Challenges to Efficient Collaboration and Information Exchange*

*Business Drivers for Improving Collaboration and Information Exchange*

*Technologies for Efficient Collaboration and Information Exchange*

*Management Best Practices for Efficient Collaboration and Information Exchange*

⌘ **Download it today** ⌘

It could be claimed that the option to restore from backup is an availability option. In the case of low-usage, non-business-critical databases, it may well be acceptable to restore the database rather than spend the money to make the system more resilient. This is a business decision based on the purpose for which the database is used. If this option is chosen, the restore time must be agreed upon with the business and expectations managed correctly.

## *Database Snapshots*

A database snapshot is a read-only point-in-time view of the database. It will not refresh as the source of the snapshot changes. Snapshots can only reside on the same server as the source database. Snapshots can be used for reporting purposes.

In the event of a user error, the source database can be reverted to the state of the snapshot. Snapshots are transactionally consistent as of the time they are taken, but any changes since the time of the snapshot will be lost. In addition, it is not possible to snapshot the system databases.

## *SQL Server Hardware*

It may seem like an obvious statement, but don't try cost cutting on the servers you implement for SQL Server. Without implementing any of the techniques we will discuss later in the article, it is possible to increase the resiliency of your systems:

- ▶ Use redundant components wherever possible:
  - Power supplies and redundant fans are prime examples
- ▶ Use redundant disk controllers (HBAs or however you are connecting)
- ▶ Provide resiliency to the network cards by teaming
- ▶ Protect the data by using RAID technologies as appropriate

This concept should be extended to your network so that multiple, redundant routes exist from the clients to your databases.

## *Terminology*

Each of the techniques has its own terminology—especially when discussing the data replication technologies. In order to be consistent, I will refer to the server holding the working, production copy of the database as the source. The server to which data is replicated will be referred to as the target. The aim of this article is to give an overview of possible techniques rather than a detailed description of individual technologies. Adopting a consistent terminology preserves the meaning without generating confusing details.

## *Failover Clusters*

Failover clustering is the first thing many people think of when the subject of high availability arises. A failover cluster consists of two or more physical servers (nodes) that are running a clustering service. The Microsoft Cluster Service is included in Windows. First introduced with Windows NT, clustering has become easier to implement with each successive generation. Initially only capable of containing two nodes, Windows clusters now scale to a total of 16 nodes with Windows Server 2008. Other clustering services are available from third-party vendors.

A cluster-aware application, such as SQL Server, is installed to the cluster. It will be configured so that it can run on any one of a number of nodes. In the event of a problem affecting the node on which SQL Server is currently hosted, it will automatically move (failover) to another node.

Clustering is often thought of as only a two-node activity. We have an active node actually running SQL Server, and we have a passive node sitting “doing nothing.” It is waiting for a failure to affect the active node so that the passive node can take over the task of running SQL Server. This is where a degree of the cost of high availability comes from in that 50% of the cluster is not actively working. Scale this issue up for an organization running a number of clusters, and the cost of the passive nodes becomes significant.

One potential way around this is to run the cluster nodes in an active-active arrangement. This means we have two instances of SQL Server running on the cluster—one on each node. In the event of a failure, both instances are running on a single machine. The hardware specification of the individual nodes must be higher to be able to accommodate the extra processing.

An alternative is to run a number of active nodes and a number of passive nodes in the same cluster. For instance, you may have three or four active nodes and two passive nodes for an N+2 cluster. This arrangement minimises the passive servers and reduces overall cost. There is a risk that if more than two servers fail, the remaining nodes may not be able to cope with the extra load; in reality, the probability of this arising is very low.

The ability to install SQL Server into a clustered environment is limited to SQL Server Enterprise Edition. The cost of this license together with the cost of buying hardware that is suitable, and approved, for clustering may make the cost prohibitive for many smaller organisations. Ideally, we would want the nodes in our cluster to be identical so that we could expect the same performance from either node.

Windows 2008 makes clustering easier to configure. As long as the cluster configuration is approved by the setup routine, it will be supported. This may ease the hardware burden.

One great advantage of clustering is that there are no additional configuration steps to perform to preserve client connectivity when a SQL Server instance fails over between nodes. The clients will lose their connection, but it will usually be regained either automatically or by an application restart. In the worst case, a client restart is required.

Windows servers also have the capability to participate in Network Load Balanced (NLB) clusters. NLB is usually used for load balancing TCP/IP traffic and resiliency in Web farms or similar scenarios. It is not suitable for providing high availability to SQL Server. In theory, read-only databases could be protected using this technique, but it is not a commonly-used approach.

Clustering protects at the server level but it does not protect the data. The data is a single point of failure in our infrastructure. If the storage system suffers a failure, we lose our data and our database is no longer available. We need to look at other technologies to provide a data resiliency capability.

### ***Service vs. Data***

We need to think about the differences between clustering and data replication technologies for high availability. Clustering protects the whole service from hardware failure. If the node that the SQL Server instance resides on fails, the whole instance fails over to another node. This includes the Master and MSDB system databases, which are essential, among other things, for the login and job information they contain.

