THE MICROSOFT JOURNAL FOR IT PROFESSIONALS
TechNet
MAGAZINE

May 2010

Stardock Tweak7

Tweak7

Welcome
Resources
Security
Display
Power
Performance
Usability
Misc Tweaks
About

Security

Windows contains advanced security features.
These features can help protect against malware, but hinder usability.
Tweak7 helps you customize this protection, or turn it off altogether.
Membership of the Administators group is required to make any changes.

Overview | Advanced

Elevation prompts

● High security - require login details for elevation
○ Medium security - require user consent for elevation (default)
● Medium-low security - remove all prompts for elevation
● Low security - disable User Account Control entirely

Changes to or from low security require a reboot before taking effect
Alerts for medium-low and low security can be adjusted or disabled in Security Center

Login requirements

● Require Ctrl+Alt+Del and a password to login
○ Require a password to login (default)
● Automatically login as a particular user

For more detailed changes: User accounts Security policy Security Center

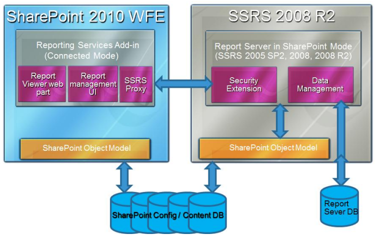

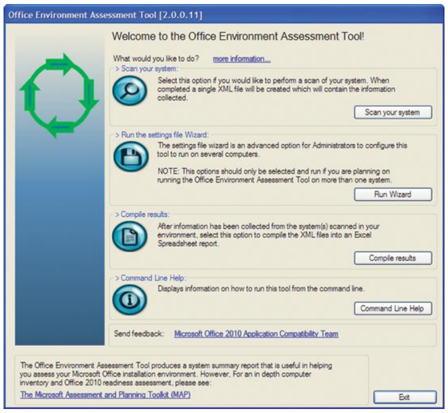

Office Environment Assessment Tool [2.0.0.11]

Welcome to the Office Environment Assessment Tool!

What would you like to do? more information...

> Scan your system:
Select this option if you would like to perform a scan of your system. When completed a single XML file will be created which will contain the information collected.
Scan your system

> Run the settings file Wizard:
The settings file wizard is an advanced option for Administrators to configure this tool to run on several computers.
NOTE: This options should only be selected and run if you are planning on running the Office Environment Assessment Tool on more than one system.
Run Wizard

> Compile results:
After information has been collected from the system(s) scanned in your environment, select this option to compile the XML files into an Excel Spreadsheet report.
Compile results

> Command Line Help:
Displays information on how to run this tool from the command line.
Command Line Help

Send feedback: Microsoft Office 2010 Application Compatibility Team

The Office Environment Assessment Tool produces a system summary report that is useful in helping you assess your Microsoft Office installation environment. However, For an in depth computer inventory and Office 2010 readiness assessment, please see:
The Microsoft Assessment and Planning Toolkit (MAP)

Exit

Editor’s Note

Keeping Tabs on Remote Offices

Lafe Low

With the myriad challenges you face every day as an IT manager, keeping a corporate network running that supports all sorts of remote offices has got to rank somewhere near the top. You have all the standard challenges—providing the right applications, ensuring certain performance and service levels, maintaining bandwidth and capacity planning, ensuring a rock-solid backup and disaster recovery plan—and now you have to do it from miles away. And there’s often no one at the remote office who “speaks IT.”

Reliability, capacity and security become even more important when your supporting branch offices are spread out all over the place. Thankfully, advances in server technology and domain controllers (DCs) make your job a little easier. Check out Paul Yu’s pair of features this month that cover the entire process of supporting remote offices with read-only DCs. Yu takes us through planning, deployment and configuration. He delves into Windows 7 and Windows Server 2008 R2, and how the new RODCs can help you tie up your organization’s network no matter how many locations you’re supporting.

Is this type of remote support something that affects your daily life as an IT manager? Are you supporting a handful of branch offices? We’d love to hear your story. How is it working? How did you develop your plan and deploy the technology? How do you maintain those connections? Let us know and perhaps we’ll provide some “down in the trenches” coverage in the coming months. Yours are the stories we want to tell.

We’re here for you. Let us know what you think. What do you want to see on the site? What do you like about our coverage? What else would you like to see? We love feedback, so feel free to visit our LinkedIn group, send us an e-mail at tnmag@microsoft.com or e-mail me directly at the address listed below. I look forward to hearing from you.

Lafe Low *is the editor in chief of* TechNet Magazine*. A veteran technology journalist, he’s also the former executive editor of 1105 Media’s* Redmond *magazine. Contact him at llow@1105media.com.*

Windows Server 2008 R2 Domain Controllers

Plan Carefully for RODCs

Paul Yu

When physical security is lacking, it becomes essential to increase the focus on data security. Windows Server 2008 and R2 provide some new ways to do so that seem uniquely tailored for environments like remote offices where physical security may not be as tight. Read-only domain controllers (RODCs) are a new feature of the Active Directory Domain Services (AD DS) in the Windows Server systems. They represent a fundamental change to how you'd typically use domain controllers (DCs).

Because many of RODCs' new capabilities impact key aspects of the design and deployment process, it's important to understand how you can leverage them in your enterprise. There are also critical design and planning considerations you must take into account before introducing them into your environment. RODCs are DCs that host complete, read-only copies of Active Directory database partitions, a read-only copy of SYSVOL, and a Filtered Attribute Set (FAS) that restricts the inbound replication of certain application data from writable DCs.

By default, RODCs do not selectively store user and computer account credentials, but you can configure them to do so. That alone typically warrants using RODCs in remote branch offices or perimeter networks lacking the physical security commonly found in datacenter intranets. RODCs also provide other less-well-known security features, such as a special Kerberos ticket-granting account that addresses ticket-based attacks associated with the RODC itself becoming compromised.

Although security concerns are the most common reasons to deploy RODCs, they also provide a number of other advantages, like enterprise manageability and scalability. Generally speaking, RODCs are meant for environments that require local authentication and authorization, but lack the physical security to safely use writable DCs. Therefore, RODCs are most common in datacenter perimeter networks or branch office locations.

One example of a good use of RODCs is a datacenter that requires AD DS, but because of security constraints can't leverage the corporate AD DS forest in the perimeter network. In this case, RODCs might meet relevant security requirements, therefore changing the infrastructure scope of the corporate AD DS implementation. This type scenario will likely become more frequent. It also reflects current best practice AD DS models for perimeter networks, such as the extended corporate forest model.

Branching Out with RODCs

The most common environments for RODCs using AD DS are still branch offices. These types of environments are typically end points in a hub-and-spoke network topology. They are widely distributed geographically, in large numbers, and they individually host small user populations, connect to hub sites by slow, unreliable network links, and often lack local, experienced administrators.

For branch offices already hosting writable DCs, it's probably unnecessary to deploy RODCs. In this scenario, however, RODCs may not only meet existing AD DS-related requirements, but also exceed them with regard to tighter security, enhanced management, simplified architecture and lower total cost of ownership (TCO). For locations where security or manageability requirements prohibit using DCs, RODCs can help you introduce DCs into the environment and provide a number of beneficial, localized services.

Although the new features and benefits make evaluating RODCs compelling, there are additional factors to consider, like application compatibility issues and service impact conditions. These could render RODC deployments unacceptable for certain environments.

For example, because many directory-enabled applications and services read data from AD DS, they should continue to function and work with an RODC. However, if certain applications require writable access at all times, an RODC may not be acceptable. RODCs also depend on network connectivity to a writable DC for write operations. Although failed write operations may be the cause of most well-known application-related issues, there are other issues to consider, such as inefficient or failed read operations, failed write-read-back operations, and general application failures associated with the RODC itself.

Besides application issues, fundamental user and computer operations can be affected when connectivity to a writable DC is disrupted or lost. For example, basic authentication services may fail if account passwords are not both cacheable and cached locally on the RODC. You can easily mitigate this issue by making accounts cacheable through an RODC's Password Replication Policy (PRP), and then caching the passwords through pre-population. Performing these steps also requires connectivity to a writable DC.

Along with other authentication issues, password expirations and account lockouts are significantly impacted when connectivity to a writable DC is unavailable. Password change requests and any attempts to manually unlock a locked account will continue to fail until connectivity to a writable DC is restored. Understanding these dependencies and subsequent changes in operational behavior is critical to ensuring your requirements and any service level agreements (SLAs).

There are several general scenarios in which you can deploy RODCs. They're useful in locations that don't have existing DCs, or in locations that currently host DCs that will either be replaced or upgraded to a newer version of Windows. Although there are comprehensive planning considerations specific to each scenario, we'll focus here on non-specific approaches. They are, however, distinct to RODCs, rather than to traditional writable DCs.

Preplanning

Before you start any formal RODC planning, you should conduct an appropriate level of due diligence and fundamental AD DS preplanning. This should include key tasks like validating the existing AD DS logical structure, and ensuring the administration model and the AD DS physical structure supports existing business and technical requirements. You'll also have to consider hardware requirements, software upgrade strategies, and applicable operating system known issues, and to evaluate RODC AD DS prerequisites. This information will be critical to the planning and deployment processes. You'll find it's well-documented in detailed deployment check lists.

Installation and Management

There's a substantial manageability feature in RODCs called Administrator Role Separation (ARS). This delegates to non-service administrators the ability to install and administer RODC servers, without granting Active Directory rights. This is a significant change to the traditional considerations with respect to DC server design, delegation of administration, and deployment procedures. This role separation becomes increasingly important with critical applications requiring direct installation on a DC, or for locations that host single, multi-purpose servers.

Additional Server Roles

As a general rule, you should eliminate from the server all roles not required for the RODC to function. Therefore, the only roles you should add to RODCs are the DNS and global catalog server roles. You should install the DNS server role on each RODC so local DNS clients can perform DNS resolution when network connectivity to a writable DC is unavailable. However, if the DNS server role is not installed through Dcpromo.exe, you'll have to install it afterward. You have to use Dnscmd.exe to enlist the RODC in the DNS application directory partitions that host the Active Directory integrated zones. You should also configure RODCs as global catalog servers so they can perform authentication and global catalog queries using just the RODC. From an authentication standpoint, if the global catalog role is not an option, you can use universal group caching. Successful authentication to an RODC may ultimately be dependent on the RODC's PRP configuration.

RODC Placement

DC placement has changed considerably since the introduction of the RODC PRP. RODCs must be able to replicate the domain partition from a writable DC running Windows Server 2008 or Windows Server 2008 R2 in the same domain, because only these DCs can enforce the PRPs for RODCs. To ensure proper replication, the writable DC should be placed in the AD DS site that has the lowest cost site link to the site containing the RODC.

If you can't establish this configuration, the RODC replication will need to depend on the Bridge all site links option (meaning site link transitivity) or site link bridges between whichever site links contain the RODC site and the site of the writable DC. If site transitivity or site link bridges aren't an option, you can create new site links to directly connect the RODC site and the site of the writable DC.

As a general best practice, you shouldn’t place other DCs in the same AD DS site as the RODC, because client operations may become inconsistent, making client behavior unpredictable. Basic operations like authentication, LDAP reads and writes, and password changes can all behave differently depending on disparate RODC configurations, the Windows version of a writable DC, and whether or not network connectivity to other writeable DCs is available. You should also keep all users and resources from a single domain in an RODC site. RODCs don’t store trust passwords and they require cross-domain authorization to forward authentication requests to different writable DCs in each domain. Assuming writable DCs reside in separate sites, all cross-domain authentication requests would be dependent on network connectivity and wouldn’t work in the event of a network connectivity failure.

Scalability and Replication

RODCs also provide scalability benefits that are helpful for larger or more complex AD DS implementations. For example, RODCs provide unidirectional replication. Therefore, deploying RODCs in branch offices reduces the performance load on hub site bridgehead servers that normally process inbound replication for branch office DCs. From a TCO perspective, this reduces the number of connection objects you need to create and manage. It may also reduce the required number of hub site DCs.

RODCs also provide load-balancing improvements that help automatically distribute outbound connection objects evenly across hub-site bridgehead servers. With previous versions of Windows, this required routine manual intervention. Now, when the knowledge consistency checker (KCC) on an RODC detects that a bridgehead server is added or removed in a hub site, it determines whether to switch replication partners to a new bridgehead. It does this by running an algorithm and probabilistic load-balancing. If RODCs are added in branch office locations, the KCC will also balance inbound connections across existing hub-site bridgehead servers.

Credential Caching

An RODC’s PRP determines whether accounts are cacheable on that particular RODC. By default, the “allow” list in the PRP specifies that you can’t cache any account passwords. Also, it explicitly denies certain accounts from being cached. This takes precedence over manually configured “allow” configurations. As mentioned earlier, you may need to configure the PRPs on each RODC to allow the passwords for accounts to be cacheable.

Take this step cautiously, as PRP modifications have both security and service availability impacts. For example, the default scenario of no accounts cacheable results in high security, but no offline access in the event network connectivity to a writable DC becomes unavailable. Conversely, when a large percentage of accounts are cacheable (e.g., the domain users group), security is much lower if the RODC is compromised, but there’s a higher level of service availability for the cacheable accounts. Due to the unique business and technical requirements across various infrastructure environments, PRP designs will differ from organization to organization.

Once you've established a PRP model, you'll have to configure the PRP on each RODC so you may cache the appropriate accounts. As a best practice, configure the PRPs with explicit allows and don't modify the default deny list. The deny list is critical, because it prevents critical account credentials (such as AD DS service administrators) from ever being cached on RODCs.

Another key aspect of PRP design is determining whether cacheable accounts will be pre-populated with passwords. By default, the credentials of cacheable accounts aren't actually cached until after the initial logon to an RODC when the authentication request is forwarded to a Windows Server 2008 or Windows Server 2008 R2 writable DC and the credentials are replicated to the RODC. This means that if network connectivity to a writable DC becomes unavailable before cacheable accounts are authenticated against an RODC, successful logon will fail even though the accounts have been configured as cacheable.

To address this issue, you can manually pre-populate the password cache as soon as the PRP is configured and accounts are marked cacheable. This operation also requires network connectivity between a Windows Server 2008 or Windows Server 2008 R2 writable DC and the RODC. You can do this ahead of time during the deployment process, well before cacheable users log in for the first time.

You can use this fundamental architectural design guidance as a foundation for your RODC planning. By covering key design considerations, this article provides an effective starting point for designing a detailed and comprehensive RODC solution. This is not a simple process and requires substantial time to reconcile new features and design considerations against your organization's unique environment and requirements.

Paul Yu *(Paul.Yu@microsoft.com) is a senior consultant within Microsoft Consulting Services and has worked for Microsoft for ten years, providing enterprise infrastructure solutions to commercial corporations and public sector organizations.*

Related Content

- BranchCache and DirectAccess: Improving the Branch Office Experience (December 2009)
- Export, Compare and Synchronize Active Directory Schemas (April 2009)
- Understanding Proxy Authentication in AD LDS (December 2008)

Windows Server 2008 R2 Server Core

Securing the Branch Office Connection

Paul Yu

Few environments are more challenging to support than the remote office. Establishing and maintaining secure connections presents a host of technical and procedural challenges. The same can be said of any datacenter to which a remote office is connected.

Remote offices are often spread out across a wide area, there are typically quite a few of them, each has a relatively small user population, they connect to datacenter sites by slow network links, and there's rarely a skilled IT manager on site. Each factor by itself would pose a challenge. When you combine them, you have a recipe for an IT headache.

Although most branch office environments typically exhibit a similar set of characteristics (like those listed here), there are also organization-specific requirements that determine how the environment operates. The importance of security or the degree to which technology resources are centralized may differ from one remote office framework to another. Regardless of how any branch office operates, Windows Server 2008 R2 presents new opportunities to deploy updated technology that makes fundamental changes to how you can deploy and leverage domain controllers (DCs) in remote environments.

A Secure Solution

Securing branch office deployments is a common scenario for Active Directory Domain Services (ADDS). The unique characteristics and constraints of remote environments are well-suited for ADDS. Windows Server 2008 R2 has a considerable set of new features and updated approaches for deploying ADDS. Let's explore how to deploy Windows Server 2008 R2 in this type of environment in a secure manner. We'll cover typical branch office characteristics, key architectural design elements and recommended deployment options specific to deploying Windows Server 2008 R2.

The most noteworthy aspect of Windows Server 2008 R2 involves read-only domain controllers (RODCs). Generally speaking, RODCs are deployed in locations that require DC services, but lack the appropriate physical security measures. Because of the improved security and enhanced management features, replacing existing DCs with RODCs may exceed existing ADDS-related requirements in many branch office environments. For locations where there are currently no DCs, RODCs represent a substantial opportunity to introduce ADDS into the environment.

Another important factor here is the updated version of Server Core, an installation option in Windows Server 2008 R2 that provides a minimal environment for running specific, supported server roles such as the DC role. Selecting Server Core only installs the subset of binary files required by the installed server role, which in this case would be the RODC. This minimal installation results in reduced attack surface, reduced maintenance and easier management. It also helps the RODC server run on fewer resources.

Because many of these factors can help alleviate the constraints commonly associated with branch office environments, Windows Server 2008 R2 Server Core and RODCs are a compelling option for all remote environments. The following deployment considerations go into detail on the deployment and configuration options specific to Server Core and RODC components.

This exact configuration is a fairly common setup, although it may not meet every organization's branch office requirements. or instance, most major aspects of this configuration are already highlighted in the current best practices guide for the Windows Server 2008 Security Compliance Management Toolkit.

Here's a run-down of the critical deployment considerations associated with establishing a secure Windows Server 2008 R2 Server Core RODC solution for branch office environments. This covers solution assumptions, important design elements, infrastructure prerequisites and other detailed deployment options.

Preliminary Design Planning

As with any DC deployment, you have to make a number of critical design decisions prior to deployment. Some of these decisions include assessing hardware requirements, deciding the software upgrade strategy, determining the RODC server build and identifying the DC upgrade order. These decisions will dictate the overall deployment process and available configurations.

