Backup and Recovery for Microsoft Hyper-V Using Best Practices Planning

Brien M. Posey

Introduction

There usually isn't anything overly complicated about backing up a physical datacenter. While it's true that some applications require special consideration, the backup process itself is usually very straightforward. After all, backups have been a common practice for decades. Even so, backups take on a new level of complexity when they are performed in a virtual datacenter. This inherent complexity requires much more planning in order to ensure that virtual servers (and the data residing on them) are recoverable.

Host Level Backups

In a Microsoft Hyper-V environment there are two main options for backing up your servers. The first option is to perform a host level backup. A host level backup is a backup that is made within the parent partition. This type of backup backs up the host operating system and all of the virtual machines that reside on the host.

There are two main benefits to host level backups. The first benefit is simplicity. A host level backup makes it easy to backup all of your virtual machines without having to worry about backing up each virtual machine individually.

The second benefit to performing a host level backup is that depending upon which backup application you are using, you may be able to decrease your licensing costs by performing a host level backup. Some (but not all) backup applications are licensed according to the number of agents that you use. A host level backup typically requires only a single agent, rather than requiring a backup agent to be installed to every virtual machine. For backup applications that are licensed on a per agent basis, this approach could result in significant cost savings.

Guest Level Backups

The other approach to creating Hyper-V backups is to perform guest level backups. A guest level backup is simply a backup of an individual virtual machine. Although this approach tends to be more complex and potentially more expensive than a host level backup, it has the distinct advantage of allowing more granular control over your backups. When you perform a host level backup, the virtual machines are backed up in their entirety (although it is sometimes possible to perform incremental or differential host level backups). In contrast, guest level backups give you the ability to exclude any data that you do not want to backup.

The Limitations of Windows Server Backup

Regardless of whether you choose to perform host level backups or guest level backups, there tend to be some significant limitations to the backup process. Understanding these limitations is critical to the backup planning process. This is particularly true for Windows Server Backup (many of the third party backup applications also have similar limitations).

The reason why it is so important to understand the limitations of Windows Server Backup is because it's too easy to assume that any data you backup can be restored. However, Windows Server Backup sometimes limits the ways in which data can be restored. For example, in some cases it may be impossible to restore an individual file without performing a full restoration of an entire virtual machine. The actual limitations that you may encounter vary depending upon whether you are performing a host level backup or a guest level backup.

Host Level Backup Limitations

When you use Windows Server Backup to create a host level backup, the backup process will back up the host operating system, all of the virtual machines residing on the host, the configuration data for each virtual machine, and any virtual machine snapshots that may exist. This type of backup is designed to use the Volume Shadow Copies Services (VSS) so that the virtual machines can continue to run while they are being backed up.

Although this approach may at first seem ideal, it is not without its disadvantages. For example, there are potential issues that may be encountered because of the backup process' reliance on VSS. As you probably know, VSS is a Microsoft technology, and tends not to be supported by non-Microsoft operating systems. As such, a host level backup may fail to properly backup virtual machines running Linux or other non-Microsoft operating systems.

For the vast majority of virtual machines in production today, Windows Server technologies are the operating system installed. This allows VSS to be a critical part of the data protection strategy in that ensures application consistency. VSS is a Windows framework for providers (operating system or storage device), requestors (backup applications) and writers (applications) to ensure proper backups.

Even if all of your virtual machines are running Windows, there are still a number of criteria that must be met in order for a Windows Server Backup to properly backup all of your virtual machines. First, all of the virtual machines must have the Hyper-V Integration Services installed. As you may know, the Hyper-V Integration Services are made up of a number of subcomponents that can be enabled or disabled individually. As such, simply installing the Hyper-V Integration Services alone is not enough. You must ensure that the Backup component of the Integration Services is enabled for each virtual machine, as shown in Figure A.