If we use a data replication technology, we can duplicate the contents of the individual databases on another server. We cannot apply these techniques to the system databases. Other methods must be utilized to ensure the required login and job information is duplicated on the target server, for instance, using SSIS. This step is often forgotten until it becomes necessary to initiate a failover to the replication target, at which point the recovery time is extended. Include regular updates of this information as part of the replication schedule.

## Replication

All recent versions of SQL Server have the ability to replicate data between servers. Numerous replication scenarios exist with the possibility of replicating subsets of the data in the database, but here we are concerned with replicating the whole of the database to another server. There are two replication techniques we can consider: snapshot replication and transaction replication. Neither of these techniques provides a capability for automatic failover between the source and the target servers. However, these techniques can be used to replicate the database to a number of targets.

Snapshot replication takes a copy of the whole database and replicates it to the target server. The copy of the database on the target server is overwritten each time replication occurs. This technique can be thought of as an automated backup and restore. It is suitable for relatively small databases that do not change frequently. This is not the sort of technique that can really be run numerous times during the day!

Transactional replication is a much more granular approach to data replication. When a transaction is performed against the source database, the information in the database log about that transaction is replicated to the target server. The transaction is replayed so that the same change is performed on the target database. A schedule can be established to determine how frequently transactions are replicated to the target server. The data can be replicated immediately or delayed to ensure rogue transactions do not corrupt both copies of the data.

Replication is not limited to a single target server. It is possible to replicate to a number of servers to ensure multiple copies are available, preferably at different locations. A distribution server is used to control the replication process. This can be the source server, a target server, or a totally separate server. One potential issue exists because of the reliance on the distribution server. If it fails, replication stops, which removes our high availability! We need to ensure that the distribution server is protected (by clustering possibly) if we are going to rely on this technique.

Replication is deemed to be complex to set up and configure. This reputation is mainly derived from the configuration of merge replication and the need for conflict resolution. Transactional replication doesn't have this requirement and is therefore relatively straightforward to configure.

One advantage of using replication is that the target database can be accessed. It could be used for reporting purposes or backups could be taken against the target server rather than the source server.

## Log Shipping

Log shipping is similar to replication in some ways. It is based on the backup and restore mechanism within SQL Server. A backup of the transaction logs is taken periodically. The backup file is shipped (copied) to one or more target servers. Once it has been delivered, the backup is restored on to the target database. It is possible for multiple source servers to log ship to a single target server enabling a number of databases to be protected by a single server.

The mechanism of backup, copy, and restore is managed by the SQL Server Agent Service as a set of jobs. There is no mechanism for an automatic failover from the source server to a target server. Log shipping supplies a warm standby capability. A manual process is required to bring the target database online and make it available to the user population. Any outstanding backup files must be restored. If possible, a backup of the active log on the source database must be taken and applied to the target. The target database can then be brought online.

Queries can be executed against the target database if it is configured in STANDBY mode. However, the users will be disconnected when a restore occurs. If a delay in the restoration process is acceptable, the database can be used for querying. A backlog of backup files will accumulate until the users have disconnected and the restore process can be performed.

One drawback to log shipping is that it is based on the periodic backup and restore of the transaction logs. The backups must be applied in the correct sequence. If the sequence is broken for any reason, no further restores can occur. Log shipping will have to be re-initialized to recover from this situation. While this is happening, the level of protection available to the source database is reduced.

## *Mirroring*

Database mirroring was introduced with SQL Server 2005 and further enhanced in SQL Server 2008. Mirroring is a technology that protects the data rather than the server. It has to be implemented on individual databases. The system databases cannot be mirrored, so other techniques have to be used to provide the login information and so on to the target server.

The mirroring process maintains a duplicate of the source database automatically. The target can be configured to be a hot or warm standby. Correctly-written applications can be aware of the mirroring and fail over to the mirror target in the event of the primary database being unavailable.

There are a number of ways that mirroring can be configured:

- ▶ **Asynchronous mirroring.** This gives maximum performance but does not provide the maximum level of availability. Transactions are committed on the source server independent of the mirroring process. It is possible for the target and source databases to become out of step if there is a significant latency in the mirroring process. If the source database is lost, the target is available as a warm standby. A manual process is involved in bringing the target online and pointing the clients at the target database. This version of mirroring is ideal for a disaster recovery scenario where a delay in bringing the standby database online is acceptable.
- ▶ **Synchronous mirroring.** This provides the maximum protection but may incur a performance penalty. Once the initial synchronization has occurred between the source and target database, any transaction performed on the target database must be transmitted to the mirror target database and an acknowledgement received before it will be committed. This can impact performance on the source database if there is a significant delay in receiving the acknowledgement. The advantage is that all committed transactions are guaranteed to be written to disk on the mirror target.

When synchronous mirroring is performed, failover between the nodes can be manual or automatic. If automatic failover is required, a witness server must be brought into the scenario to establish a quorum of availability. The witness helps establish which database is available, in other words.

Mirroring will work with other high-availability technologies including clustering, log shipping, and replication. One possible example is to utilize a cluster to provide availability within the data center and to enhance that by using database mirroring to duplicate the database in an offsite location. This provides another layer of protection to the availability and provides the basis of our disaster recovery solution.