From a hardware assessment perspective, you must determine if your existing DC hardware complies with the recommended requirements. There are well-documented specifications, so these will probably not pose an issue for most organizations. With regard to supported software upgrade paths, there are a number of options that vary across versions, editions and different installation options of Windows.

Server Core requires a clean installation of branch office DCs. There is no way to upgrade to this installation option. With regard to installed server roles, for security purposes, it's generally recommended that all server roles and services not required for the RODC to function be eliminated from the server build. This RODC solution stipulates that Windows Server 2008 R2 Server Core RODCs host no additional services or server roles other than that of global catalog and DNS server, which is an increasingly common approach among enterprise organizations.

The order in which you deploy your DCs is another critical aspect of the process. The recommended order involves first installing ADDS in datacenters on pristinely built Windows Server 2008 R2 member servers for each domain, starting with the forest root, and then transferring all applicable operations master roles for each domain to these DCs. Continue deploying datacenter locations and fully decommissioning all legacy DCs in these sites. This helps stabilize ADDS in a large, well-administered location. It also contributes to simplifying the RODC deployment process itself. Once you've replaced the datacenter DCs, you can start on the branch office DCs.

Windows Server 2008 R2 Forest and Domain Preparation

Before deploying a single Windows Server 2008 R2 DC to the existing environment, you must prepare the Active Directory forest and domains by running Adprep.exe. First, update the forest schema on the DC hosting the schema operations master role with the **adprep /forestprep** command. At this point, you can update the forest to allow RODC installation with the **adprep /rodcprep** command. To prepare each child domain, you must run the **adprep /domainprep /gpprep** command on the DC hosting the infrastructure master role.

Last, you must deploy at least one writable DC running Windows Server 2008 R2 in the same domain where the RODCs will reside. As a side note, for environments where you ran the Windows Server 2008 version of the Adprep.exe commands, upgrading to Windows Server 2008 R2 still requires that you re-run all commands with the R2 version, with the exception of **adprep /rodcprep**, which makes no changes from the Windows Server 2008 version.

RODC Placement

From an architectural design perspective, RODC placement considerations have changed with the introduction of the Password Replication Policy (PRP). For example, RODCs must be able to establish domain partition replication with a Windows Server 2008 R2 writable DC. Because most branch office environments subscribe to hub-and-spoke network topologies, RODC ADDS sites are most likely to be separated by a single, lowest cost site link to the datacenter site, where Windows Server 2008 R2 writable DCs are located.

For instances where this might not apply, you'll be able to deploy additional writable DCs in intermediate sites, deploy new site link bridges, or create new site links to control how you create replication connections. You also have to ensure that other DCs aren't placed in the same ADDS site as the RODC. This condition is expected to be a non-issue for most branch office environments, which typically host a minimal number of servers. For locations hosting multiple DCs, deploying RODCs may simply not be an acceptable solution.

Credential Caching

Perhaps the most important security element here is the credential-caching model. This is a critical design component you'll have to carefully establish before you start branch office deployment. For many environments, the "few accounts cacheable" model will probably be the most common and appropriate model.

This approach, where only local accounts in an RODC site are configured as cacheable, provides suitable conditions with respect to principle of least privilege and service availability. One disadvantage to this approach is that it results in increased administrative responsibilities, because each RODC's PRP will be unique and require operational provisioning and de-provisioning of accounts as necessary.

How you deal with traveling users is another common design issue expected in many branch office environments. Often branch office environments include users and resources that may require their accounts to be cached on certain RODCs for service availability purposes. Ideally, these accounts would be provisioned in advance, similar to newly hired users. However, due to the random and unpredictable nature of travel behavior, this option is often not possible, especially for a large number of resources that travel across many RODC locations.

To address this issue, you could add additional accounts to the appropriate RODC PRPs. For extreme cases where a group of accounts require allowed access on all RODC PRPs, you can also leverage the default group "Allowed Read-Only Domain Controller Password Replication Group." However, this group should be used carefully, as membership in this security group makes all members cacheable on all RODCs.

Another consideration involves when cacheable accounts are actually cached. By default, this wouldn't occur until after initial logon to an RODC when the authentication request is forwarded to a Windows Server 2008 R2 writable DC and the credentials replicated to the RODC. Because branch office environments hosting existing DCs most likely possess preexisting requirements regarding service availability, the option to pre-populate credentials on RODCs may be critical.

This might be particularly important during the initial deployment of an RODC when all accounts in the RODC site have yet to cache their credentials. As soon as the PRP is configured and accounts are marked cacheable, you can use pre-populated passwords on an RODC. However, it's important to note that using the two traditional means for pre-populating passwords has some limitations. Currently, using the Active Directory Users and Computers console or the repadmin command does not allow for the usage of security groups.

Because pre-populating passwords one account at a time or in small batches based on organizational units may not be practical, you can use security groups in a scripted manner. For instance, in order to utilize the same security group that authorizes credential caching on a particular RODC, the following may be used:

```
For /F %%a in ('"dsquery group dc=corp,dc=contoso,dc=com -name
<Groupname>| dsget group -members"') do (Repadmin /rodcpwdrepl
<RODCname> <RWDCname> %%a)
```

RODC Staged Installation

Of the two DC installation methods provided by Windows Server 2008 R2, the staged installation option is preferable to direct installation. The direct alternative method equates to the traditional process available in previous versions of Windows. Staged installation uses Administrator Role Separation (ARS), a feature in Windows Server 2008 R2 that delegates non-service administrators the ability to install and administer RODC servers without granting Active Directory rights.

From a security perspective, staged installation removes the requirement to use highly elevated credentials in branch office locations that may not be secure. For this reason, arguments recommending DCs be staged in the data center to support remote offices may no longer apply. Staged installation separates the RODC installation process into two stages.

The first stage requires that an ADDS service administrator pre-create a computer account for the RODC and provide configurations such as computer name, the appropriate ADDS site, which server roles to install, the PRP configuration and ARS delegation. As best practice, the delegated administrator should be represented by a security group and each member should have his credentials cached on the RODC. The second stage involves the delegated RODC server administrator using his non-ADDS service administrator credentials to join a workgroup server to the pre-created RODC account and completing the RODC promotion process.

RODC Promotion Source

For the RODC promotion source, this configuration uses the Install from Media (IFM) installation option in conjunction with staged installation. This option significantly reduces the amount of data replicated to the RODC during the installation of ADDS. Using Ntdsutil.exe on a Windows Server 2008 R2 writable DC, there are four types of installation media available. Of these four, only two are pertinent here—RODC and RODC with SYSVOL.

The RODC media is similar to the Full installation media, but does not contain cached secrets like passwords. This is an important feature from a branch office security perspective. The only requirement to produce RODC media is that you must have a Windows Server 2008 R2 writable DC installed. However, the second installation media, RODC with SYSVOL, requires much more from an infrastructure perspective. Although Ntdsutil.exe will create the RODC with SYSVOL media, using that media during installation requires that Distributed File System Replication (DFS-R) for SYSVOL replication, which requires a domain functional level of Windows Server 2008 R2.

Given most organizations with branch office environments most likely will not meet this condition, that option to use the media will most likely be unavailable until initial deployment is finished. However, once this milestone is achieved, migrating to DFS-R and using both the RODC and RODC with SYSVOL installation media will significantly minimize the directory replication. It will also make installing future branch office DCs much more efficient.

Running DCPROMO

The Active Directory Domain Controller Installation Wizard will be unavailable as you deploy this configuration because it uses RODCs running Windows Server 2008 R2 Server Core. You can't use this during the actual RODC promotion. Therefore, in addition to staged installation with IFM, an unattend file with Dcpromo.exe will install the DC role. From a security and manageability standpoint, this aspect of the solution promotes secure and consistent DC build practices, which helps sustain ADDS security and configuration across the branch office environment.

In addition, automated, predictable, and repeatable build practices might minimize the possibility of unauthorized software, services, and configurations being introduced into the build process through manual intervention. The following is an example of the dcpromo command and a simple example answer file:

```
DCPROMO /unattend:c:\unattend.txt
    [DCINSTALL]
    ReplicaDomainDNSName=corp.contoso.com
    UserDomain=corp
    UserName=corp\<delegated RODC security
    group>
    Password=*
    ReplicationSourcePath=C: \IFM
    Safemodeadminpassword=<password>
```

It's important to note that if any manual PRP configurations are included in the answer file and not during the RODC account pre-creation section of staged installation, you must explicitly add all default PRP values. Manually adding explicit PRP configurations in the answer file essentially replaces the default PRP configuration with whatever configurations are specified in the answer file.

Replication

Because Windows Server 2008 R2 RODCs provide unidirectional replication, replacing existing branch office DCs with RODCs reduces the performance load on datacenter bridgehead servers that normally process inbound replication for branch office DCs. For branch office environments, this is significant. It increases scalability and can reduce the overall number of servers required in the datacenter.

RODCs also provide automatic distribution of outbound connection objects evenly across hub site bridgehead servers, something in Windows Server 2003 that required an additional tool such as Adlb.exe. For this reason, it's recommended that you upgrade all datacenter DCs to Windows Server 2008 R2 before deploying any RODCs. This ensures that inbound replication connections are evenly load-balanced and prevents the need for alternative measures that address issues associated with datacenter bridgehead server overload during RODC deployment.

This detailed design and deployment guidance can help with secure branch office deployment of Windows Server 2008 R2 Server Core RODCs. By covering key aspects of branch office characteristics, important architectural design elements and recommended deployment options, you can use this as a foundation for future best practices with RODC branch office deployments.

Paul Yu *(Paul.Yu@microsoft.com) is a senior consultant within Microsoft Consulting Services and has worked for Microsoft for 10 years providing enterprise infrastructure solutions to commercial corporations and public sector organizations.*

Related Content

- PKI Enhancements in Windows 7 and Windows Server 2008 R2 (May 2009)
- Inside SharePoint: Security and Compliance with AD RMS (April 2009)
- Geek of All Trades: Server Core in Windows Server 20008 (February 2009)

Desktop Virtualization

Care and Feeding of Virtual Environments

Wes Miller

Virtualization has evolved from an anomaly that required explanation to a viable technology most of us can't live without. Perhaps you're using it for quality assurance testing, development, Web design or training. Maybe you're part of the vanguard—setting the trend by deploying a virtual infrastructure, or even one of the masses using "cloud" virtualization from Amazon.com, Rackspace Inc. or another cloud vendor.

No matter how you're using virtualization, if you've used it for any length of time, you're no doubt realizing that it comes with its own set of challenges—just as maintaining physical hardware has its own dilemmas. Many issues are different; others are similar.

Handling Hypervisors

You've probably heard the word "hypervisor" bandied around for a while. It has become the cool term in virtualization. Hypervisors aren't new, however. We've been using them as long as we've been using virtual machines (VMs). In fact, IBM coined the term hypervisor in the 1970s.

The hypervisor is the software that presents the guests running "virtually" on a system with a set of virtualized hardware. It abstracts the physical hardware for the guest OSes. The confusion comes about with the big push to "type 1 hypervisors" running on the x86 platform over the last several years, including Microsoft Hyper-V and VMware ESX Server. The hypervisor most people use—especially for client systems—is referred to as a "type 2 hypervisor." What's the difference?

1. A type 1 hypervisor runs directly on the host hardware and does not require a "host OS." Microsoft Hyper-V and VMware ESX Server are common examples of a type 1 hypervisor.
2. A type 2 hypervisor requires a host OS to function. Generally, a type 2 hypervisor runs principally as a user-mode application on its host OS. Microsoft Virtual PC and VMware Workstation are common examples of a type 2 hypervisor.

More often than not, you would want to use a type 1 hypervisor for any “always-on” workload, such as a virtualized SQL or file server. At a minimum, it will use fewer resources than a type 2. Depending on the host, however, it may require a user logon in order to start, which isn’t a good option for a mission-critical system. A type 2 hypervisor, on the other hand, makes more sense for “on-demand” VMs. This type of role includes VMs for testing, application compatibility or secure access.

What Does Virtualization Save?

The obvious answer is that virtualization saves money on hardware, but it’s not quite that simple. Sure, if you have two server systems in rack-mountable 1U form factors, and you take those two same workloads and load them on one 1U system, you’ve saved on up-front hardware costs—but there’s a trick to it. When you take those same two server systems, both operate happily on two individual 1U servers, where each one has dual-core CPUs, 2GB of RAM and a 160GB SATA hard disk.

Now, when you put both of those onto one server, with the same hardware configuration, you’ll have to split the resources down the middle—or will you? You’ll generally need more resources for a type 2 hypervisor.

Then take the CPU, RAM and HDD costs necessary into account when figuring out how to consolidate workloads from physical to virtual. Virtualized consolidation is often called “stacking systems vertically instead of horizontally,” because you’re removing dependence upon *n* physical systems from an OEM. In turn, you’re asking far more of one individual system than you might have been prior to virtualization. This creates a systems-management ricochet many organizations don’t take into account as they rush headlong into virtualization.

What Does Virtualization Cost?

Once upon a time, good virtualization software cost quite a bit of money. Over time, the market has heated up, and you can get many types of virtualization software for fairly low cost. Most of the critical enterprise features still cost money, however, including the host OS or hypervisor.

Depending on the workload you’re planning to run on a VM, you may need to investigate failover. Guests get corrupted sometimes, and host hardware can fail. Virtualization doesn’t make hardware more reliable. It just changes the odds. For mission-critical systems, you *still* need to come up with a strategy for backing up the guest OS whether you’re backing up the VM container itself (which is absolutely recommended) or the file system contained within.

Even if you’re just virtualizing a bunch of guest OSes for testing or development on a type 2 hypervisor, you still need to allocate enough RAM to run one or more of those guests at a time (on top of the host OS). The most often overlooked issue in virtualization management is disk space consumption.

I've used virtualization for some time as a security test bed. Nothing beats running a potential exploit on a VM, seeing it work, and rolling back to an earlier version using your hypervisor's undo or snapshot functionality, only to retest it again. The real beauty of stacking these undo changes one on top of the other is that disk space can rapidly get out of control. It may end up far exceeding the actual size of the hard disk within the guest OS itself.

One of the VMs I use regularly has a 50GB hard disk image—I didn't realize how out of control it had gotten until I attempted to move it (it had six VMware snapshots), and the disk was well over 125GB.

Here are a few best practices to minimize the impact/cost of virtualization:

- If you're using a Windows client OS on a type 2 hypervisor with "undo" functionality, then by all means, *disable Windows System Restore*. Otherwise, you'll have disk growth every time you make a system change.
- If you perform step 1, be religious about demarcating when you do want to create an undo point.
- If you're doing security/exploit testing—do *not* rely on Windows to roll you back to an earlier point in time. Use your hypervisor's undo functionality, as it can't generally be tainted in the way restore points can be.
- Run guest OSes with the minimal amount of resources necessary.
- Ensure you've allocated enough RAM so client OSes aren't swapping RAM to disk all the time. This can slow down your host *and* all of your guests.
- Defragment your guests internally, and then defragment them externally (see the section on defragmentation later). Do both with some regularity.

VM Proliferation

As you can see, managing VMs can quickly become a problem. The ease of duplicating a VM can be a great benefit, but it can also create huge problems with managing and securing guests, keeping track of OS licenses with Windows (pre-Windows Vista, where new key management can actually be of benefit here), and making sure trade secrets don't escape from your grasp. It's a heck of a lot easier for a rogue employee to take a VM out via a USB flash drive or USB hard drive than it is for them to try and take an entire desktop system.

VM proliferation is much more of a problem among the highly technical (who understand the innards of virtualization). Generally, it's also more prevalent among client guests, not among virtualized server guests.

Systems Management

Entire companies have started to focus on helping to regain control over virtualized systems. Both Microsoft and VMware have consciously focused less on the value of virtualization itself and more on systems management. This is important because you aren't getting rid of systems—you're just virtualizing them.

Many systems management products can perform perfectly fine on VMs—but some newer functionality allows more intelligent management of virtualized systems, including waking and updating of guests that would otherwise fail to be updated. In the era of zero-day exploits, that's critical. The last thing you need is an infrequently used VM becoming the local botnet representative on your corporate network.

Your systems management approach needs to take into account that you have hosts and guests, ensuring they're updated accordingly, and that it knows the roles of each. The last thing you want is a poorly designed patch management solution updating your hypervisor, and tearing it down in the middle of the day for a reboot, taking four mission-critical guest servers with it.

You also need to be approaching recovery of these systems in the same way you would have historically. Just because a system is virtualized doesn't mean you can't lose it due to registry corruption or corruption of the entire VM—you can. Back up with the same fervor you apply to your physical systems today.

One extra consideration is whether your hypervisor does undo functionality. Bear this in mind when patch management comes into account. It's easy to update a guest on the Wednesday after Patch Tuesday, have it rolled back to Monday's undo point, only to get hit by a zero-day that it was theoretically "protected" against. This is a big problem, given undo technologies work by rolling back to an earlier point of the entire disk presentation from the hypervisor—meaning you will lose any Windows and application patches, as well as any antivirus signatures.

Security Software

Undo functionality aside, you need to be providing the same security protection to virtualized guests as you would to physical machines, and then some. When it comes to inbound threats, VMs are just as susceptible as physical machines—it makes no difference. But the big difference is that non-critical VMs (those that are not always on) often have latency for patching and AV updates. As a result, these can become a much bigger, untrackable target for zero-day exploits. This is all the more reason to ensure you're using a mature systems management solution that can take this into account and patch virtual systems as well.

Outbound threats are a different matter. VMs can be a doorway to the theft of intellectual property. This is critical to understand because VMs running on uncontrolled hosts can create a loophole for your data. First, if the virtual environment can be copied easily, that's a problem—especially if you're dealing with any compliance requirements that control access to data (as I discussed in an article back in 2008, technet.microsoft.com/magazine/2008.06.desktopfiles).

