



Windows[®]IT Pro

The Top Ten Things Every IT Pro Using Hyper-V Needs to Know

By John Savill

Sponsored by  **VEEAM**

Contents

10. Cluster, cluster, cluster	2
9. Prepare for the worst	3
Figure 1	4
8. How to manage Hyper-V the Enterprise way	5
7. Enabling a big red button with PowerShell.....	6
Figure 2	7
6. How to back up with Hyper-V.....	8
5. Know what is happening in your environment.....	9
4. Maximize mobility with Hyper-V.....	9
3. Does Hyper-V support Betamax?	10
2. VLANs in a Hyper-V world	10
1. Using the New Generation	11
About Veeam Software	12

Sponsored by 

The Top Ten Things Every IT Pro Using Hyper-V Needs to Know

Microsoft's Hyper-V has cemented itself in many organizations and, in those organizations where Hyper-V is not the hypervisor of choice, Hyper-V is being used in a lesser capacity or is being evaluated. For many organizations, Hyper-V appeared quickly and the virtualization IT pros are new to Hyper-V and may face a steep learning curve. This whitepaper will arm readers with the most important ten things every IT pro using Hyper-V needs to know.

10. Cluster, cluster, cluster

Virtualization brings a number of benefits, including greater utilization of resources and faster provisioning. These benefits are enabled as multiple operating systems run on the same physical piece of hardware, each OS in its own virtual machine which partitions the resources of the physical server. However, this consolidation greatly impacts the availability of services running on the reduced number of physical hosts. If a physical host fails instead of one workload on one OS instance being affected because of a server failure, every VM on that server will be impacted. It is therefore critical that your Hyper-V hosts are part of a cluster. This allows virtual machines to be moved between nodes in planned maintenance scenarios, such as patching, and also ensures VMs are restarted quickly in the event of unplanned incidents such as a hardware failure. For those people who have had a bad experience with cluster based on experiences in Windows Server 2003, take heart that

the cluster creation process and troubleshooting changed radically in Windows Server 2008 and has continued to improve so give it another try.

There will certainly be times where single Hyper-V hosts are required, such as in a branch scenario, and in those cases it's important to provide some resiliency using other means such as application-level replication. But, where possible, always try to have at least two hosts and cluster them. You don't always need a SAN for the storage as SMB is supported for Hyper-V storage and with an external storage enclosure connected to both hosts, a cluster Storage Space can be created enabling access for both hosts.

9. Prepare for the worst

Given the number of natural disasters in recent years and with IT infrastructure a necessity for nearly any business to operate today, every organization is looking at how to enable Disaster Recovery. Many Microsoft technologies have a native replication technology that can help in Disaster Recovery preparation, such as SQL Server AlwaysOn replication, Exchange Database Availability Groups, Active Directory Multi-master replication, and more. These application-level replication technologies should always be used as the primary technology to enable replication and site resiliency. However, there are many applications that do not have a native replication technology. In these circumstances, the new Hyper-V Replica technology introduced in Windows Server 2012 can provide a solution. Hyper-V Replica provides an asynchronous replication of virtual machine virtual hard disks from one Hyper-V host or Hyper-V cluster to another Hyper-V host or cluster. In Windows Server 2012 this asynchronous replication occurred every 5 minutes. Additional configurability was added in Windows Server 2012 R2, adding 30-second and 15-minute replication intervals to the existing 5-minute interval. Because the replication is asynchronous, there is negligible processor and memory performance impact on the primary virtual machine but there is additional disk IO and network traffics but the disk IO was reduced in 2012 R2, specifically on the replica server. Hyper-V replica does not place any limits on the virtual machine, such as maximum number of virtual processors or memory. Additionally, there are no hard network requirements other than ensuring the network connection is sufficient to ensure the changes to the virtual hard drives can be sent between locations in a timely manner. The exact amount of bandwidth would depend on the rate of change on the virtual hard disks that are being replicated so it is important to calculate this rate of change to ensure required network bandwidth is available. Microsoft has a calculator tool that monitors the virtual machines you wish to replicate and then reports on the bandwidth that would be required to replicate them. This is available at <http://www.microsoft.com/en-us/download/details.aspx?id=39057>.