Database mirroring works in pairs across a mirroring session. A single source database is mirrored to a single target database. While the mirroring is operational, it is not possible to query the target database or to use it for other purposes such as backups. If a database snapshot is taken of the mirror target, clients can connect to it via the snapshot. However, it must be remembered that the snapshot is a point in time event, so the data will not be refreshed until another snapshot is generated.

## Virtualization

Virtualization is currently seen as one of the main future directions of server-based computing. The ability to host a number of virtual machines on one physical server provides a number of gains:

- ▶ Increased server utilization provides an improved return on our capital outlay
- ▶ The increase in server density reduces the amount of rack space, power, and air conditioning required by our data center operations
- ▶ Licensing costs may be reduced
- ▶ Operational flexibility is gained with the ability to easily move virtual machines between physical hosts

The automated failover of a virtual machine between two physical hosts provides a degree of resiliency. It is analogous to clustering in that we are protected from failures in the underlying hardware. However, there is no protection provided for the applications, such as SQL Server, or the data.

Virtualization technologies have limitations around the number of virtual CPUs that can be allocated to the virtual machine and the amount of memory that can be allocated. A database with a high I/O requirement may not be suitable for virtualization.

This leaves us in the position that virtualization becomes another option in our resiliency toolkit. It is an ideal for test and development environments, especially if there is a need to frequently create new machines. It may be ideal for small databases that do not require a high level of resources. It can be combined with other techniques; for instance, a mirror target could be hosted as a virtual machine rather than being a physical server.

## Comparison

We have looked at the individual techniques, but how do they compare?

| Technique                 | Automatic Failover            | Server Resiliency | Data Resiliency | Target Available for Other Uses |
|---------------------------|-------------------------------|-------------------|-----------------|---------------------------------|
| Clustering                | Y                             | Y                 | N               | n/a                             |
| Transactional Replication | N                             | N                 | Y               | Y                               |
| Log Shipping              | N                             | N                 | Y               | Y\N<br>(some reporting options) |
| Database Mirroring        | Y<br>(certain configurations) | N                 | Y               | N                               |
| Virtualization            | Y<br>(at server level)        | Y                 | N               | n/a                             |



Which high-availability technique is best for you? Unfortunately, only your organization can decide based on answers to these questions:

- ▶ What level of availability do you need?
- ▶ What level of availability can you afford?
- ▶ How will it fit with your disaster recovery plans?
- ▶ Is there availability around the rest of the business process? There is no point in making the database highly available if there are no clients to connect to the database.
- ▶ Can I test my high-availability solution?

SQL Server and high availability are not a closed and static subject. We see changes with each new release of SQL Server, either through the enhancement of existing technologies or the introduction of new options.

### **Future**

What does the future hold for SQL Server high availability? Exchange 2010 has abandoned using clustering as a high-availability technique. It now utilizes replication\mirroring-type technologies. Will future versions of SQL Server follow this path? It has advantages in terms of protecting the data and making it available, but Exchange utilizes Active Directory (AD) for a lot of the tasks the system databases perform in SQL Server. At the moment, the system databases are confined to a single server. Will there be a move to make SQL Server a more distributed environment by replicating information in the system databases or moving some of it into AD?

The advent of cloud computing brings other questions. Should your critical databases be hosted in the cloud rather than onsite? Disaster recovery and high availability become someone else's problem to manage, but what are the contractual issues? How do you test your high availability in these scenarios? Is your data secured when it is hosted by a third party?

The answers are unknown but are sure to add more interest and debate to this subject when they arrive. ♦

*Richard Siddaway is a Technical Infrastructure Architect with SCH plc. With over 20 years experience in various aspects of IT Richard is currently concentrating on the Microsoft environment at an architectural level especially around Active Directory, Exchange, SQL Server and Infrastructure Optimisation. Much of his recent experience has involved migrating enterprises from Windows NT\Exchange 5.5 to Windows 2003\Exchange 2003. Richard has administration experience with several database systems especially SQL Server. Richard is a PowerShell MVP who founded and currently leads the UK PowerShell User Group. Richard has presented to the Directory Experts Conference 2007, at various events at Microsoft UK, and for other UK User Groups. Richard can be contacted via his blog at <http://richardsiddaway.spaces.live.com/>*

# The Deep Dive

## A Business Perspective on SharePoint

---

by Robert Emmert

An application such as MOSS 2007 has the breadth of functionality to manage an enterprise's information, automate operations, integrate line of business systems, or serve as the company Web portal—and it all integrates into the Microsoft desktop most workers already know and use. With limitless possibilities, the question is not what *can* be implemented in SharePoint, but what *should* be implemented in SharePoint. The answer to that question depends on what your business is and who in your business you ask. This article looks at the question from the business view of the executives who manage the company, control the budgets, and decide whether your SharePoint project gets approved.

### *Working with Decision Makers*

In order to begin a SharePoint implementation, or extend an existing one, someone must authorize the project and approve its budget. Where IT managers used to control the expenditures for projects such as SharePoint, the current economic climate has increased financial oversight and moved approvals to higher- and higher-level executives. Understanding what motivates these decision makers is not only the key to moving forward at all but also to generating the strong sponsorship you need to sustain your efforts.

