



The do's and don'ts of a successful Hyper-V deployment

Microsoft's latest release of Hyper-V delivers a number of new improvements and features that are exciting for new and old adopters of this platform alike. If you've been waiting for the "right time" to join the Hyper-V party, R2 is it! However, like any virtualization platform, maximizing the benefits you can get from this technology is contingent upon a successful deployment. In this expert e-guide from SearchWindowsServer.com, gain expert insight into the "dos and don'ts" of a Hyper-V deployment. Learn what common mistakes IT shops often fall victim to and find out which key considerations cannot be overlooked.

Sponsored By:



Microsoft



SearchWindowsServer.com

E-Guide

The do's and don'ts of a successful Hyper-V deployment

Table of Contents

[Five mistakes to avoid when deploying Hyper-V virtual machines](#)

[Key considerations for Hyper-V virtual machine deployments](#)

[Resources from Dell and Microsoft](#)



Five mistakes to avoid when deploying Hyper-V virtual machines

By Eric Beehler, Contributor, SearchWindowsServer.com

Microsoft has made working with Hyper-V so easy that it no longer takes a specialized skill set to get a virtual machine (VM) up and running. The downside, however, is that there is plenty of room for mistakes. Even though the wizards and setup verbiage try to move you in the right direction for best practices, I still see plenty of people introducing risk into their virtual environments without even realizing it.

Here are the five most common mistakes I see with Microsoft Hyper-V deployments and how you can avoid them.

#1. Ignoring the management network

When you first add the Hyper-V virtualization role, you should dedicate a single network interface card (NIC) for management. Many people skip this because it wastes a network port. After all, without a dedicated management interface you are still able to manage the host, so why waste the port?

Well, consider the security implications. The host, otherwise known as the parent partition, is hosting several virtual machines, all with their own workloads and data. When you have access to the host, you also have direct console access to those virtual machines and the virtual hard disks. Would you allow your DMZ and internal network to operate in the same subnet? Of course not -- there would be too many risks.

Consider the host to be of a different security level than your virtual machines. The parent partition should only be manageable from a separate network interface on a network dedicated to administrative access. Without it, you open yourself up to security risks.

#2. Using the wrong kind of disk

When you setup a new virtual machine, you also setup a virtual hard disk. This dynamically expanding disk is a file on the hard drive of the host, which you have the option to make any size. This is a great option because it starts you off with a smaller file that grows only as you need it. Even if you specify a 250 GB hard disk, it will only use the space it needs. This means you will end up with a true VHD file size that is usually much smaller.

With this convenience there is a trade-off, as the performance of a dynamically expanding disk takes a hit. It not only has to expand the file as you need it, it also requires additional maintenance if you add and remove large amounts of data using the compact function. It can also pose a problem if you do not keep track of your configuration and run out of disk space. With fixed disks, which reserve the space by creating a right-sized VHD file up front, your performance is roughly consistent with the hardware and you avoid the possibility of running out of disk space. If you already have dynamic disks you can use the Convert operation to transform them into physical disks.

#3. Incorrect configurations of snapshots

One of the best reasons for systems administrators to use Microsoft Hyper-V for virtualization is the snapshot feature. It's an easy way to revert and get yourself out of a jam, allowing you to avoid a career limiting move. Still, there are a few problems with using snapshots.

First, a snapshot is not a backup. It seems counter-intuitive because the magic of snapshots makes it seem like a perfect backup, but it doesn't give you file-level restores nor does it protect you against issues on the Hyper-V host. It's a system state backup, so certain applications like Microsoft Exchange Server will have problems with in-flight data or connections. Next, snapshots are stored by default in the same location as the VHD file, so the snapshot files can build up and choke the storage of a disk with limited available space. Your first inclination may be to delete those snapshot files if you don't need them using the Hyper-V Manager. This won't actually get rid of those snapshot files; it will merely mark them for a merge into your main VHD. The next time you shut down the virtual machine the merge will occur, so if you have many snapshots this can take a bit of time. Thus, there

aren't any quick ways to relieve disk space issues with snapshots without a complete shutdown of a virtual machine, so be sure to plan your snapshots appropriately and make time for maintenance to avoid future issues.

#4. Too many CPUs

Multi-core processors are commonplace now, with a typical server having eight CPU cores as a starting point. Most people equate more cores to better performance, and Microsoft Hyper-V does allow you to assign up to four processors (or 32 processors in Hyper-V R2) to a virtual machine.