Most disaster recovery locations use a different IP configuration from that of the primary site and to automate the reconfiguration of the IP stack within virtual machines during

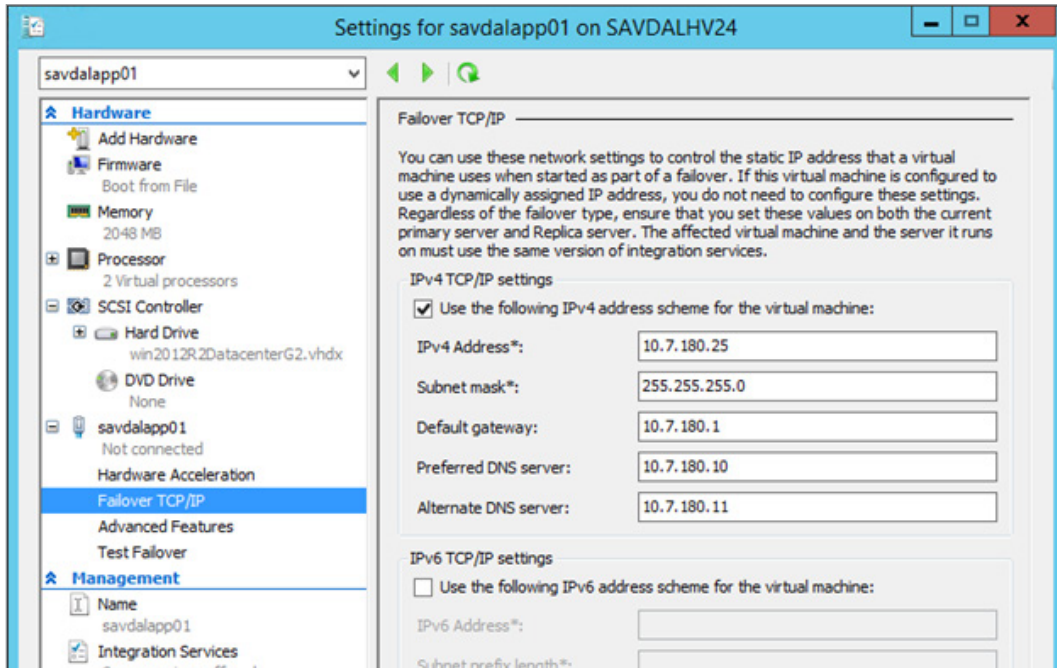


Figure 1

failover Hyper-V supports a failover IP configuration which is injected when the virtual machine is starting during a failover. This injection of IP configuration works for Windows and Linux guests.

While the original vision for Hyper-V Replica was to provide disaster recovery capability between locations, a large number of organizations leveraged Hyper-V Replica within a single location, almost like a poor man's clustering without any shared storage. The new 30-second replication interval helped greatly in those scenarios however many of those organizations still wanted to also have a replica at another location. Windows Server 2012 R2 adds the ability to extend replication which creates a replica of the replica which has its own independent replication interval which can be every 5 minutes or 15 minutes (but not every 30 seconds).

For organizations that wanted to take this a step further with a full automated site failover, Microsoft has a cloud-based service, Azure Site Recovery (formerly known as Hyper-V Recovery Manager), which provides an orchestrator failover between sites by integration with on-premises System Center Virtual Machine Manager 2012 R2 deployments that allows not only the failover of virtual machines but can also run scripts which are stored in the SCVMM library. A brand new capability that was revealed at TechEd 2014 is the ability to

also replicate from on-premises Hyper-V to Azure IaaS which enables Azure to be used as the disaster recovery site for your organization which at time of writing is in preview.