Second, as you might recall from my article on RMS and IRM (technet.microsoft.com/magazine/2008.11.desktopfiles), these controls rely upon the OS to prevent screen capture, printing and so on. However, those controls don't stretch out to the hypervisor—meaning that if RMS-protected content is displayed on a guest OS, the host OS can still print individual screenshots or create a video capture of the screen.

Though it isn't technically analog, this isn't entirely different from the "analog hole." I'm not aware of any way to protect DRM-controlled content from being exploited in this manner. Realistically, even if you could, then you're back at the problem of protecting from users with cameras or video cameras who can perform the same "exploit."

Disk Defragmentation

Disk defragmentation is a unique challenge on VMs, for several reasons:

1. You generally will have two levels of fragmentation—within the virtualized disk container itself (fragmentation the guests each see on their behalf)—what I refer to as "primary fragmentation," and fragmentation of the actual file containing the virtualized disk across the disks of the host OS, or "secondary fragmentation."
2. Virtualization products with disks that are the minimum size required and grow "on-demand" can lead to secondary fragmentation.
3. Undo functionality can rapidly lead not only to disk bloat, but massive secondary fragmentation—because as it consumes additional host OS disk space, each guest begins competing for available sectors.
4. With disks that grow on-demand, most do not have the capability to shrink when demand wanes. If you allocate 40GB, only use 10GB initially, but grow to require 35GB, the disk will not recover on its own—meaning you've got a large file that's much more likely to have secondary fragmentation.

The sheer size of virtual disks, the velocity with which they can change, shrink or grow, and the fact that they're susceptible to two types of fragmentation means you should treat them even more seriously than you would physical systems.

Here's one approach to protecting your files:

1. Minimize the use of any undo technology, as it will cause undue growth of the overall disk files, and can't readily be defragmented in the guest, though the host can defragment the files that comprise the virtual disk.
2. Use a good disk defragmentation product on your guests to begin with, and run it regularly.
3. If you're using on-demand disk expansion technology:
    a. Use the Sysinternals sdelete.exe utility as follows: **sdelete –c *drive_letter*** where drive_letter is the volume you want to zero-out. For example **sdelete –c C:** will zero-out all unused disk space after defragmentation.
    b. Use any virtual disk shrinking technology (if provided by your vendor) to reduce the virtual disk container to its minimum size.
4. Defragment the host OS's volumes containing the VMs.

Many people disregard disk defragmentation. The sheer volume of reader mail I've received from my article on disk defragmentation in 2007 (technet.microsoft.com/magazinebeta/2007.11.desktopfiles) proved it's often a misunderstood topic, but shouldn't be ignored—even with virtualized systems.

As virtualization continues to explode in importance and use, it can become too easy to get swept into the "consolidate" message without understanding the costs, and its inherent unintended complexities. This should help you discover some of the additional costs you need to consider when migrating to, or living with, virtualization.

Wes Miller *is the Director of Product Management at CoreTrace (CoreTrace.com) in Austin, Texas. Previously, he worked at Winternals Software and as a program manager at Microsoft. Miller can be reached at technet@getwired.com.*

Related Content

- Add Hyper-V to RDS for Inexpensive Virtual Desktops
- Manage Your Virtual Environments with VMM 2008
- Getting Started with Microsoft Application Virtualization

Toolbox

New Products for IT Professionals

Greg Steen

The opinions expressed in this column are solely those of the author and do not necessarily reflect those of Microsoft. All prices were confirmed at the time of writing and are subject to change.

SmartDeploy Enterprise

smartdeploy.com

The larger your organization, the greater variance there will be in the hardware and software configurations you'll need to support. As a systems administrator, you most likely have images of the different configurations for your environment to speed the deployment of new or updated systems across the organization. Maintaining an image for every different hardware system when the application stack installed on those systems is so similar seems like a waste of space and time for creating, updating and managing those system images.

There is a tool that can ease system image deployment, as well as reduce the number of images you need to have on hand: SmartDeploy Enterprise from Prowess Consulting. After getting the software installed, you can quickly create and deploy your first image. SmartDeploy distills the deployment process down into five steps:

1. Building a master installation
2. Capturing the image
3. Creating a "Platform Pack"
4. Creating the deployment boot media
5. Deploying the image

The master installation is a virtual machine (VM) you create with any of the major virtualization platforms, like Windows Virtual PC or VMWare. Once you have the system up-to-date with all the programs and patches you want to load on that VM, run sysprep.exe to prepare the system for image capture, and then shut it down.

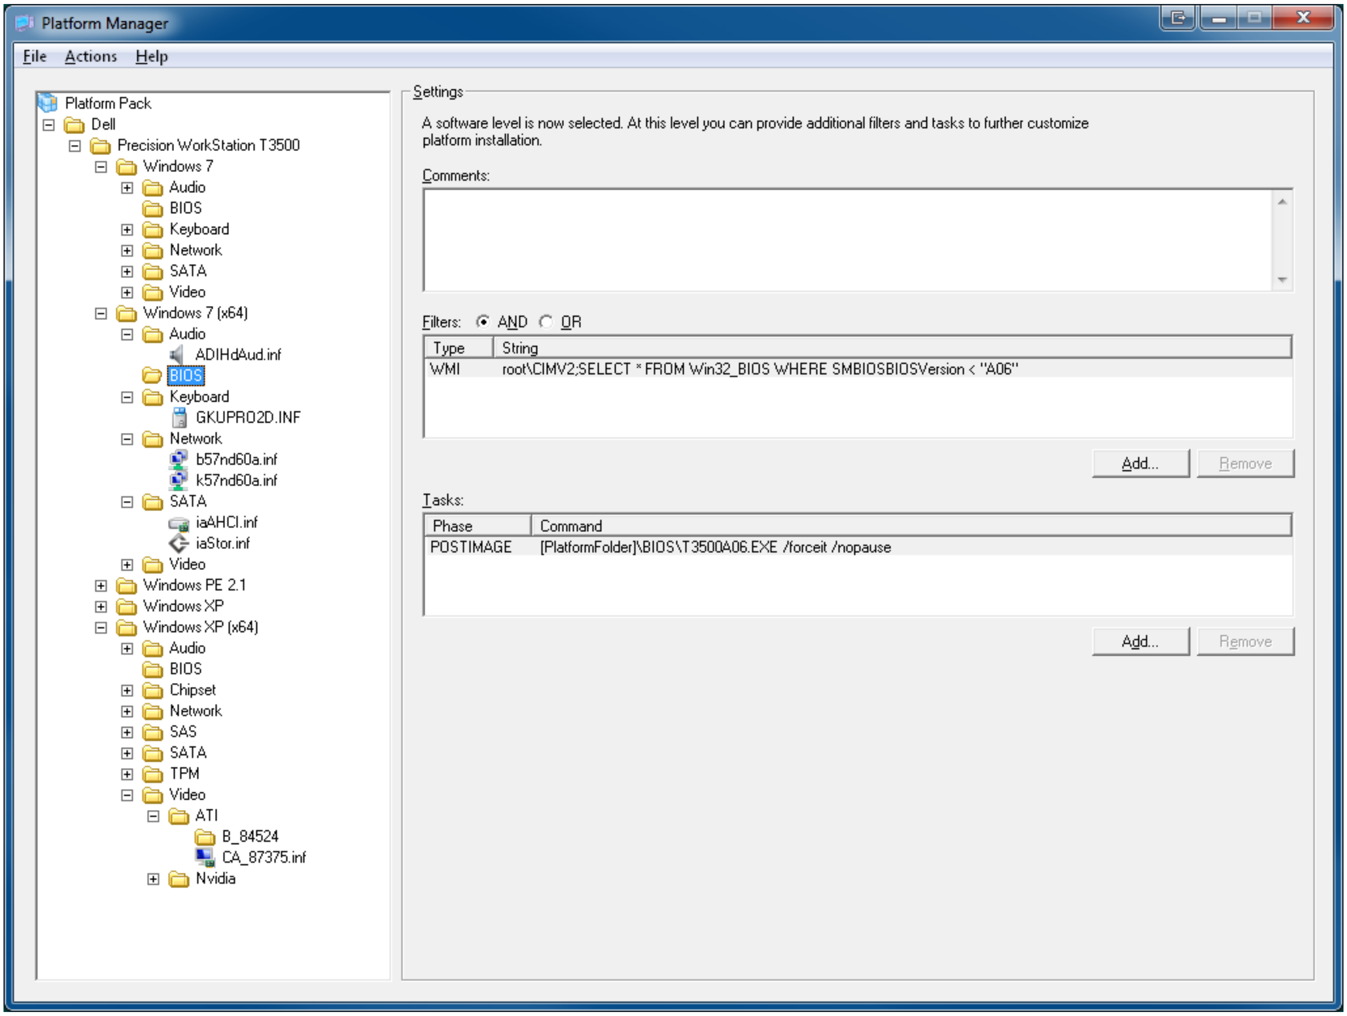

A software level is now selected. At this level you can provide additional filters and tasks to further customize platform installation.

Comments:

Filters: ● AND ○ OR

Type	String
WMI	root\CIMV2;SELECT * FROM Win32_BIOS WHERE SMBIOSBIOSVersion < "A06"

Tasks:

Phase	Command
POSTIMAGE	[PlatformFolder]\BIOS\T3500A06.EXE /forceit /nopause

SmartDeploy Enterprise

You won't need to install any machine-specific drivers or software as the Platform Pack you build (or download) will include all the necessary files. This is what lets you have one software image for a variety of target machines. Because you have a smaller set of images to manage and they're all VMs, it's much easier to keep them up-to-date with the latest software and security patches. This factor alone could greatly reduce maintenance times.

The next step is to capture an image of the VM using the SmartDeploy Capture Wizard. Point the wizard to your VM file, enter volume license information if applicable, and enter a destination for the resultant prepared Windows Image File (.wim). The SmartDeploy Capture Wizard also lets you create a Differencing Image, which helps with updates and add-ons to existing images because it creates a file of only the differences between the two VM references.

Next, use the Platform Manager to create or modify your Platform Pack files. The Platform Pack file contains all the machine-specific drivers and software necessary to support a target system type. You could also create a Platform Pack file that supports all of the different target machines in your environment. This is probably the most complex part of the whole process. Fortunately, Prowess Consulting has created a number of platform packs for the major hardware vendors. These packs are available on the company's Web site.

For manual packs, the documentation shows you how to create and lay it out. Then, use the SmartDeploy Media Wizard to create the boot media. You'll use this to start up the target machines to accept your newly created images. The wizard will create a disk containing the SmartDeploy Preinstallation Environment (SmartPE) you'll launch to deploy your image. Deploying the image, the last step, is also wizard-driven via the SmartPE environment's Deploy Wizard.

SmartDeploy Enterprise starts at $1,995 per operating technician (so you can deploy to or from as many machines as needed) and supports deployment of Windows systems from Windows 2000 through Windows 7/Windows Server 2008 R2 for both 32- and 64-bit versions. There are three support options available depending on your requirements, ranging from $299 to $1,300 per year. A free trial is also available for registered users via the product Web site.

If your deployment images are taking up way too much space, or your deployment process for new systems is taking way too long, or you're just looking to simplify the management of your whole deployment process, you might want to consider adding SmartDeploy to your toolbox.

Tweak7

tweaks.com/software/tweak7

As an IT Pro and power desktop user, you have undoubtedly tweaked and tuned your system to eek out as much performance and streamline your Windows experience to suit your needs. Under the hood, Windows 7 has a ton of features that you can tune and tweak in that regard, but finding them can be difficult. There's one tool out there that centralizes your control over some of these features to help get your system up to your specifications—the Tweak7 utility from Stardock.

This latest version of the company's Windows "tweaking" tool has many of the previous versions' features plus new features for Windows 7. The straightforward GUI splits the options into categories: Welcome, Resources, Security, Display, Power, Performance, Usability and Miscellaneous. To start tweaking, you simply pick a category and then pick a tab.

One nice thing about the interface is that there are numerous help icons and "hover-overs" that give you enough information about the tweak or item in question to determine whether you really want to proceed. The initial Welcome tab also gives you recommendations right away for tuning your system, including removing programs from your context menus and checking to make sure your display drivers are the most current available. Clicking a recommendation takes you to a detailed view on one of the other categories.

The Resources category is split into five tabbed sub-categories. What's Running shows you, much like the Task Manager, the running processes on your system and the executable's location. It also gives more detailed process information like kernel handles and GDI objects, descriptions when available, and one-click termination and Explorer to location.

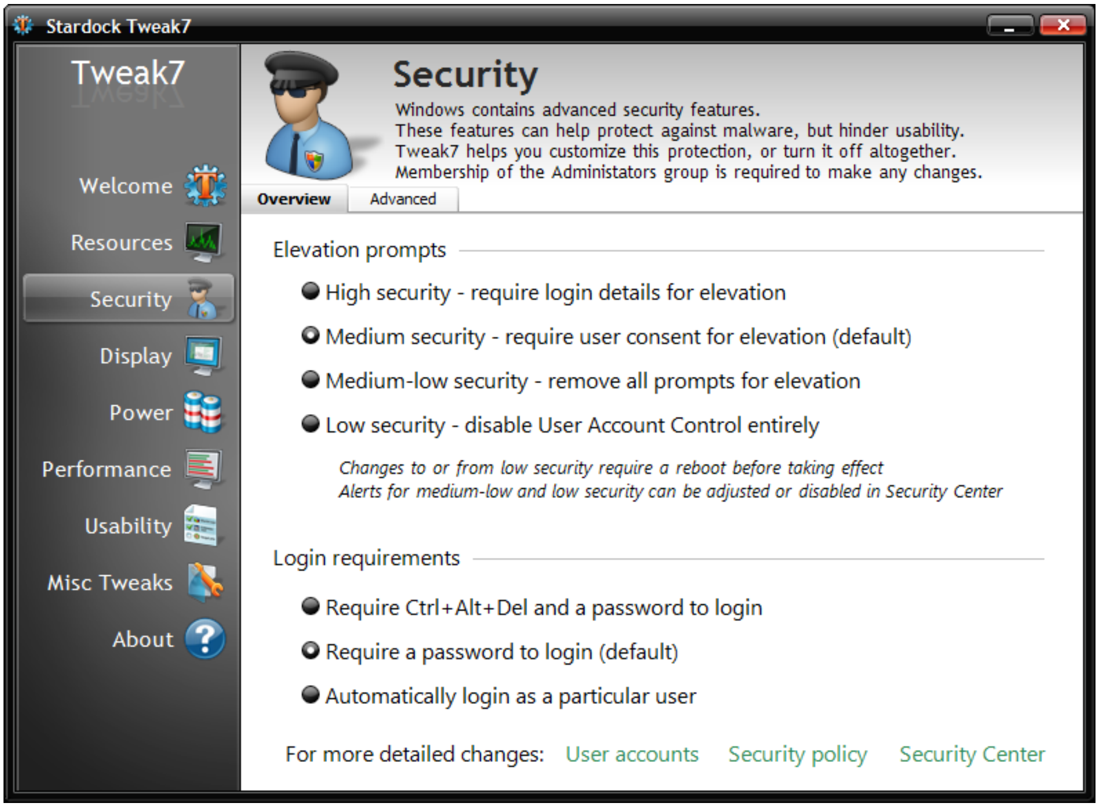

Security

Windows contains advanced security features.
These features can help protect against malware, but hinder usability.
Tweak7 helps you customize this protection, or turn it off altogether.
Membership of the Administators group is required to make any changes.

Elevation prompts

● High security - require login details for elevation
◉ Medium security - require user consent for elevation (default)
● Medium-low security - remove all prompts for elevation
● Low security - disable User Account Control entirely

Changes to or from low security require a reboot before taking effect
Alerts for medium-low and low security can be adjusted or disabled in Security Center

Login requirements

● Require Ctrl+Alt+Del and a password to login
◉ Require a password to login (default)
● Automatically login as a particular user

For more detailed changes: User accounts Security policy Security Center

Tweak7

The next tab, Start Up Programs, shows you all the programs that launch on system startup and lets you deactivate those you deem unworthy. Here, too, you'll get detailed program information and file locations, as well as uninstall options and deactivation recommendations.

The next three tabs—System Profile, Service Profile, and Services—let you customize what native services will run on your system based on your usage. Here you can choose a basic option, such as "I want to play games," or you can manually check on and off services at your discretion. The Security category lets you control elevation prompting and login requirements. Also, you can specifically disable UAC or non-administrator registry stores. The Display shows detailed information about your graphics cards, such as driver versions, Aero test results and Direct3D capabilities, in addition to links to tasks like changing font settings, color calibration or even disabling the Desktop Window Manager on startup.

The Power category gives you detailed battery information, such as rate of charge and "low battery" thresholds. You can enable or disable Hibernation or build an estimate of the machine's carbon footprint. The Performance category is split into CPU, Memory, Disk and Internet sub-categories. Besides showing tons of details about your hardware, you can also set Prefetch and SuperFetch settings, disable paging for the kernel and drivers, check for memory problems and tweak your Internet connection for better performance based on your typical usage.

The Usability category lets you change context menu items, disable shell extensions, move your Windows profile information to an alternate location, and disable confirmation notices, such as the warnings you get before emptying the recycle bin. Finally, the Miscellaneous category wraps it up with six sections:

- Environment lets you tweak your system and user environment variables.
- SMARTGuard shows you detailed information about your drives if they support the Self-Monitoring, Analysis and Reporting Technology.
- File Locks lets you drag a locked file onto the UI and then shows you what program or process has grabbed onto the file.
- Events shows you the detailed timing of your startup, shut down and resume cycles, broken down into the different areas such as "kernel startup" and "prefetch duration."
- Search Engine lets you pick your default Internet Explorer search box reference.
- System Tools gives you quick links to the Task Manager, Resource Monitor, System Restore, and the Reliability and Performance Monitor.

There's an almost full-featured (you can't move your Windows profile nor skip confirmation dialogs), free 30-day trial available for download. A single license for Tweak7 will run you $19.95 direct. If you want a central place to do your basic system tweaking on your Windows 7, you might want to take a look at adding Stardock's Tweak7 to your system.