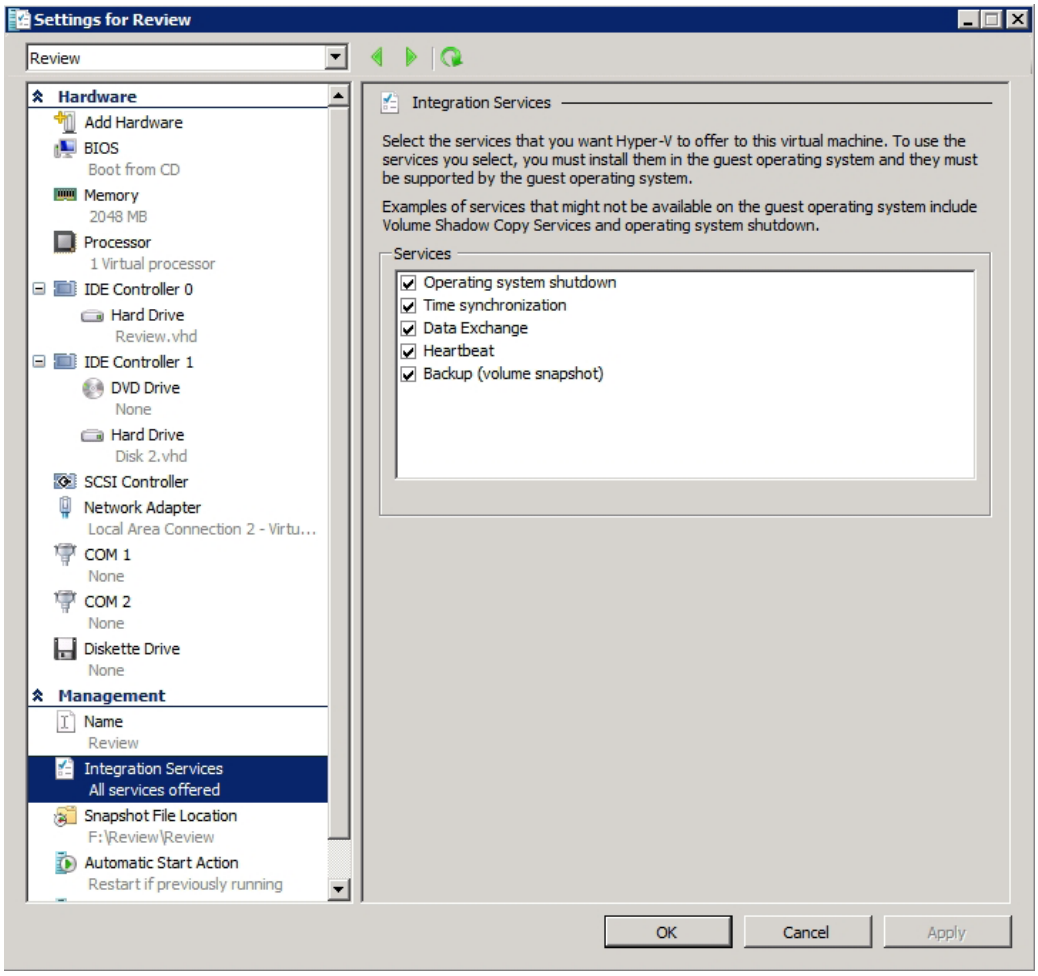

Figure A. The Backup component of the Integration Services must be enabled for each virtual machine.

There are also some requirements that must be met with regard to the types of storage that the individual virtual machines use. First, all of the virtual machines must use NTFS volumes realizing that FAT-16 and FAT-32 volumes are not supported.

Another requirement is that the guest operating systems cannot use dynamic hard disks. It is important to understand that when you create a virtual machine, Hyper-V is configured to use thin provisioning by default for the system drive. This thin provisioning makes use of a mechanism known as a dynamically expanding virtual hard disk file. This is different from a dynamic hard disks. A dynamic hard disk is a Windows logical disk structure that exists independently from any virtualization that may be in place. Because of this reason, dynamically expanding virtual hard disks are fully supported for use with host level backups, but all of the guest operating systems must treat the underlying virtual disk resources as basic disks rather than dynamic disks, as shown in Figure B.