While Hyper-V Replica is a great technology, you should always use the application native replication and DR technologies as the first preference, then use Hyper-V Replica where no application-native solution is available. Also there are a number of 3rd party replication solutions available that offer additional capabilities both from within the virtual machines and externally which should be investigated.

8. How to manage Hyper-V the Enterprise way

Hyper-V Manager is a nice tool for basic VM and host management when you have one or two Hyper-V hosts, however if you want to do any of the following, then you will need something more:

- manage a larger scale Hyper-V or even hybrid environment
- use templates
- get insight into the storage and network
- dynamic optimization so virtual machines are automatically moved between hosts based on utilization
- software-defined networking
- deploy entire multi-tiered services
- enable Private Cloud

For these activities, you need System Center and specifically System Center Virtual Machine Manager, SCVMM.

SCVMM is really poorly named. In its first version it focused on virtual machine management. But since then it's really grown in scope and the naming does not do the product justice. Yes, SCVMM manages virtual machines running on Hyper-V, and ESX and XenServer, but it also performs host management and even host deployment for Hyper-V and supporting services such as file servers. It connects to storage services using industry standard protocols such as SMI-S, enabling generic classifications of storage to be assigned to the various arrays in the environment and will automatically use LUNs of the desired classification during virtual machine provisioning.

Consider setting up the networking for a virtual machine manually. You have to identify the right virtual switch on the Hyper-V server that has physical connectivity to the network you wish to connect to, then make sure you use the right VLAN for that network for each specific location and then likely identify an available IP address. With SCVMM all that networking configuration is performed automatically as the network architecture is architected inside

SCVMM, enabling SCVMM to make the right decisions based on connectivity requests for a VM to a certain logical network such as Production or DMZ.

SCVMM allows for the definition of clouds which contain certain compute, storage, and networking resources, but also specific capacities, capabilities and other assets that are usable, along with specific tenants who can use the cloud and how much the tenant and each individual in the tenant can use. If desired, these clouds can be made available to end users in a self-service model via the App Controller component which provides a web browser-accessible portal for provisioning and utilization of both on-premises virtual environments and public cloud environments such as Microsoft Azure. For organizations wanting more powerful workflow capabilities, the Service Manager service catalog can be used in conjunction with Orchestrator. All of these are part of System Center. For those wanting to bring an Azure experience on-premises, then the Windows Azure Pack is available which runs on top of System Center.

If you are looking to adopt Hyper-V in a serious way, you really need System Center. Great things can be achieved using custom code and PowerShell and other technologies, but you will be re-inventing the wheel and likely won't have an equal solution.

7. Enabling a big red button with PowerShell

Even the most powerful graphical management tool is still something you are using in a very interactive fashion, which is exactly what is required for certain types of operation. However there are also many types of action that you perform over and over again and even actions you want to be performed on a schedule or in response to some other condition. For any type of real automation you need to be looking at utilizing PowerShell, which is Microsoft's command line interface (CLI) and scripting interface of choice. With Windows Server 2012, anything you can do using the graphical interface you can do with PowerShell and the reality is there are things you can only do in PowerShell with no graphical equivalent available.

If you are looking to invest in one area of learning today, then make it PowerShell. Great strides were taken in Windows Server 2012 to make PowerShell easier to use with concepts like "least cognitive distance," think of what you want to do and that should equate to the verb-noun naming of PowerShell cmdlets (commands). For example, if I want to create a new VHD file, then the cmdlet is New-VHD. Even better is the PowerShell Integrated Scripting Environment, ISE. The PowerShell ISE is a great environment to create PowerShell scripts but also for running interactive commands. It features IntelliSense, which automatically recommends cmdlets and parameters as you type and also has a command window enabling commands to be searched for with a forms-based interface

to complete required parameters as shown. If you are starting with PowerShell, then use the PowerShell ISE.