Greg Steen *is a technology professional, entrepreneur and enthusiast. He's always on the hunt for new tools to help make operations, QA and development easier for the IT professional.*

Working Together

SQL Server 2008 R2 Reporting Services Integration in SharePoint 2010

Alan Le Marquand

SQL Server and SharePoint have always worked together well. When SharePoint Server 2010 and SQL Server 2008 R2 were released, there were some significant improvements to the integration between SharePoint and SQL Server 2008 R2 Reporting Services (SSRS). Here's a look at how to configure and use the latest enhancements.

Server Integration Architecture

The Reporting Services Add-in for SharePoint is what truly drives the integration between the two servers. You install the add-in, which is available as a free download from the Microsoft Download Center, on all SharePoint 2010 Web Front End (WFE) servers that require integration with a Report Server. **Figure 1** shows the architecture of the integration components.

On the SharePoint 2010 WFE, the add-in installs three components: the SSRS Proxy, a Report Viewer Web Part, and the application pages that allow you to view, store, and manage report server content on a SharePoint site or farm. The SSRS Proxy facilitates communication between the WFE and the Report Server. On the Central Administration Reporting Services pages within SharePoint, you configure the proxy with the Report Server you want to access, as well as the authentication method and credentials to access the server. For the integration to work, you must configure the Report Server to run in SharePoint Integrated mode.

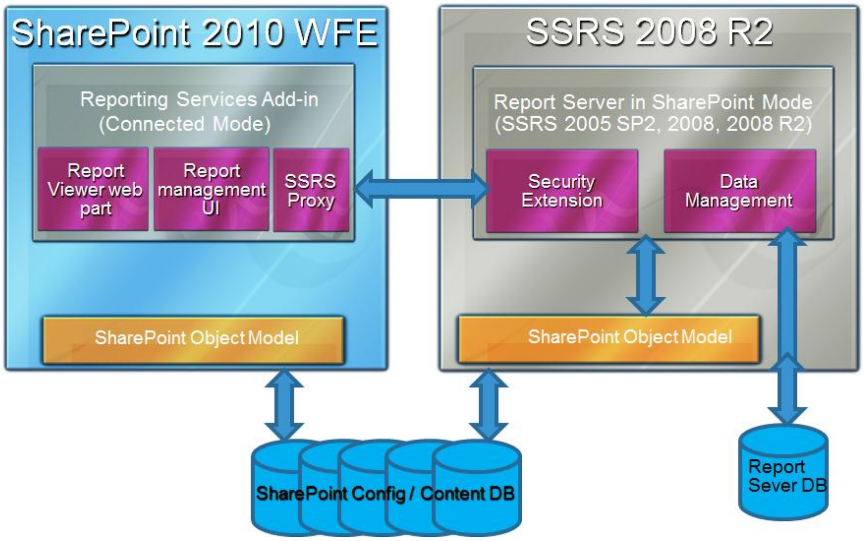

Figure 1 **Server Integration Architecture**

One item to note in **Figure 1** is the SharePoint Object Model component on the Report Server. For the Report Server to understand the reporting information stored in SharePoint and to be able to secure it, the Report Server has to interact with the configuration and content databases on your SharePoint site or farm. You can achieve this by installing a minimum copy of SharePoint on the Report Server and joining it to the farm.

The version of SharePoint you install on the Report Server must be the same as the version used throughout the farm. You need to do this only if you're running your Report Server on a separate machine. If you are running both SharePoint and Reporting Services on the same machine, you only need to install the add-in.

Configuring Integration

Overall, configuring integration has been simplified with SharePoint 2010 and SQL Server 2008 R2. The order in which you perform the configuration depends on what you've already installed. Even if you're starting from scratch or from an existing installation, the key is to have all the main components installed before you configure the SSRS Proxy in SharePoint. For best results when integrating SQL Server Reporting Service 2008 R2 with SharePoint 2010, the recommended order if starting from scratch is:

1. Run the SharePoint 2010 prerequisite installer—this will install the SSRS 2008 R2 Add-in for SharePoint.
2. Install and configure SharePoint 2010 in a farm configuration.
3. Repeat steps 1 and 2 on the Report Server machine if it is separate from the SharePoint WFE machine, and configure it to join the SharePoint farm created in step 2.
4. Install SQL Server Reporting Services in SharePoint Integrated mode.
5. Configure the SSRS Proxy via the Reporting Services Integration page and activate the Reporting Services feature.

6. If you don't see the Reporting Services content types in your site under the Document| New menu, you'll need to manually add them. I describe how to add the Report Server Content types later in this article under *Integration with Report Builder 3.0*.

In this scenario, I'd use the SQL Server for the SharePoint database, rather than the embedded edition that SharePoint defaults to. If you plan to install all the components on one machine, step 5 is redundant. Steps 1 and 2 can be combined within the SQL Server installation process.

If you have an existing SharePoint installation, you can download and install the add-in at any time. The add-in installation process adds the necessary pages to the SharePoint Central Administration as well as the new Report Server content types for existing SharePoint Libraries in sites using the Business Intelligence (BI) Center site template.

On the SharePoint side, you can configure integration on either SharePoint Server 2010 or SharePoint Foundation 2010. Both support the installation of the Reporting Services Add-in. If you install SharePoint and Reporting Services on different machines, you must install the same version of SharePoint on the Report Server. For example, you would not be able to install SharePoint Foundation 2010 on the Report Server if you were using SharePoint Server 2010 as your Web Front End.

The add-in installation is very simple; besides entering you name and company, no other configuration is required. If you're installing SharePoint for the first time, you install the add-in before you install SharePoint; this is done automatically when you run the SharePoint 2010 prerequisite.

Configuring the Report Server is straightforward. The key considerations are:

- The edition of SQL Server has to be Standard, Enterprise or higher.
- The Report Server database must be created for SharePoint Integrated mode.
- If you are using separate machines for SharePoint and Report Server, you'll need a minimal installation of SharePoint and this must be joined to the farm on the Report Server.

A Report Server is implemented as a single Windows service that runs under a built-in account or a local or domain Windows user account. In SharePoint Integrated mode, the Report Server service account is provisioned appropriately to access to the SharePoint configuration and content database as well as SharePoint object model resources. This happens when configuring the Reporting Services integration with SharePoint via the Reporting Services Integration page.

When the authentication mode is "Windows Integrated," the Windows user logged into SharePoint will be impersonated when connecting from the WFE to the Report Server. When the authentication mode is a trusted account, the SharePoint user context of the user logged into SharePoint is passed on to Report Server in the form of the SharePoint user token. The SharePoint WFE's application pool account is used to make the connection from the WFE to Report Server. You'll find a summary of the Service Account Configuration in the TechNet article "Configuring Reporting Services for SharePoint 2010 Integration."

If you've already installed Reporting Services using the default settings, the Reporting Services database will be in Native mode. To operate in SharePoint Integrated mode, you'll need to return to the Reporting Services Configuration tool and, from the Database Settings page, change the mode from Native to SharePoint Integrated.

You can change the Report Server mode from Native to SharePoint Integrated at any time; however, this does not convert the existing database. Each time you switch you must either create a new database or connect to an existing one.

Before configuring the Reporting Services Proxy options within SharePoint, there's one other configuration you need to make. You should ensure anonymous access has not been enabled on the Web application. While this will not stop you from configuring the Reporting Services Proxy settings, your users will get an error when they run reports. You can choose to use Windows or any claims-based authentication from the other authentication providers, and if you are configuring integration between a report server and a SharePoint farm, each SharePoint Web application in the farm can be configured to use different authentication providers.

The Add-in creates a new Reporting Services section within the SharePoint Central Administration General Application Settings page. On the Reporting Services Integration page you enter the Report Server URL and the authentication details and activation of the Reporting Services feature on all or selective site collections in the farm.

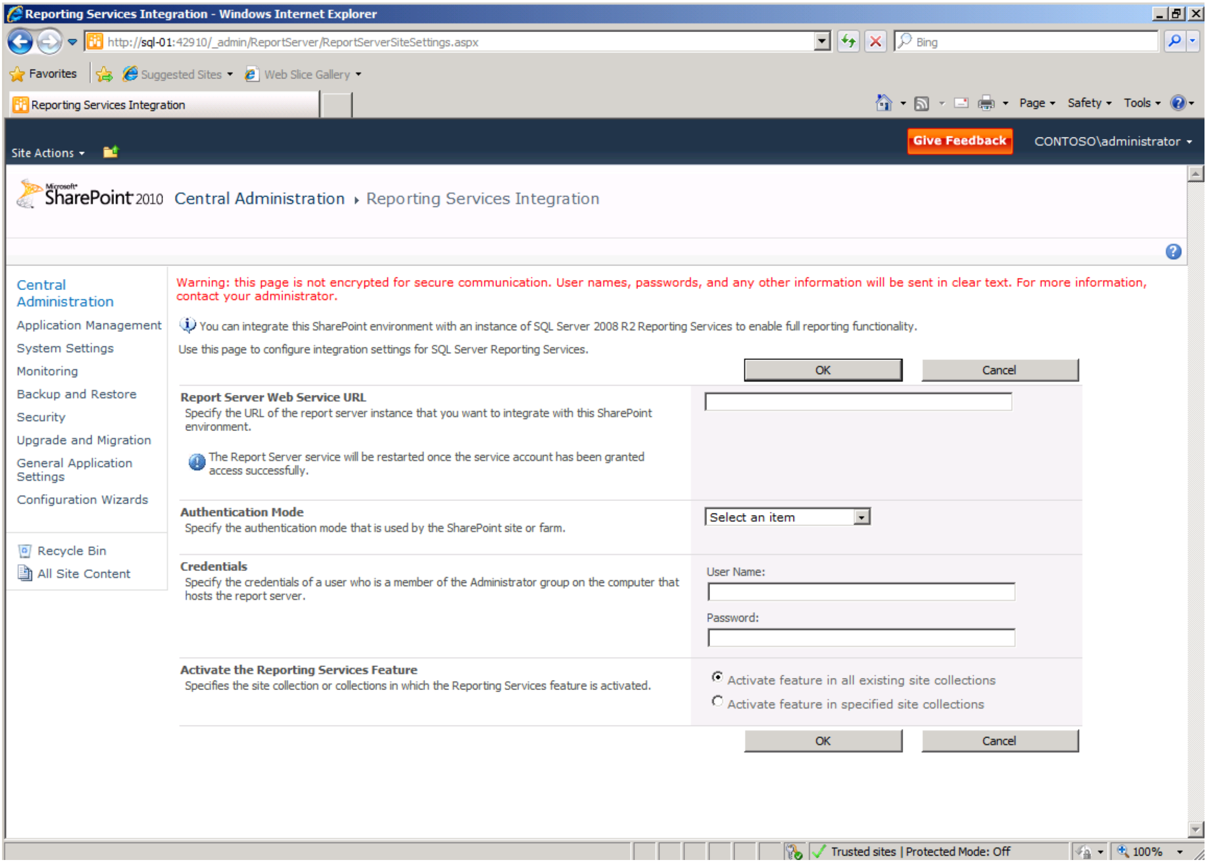

Warning: this page is not encrypted for secure communication. User names, passwords, and any other information will be sent in clear text. For more information, contact your administrator.

You can integrate this SharePoint environment with an instance of SQL Server 2008 R2 Reporting Services to enable full reporting functionality.

Use this page to configure integration settings for SQL Server Reporting Services.

Report Server Web Service URL
Specify the URL of the report server instance that you want to integrate with this SharePoint environment.

The Report Server service will be restarted once the service account has been granted access successfully.

Authentication Mode
Specify the authentication mode that is used by the SharePoint site or farm.

Select an item

Credentials
Specify the credentials of a user who is a member of the Administrator group on the computer that hosts the report server.

User Name:

Password:

Activate the Reporting Services Feature
Specifies the site collection or collections in which the Reporting Services feature is activated.

● Activate feature in all existing site collections
○ Activate feature in specified site collections

Figure 2 **Configuring the Reporting Services proxy.**

Once you complete the page shown in **Figure 2**, the integration configuration process is complete.

Integration with Report Builder 3.0

The main benefit of the integration between SharePoint and Reporting Services is that it allows users to create, modify and publish reports from within SharePoint. Reporting Services provides some predefined content types that are used to manage various files, including the shared Report data source (.rsds) files, the Report Builder model (.smdl), and the Report Builder report definition (.rdl) files. After you have configured integration to allow users to create and manage these new content types from the ribbon and the context menus, you need to enable the new content types on those libraries.

If you are using the BI Center site template, you do not have to do anything; the content types are automatically enabled with the template and for all sites created using this template. For all other sites and document libraries, you'll need to carry out a two-step configuration process. First, you need to enable Content Type Management within the libraries; by default it is off. Then you need to enable the content types for the library. To enable content type management for a document library, follow the procedure in the TechNet article "How to: Add Report Server Content Types to a Library (Reporting Services in SharePoint Integrated Mode)."

Once these new content types have been added to a library, three new options will appear from the New Document drop-down on the Documents tab. If you now select the Report Builder Report option, Report Builder 3.0 will be downloaded to the client and run. You can alter this behavior from the SharePoint Central Administration. The Reporting Services Server Defaults allows you to turn off this option as well as configure an alternate URL for Report Builder.

Using the Report Viewer Web Part on a SharePoint Site

The Report Viewer Web Part is a custom Web Part that is installed by the Reporting Services Add-in. You can use the Web Part to view, navigate, print and export reports on a report server. To add this Web Part to a page, you can use the steps in the TechNet article "How to: Add the Report Viewer Web Part to a Web Page (Reporting Services in SharePoint Integrated Mode)."

Each Report Viewer Web Part renders one report at a time based on the absolute URL to the report file (.rdl) specified in the Report property. The URL must be the fully qualified path to a report on the current SharePoint site or on a site within the same Web application or farm. The URL must resolve to a document library or to a folder within a document library that contains the report. The report URL must include the .rdl file extension. If the report depends on a model or shared data source files, you don't need to specify those files in the URL. The report contains references to the files it needs.

Claims Authentication and Reporting Services

One of the new features introduced with SharePoint Server 2010 is support for claims-based authentication. In claims-aware applications, clients present "claims" to the application. These claims are pieces of information about the user, such as user name, e-mail address, or manager name. This provides the application with more information than it would receive using Kerberos. Take, for example, a purchasing application: Two of the claims passed to the application could be the user's manager's e-mail address and the user's purchasing limit. In a non-claims-aware application, this information would have to be managed by the application.

In the SharePoint world, claims authentication solves the problem of sharing SharePoint sites across organizations. Using a product like Active Directory Federation Services (AD FS), two organizations with different authentication methods can set up claims that allow SharePoint to identify a user and assign the correct permissions.

Because this functionality is built into SharePoint 2010 products, Reporting Services can work with this authentication model. Reporting Services is not claims-aware; instead it communicates with SharePoint through a trusted account. The proxy service within the SQL Server 2008 R2 add-in uses the SharePoint object model to convert the claims token into a corresponding SharePoint user context in the form of a SharePoint user token that the Report Server can understand and use to validate against the SharePoint database. In a nutshell, the process works like this:

1. SharePoint performs the appropriate claims authentication and, using the SharePoint Secure Token Service, communicates the claims token to the Reporting Services proxy.
2. The Reporting Services proxy then uses the claims token to communicate with the SharePoint object model and generate a corresponding SharePoint user token that it forwards to the Report Server.
3. The Report Server uses the SharePoint user token against the local SharePoint object model to generate the correct SharePoint user context.
4. If the user has the required permission, Report Server sends the requested information back to SharePoint using the appropriate SharePoint user context as it would normally.

Native List Reporting

SQL Server 2008 R2 Reporting Services now supports SharePoint lists as a data source. This support allows you to retrieve list data from SharePoint Foundation 2010, SharePoint Server 2010, Windows SharePoint Services 3.0, and Office SharePoint Server 2007. The ability to access list data is not reliant on the add-in or running Report Server in Native or SharePoint Integrated mode. The functionality is built into Report Server. What changes in the different configurations is the method of access.

There are two methods by which SharePoint list data is accessed. One is via the lists.asmx web service and the other is via the SharePoint object model APIs. On any SharePoint installation, if you enter the URL http://<sharepoint_server_name>\lists.asmx, you'll get an XML list of all the lists on the SharePoint site that you're able to access. By using this method, Report Builder 3.0 is able to retrieve the lists. A Report Server configured in Native mode also uses this method.

The SharePoint object model API method can be used in two scenarios. One is where a Report Server is configured in SharePoint Integration mode and the list exists in the same SharePoint farm Reporting Services is integrated with, and this is all on the same machine; remember that in this scenario there is a copy of SharePoint running on the Report Server that gives it access to the API set. The other scenario is where you have SharePoint 2010 installed along with the add-in, but you have no Report Server. This is called *local mode* and is covered later in the section "Reporting Without Reporting Services."

To use data obtained from a SharePoint list within a report first requires you create a data source, then a dataset that uses that data source. In Report Builder 3.0, there is a new connection type on the Data Source properties page called Microsoft SharePoint List, as shown in **Figure 3**. Along with this option, you enter the URL of your SharePoint site—no need to add lists.asmx to the URL. The data source can also be configured with different credentials to use when accessing the SharePoint server.

