CONTENT PROVIDED BY

Redmond VIRTUALIZATION REVIEW

SPONSORED BY

VEEAM IT JUST WORKS!™

Hyper-V Evolves

How Microsoft's virtualization platform is developing and what it means for technology professionals

- First Look: Hyper-V Server Technical Preview *Page 1*
- First Look: Microsoft Azure RemoteApp *Page 8*
- Disaster Recovery as a Service *Page 15*
- Hyper-V Replica for Disaster Recovery *Page 22*
- 7 DRaaS Platforms Gaining Speed *Page 29*
- First Look: Microsoft Azure Site Recovery *Page 34*
- Manage Mobile Devices and Policies in Active Directory *Page 42*
- To Join or Not to Join? *Page 50*

1105 MEDIA

VEEAM
IT JUST WORKS!™

Veeam Availability Suite **v8**

AVAILABILITY™
for the Modern Data Center

RTPO <15 min for **ALL** applications and data

Veeam® bridges the availability gap by providing Availability for the Modern Data Center™, which delivers RPOs and RTOs (RTPO™) of < 15 minutes for ALL applications and data.

www.veeam.com

First Look:

Hyper-V Server Technical Preview

A look at the first test build of the upgraded Microsoft hypervisor server shows improved migration, upgraded Hyper-V Manager and new VM versions. **By Brien M. Posey**

Microsoft last month delayed the release of the next major Windows Server upgrade until next year. Having last upgraded Windows Server in late 2013, Microsoft had originally suggested the next version would arrive this year. What little we've seen of the next Windows Server version, often referred to as vNext, were some major new features slated for its hypervisor, Hyper-V.

Hyper-V is a key component of the Microsoft Cloud OS.

Indeed, Hyper-V is a key component of the Microsoft Cloud OS, the company's next-generation datacenter platform enabling hybrid cloud architectures. The company said it will release a second preview of Hyper-V this year. Having spent several months with the first Hyper-V Server Technical Preview, there are a number of key new features that will benefit systems administrators. So far Microsoft has disclosed roughly a dozen new features as part of the technical preview, which include:

- Rolling Hyper-V Cluster Upgrade
- Storage Quality of Service (QoS)
- Virtual Machine Configuration Version
- New Virtual Machine Configuration File Format
- Production Checkpoints

- Hyper-V Manager Improvements
- Integration Services Delivered Through Windows Update
- Hot Add and Remove Network Adapters and Memory
- Linux Secure Boot
- Compatible with Connected Standby

At first glance, it's easy to scoff at the list of new features. After all, the last two Hyper-V releases each boasted hundreds of new features, while this list of new features in the Hyper-V Server Technical Preview is comparatively quite modest. However, there are a couple of things to consider.

One way Microsoft is helping to eliminate the down time commonly associated with the migration process is through the Rolling Cluster Upgrade feature.

First, this is a preview release—Microsoft could announce additional features. Also, many of the new features are geared toward eliminating downtime that would otherwise occur as a result of maintenance and migration processes. This is extremely important because so many organizations now run mission-critical workloads on Hyper-V (see the June 2014 feature, "Hyper-V Moves into the Fast Lane," at Redmondmag.com/HyperV0614).Such an organization would presumably like to avoid taking a mission-critical workload offline as part of the migration to the next version of Hyper-V.

Rolling Cluster Upgrade

One way Microsoft is helping to eliminate the down time commonly associated with the migration process is through the Rolling Cluster Upgrade feature. Perhaps a better name for it would be Mixed Clusters, because the feature allows a clustered Hyper-V deployment to include a mixture of legacy cluster nodes and new cluster nodes.

An example is the Failover Cluster Manager (see **Figure 1**) as it appears in Windows Server 2012 R2. Prior to the release of the Windows Server Preview, I had created a Hyper-V cluster that was

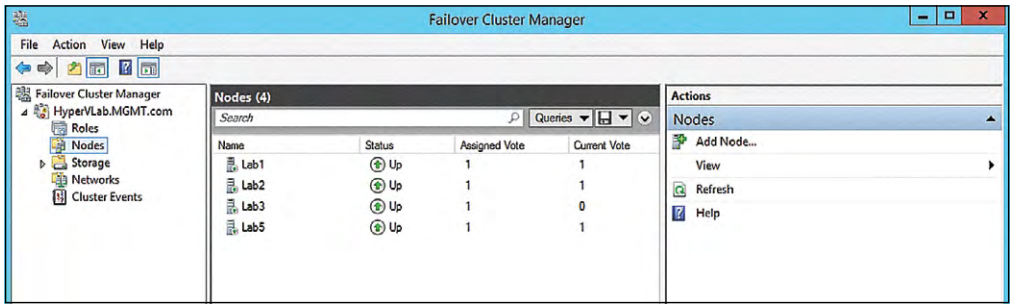

Figure 1. Joining a Preview release server to an existing Windows Server 2012 R2 cluster.

based solely on Windows Server 2012 R2 servers. Once Microsoft released the new Windows Server Technical Preview, I was able to join it to my existing Windows Server 2012 R2-based cluster. As you can see in **Figure 1**, p. 2, servers Lab1, Lab2, and Lab3 are all running Windows Server 2012 R2, whereas server Lab5 is running the Windows Server Technical Preview. Note: I didn't have to do anything special in order to join the cluster.

This is what Microsoft means when it refers to rolling cluster upgrades. You don't have to upgrade your cluster nodes all at once, nor do you have to create a brand-new cluster in order to migrate to the next version of Windows Server. You can gradually update the servers in your cluster (or bring new nodes into the cluster) as you see fit. The cluster can run a mixture of Windows Server versions for an indefinite period of time.

You don't have to upgrade your cluster nodes all at once, nor do you have to create a brand-new cluster in order to migrate to the next version of Windows Server.

Virtual Machine Migrations

Mixing Windows Server versions within a cluster is nice, but you might have concerns about how doing so impacts your virtual machines (VMs). A VM running on Windows Server 2012 R2 Hyper-V can be live migrated to a Hyper-V Server Technical Preview using

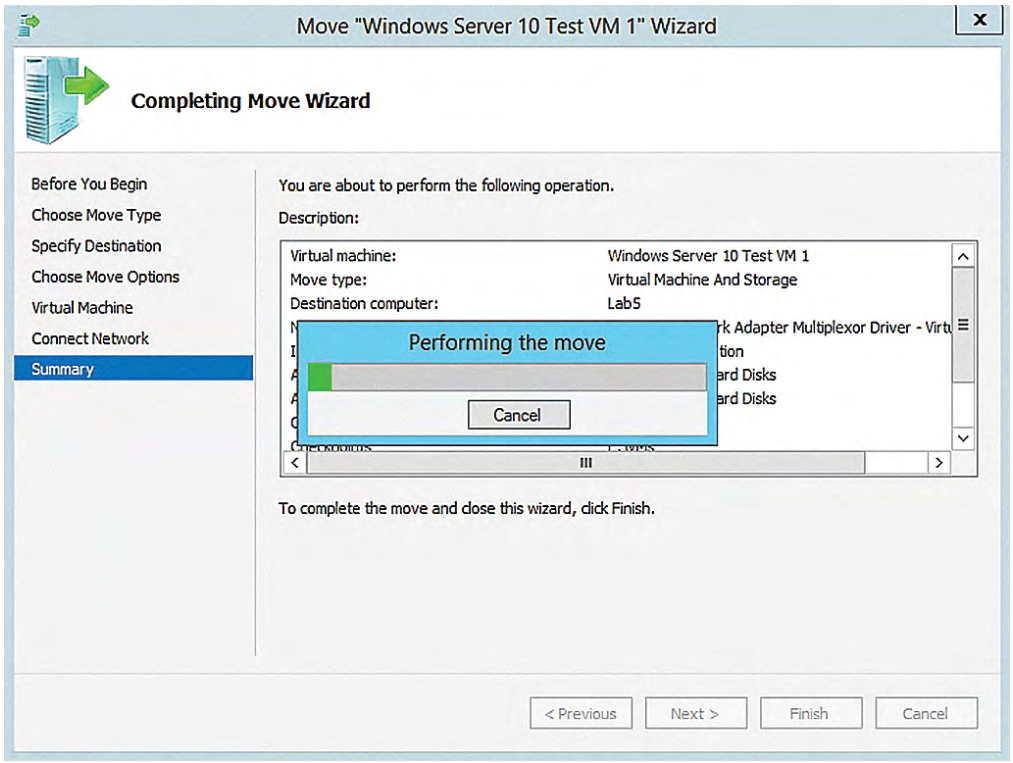

Figure 2. Virtual machines can be live migrated from one version of Hyper-V to another.

exactly the same method that would be used if that VM were being migrated to another Windows Server 2012 R2 Hyper-V (see **Figure 2**, p. 3). This holds true whether or not you're using failover clustering.

It's worth noting live migrations to Hyper-V Server Technical Preview aren't necessarily permanent actions. You can live migrate back to a legacy Hyper-V server without any problems.

As great as it is to be able to mix server versions within a failover cluster and to live migrate a running VM between different versions of Hyper-V, it would be irresponsible of me not to point out that your ability to live migrate a VM from a Hyper-V Server Technical Preview to a legacy Hyper-V server hinges on the version of the VM.

Your ability to live migrate a VM from a Hyper-V Server Technical Preview to a legacy Hyper-V server hinges on the version of the VM.

Virtual machines created on a server running Windows Server 2012 R2 Hyper-V have a version number of 5. As long as the version number remains unchanged, the VM can be live migrated across different versions of Hyper-V servers to your heart's content. However, if you want to use any of the new VM-specific features (such as production checkpoints), then you'll have to upgrade the VM version to 6. Once the VM version has been updated, it can no longer be hosted on a legacy Hyper-V server.

This brings up an important point. If a failover cluster is running a mixture of new and legacy Hyper-V servers, then you won't be able to upgrade the VM version number until all of the failover cluster nodes are running the new version of Hyper-V and the cluster functional level is also upgraded. At that point, it's impossible to add any legacy nodes to the cluster.

Incidentally, if you have a cluster containing multiple Hyper-V versions, then it's important to use the Hyper-V Server Technical Preview to manage the cluster. The Failover Cluster Manager that's included with Windows Server 2012 R2 might not always work properly with the Windows Server Technical Preview.

Hyper-V Manager Improvements

Microsoft has added three noteworthy improvements to Hyper-V Manager. The first is support for alternate credentials. This makes it a lot easier to connect to remote Hyper-V servers that use a different set of administrative credentials.

Second, the console now communicates with Hyper-V the WS-MAN protocol. This allows for CredSSP, Kerberos, or NTLM Authentication. The third improvement is related to backward compatibility. The Hyper-V Manager console included with the Technical Preview supports the management of legacy Hyper-V servers. Take an implementation where the console is connected to five different Hyper-V servers. Lab1, Lab2, Lab3, and Lab4 all run Windows Server 2012 R2 Hyper-V, while Lab5 is running the Hyper-V Technical Preview (see **Figure 3**). Hence, multiple versions of Hyper-V can be managed through a common console.

Virtual Machine Versions

The console now communicates with Hyper-V the WS-MAN protocol.

As noted, legacy VMs are based on version 5 and new VMs use version 6. But what does this really mean in a practical sense?

You'll notice in **Figure 3** that there are two VMs on the server Lab5. One of these VMs (Windows Server 10 Test VM 1) was live migrated from a legacy Hyper-V server. The other VM (Windows Server 10 Test VM 3) was created directly on the Hyper-V Server Technical Preview server. With that said, take a look at the summary information for the currently selected VM. As you can see in the figure, the VM version is 6.0.

If you need to check VM versions in bulk, you can do so through Windows PowerShell by entering the following command:

```
Get-VM * | Select-Object Name, Version
```

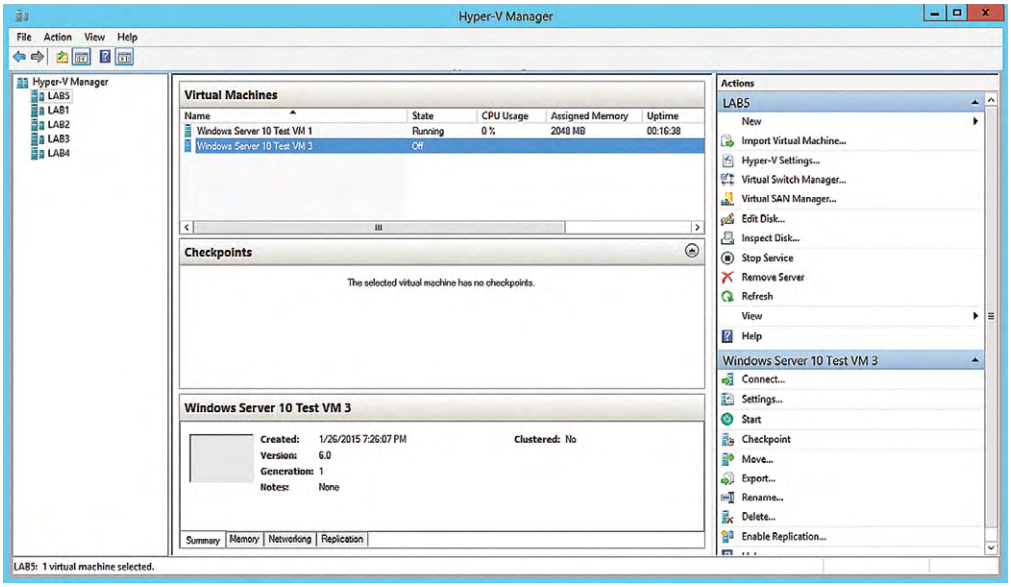

Figure 3. The Hyper-V Manager console can be used to manage multiple versions of Hyper-V.

Virtual Machine Structure

It isn't just the versions of the VMs that are dissimilar, but also the VM structures. That's because there's more to a VM than its virtual hard disks. Hyper-V uses several different file types. Some of the more important file types used by Windows Server 2012 R2 Hyper-V (at least for the purposes of this discussion) include:

- .VHD and .VHDX: virtual hard disk files
- .XML: VM configuration data
- .VSV: VM save state file
- .BIN: VM runtime state data (memory contents)

Microsoft has changed some of the file types that are used for version 6 VMs.

Every version 5 VM has a configuration data file and there is also usually at least one virtual hard disk, The .VSV and .BIN files may or may not exist, depending on the VM's current state.

Microsoft has changed some of the file types that are used for version 6 VMs. This was done as a way of improving efficiency and of reducing the chances of data corruption in the event of a storage failure.

Version 6 VMs still use .VHD or .VHDX based virtual hard disks. The folder structure used by the VM is still the same as well. Virtual hard disks and the other VM components are stored in subfolders beneath the VM name. For example, I created a version 6 VM named Windows Server 10 Test VM 3 and I stored it at F:\VMs. As such, the main paths used by the VM are:

- F:\VMs\Windows Server 10 Test VM 3\Virtual Hard Disks
- F:\VMs\Windows Server 10 Test VM 3\Virtual Machines

These are the same paths that are used by Windows Server 2012 R2 Hyper-V. You really don't begin to see the differences until you dig into the Virtual Machines folder.

If you open the Virtual Machines folder for a version 5 VM, you'll see a folder whose name matches the VM GUID. There's also an .XML file that uses the VM GUID as its file name (see **Figure 4**, p. 7). If a .VSV and a .BIN file exist for the VM, they'll be in the ‹GUID› folder.