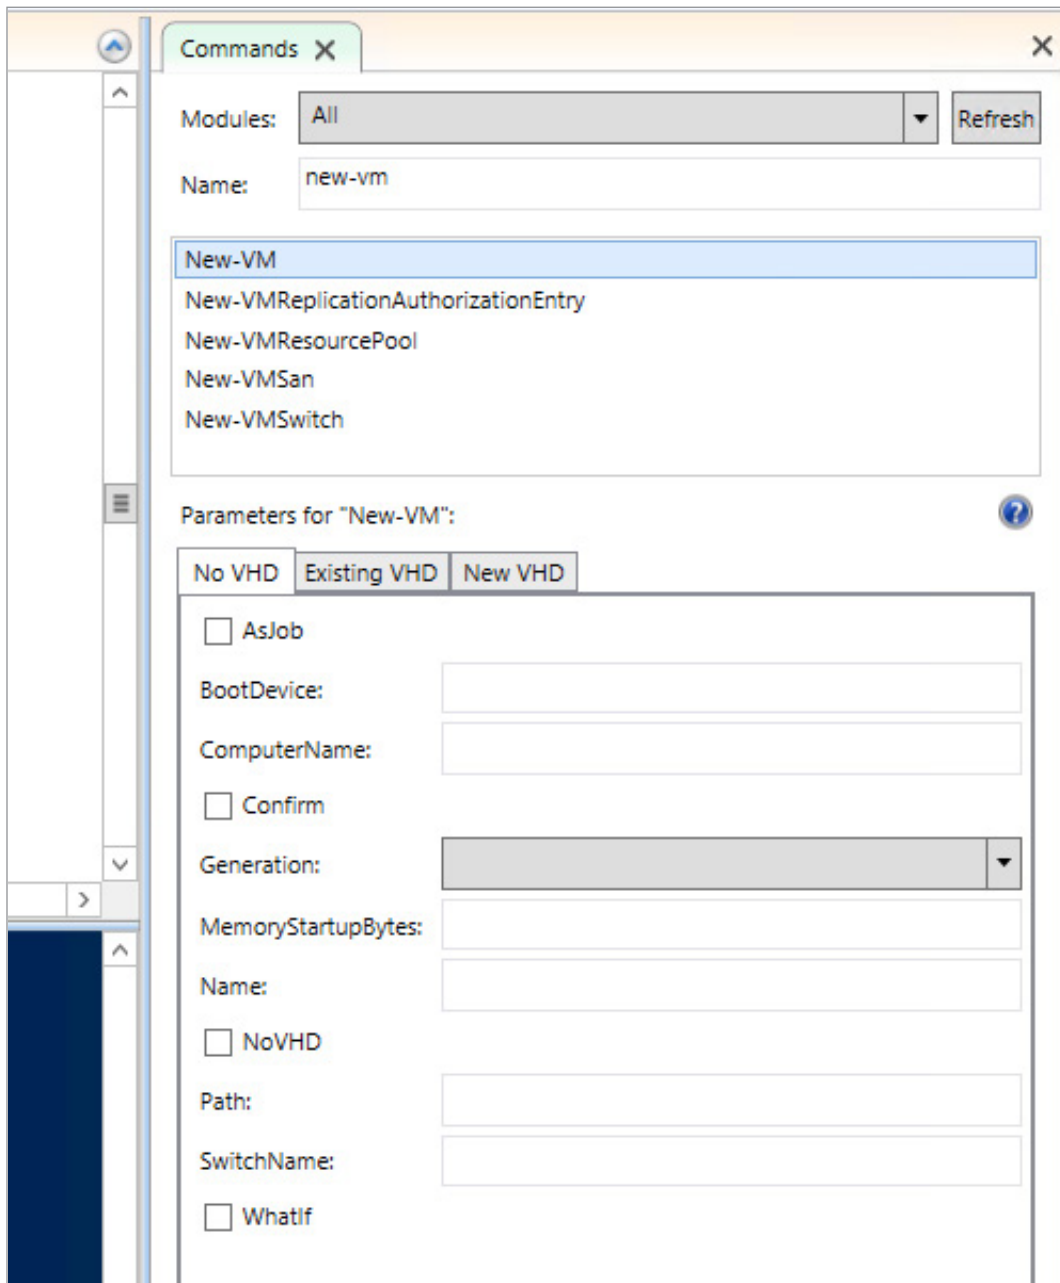


Figure 2

6. How to back up with Hyper-V

There are really two types of backup for a Windows operating system. The first is an application-consistent backup, which means special application-aware modules are utilized during backup operations to ensure required application data is flushed out to disk during a backup and other changes are quiesced during the backup which gives confidence that, if needed, the backup could be used in a recovery scenario and the data would be good to use. The second type of backup is a crash-consistent backup, which as the name implies, is not as appealing and essentially just means the current state of the file system is captured “as is.” This may mean some files are mid-way through a transaction, which means if you used that data as part of a restore there may be integrity checks required and repair operations for the data. Wherever possible you want the application-consistent backup, which fortunately is readily available.

When you think about backups in a virtual environment, once again there are two choices: You can back up at the Hyper-V host level and backup the entire VM or you can run the backup from within the virtual machine. The actual backup you get in both cases is the same, application-consistent, and this is possible as when a backup is taken at the Hyper-V host the backup request is passed through to the virtual machine via the backup integration service, which tells the guest OS to perform a backup. This calls all the registered VSS writers (those application specific modules) to ensure data on the disk is backup-ready. When the virtual hard disk files are backed up at the host, the content is known to be application-consistent.

Why would you ever back up from within the virtual machine then? It just seems like extra work deploying agents into every guest OS instead of a simple backup at the Hyper-V host level. The answer is another question: What do you want to restore? When you back up a virtual machine at the host level, you are backing up the entire virtual machine. You can restore the entire virtual machine or files from the file system but that is the level of granularity. The host-level backup has no knowledge of any application specific recovery units such as databases or items. If you back up a virtual machine within the guest OS, then the backup agent can identify the