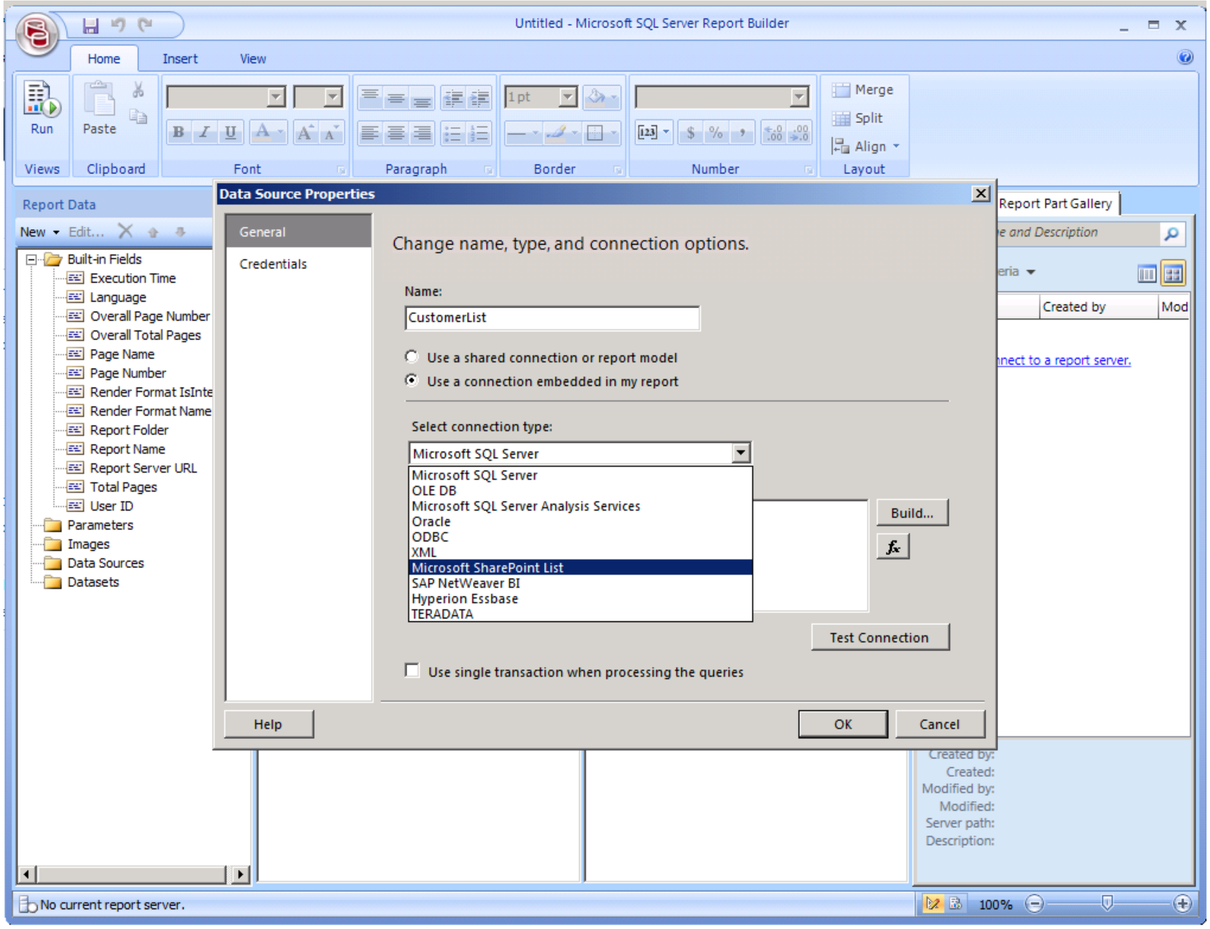

Data Source Properties

General
Credentials

Change name, type, and connection options.

Name: CustomerList

○ Use a shared connection or report model
● Use a connection embedded in my report

Select connection type: Microsoft SQL Server

- Microsoft SQL Server
- OLE DB
- Microsoft SQL Server Analysis Services
- Oracle
- ODBC
- XML
- Microsoft SharePoint List
- SAP NetWeaver BI
- Hyperion Essbase
- TERADATA

Build...
Test Connection

□ Use single transaction when processing the queries

Help
OK
Cancel

Figure 3 **SharePoint List Connection Type**

When you create a new dataset based on this data source, you will be given a list of all the SharePoint lists on the site you have access to. You can then drill down into a list and access the individual list items, create filters, create parameters and create reports just as if this were a SQL database table.

Alternate Access Mapping Support

Another integration enhancement is support for Alternate Access Mapping (AAM). AAM has been in SharePoint since the 2007 version, but Reporting Services did not support it. Now if you configure an alternate access mapping within SharePoint Central Administration, the Reporting Service Add-in will maintain the URL structure, as shown in the very simple report in **Figure 4**. Both http://sql-01 and http://www.contoso.com render the same report.

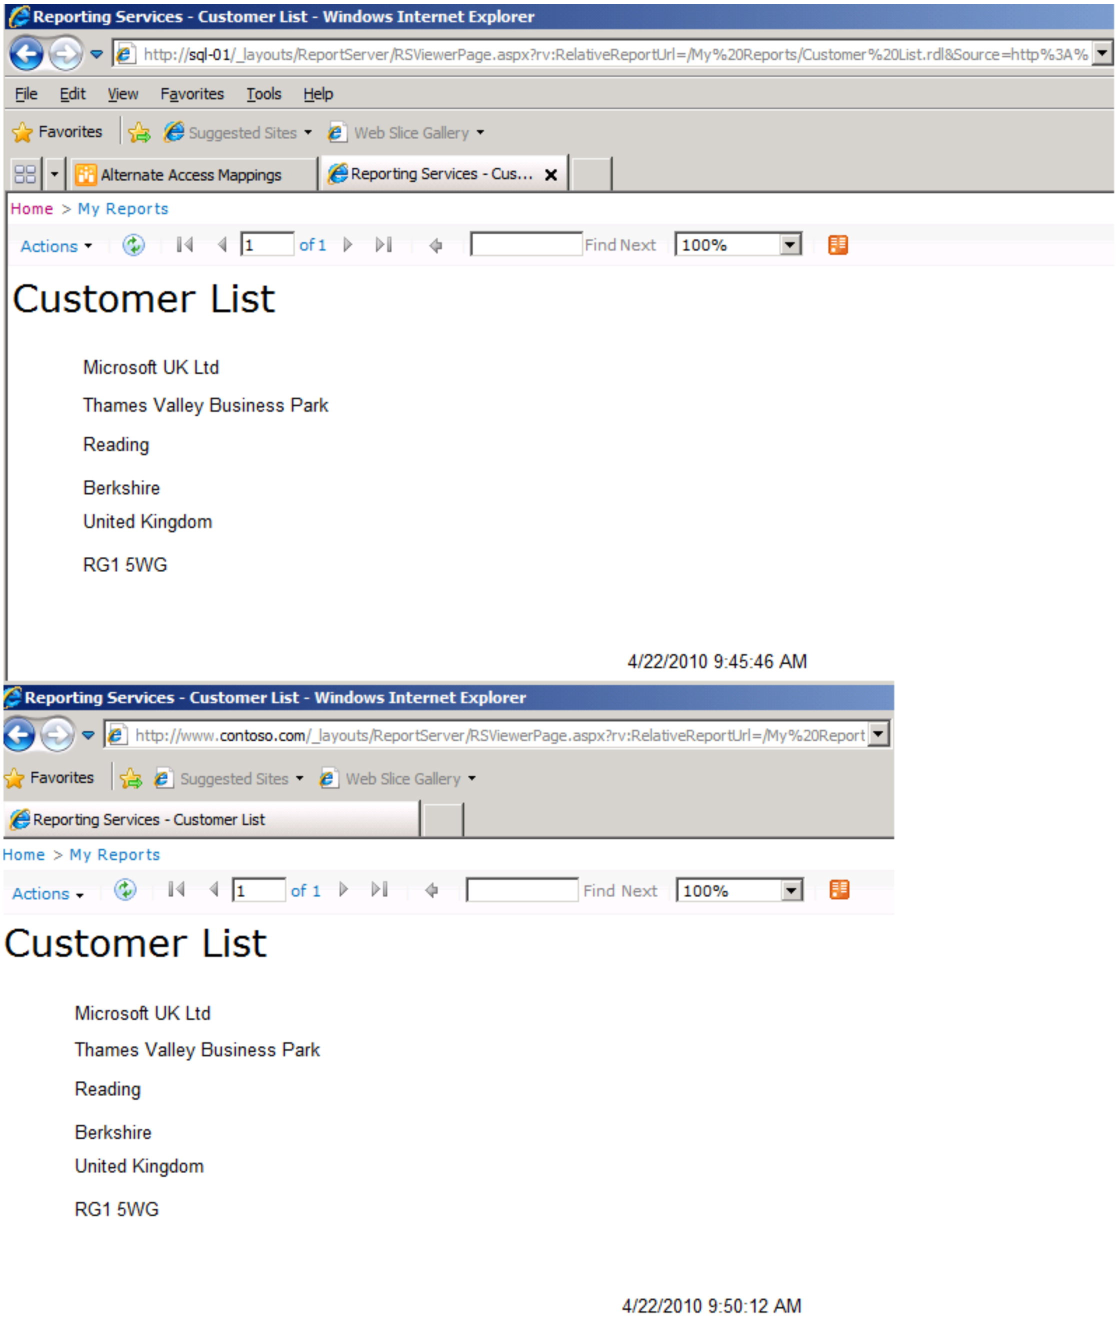
Reporting Services - Customer List - Windows Internet Explorer
File
Edit
View
Favorites
Tools
Help
Favorites
Suggested Sites
Web Slice Gallery
Alternate Access Mappings
Reporting Services - Cus...
Home > My Reports
Actions
of 1
Find Next
100%
Reporting Services - Customer List - Windows Internet Explorer
Favorites
Suggested Sites
Web Slice Gallery
Reporting Services - Customer List
Home > My Reports
Actions
of 1
Find Next
100%

http://sql-01/_layouts/ReportServer/RSViewerPage.aspx?rv:RelativeReportUrl=/My%20Reports/Customer%20List.rdl&Source=http%3A%

Customer List

Microsoft UK Ltd

Thames Valley Business Park

Reading

Berkshire

United Kingdom

RG1 5WG

4/22/2010 9:45:46 AM

http://www.contoso.com/_layouts/ReportServer/RSViewerPage.aspx?rv:RelativeReportUrl=/My%20Report

Customer List

Microsoft UK Ltd

Thames Valley Business Park

Reading

Berkshire

United Kingdom

RG1 5WG

4/22/2010 9:50:12 AM

Figure 4 **Alternate Access Mapping**

Reporting Without Reporting Services

So far, all the information in this article has pertained to what is called *connected mode*. In the previous versions of Reporting Services, this was the only mode available and meant that SharePoint had to be connected to a Reporting Services report server configured in SharePoint Integrated mode in order to render reports using the Report Viewer.

With the release of SQL Server 2008 R2, you can render reports without integrating your SharePoint site or farm with a Reporting Services report server. Instead, you can use the Report Viewer to directly render reports from SharePoint when the data extension supports *local mode* reporting. Out of the box, only the SharePoint List and the Microsoft Access 2010 reporting extension support this.

When you're in local mode, you can also render a report that has an embedded data source or a shared data source from an .rsds file. However, you can't manage the report or its associated data source as this is not supported in local mode.

Supported Combinations of the SharePoint Add-in and Report Server

With the release of SQL Server 2008 R2 and SharePoint Server 2010, there are now three versions of SQL, three versions of the add-in, and two versions of SharePoint. The integration components can work on any of these releases, but you have to mix and match the right versions. The table in **Figure 5** provides the supported combinations of products.

Figure 5 **Supported combinations of the SharePoint add-in and Report Server.**

Report Server	Add-in	SharePoint	Supported
SQL Server 2008 R2	SQL Server 2008 R2	SharePoint 2010 Products	Y
SQL Server 2008 R2	SQL Server 2008 SP2	SharePoint 2007 Products	Y
SQL Server 2008 SP1 CU8	SQL Server 2008 R2	SharePoint 2010 Products	Y
SQL Server 2008	SQL Server 2008	SharePoint 2010 Products	N
SQL Server 2008	SQL Server 2008	SharePoint 2007 Products	Y
SQL Server 2005 SP2	SQL Server 2005	SharePoint 2007 Products	Y

Alan Le Marquand *is an IT Pro Content Architect for Microsoft based in the United Kingdom. You can read more from Le Marquand on his blog Alan's World of IT*.

Related Content

- Reporting Services with SharePoint Integration
- SharePoint Security: The Fundamentals of Securing SharePoint Deployments
- Protect SharePoint Data

Windows PowerShell

Writing Cmdlets in Script

Don Jones

One of the cool new features in Windows PowerShell v2 is the ability to write greatly improved functions. These functions, written entirely in script, have the same capabilities as a “real” cmdlet written in C# or Visual Basic and compiled in Visual Studio. These *advanced functions* (they were originally called “script cmdlets” early in the v2 development cycle) help you write more flexible functions you can then use seamlessly alongside regular cmdlets.

It’s All in the Binding

The real difference between a mere function and a full cmdlet is that cmdlets support powerful parameter bindings. You can use positional parameters, named parameters, mandatory parameters, and even do basic parameter validation checks—all by simply describing the parameter to the shell. Here’s an example:

```
function Get-Inventory {
    [CmdletBinding()]
    param (
        [parameter(Mandatory=$true,ValueFromPipeline=$true)]
        [string[]]$computername,

        [parameter(Mandatory=$false)]
        [alias("PF")]
        [switch]$pingfirst,

        [parameter(Mandatory=$true,Position=0)]
        [AllowEmptyString()]
        [string]$class

    )
    PROCESS {
    }
}
```

In this statement, I’ve declared three parameters:

- -computername is a single string or an array of strings. It is mandatory, and it accepts string pipeline input—meaning if you pipe in a bunch of strings, they’ll be automatically dropped into the $computername variable.
- -pingfirst is not mandatory, but if you do use it, you should use the -PF alias. It will save a little typing. This is a switch parameter, meaning it does not accept a value. It’s either on or off.
- -class is also mandatory, but you don’t even have to type the -class parameter name. Just give it the appropriate value as the first-position value when you run the function. Although this is mandatory, it will accept an empty string.

There are many more attributes—and plenty of examples—in the online help. Run **help about_Functions_Advanced_Parameters** to see them all.

Accessing Common Parameters

The shell defines several common parameters shared by all cmdlets. One of them is -verbose, which is intended to tell a cmdlet to output more information than usual about what it's doing. The following function definition, however, will result in an error:

```
function Test-Something {
    [CmdletBinding()]
    param (
        [switch]$verbose
    )
    PROCESS {
    }
}
```

That's because you can't re-define one of the common parameters like -verbose. So how can you tell if your function was run with -verbose or not? Well, it turns out to be unnecessary. Windows PowerShell keeps track of it for you. You simply call **Write-Verbose**, and Windows PowerShell will ignore those calls if -verbose wasn't used:

```
function Test-Something {
    PROCESS {
        Write-Verbose "Starting cmdlet"
    }
}

test-something –verbose
```

Confirming Impact

Another pair of common parameters is -whatif and -confirm. Any cmdlet that makes some kind of change to the computer is supposed to recognize those. They give you the option of having the cmdlet display what it would normally do (-whatif), or have it individually confirm each action (-confirm). Together, these parameters are called ShouldProcess, and you can declare a function that supports them, as shown here:

```
function Delete-Things {
    [CmdletBinding(
        SupportsShouldProcess=$true,
        ConfirmImpact="Medium"
    )]
    PROCESS {
    }
}
```

This declaration enables both -whatif and -confirm as parameters for your function. It also specifies that your function has a "Medium" impact level on the OS. There are no strict guidelines for what "Medium" means—I'd imagine it's something less than the possibility for total disaster. The real trick is that the shell's $ConfirmPreference variable defaults to "High."

When a cmdlet’s impact is lower than $ConfirmPreference, then the cmdlet will run without confirmation unless -whatif or -confirm are specified.

If the cmdlet’s impact is equal to or higher than $ConfirmPreference, then every time you run the cmdlet, it will act as if you had specified -confirm, even if you forgot to do so. Therefore, if your function is about to do something *really dangerous,* specify a ConfirmImpact="High" so your cmdlet will always request confirmation. Your other choices are "None" and "Low".

None of the shell’s built-in help actually shows you how to ask for confirmation—and it isn’t automatic. The help refers you to the online MSDN help, which is intended for Microsoft .NET Framework developers and doesn’t reference the shell’s scripting language at all. So I’ll show you here:

```
function Delete-Things {
    [CmdletBinding(
        SupportsShouldProcess=$true,
        ConfirmImpact="High"
    )]
    Param ($param1)
    PROCESS {
        if ($pscmdlet.ShouldProcess($param1)) {
            Write "Deleting..."
        }
    }
}

Delete-Things "organizationalunit"
```

$pscmdlet is a built-in variable you can use within the PROCESS script block to access cmdlet-level functionality—including the ShouldProcess method. You pass in a description of what you’re about to modify, and the shell will take care of displaying the actual confirmation or “what if” message.

If ShouldProcess returns $True, then it lets you continue. If it returns $False, then you should not do whatever it is you were about to do. Once you know about the $pscmdlet variable, making sense of those MSDN developer docs is easier. They do accurately describe the different ways you should use ShouldProcess and its companions—like ShouldContinue.

Help! Help! Help!

Don’t forget that functions—even advanced ones—can contain their own built-in help in specially formatted comments, as I described in my March 2010 column. Typically, I list the comment-based help first, then the CmdletBinding statement, then my parameters, and then the BEGIN{}, PROCESS{}, and END{} scriptblocks. It’s *always* a good idea to include help within your functions—you never know who might benefit from it.

If you've written pipeline functions before (also called "filtering functions"), then you already know everything else you need to know in order to write a "script cmdlet." The PROCESS{} script block is where your code goes, and it will execute once for each object piped into your cmdlet. Everything else about these advanced functions is also just like their somewhat-simpler counterparts.

Windows PowerShell v2 Is Available Now

Although it shipped pre-installed with Windows Server 2008 R2 and Windows 7, Windows PowerShell v2—and its companion Management Framework components—is now available for Windows XP, Windows Server 2003, Windows Vista and Windows Server 2008. Visit support.microsoft.com/kb/968929 to get the download link for whatever OS you're using. This should be compatible with your v1 scripts, and all of my future columns will assume that you're using 2.0.

A Range of Audiences

The Windows PowerShell team rightly prides itself on making Windows PowerShell useful to a wide range of audiences with varying skill levels. Advanced functions are definitely an example of something that only an *advanced* audience will find useful.