Although flexible, there is a cost to using multiple processors on a single virtual machine. You shouldn't exceed a 2:1 ratio for virtual processors to physical cores. In the case of a quad-core server, this means you shouldn't allocate more than eight virtual processors for performance reasons. Additionally, you should not assign four processors to a virtual machine because a virtual processor doesn't map to a single physical core. Instead, that work is distributed to all of the physical cores. Therefore, you should have a good reason for taking your virtual machine to multiple cores, like a CPU-intensive server such as a database server. If you are working with something like a WSUS server that sits idle most of the time and doesn't stress the processor, you'd be wasting resources by assigning too many virtual CPUs.

#5. Not taking advantage of virtual switches

Virtual switches are an extension of your network topology. Network administrators set up virtual LANs, or VLANs, and use 802.1Q trunking to make the network more efficient and easier to manage.

When we plug into the switch port, we still consider the host as the endpoint of that network connection, but that's no longer the case. If you use VLANs and trunking in your network, you can extend that functionality to your virtual machine host. Have your network administrator give you the rundown of the switch configurations and how they are set up. Instead of configuring a dedicated NIC for each subnet, you may find that you can host a variety of virtual machines all bound for different networks while conserving network ports.

This is especially helpful when you have limited NIC ports from space-saving designs such as blade servers.

Microsoft Hyper-V has helped move virtualization under the wing of the Windows administrator without the need for extensive specialized training. Just because it's easy to setup doesn't mean there aren't more complex options to consider. The option to virtualize comes with so many positives that it's easy to forget about the drawbacks until they occur in a production emergency. To avoid this scenario, make the right decisions up front and you'll be well on your way to a fully functional Hyper-V environment.



Key considerations for Hyper-V virtual machine deployments

By Eric Beehler, Contributor, SearchWindowsServer.com

Microsoft Hyper-V holds the promise of consolidation and better utilization, but how do you know if certain servers belong on a Hyper-V host? In many companies, Hyper-V starts in the lab, so it's hard to get a real-world understanding of how production systems will react when running on a shared host.

There are plenty of opinions out there, largely from database administrators and developers concerned about available resources to their applications. There are also several stories involving substandard performance of virtualized applications, but that is usually because the process wasn't fully assessed.

Previously, we looked at some of the most common mistakes involved with deploying virtual machines, but that mainly covered what you shouldn't do. So what about what you need to do? Here are some of the key considerations to keep in mind regarding Hyper-V virtual machine deployments.

Processor needs

If you are coming from the physical world, the mantra is "more is better." When attempting to share resources across several machines, however, be sure to take a closer look at your needs. Servers that are reported to have intensive CPU requirements should be looked at with scrutiny. While you can expect a 2% to 12% performance sacrifice when migrating to a Hyper-V virtual machine, this is not normally a breaking point for many systems, especially if you are moving to better processors on a virtual host.

The simple fact is that many physical systems are over-engineered. For instance, take the PerfMon readings from the physical server and determine how much utilization the server really sees. You'll have up to four processors to assign to the virtual machine, but you'll need to balance that with the overall utilization on the host.

When judging the performance of the virtual machine, you'll want to use the performance counters to determine how the machine is performing for that application. To determine its impact on the host, load up PerfMon on the host machine and look at the following counters:

Hyper-V Hypervisor Virtual Processor\% Guest Run Time and the Hyper-V Hypervisor Virtual Processor\% Hypervisor Run Time

These counters show how much of the processor the guest is actually using and how much is going toward hypervisor management.

When assigning processors to virtual machines, a 1:1 processor assignment of available physical core to virtual processors will provide the absolute best performance, but will usually leave processors sitting idle. A 2:1 virtual CPU to physical core ratio is more common for better utilization of resources, but can affect performance for those systems that require intensive CPU. If you cannot give a machine the right CPU resources or find them too expensive to dedicate to a single machine, this may be a reason to eliminate that machine from your virtual environment.

Memory requirements

While some might consider virtual memory management both a blessing and a curse, remember that Hyper-V doesn't have the same decision points about memory compared to VMware and its ability to over-commit memory. When looking at it from a performance standpoint, the assigned memory is also the available memory. It is important to allow the host operating system at least 512 MB, but 2 GB is recommended since you need to add 32 MB per virtual machine for the first 1 GB of RAM assigned as overhead, and 8 MB for each additional 1 GB.

The next step is to assign memory as you would a physical machine based on a simple calculation. If the physical server requires 8 GB of RAM, assign 8 GB to the virtual server. The option to assign additional memory will not be available until the release of Windows Server 2008 R2 Service Pack 1, which allows more memory management flexibility. For

now, the decision to virtualize comes down to available memory in your virtual environment.