Early in my career, an older and wiser friend told me a story of how he learned to communicate with management. He was presenting his solution to a critical problem with a large-scale software system. He thoroughly explained the technical justifications and implementation details from his rigorous analysis, but none of the managers seemed interested. Then he noted that the alternative was to purchase two mainframes at seven and a half million dollars each. At that point, he had everyone's attention. What had been an obvious consequence to him wasn't so obvious to the non-technical managers. His primary concerns weren't their primary concerns.

Most of us have learned to frame IT requests in terms of reduced costs and efficiency gains, especially with financially-focused managers. Presenting a cost of money over time, or net present value analysis, can go a long way to convincing your CFO your request is well thought out. Although it is possible to quantify the cost savings of some productivity improvements, these analyses are largely subjective. Experienced managers have also been promised these types of results before, only to end up with projects that exceed budget, don't work as planned, or don't even deliver. Financial justification alone is not sufficient.

Instead of trying to better explain or justify productivity benefits, the key is to find capabilities that get decision makers excited. Capabilities that solve their tough problems or give them tools they have always wanted will resonate and get their attention. If the decision makers find your proposal highly valuable, they will find the money to pay for it.

It comes down to value, but your management has a different perception of value. Top-level executives are concerned with overall business strategies and balancing resources to achieve their goals. Most are non-technical. They have to manage revenue growth, build efficient operations, and allocate scarce financial resources. Instead of performing these individual tasks, they manage the people who do the tasks. Actually, they manage the managers who manage the people who do the tasks. To understand what they will value, you must consider their world.

Organizations are unique, and so are individual decision makers. Some may respond to SharePoint's basic productivity features and efficient intranet portal. Although the higher you get in a company, the more managers become concerned with overall business strategy and not technical details. The important issue is to understand the challenges and concerns of the decision makers in your organization from their perspective. Framing your project with that view demonstrates business value that opens up

budgets. Taking this approach has the added benefit of changing the way executives look at the IT function. Solving business problems and furthering business goals makes IT a strategic partner in the business plan—not just a necessary expense.

### ***SharePoint Capabilities that Resonate with Business-Focused Managers***

As IT professionals, we must deal with day-to-day support systems and the problems people have with them. We see first-hand how functionality, such as better document management and easy access to company information, will increase everyone's productivity. Executive managers live in a different world. Although they may still be exposed to daily efficiency issues, those activities are not their chief concern. As much as decreasing the time it takes to locate contracts or automating a tedious process would reduce effort and costs, decision makers don't intuitively perceive a compelling value. They can usually solve those problems by just making people work more hours. Without a direct, overwhelming cost savings, general productivity features won't get your SharePoint project authorized.

During a recent planning session with an executive of a small company, I discussed using SharePoint as an alternative to file shares. One of the most useful capabilities in SharePoint is the ability to replace file shares with a real document management system. Eliminating duplicate copies of documents, controlling versions, automating retention policies, and using an effective search system are a significant productivity boost over arbitrary and uncontrolled directories. He wasn't interested. What did get him excited was the simple ability to collect industry-specific links and information on role-based pages for his business analysts. That capability was the core of his business need and what got traction.

So what SharePoint features do resonate with business users and decision makers? Each manager is unique, so no formula will tell you what to propose, but there is an approach to help you find the answer. Be prepared to listen and build an understanding of your management's specific problems and concerns. Learn what their day encompasses; that knowledge can spark ideas. Think about their perspective and the problems they need to solve.

One highly fruitful area is SharePoint's business intelligence functionality. Executives are paid to develop and execute business strategies, their plan for the organization. SharePoint has exceptional abilities to collect and present information. Being able to more effectively communicate their vision and influence the entire organization can be a motivating feature. With the addition of key performance indicators, executives can also communicate and track progress towards implementing their strategies. Another potentially valuable feature comes from asking an executive what information he or she would like but isn't currently getting. This conversation can uncover data integration and reporting needs. Features like these and others that address higher-level business objectives can be the path to creating demand for SharePoint.

Understanding the desires of the business users and developing business-focused features doesn't, however, relieve you of the responsibility to design a SharePoint system that can serve the entire organization. You still need to architect a system that delivers availability, performance, and general utility. Your CFO may not be excited by SharePoint's search capability, but it's still a critical piece of your portal's infrastructure.

### ***Conclusions***

SharePoint offers an incredible collection of enterprise capabilities that can streamline business operations, improve team productivity, and generate timely insight. All this potential is irrelevant unless you can get your organization's decision makers to understand the value. By looking at SharePoint from their perspective and solving problems that help overall business objectives, you can not only get your projects approved but also become an important part of your organization's strategic plan. In tough economic times, going beyond technical expertise to differentiate yourself as someone who understands the business and contributes to its growth makes you more valuable as well. The job you save may be your own. ♦

*Robert Emmert is a business process and technology consultant with more than 20 years of IT and management experience. His SharePoint focus includes planning, architecture, business intelligence, workflow automation, and custom Web part development.*

# Practical PowerShell

## Excel-lent PowerShell

---

*by Jeffery Hicks*

You can download a zip file with all these scripts from [http://www.realtime-windowsserver.com/code/v2n7\\_Practical\\_PowerShell.zip](http://www.realtime-windowsserver.com/code/v2n7_Practical_PowerShell.zip).