If you're just getting started with the shell and have to remind yourself to run **Help**, then advanced functions are probably still off in the future. You can successfully use the shell without ever writing an advanced function. Once you start getting more advanced, and start writing re-usable components, you'll find that advanced functions are a great way to do so.

The blog post here is a great example of this: It starts with a simple command that accomplishes a valuable task—something any administrator might write. Then the author gradually expands his command's functionality into a function, then a filtering function, and finally an advanced function, showing how the shell can scale up as your needs and skills grow.

Don Jones *is a founder of Concentrated Technology, and answers questions about Windows PowerShell and other technologies at ConcentratedTech.com. Jones is also an author for Nexus.Realtimepublishers.com, which makes many of his books available as free electronic editions.*

Geek of All Trades

iSCSI Is the Perfect Fit for Small Environments

Greg Shields

Do you remember the good old days of SCSI? Back then, figuring out your SCSI connections, speeds, acronyms and interfaces practically required a secret decoder ring. Determining whether your server needed Fast SCSI or Ultra SCSI or Ultra2 Wide SCSI or any flavor in-between made working with SCSI disk drives a complicated task.

As a result, more than a few of us threw up our hands in frustration. In those days, we often saved the SCSI work for outside consultants, or found ourselves favoring slower disks that operated under frameworks we understood.

Things have changed a lot since then, thanks to greater levels of standardization. Today, we can find SCSI pretty much everywhere. Our server manufacturers now deliver their equipment with internally preconfigured SATA or secure attention sequence (SAS) drives. No more worrying about acronyms, connectors or secret decoder rings. Just add your disks and go.

However, this increasing standardization of direct-attached SCSI still doesn't get around the fact that its storage devices must be directly attached to a server. Using traditional direct-attached SCSI, there are no clean ways to connect multiple servers to centralized storage over your existing network.

This need for a networkable solution is why iSCSI exists. The "i" in iSCSI replaces SCSI's multitude of connections with everyday, run-of-the-mill Ethernet cabling. By consolidating SCSI's per-server storage into a single and shared device, and then connecting servers to storage through your existing network, your small environment can make better use of storage as you discretely provision it to servers and file shares as needed.

Replacing SCSI's physical connections with copper Ethernet makes managing the physical layer easier. However, correctly incorporating iSCSI into your small environment still requires some techniques and a few extra protocols that may not be immediately obvious. With low cost and easy management, iSCSI can be a perfect fit for the storage needs of the Jack-of-all-Trades IT professional. Read on to learn more so you can use it successfully.

Goodbye Physical Connections, Hello Logical

Understanding iSCSI's connection potential is best done by examining the options. Take a minute to call up a remote console on one of your favorite servers. Once there, open its iSCSI Initiator administrative tool. This console was given a much-needed facelift in Windows Server 2008 R2, most notably adding the Quick Connect dialog box you'll see in **Figure 1**. You can easily create a basic connection between this server and an exposed iSCSI LUN using the Quick Connect dialog box.

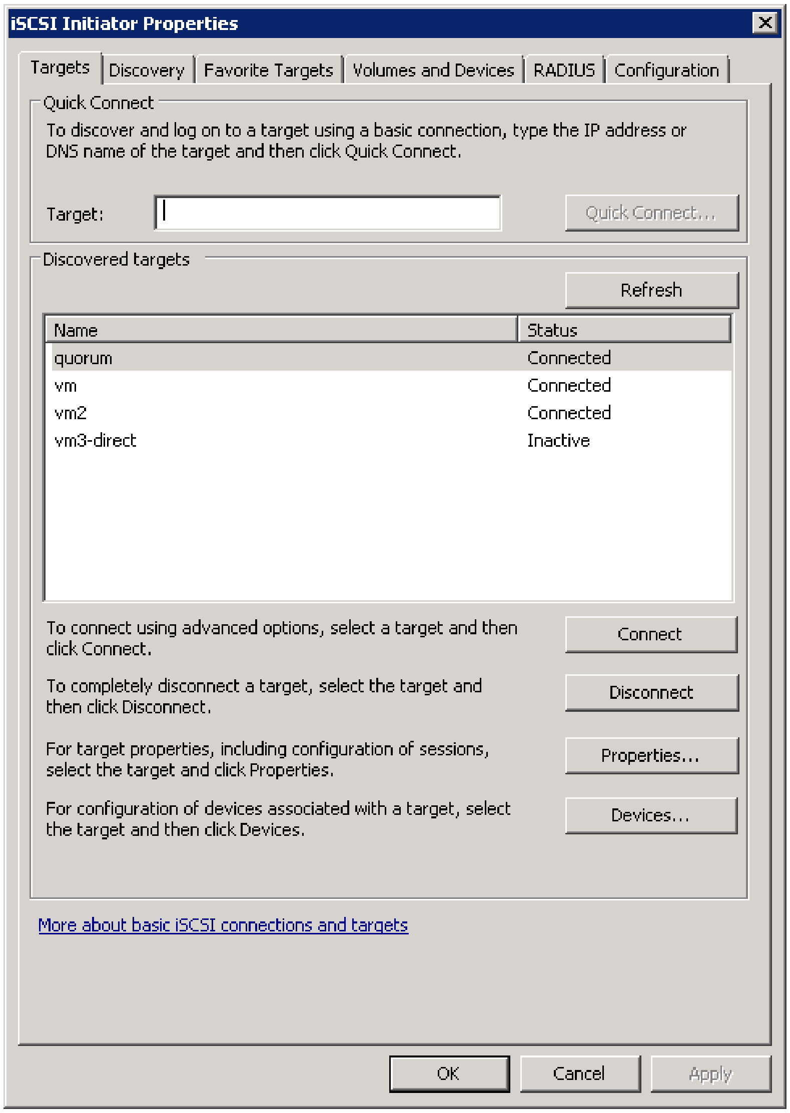

To discover and log on to a target using a basic connection, type the IP address or DNS name of the target and then click Quick Connect.

Target:

Name	Status
quorum	Connected
vm	Connected
vm2	Connected
vm3-direct	Inactive

To connect using advanced options, select a target and then click Connect.

To completely disconnect a target, select the target and then click Disconnect.

For target properties, including configuration of sessions, select the target and click Properties.

For configuration of devices associated with a target, select the target and then click Devices.

More about basic iSCSI connections and targets

Figure 1 **The iSCSI Initiator administrative tool.**

At this point, a few definitions might be helpful. As in the physical world, every iSCSI connection requires two devices. At your favorite server is what is called the iSCSI Initiator. The iSCSI Initiator represents the “client” that will be requesting services from a set of iSCSI disks.

At the other end is the iSCSI Target. The iSCSI Target has the disks to which you want to connect. It operates as the “server” that services requests from one or more client initiators. Once the connection between a target and an initiator is established, you can discover and initialize one or more LUNs as disk volumes on your server.

To use iSCSI, you'll obviously require a device that supports the iSCSI protocol. This may be an existing SAN on your network. It can also be a regular Windows server that runs an instance of iSCSI Target software.

Microsoft has its own iSCSI Software Target software. However, this software is intended for use atop Windows Storage Server. Third parties also distribute software that installs as a service to a regular Windows server. That service exposes the server's direct-attached disks to locations anywhere on your network.

Once installed, the next step always starts at the iSCSI Target. Prior to connecting any "client" computers to this device's exposed LUNs, you'll first need to create and expose those LUNs to the network. The details of this process will vary greatly and depend on your device and your software. Consult your iSCSI Target's manual for the details. At the very least, you'll have to carve out a quantity of disk space as a LUN, connecting that LUN to the proper network and network interfaces, and adding any security or authentication options.

Once you've completed this first step, creating a basic connection requires only a few clicks. First, enter the IP address or DNS name of the server or device that runs the iSCSI Target software into the Target field shown in **Figure 1**and click Quick Connect. If you've correctly created and exposed your LUNs to this server, you'll see them appear in the list of Discovered targets.

Figure 1 shows four discovered targets, three of which have been connected. Discovered targets always appear first in an Inactive state. This ensures that you can connect to them only when you're ready.

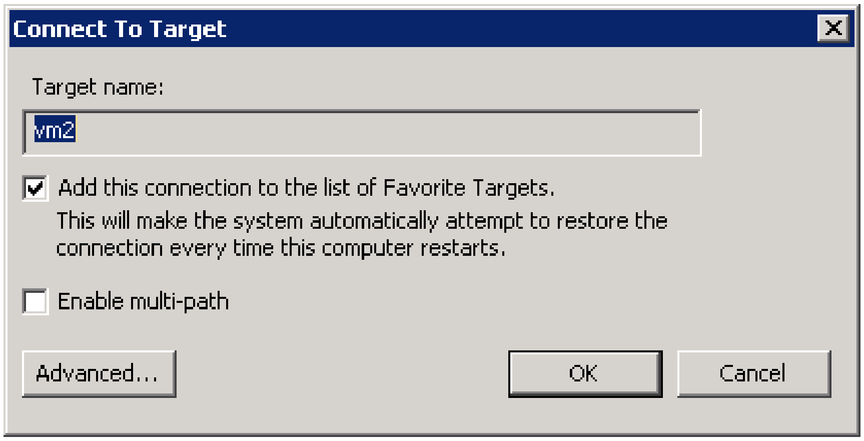

Connect To Target

Target name: vm2

☑ Add this connection to the list of Favorite Targets.
This will make the system automatically attempt to restore the connection every time this computer restarts.

□ Enable multi-path

Figure 2 **The Connect to Target wizard.**

Select a target and click the Connect button. You'll see a window similar to **Figure 2**. For most basic connections, ensure that the top checkbox is marked and click the OK button. Marking the top checkbox instructs the system to automatically restore the connection after every reboot.

There's also an Advanced button in this window. As you'll discover in a minute, all but the most basic of connections will require a set of advanced configurations, such as identifying the Initiator IP and Target portal IP—more on that shortly.

For your basic connection, there are two steps remaining to prepare your drive. First, select the iSCSI Initiator Volumes and Devices tab and click the Auto Configure button. This step automatically configures all available devices, further binding them so they're ready for use at the next system restart.

After this step, you'll find the disk is visible within Disk Management. Simply bring it online, initialize, and format the disk (if necessary). Your disk is now available for use just as if you had a direct-attached disk.

MPIO/MCS: Must-Have High Availability and Load Balancing

While connecting iSCSI disks to servers over your existing network is great for pervasive connectivity, your network's interconnections can and will create points of failure. Someone could accidentally unplug a cable, misconfigure a router, or any of the myriad of issues that happen on a traditional network. Thus, any iSCSI production use really requires redundant connections.

The seemingly easy answer might be to use NIC "teaming" like you do with your production network connections. However, classic NIC teaming for iSCSI connections is neither supported nor is it considered a best practice. Don't do it.

Connecting with iSCSI leverages its own set of protocols that handle high-availability and load balancing. You'll also find that iSCSI's Multipath Input/Output (MPIO) and Multiple Connected Sessions (MCS) protocols are superior in many ways to classic NIC teaming, as each protocol has a greater range of failover and load balancing capabilities.

MPIO is a much different protocol than MCS. Using MPIO requires a Device-Specific Module (DSM) connected to the server that runs the iSCSI Initiator. Microsoft includes its own "default" DSM with the Windows OS, installed as the Multipath I/O Feature within Server Manager.

Many storage devices can use that DSM with no additional software installation. Others require their own specialized DSM from the manufacturer. Consult your device manufacturer to determine if special driver installation is required or if the in-box Microsoft DSM is acceptable.

MCS requires no such DSM installation to the server. However, in order to use MCS, your storage device must support its protocol. Not all devices are MCS-enabled, which means you'll need to do a little research to determine which protocol is appropriate for your situation.

While different in their underlying code, managing their multipathing is fairly similar. Both MPIO and MCS provide a way to create multiple, parallel connections between a server and an iSCSI target. Most of what's needed for either is to specifically identify the NICs and networks you want to use.

Because MCS involves the fewest steps of the two protocols, I'll show its setup process. What you learn here will translate well to using MPIO. Remember the earlier example of a basic connection attached a server to storage through a single network connection. That connection existed between the single IP address on the server and the single IP address on the storage device.

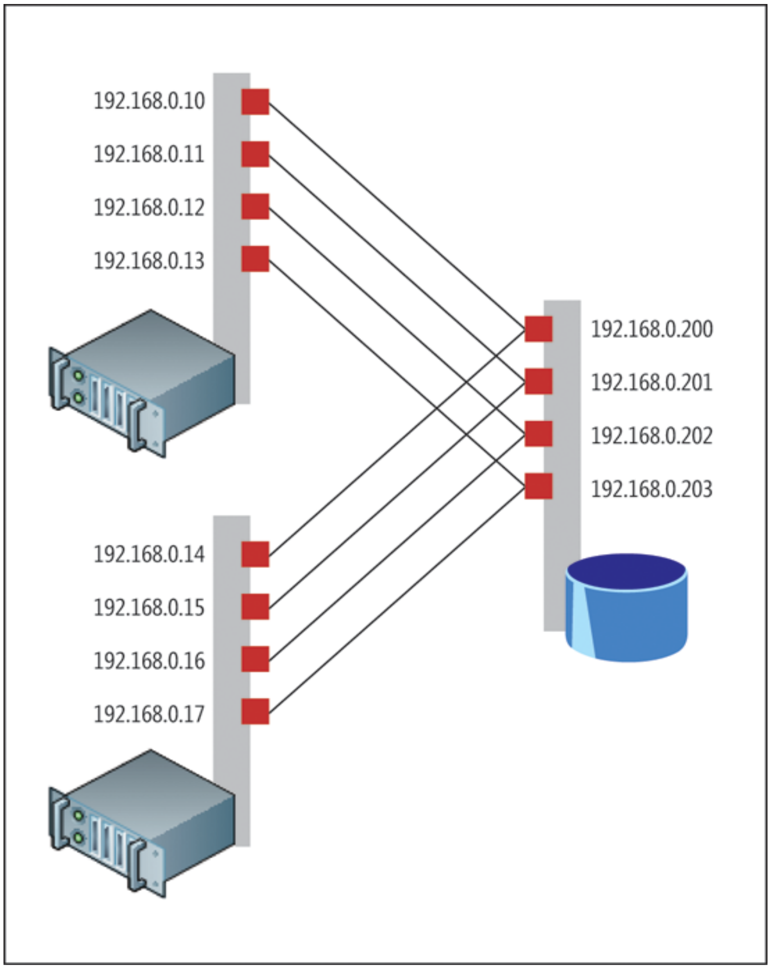

Figure 3 **Two servers, each with four network interfaces, connect with four network interfaces on a storage device.**

The "M" in MCS relates to augmenting that basic setup with multiple connections. Each connection relates to a network card and its associated IP address, with each iSCSI Target and Initiator using multiples of each. **Figure 3** shows how this might look when two servers, each with four network interfaces and associated IP address, are connected to four network interfaces and IP addresses on the storage device.

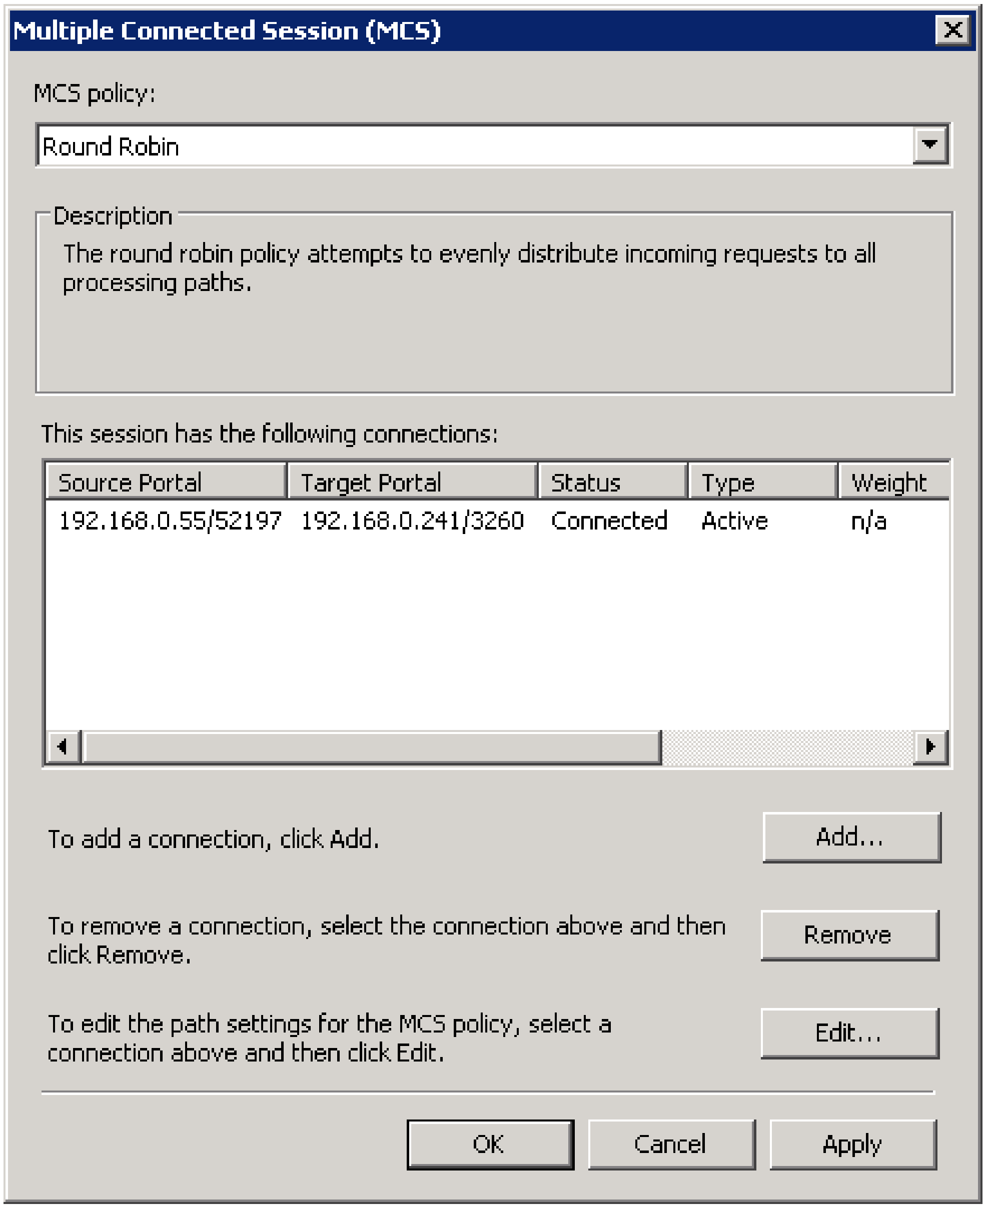

Multiple Connected Session (MCS)

MCS policy: Round Robin

Description

The round robin policy attempts to evenly distribute incoming requests to all processing paths.

