Microsoft

Optimizing and Troubleshooting

Hyper-V Storage

Mitch Tulloch with
the Windows Server Team

PUBLISHED BY
Microsoft Press
A Division of Microsoft Corporation
One Microsoft Way
Redmond, Washington 98052-6399

Library of Congress Control Number (PCN): 2013938808
ISBN: 978-0-7356-7898-9

Printed and bound in the United States of America.

First Printing

Microsoft Press books are available through booksellers and distributors worldwide. If you need support related to this book, email Microsoft Press Book Support at mspinput@microsoft.com. Please tell us what you think of this book at *http://www.microsoft.com/learning/booksurvey*.

Acquisitions Editor: Anne Hamilton
Developmental Editor: Karen Szall
Project Editor:
Editorial Production: Jean Trenary
Technical Reviewer:
Copyeditor: Megan Smith-Creed
Indexer:
Cover:

Contents

Introduction

Troubleshooting is a difficult art to learn because it requires deep knowledge of the subject of study, familiarity with a wide variety of tools, and thinking that can be both sequentially logical and inspirationally outside the box. Perhaps the best way of learning such arts is by watching experts demonstrate their skills as they are exhibited in different situations.

Optimizing how something performs can also be quite difficult to master. If you've ever used an old-fashioned radio where you had to find your station using a dial, you'll realize that a certain degree of fiddling is required to tune things just right. Now imagine a device that has dozens of dials, each tuning a different variable, with all the variables related to one another so that tuning one affects the settings of the others. Tuning an information technology system can often be just like that...or worse!

Optimizing and Troubleshooting Hyper-V Storage is all about watching the experts as they configure, maintain, and troubleshoot different aspects of storage for Hyper-V hosts and the virtual machines running on these hosts. And when I use the word "expert" here, I really mean it, because the contributors to this book all work at Microsoft and have first-hand knowledge and experience with the topics they cover. The different sections in this book range from how to automate configuration using Windows PowerShell to get it right the first time so you won't have to troubleshoot, to step-by-step examples of how different problems were identified, investigated, and resolved. Of course there's no way to exhaustively or even systematically cover the subject of optimizing and troubleshooting Hyper-V storage in a short book like this. But I hope that by reading this book (or by referring to certain topics when the need arises) your own troubleshooting skills will become more finely honed so you will be able to apply them more effectively even in scenarios that are not described in this text.

This book assumes that you are a moderately experienced administrator of the Windows Server virtualization platform. You should also have at least a basic understanding of Windows PowerShell and familiarity with tools and utilities for managing Windows servers, Hyper-V hosts, virtual machines, and the various components of an enterprise storage infrastructure. The main focus of this book is on the Windows Server 2012 version of Hyper-V and associated storage technologies, including version 3.0 of the Server Message Block file-sharing protocol (SMB 3.0). Some content in this book will also apply to earlier versions of Hyper-V and Windows Server, and we've tried to indicate this where applicable.

Good luck in mastering this arcane art!

—Mitch Tulloch, Series Editor

About the contributors

Carlos Mayol Berral is a Microsoft Premier Field Engineer born in Majorca and based in Madrid, Spain. He is a specialized engineer for Clustering, Hyper-V, and Directory Services. Before working for Microsoft, Carlos worked for more than 12 years at TIC where he was involved in design, administration, and management areas. Now Carlos does technical and health assessments in the field and conducts workshops for Microsoft Premier customers in Spain and around the EMEA Region. You can follow his activities on the PFE Spain TechNet blog at http://blogs.technet.com/b/pfespain/. His LinkedIn profile can be found at http://es.linkedin.com/in/carlosmayol.

Chuck Timon has been with Microsoft for 15 years and is a Senior Support Escalation Engineer with Microsoft Commercial Technical Support (CTS) in Charlotte, North Carolina, US. He specializes in High Availability (Failover Clustering) and Virtualization (Hyper-V, System Center Virtual Machines Manager, App-V) technologies. Chuck has credits in Microsoft Press books, and he authors manuals for and provides training to Microsoft employees. He is a frequent contributor to the "Ask the CORE Team" TechNet blog at http://blogs.technet.com/b/askcore/ and is one of the moderators for the High Availability (Clustering) Windows Server TechNet forum found at http://social.technet.microsoft.com/Forums/en-US/winserverClustering/.

Jose Barreto is a Principal Program Manager with the File Server and Clustering Team at Microsoft, currently working on Windows Server (including several SMB features and the Hyper-V over SMB scenario). His work for the last 10 years has been focused on Microsoft storage-related technologies like SMB, DFS, the Microsoft iSCSI Software Target, SQL Server, SharePoint Server, and Data Protection Manager. He graduated with a degree in Computer Science from the Universidade Federal do Ceara in Brazil in 1989, moved to the United States in 2000, and joined Microsoft in 2002. His blog can be found at http://smb3.info, and he is also on Twitter at *@josebarreto*.

Manjnath Ajjampur has 30 years of experience in the IT industry and has spent the past 16 years at Microsoft. He is currently a Principal Datacenter Technologist at Microsoft, focusing on Systems Management and Virtualization. Follow him on Twitter (@*inadatacenter*) and LinkedIn at http://www.linkedin.com/in/manjnath.

Mark Ghazai is a Data Center Specialist with the Microsoft US State and Local Government (SLG) team. His goal is to address challenging issues within SLG customer datacenters and their journey toward private and public cloud adoption. Assisting customers to get a deeper understanding of managed and consolidated datacenters powered by Windows Server 2012, Windows Server 2012 Hyper-V, Remote Desktop, VDI and System Center 2012 suite, along with Microsoft Identity Management Solutions (FIM, UAG, TMG) is his main area of focus. Before this role, he was a Senior Premier Filed Engineer (PFE) and Senior Support Escalation Engineer for several years. His TechNet blog can be found at http://blogs.technet.com/mghazai.

Satya Ramachandran works as a Premier Field Engineer at Microsoft and is based out of Bengaluru (Bangalore), India. Satya helps customers deploy and troubleshoot issues with Windows Server virtualization solutions and clients. He specializes in areas of capacity planning and server performance.

Subhasish Bhattacharya is a Program Manager for Clustering and High Availability at Microsoft. He has worked at Microsoft at for seven years in multiple teams including High Availability and Clustering and Core Networking (DNS). His LinkedIn profile can be found at http://www.linkedin.com/pub/subhasish-bhattacharya/1/a75/b0.