applications installed and back up in an application-specific way, for example, if you are backing up specific SQL databases or Exchange mailboxes or SharePoint items. During a recovery it's possible to not only restore the entire OS, but also a specific SQL database or Exchange mailbox or an item from SharePoint. Therefore, the decision of where to back up should be based on the granularity and application awareness of the restorations. It is common for many virtual machines to be backed up at the host level and then backup specific application workloads that are supported by the backup agent from within the virtual machine. One size does not fit all and I should point out there are some third party

backup applications that provide application aware restorations even when you backup at the Hyper-V host level.

5. Know what is happening in your environment

As I've mentioned already, with consolidation it's critical to take good care of the Hyper-V hosts as any problem will affect many workloads. This extends to ensuring the ongoing health of the environment and its utilization, which includes items like patching and back-ups, but my focus here is monitoring. Knowing the utilization of your environment, growth trends, and your capacity is critical to ensuring a performant environment that does not "get full."

Utilize a monitoring solution that goes beyond just telling you when something is broken (your users will do that for you for free) but rather can identify when resources are trending to capacity, when configuration is non-optimal, and can monitor all aspects of the environment, such as the hardware, network, storage, and the actual operating systems. Capacity planning can be very difficult and there are tools from Microsoft such as the Microsoft Assessment and Planning Toolkit, but also tools like System Center can help. System Center can also help "balance" your environment by looking at utilization of hosts in a cluster and if the resources are not distributed evenly or a specific host is overloaded then it can automatically live migrate VMs between hosts, which means no downtime to the VM. A great feature of System Center is its extensibility and rich partner ecosystem and there are great extensions from partners that bring a new level of depth to the monitoring of Hyper-V, a great example is the Hyper-V MP from Veeam who are actually sponsoring this paper.