I know many Windows administrators find Microsoft Excel to be an integral part of their administrative toolbox. Very often, it is used as a quick and dirty database or storage facility for all types of information. Perhaps you've even gone so far as to publish Excel information to a Web page on your intranet. How does PowerShell factor in? Yes, you can integrate Excel into PowerShell scripting by using the Excel COM object, just as you might have done in VBScript. However, this month, let me show you another way to integrate PowerShell and Excel that you might not have thought of.

The steps I'll walk through assume you are using Microsoft Office 2007, although I believe most of the functionality is available in older versions as well. You'll have to adjust accordingly. All of the sample files and scripts I'm going to discuss [are available for download](#) from the Realtime Web site.

If you've used PowerShell for any length of time, I'm sure you've realized how much information it can gather. Why not use Microsoft Excel as a vehicle for displaying that data? Sure, you could write a script to take data and create a worksheet, but working with COM can get a little tedious. Wouldn't it be nicer to simply toss the data into Excel so that you could format it and create charts from it or whatever without having to repeat the steps every time you wanted to update the data?

Microsoft Excel is a veritable Swiss army knife. One feature we can take advantage of is Excel's ability to display data from different data sources, like databases and the Web. We will use PowerShell as a data source and Excel as the display mechanism. Well, technically, we'll use PowerShell to create a data source that Excel can connect.

One data source that Excel can leverage is a simple CSV file. PowerShell can easily create rich CSV files. Put them together, and you can create a PowerShell reporting or monitoring solution that uses Excel as its display interface. Here's how you might do this.

First, we need a few PowerShell expressions to generate some data. Gathering service information is always helpful:

```
Get-Service | select name,displayname,status,@{Name="Last Check";Expression={{(get-date)}}
```

This is essentially the output from **Get-Service** with the addition of a custom property that captures the date and time when the expression is executed.

What about drive information? I have a short script that takes a computer name as a parameter:

```
#Get-DriveData.ps1
Param([string]$computername=$env:computername)

Get-WmiObject win32_logicaldisk -filter "drivetype=3" `
-computername $computername |
select SystemName,DeviceID,Volumename,Size,FreeSpace,@{
Name="Free";Expression={
"{0:P2}" -f ((($_.freespace -as [int64])/($_.size -as [int64]))}
},@{Name="Last Check";Expression={(get-date)}}
```

Finally, I'll capture some process information:

```
Get-Process | where {($_.workingset -gt 10mb -and $_.id -ne $pid} | sort workingset -desc
| select -first 10 |
select Handles,NPM,PM,WS,VM,CPU,ID,ProcessName,Path,@{Name="Last
Check";Expression={(get-date)}}
```

This expression will return the top-10 processes based on workingset greater than 10MB. It will exclude the PowerShell process. The `$pid` variable is an intrinsic PowerShell variable that returns the process ID for the current console.

Once I've verified that all these expressions return information to the console, I'll modify the expressions to pipe the output to **Export-CSV**, sending each information to a separate CSV file. In fact, I'm going to plan ahead and put all these commands in a single PowerShell script:

```
#Update-ExcelData.ps1
Write-Host "Updating service data" -foregroundcolor CYAN
Get-Service | select name,displayname,status,@{Name="Last Check";Expression={(get-date)}}
| Export-Csv c:\test\servicedata.csv -notypeinformation

Write-Host "Updating drive information" -foregroundcolor CYAN
c:\scripts\posh\get-drivedata.ps1 | Export-Csv c:\test\drivedata.csv -notypeinformation

Write-Host "Updating process information" -foregroundcolor CYAN
Get-Process | where {($_.workingset -gt 10mb -and $_.id -ne $pid} | sort workingset -desc
| select -first 10 |
select Handles,NPM,PM,WS,VM,CPU,ID,ProcessName,Path,@{Name="Last
Check";Expression={(get-date)}} |
Export-Csv c:\test\procddata.csv -notypeinformation
```

The `-notypeinformation` parameter is important, as I don't need the PowerShell type information. Run this script to create the CSV files and verify they exist.

**CSV Limits**—The CSV format is limited to simple values like strings, dates, and numbers. It won't help much if one of the serialized object's properties is another object. Still, I think for the most part, the information you want to monitor fits this limitation.

Now let's set up Excel. Launch it, and create a new workbook. My workbook will have three worksheets that will hold data from my three PowerShell expressions. I saved the file with the Excel 2007 xlsx extension.

On the first worksheet, select Data from the ribbon. In the Get External Data panel, select "From Text." Browse to the appropriate CSV file. On my first tab, I'm displaying service information, so I loaded servicedata.csv. You'll likely need to change the delimiter to a comma and import the file. Adjust column types if necessary. You can insert the data where you want into your worksheet, but I prefer to start in cell A1.