With the Virtual Machines folder for a version 6 VM, the GUID is still used, but the .XML-based configuration file has been replaced by a .VMCX and a .VMRS file (see igure 5). The .VMCX file is a binary

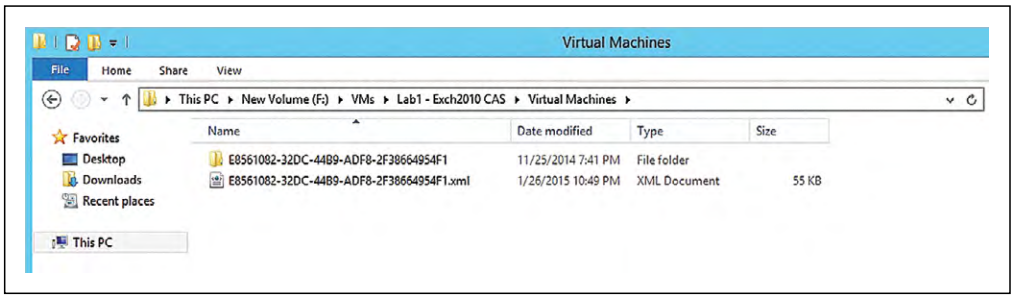

Figure 4. This is what a version 5.0 Virtual Machines folder looks like.

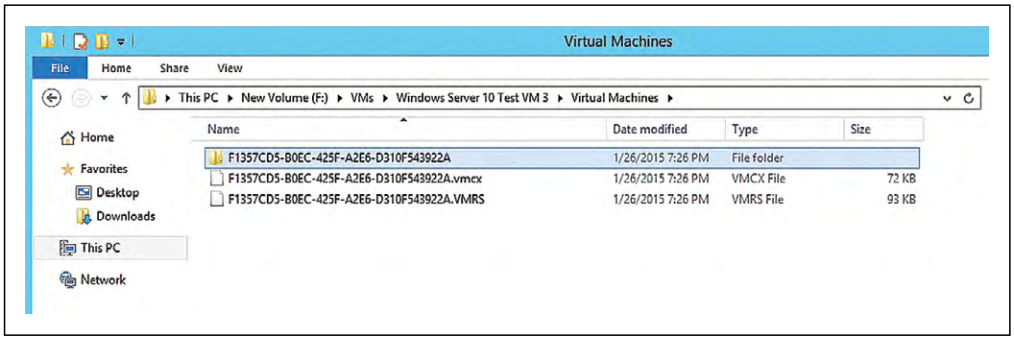

Figure 5. Version 6 virtual machines use a different structure than version 5 virtual machines.

Microsoft has designed the first Hyper-V Server Technical Preview in a way that allows it to fit seamlessly into your existing Hyper-V infrastructure.

file used to store the VM configuration data. The .VMRS file contains runtime state data.

Upgrading a Virtual Machine

While you can upgrade a VM from version 5 to version 6 in order to unlock the new features for the VM, keep in mind once you upgrade the VM you can't host it on a legacy Hyper-V server. With that caveat, you can convert a virtual machine to version 6 by entering the following command into Windows PowerShell:

```
Update-VmConfigurationVersion <virtual machine name>
```

Microsoft has designed the first Hyper-V Server Technical Preview in a way that allows it to fit seamlessly into your existing Hyper-V infrastructure. This allows Hyper-V admins to transition to the new version in a way that makes sense for their own organization. VR

Brien M. Posey is a seven-time Microsoft MVP with more than two decades of IT experience. He's written thousands of articles and several dozen books on a wide variety of IT topics. Visit his Web site at brienposey.com.

First Look: Microsoft Azure RemoteApp

The new Microsoft Remote Desktop Application as a Service is now available and takes traditional RDP to a new level.

By Brien M. Posey

As users make the transition away from traditional desktop PCs to all manner of mobile devices, application delivery and application compatibility become increasingly challenging issues. Organizations still largely depend on a variety of Windows-based line-of-business applications. However, these applications are typically not natively compatible with mobile platforms such as the Windows Runtime, iOS, Android or Windows Phone. Because the adoption of mobile consumer electronics devices in the enterprise has become such a pervasive trend, a number of different application delivery solutions are now available from a variety of players.

Azure RemoteApp is similar to Windows RemoteApp in that it allows organizations to stream applications to devices through a Remote Desktop Protocol (RDP) session, even if the device isn't natively capable of running the application.

Microsoft's latest attempt to address the challenge of delivery of Windows applications to non-Windows devices is the company's new Azure RemoteApp. Introduced last spring, Microsoft released Azure RemoteApp at the end of the year. It's similar to Windows RemoteApp in that it allows organizations to stream applications to devices through a Remote Desktop Protocol (RDP) session, even if the device isn't natively capable of running the application. All the device needs is a native RDP client.

Azure RemoteApp is available on a subscription basis. The $10 per user per month Basic Plan is for employees typically dedicated to a specific task using simple Web apps, while the $15 Standard plan is for knowledge workers who use the likes of Microsoft Office. The company issued a detailed price plan, which is available at **bit.ly/1D4Dg82**.

Planning Your Deployment

As you consider using Azure RemoteApp to deliver applications to devices, you must decide on the type of deployment you want to create. Microsoft gives you two options: You can create a cloud-based deployment or you can create a hybrid implementation.

A cloud-based deployment is the easier of the two options, but it isn't appropriate for every organization. A cloud-based deployment is less complex than a hybrid deployment because the deployment exists solely within the Azure cloud. You don't have to involve back-end infrastructure on your private network. When you create a cloud deployment, Microsoft provides you with Office 2013 apps that you can make available to your users. Of course, most organizations will likely have additional applications they need to publish. For that you have the ability to build a custom template image.

A cloud-based deployment is less complex than a hybrid deployment because the deployment exists solely within the Azure cloud.

Like cloud deployments, hybrid deployments allow you to make applications available to your users. However, hybrid deployments offer the distinct advantage of allowing the applications to run in a domain-joined environment, which makes it possible for users to access data on your private network through the application.

Considerations for Custom Templates

The custom template lets you make custom applications available as Azure RemoteApps. The template creation process is relatively straightforward, but there are a number of considerations you must take into account.

A template is really nothing more than a Windows virtual hard disk (VHD) that contains a sysprepped Windows Server 2012 R2 installation and the application you want to make available to your users. Although this might sound really simple (and it is), there are a number of requirements to which you must adhere. Otherwise, you risk the template being rejected.

By far the most critical of these requirements involves the structure of the VHD itself. The VHD must be a VHD-format disk. Microsoft Azure doesn't support VHDX-based disks. The disk can be dynamically expanding or it can be of a fixed size, but Microsoft recommends using a dynamically expanding VHD. The virtual disk

can contain multiple volumes, but it must use MBR partitions (Azure RemoteApp doesn't currently support GUID partitions) and the VHD must contain only a single copy of Windows Server.

There are also some very rigid requirements regarding the size of the VHD. It must be at least 40GB, but can't exceed 127GB. The size you should use depends on your application's requirements, plus the requirements of any data you plan to store alongside the application (such as in a separate partition within the VHD). It's a good idea to overestimate the space requirements because adjusting the size later on is likely to be a disruptive process.

There are some very rigid requirements regarding the size of the VHD.

The most important thing you need to know about the VHD is its size must be an exact multiple of megabytes. For example, if you wanted to create a 40GB VHD, it would need to be 42,949,672,960 bytes in size, because that number is exactly 40,960MB. You can't just round the number to 42,000,000,000 because that number is not evenly divisible by 1,048,576 (1MB). If you get the size wrong, the upload process will fail.

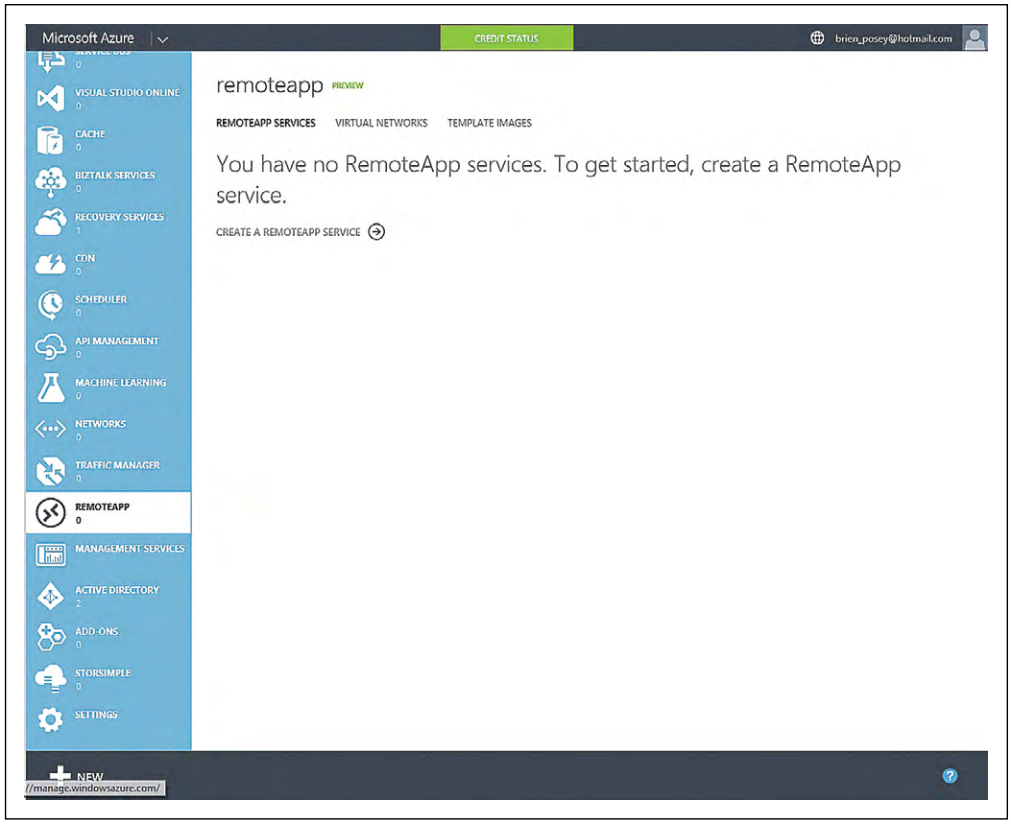

Figure 1. To configure an Azure RemoteApp environment, you begin by clicking on the Create a RemoteApp Service link.

There are a few other important considerations around the creation of a custom template. The template will need to contain the Remote Desktop Session Host role and the Encrypting File System will need to be disabled. The image also requires the Desktop Experience feature to be installed (you can't perform a server core deployment). As previously noted, the image will need to be sysprepped. The instructions are available at **bit.ly/1xFuUBu**.

Configuring RemoteApp

Whether you're performing a cloud or a hybrid deployment, the first step in the process is to create a RemoteApp service. This is a really simple process that requires you to go to the Azure Management Portal and then go to the RemoteApp page. From there, you would click on the Create a RemoteApp Service link, shown in **Figure 1**, page 10.

Whether you're performing a cloud or a hybrid deployment, the first step in the process is to create a RemoteApp service.

If you plan to perform a cloud-based deployment, you can use the Quick Create option (see **Figure 2**). It allows you to create a

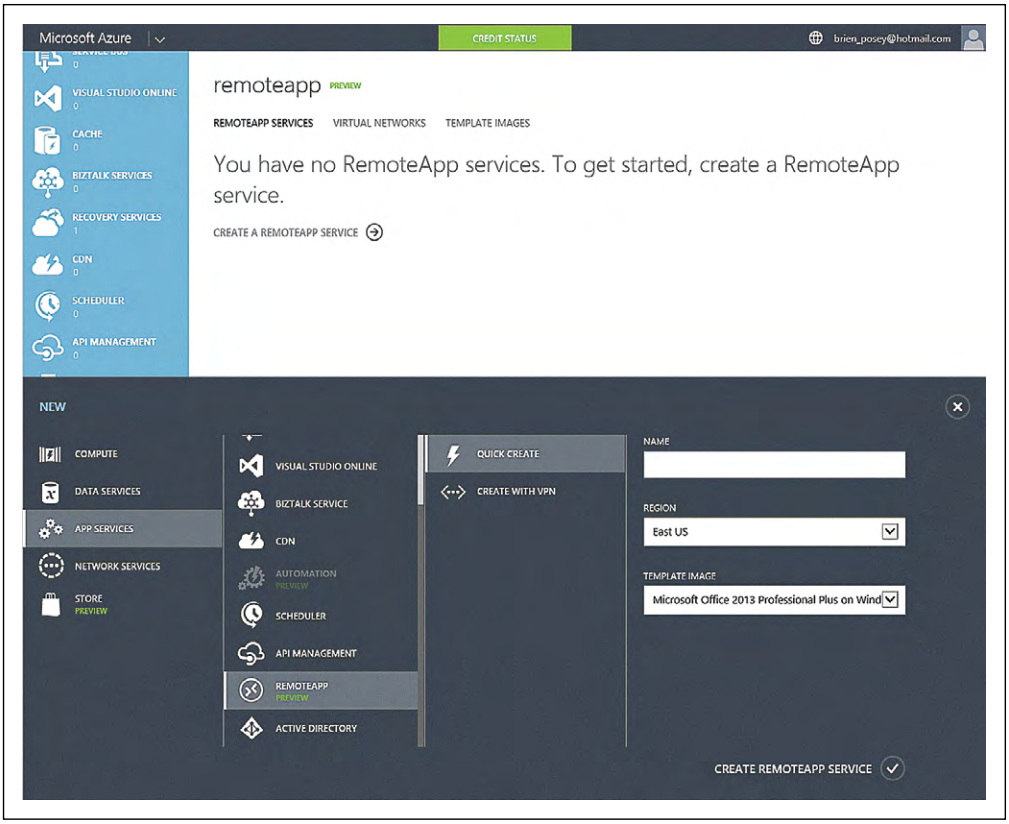

Figure 2. The Quick Create option includes a predefined template for Microsoft Office 2013 Professional Plus.

RemoteApp service that's based on the Microsoft Office 2013 Professional Plus template. If you're performing a hybrid deployment then you would use the Create with VPN option. Either method can take up to half an hour to provision the service.

Neither cloud-based nor hybrid deployments use a single-step configuration process. Both types of deployments require you to perform some additional tasks to get Azure RemoteApp up and running. Most of these tasks are the same for both types of deployments, although the order in which you must perform the tasks will vary depending on the type of implementation. Furthermore, a hybrid infrastructure requires you to perform some tasks not required (or in some cases are optional) for cloud deployments.

The virtual network's job is to bridge the gap between Azure and your private network.

Linking to a Template Image

It should come as no surprise that one of the tasks you must perform is to link the Azure RemoteApp service to a template image. If you're performing a cloud-based deployment, you perform this step in concert with the creation of the Azure RemoteApp service. The wizard will ask you which template you want to use. You can choose to use the built-in template that contains the Office 2013 applications or you can choose your custom template as an alternative. In the case of a hybrid deployment, you wouldn't provide the template image until a little bit later in the process.

Link to a Virtual Network

If you're performing a hybrid deployment then the step that must be performed prior to uploading your template image involves linking Azure to a virtual network. The virtual network's job is to bridge the gap between Azure and your private network. That way, your users will be able to access data that's stored on your private network while accessing an Azure RemoteApp.

In addition to creating a virtual network, you'll need to create a site-to-site VPN that can provide connectivity between Azure and your private network. Your private network will require a Windows Server 2012 (or 2012 R2) server running the Routing and Remote Access Service (RRAS) or a compatible VPN device (a list of compatible devices is available at: **bit.ly/1D6D4nG**).

In either case, you must provision your VPN with an externally accessible IPv4 address (you can find instructions for setting up the VPN at **bit.ly/149uAoh**.

Active Directory Synchronization

Another task you might need to complete is Active Directory synchronization, which lets you base access permissions to Azure RemoteApps on the same Active Directory user accounts and security groups established on your local network.

In the case of a cloud deployment, Active Directory synchronization is optional.

In the case of a cloud deployment, Active Directory synchronization is optional. There's nothing stopping you from using Active Directory objects to control access to Azure RemoteApps, but you don't have to use this approach. You can use local users and groups (within the custom template) as an access control mechanism instead.

In the case of a hybrid deployment, Active Directory synchronization is an absolute requirement. Remember, the advantage of performing a hybrid deployment is that you can provide users with access to a RemoteApp that's hosted on Azure, while also allowing them to access application data that's stored locally. The only way to achieve this level of access is to make use of Active Directory synchronization so the user's credentials remain valid for access control on the local network and on Azure.