Thomas Roettinger is a Program Manager at Microsoft. Thomas is part of the Partner and Customer Ecosystem Team and works with technologies like Hyper-V and System Center Virtual Machine Manager. His team runs the Windows Server TAP Program and collects very early technology best practices. Before he joined the Product Group he was the EMEA Virtualization Lead in Microsoft Premier Field Engineering. During this time he was responsible for various services including the Hyper-V Risk Assessment Program and the Implementing Hyper-V Workshop. He has rich experience in cloud implementations across various business segments such as Hosters and Enterprises. Thomas maintains a personal blog at http://blogs.technet.com/b/cloudytom and also contributes to his team blog at http://blogs.technet.com/b/wincat.

About the companion content

The companion content for this book consists of a zip file containing the Windows PowerShell scripts found in certain sections of this text. This companion content can be downloaded from the following page:

http://aka.ms/TroubleshootHyper-VStorage/files

Acknowledgments

Thanks to Anne Hamilton and Karen Szall at Microsoft Press, to Megan Smith-Creed our copy editor, and to Jean Trenary for her production services.

Errata & book support

We've made every effort to ensure the accuracy of this content and its companion content. Any errors that are reported after this content is published will be listed on our Microsoft Press site at oreilly.com:

http://aka.ms/TroubleshootHyper-VStorage/errata

If you find an error that is not already listed, you can report it to us through the same page.

If you need additional support, email Microsoft Press Book Support at mspinput@microsoft.com.

Please note that product support for Microsoft software is not offered through the addresses above.

We want to hear from you

At Microsoft Press, your satisfaction is our top priority, and your feedback our most valuable asset. Please tell us what you think of this book at:

http://aka.ms/tellpress

The survey is short, and we read every one of your comments and ideas. Thanks in advance for your input!

Stay in touch

Let's keep the conversation going! We're on Twitter: http://twitter.com/MicrosoftPress.

Hyper-V storage fundamentals

Before you can achieve any success troubleshooting you need to know the fundamentals. In this section Thomas Roettinger summarizes some basic information about key Hyper-V storage concepts and technologies that you need to know before you can effectively troubleshoot Hyper-V storage problems.

Virtual storage controllers

The first release of Hyper-V in Windows Server 2008 introduced two different types of storage controllers. The first and most important one is the IDE controller. This controller is required for booting a virtual machine. Your virtual hard disk must be attached to this controller in order to boot. The reason for this is that Hyper-V has the concept of two different device types: emulated and synthetic devices.

The IDE controller is an emulated device for the virtual machine bios and helps Windows to boot. As soon as a modern operating system that supports the installation of Hyper-V Integration components is booted, the synthetic driver for IDE is loaded as well. This ensures the same performance for a virtual hard disk attached to IDE or SCSI.

The second controller type is SCSI. A virtual machine can have up to four SCSI controllers, and each SCSI controller can handle 64 virtual hard disks. With Windows Server 2008 R2, Microsoft introduced *hot add* storage for the SCSI controller. This enables an administrator to increase available storage without the need to shut down virtual machines.

Virtual disk file formats

The file format of virtual hard disks is called *VHD* and was carried over from Virtual Server. Virtual hard disks are limited to 2 terabytes, and they come in three different flavors.

The first type is a *fixed size VHD*, meaning that when you specify the size at the point of creation the whole VHD is zeroed out. The file that gets created takes up the specified size on your physical disk. This type of VHD offers the best performance but is not very efficient in storage usage. Most administrators get around this problem by levering a storage array capability called *deduplication*.

> ***CAUTION*** **Microsoft does not recommend configuring deduplication on volumes supporting live virtual machines. For more information, see** http://technet.microsoft.com/en-us/library/hh831700.aspx**.**

The second type of virtual hard disk is called *dynamic expanding VHD*. As the name suggests, when you create a dynamic expanding VHD with 127 gigabytes (GB) it takes only a minimum of capacity on your physical disk and grows when data is stored inside it. This format is very efficient in terms of storage usage but has a performance penalty, especially on write. In real-world deployments, when you use a dynamic expanding VHD, the most important consideration is free disk space monitoring for your physical disk. If your physical disk is running out of free space, all virtual machines will go into save state.

The third type of virtual hard disk is called *differencing VHD*. This type requires a parent VHD to function and is based on dynamic expanding disks. This type is often used for virtual desktop infrastructure deployments where you have a parent VHD that contains the base operating system and one differencing disk for each virtual desktop instance.

The VHD file format is an open file format developed by Microsoft, and the specifications are available to everyone. You can find the specifications here:

http://www.microsoft.com/en-au/download/details.aspx?id=23850

Microsoft also published a performance whitepaper to show different workloads using different VHD types. This document was written to help administrators to better understand which type they should choose based on different scenarios. You can find this whitepaper here:

http://download.microsoft.com/download/0/7/7/0778C0BB-5281-4390-92CD-EC138A18F2F9/WS08_R2_VHD_Performance_WhitePaper.docx

Hyper-V also offers the possibility to pass through a physical attached disk from a Hyper-V host to a virtual machine. These disks are called pass-through disks and are intended for use with workloads that require a lot of I/O. In real-world environments, administrators also use pass-through disks to expand or shrink logical unit numbers (LUNs) online by leveraging storage array capabilities.

Storage improvements in Windows Server 2012

In October 2012, Microsoft released Windows Server 2012, which includes the newest version of Hyper-V. This new version introduced several new storage-related features.

The following are some of the storage features related to Hyper-V introduced with Windows Server 2012:

- VHDX file format
- Native support for 4 KB disks
- Support for SMB 3.0 file shares
- Virtual FC adapter
- ODX offloading
- Live storage migration

The new VHDX file format now supports up to 64 terabytes. In addition, another benefit is a log file to ensure resiliency of the VHDX, in case of a power outage for example. The new file format also includes several performance adjustments and supports 4 KB disks.

Detailed information about the new VHDX format specifications can be found here:

http://www.microsoft.com/en-us/download/details.aspx?id=34750

The most widely requested feature since the first version of Hyper-V was the ability to use file shares to store a virtual machine. The file share must be on storage backend that supports SMB 3.0, for example Windows Server 2012 with its scale out file server feature. But the industry is also working hard to get the SMB 3.0 protocol implemented in their storage solutions. To ensure resiliency, availability, and enhanced performance, SMB 3.0 supports SMB Multichannel. This requires at least two different network connections between hosts but not the storage backend.

For more information about SMB Multichannel, see the following article:

http://blogs.technet.com/b/josebda/archive/2012/06/28/the-basics-of-smb-multichannel-a-feature-of-windows-server-2012-and-smb-3-0.aspx

This new capability is comparable with what you do today with Fibre Channel (FC) or ISCSI with the help of MPIO. Frequently, when you talk to storage administrators about implementing an SMB-based storage solution, they roll their eyes and start talking about latency. That's why in Windows Server 2012, we support using SMB Direct. People often use the term *RDMA*, which stands for *remote direct memory access*. This type of network card bypasses the network stack and offers low latency and high bandwidth.