Of course, you aren't limited to these simple values. Format the cells however you prefer. I formatted the column headings, for example, and used conditional formatting so that if the service was stopped, the service status is displayed in a red font, as shown in Figure 1.

|    | A                              | B  | C       | D               | E |
|----|--------------------------------|--|---------|-----------------|---|
| 1  | Name                           | DisplayName                                  | Status  | Last Check      |   |
| 2  | AeLookupSvc                    | Application Experience                       | Running | 5/18/2009 12:40 |   |
| 3  | ALG                            | Application Layer Gateway Service            | Stopped | 5/18/2009 12:40 |   |
| 4  | Appinfo                        | Application Information                      | Stopped | 5/18/2009 12:40 |   |
| 5  | AppMgmt                        | Application Management                       | Running | 5/18/2009 12:40 |   |
| 6  | AudioEndpointBuilder           | Windows Audio Endpoint Builder               | Running | 5/18/2009 12:40 |   |
| 7  | AudioSrv                       | Windows Audio                                | Running | 5/18/2009 12:40 |   |
| 8  | BFE                            | Base Filtering Engine                        | Running | 5/18/2009 12:40 |   |
| 9  | BITS                           | Background Intelligent Transfer Service      | Running | 5/18/2009 12:40 |   |
| 10 | Browser                        | Computer Browser                             | Running | 5/18/2009 12:40 |   |
| 11 | CertPropSvc                    | Certificate Propagation                      | Stopped | 5/18/2009 12:40 |   |
| 12 | clr_optimization_v2.0.50727_32 | Microsoft .NET Framework NGEN v2.0.50727_X86 | Stopped | 5/18/2009 12:40 |   |
| 13 | clr_optimization_v2.0.50727_64 | Microsoft .NET Framework NGEN v2.0.50727_X64 | Stopped | 5/18/2009 12:40 |   |
| 14 | COMSysApp                      | COM+ System Application                      | Stopped | 5/18/2009 12:40 |   |
| 15 | CryptSvc                       | Cryptographic Services                       | Running | 5/18/2009 12:40 |   |
| 16 | CscService                     | Offline Files                                | Stopped | 5/18/2009 12:40 |   |
| 17 | DcomLaunch                     | DCOM Server Process Launcher                 | Running | 5/18/2009 12:40 |   |
| 18 | Dhcp                           | DHCP Client                                  | Running | 5/18/2009 12:40 |   |

Figure 1: An example services report.

I encourage you to try out the different conditional formatting options. Repeat this process for each CSV file. Figure 2 illustrates what I did with my disk usage information.

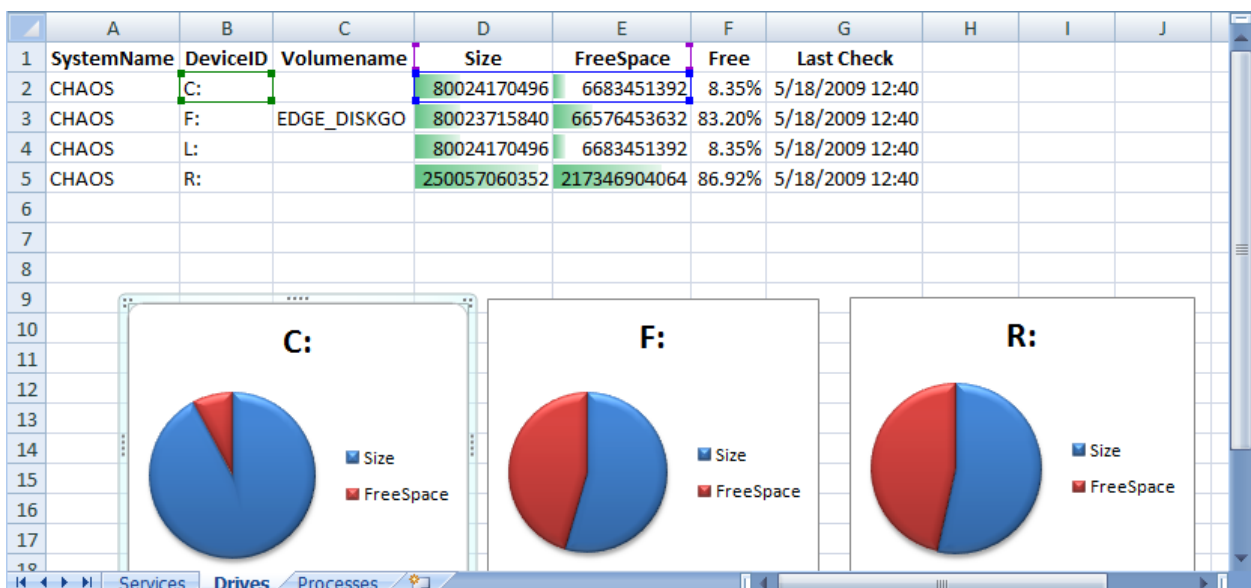


Figure 2: An example drive report.