Establishing directory synchronization between Azure and your private Active Directory forest involves a lot of work, described in Microsoft documentation at **bit.ly/1wOl3SW**.

Publish RemoteApp Programs

Whether you're configuring a hybrid or a cloud deployment, the last part of the process involves publishing RemoteApp programs and then granting user access.

The first step in making an Azure RemoteApp available to your users is to publish the app. You can do this by going to the Azure RemoteApp page and clicking Publish. You can choose to publish the RemoteApp through the template image's Start menu or you can specify the application's path within the template image.

Azure RemoteApp Resources

- Cloud deployment of Azure RemoteApp collections: **bit.ly/TMnvOi**
- Create hybrid Azure RemoteApp collections: **bit.ly/1hwSxyZ**
- Detailed price plan: **bit.ly/1D4Dg82**
- Compatible devices: **bit.ly/1D6D4nG**
- VPN setup instructions for connecting external devices: **bit.ly/149uA0h**
- Build custom template images for RemoteApp: **bit.ly/1xFuUBu**

Once the application is published via Azure RemoteApp, you must grant your users access to it. You can do this from the Quick Start page.

Once the application is published via Azure RemoteApp, you must grant your users access to it. You can do this from the Quick Start page by clicking on Configure Access. After doing so, simply specify the user or group to which you want to grant access.

Client Components

One last thing you need to know about Azure RemoteApp is that your users probably won't be able to access the remote application through a standard RDP client. Microsoft publishes a version of the client that's designed for use with Azure RemoteApp. You can download the new client from the Azure RemoteApp Quick Start page.

Recommendations and Considerations

Some might argue Azure RemoteApp isn't really anything new because Windows Server has included a RemoteApp feature for quite some time. The nice thing about the Azure RemoteApp feature, however, is that it allows applications to be hosted in the cloud so that users can access those apps without having to consume the organization's bandwidth (except for in the case of a hybrid deployment). Furthermore, it's now available as a service and the Azure infrastructure makes it easy to scale application delivery as your organization grows. VR

Brien M. Posey is a seven-time Microsoft MVP with more than two decades of IT experience. He's written thousands of articles and several dozen books on a wide variety of IT topics. Visit his Web site at brienposey.com.

Disaster Recovery as a Service

Microsoft has made a big bet this year that DRaaS will be a killer application for its Cloud OS and Azure service. Numerous players are making similar gambits.

By Jeffrey Schwartz

In today's new age of "always-on" business, prolonged downtime or even brief outages are no longer acceptable. Whether it's at a global enterprise with thousands of employees, a 200-person organization or even a small office, all are expected to have their core information systems up and running all the time. Providing the ability to recover from downtime—scheduled or unplanned—is becoming easier and more affordable thanks to a growing number of emerging enterprise-grade cloud-based Disaster Recovery-as-a-Service (DRaaS) options.

Over the past year, the sheer number and scope of options has started to amass.

Many such DRaaS offerings, where organizations replicate snapshots of their data, system settings and applications to either a local or major cloud provider or dedicated hosting operator, have been around for some time from specialists such as SunGard or Verizon Communications and a variety of high-end solutions. But over the past year, the sheer number and scope of options has started to amass, and many more are building out cloud-based disaster recovery service operations with varying types of capabilities, architectures and costs.

In 2014 Microsoft made a huge splash launching an extensive new portfolio of cloud-based disaster recovery options, recognizing and emphasizing disaster recovery as a key driver for its hybrid and Infrastructure-as-a-Service (IaaS) offerings. The Microsoft disaster recovery thrust came on the heels of last year's release of Windows Server 2012 R2, which included the second version of Hyper-V Replica, providing point-to-point replication of Hyper-V virtual machines (VMs) via either a LAN or WAN connection (see "Hyper-V Replica for Disaster Recovery," p. 22).

In 2014 Microsoft made a huge splash launching an extensive new portfolio of cloud-based disaster recovery options.

Building on that, Microsoft this year made it possible to use its Microsoft Azure cloud in lieu of a secondary datacenter for disaster recovery. At the core is Azure Site Recovery, which Microsoft announced in May at its TechEd conference in Houston. Azure Site Recovery, which became generally available in October (see "First Look: Azure Site Recovery" on p. 34), is a service enabling the replication of VMs between two datacenters or from an organization's site to Azure datacenters. The service, which unlike Hyper-V Replica also supports VMware VMs and Linux servers, offers automated protection of VMs, which Microsoft backs with a service-level agreement.

The July acquisition of InMage gave Microsoft an on-premises appliance that offers real-time data capture on a continuous basis, which simultaneously performs local backups or remote replication via a single data stream. Microsoft is licensing Azure Site Recovery with the Scout technology on a per-virtual or per-physical instance basis.

At its recent TechEd conference in Barcelona, Microsoft introduced some additional capabilities including support for its Azure Automation, a runbook automation service now in preview that lets customers automate Azure Site Recovery through planned support for Windows PowerShell scripting. Microsoft sees DRaaS as a key steppingstone to offering IaaS. DRaaS especially has appeal to customers because it delivers what for many is a much-needed capability that can be out of reach, and certainly far less expensive for those using secondary datacenters or operating co-location facilities.

Extending Azure Disaster Recovery via Cloud OS

In the same way Office 365 might not suffice for all Exchange and SharePoint users, Microsoft realizes its own Azure service won't cut it for all seeking DRaaS, either, especially those with data

"Where we see people using Disaster Recovery as a Service from us are those who need a recovery time objective or a recovery point objective that's measured in minutes, rather than hours or days."

Monty Blight, Vice President, Peak 10 Inc.

sovereignty requirements. As such, many Microsoft managed services and cloud hosting partners are delivering DRaaS, while some are building up to that point. One such partner is Peak 10 Inc., a hosting and managed services provider with 25 datacenters in 10 markets, whose clients include Chiquita Brands, Magazines.com, Meineke, Pergo and the PGA of America. A longtime Microsoft partner, the Charlotte, N.C.-based company has seen significant growth in its DRaaS offering this year, says Monty Blight, a vice president at Peak 10.

"Where we see people using Disaster Recovery as a Service from us are those who need a recovery time objective [RTO] or a recovery point objective [RPO] that's measured in minutes, rather than hours or days," Blight says. "The big key component of that is the replication piece between the two."

Of course, not all organizations need, or can justify the cost, of the RTOs and RPOs of mere minutes and most commonly, it depends on the application and business function. "Where this allows the customer to have private cloud, as well as data backup to a second site, it also means you look at integrating that file-level restore, which we do for Microsoft servers all day long," Blight says. "So it integrates in with our existing backup and restore and [DRaaS] option, but also specifically on the Cloud OS it gives them a second site to ensure their data is there."

DRaaS Considerations

Indeed, while Microsoft and all of its rivals including Amazon Web Services Inc. (AWS) and VMware Inc., as well as thousands of local and regional managed services providers and hosting operators have similar designs on DRaaS. Whether or not you use all or part of the Microsoft DRaaS or Cloud OS stack, customers have no shortage of options. At the same time, not all are created equal and IT architects need to consider numerous scenarios, requirements and capabilities, warns Enterprise Strategy Group analyst Jason Buffington.

"Providers and IT decision makers need to beware of over promising on what disaster recovery means," Buffington says. "Real disaster recovery—even in the cloud—still means I've got to have orchestration, I've got to build a sandbox so I can do testing, it means I've got

to be able to define policies, so the right [VMs] come up in the right order. Based on priority and based on dependencies of those VMs, there's a lot more to it than, 'I'm going to make a copy of my VMs and put them someplace else and when something bad happens I'm going to turn them on.'"

Among those large enterprises using Hyper-V Relica to connect to secondary datacenters and the Azure cloud is ABM Industries Inc., the largest United States provider of facility management services ranging from HVAC repair, security and landscape maintenance with 100,000 employees and nearly $5 billion in annual revenues. Andre Garcia, ABM's assistant vice president of global technology, referred to the disaster recovery scenario during a panel session on Hyper-V migration at the August TechMentor Redmond conference, which, like Redmond magazine is produced by 1105 Media Inc.

One shortcoming of Hyper-V Replica is that it's synchronous.

"Hyper-V Replica is just a feature of Hyper-V that's on by default—you just have to right-click and tell VMM [Virtual Machine Manager] what the target is for that source," Garcia said during the panel discussion. "It's a phenomenal capability," added panel participant Matt McSpirit, a Microsoft technical product manager focused on Hyper-V. "It has enabled organizations to replicate changes up to every five minutes, between Site A and B. It's well-received, with a PowerShell layer for automating it."

One shortcoming of Hyper-V Replica is that it's synchronous. Microsoft has said it's developing an answer to that with a new tool called Storage Replica (see more **bit.ly/1tDpmyH**).

Alternative Services Emerge

Yet numerous other software and services providers—many point out they're Microsoft partners—say organizations need better automation and replications than Hyper-V replica can offer. Many of them point to better recovery times, links to multiple clouds and faster continuous data protection (CDP), compression and data deduplication algorithms. Most dismiss Microsoft Hyper-V Replica as a suitable base-level replication mechanism for creating Windows Server Hyper-V clusters, but not sufficient for providing complete DRaaS.

There's no shortage of those who have stepped up their DRaaS offerings and market presence this year. Among them in various stated of delivering new DRaaS capabilities are Acronis International GmbH, ArcServe (spun off from CA Technologies), Asigra Inc., Axcient Inc., Dell Inc. (AppAssure), Hewlett-Packard Co. (via its Helion cloud platform), Nasuni Corp., Symantec Corp., Vision Solutions Inc., Unitrends, Veeam Software and Zerto, while CommVault is said to have new DRaaS capabilities in the works.

"Hyper-V Replica definitely has a place for the lower tier workloads," says Tim Laplante, a senior product director at Vision Solutions Inc., supplier of DoubleTake. "But where you need true high availability or you need to replicate it to something other than Hyper-V, you're going to need a solution like ours, where you need the real time and the flexibility from a target perspective."

"Hyper-V Replica definitely has a place for the lower tier workloads,"

— Tim Laplante, a senior product director at Vision Solutions Inc.

Laplante points to Peak 10 as a provider that subscribes to that model. Peak 10's Blight says while using Hyper-V Replica is suitable in certain scenarios is suitable, in others he sees the need for third-party solutions, notably Double Take and Zerto.

"The customer who needs DoubleTake requires real-time replication," Blight says. In cases where CDP is necessary, Peak 10 has also been working with Zerto, whose namesake software has long-offered that capability for VMware environments and last month gained Hyper-V support.

Many providers of backup software are making big pushes into DRaaS. Veeam, the rapidly growing provider of VM backup and disaster recovery software for midsize organizations, in October kicked off a major push into DRaaS, adding a component to its newly branded suite called Veeam Backup and Replication v8. A key new component in its new release, Cloud Connect, offers an interface that lets users search a network of partner cloud providers and MSPs. The initial Cloud Connect supports just backup and recovery. Next year providers will also be able to deliver DRaaS using Cloud Connect.

"We believe that next year will be the year where disaster in the cloud will start to become mainstream," says Veeam CEO Ratmir Timashev, "and we will be one of the driving forces for that, because

we have a better license base and we provide this very easy out-of-the-box experience for end customers and for our service providers."

The MSP Azure Connection

Veeam is also enabling its MSP partners to use the back-end services of Azure. The company has made Cloud Connect available in the new Azure Marketplace. "Veeam cloud providers who want to offer Veeam Cloud Connect [can] leverage Azure to provide the underlying core infrastructure—network, compute and storage in the form of VMs,"says Rick Vanover, a Veeam product strategy specialist. Selecting the Veeam Cloud Connect option in the [Azure] Marketplace will let that Veeam partner run the Cloud Connect infrastructure in Azure."

Not all DRaaS providers see the benefits of using a larger cloud provider.

Unlike Veeam, Unitrends operates its own cloud and argues it offers higher service levels than what's available by larger cloud services like Amazon EC2/S3 and Azure. In addition to integrating its on-premises appliance with its cloud, Unitrends offers its own DRaaS and touts a tool called Reliable DR, which offers governance and compliance auditing. The company says its DRaaS has grown 180 percent this year to hundreds of customers. "They have the advantage of our software to build out similar services that we have," says Ubo Guha, Unitrends vice president of product management. Unitrends is still considering whether to forge ties with Azure Amazon or another major cloud network.

Not all DRaaS providers see the benefits of using a larger cloud provider. "Public clouds are generally not purpose built, so they're good at many things, not great at any one application layer," says Justin Moore, CEO of Axcient, which provides a turnkey replication appliance and runs its own multi-petabyte cloud for DRaaS. "If you think of disaster recovery as a service, it's more of an application layer offering than it is an infrastructure."

The City of Williamsburg in Virginia is among those who have deployed a DRaaS solution using the Axcient service, where it backs up 10TB of data including its Novell GroupWise server, SQL Server databases and file systems, all running on 22 servers tied to VMware-based VMs. The replication is performed overnight, meaning in a worst-case scenario, the city's data would be 24 hours old. "We're

pretty small so that's a pretty good recovery time objective," says the city's IT manager Mark Barham. "I could knock it down to 30 minutes if I wanted to."

The Outdoor Group LLC, which supplies sporting goods gear—mainly high-end archery equipment—has started using the Veeam Cloud Connect tool through DR provider Offsite Data Sync to replicate its Exchange e-mail system, SQL Server databases, and various application servers. "If we lose that information we're basically starting over from scratch," says IT Director Jim Klossner.

Gartner Inc. analyst Pushan Rinnen warns customers that Backup as a Service shouldn't be confused with DRaaS.

TBG Partners, a landscape architecture firm uses Nasuni's replication service. With the Nasuni appliances, CTO Greg Nichols says his company can replicate large CAD files that could be gigabytes in size each. Nasuni offers customers a choice of AWS or Azure to host their backed-up data. Nichols says data is backed up more frequently for the firm's architects. "Having it backed up every five minutes is great for our users, because they literally don't lose anything," he says.

Buyer Beware

Gartner Inc. analyst Pushan Rinnen warns customers that Backup as a Service shouldn't be confused with DRaaS, even as many of the same companies offer both. "Disaster recovery involves not just the bits of the data, a copy of the storage part, but a lot of the business processes in the servers, applications and the consistency of the data," she says. "It's a lot more complex than backup."

If you're not using DRaaS yet, you're not alone. Many of these services are in their evolutionary state, Rinnen says. "We are definitely seeing more implementations of Disaster Recovery as a Service," she says. "But we're still very early at the beginning stage." **VR**

Jeffrey Schwartz is editor of Redmond.

Hyper-V Replica for Disaster Recovery

The replication feature Microsoft introduced in Windows Server 2012 provides business continuity. Though no substitute for failover clustering, it's an affordable option. **By Brien M. Posey**

Although many small and midsize businesses run their workloads on virtualized servers, they haven't been able to take advantage of the fault tolerant capabilities of virtualization such as failover clustering. The licensing and hardware costs and technical complexity involved in building a clustered Hyper-V deployment tend to put failover clustering out of reach for smaller organizations. Fortunately, Hyper-V offers a replica feature that's well suited for helping smaller organizations improve their disaster readiness.

Although there are similarities between replication and failover clustering, failover clustering is the preferred method for protecting your virtual machines (VMs).

Appropriately called Hyper-V Replica, Microsoft introduced it with Windows Server 2012 R2 and upgraded it in the subsequent release. While it provides replication designed to ensure business continuity, Hyper-V Replica is not a substitute for failover clustering. If your organization has the budget to build a clustered Hyper-V deployment, you should definitely do so. Although there are similarities between replication and failover clustering, failover clustering is the preferred method for protecting your virtual machines (VMs).

Of course, that isn't to say the Hyper-V Replica feature is inadequate—quite the contrary. I use Hyper-V Replica to protect my own VMs. I recommend the use of failover clustering whenever possible because a failover cluster's job is to make sure critical workloads never go offline. Replication won't guarantee that your VMs stay running in the event of a disaster, but it will give you at least one "spare copy" of your VMs, which you can launch at a moment's notice.