If you want to read more about SMB Direct, I recommend these blog posts:

- http://blogs.technet.com/b/josebda/archive/2012/07/31/deploying-windows-server-2012-with-smb-direct-smb-over-rdma-and-the-mellanox-connectx-2-connectx-3-using-infiniband-step-by-step.aspx
- http://blogs.technet.com/b/josebda/archive/2012/07/31/deploying-windows-server-2012-with-smb-direct-smb-over-rdma-and-the-chelsio-t4-cards-using-iwarp-step-by-step.aspx

To enable virtual machine guest clustering with FC LUNs, a virtual host bus adapter (HBA) is now part of Hyper-V. This allows you to add a virtual FC adapter to a virtual machine. This requires a physical FC adapter that is NPIV compatible in your Hyper-V host. Before Windows Server 2012, the only option to create a guest cluster was to use ISCSI-based storage. A guest cluster is a Windows failover cluster between at least two virtual machines. This is often used for SQL Server and other workloads that are cluster aware.

If your storage backend does support offloaded data transfer (ODX), you can now leverage this capability from your Hyper-V host. This is very useful when deploying virtual machines because the data moves through the high-speed storage fabric instead of using client-server network traffic and CPU time by sending an offload write including a token to request data

movement. That token is first received by starting the copy operation with an offloaded read and represents data from the storage device.

MORE INFO **To explore offloaded data transfer, see the article at** http://msdn.microsoft.com/en-us/library/windows/desktop/hh848056(v=vs.85).aspx.

For storage migration System Center Virtual Machine Manager (SCVMM) offers a feature called Quick Storage Migration. This feature takes a snapshot, and from that point on all writes go to the AVHD file and the original VHD is read only. The VHD then is transferred to the new storage location while the virtual machine is online. When the VHD transfer is finished, the virtual machine goes into save state and the AVHD is transferred. The content from the AVHD then is merged into the VHD, and the virtual machine is resumed. As you might guess, this feature causes service interruption. In Windows Server 2012, live storage migration is built into the platform without the need to leverage Virtual Machine Manager. This new feature allows you to move the storage between different locations without a service interruption. Even better, it doesn't even matter if you move the VHD or VHDX between local storage, FC, ISCSI, or SMB-based storage. When you start a storage live migration, a new virtual hard disk is created, and the content is synced between source and destination VHD. During the sync all reads go to the source VHD, but the writes go to both source and destination VHD. So if there is a failure or you abort the live migration, no damage will happen.

If you have ever done a storage migration project in a hosted or enterprise environment where the SAN is shared across multiple tenants, then you know the value of being able to live migrate storage. This is a huge cost- and time-saving functionality. But even for small or medium businesses, when a local disk runs out of free space you can just hot add new physical storage to your server and live migrate your virtual hard disks or the entire virtual machine to the new physical disk.

—Thomas Roettinger, Program Manager, Partner and Customer Ecosystem Team

Additional resources

Here are a few additional resources concerning this topic:

- What's New in Windows Server 2012 (TechNet) at: http://technet.microsoft.com/en-us/library/hh831769.aspx
- Windows Server 2012 Storage Evolved For Hyper-V (TechNet Video) at: http://technet.microsoft.com/en-us/video/windows-server-2012-storage-evolved-for-hyper-v.aspx

Storage sizing

The first step in successful optimization of virtual machine storage for Hyper-V hosts is planning. Thomas Roettinger examines a free tool from Microsoft you can use for sizing the storage needs for your virtual environment.

Using MAP

This section introduces a tool that helps you to correctly size your storage for a virtual environment. This could be either migrating physical machines to a new virtual environment or replacing an existing SAN with a new storage project.

In the section of this book titled "Monitoring storage performance," I discuss some of the different storage counters you can use for capturing all the required information via Windows Performance Monitor. As you can imagine, this requires a lot of manual effort to aggregate all the information from each machine. When you talk to storage vendors, they are all interested in the total number of IOs you need and some are interested in how many megabytes per second you need.

They need these numbers to calculate the amount of spindles you will need. The world changed and we moved into the cloud. In the cloud we use different storage tiers that are made of SSD, SAS, SATA, or a combination of these disk types.

So let me be clear on this point: *storage sizing is never an easy task*. Even after you gather the required numbers, you still need to decide whether to go for a traditional SAN with FC or ISCSI or to choose a new path with an SAS technique with SAS HBAs and JBODs to build a Windows Storage Solution. Windows Server 2012 offers many great storage features such as storage space, deduplication, SMB 3.0, and many more. This allows you to buy inexpensive storage, to purchase only what you need, and to scale.

Let's have a look at a Microsoft Solution Accelerator called the Microsoft Assessment and Planning (MAP) Toolkit that helps with several scenarios. It's a free download you can grab from here:
http://technet.microsoft.com/en-us/solutionaccelerators/dd537566.aspx

When you open MAP, the first step you need to do is perform an inventory of your servers. This can be physical machines or virtual machines—it doesn't matter. There are several ways to discover these machines such as Active Directory, IP Range, and many others.

You should select the machines that will utilize the new storage subsystem that you plan:

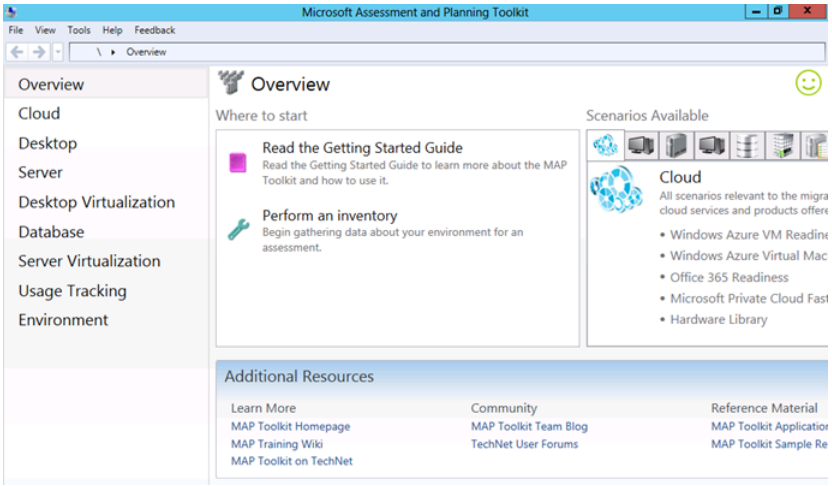

After you perform an inventory, you can gather the performance data for these servers. You need to select a time range for collecting data. I recommend you do multiple runs but at least 24 hours. Do multiple runs on different work days as well as the beginning and ending of the month. Then take the average result of the different values you receive:

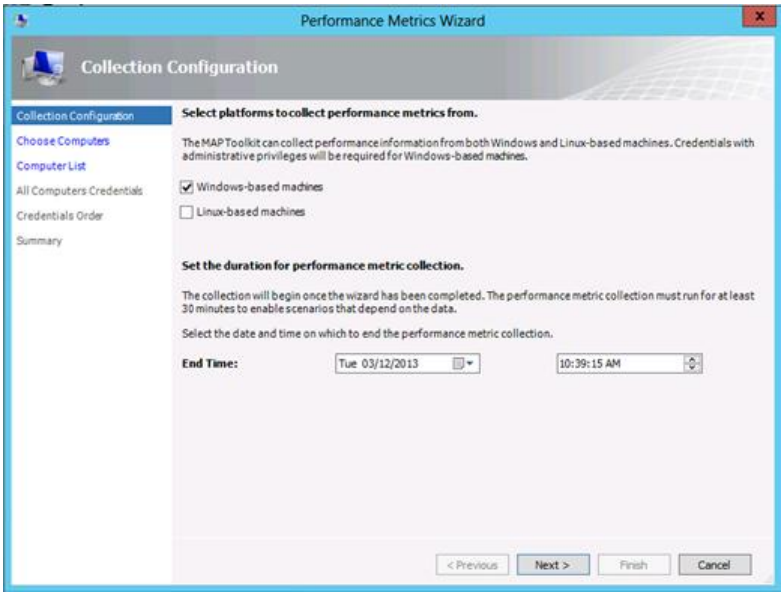

When the performance data collection finishes, click Performance Metrics:

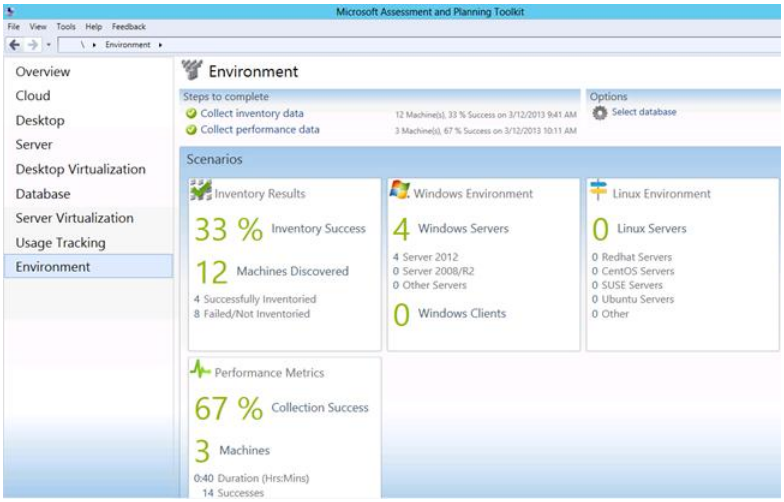

The screen will refresh and at the upper-right of the window, you can click Performance Metrics Report. This will create the report that you are looking for:

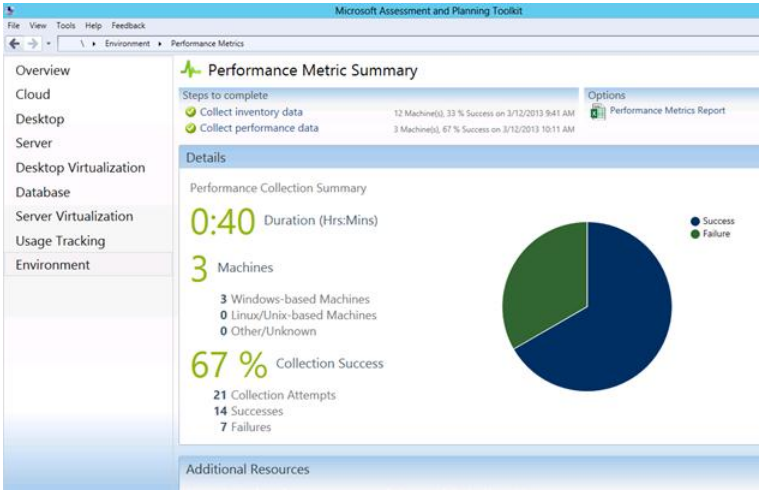

After the report is created, you will find it in your documents folder as a Microsoft Office Excel file:

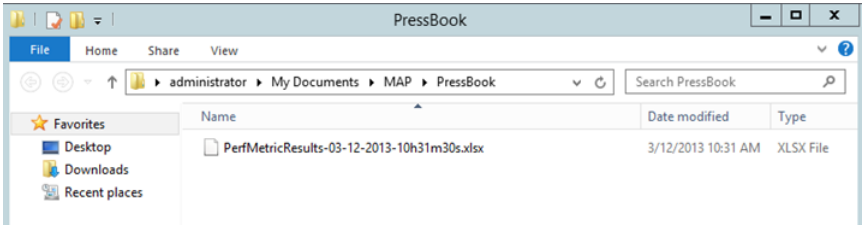

The report will contain the Average Disk IOPS, Maximum IOPS, and more:

Average Disk IOPS	Maximum Disk IOPS
1.08	3.57
0.75	2.05

The only thing you need to do is to let Excel sum up the values. This will give you the necessary values you need for a proper sizing.

Don't forget to add spare capacity for growth.

—Thomas Roettinger, Program Manager, Partner and Customer Ecosystem Team

Additional resources

Here are a few additional resources concerning this topic:

- The Microsoft Assessment and Planning (MAP) Solution Accelerator (TechNet) at: http://technet.microsoft.com/en-us/solutionaccelerators/dd537566.aspx
- Planning for Disks and Storage (TechNet) at: http://technet.microsoft.com/en-us/library/dd183729(v=WS.10).aspx

Pass-through disks

Pass-through disks are a feature of Hyper-V that allow virtual machines to access storage that is directly mapped to the Hyper-V host without requiring that the volume be configured. This storage can be any of the following:

- Physical disk internal to the host
- Direct-Attached Storage (DAS) device attached to the host
- Storage Area Network (SAN) Logical Unit Number (LUN) that is mapped to the host

In this section Chuck Timon demonstrates how to troubleshoot some scenarios involving pass-through disks for highly available virtual machines running on Hyper-V hosts.