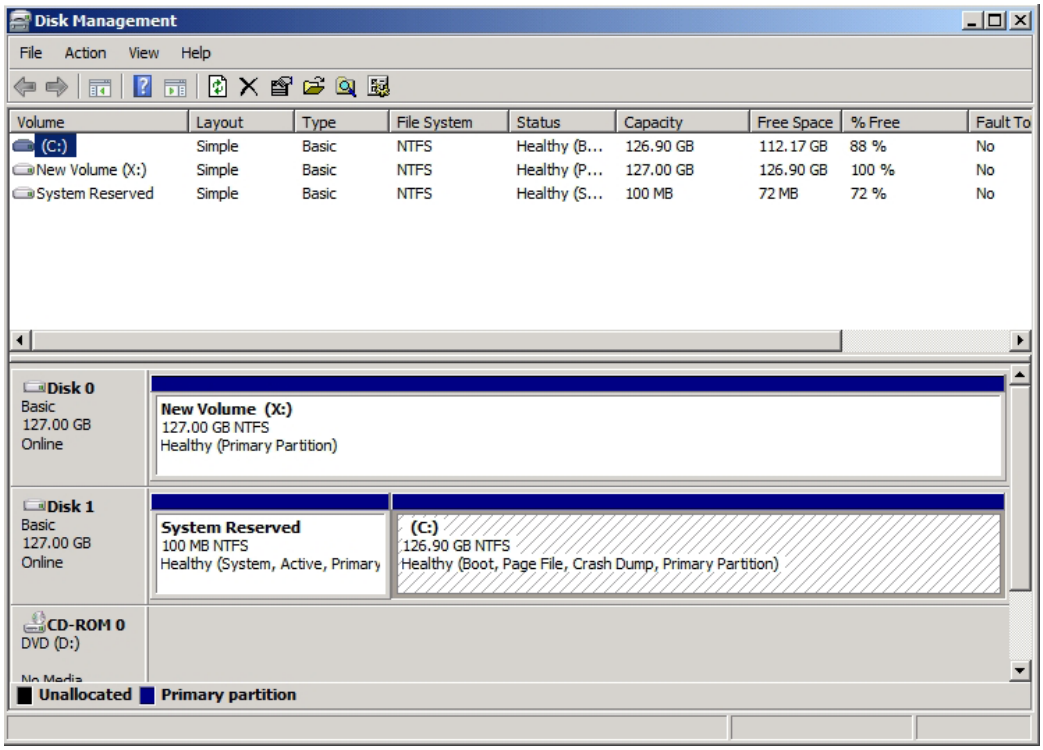

Figure B. Virtual machines must be configured to use basic disks with NTFS partitions.

At the host operating system level, VSS must be enabled for any volume that contains virtual server components. This includes virtual hard disk files, configuration files, snapshots, etc. Furthermore, each volume must be configured to store shadow copy data locally. For example, shadow copies of F: must reside on F:.

Obviously, not all host servers will meet all of these criteria. Sometimes operational requirements mean running non-Windows operating systems, using non-NTFS partitions, or using some other mechanism that is not supported by VSS. In these types of situations, host level backup using Windows Server Backup are still possible, however the noncompliant virtual machines must be shut down or suspended prior to running the backup.

Unfortunately, there are some issues with using Windows Server Backup even if your host operating system and all of your virtual machines fully comply with the requirements outlined above. The first such issue that you need to be aware of is that Windows Server Backup does not backup your virtual network configuration. Therefore, if you restore a virtual machine you will have to manually re-create the virtual network after the restore operation completes. As such, it is important to document all of your virtual network settings.

The inability of Windows Server Backup to restore virtual network settings is not the only major limitation associated with the restoration process. The biggest limitation stems from the inability to perform granular restorations. For example, if you perform a host level backup you can restore an individual virtual machine (in a roundabout way), but you can't restore files, folders, or applications within an individual virtual machine.