This session has the following connections:

Source Portal	Target Portal	Status	Type	Weight
192.168.0.55/52197	192.168.0.241/3260	Connected	Active	n/a

To add a connection, click Add.

To remove a connection, select the connection above and then click Remove.

To edit the path settings for the MCS policy, select a connection above and then click Edit.

Figure 4 **The MCS configuration console.**

To manage MCS, select one of the targets in **Figure 1**, then click the Properties button, then the MCS button. You'll see a console that looks similar to **Figure 4**. The earlier example's "basic" setup configures a single connection between the Source Portal on the local server and the Target Portal on the storage device.

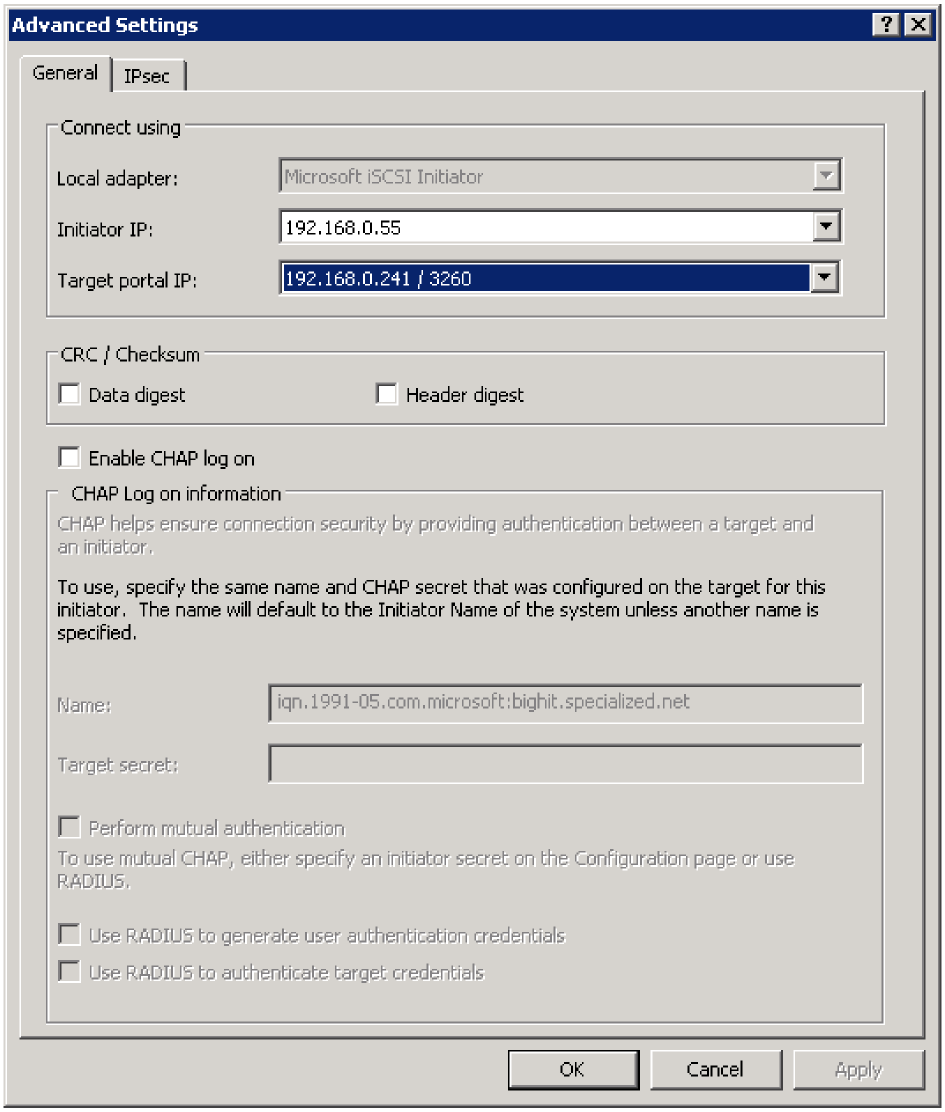

Connect using

Local adapter: Microsoft iSCSI Initiator

Initiator IP: 192.168.0.55

Target portal IP: 192.168.0.241 / 3260

CRC / Checksum

□ Data digest □ Header digest

□ Enable CHAP log on

CHAP Log on information

CHAP helps ensure connection security by providing authentication between a target and an initiator.

To use, specify the same name and CHAP secret that was configured on the target for this initiator. The name will default to the Initiator Name of the system unless another name is specified.

Name: iqn.1991-05.com.microsoft:bighit.specialized.net

Target secret:

□ Perform mutual authentication

To use mutual CHAP, either specify an initiator secret on the Configuration page or use RADIUS.

□ Use RADIUS to generate user authentication credentials

□ Use RADIUS to authenticate target credentials

Figure 5 **Advanced Settings for adding a connection.**

To add a connection, click the Add button followed by the Advanced button. This will bring you to the Advanced Settings console, as shown in **Figure 5**. In this console, you should designate the IP addresses for a second local NIC in the Initiator IP box, along with the IP address for a second remote target portal on your SAN.

If there are no additional target portal IP addresses available here, you'll need to discover them within the main console Discovery tab. Repeat this process for each initiator and target portal combination.

MCS Policies Define Behaviors

By creating these multiple connections, you provide more than one path (both physical and logical) through which you can transfer your storage network traffic. These multiple paths function as failover targets should you lose a connection. They can also load balance traffic, adding greater network capacity with each new connection.

Yet with these multiple connections must come some manner of defining how failover and load balancing behaves. With MCS, you can configure five policies:

1. **Fail Over Only:** Using a Fail Over Only policy, there is no traffic load balancing. It will only use a single path, with others remaining in standby mode until the first connection's path is lost.
2. **Round Robin:** This is the simplest policy that includes load balancing. Using this policy, traffic is rotated among available paths in order.
3. **Round Robin with a subset of paths:** This policy works similarly to Round Robin, except one or more paths are kept out of load balancing. These paths are used as standbys in the event of a primary path failure.
4. **Least Queue Depth:** Also similar to Round Robin, this policy load balances traffic by identifying and using the path with the least number of queued requests.
5. **Weighted Paths:** This policy lets you weight paths in situations where some may enjoy greater capacity than others. Traffic is balanced among paths as determined by the assigned weight.

Because MCS operates on a per-session basis, each individual session and connections can have its own MCS policy. Pay careful attention to your selected policy, as it can have a dramatic effect on the overall performance and availability of your storage connections.

The Perfect Fit for Small Environments

There's a saying about the iSCSI Initiator console—"work from left to right." Looking back at **Figure 1**, you can see six different tabs: Targets, Discovery, Favorite Targets, Volumes and Devices, RADIUS, and Configuration.

While many connections don't require edits to the RADIUS or Configuration settings, creating connections with this console works best when you start with Targets and continue configuring through Volumes and Devices. Conversely, getting rid of a connection involves reversing these steps and working from right to left.

While these extra steps in configuring iSCSI's high-availability options might seem cumbersome, remember you only need to do them as you're adding new disks to servers. Once connected, those disks will reconnect with every server reboot and reconfigure automatically should a path fail.

Because of iSCSI's reliance on existing networking infrastructure, the way it augments traditional SCSI can be a perfect fit for a small environment. An iSCSI SAN is a relatively inexpensive purchase, nowhere near the cost of yesteryear's refrigerator-sized SAN chassis. Without needing the arcane knowledge required by other storage mediums, iSCSI's ease of management makes it a great fit for the Jack-of-all-Trades IT professional.

Greg Shields, *MVP, is a partner at Concentrated Technology. Get more of Shields' Jack-of-all-Trades tips and tricks at ConcentratedTech.com.*

Related Content

- GOAT: A Case for a Layered Approach to Deploying Windows Desktops (December, 2009)
- GOAT: How to Install Non-Microsoft Patches Using System Center Essentials (September, 2009)
- GOAT: The Benefits of Big, Cheap Disks (June, 2009)

Windows Confidential

The Third Rail of Keyboard Shortcuts

Raymond Chen

In Windows 95, Microsoft introduced the Windows+E hotkey. You could use this keystroke combination to open My Computer. That's what it has done ever since, though not because people haven't been tempted to change its function. Every so often, someone tries to change what the Windows+E hotkey does based on user research. That degree of change is always met with strong resistance from the old-timers.

For example, during one of the beta cycles, developers changed the function so the Windows+E hotkey opened your user files. That's the same folder that opens when you click your name on the Windows Vista Start menu. I'm guessing the idea was that typical users spend time manipulating their files, not messing around with their hard drives.

A few months later, the target of the Windows+E hotkey in an internal build of Windows changed again so that it opened your Libraries, the same folder you get if you click on the word Libraries in the Windows 7 navigation bar. For one thing, the old behavior of opening the user files folder was redundant. For another thing, research data showed that when people opened My Computer, most of the time they just started clicking through their drives and folders looking for their stuff. In other words, it was the first step in what promised to be a long and painful journey.

Libraries provide an easier way to find your stuff, since they aggregate your documents into one place rather than making you hunt all over for them. Sending users to Libraries told them, "Hey, there's an easier way to do this. Let me help you find what you're looking for."

Fast-forward another few months and the target of the Explorer window that opens when you press the Windows+E key returned to opening My Computer. The request to change it back came from a programmer who worked with removable storage devices and file system filters—that sort of low-level stuff. He relied heavily on the old Windows+E hotkey opening My Computer so he could check on the status of the hard drives in his system to see if his driver was working correctly.

OK, it didn't change back just because of a low-level driver developer's feedback, but that triggered a re-evaluation of the hotkey. I think what may have finally tipped the balance back to having Windows+E send you to My Computer is all the help content and step-by-step instructions on the Internet and in printed materials that tell users to type Windows+E to open My Computer.

Long Road to Nowhere

It was a long journey that resulted in a net change of nothing, but we learned a lot along the way. For one thing, it shows that even though people complain that Windows is stale and never tries anything new, when you actually do something new—even if that new way is backed by research that demonstrates it's an improvement for the majority of users—people will tell you to change it back to the old comfortable way.

The Windows Explorer shortcut on the Start menu is not as heavily encumbered. In Windows 95, it opened the root of your C: drive. In Windows 2000, it changed to opening your My Documents folder (later named simply Documents). In Windows 7, the destination folder changed yet again, this time to your Libraries folder.

Here's a deployment tip: You can create a shortcut that opens an Explorer window on a folder of your choosing by simply creating a shortcut to that folder. You can give the shortcut a name of your choosing and deploy it to your users. If you're really sneaky, you can even call that shortcut Windows Explorer.

Raymond Chen*'s Web site,* The Old New Thing*, and his book of the same title (Addison-Wesley, 2007) deal with Windows history, Win32 programming and microwave popcorn.*

Related Content

- Windows Confidential: A Look at the Evolution of Shut Down on the Start Menu (March 2010)
- Windows Confidential: Tracking Shortcuts (October 2009)
- Windows Confidential: The Forgotten Hotkey (March 2009)

Utility Spotlight

Microsoft Office Environment Assessment Tool

Lance Whitney

If you're considering a migration to Microsoft Office 2010, you'll want to take a close look at this month's free utility—the Microsoft Office Environment Assessment Tool (OEAT). This tool checks a PC's overall configuration to see how it will fare during an upgrade to Office 2010. Beyond reporting on memory, disk space and similar factors, the tool also scans all add-ins installed in the current version of Office and gives you details on any Microsoft or third-party applications that interact with Office.

You can grab OEAT from the Microsoft Download Center. Download and run the self-extracting Office14EnvAssessment.exe file. The installation will prompt you for where you want to store the program's two extracted files:

- READ_ME.htm: this has links to the Quick Start Guide and User Manual
- OEAT.exe: the executable you trigger to run the scan

Upon launching OEAT.exe, you'll see the utility's main screen, as shown in **Figure 1**. To test the tool on your current PC, select the first option to scan your system. OEAT gives you two different scan options—a quick scan and a passive scan. What's the difference? The quick scan checks a default list of folders and Registry keys for add-ins and only takes a few seconds to run. The passive scan monitors specific Registry keys created or modified when an application calls for an Office API. This scan runs silently in memory and takes at least an hour.

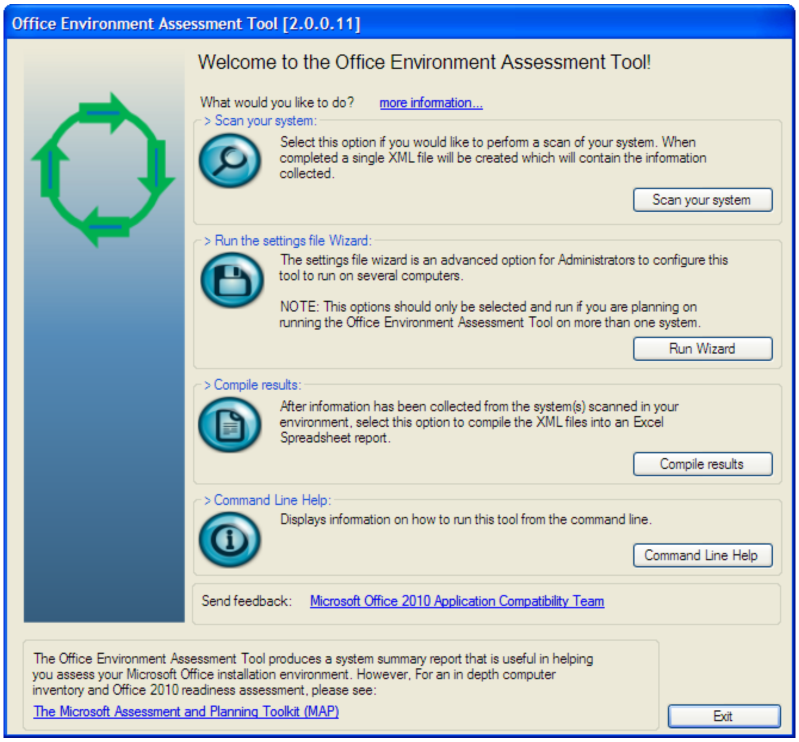

Office Environment Assessment Tool [2.0.0.11]
more information...
> Scan your system:
Scan your system
> Run the settings file Wizard:
Run Wizard
> Compile results:
Compile results
> Command Line Help:
Command Line Help
Microsoft Office 2010 Application Compatibility Team
The Microsoft Assessment and Planning Toolkit (MAP)
Exit

Welcome to the Office Environment Assessment Tool!

What would you like to do?

Select this option if you would like to perform a scan of your system. When completed a single XML file will be created which will contain the information collected.

The settings file wizard is an advanced option for Administrators to configure this tool to run on several computers.

NOTE: This options should only be selected and run if you are planning on running the Office Environment Assessment Tool on more than one system.

After information has been collected from the system(s) scanned in your environment, select this option to compile the XML files into an Excel Spreadsheet report.

Displays information on how to run this tool from the command line.

Send feedback: Microsoft Office 2010 Application Compatibility Team

The Office Environment Assessment Tool produces a system summary report that is useful in helping you assess your Microsoft Office installation environment. However, For an in depth computer inventory and Office 2010 readiness assessment, please see: The Microsoft Assessment and Planning Toolkit (MAP)

Figure 1 **The Office Environment Assessment Tool main screen.**

For the passive scan to be truly effective, you'd want to use Word, Excel, PowerPoint, Outlook and various third-party applications as the scan runs so that OEAT can monitor which apps make the Office API calls. It's a wise idea to run the quick scan first. You can always run the passive scan later to see if it catches more information.

The scan will capture the associated data and store it in two XML files, one to store the settings used in the scan and another to store the actual results. After the scan is complete, select Compile Results from the main screen and OEAT will organize the data into individual worksheets within an Excel workbook, which automatically opens in Excel, as shown in **Figure 2**.