Storage options for Hyper-V

There are several options available to Hyper-V administrators for attaching storage to virtual machines. The most frequently used option is virtual hard disks (VHD\VHDX). Another is pass-through disks. When Hyper-V first showed up on the scene as an out-of-band release for Windows Server 2008, only the VHD format existed. There was a limit on the size a VHD could be (2 terabytes), and it did not perform as well as administrators would have hoped. As a result, pass-through disks were used to get the larger disk sizes and to get the storage performance needed for virtualized application workloads.

With the introduction of a new virtual hard disk format (VHDX) in Windows Server 2012 and the capability to access SAN storage directly inside of a virtual machine by way of a synthetic fibre-channel adapter, pass-through disks are no longer needed to achieve high storage performance or to gain access to the larger disk sizes needed in the enterprise. A brief overview of some of the new features in the VHDX format can be found on TechNet (http://technet.microsoft.com/en-us/library/hh831446.aspx). The full VHDX specification is also available for download (http://www.microsoft.com/en-us/download/details.aspx?id=34750).

Pass-through disk quick review

Virtual machines can be connected to storage attached directly to the Hyper-V server. The storage can be disks internal to the Hyper-V server itself or attached externally to the server. An example of externally attached storage would be a fibre channel connection to a storage area network (SAN) using a host bus adapter (HBA). There are only two basic requirements for this configuration to work:

1. The disk is registered in the host (for example, the disk is visible in the Disk Management interface).
2. The disk is offline. The disk must be offline before it is attached to a virtual machine. Once attached, the disk is never brought online again in the host operating system. Bringing a pass-through disk online outside of the operating system in the guest could result in data corruption. In Windows Server 2012, attempting to bring a pass-through disk online is blocked.

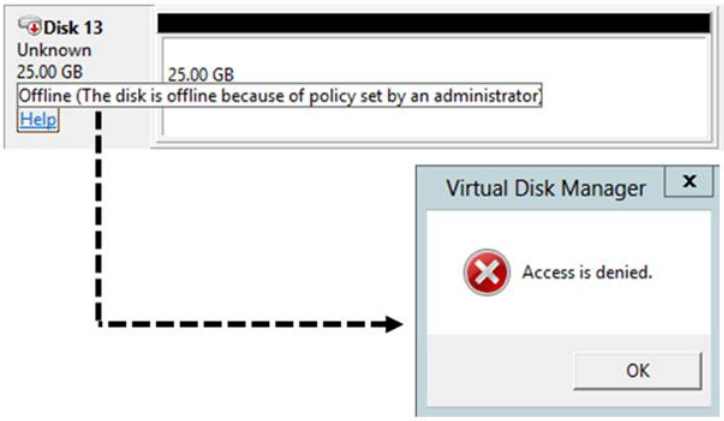

A pass-through disk is attached using either an IDE or SCSI controller in a virtual machine. There are a limited number of IDE attachment points (four). Using SCSI, 256 disks (distributed among four virtualized SCSI controllers) can be attached to a virtual machine. If an administrator wants to hot-add storage to a virtual machine (i.e., add storage while the virtual machine is up and running), the SCSI controller is the only option. A pass-through disk is configured in a virtual machine by choosing the Physical Hard Disk option when configuring a hard drive:

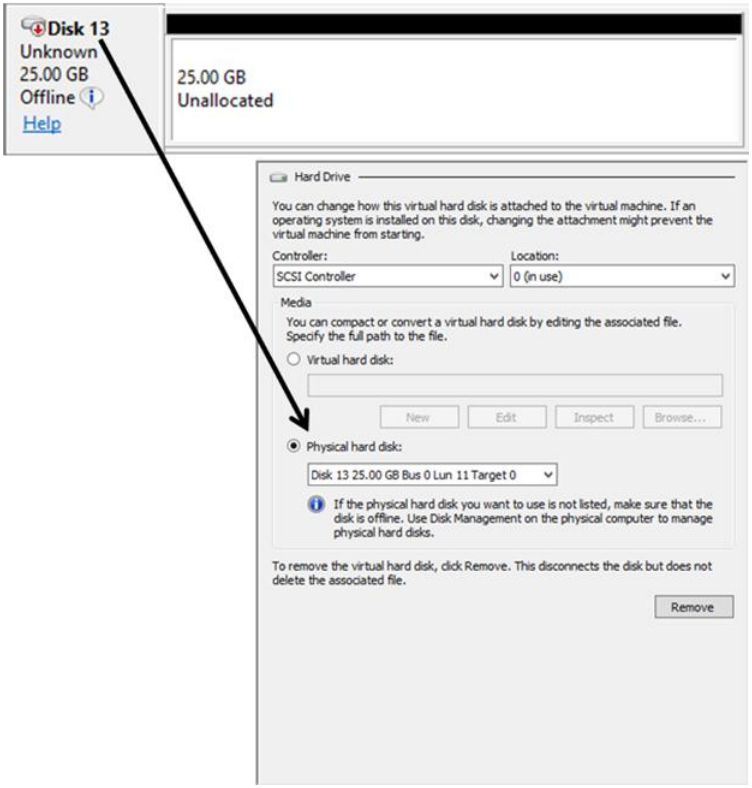

Once the settings are applied, the disk appears in Disk Management in the virtual machine and can be configured for use there. If the disk was not brought online and initialized in the Hyper-V host, when it is brought online for the first time in the virtual machine, it will need to be initialized before it can be configured further:

The standard limitations apply when using pass-through disks in a virtual machine, that is no snapshots and no backups at the host level (outside of the guest operating system). One additional limitation has been added in Windows Server 2012—Hyper-V Replica does not support replicating pass-through disks attached to a virtual machine. Hyper-V Replica supports only file replication (VHD, VHDX, AVHD, AVHDX) between primary and replica sites.

Pass-through disks can be used in virtual machines running in standalone Hyper-V servers or in highly available virtual machines running in Hyper-V failover clusters. Pass-through disk behavior as it applies to a standalone Hyper-V host has not changed in Windows Server 2012. The same cannot be said for pass-through disks attached to virtual machines in a Hyper-V failover cluster. There are several pass-through disk behaviors that have changed with respect

to highly available virtual machines in Windows Server 2012. To put it in perspective, let's flash back in time a little for a quick refresher on how things were, and still are, in Windows Server 2008 R2 Failover Clusters.