Another major consideration to take into account when performing host level backups using Hyper-V is that Windows Server Backup makes it difficult to restore your virtual machines to an alternate host server. Suppose for example that a non-clustered host server were to fail and that you needed to get the virtual machines back up and running as quickly as possible. A common response to this type of situation might be to restore the virtual machines to a different host server until you can get the failed host server fixed.

Windows Server Backup will allow you to restore a backup that was created on a different server, but the outcome of the restoration varies depending on what you attempt to restore. If for example, you try to perform a Hyper-V restoration (an application level restoration), Windows Server Backup will inform you that the restoration process will copy the application data (your virtual machines) to the new location, but will not recover the application (Hyper-V) itself. You can see an example of this in Figure C. Hence if you attempt to perform this type of recovery you will have to manually recreate your virtual machines using the virtual hard disk files that were recovered during the restoration process.

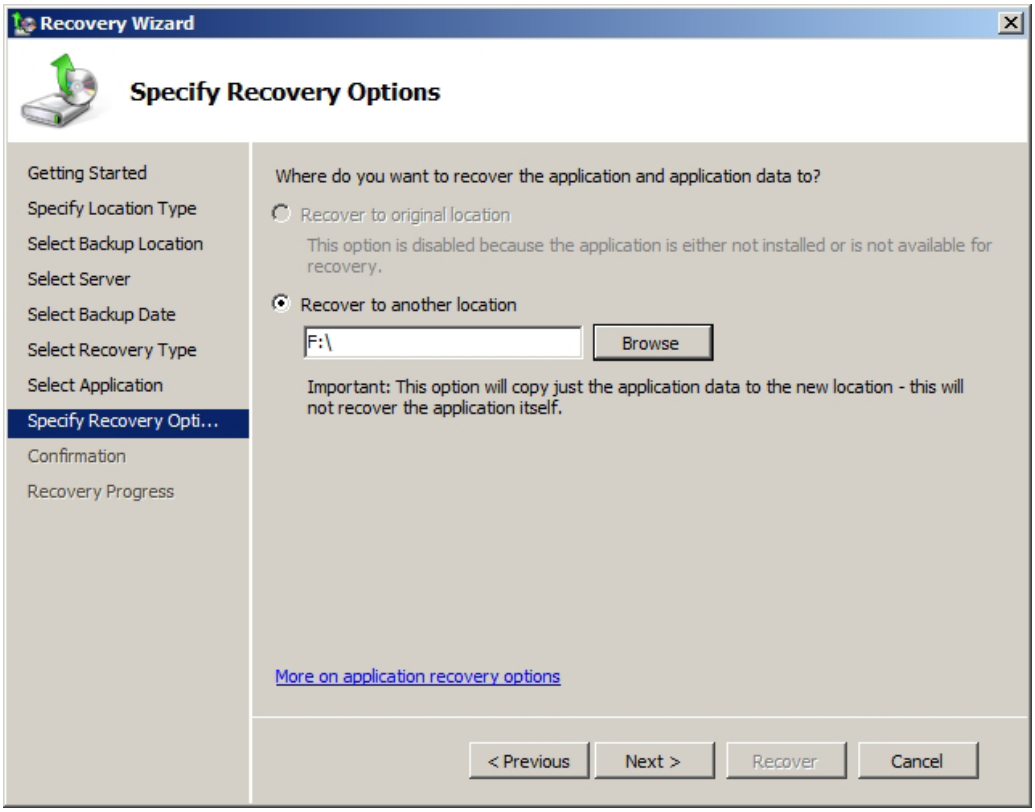

Figure C.

If you are forced to perform this type of recovery you will need to know what resources should be allocated to the virtual machines that you are trying to restore. There is also a high probability that virtual machines running Windows Server will need to be reactivated once they have been brought online using this method.