The Hyper-V Replica feature is based on the idea of asynchronously replicating a virtual disk from a primary site to a replica site. Although Microsoft refers to the source and target in terms of sites, it's important not to confuse the concept with Active Directory sites or geographic sites. In my own organization, for instance, my primary and replica "sites" exist within the same rack and on the same network segment.

The replication process occurs at the virtual hard disk level on an asynchronous basis. Once the initial copy process has been completed, replication occurs on a scheduled basis. In the version of Hyper-V Replica delivered with Windows Server 2012 R2, it's now possible for administrators to adjust the replication frequency. Replication can be scheduled to occur at 30-second, five-minute or 15-minute intervals. Intervals of 30 seconds do the best job of keeping the replica up-to-date, but aren't always appropriate. If the primary server is heavily utilized or if there's a slow link between the primary and the replica servers, then a longer duration replication frequency might work better.

The Hyper-V Replica feature is based on the idea of asynchronously replicating a virtual disk from a primary site to a replica site.

Another improvement is the addition of Hyper-V Extended Replication. Extended Replication allows for the creation of a secondary replica. The most common use for this feature involves placing one replica within the local datacenter (so that it's easily accessible) and placing the secondary replica in a remote location (so that it's protected against datacenter-level disasters).

Planning Considerations

First, the server that will store your replica doesn't need to be 100 percent identical to your source server, but it needs to be capable of hosting your VMs if necessary. As such, you'll need to make sure the replica server has adequate hardware resources to ensure a good UX in the event that it ever has to be put into use.

Another important consideration is the authentication type that's used by the replication process. By default the replication process is based around the use of Kerberos and the HTTP protocol. If you require encryption, however, you might be better off using certificate-based authentication, which is based on HTTPS.

You'll also need to consider the initial synchronization process.

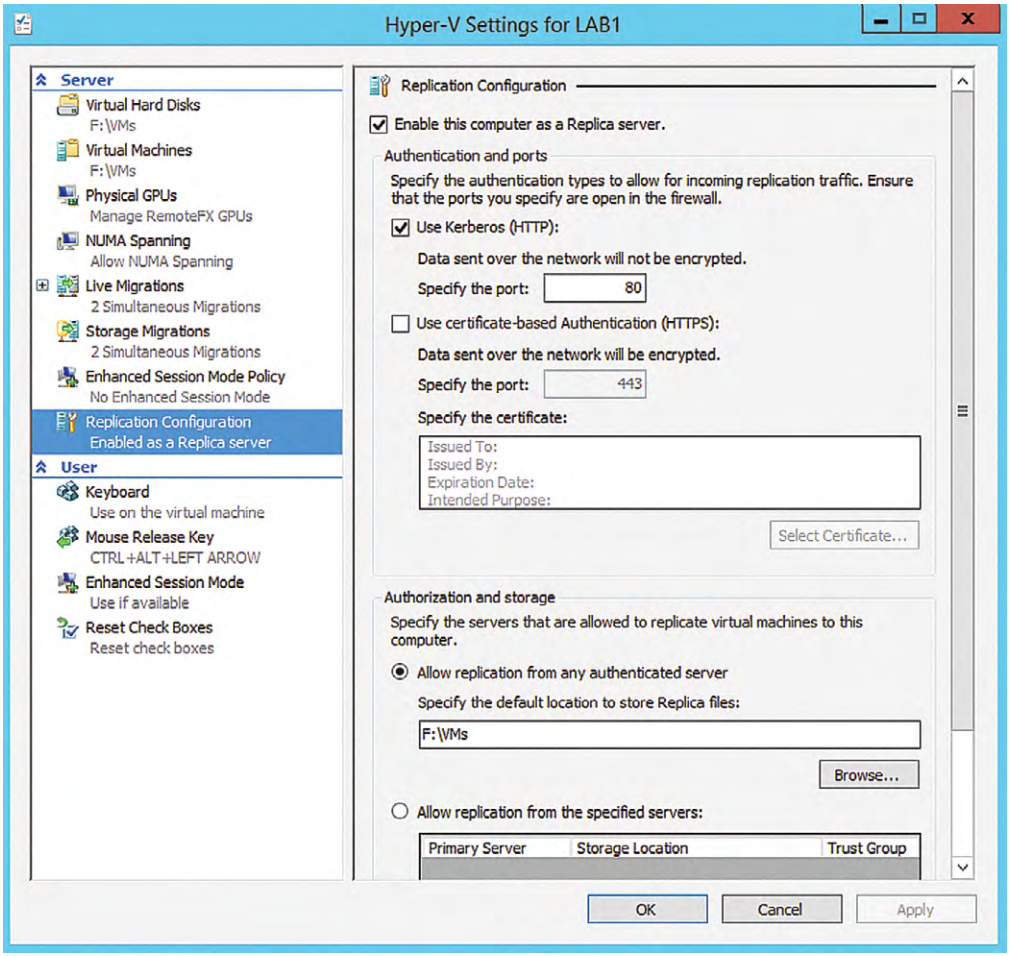

Figure 1. Select the Replica Configuration container.

Normally, you should be able to perform the initial synchronization process across the network.

Normally, you should be able to perform the initial synchronization process across the network. In the case of excessively large VMs, you're often better off using removable media to create the initial replica.

In addition, you'll need to consider other aspects of the replication process, such as the most appropriate frequency and whether you'll require extended replication.

Enabling Hyper-V Replication

The process of enabling Hyper-V replication involves performing various tasks on both the source server (the primary site) and the destination server (the replica site). Incidentally, the focus here is on Hyper-V replication in terms of a source server and a destination server, but you can replicate a VM to or from a cluster, or even between clusters so long as the Replication Broker is installed.

The destination server must be configured first. Open the Hyper-V Manager, select the listing for the destination host server and then

click on the Hyper-V Settings link, found in the Actions pane. When the Host Server Settings dialog box opens, select the Replica Configuration container (see **Figure 1**, p. 24).

Next, select the Enable this Computer as a Replica Server checkbox. You'll also need to select the type of authentication you want to use: allowing replication from any authorized server or specifying a list of Hyper-V servers from which you want to allow replication. Finally, click the Browse button and specify the location where you want to store the VMs. Click OK to complete the process. You might receive a warning message saying you need to configure your firewall to allow replication traffic.

When prompted to enter an authentication type, make sure to specify the same authentication method you used on the destination server and click Next.

The next thing you need to do is to open the Hyper-V Manager on the source server. Next, right click on the VM you want to replicate and select the Enable Replication command from the shortcut menu. You can replicate multiple VMs, but you'll need to enable replication separately for each VM.

At this point, Windows will launch the Enable Replication Wizard. Click Next to bypass the wizard's Welcome screen and you'll see a screen prompting you to enter the name of the replica server. Enter your destination server's name and click Next. When prompted to enter an authentication type, make sure to specify the same authentication method you used on the destination server and click Next. You'll be asked if you want to compress the data sent across the network. Compression reduces bandwidth consumption, but slightly increases CPU utilization. It's usually a good idea to use compression. Make your selection and click Next.

The next screen you'll see asks you to specify the virtual hard disks you want to replicate. Remember, replication works on a per-virtual hard disk (not a per-VM) basis. Click Next and you'll be asked to specify your replication frequency. After doing so, click Next.

The following screen asks you to choose the number of recovery points you want to store for the VM. Creating recovery points allows you to revert the replica to an earlier point in time. Windows Server 2012 R2 allows up to 24 hours' worth of recovery points to be maintained (the previous limit was 15 hours). It's worth noting that the replica's storage requirements increase as you add recovery points.

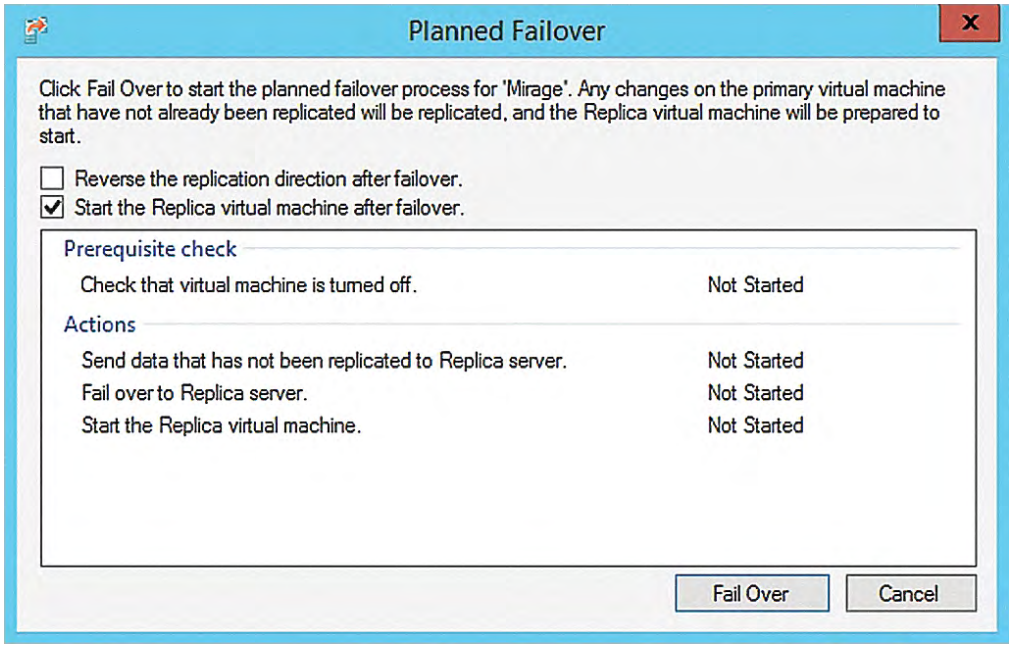

Figure 2. The Planned Failover dialog box.

Replicas exist for disaster recovery purposes.

Click Next and you'll be prompted to select the method you want to use for the initial synchronization process. After doing so, click Next. Assuming you're synchronizing across the network, you'll be asked when you'd like the replication process to begin. Make your selection and click Next. You should now see a summary screen displaying the replication options you've chosen. Take a moment to make sure everything is correct and click Finish. When you do, the VM Status should change to Initial Replication.

Replica Failover

As previously noted, replicas exist for disaster recovery purposes. As such, you can perform a planned failover or an unplanned failover. You can also perform a test failover.

A planned failover is useful in situations in which you need to take the primary host offline for maintenance. To do a planned failover, however, you need to first power down the VMs being replicated.

To perform a planned failover, right-click on the VM and select the Replication | Planned Failover commands from the shortcut menu. You'll see the dialog box in **Figure 2**. You can complete the failover by simply clicking on the Fail Over button. However, it's usually a good idea to select the Reverse the Replication Direction After Failover checkbox first. This checkbox causes the source VM to become the replica and the replica to become the primary.

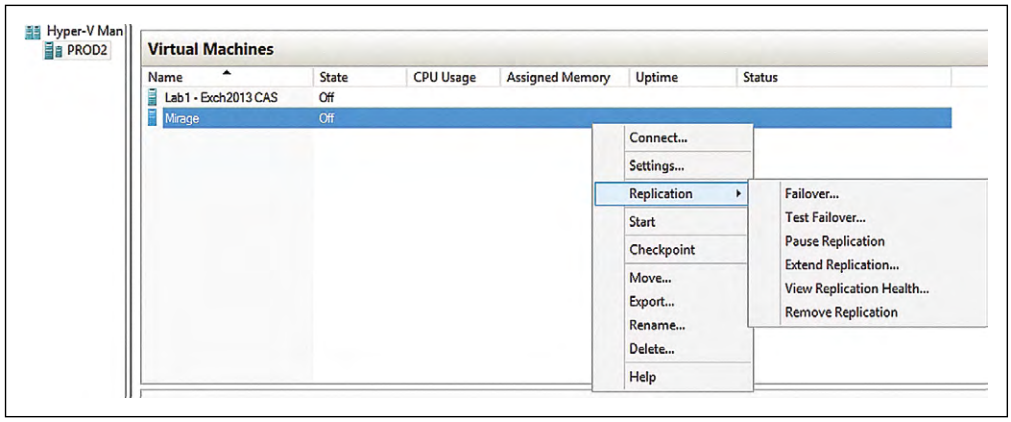

Figure 3. The Replication | Failover commands from the shortcut menu when right-clicking the VM replica to perform an unplanned failover.

You can safely perform a planned failover at any time.

You can safely perform a planned failover at any time. An unplanned failover should only be performed in the event that your primary VM has suffered a catastrophic failure. The reason for this is that an unplanned failover does not perform a synchronization as part of the failover process. Consequently, any data not already synchronized will be lost. The amount of data lost depends on the length of your replication cycle and the volume of data that was added to the primary VM since the last successful replication cycle.

To perform an unplanned failover, open the Hyper-V Manager on the server that contains your VM replica. Right-click on the replica and select the Replication | Failover commands from the shortcut menu (see **Figure 3**). Next, choose the recovery point that you want to use for the failover and then click the Failover button.

It's a good idea to perform a test failover. A test failover doesn't actually result in a failover. Instead, the process creates a brand-new test VM. This test VM lacks network connectivity, so it can be safely powered on and tested. There's a VM named Mirage-Test (see **Figure 4**), which is a test VM.

You can perform a test failover by going to the replica server, right-clicking on the VM, and selecting the Replication | Test Failover commands from the shortcut menu. Upon doing so, you'll be asked to select the recovery point you want to test. Make your selection and click the Test Failover button.

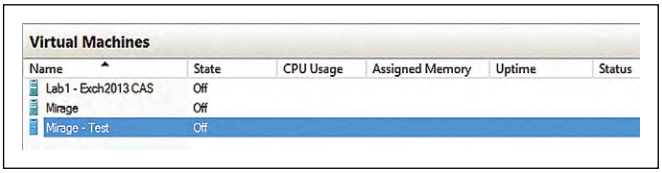

Figure 4. The test virtual machine.

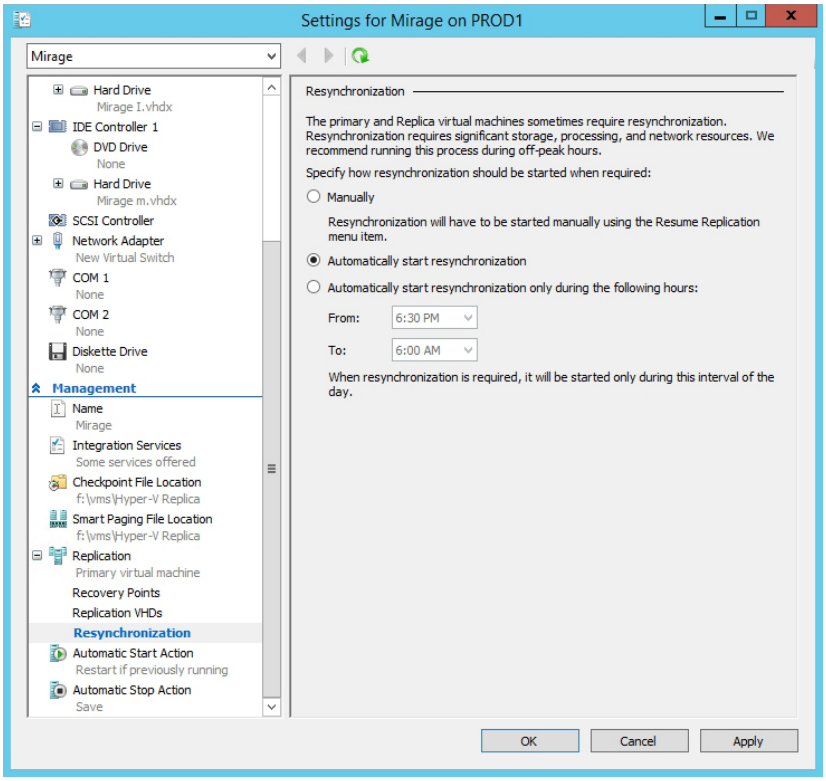

Figure 5. It's a good idea to enable automatic resynchronization.