For those of us who are veterans when it comes to working with failover clusters, we know that, in a cluster, it is all about resources and the cluster having control of those resources. It does not matter what the resource is, the cluster needs to be in control of it in a high availability scenario. Working with pass-through disks in virtual machines translates into working with physical disk resources in a cluster. What this means is that before you configure a pass-through disk in a virtual machine that is either already highly available or is in the process of being made highly available, the disk, mapped to the host, which will be configured as the pass-through disk in the virtual machine, must first be added to the cluster as a physical disk resource. If not, all kinds of alarms start going off when making a virtual machine highly available. Let's take a look....

The first indication that something is wrong is when the process for making the virtual machine highly available completes. A pop-up indicates the refresh virtual machine configuration process completed with warnings, and the administrator is provided an opportunity to view a report:

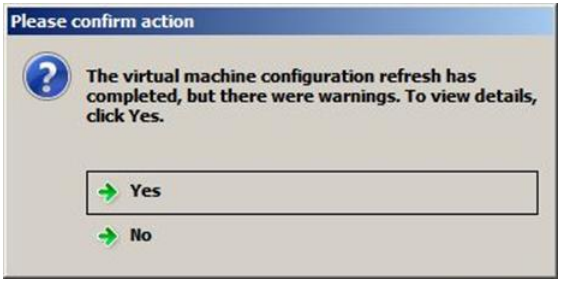

The available report indicates a failure has occurred:

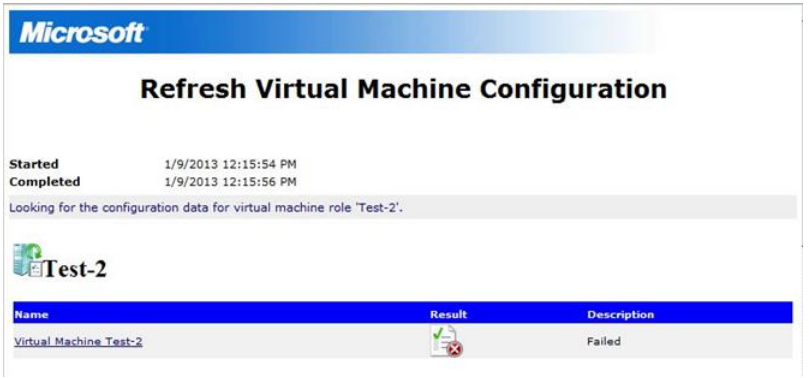

The actual details of the failure message are not clear as to why the failure occurred (i.e., "Element not found"):

A review of the additional information contained in the report clarifies the failure. The information also provides help with respect to the corrective action that can be implemented to fix the issue (i.e., add the disk to the cluster):

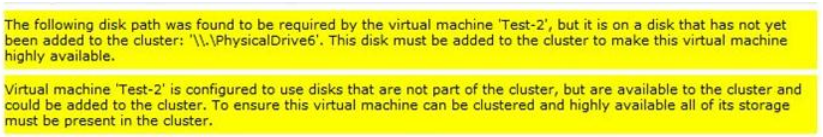

The reported failure does not prevent the pass-through disk from being added to the virtual machine configuration. However, this issue must be addressed or live migrations of the virtual machine may fail.

There are relevant events registered in both the FailoverClustering-Manager and Hyper-V-High-Availability logs:

```
Event ID: 4649

Source: FailoverClustering -Manager
Level: Warning
"Failover Cluster Manager detected that virtual machine <VM_Name> is configured to
use one or more disk that are not yet added to the cluster. Please add all
required disks to the cluster before making this virtual machine highly
available".

The Hyper-V-High-Availability log registers an event as well -

Event ID: 21105

Source: Hyper-V-High-Availability
Level: Error
"'Virtual machine Configuration <VM_Name>' failed to update the configuration data
of the virtual machine: Element not found. (0x80070490)."
```

Even with all the pop-ups and events registered in the logs, the pass-through disk is still visible in Disk Manager in the virtual machine and can be manipulated as if it had been properly configured. It is up to the administrator to heed the pop-ups and implement corrective action before actually placing the virtual machine in production.

In the end, when the virtual machine is configured properly, the pass-through disk appears as a normal disk (physical disk resource) in Failover Cluster Manager. The disk is placed in the resource group with the virtual machine it is associated with. Pass-through disks also co-exist alongside cluster shared volumes (CSV). The major difference being the physical disk resource representing a pass-through disk must be taken offline as part of a virtual machine migration process. CSV volumes, on the other hand, do not have to move with the virtual machine(s) they support:

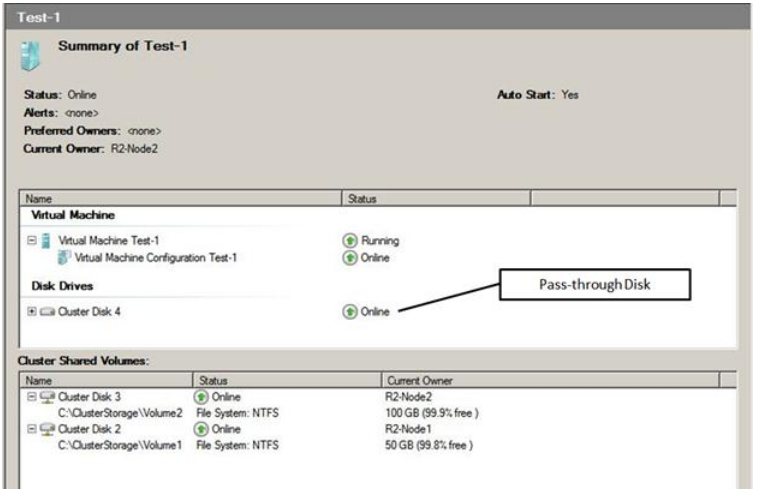

Let's switch gears and examine this same behavior in a Windows Server 2012 failover cluster.

In Windows Server 2012 Hyper-V failover clusters, pass-through disk configurations are still supported, however the check-and-balance mechanism that was in place in Windows Server 2008 R2 is no longer available.

Scenario 1: Adding a pass-through disk to an already highly available virtual machine

When configuring a pass-through disk in an already highly available virtual machine using the Failover Cluster Manager interface in Windows Server 2012 when the disk is not a cluster resource, there is no report generated as part of a Refresh Virtual Machine process. Therefore,

information that could make an administrator aware of a misconfiguration is not immediately available. The event documented above for the Hyper-V-High Availability Log (Event ID: 21105) is still registered (the FailoverClustering-Manager Event ID: 4649 is not). Even with the misconfiguration, the disk(s) can still be manipulated in the virtual machine and the virtual machine role moves freely (migrates) between nodes in the cluster without error. A vigilant administrator may eventually notice that the improperly configured pass-through disk is not listed in the Resources tab for the virtual machine and could then correct the misconfiguration.