Not only am I using some fancy conditional formatting, I've also created a few pie charts from the data. Figure 3 shows my last PowerShell drive report.

|    | C         | D         | E         | F           | G     | H            | I  |
|----|-----------|-----------|-----------|-------------|-------|--------------|--|
| 1  | PM        | WS        | VM        | CPU         | Id    | ProcessName  | Path   |
| 2  | 192356352 | 195125248 | 314544128 | 1016.471316 | 1092  | svchost      | C:\Windows\System32\svchost.exe                    |
| 3  | 117669888 | 103301120 | 566398976 | 49.2651158  | 1112  | svchost      | C:\Windows\system32\svchost.exe                    |
| 4  | 97673216  | 101462016 | 425668608 | 492.4639568 | 3860  | explorer     | C:\Windows\Explorer.EXE                            |
| 5  | 89837568  | 98705408  | 224952320 | 53.04034    | 10928 | thunderbird  | C:\Program Files (x86)\Mozilla Thunderbird\thunde  |
| 6  | 72482816  | 82259968  | 184578048 | 16.1461035  | 10184 | firefox      | C:\Program Files (x86)\Mozilla Firefox\firefox.exe |
| 7  | 92020736  | 72437760  | 362770432 | 4609.314747 | 3476  | PrimalScript | C:\Program Files (x86)\SAPIEN Technologies, Inc\Pr |
| 8  | 34037760  | 59252736  | 254685184 | 11.8248758  | 8180  | EXCEL        | C:\Program Files (x86)\Microsoft Office\Office12\E |
| 9  | 45961216  | 59052032  | 266842112 | 1813.776827 | 3636  | Moe          | C:\Users\Jeff\AppData\Local\Microsoft\Live Mesh\   |
| 10 | 24338432  | 36761600  | 177754112 | 108.5298957 | 5160  | MoeMonitor   | C:\Users\Jeff\AppData\Local\Microsoft\Live Mesh\   |
| 11 | 31870976  | 35414016  | 175083520 | 63.5080071  | 4344  | SnagitEditor | C:\Program Files (x86)\TechSmith\Snagit 9\snagitec |
| 12 |           |           |           |             |       |              |  |
| 13 |           |           |           |             |       |              |  |
| 14 |           |           |           |             |       |              |  |
| 15 |           |           |           |             |       |              |  |
| 16 |           |           |           |             |       |              |  |
| 17 |           |           |           |             |       |              |  |
| 18 |           |           |           |             |       |              |  |

Figure 3: An example process report.

Up to this point, we have a static display. To make it livelier, on the Data ribbon, click connections. Find the data connections you just created, and select Properties. You'll want to uncheck "prompt for filename on refresh." You can configure other refresh settings that you feel are appropriate. If your CSV files will be updated on a regular basis and you intend to keep Excel open as a monitoring tool, you might want to configure the "Refresh Every X minutes" option. Figure 4 shows a sample configuration.

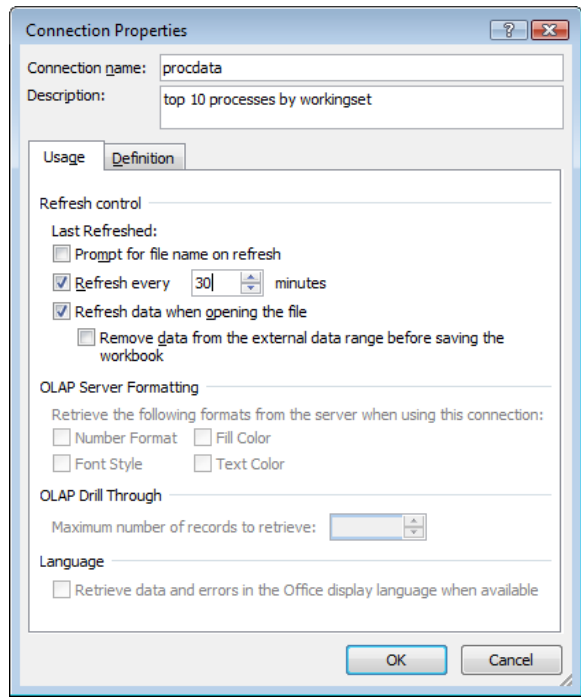


Figure 4: An example connection configuration.

You'll also likely need to configure security settings to allow updating from external content. Otherwise, you'll get a warning every time you open the spreadsheet. The easiest approach is to edit the Trust Center settings under Excel Options. Navigate to Trusted Locations, and add the folder that contains your CSV files.

Now, let's see this in action. Close Excel, saving your changes. In PowerShell, run the script to update the CSV files. Open your Excel reporting workbook, and you should see that the data has been updated. If disk usage has changed, you'll also see the graphs have been updated. You can always manually refresh the data connections in Excel if you've updated the CSV files.

Certainly, you can take this idea further. For example, you could create a macro-enabled version of this workbook. Create an Excel macro to run your PowerShell script and refresh all the data connections. Now all your functionality is contained in a single Excel file.