If you're going to use the replication feature, I strongly recommend enabling the automatic resynchronization of replicas.

When you're done with your tests, right-click on the destination VM (not the test VM) and select the Replication | Stop Test Failover commands from the shortcut menu. This will cause the test VM to be deleted and everything will be put back to normal.

Replica Resynchronization

If you're going to use the replication feature, I strongly recommend enabling the automatic resynchronization of replicas. Replicas occasionally fall out of sync, and the resynchronization feature can fix the problem whenever necessary. You can access this feature by right-clicking on your VM and selecting the Settings command from the shortcut menu. When the Settings dialog box appears, expand the Replication container to reveal the Resynchronization container. You can choose to manually resynchronize, automatically resynchronize or automatically resynchronize during a scheduled time.

The Hyper-V replica feature is relatively easy to use, but there are loads of features not covered here, which you should explore. VR

Brien M. Posey is a seven-time Microsoft MVP with more than two decades of IT experience. He's written thousands of articles and several dozen books on a wide variety of IT topics. Visit his Web site at brienposey.com.

7 DRaaS Platforms Gaining Speed

There's no shortage of software, hardware and cloud providers adding Disaster Recovery as a Service (DRaaS) if Hyper-V Replica isn't enough for your requirements.

By Jeffrey Schwartz

If Microsoft's Hyper-V Replica doesn't meet your service-level requirements, there's no shortage of providers of software, hardware and appliances that suppliers are making available for cloud-based Disaster Recovery as a Service (DRaaS). Many are offered as appliances, others as pure software and services solutions.

There's no shortage of providers of software, hardware and appliances that suppliers are making available for cloud-based Disaster Recovery as a Service (DRaaS).

Some suppliers run their own cloud services, others are in the process of enabling partner networks of local and regional managed services and hosting providers to deliver those services. A number now also offer the option to use both local services providers and large ones such as Amazon Web Services (AWS) and Microsoft Azure. Others are still looking into doing so. Here are seven providers that have recently updated their offerings:

DRaaS Coming to Veeam Availability Suite in 2015 via Cloud Connect

The newly released Veeam Software Data Availability Suite v8 looks to enable customers who have used its virtual machine-focused backup and recovery software to implement disaster and recovery capabilities via secondary datacenters or using a cloud services provider. CEO Ratmir Timashev says that Veeam is on pace to post $500 million in booked revenue (non GAAP) this year and is aiming to double that to $1 billion by 2018. To get there, Timashev sees the growing DRaaS business as a key catalyst of that growth.

Timashev says Veeam can reach those fast-growth goals without deviating from its core mission of protecting virtual datacenters. The new Data Availability Suite v8 incorporates the company's new Cloud

Connect interface that will let customers choose from a growing network of partners that are building cloud-based and hosted backup and disaster recovery services.

Released last month, the Cloud Connect component initially only supports backup and recovery with DRaaS replication promised early next year, Timashev says. "From the user perspective, they are just going to see in the interface, 'Do you want to also backup up to cloud?,' and then they can select, 'Yes,' and then they can go directly to our Web site for the services provider they want to use. We have a simple registration and certification process for them to become a services provider who is using the Cloud Connect. So customers will be able to select in different countries the services providers in their cities."

The new v8 suite offers a bevy of other features including what it calls "Explorers" that can now protect Microsoft Active Directory and SQL Server.

Because Veeam Cloud Connect just became available, the company has only formally announced a handful of providers offering the service. They Include Cirrity LLC, iLand, NewCloud Networks, OffisteDataSync and Phoenix NAP. Veeam says it aims to have 1,500 services providers available in the coming year.

The new v8 suite offers a bevy of other features including what it calls "Explorers" that can now protect Microsoft Active Directory and SQL Server, and provides extended support for Exchange Server and SharePoint. Also added is extended WAN acceleration introduced in the last release to cover replication and a feature called Backup IO, which adds intelligent load balancing.

Unitrends New Offering Links Appliances and Cloud Service

The new Unitrends DRaaS offering uses the company's own cloud network, which it believes offers higher service levels than larger cloud services providers such as AWS Inc., Microsoft and Google. Though the company hasn't ruled out partnering with such players or others in the future for certain capability, the DRaaS offering lets customers use its appliances to conduct on-site backups of servers and virtual machines (VMs) and utilize its continuous data replication technology for data, systems and applications to the company's No Limits Cloud service, which the company says offers 24x7 telephone services and the use of its newly acquired optional Reliable DR disaster recovery testing tool to meet compliance and governance requirements.

Either live VMs or physical servers are spun up in real time to the cloud, providing recovery of those systems in the event of unplanned downtime or a disaster. On-premises appliances range in configuration from 1TB to 97TB and the company also offers software-based virtual appliances for instant recovery of both physical and VMs.

"We take it one step further and provide what we call deep virtualization, meaning we can go into the application that sits on the virtual machine," says Ubo Guha, Unitrends vice president of product management. "There may be an application like Exchange or custom apps that need to have a lot more deeper management of the operating system, the application, and you might want to adjust things."

"We provide what we call deep virtualization, meaning we can go into the application that sits on the virtual machine,"

— Ubo Guha, vice president of product management at Unitrends.

Vision Solutions Adds DRaaS to DoubleTake

The new DoubleTake 7.1, released last month from Vision Solutions Inc., dons a number of improved migration and high-availability features, but also provides disaster recovery for Windows hybrid cloud environments. It's suited for DRaaS, thanks to a new metered usage feature available for cloud and managed services providers deploying the product.

DoubleTake 7.1 is also now fully API-enabled and designed with full server data replication and is container-based rather than volume-based. It supports the new Microsoft virtual hard drive format VHDX and its Volume Shadow Copy Service (VSS), says Tim Laplante, director of product strategy at Vision Solutions.

"This provides more granular level of control and gives you that near CDP [continuous data protection), which is nice because it gives you the best of both worlds," Laplante says. "If there's a disaster and you need to execute your DR plan, it gives you the option at that point to say, 'Do I need to go back to that exact point in time, or do I need to go back to 15 minutes ago because it was really just a virus or data corruption that happened, so I need to step back for a couple of minutes to the point that happened before then?"

Besides the metered usage, it's suited for DRaaS in that the DoubleTake 7.1 repository can replicate both physical machines and VMs on-premises to another datacenter, private cloud or public cloud. Likewise, recovery service can be anywhere in the physical, virtual

and cloud mix, as well. Administrators can specify discrete repository server targets, so customers know exactly where a specific system and data is, which should appeal to those who have sovereignty requirements. "It's not that your data is in multiple zones," Laplante says. "You know exactly where that data is when you need it for compliance purposes."

With the new disaster recovery feature in DoubleTake, LaPlante says Vision Solutions will step up working with services providers to offer DRaaS. "It's a huge piece of where we see our growth," he says.

Zerto Virtual Replication Now Supports Hyper-V

"We took real hard enterprise-class replications from storage and moved it into the hypervisor."

— Gil Levonai, president of marketing, Zerto

Zerto, a 4-year-old company with headquarters in Israel and the United States that provides disaster recovery and replication software, until now has a following among VMware Inc. shops. The company has recently entered the Hyper-V world. The Zerto Virtual Replication now supports replication of Hyper-V hypervisors to other Hyper-V targets, as well as to vSphere and vice versa.

In short, the company says its CDP-based replication tool is now hypervisor-agnostic. Gil Levonai, the company's president of marketing, says its software offers recovery point objectives (RPOs) of seconds, and said it can provide consistent recovery of multiple VM applications. It doesn't use snapshots, just CDP, automatically orchestrates disaster recovery processes ensuring the consistency of applications and data, and generates reports.

"We took real hard enterprise-class replications from storage and moved it into the hypervisor," Levonai says. "You don't have to worry about where the VM is and you don't care about where the data is. You can move it between storage. We are agnostic to storage because we are replicating virtual objects, which can be VMs or volumes."

Dell Combines Backup and DRaaS in New AppAssure Suite

Dell Inc. was one of the earliest players to offer DRaaS to enterprises and earlier this year said it has more than 1,000 managed services providers (MSPs) offering its AppAssure replication software. The latest release, AppAssure 5.4, offers multi-target and multi-hop replication, which the company claims makes it suited for multi-tier disaster recovery.

AppAssure 5.4 also lets customers set multiple data retention policies both for on-premises and off-site cloud and MSP facilities. Customers can customize replication schedules for each target, enabling them to throttle when needed and restrict speed in bandwidth-limited situations.

Dell is offering AppAssure as part of a new data protection that includes NetVault Backup and vRanger backup and recovery offerings. The company is also now offering a capacity-licensing model with a range from 1TB going as high as 250TB of data.

Acronis Enters DRaaS with nScale Deal

Acronis says users of its Hosted Backup as a Service offering will be able to use nScale to extend that into a cloud-based disaster recovery offering.

Known for its protection of Windows physical and virtual file server data protection wares, including specialty versions for SharePoint, Exchange, SQL Server and VMware environments, Acronis International GmbH in September jumped into the DRaaS mix with the acquisition of San Francisco-based nScaled.

Acronis says users of its Hosted Backup as a Service offering will be able to use nScale to extend that into a cloud-based disaster recovery offering. The company will enable its partners to offer the nSCaled DRaaS offering, which is designed to enable remote and local sites to failover via the cloud to ensure recovery within minutes of an outage.

Nasuni Adds Azure to DRaaS

Until recently Nasuni Corp. has relied on AWS as the cloud provider for its DRaaS offering, now the company has added the Microsoft Azure service as an option. Customers can now choose which provider they want their data replicated to, or if they prefer, can use both for contingency.

The latest version of its offering was released this summer. It includes the 6.0 release, which the company says adds file data virtulization that separates file data from storage hardware. It adds global file locking to utilize cloud storage architectures. With it is the new Nasuni Filer NF-100 appliance, the company says service is suited for providing recovery of blocks of data including CAD and BIM files. VR

Jeffrey Schwartz is editor of Redmond magazine.

First Look: Microsoft Azure Site Recovery

There's no shortage of software, hardware and cloud providers adding Disaster Recovery as a Service (DRaaS) if Hyper-V Replica isn't enough for your requirements.

By Jeffrey Schwartz

The best way to truly protect your data is to have at least three copies of it. First, there's the original copy—the live data, of course. Next, you need a backup copy of the data that you can quickly and easily restore. The third copy is the alternate backup that resides outside your datacenter. Once upon a time you could fulfill these requirements by writing a nightly backup to redundant tapes and keep one tape on-site and ship the copy off-site for safe keeping.

Organizations have come to expect near-real-time data protection.

This tried-and-true backup technique is now outdated. Nightly backups have largely become inadequate. Organizations have come to expect near-real-time data protection. In the scramble to provide top-notch protection in the virtual datacenter, a number of competing solutions have evolved. Even Microsoft provides several different ways of protecting Hyper-V virtual machines (VMs).

At first glance, one of Microsoft's solutions would seem to be ideal: Hyper-V Extended Replication. If you aren't familiar with Hyper-V Extended Replication, it's a feature that was introduced with Windows Server 2012 R2 that allows you to create two separate replicas of a VM. One of these replicas can reside in the local

datacenter, while the other can reside outside the datacenter. As such, the Hyper-V Extended Replication feature provides near-real-time protection, while also meeting the requirements of my three-copy rule. When you consider that Hyper-V replicas can be configured to provide point-in-time rollback capabilities, Hyper-V replicas appears to be an ideal solution.

There's just one problem with protecting your VMs using Hyper-V Extended Replication. The feature was designed for small and midsize businesses and simply doesn't scale well enough to make it a viable option for protecting large, enterprise-class organizations. So what's a company to do?

While native Hyper-V replication is designed to replicate individual VMs (or even individual virtual hard disks), Azure Site Recovery is focused on private cloud replication.

Enter Microsoft Azure Site Recovery—a new disaster recovery feature in Azure that can replicate Hyper-V VMs in a way that can provide better scalability.

While native Hyper-V replication is designed to replicate individual VMs (or even individual virtual hard disks), Azure Site Recovery is focused on private cloud replication. In other words, if you have a System Center Virtual Machine Manager private cloud, you can replicate your Hyper-V VMs to another private cloud that's running in another datacenter. As an alternative, you can replicate VMs to Azure.

Although enabling protection for VMs involves a little bit of work up front, the process is surprisingly straightforward. The key to making the process work is ensuring the certificates are configured correctly. The certificates are used to positively identify your Virtual Machine Manager server to Azure.

Creating a Self-Signed Certificate

In order to use Azure Site Recovery, you need to generate a certificate. A self-signed certificate will work fine. There are a few different ways of generating the necessary certificate, but Microsoft recommends using a tool found in the Windows SDK for Windows 8.1 called MakeCert.exe (**bit.ly/1DrOjTG**). The SDK has a lot of different components, but the only component you have to install is the Windows Software Development Kit.

After installing the MakeCert utility, open an elevated command-prompt window and navigate to C:\Program Files (x86)\

Windows Kits\8.1\Bin\x64 and run the following command:

```
makecert.exe -r -pe -n CN=AzureBackup -ss my
-sr localmachine -eku 1.3.6.1.5.5.7.3.2 -len 2048 -e
01/01/2016 AzureBackup.cer
```

Azure is very picky about the way you create the self-signed certificate. If you deviate from the command here, MakeCert may tell you that you've entered too many parameters, or you could end up creating a certificate that Azure won't accept. Both are common problems you want to avoid, so be sure to correctly type the command.

Azure is very picky about the way you create the self-signed certificate.

Importing the Certificate

Now that the self signed-certificate has been created, you need to import it into the computer on which Virtual Machine Manager is running. To do so, enter the Microsoft Management Console (MMC) command at the server's Run prompt. Then, choose the Add/Remove Snap-in command from the shortcut menu. When the list of snap-ins appears, choose the Certificates option and click Add. When prompted, make sure to choose the Computer Account option, and then click Next. After that, choose the Local Computer option and click Finish, followed by OK.

Right-click on the Personal container and select the All Tasks | Import commands from the shortcut menus. This will cause Windows to launch the Certificate Import Wizard. Click Next, and then browse to and select the certificate you created earlier. Now, complete the wizard. When you're prompted to specify the certificate store, be sure to put the certificate in the Personal store.

Exporting the Certificate

Now you need to export the certificate in PFX format. To do so, navigate through the Certificates console tree to Certificates (Local Computer) | Personal | Certificates. Right-click on the certificate and select the All Tasks | Export commands from the shortcut menus. This will cause Windows to launch the Certificate Export Wizard. Click Next and you'll be asked if you want to export the private key. Choose Yes and click Next. Make sure the wizard is set to export the certificate in PFX format and then click Next. On the following screen, you must enter and confirm a password that can be used to encrypt the private key. Click Next and you'll be prompted for a path

and filename to use for the exported certificate. Click Next, followed by Finish to complete the process.

Now you need to import the certificate on your Virtual Machine Manager servers. If you only have a single Virtual Machine Manager server and you already imported the certificate on that server, then you can skip this step. Otherwise, open the Certificates console on your Virtual Machine Manager server and import the PFX file you just created.

Create a Site Recovery Vault

When you run the executable file, Windows will display the Microsoft Azure Site Recovery Provider Setup wizard.

The next step in the process is to create a Site Recovery Vault. You'll need to log in to the Azure Management Portal. Now, click New and then click on Data Services | Recovery Services | Recovery Site Vault | Quick Create. You'll need to enter a name for the vault you're creating, and you must specify the region in which the vault is to be created, as shown in **Figure 1**, p. 38. Click Create Vault to complete the process.

Now that you've created the vault, it must be configured. Click on the Recovery Services tab and then click on the vault you just created. The first thing you'll need to specify is whether site recovery will occur between a Hyper-V site and Azure, or between two on-premises Hyper-V sites (see **Figure 2**, p. 38).

Next, click on the Manage Certificates link. When prompted, provide the certificate (the .CER file) that you created earlier. Once the certificate has been uploaded, click on the Get the Vault Key link. Be sure to make a note of the key.