Guest Level Backup Limitations

With host level backups having so many crippling limitations, it may seem as though guest level backups might be a better choice for protecting your virtual machines. In some cases guest level backups may indeed be a better choice. After all, a guest level backup offers the same level of granularity as could be achieved by backing up a physical machine.

As previously discussed, guest level backups are performed within each individual guest machine rather than at the host operating system level. Although this may sometimes increase the backup licensing costs because of the number of backup agents that must be licensed, this concept becomes a moot point when the discussion is turned to Windows Server Backup because it is a part of the Windows Operating System and is not licensed separately. Even so, there are some issues that prevent guest level backups using Windows Server backup from being an ideal solution.

Although not necessarily a barrier to achieving effective backups, one important issue that must be considered is the impact that the backup process will have on the server's physical hardware resources. The very nature of the backup process means that disk I/O cycles will be produced. If the backup is being written to a network location then network bandwidth will also be consumed. Normally the increased load that the backup process places on a virtual machine should be well within the limits of what the underlying hardware can support. However, if multiple virtual machines on a common host are backed up simultaneously then the combined I/O and bandwidth demands placed on the host by the various backup operations can deplete the server of resources to the point that critical business processes are impacted.

While it is easy to focus on the increased workload that simultaneous guest backups can place on the physical hardware, it is also important to consider that sometimes hardware compatibility issues can make guest level backups all but impossible. This is particularly true for environments in which the backup hardware is connected directly to the physical host (as opposed to the host accessing the backup hardware from across the network).

The reason why such compatibility issues might exist has to do with the way in which virtual machines communicate with the underlying physical hardware. If the guest machine is enlightened, then the Hyper-V Integration Services can act as a collection of drivers that facilitate communications between the guest machine and the physical hardware. If the guest machine is not enlightened, the Hyper-V uses hardware emulation as a means for enabling guest machines to interact with physical hardware.

In either case, Hyper-V only allows guest machines to access the most basic server hardware (such as optical drives, network cards, and CPU cores). Depending on the type of interface that the backup hardware uses to connect to the server, the guest machines may lack the ability to communicate with the backup hardware. For example, USB and PCI based backup hardware is almost always inaccessible to Hyper-V guests. SCSI based backup devices may or may not be accessible

to guest machines depending on the particulars of the individual device. As a general rule, if you plan to perform guest level backups o Hyper-V using Windows Server Backup, you should plan to write the backups to network based backup hardware.

Once you get past the resource contention issues and the potential hardware compatibility problems, most of the issues surrounding guest level backups center around the data that is not backed up. In some ways the omission of certain data from a guest level backup is common sense. Typically when you perform a guest level backup the backup software is completely unaware that it is backing up a virtual machine. As such, the backup software treats the virtual machine as if it were running on physical hardware and skips backing up anything that is specific to the virtualization infrastructure.

More specifically this means that when you perform a guest level backup the contents of the virtual machine's virtual hard disks are backed up, but the virtual machine's configuration is not backed up. This means that guest level backups are great for allowing the granular restoration of data, but you can't use a guest level backup to restore a virtual machine as a whole.

If you find yourself needing to recover an entire virtual machine and all you have is a guest level backup then you will need to manually create and configure the a new virtual machine before you begin the restoration process (assuming that the virtual machine that you are trying to recover no longer exists). Once the new virtual machine is in place, you can perform a full system state restore just as you would do if you were attempting to perform a bare metal restoration of a physical server.

So does this mean that you can fully recover from any disaster that comes your way so long as you have a guest level backup of each virtual machine and have documented each virtual machine's configuration? Not quite. This approach still leaves a couple of things unprotected.

One of the items left unprotected by such a disaster recovery plan is the host operating system. Of course many organizations consider this to be trivial since Microsoft's recommended best practices for Hyper-V stipulate that nothing should run on the parent operating system except for the Hyper-V role and any required management agents. Still, it is worth considering that in the event of a full server failure you would have to manually deploy the host operating system and then recover each guest machine individually.