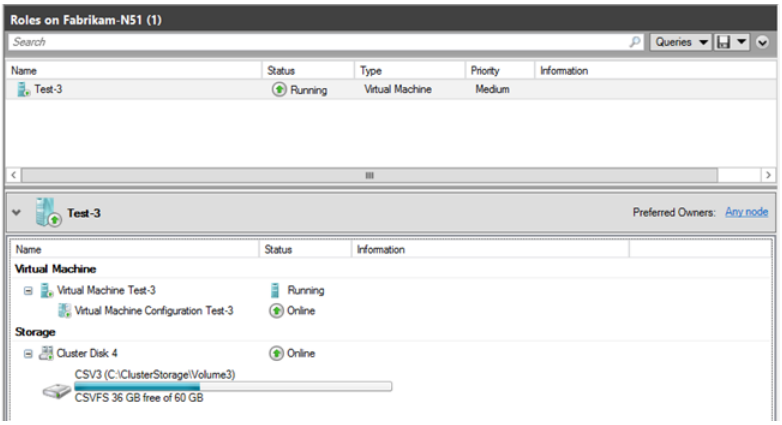

To correct the misconfiguration, the virtual machine must be shut down, the pass-through disk must be added to the cluster as a physical disk resource, and then the virtual machine must be restarted. Once that is accomplished, the pass-through disk appears in the Resources tab for the virtual machine and is properly identified as a pass-through disk in Failover Cluster Manager:

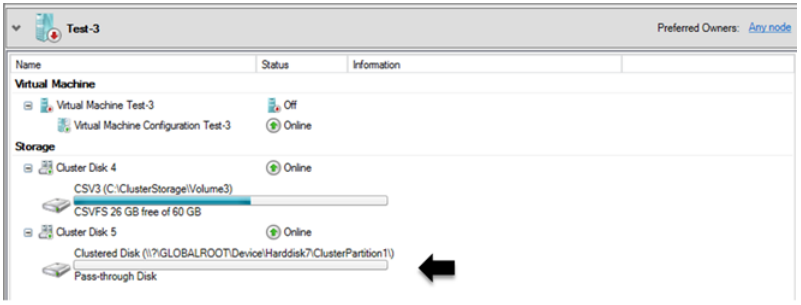

Migration of the virtual machine role continues to function properly.

Scenario 2: Adding a pass-through disk to a virtual machine before making it highly available

If a virtual machine is configured with a pass-through disk before it is made highly available (i.e., it is configured in Hyper-V Manager outside of the cluster), and the disk(s) has not been added as a cluster resource, the Configure Role process in Failover Cluster Manager detects the misconfiguration, and the process fails:

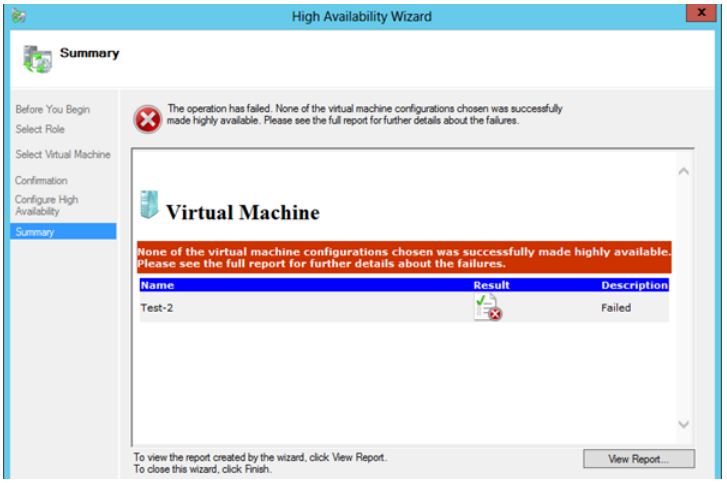

The information contained in the generated report provides actionable information to the administrator to help resolve the issue. If the disk had *not* been initialized in the host, the information detail would state that a disk is "not a path to storage in the cluster or to storage that can be added to the cluster." In other words, the cluster is not aware of the disk at all:

Disk path '\\.\PhysicalDrive6' is not a path to storage in the cluster or to storage that can be added to the cluster. You must ensure this storage is available to every node in the cluster to make this virtual machine highly available.

You can change the location of the virtual machine and its files to a valid cluster-managed storage location, or use the Move Virtual Machine Storage dialog in Failover Cluster Manager (or the Move-VMStorage Windows PowerShell cmdlet) to move the virtual machine. To open the Virtual Machine Storage dialog, select the virtual machine in the Roles view, then select the Move action, and then the Virtual Machine Storage option. The virtual machine and its files can be moved to a different location while the virtual machine is running.

If the disk had been initialized in the host, but not yet added to the cluster as a resource, the information provided in the report is different and states that the disk has not yet been added to the cluster:

The following disk path was found to be required by the virtual machine 'Test-2', but it is on a disk that has not yet been added to the cluster: '\\.\PhysicalDrive6'. This disk must be added to the cluster to make this virtual machine highly available.

Virtual machine 'Test-2' is configured to use disks that are not part of the cluster, but are available to the cluster and could be added to the cluster. To ensure this virtual machine can be clustered and highly available all of its storage must be present in the cluster.

Both of these examples provide actionable information to the administrator so he or she can remedy the situation. In Windows Server 2012, the virtual machine is not made highly available until the problem is corrected. Once the discrepancy is fixed, the High Availability

Wizard completes successfully, and the virtual machine role is properly configured. As in Windows Server 2008 R2, the application log entries both for Failover Cluster Manager and Hyper-V Manager provide information the administrator can use to resolve this problem.

Scenario 3: Adding a pass-through disk to a virtual machine that is already running

The final scenario in Windows Server 2012 failover clusters involves adding pass-through disks (hot-add storage) to a virtual machine that is already running (online and providing services to end users). Ever since Windows Server 2008 R2, administrators have been able to, on demand, add storage to a virtual machine while it is running (a.k.a. hot-adding storage) provided the storage was connected to a virtual SCSI controller. The storage could be file based (a VHD\VHDX) or a pass-through disk.