4. Maximize mobility with Hyper-V

A great benefit of virtualization is abstraction of the OS from the physical hardware, which enables organizations to use different types of hardware without having to make modifications to the guest OS. The only requirement is the same hypervisor is used across all the various hardware platforms in the organization. Consider a typical cluster with 8 nodes in it. It is common to want to perform maintenance on a node which could be as rudimentary as monthly patching and during those maintenance operations it is important the virtual machines are not impacted. Windows Server 2008 R2 introduced Live Migration which is the ability to migrate virtual machines between nodes in a cluster with no downtime by copying the memory and device state of a VM between nodes while it's still running. This also enables the optimization of clusters by moving VMs between nodes in a cluster to rebalance workloads. However, Windows Server 2012 took mobility one step further with Shared Nothing Live Migration, SNLM. SNLM enables virtual machines to be moved between any two Hyper-V hosts with no downtime to the virtual machine. No need for

shared storage, no need for a cluster, the only requirement is a 1 Gbps network connection. A SNLM consists of:

- Perform a storage migration from the storage the VM currently uses to storage on the target Hyper-V host (which could be cluster storage) and then keep the storage synchronized
- Perform a Live Migration of the memory and device state
- Perform a flip of the active VM so it is now running on the target Hyper-V host

SNLM is a great feature for stand-alone hosts, but really stretches the mobility capabilities, which really removes all walls for VM mobility. VMs can be moved between clusters with no downtime to the virtual machine. Cross-version migration is possible enabling a VM to be live-migrated from a Windows Server 2012 Hyper-V host to a Windows Server 2012 R2 Hyper-V host, which is great for migration scenarios.

I alluded earlier to storage migration when I described the SNLM process, and to be 100% clear, it is possible to migrate the storage of virtual machines between supported storage types with no downtime to the virtual machine. You can really move virtual machines in any way you can think of.

3. Does Hyper-V support Betamax?

No, but it supports pretty much any other type of storage commonly seen in a datacenter. Pretty much any type of block storage can be used, including Fibre Channel and iSCSI but also file-level protocols in Windows Server 2012 with the support of SMB 3 as a supported storage medium for Hyper-V virtual machines. The new support for SMB enables Hyper-V to be used in new scenarios where a SAN may not be available but there may be a NAS device available that supports SMB 3. It's important to also not forget about utilizing direct attached storage with the Storage Space technology that when combined with certified external storage enclosures enabling storage in the enclosure to be seen by all connected nodes and then used as actual cluster storage. This is a great solution for smaller locations that have two servers and an external storage enclosure connected to both of them and you have an instance Hyper-V cluster!

2. VLANs in a Hyper-V world

Most organizations use VLANs to help manage broadcast traffic and provide network isolation and wonder how this can be used with Hyper-V. The good news is, it's actually very simple. First, do not configure VLAN tagging on the physical NIC that is associated with a Hyper-V virtual switch. The VLAN configuration for virtual machine vmNICs is configured on a per-vmNIC level enabling different vmNICs on different VMs to connect to different VLANs providing that the port the physical NIC in the host is connected to is configured correctly, e.g. a trunk port. The VLAN ID that is available in the virtual switch configuration

does not apply to vmNICs configured to connect to the switch but is used if you check the option to share the switch with the management partition and the VLAN ID specified is used by the vNIC created on the host OS which connects to the virtual switch for host level connectivity.

While the use of VLANs is certainly possible in Hyper-V, and will be used by most organizations, it is worth looking at Network Virtualization, which provides a true abstraction of the networks seen by virtual machines and the underlying network fabric. VLANs can become troublesome to manage and have finite limits whereas Network Virtualization is completely separate from the physical network fabric, freeing administrators from physical network switch management and a complete software-defined-networking experience.