Azure Site Recovery Provider

Now it's time to download the Azure Site Recovery Provider and install it on your Virtual Machine Manager servers. Select the Download Microsoft Azure Site Recovery Provider and Install it on the Virtual Machine Manager servers link. When prompted, save the file to a centrally accessible location. Now, shut down the Virtual Machine Manager service and then run the executable file on each of your Virtual Machine Manager servers.

When you run the executable file, Windows will display the Microsoft Azure Site Recovery Provider Setup wizard. Click Install to begin the installation process.

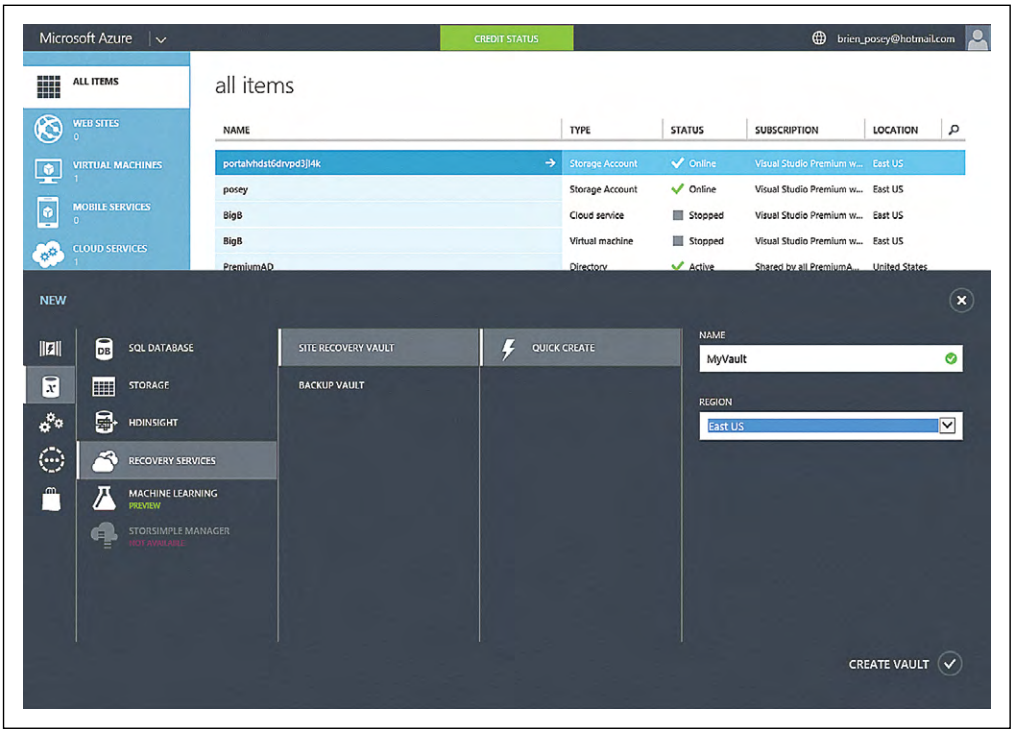

Figure 1. Creating a Site Recovery Vault in the Microsoft Azure Management Portal.

You'll need to select your certificate and then specify your vault and your vault key.

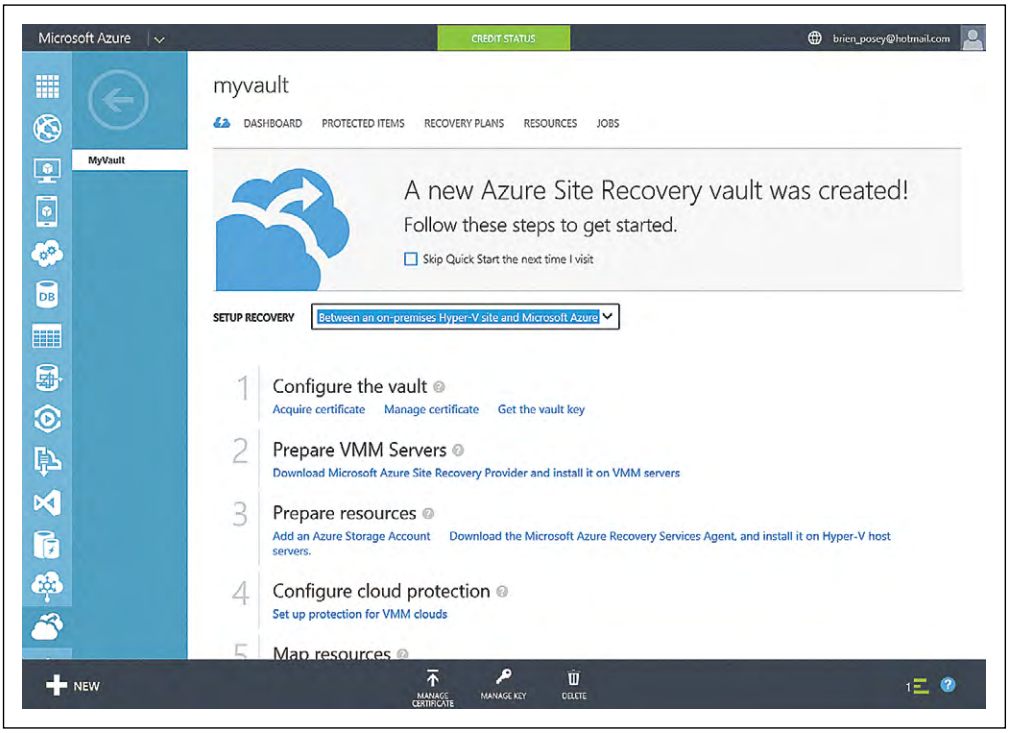

Figure 2. Specifying the type of site recovery.

After a few seconds, you should see a message telling you that Setup completed successfully. Click Next and you'll be prompted for your Internet connection settings. Click Next again and you'll be taken to the Vault Registration screen. You'll need to select your certificate and then specify your vault and your vault key (see **Figure 3**, p. 39).

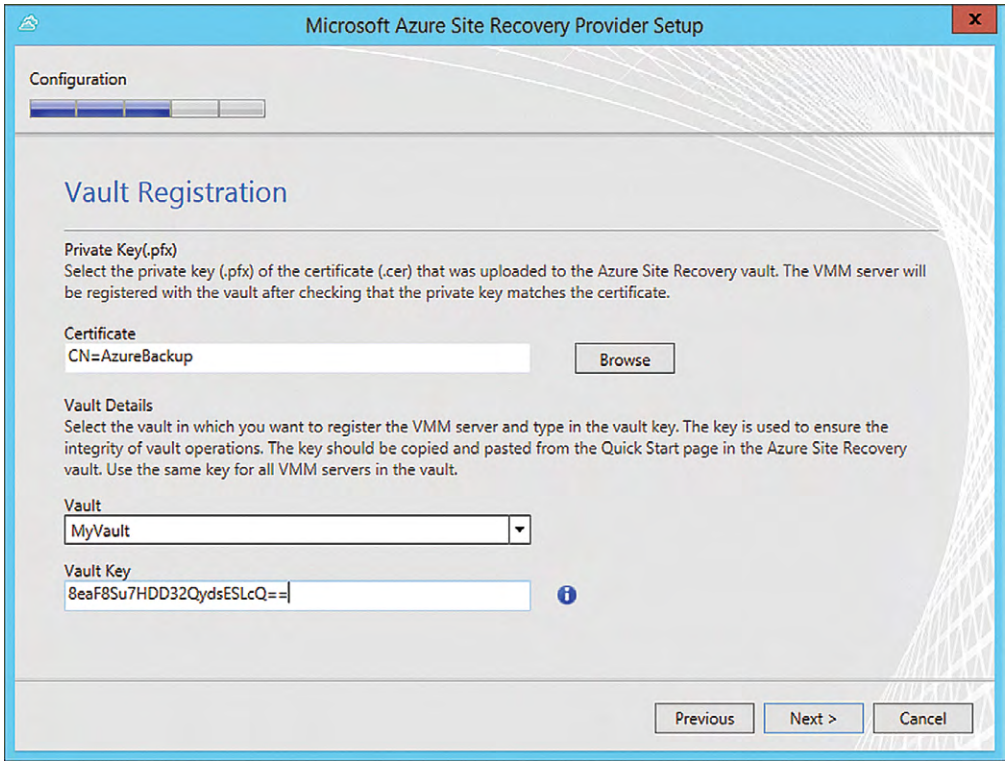

Figure 3. You must specify your certificate, vault and vault key.

If you allow this option, an encryption certificate will be automatically generated.

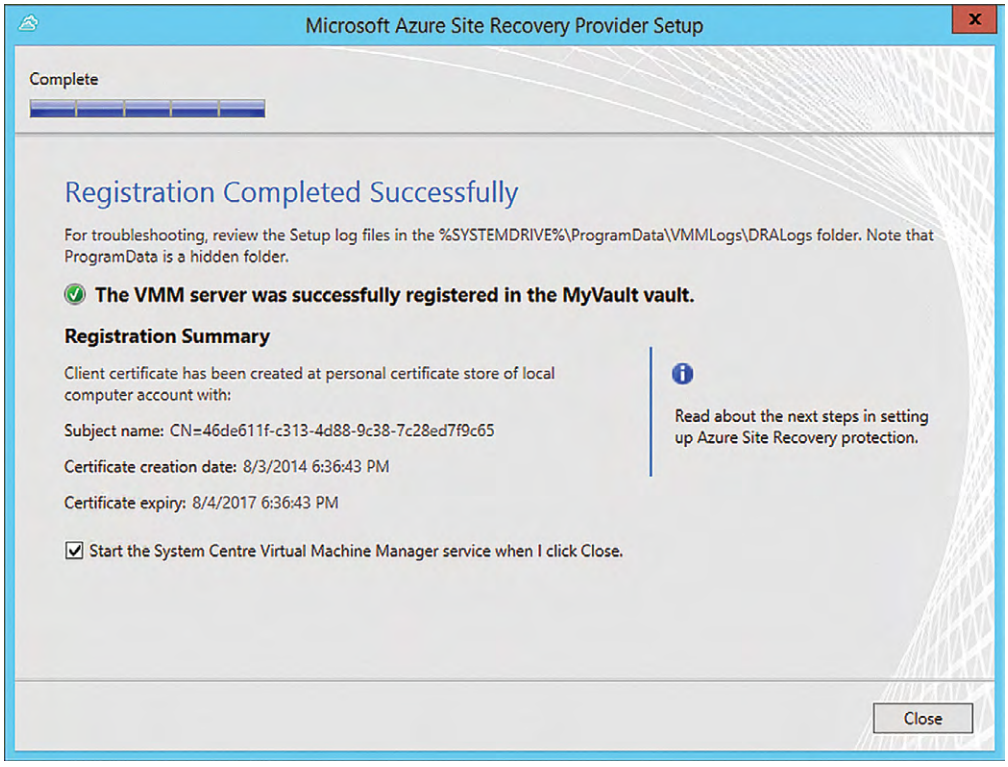

Figure 4. Confirmation of a successful registration.

Click Next and you'll see a prompt asking you if you want to encrypt replicated data. If you allow this option, an encryption certificate will be automatically generated. You'll have to provide this certificate whenever you fail over VMs. Click Next, followed by Register to

complete the process. When the process completes, you should see a message confirming you've successfully registered the Virtual Machine Manager server with your vault (see **Figure 4**, p. 39).

Protecting a Cloud

At this point, you've created a vault on Azure and associated the vault with Virtual Machine Manager. Usually, the next step in the process is to protect a private cloud. This will vary depending on your goals and whether you're replicating to Azure Storage or to a private cloud.

To protect a private cloud, you must right-click on the private cloud within the Virtual Machine Manager console.

To protect a private cloud, you must right-click on the private cloud within the Virtual Machine Manager console (assuming the cloud isn't already being synchronized) and select the Properties command from the shortcut menu. When the cloud's properties sheet appears, go to the General tab and select the Send Configuration Data About this Cloud to the Azure Hyper-V Recovery Manager checkbox, and click OK. After doing so, go into Azure, click on your vault, and select the Protected Items tab. You should see your cloud listed in the vault, as shown in **Figure 5**.

Click on the cloud and select the Configure Protection Settings link. You can now complete the process by answering questions about the protection you want. For instance, you're initially asked to select

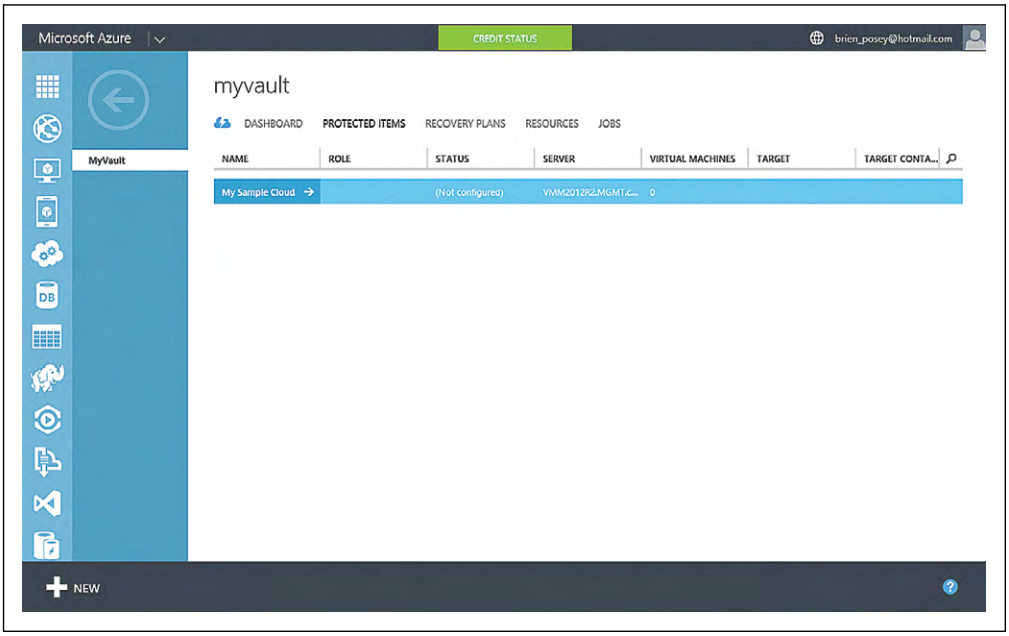

Figure 5. The private cloud now appears in the vault.

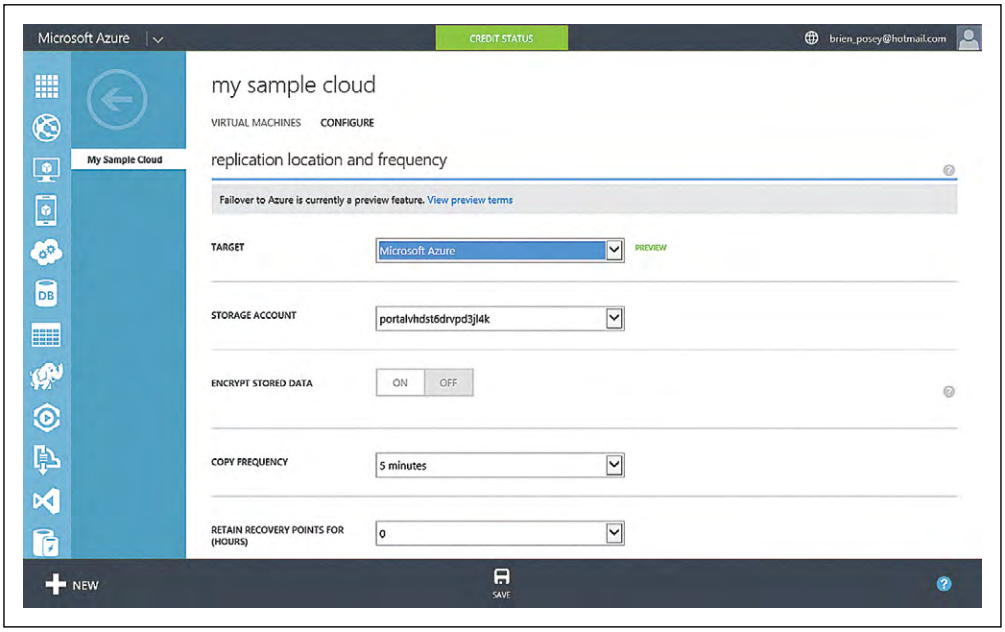

Figure 6. Configuring the replication parameters.