The other item that is left unprotected by guest level backups is any snapshots that may exist of the virtual machines. Although snapshots exist as virtual differencing disks, the guest operating system is unaware of their existence so they are not backed up.

Important VSS Considerations

When deciding whether to back up your Hyper-V servers using guest level or host level backups, it is important to remember that host level backups require the use of VSS. In the case of compatible Windows operating systems this means that backups occur while the virtual machines are running. This isn't really a problem for backing up file servers or application servers relying on remote data sources, but if your guest machines are running database applications then those applications must be VSS compliant.

The reason for this requirement is that database transactions actively occur while the backup is being made. The Volume Shadow Copy Services use VSS snapshots to ensure that the data that is written to your backup remains in a consistent state. Otherwise the data would change before the backup completes and the resulting backup would likely be corrupt.

The Roll of Snapshots

Hyper-V offers a snapshotting mechanism that allows a snapshot to be taken of a virtual machine prior to a configuration change. That way if something were to go wrong the snapshot can be used to roll the virtual machine back to its previous state. Although Hyper-V snapshots can be handy, they should be used sparingly for three main reasons.

First, snapshots are not true backups. Snapshots reside on the Hyper-V server. If the drive containing the snapshots were to fail then the snapshots would not be able to be used to roll back the virtual machine.

The second reason for using snapshots sparingly is that snapshots hurt virtual machine performance. Snapshots are nothing more than differencing disks. When you create a snapshot you are really just creating a special type of virtual hard disk file. All future writes are performed against this file rather than against the original virtual hard disk. This is what ensures that the virtual machine's original contents remain unchanged.

The problem is that snapshots impact read operations. When a read occurs, Hyper-V searches the differencing disk first since it contains the most recent data. If the requested data does not exist on the differencing disk then Hyper-V checks the original virtual hard disk file. The performance impact is compounded if a virtual machine has multiple snapshots.

The third and most important reason for using snapshots sparingly is that snapshots are not supported by all applications. Snapshots are fine if the server is running a relatively static application, but rolling back a snapshot has the potential to corrupt databases. In fact, Microsoft even says not to use snapshots with certain applications, such as Exchange Server.

Conclusion

So is it better to back up Hyper-V using host level backups or guest level backups? If you are using Windows Server Backup then there are tradeoffs with both backup types. Host level backups offer more comprehensive protection, but restoring exactly the data that you need can be difficult or perhaps even impossible. Guest level backups make the restoration process a lot easier, but omit important configuration data.

If you want to avoid having to deal with these tradeoffs then one option is to perform both guest and host level backups. Of course doing so will cause your backups to take longer and will increase your backup costs because of the additional storage requirements. More importantly however, host level backups are typically impractical for organizations that operate non-Windows or legacy Windows operating systems on their virtual machines.

For the time being the most practical way to ensure that your Hyper-V environment is being properly backed up and that you will be able to perform restorations with the required level of granularity is to use a third party backup application that is specifically designed to work with Hyper-V.

Veeam's solution for Hyper-V Backups

Veeam Backup & Replication provides full-featured agentless backup and replication for Hyper-V environments. Application consistency with VSS is ensured with Veeam's application processing engine for both backup and replication jobs. Hyper-V backups with Veeam Backup & Replication provide additional features such as:

- Compression and deduplication on backup jobs
- Host based replication of Hyper-V virtual machines
- Changed block tracking for super-fast backup and replication jobs

Veeam Backup & Replication version 6 can solve your Hyper-V backup challenges. Download a trial today at http://www.veeam.com.

About the Author

Brien Posey is a freelance technical writer who has received Microsoft's MVP award six times for his work with Exchange Server, Windows Server, IIS, and File Systems Storage.

Brien has written or contributed to about three dozen books, and has written well over 4,000 technical articles and white papers for a variety of printed publications and Web sites.

In addition to his writing, Brien routinely speaks at IT conferences and is involved in a wide variety of other technology related projects.