1. Using the New Generation

Looking back at server virtualization, it was built on the premise that the operating system running within the virtual machine should not have to be modified to be able to run inside the virtual machine, although where possible the operating system would be “enlightened” to more efficiently run inside the virtual machine using special drivers for network and storage communication and other types of device. But this enlightenment should not be required. This meant that the virtual hardware exposed to the guest operating system had to emulate the hardware that was expected by the operating system and that could be used without additional drivers, such as booting from the IDE controller, supporting PXE boot using the emulated network adapter, and a large number of system devices, such as the programmable interrupt controller. This logic made sense 10 years ago when operating systems did not natively understand virtualization, but if you fast forward to today, all Windows operating systems and most Linux distributions natively understand virtualization and have the Hyper-V enhanced drivers built-in (yes, even Linux has the Hyper-V drivers built-in) which means the requirement for the virtual machine to pretend it’s a physical system is no longer valid. For this reason the Generation 2 virtual machine was introduced in Windows Server 2012 R2 Hyper-V.

Generation 2 virtual machines remove the legacy, emulated hardware that is required for operating systems that do not natively understand a pure virtual environment, specifically Windows 8/8.1 64-bit and Windows Server 2012/2012 R2 and some new Linux distributions. A Generation 2 virtual machine is UEFI-based instead of BIOS-based, which means Secure Boot is supported. A Generation 2 virtual machine does not have an IDE controller nor does it support the legacy network adapter, which means a Generation 2 virtual machine can boot from the SCSI controller and PXE boot from the synthetic network adapter. A Generation 2 virtual machine will install the OS faster and boot faster, but once it is up and running workloads will not run faster and any overhead reduction is minimal (although with the legacy hardware

removed its potential vulnerability surface is reduced.) The pre-Generation 2 virtual machine format, Generation 1, is still available and should be used for operating systems or scenarios that don't support Generation 2 virtual machines and also if you need compatibility with pre-Windows Server 2012 R2 Hyper-V and Microsoft Azure IaaS. ■

John Savill is a Windows technical specialist, an 11-time MVP, and an MCSE for Private Cloud and Server Infrastructure 2012. He's also ITIL certified and a CISSP. John is the author of the popular [FAQ for Windows](#) and a senior contributing editor to *Windows IT Pro*, as well as the author of *The Complete Guide to Windows Server 2008* (Addison-Wesley) and *Microsoft Virtualization Secrets* (Wiley). John's blog is available at <http://www.savilltech.com/blog>.

About Veeam Software

Veeam® enables the Always-On Business™ by providing solutions that deliver Availability for the Modern Data Center™, which provides recovery time and point objectives (RTPO™) of less than 15 minutes for all applications and data. Veeam recognizes the challenges in keeping a business up and running at all times and addresses them with solutions that provide high-speed recovery, data loss avoidance, verified protection, risk mitigation and complete visibility. Veeam Backup & Replication™ leverages technologies that enable the modern data center, including VMware vSphere, Microsoft Hyper-V, NetApp storage, and HP 3PAR StoreServ and StoreVirtual Storage, to help organizations meet RTPOs, save time, mitigate risks, and dramatically reduce capital and operational costs. Veeam Availability Suite™ provides all of the benefits and features of Veeam Backup & Replication along with advanced monitoring, reporting and capacity planning for the backup infrastructure. Veeam Management Pack™ for System Center the most comprehensive, intuitive and intelligent extension for app-to-metal management of Hyper-V and vSphere infrastructures, and includes monitoring and reporting for Veeam Backup & Replication. The Veeam Cloud Provider Program (VCP) offers flexible monthly and perpetual licensing to meet the needs of hosting, managed service and cloud service providers. VCP currently has more than 4,500 service provider partners worldwide. Monthly rental is available in more than 70 countries from more than 50 Veeam aggregators. Founded in 2006, Veeam currently has 25,000 ProPartners and more than 101,000 customers worldwide. Veeam's global headquarters are located in Baar, Switzerland, and the company has offices throughout the world. To learn more, visit <http://www.veeam.com>.