Replicating a Virtual Machine Manager to the Microsoft cloud using Azure Site Recovery is a fairly straightforward process.

a target. This is where you would specify whether you want to replicate the cloud to Virtual Machine Manager or to Azure. After making this selection, you can specify your storage account (if you're synchronizing to Azure), as well as your copy frequency, recovery point retention period, and the frequency of application consistent snapshots (see **Figure 6**). Click Save to save your changes.

And that's it! Replicating a Virtual Machine Manager to the Microsoft cloud using Azure Site Recovery is a fairly straightforward process. The key to making the process work is to generate the certificates correctly. VR

Brien M. Posey is a seven-time Microsoft MVP with more than two decades of IT experience. He's written thousands of articles and several dozen books on a wide variety of IT topics. Visit his Web site at brienposey.com.

Manage Mobile Devices and Policies in Active Directory

Introduced in Windows Server 2012 R2, Workplace Join lets otherwise incapable mobile devices participate in an Active Directory domain, but doesn't provide comprehensive security.

By Brien M. Posey

One of the major challenges facing organizations today is the proliferation of mobile devices. Mobile device use raises a number of concerns around issues such as security and privacy. Although there are a number of different solutions available from Microsoft and from third-party vendors for mobile device management, IT shops are increasingly finding that achieving the desired level of security requires them to adopt multiple solutions.

In order to understand the advantages and disadvantages of the available solutions, it's important to understand why mobile devices pose such a challenge in the first place. Mobile devices are different from desktop PCs in that they typically cannot be joined to an Active Directory domain. When a PC is joined to the Active Directory domain, a computer account is created as a means of authenticating and positively identifying that computer as it participates on the network. Furthermore, applying Group Policy settings can secure the computer's OS, and access control lists provide a mechanism for controlling access to network resources at the computer level.

Mobile devices are the complete opposite. They can't be joined to the Active Directory domain (at least not in the traditional sense) and, therefore, you can't apply Group Policy settings to a mobile device.

Previously, one of the best options for securing mobile devices that participate on a corporate network was the use of ActiveSync policies. For those who might

not be familiar with ActiveSync policies, they were first introduced in Exchange Server as a mechanism for pushing mail to mobile devices. Eventually, Microsoft extended the ActiveSync protocol in an effort to allow security settings to be applied to mobile devices.

The Trouble with ActiveSync

ActiveSync has become an industry standard for providing synchronization between Exchange and mobile devices. Every major device manufacturer supports the use of ActiveSync policies as a way of locking down their devices. However, there are a couple of issues that limit ActiveSync as a comprehensive mobile device security solution.

ActiveSync has become an industry standard for providing synchronization between Exchange and mobile devices.

The first issue is that mobile devices support ActiveSync in varying degrees. ActiveSync policies are made up of a collection of individual policy settings. Most of the mobile devices available today do not support every available policy setting. As such, administrators aren't

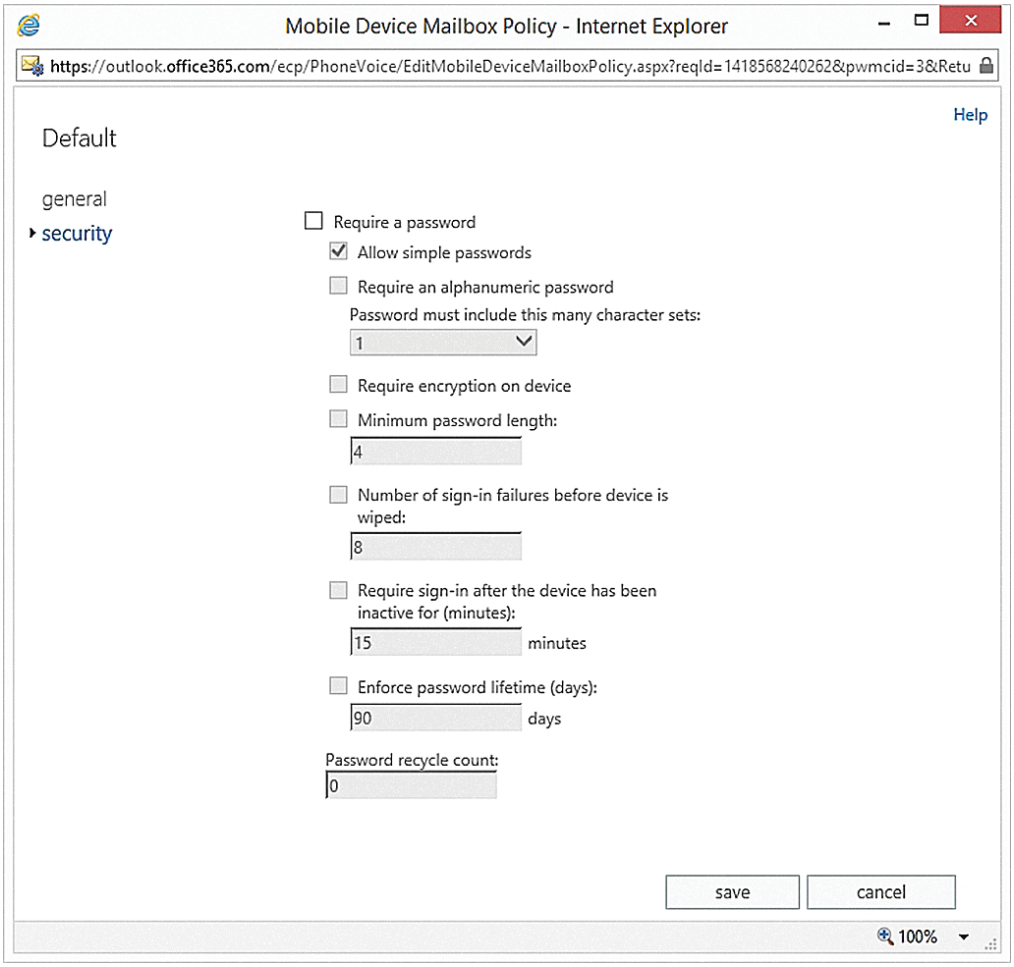

Figure 1. Office 365 exposes only the most basic ActiveSync policy settings.

assured that a policy setting will apply to every device that's in use on their network unless they restrict access to devices that do not fully support all policy settings (which includes most device types). The various ActiveSync policy settings and the devices that support them are published in a Wikipedia post at **bit.ly/1rHOUk3**.

Another issue with using ActiveSync for mobile device security is that it has become increasingly difficult to fully implement. Exchange Server 2013 and Microsoft Office 365 expose ActiveSync policy settings through mobile device mailbox policies. Microsoft makes it relatively easy to create mobile device mailbox policies, but the Exchange Server 2013 Administrative Center and the Office 365 Exchange Admin Center only expose the most basic policy settings (see **Figure 1**, p. 43).

Organizations that create comprehensive mobile device mailboxes usually find that those policies work really well for securing mobile devices.

These are the policy settings that are used to require a password and to set password-related attributes, such as the minimum password length. There are dozens of other ActiveSync settings available, but using them means delving into Windows PowerShell (or reverting to Exchange Server 2010).

Organizations that create comprehensive mobile device mailboxes usually find that those policies work really well for securing mobile devices. Even so, ActiveSync has its limits. For instance, ActiveSync doesn't allow for device-level access control for resources that exist outside of Exchange Server.

Pluses and Minuses for Workplace Join

One of the new features in Windows Server 2012 R2, called Workplace Join, is a way for allowing otherwise incapable mobile devices to participate in the Active Directory domain. Even though a device such as an iPad or a Windows RT tablet can't join to an Active Directory domain in the same way a Windows desktop PC can, the Workplace Join feature lets these types of devices participate in Active Directory in other ways.

The Workplace Join feature accomplishes three main things. First, it allows a device to be positively identified on the network. When a user enrolls a mobile device, a certificate is assigned to that device. This certificate is used as an identification mechanism. The device is also added to Active Directory and is listed in the Enrolled Devices container.

Second, the Workplace Join feature enables device-level access control. It's possible to base access to Web applications on whether the device has been enrolled in Active Directory. This isn't done through standard access control lists, but rather as a function of Active Directory Federation Services and a relaying-party trust. Third, Workplace Join provides single sign-on capabilities for certain network resources.

The Workplace Join feature enables device-level access control

Unfortunately, Workplace Join is inadequate by itself. Although it does provide some interesting capabilities, it also lacks some keybfeatures that are required for organizations to effectively manage mobile devices. Keep in mind Workplace Join is designed to simplify resource access. As such, it's not a true mobile device security feature. For instance, Workplace Join doesn't provide a collection of Group Policy settings that can be applied to mobile devices. Similarly, it doesn't provide the types of device security controls that are available through ActiveSync policies.

Workplace Join is also inadequate as an access control mechanism. While it's true Workplace Join can be used to grant or deny access to network resources based on device enrollment, there are significant limitations. You cannot, for instance, use an access control list to block devices that aren't enrolled. Most often Workplace Join is used to control access to browser-based apps.

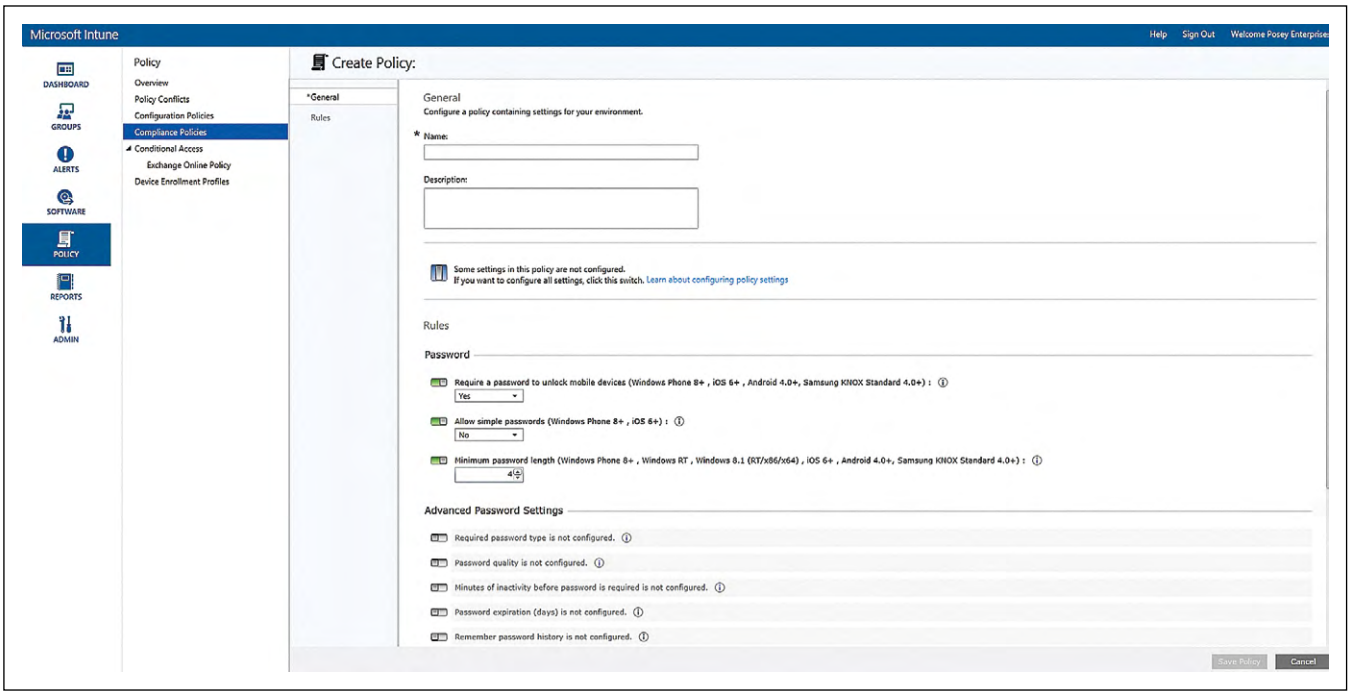

Figure 2. Microsoft Intune policies are similar to ActiveSync policies.

Moving to a Primary Solution

Workplace Join and ActiveSync policies both have their place, but the primary Microsoft solution for securing and managing mobile devices is Microsoft Intune (formerly known as Windows Intune). Microsoft Intune is a cloud-based service that can apply security controls to user devices. These controls are very similar to those that are available through ActiveSync (see **Figure 2**, page 45). Unlike raw ActiveSync, however, Intune is also able to create template-based policies that can be applied to various types of mobile devices (see **Figure 3**).

Workplace Join and ActiveSync policies both have their place, but the primary Microsoft solution for securing and managing mobile devices is Microsoft Intune.

Microsoft Intune is also able to manage applications on mobile devices. For example, the organization's approved applications can be made available to mobile devices. Administrators also have the ability to blacklist undesirable applications.

So what should you be using to manage your network—Microsoft Intune, Exchange Server, Workplace Join or perhaps something else? For the best overall management experience, you might want to use a combination of all three. All three technologies have both abilities and limitations. However, the three technologies complement one another rather nicely.

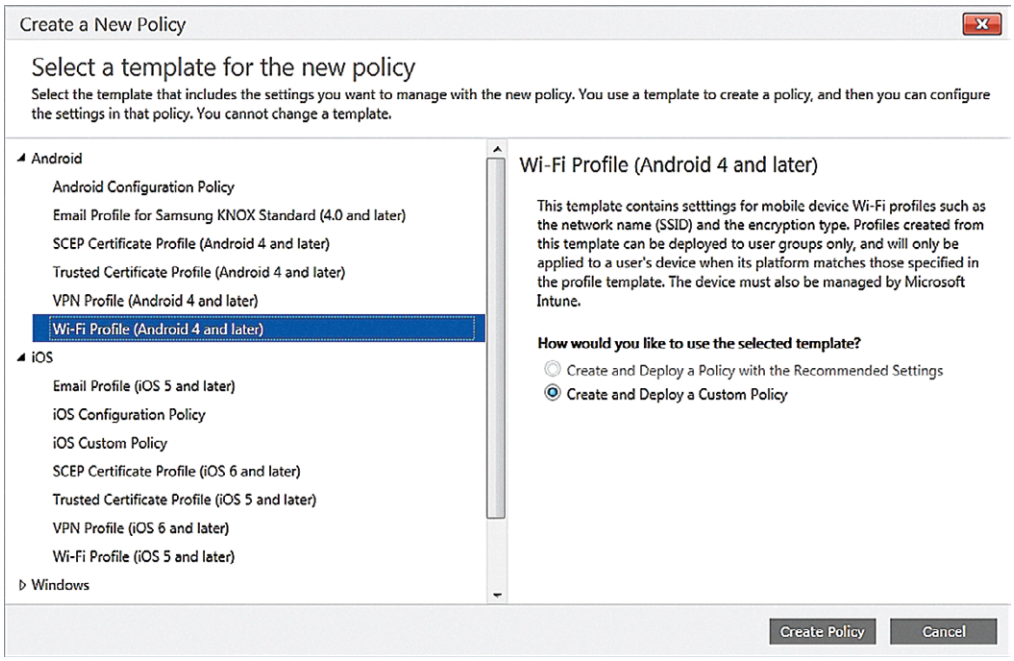

Figure 3. Microsoft Intune can create template-based policies for mobile devices.

On the surface, there seems to be a high degree of overlap between the three solutions. Microsoft Intune, for example, provides the ability to enforce password policy settings on mobile devices, but so does ActiveSync. Similarly, the Workplace Join feature provides access control to Web-based applications, but Microsoft Intune also provides application access control capabilities.