About Veeam Software

Veeam Software, an Elite VMware Technology Alliance Partner, develops innovative software to manage VMware vSphere®. Veeam vPower™ provides advanced Virtualization-Powered Data Protection™ and is the underlying technology in Veeam Backup & Replication™, the #1 virtualization backup solution. Veeam nworks extends enterprise monitoring to VMware and includes the nworks Management Pack™ for VMware management in Microsoft System Center and the nworks Smart Plug-in™ for VMware management in HP Operations Manager. Veeam ONE™ provides a single solution to optimize the performance, configuration and utilization of VMware environments and includes: Veeam Monitor™ for easy-to-deploy VMware monitoring; Veeam Reporter™ for VMware capacity planning, change management, and reporting and chargeback; and Veeam Business View™ for VMware business service management and categorization. Learn more about Veeam Software by visiting www.veeam.com.

VEEAM

#1 VM Backup

Veeam Backup & Replication

Backup | **Replicate** | **Recover**

100% Reliability | Best RTOs and RPOs | Fast and Flexible

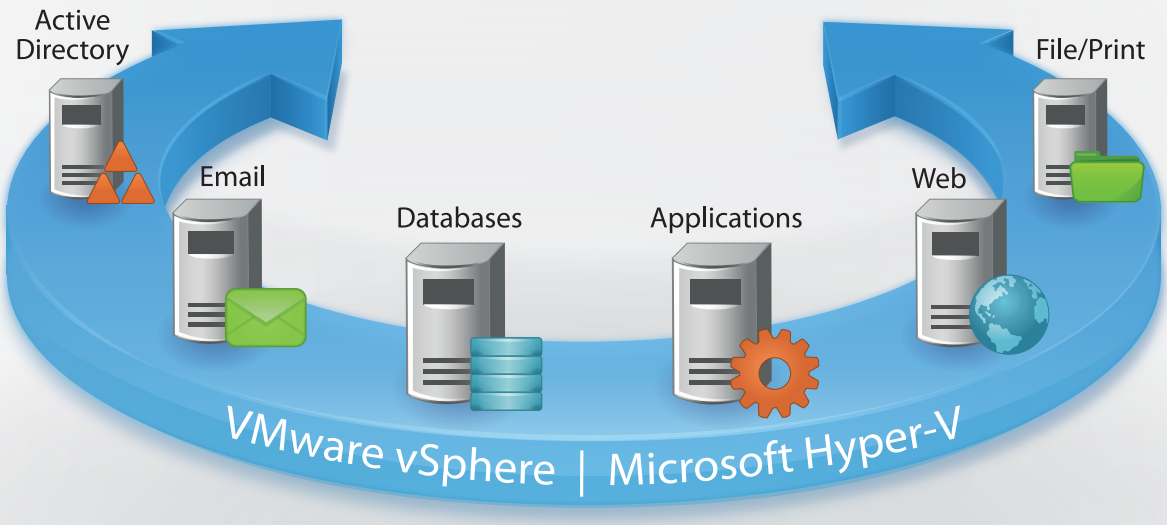

NEW! Veeam Backup & Replication™ v6

Extending the lead in VM backup with these new capabilities:

- **Enterprise scalability:** new distributed architecture streamlines deployment and maintenance of remote office/branch office (ROBO) and large installations.
- **Advanced replication:** accelerates replication by 10x, streamlines failover, and provides real failback with delta sync.
- **Multi-hypervisor support:** brings Veeam's award-winning data protection to Microsoft Hyper-V and lets you protect all your VMs—VMware and Hyper-V—from one console.
- **and more!**

BEST OF vmworld 2011 GOLD AWARD NEW TECHNOLOGY

BEST OF vmworld 2010 BEST OF SHOW

BEST OF vmworld 2010 GOLD AWARD NEW TECHNOLOGY

2010 Products of the Year GOLD

WINNER 2011 VIRTUALIZATION REVIEW

Products of the Year

TECHWORLD AWARDS

To learn more, visit
http://www.veeam.com/backup