Or perhaps you'd like to publish the Excel file. You could use Excel Services and SharePoint 2007 or publish as a static html file, which you could push to an intranet server. The technical details of these enhancements go beyond the scope of this column, but I hope you are inspired. ♦

*Jeffery Hicks (MCSE, MCSA, MCT) is a Microsoft PowerShell MVP and Scripting Guru for SAPIEN Technologies. Jeff is a 17 year IT veteran specializing in administrative scripting and automation. Jeff is an active blogger, author, trainer and conference presenter. His latest book is Managing Active Directory with Windows PowerShell: TFM (SAPIEN Press). Follow Jeff at [Twitter.com/JeffHicks](http://Twitter.com/JeffHicks) and [blog.sapien.com](http://blog.sapien.com). You can contact Jeff at [jhicks@sapien.com](mailto:jhicks@sapien.com).*

**World's hottest IT topics**

- Windows PowerShell 2nd Edition
- Windows PowerShell 3rd Edition  
(covers Windows PowerShell v2.0)
- ADSI Scripting
- WSH and VBScript Core
- Windows Server 2008: What's New/What's Changed
- Exchange Management Shell
- Managing Active Directory With Windows PowerShell
- Managing VMware Infrastructure With Windows PowerShell

**SAPIEN PRESS**

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

The advertisement features a collage of book covers from SAPIEN PRESS, including titles like 'ADSI Scripting', 'WSH and VBScript Core', 'Managing Active Directory With Windows PowerShell', 'Exchange Management Shell', 'Windows PowerShell 2nd Edition', and 'Windows Server 2008: What's New/What's Changed'.



# Exclusively Exchange

## Exchange Server 2007 SP2 Coming 3Q

by J. Peter Bruzzese

The last service pack for Exchange brought with it a slew of new features, including an additional high-availability option in the form of Standby Continuous Replication (SCR), a Public Folder management tool so we didn't have to manage folders through the Exchange Management Shell if we didn't want to, and many more ActiveSync policy settings. So, what will this new service pack bring our way?

The following list highlights customer benefits within the upcoming SP2:

- ▶ **Enhanced Auditing**—This addition is certainly a beneficial feature for Exchange administrators looking to audit their Exchange servers
- ▶ **Exchange Volume Snapshot Backup Functionality**—Before you get too excited, this feature has some serious stipulations that are discussed later in this article. However, it is just nice to see it part of the lineup considering the fact that you couldn't back up Exchange 2007 off Server 2008.
- ▶ **Dynamic Active Directory Schema Update and Validation**—For those of you who have to deal with Active Directory (AD) schema updates, this feature will be useful in preventing conflicts and helps to dynamically deploy future schema updates.
- ▶ **Public Folder Quota Management**—SP2 improves the PowerShell cmdlets needed to perform quota management.
- ▶ **Centralized Organizational Settings**—This is another PowerShell option to enable centralized management of Exchange organizational settings.
- ▶ **Named Properties cmdlet**—This cmdlet enables a new way to monitor named properties per database.
- ▶ **New User Interface for Managing Diagnostic Logging**—Found with the EMC, this feature will allow for an easier logging experience for Exchange Administrators to configure.

One important item to keep in mind is that Exchange 2010 is coming out later this year. You should know that you need to deploy SP2 for your existing environment in order to interoperate with Exchange Server 2010 and transition services.

Let's return to the backup solution coming with SP2. Here is what the Exchange Team listed out for stipulations as to its use:

- ▶ It is only VSS-based, no streaming ESE backups (which is fine, actually). The VSS backup must be run locally as well, although you can back up to a remote network share if you prefer.
- ▶ Backups are at the volume level. Thus, you have to back up the entire volume that is holding the storage groups—not a storage group or database or mailbox (which would be nice granularity to a free backup solution, even I'll admit).
- ▶ You can perform only full backups and only perform full backups of the active copy of a database (no passive backup, which is a bit odd considering the VSS solution is really perfect for an LCR or CCR passive side backup).
- ▶ You can restore only the Exchange data but you must restore all backed up storage groups and databases in a single movement. You cannot restore to a Recovery Storage Group, but you can restore to an alternative location and manually move the data to an RSG.

I'm sure there are more stipulations, but those are the main points. Ultimately, if you really want a backup solution for your Exchange 2007 Server running on Server 2008, you will be better off looking for a third-party solution. Microsoft would love to offer you the System Center: Data Protection Manager, which performs the backup quite beautifully, as it rightfully should coming from the same vendor that creates the Exchange messaging application.

So, we look forward to this SP2 for Exchange 2007. And even more so, we look forward to Exchange 2010 with the bevy of new features coming our way with that release as well. ♦

*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). His most recent book "Exchange 2007 How-To" was published by Sams in January 2009. 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).*

#### ExclusivelyExchange.com Free Training Videos

Would you like to learn more about Exchange 2007 and 2010? Check out the free training videos at [www.exclusivelyexchange.com](http://www.exclusivelyexchange.com). And if you want to learn about other subjects like SharePoint, Server Core, Hyper-V, and more...check out [www.cliptraining.com](http://www.cliptraining.com).



### 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](http://www.ClipTraining.com)

Email: [info@ClipTraining.com](mailto:info@ClipTraining.com)

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

## Copyright Statement

© 2009 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). ♦