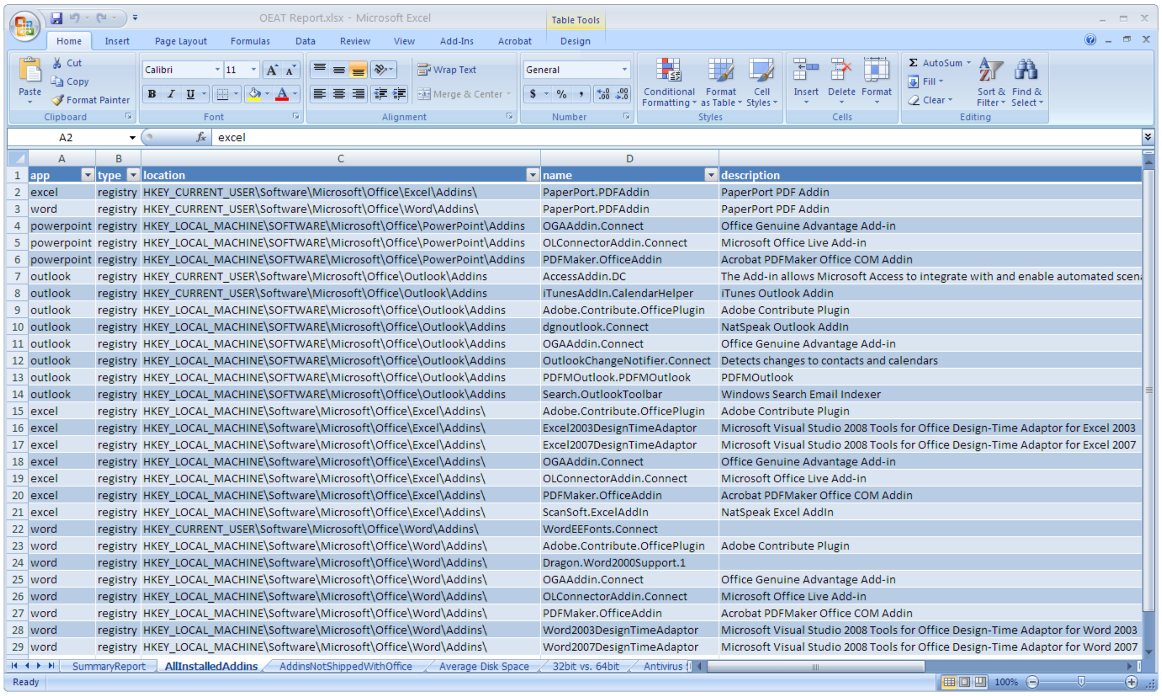

	app	type	location	name	description
2	excel	registry	HKEY_CURRENT_USER\Software\Microsoft\Office\Excel\Addins\	PaperPort.PDFAddin	PaperPort PDF Addin
3	word	registry	HKEY_CURRENT_USER\Software\Microsoft\Office\Word\Addins\	PaperPort.PDFAddin	PaperPort PDF Addin
4	powerpoint	registry	HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Office\PowerPoint\Addins	OGAAddin.Connect	Office Genuine Advantage Add-in
5	powerpoint	registry	HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Office\PowerPoint\Addins	OLConnectorAddin.Connect	Microsoft Office Live Add-in
6	powerpoint	registry	HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Office\PowerPoint\Addins	PDFMaker.OfficeAddin	Acrobat PDFMaker Office COM Addin
7	outlook	registry	HKEY_CURRENT_USER\Software\Microsoft\Office\Outlook\Addins	AccessAddin.DC	The Add-in allows Microsoft Access to integrate with and enable automated scena
8	outlook	registry	HKEY_CURRENT_USER\Software\Microsoft\Office\Outlook\Addins	iTunesAddIn.CalendarHelper	iTunes Outlook Addin
9	outlook	registry	HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Office\Outlook\Addins	Adobe.Contribute.OfficePlugin	Adobe Contribute Plugin
10	outlook	registry	HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Office\Outlook\Addins	dgnoutlook.Connect	NatSpeak Outlook AddIn
11	outlook	registry	HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Office\Outlook\Addins	OGAAddin.Connect	Office Genuine Advantage Add-in
12	outlook	registry	HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Office\Outlook\Addins	OutlookChangeNotifier.Connect	Detects changes to contacts and calendars
13	outlook	registry	HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Office\Outlook\Addins	PDFMOutlook.PDFMOutlook	PDFMOutlook
14	outlook	registry	HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Office\Outlook\Addins	Search.OutlookToolbar	Windows Search Email Indexer
15	excel	registry	HKEY_LOCAL_MACHINE\Software\Microsoft\Office\Excel\Addins\	Adobe.Contribute.OfficePlugin	Adobe Contribute Plugin
16	excel	registry	HKEY_LOCAL_MACHINE\Software\Microsoft\Office\Excel\Addins\	Excel2003DesignTimeAdaptor	Microsoft Visual Studio 2008 Tools for Office Design-Time Adaptor for Excel 2003
17	excel	registry	HKEY_LOCAL_MACHINE\Software\Microsoft\Office\Excel\Addins\	Excel2007DesignTimeAdaptor	Microsoft Visual Studio 2008 Tools for Office Design-Time Adaptor for Excel 2007
18	excel	registry	HKEY_LOCAL_MACHINE\Software\Microsoft\Office\Excel\Addins\	OGAAddin.Connect	Office Genuine Advantage Add-in
19	excel	registry	HKEY_LOCAL_MACHINE\Software\Microsoft\Office\Excel\Addins\	OLConnectorAddin.Connect	Microsoft Office Live Add-in
20	excel	registry	HKEY_LOCAL_MACHINE\Software\Microsoft\Office\Excel\Addins\	PDFMaker.OfficeAddin	Acrobat PDFMaker Office COM Addin
21	excel	registry	HKEY_LOCAL_MACHINE\Software\Microsoft\Office\Excel\Addins\	ScanSoft.ExcelAddIn	NatSpeak Excel AddIn
22	word	registry	HKEY_CURRENT_USER\Software\Microsoft\Office\Word\Addins\	WordEEFonts.Connect	
23	word	registry	HKEY_LOCAL_MACHINE\Software\Microsoft\Office\Word\Addins\	Adobe.Contribute.OfficePlugin	Adobe Contribute Plugin
24	word	registry	HKEY_LOCAL_MACHINE\Software\Microsoft\Office\Word\Addins\	Dragon.Word2000Support.1	
25	word	registry	HKEY_LOCAL_MACHINE\Software\Microsoft\Office\Word\Addins\	OGAAddin.Connect	Office Genuine Advantage Add-in
26	word	registry	HKEY_LOCAL_MACHINE\Software\Microsoft\Office\Word\Addins\	OLConnectorAddin.Connect	Microsoft Office Live Add-in
27	word	registry	HKEY_LOCAL_MACHINE\Software\Microsoft\Office\Word\Addins\	PDFMaker.OfficeAddin	Acrobat PDFMaker Office COM Addin
28	word	registry	HKEY_LOCAL_MACHINE\Software\Microsoft\Office\Word\Addins\	Word2003DesignTimeAdaptor	Microsoft Visual Studio 2008 Tools for Office Design-Time Adaptor for Word 2003
29	word	registry	HKEY_LOCAL_MACHINE\Software\Microsoft\Office\Word\Addins\	Word2007DesignTimeAdaptor	Microsoft Visual Studio 2008 Tools for Office Design-Time Adaptor for Word 2007

Figure 2 **Office Environment Assessment Tool compiled results open automatically in Excel.**

The first worksheet—SummaryReport—gives you information about your PC, including hard disk space, memory and processor type. It also displays the version of Windows and Office, along with details on your antivirus software.

The second worksheet, AllInstalledAddins, shows you the Registry keys for any add-ins that interact with Office. On one of my scans, for example, OEAT displayed the keys of Office add-ins for Adobe PDF Maker, Dragon NaturallySpeaking and Nuance PaperPort scanning software. This list also covers Microsoft's own Office add-ins.

The third worksheet, AddinsNotShippedWithOffice, shows the add-ins not included with Office. The next several worksheets (Average Disk Space, 32bit vs. 64bit, Antivirus Status and more) display pie charts to visually illustrate some of the data included in the other worksheets. The final worksheet, RawData, sums up the information from the other worksheets and adds further details, including computer name, IP address, user account, domain name and others.

This utility automatically saves the Excel workbook with the filename OEAT Report.xlsx. That file and the XML files are stored in the same folder where you extracted the OEAT.exe file and the READ_ME file. An errors.log is also generated in the event the scan fails to run properly.

Plan on This Scan

If you're planning to deploy Office 2010 throughout your organization, you'll want to run OEAT on all of your client PCs via an automated script. To set the script up, select the Run Wizard option from the main screen. You'll be able to choose the different options for the scan, such as a quick scan versus a passive scan. You can opt to run the quick scan silently (the passive scan will automatically run in silent mode). You can even choose the duration for the passive scan, anywhere from one hour to 23 days.

Finally, you'll need to specify a mapped drive or UNC path in which to save the resulting XML files generated by your client scan. Make sure your users have read, write and execute privileges to this location. Your users must also be local administrators on their PCs so that OEAT can audit the appropriate Registry keys. However, Microsoft does offer a workaround if you can't grant your users local admin rights.

After you complete the Wizard, OEAT writes a settings.xml file to the folder from which you launched OEAT. This file stores all the options you've specified. You can always open this file in an XML reader to review and confirm the settings.

Your next step is to create an OEAT script to run on all your workstations via a login script, Group Policy or some other means. To view the parameters for this script, select the Command Line Help button from the OEAT main screen. The directions are short and simple. Make sure the OEAT.exe file and the settings.xml file are stored in the same network folder. Then simply use the oeat -scan string in your script.

After your users run the script, OEAT will store a separate XML file for each computer in the network location that you specified. Once all or most of the results are in, you can tell OEAT to compile them into an Excel workbook. To do this, simply open up a command prompt to the network location and type oeat -compile.

Excel displays the same data you saw when you ran the scan on a single PC, but now each worksheet contains details on multiple computers. Depending on the number of PCs in your organization, some worksheets could have tens of thousands of lines. Because the workstation name, IP address, domain name and other details are included, you should be able to sort or filter the results in Excel to better organize them.

As you review the results, you can see which Office add-ins are installed among your users. You can then research those add-ins to find out if they'll be compatible with Office 2010. Each add-in includes its manufacturer and version number to help you in your research.

OEAT can scan the past several versions of Microsoft Office, including Office 97, Office 2000, Office XP, Office 2003, and Office 2007. Because the report is generated as an Excel spreadsheet, you'll also need that to compile the results—preferably Excel 2007 or Excel 2010.

Lance Whitney *is a writer, IT consultant and software trainer. He's spent countless hours tweaking Windows workstations and servers. Originally a journalist, he took a blind leap into the IT world in the early '90s.*

Related Content

- Utility Spotlight: The Compatibility Question (April 2010)
- Assessment Tools for Office 2010
- Office Environment Assessment Tool User Guide

SQL Q&A

Moving, Performance Tuning, Backing Up and Mirroring Databases

Paul S. Randal

New Array Moving Day

Q: Our current RAID is filling up quickly, so we need to move some SQL Server 2005 databases elsewhere. The new array is ready and I've been preparing to move the databases. I've just discovered that one of the databases is a transactional replication publisher and I know that means I can't move the database. What should I do?

A: There is good news for you—only SQL Server 2000 (and earlier) had the limitation that restricted moving a publication database without reinitializing transactional replication or directly altering the various system tables.

With SQL Server 2005 and SQL Server 2008, there is a documented process that lets you move a database without having to do anything with transactional replication, as long as the database remains attached to the same instance of SQL Server. You have to accept some downtime, as there's no way to move a database file while it's still online. The procedure is as follows:

First, take the database offline using the code below. If there are users connected to the database, you'll need to drop them first for this process to succeed:

```
ALTER DATABASE MyDatabaseName SET OFFLINE;
```

Next, copy the data files to the new location. Use copy instead of move to allow for a fast rollback in case anything goes wrong (otherwise, you will have to perform a restore). Then let SQL Server know the new location of each file with the following code:

```
ALTER DATABASE MyDatabaseName
MODIFY FILE
   (NAME = N'LogicalFileName',
   FILENAME = N'pathname\filename');
```

Once you've physically copied all the files and updated their locations in SQL Server, bring the database back online with the code:

```
ALTER DATABASE MyDatabaseName SET ONLINE;
```

SQL Server will automatically use the new file locations and continue running transactional replication. The physical location of the database files is orthogonal to transactional replication, so as far as replication is concerned, the only thing that happened was the publication database was briefly offline.

There are additional steps you'll need to take if you're moving the database because of a problem with the I/O subsystem. You can find more on these and other procedures in the

SQL Server 2008 Books Online topic Moving User Databases (which also links to the relevant topic in SQL Server 2005 Books Online).

Closing the Page Latch

Q: I'm having problems understanding some of the concepts around performance tuning. I've read several times that I need to prevent "page latch" issues. I don't know what it means by a "page" or a "latch," or why a page latch could even be an issue. Can you explain all this?

A: All data in a SQL Server database is stored in data files. Internally, these files are organized into sequences of 8KB chunks called *pages*. A page is the basic unit of storage and I/O that SQL Server can manage. Pages are usually in the data files on disk, and need SQL Server's cache (known as the *buffer pool*) to read them before processing any queries.

SQL Server uses different kinds of pages to store different kinds of relational data (such as rows from a table, rows from a non-clustered index or text/LOB data). There are also pages that store parts of the internal data structures necessary for SQL Server to organize and access the pages storing the relational data.

A *latch* is a lightweight internal mechanism SQL Server uses to synchronize access to a page within the cache. There are two types of page latches you need to watch—*regular page latches* and *page I/O latches*. If a SQL Server thread has to wait to acquire one of these latches, it indicates a performance issue.

When SQL Server is waiting for a portion of a data file to be read from disk, it can cause a page I/O latch wait. If a page I/O latch takes an excessive amount of time, this usually indicates a performance problem with the underlying disk subsystem (that is, it's overloaded).

When multiple threads within SQL Server are trying to access the same 8KB data file page in memory, and there is contention for access to the page, this can cause a page latch wait. The most common occurrence of this involves heavy use of small temporary objects in the tempdb database.

A deeper explanation of how to monitor and mitigate page latch waits is beyond the scope of this column, but you can find more information in:

- SQL Server 2008 Books Online, in the "SQL Server, Wait Statistics Object" section, shows how to monitor wait statistics using the System Monitor.
- SQL Server 2008 Books Online, in the "sys.dm_os_wait_stats" section, lists common SQL Server wait types and what they mean, plus how to monitor wait statistics from inside SQL Server.
- The white paper "Troubleshooting Performance Problems in SQL Server 2008," provides various troubleshooting queries and techniques, including wait statistics.

Looking Through Database Snapshots

Q: I've just discovered database snapshots. Now I'm considering using them as an alternative to the full recovery model and log backups. I'll create a snapshot every hour or so, and that way if something goes wrong, I can pull back the damaged data. They seem like a lot less hassle and a much faster way to restore. Do you see any problems with making this change?

A: Yes—database snapshots are not a practical or viable substitute for a comprehensive disaster recovery strategy. A database snapshot does not provide the same capability as a transaction log backup in terms of fully recovering from a disaster. The database snapshot does not contain a copy of all pages from the database, only those that have changed since it was first created. This means that if the database is damaged in any way, the database snapshot is useless without the underlying database. It's just a collection of disparate pages from the database and can't be used for recovery.

A database snapshot will let you pull back data that has been accidentally deleted from the database, as long as the database itself is still available. If a dropped table in the database still exists in the snapshot, for example, you can use it to recreate that dropped table.

That being said, it's not a good idea to create too many snapshots of a database (as a substitute for a one-half-hourly transaction log backup) because of the potential performance problems. Before you can exchange a database page (see the explanation in the answer in the "Closing the Page Latch section"), you must first synchronously copy the page into all existing database snapshots that do not already contain a version of that page. As you create more database snapshots, the more page copies you have to make, and performance degrades.

One other reason to not create too many database snapshots is that each will contain pre-change copies of database pages. Each of them will grow as more of the database changes. This can lead to disk space problems, as well as performance problems.

Database snapshots are not designed to be a substitute for frequent log backups. You can read a more in-depth study of the performance implications of database snapshots in the white paper "Database Snapshot Performance Considerations Under I/O-Intensive Workloads."

Also, if you're using the full recovery model and transaction log backups, then you're obviously interested in being able to recover up to the point of a disaster and/or make use of point-in-time restores. (For an explanation of these see my July 2009 and November 2009 articles, "Understanding SQL Server Backups" and "SQL Server: Recovering from Disasters Using Backups," respectively.)

Mirror, Mirror

Q: I've been asked to set up a database mirror for our database, but I'm concerned that database mirroring isn't going to help with our problem. We've had some corruption issues with our SAN, so the plan is to have database mirroring protect us from corruption. Won't any corruption automatically be sent over to the mirror? How is database mirroring going to help with this?

A: This is an issue that causes a lot of confusion. It would seem that any technology that provides a redundant copy of a database would be susceptible to corruption propagating from the principal to the mirror database (to use the database mirroring terminology)—but in reality this does not occur.

The crux of the matter lies in understanding how the mirror database is maintained. Corruption would certainly be propagated to the mirror if the underlying synchronization mechanism copied complete database pages from the principal to the mirror database. A corrupt page from the principal would then be placed in the mirror.

However, database mirroring specifically avoids this because it doesn't copy database pages from one database to the other. Database mirroring works by copying the transaction log records from the principal database to the mirror. Transaction log records describe physical changes made to database pages, and do not contain the actual pages themselves. (For a complete explanation of transaction log records, logging and recovery see my article from February 2009: "Understanding Logging and Recovery in SQL Server.")

Even if a database page is corrupted by the underlying I/O subsystem of the principal database, there's no way for that corruption to directly propagate to the mirror database. The worst that could possibly happen is if SQL Server doesn't detect a page corruption (because page checksums are not enabled), and a corrupt column value is used to calculate a value stored in the database. The resulting incorrect result would be propagated to the mirror database—a *second-order corruption effect*. As I mentioned, if page checksums are enabled, such corruption would remain undetected when the page is read from disk, and the second-order corruption would not occur.

This behavior also explains why running a consistency check on the principal database doesn't yield any information about the consistency state of the mirror database and vice-versa. They are two separate databases kept synchronized by shipping descriptions of physical changes to the database, not the actual database pages.

Editor's Note: Thanks to Kimberly L. Tripp of SQLskills.com for providing a technical review of this month's column.

Paul S. Randal *is the managing director of SQLskills.com, a Microsoft regional director and a SQL Server MVP. He worked on the SQL Server Storage Engine team at Microsoft from 1999 to 2007. He wrote DBCC CHECKDB/repair for SQL Server 2005 and was responsible for the Core Storage Engine during SQL Server 2008 development. Randal is an expert on disaster recovery, high availability and database maintenance, and is a regular presenter at conferences worldwide. He blogs at SQLskills.com/blogs/paul, and you can find him on Twitter at Twitter.com/PaulRandal.*