Storage I/O considerations

When deploying your Hyper-V virtual machines, storage I/O may cause an objection to the engine because Hyper-V puts its hard drive inside of a file by default. While convenient for management, this can translate to a loss of performance control. Just remember that storage requirements for performance don't change just because you are on a virtual machine. When it comes to high-performance applications such as Online Transaction Processing (OLTP) databases, it is still important to give that machine the I/O bandwidth it deserves.

Implementing the right I/O means presenting the right kind of storage to the application. In the SQL Server example, it is ideal to have a separate RAID1 set for the transaction log files and a separate RAID5 equivalent set for the data files. Using storage that is reserved for the Hyper-V server without considering the ramifications of its configuration is bound to cause problems, so plan storage for each virtual machine just as you would for each physical machine.

For high-performance I/O, use a dedicated volume that holds a single fixed-size virtual hard disk (VHD) file. Otherwise, use a pass-through disk for direct access to the storage system. The performance differences between these two options are minor. By using a pass-through disk you will lose the ability to create snapshots or perform host-based backups, but the upside is that you are accessing existing volumes and using existing logical unit numbers (LUNs).

A larger LUN can be presented to a Hyper-V host by using striped metaLUNs to push the load across multiple RAID groups on the storage area network (SAN). Avoid using dynamically expanding disks when performance is of the utmost concern. With these considerations addressed, there is little standing in the way of virtualizing a server, even when high-performing disk I/O is required.

Other concerns with Hyper-V virtual machine deployments

On some servers, migrating to Hyper-V virtual machines will not be so easy, as specific hardware connections don't play well with Hyper-V. For example, accessing a SCSI interface for a directly connected tape library is generally not supported. Although it can be configured through some trickery with an iSCSI initiator, don't expect support from Microsoft or your backup software support line.

Legacy machines -- those with Windows 2000 or earlier -- may also be cause for concern. If you are employing a legacy operating system during virtual machine deployment, use a legacy network adapter. Note this adapter can be quite expensive to your host depending on the other virtual machines sharing that host. Also consider leaving Integration Services disabled to avoid known stability problems with some Windows updates on older operating systems. If you are managing a larger number of virtual machines, however, this may be more trouble than it's worth. Instead of using Hyper-V to extend the life of legacy software, take the opportunity to move to software that will run on a supported operating system.

This also applies to non-Microsoft operating systems. For instance, be aware of driver issues with video, network and clock sync, as well as other concerns that may arise when running a flavor of Linux that is not on the Microsoft supported list. Managing a basic Linux installation that isn't fully compatible may not be a big deal due to the command-line nature of many installations and the light-weight quality of some Linux software, but test the system thoroughly before committing to hosting a Linux virtual machine that does not support Integration Services.

Some servers don't belong on a Hyper-V machine due to their function. For example, best practices dictate at least one Active Directory domain controller should be a physical server, as the potential for mistakes are all too real when you have the ability to take snapshots and move VHD files. There is also the possibility of hosting all of your DCs on a single host. It's not necessarily the performance that's the problem, but the ability to cause site-wide corruption of the Active Directory database through VHD mistakes or single points of failure.

Also consider security requirements when performing Hyper-V virtual machine deployments. If you have an at-risk server, in the demilitarized zone (DMZ) for example, you'll want to avoid hosting that server on the same host as your database server. If you don't have a host to dedicate to servers that belong to a lower security level, leave them physically dedicated to avoid potential security issues from sharing the same host.

The bottom line

Although many people take a contrarian look at virtualizing their servers because of fear of the unknown, it is important to understand the performance and visibility-based opportunities virtualization provides. By addressing systems uniquely and discussing server concerns with the proper respect and information, you can sell admins on the possibilities that don't exist with the systems they have today.

Resources from Dell and Microsoft



[Leveraging Windows 7, Dell Desktops for Virtualization, Cloud Infrastructures and Beyond](#)

[Power Management in Microsoft Windows Server 2008 R2 and 11th-Generation Dell PowerEdge Servers](#)

[Presentation Transcript: The Real Deal About Windows Server 2008 R2, Part 2 - Best New Features for SMBs](#)

About Dell and Microsoft

Dell™ and Microsoft® work together to attain one primary goal - simplify IT.

The ongoing partnership has led to a showcase of jointly developed products that are:

- Easy to manage
- Integrate seamlessly into existing environments
- Deliver immediate business value

Access to key Microsoft architecture enables Dell to develop and deliver solutions that are business-ready. Dell can be your single point of contact for everything you need - from software and hardware to infrastructure consulting and lifecycle systems support.