Although there are some similar capabilities exposed through the various management platforms, the overlap helps to provide good, comprehensive security for mobile devices. The reason why overlapping features may be necessary is because mobile devices can exist in different states.

Workplace Join is designed to simplify resource access. As such, it's not a true mobile device security feature.

Imagine for a moment that a user wants to access e-mail and the calendar running on a mobile device. Because ActiveSync is an Exchange Server feature, a user could easily synchronize the device to Exchange Server without ever enrolling the device in Active Directory. Doing so would provide the user with access to his mail, calendar, contacts and even his task list.

In this type of situation, the ActiveSync policy settings are the only security mechanisms readily available to prevent sensitive information from the user's mailbox from being exposed should the user's device be lost or stolen. Remember, the user is able to access the full contents of his mailbox without ever performing a Workplace Join.

Although the user can access quite a bit of data without ever performing a Workplace Join, the user's abilities are also somewhat limited. Sure, the user can access the contents of his mailbox (including the calendar, contacts and task lists), but can't access any other network resources. This is where the Workplace Join feature comes into play.

Oftentimes administrators use the Workplace Join feature as an access control mechanism for users who are accessing Web applications from outside the organization. Organizations with large SharePoint deployments in place are probably the best example of this. If a user attempts to log into SharePoint from a desktop computer that exists inside the organization's firewall, the user's Active Directory credentials are normally used to authenticate the user directly into SharePoint.

In the past, connectivity to SharePoint Sites has worked similarly for a user connecting from outside the organization. The user would enter a URL for the SharePoint Site, and then be asked to enter her Active Directory credentials to gain access to the site.

The Workplace Join feature provides an extra layer of security. The network can be configured to deny access to SharePoint content for devices that aren't enrolled in Active Directory. This is a great way of preventing users who are working on unauthorized devices from accessing SharePoint resources or other Web-based resources.

Redmond's recommendations for Microsoft Intune have changed considerably over the years.

So what about Microsoft Intune? Microsoft Intune is a management tool that provides a single pane of glass for managing a variety of mobile devices. Microsoft Intune offers a self-service portal where a user can access applications or perform functions such as a device wipe or the deployment of an application to his mobile device.

Redmond's recommendations for Microsoft Intune have changed considerably over the years. The company's current recommendation is that you use ActiveSync to secure any device that isn't enrolled in Active Directory via Workplace Join. Microsoft Intune is primarily designed to manage devices that have been enrolled in the Active Directory.

Remember, though, Microsoft Intune is designed to act as an organization-level management solution for mobile devices, and it would be counterproductive to use two separate solutions for mobile device

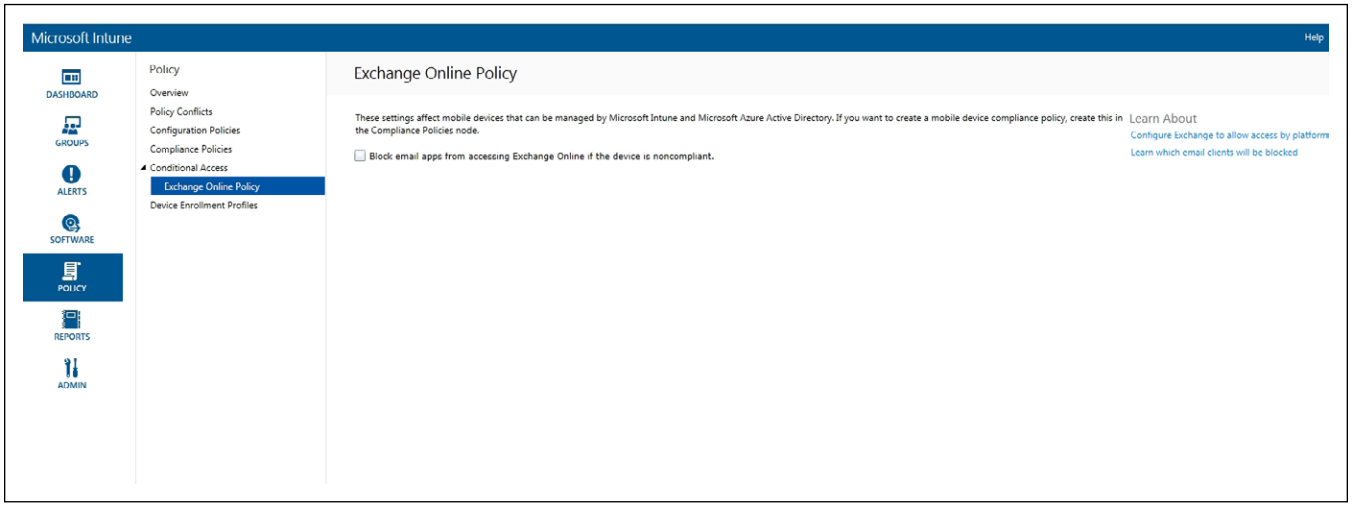

Figure 4. You can block non-compliant devices from accessing Exchange online.

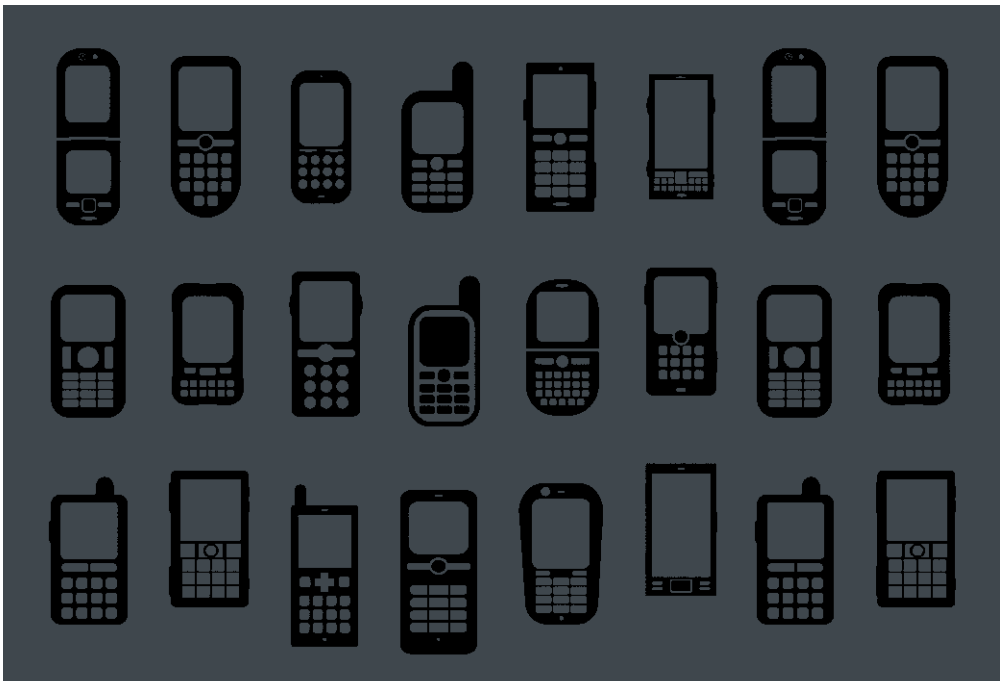

The fact that mobile devices aren't domain-joined in the traditional sense makes mobile device management tricky.

management (ActiveSync for devices not enrolled and Microsoft Intune for devices that have been enrolled).

Fortunately, it's possible to manage non-enrolled devices using Microsoft Intune. The trick to doing so is to use the Microsoft Exchange Connector. This connector ties Microsoft Intune to Exchange Server so that Intune can be made aware of devices that are being used to access Exchange Server resources, but aren't enrolled in Active Directory. That way, non-enrolled devices can be managed alongside devices that have been enrolled. In case you're wondering, it is possible to block e-mail apps on non-compliant devices from accessing Exchange.

Benefits and Limitations

The fact that mobile devices aren't domain-joined in the traditional sense makes mobile device management tricky. The tools and techniques used for mobile device management will likely change as time goes on, but for the time being, Microsoft has given us several good tools for managing device security and access control to network resources. VR

Brien M. Posey is a seven-time Microsoft MVP with more than two decades of IT experience. He's written thousands of articles and several dozen books on a wide variety of IT topics. Visit his Web site at brienposey.com.

To Join or Not to Join?

Suppliers of mobile device management and Active Directory management tools have various levels of support for the new Microsoft Workplace Join feature.

By Jeffrey Schwartz

Workplace Join allows administrators to join personal devices providing two-factor authentication and single sign-on to enterprise network resources and applications.

Microsoft introduced Workplace Join in Windows Server 2012 R2 to make it easier to connect employee-owned tablets and smartphones and other device types not designed to join an Active Directory domain—notably iPads and Android-based tablets and phones. Of course, that also includes Windows RT tablets and phones based on the Microsoft Windows Phone OS.

Workplace Join allows administrators to join personal devices providing two-factor authentication and single sign-on to enterprise network resources and applications. When enrolling a device using Workplace Join, Active Directory can retrieve the attributes of that device providing "conditional access for the purpose off authorizing issuance of security tokens for applications," according to Microsoft.

But as the article, "Manage Mobile Devices and Policies in Active Directory," p. 42, on how to implement Workplace Join warns, it has limitations. Workplace Join is only designed to simplify resource

access and is not intended as a complete mobile device security feature. It also doesn't provide Group Policy settings that can be applied to mobile devices, has limited access control mechanisms and doesn't provide the types of device security controls available with ActiveSync polices.

For its part, Microsoft doesn't market Workplace Join as a mobile device management solution, though it's enabled in the company's own new Enterprise Mobility Suite, and specifically the Intune management service, so perhaps it might become a requirement in the future.

Third-Party Options

Various suppliers of mobile device management software and Active Directory administration tools say their offerings provide more comprehensive methods of authenticating mobile and user-owned devices.

Various suppliers of mobile device management software and Active Directory administration tools say their offerings provide more comprehensive methods of authenticating mobile and user-owned devices. Those tools typically use various means of connectivity including Microsoft Exchange ActiveSync, Apple Push Notification Service (APNS) and Google Cloud Messaging (GCM). Whether Workplace Join becomes a preferred means of enrolling mobile devices in Active Directory domains remains to be seen.

So far, integration between Workplace Join and third-party tools is at a formative stage. A few offer it in some form, while others don't see a need for it at this point. Chris Ashley, a product manager at Dell Software, says customers have inquired about Workplace Join support. "Among customers I have actually talked to, they're actually excited about this feature, they're just holding off mostly because of the fact they're still running a lot Windows 7 desktops," Ashley says. "But it's something they want to introduce. A feature like this is really cool, but sometimes you do find those little shortcomings that make it a procedure."

Dell last month released a new module for its Active Administrator tool for Active Directory management. The new module aids in the management of a certificate, which is a key part of setting up Workplace Join, according to Ashley. It will install the certificate on the server so when it expires, an administrator can update it. The new module also makes it possible to assign access to resources via IP addresses. The benefit of using Active Administrator, Ashley notes, is that it supports management of Group Policy Objects.

"Certainly any policies that you use to provide access to the servers internally, we can cover," he says. "We're able to manage those policies, to recover them if they're messed up and give it confidence to change those policies because we can roll those polices back, if the change has a detrimental effect."

The new module, called Active Administrator for Certificate Management, provides DNS management capability in addition to certificate management. "The DNS management capability is important because there are two records that you have to create to make sure devices that are trying to register with Workplace Join can actually locate the machines that are required," Ashley says. "So being able to manage those records, and monitor that those records exist and that they can be reached will also be important to folks who are trying to leverage Workplace Join."

"Right now I don't see many customers adopting it."

Tomas Vetrovsky, Director of Product Management, Mobile Iron

Down the road, Ashley says Dell is evaluating how its tools, which in addition to Active Administrator include GPO Manager, might add more security to devices registered using Workplace Join. An example would be more refined policy management, but factoring into that will be new capabilities delivered by Microsoft in the next release of Windows Server and Windows 10, as well as customer demand.

Mobile Device Management

Tomas Vetrovsky, director of product management at Mountain View, Calif.-based MobileIron, a supplier of mobile device management software, says customers have inquired about integrating its namesake software with Workplace Join. "Our customers would like to use Workplace Join, mainly for single sign-on," Vetrrovsky says. "But right now I don't see many customers adopting it."

With MobileIron software, while it connects to Active Directory, once the authentication is established, it handles management and policies, he says. In and of itself, Workplace Join wasn't designed to work with GPOs, Vetrovsky adding it doesn't need to when using mobile device management (MDM). "The GPO-based approach was designed for devices that are connected on the local area network," he says. "As soon as you start talking about tablets or laptops that are spending most of the time somewhere on the Internet, just connecting from the outside, MDM provides real-time management that's better than the GPO approach."

Another major MDM supplier that will integrate with Workplace Join is Good Technology, but Eugene Liderman, the company's director, public sector technology, says users of the Good Dynamics Secure Mobility Platform don't need it.

"Good can be complementary to Workplace Join or operate completely independent of it," Liderman says. "If you look at the majority of what Workplace Join provides, which is visibility to enrolled devices, some basic device-level control and single sign-on to certain back-end resources, all of this can be provided by the Good Dynamics Secure Mobility Platform. The major difference is that Good provides this without having to upgrade the Active Directory schema like Workplace Join requires."

"An admin can set up devices so that they have zero sign-on to corporate resources, but still have precise control over what users can do."

— Paul Moore, co-founder and chief technology officer at Centrify Corp.

Liderman adds that Good Dynamics respects the user-state in Active Directory when a user requests network access. For example, if a user is removed/suspended/deactivated in Active Directory, its tools will prevent that user from gaining access to network resources/data/messages. In addition, he says while Workplace Join focuses on device-level control, Good Dynamics can support device-level control via MDM. "More important, it can also enable application-level controls and policy management, as well as single sign-on access to various back-end resources whether through Good's secure browser or through a native iOS or Android application secured with the Good SDK," he says.

Customizable Policy Management

Paul Moore, co-founder and chief technology officer at Centrify Corp., says Workplace Join is very similar to the mobile device enrollment in the Centrify Suite, but the latter offers more customizable policy management. "An admin can set up devices so that they have zero sign-on to corporate resources, but still have precise control over what users can do," Moore says. "For example, an admin can indicate that an app is only accessible from enrolled devices, at certain times of the day, from particular device types, from specified countries, etc."

Centrify also offers full device management capabilities, with features including remote wipe, lock and find for end users, and centralized policy management for administrators, he says. The centralize policy management permits the configuration of e-mail settings, the installation of applications, VPN setup, device restrictions and so on.

No Workplace Join Integration

Blake Branon, lead solutions engineer at AirWatch, acquired by VMware Inc. last year, says AirWatch provides more advanced security controls via the AirWatch Secure Email Gateway (SEG). The gateway enforces granular policies to allow or disallow access to corporate content, as well as such variables as device type, OS, encryption and whether a device is jail broken. It does make use of access control lists (ACLs), as well as Windows PowerShell support. "AirWatch and Workplace Join are mutually exclusive so a customer would either use Workplace Join and manage it with Active Directory or use AirWatch," Branon says.

"We don't actually integrate with Workplace Join and we don't recommend that."

— Ananth Vaidyanathan, product marketing manager, ManageEngine

Others who question the need to use or integrate with Workplace Join are Renee Bradshaw, senior solutions marketing manager at NetIQ Corp. Bradshaw argues its NetIQ Access Manager provides simpler single sign-on capabilities than Workplace Join. That's because Workplace Join requires Active Directory Federation Services, "which is a nightmare to use," Bradshaw says. "NetIQ Access Manage also doesn't need to bother with the management of Group Policies."

ManageEngine, which supplies an MDM tool called Desktop Central, also doesn't plan to integrate with Workplace Join, according to Ananth Vaidyanathan, a product marketing manager. "We don't actually integrate with Workplace Join and we don't recommend that," he says. "It is not a mandatory thing to have Workplace Join for managing mobile devices using the product Desktop Central," adding that customers haven't requested Workplace Join integration. But he says if it's necessary in the future the company would provide it. **VR**

Jeffrey Schwartz is editor in chief of Redmond.