> ***TIP*** **Always ensure the integration services in the virtual machine are updated and match those in the Hyper-V host. The Get-VMIntegrationService Windows PowerShell cmdlet can be used. Here is an example:**
>
> ```
> Get-VmIntegrationService -VMName Fabrikam-FS10 | Where-Object
> {$_.SecondaryOperationalStatus -eq 'ProtocolMismatch'}
> ```

In Windows Server 2012 standalone Hyper-V servers, the above statement still holds true. In Windows Server 2012 failover clusters, an administrator can hot-add storage to a running virtual machine provided that storage is file-based (a VHD\VHDX) and connected to a SCSI controller. A pass-through disk cannot be added to a running virtual machine. The pop-up error is Access Denied for the virtual machine management service account:

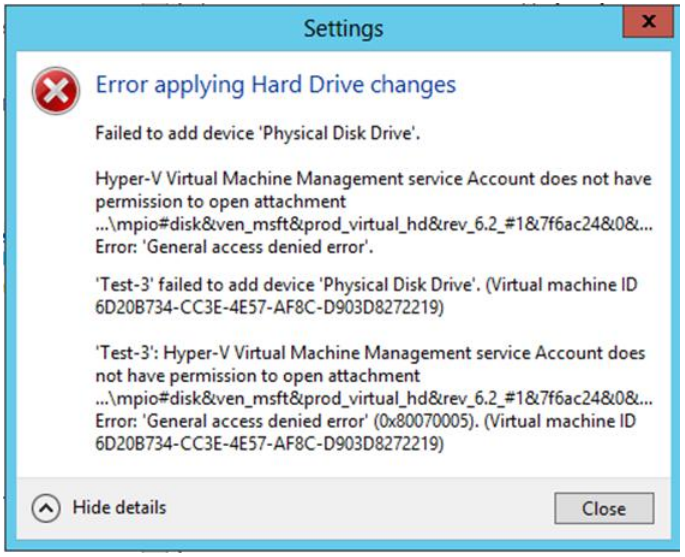

An event (Event ID: 12290) is registered in the Hyper-V-SynthStor\Admin log:

```
Event ID:  12290
Source: Hyper-V-SynthStor
Level:  Error
User: NT VIRTUAL MACHINE\<VMID>
<VM_Name>:8007005 Account does not have permission to open attachment <path to disk>.
Error: 'General access denied error' (7864368). (Virtual Machine ID <VM_guid>
```

An event (Event ID: 14140) is registered in the Hyper-VVMMs\Admin log:

```
Event ID: 14140
Source: Hyper-V-VMMS
Level: Error
User: System
'VM_Name> failed to add device 'Physical disk Drive'. (Virtual Machine ID <guid>)
```

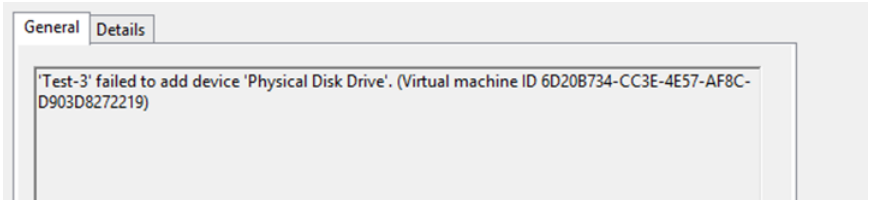

To be able to successfully add the pass-through disk to the virtual machine in this scenario requires the virtual machine first be shut down and then the disk can be configured.

That concludes my discussion on Hyper-V storage with respect to the functionality changes in Windows Server 2012 pertaining to pass-through disks. Please keep in mind what I stated earlier: there should be no reason to continue using pass-through disks in Windows Server 2012 considering the new functionality available in the new VHDX virtual hard disk format.

—Chuck Timon, Senior Support Escalation Engineer

Additional resources

Here are a few additional resources concerning this topic:

- Adding a Pass-through Disk to a Highly Available Virtual Machine at: http://blogs.technet.com/b/askcore/archive/2009/02/20/adding-a-pass-through-disk-to-a-highly-available-virtual-machine.aspx
- Configuring Pass-through Disks in Hyper-V at: http://blogs.technet.com/b/askcore/archive/2008/10/24/configuring-pass-through-disks-in-hyper-v.aspx

Virtual machine snapshots

Snapshots are a feature of Hyper-V that provide a quick and easy way to revert a virtual machine to an earlier state. This can be useful for example if you need to recreate a specific state or condition in order to troubleshoot a problem or demonstrate some functionality. Snapshots are mainly intended for use in development and test environments and are generally risky to use in production environments that use Active Directory. In this section Thomas Roettinger shows how to resolve an issue caused by a broken snapshot tree.

Understanding snapshots

Often when I talk to people they misunderstand the concept of a Hyper-V snapshot. They believe this is a backup function. But it is not! The functionality these people are looking for would be a VSS snapshot. For more information on VSS, see the "Additional resources" at the end of this section.

So what is a Hyper-V snapshot? The best way to describe it is as a point-in-time picture. The original intention for this feature was to offer a way to create a snapshot, for example before installing an update on a system. If the update was successful you could delete the snapshot, and if it was not successful, you could apply the snapshot to return to the original state.

A virtual machine consists of a configuration file and at least one virtual hard disk file. There also is a VSV file and a bin file. When you create a snapshot of a virtual machine, Hyper-V creates an AVHD/AVHDX file and saves a copy of the configuration file to the snapshot location. From that point on, the original VHD/VHDX is read only; all changes go to the AVHD/AVHDX file. A copy of the memory state file (bin) and the processor structure file (VSV) are also placed in the snapshot folder.

> ***MORE INFO*** **For more information on Hyper-V virtual machine files, see Ben Armstrong's post at** http://blogs.msdn.com/b/virtual_pc_guy/archive/2010/03/10/understanding-where-your-virtual-machine-files-are-hyper-v.aspx**.**

When you apply a snapshot, all changes that have been stored in the AVHD/AVHDX file are deleted. When you delete a snapshot, the changes from the AVHD/AVHDX are merged with the original VHD/VHDX. With Windows Server 2012, this happens online. In previous Windows Server versions, the merge did not take place until you shut down or rebooted the virtual machine.

You can create multiple snapshots of a virtual machine up to 1,024; that's the maximum chain length for differencing disks. Multiple snapshots form a snapshot tree, as you can see in the following screenshot:

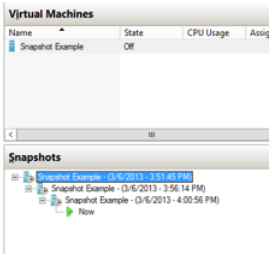

Using snapshots in a production environment is supported but is generally not recommended because of a performance penalty the longer the chain gets. In addition, Microsoft does not recommend snapshots for specific workloads because of potential data corruption issues. Examples include SQL Server and domain controllers, although Windows Server 2012 implements functionality where virtualizing domain controllers is no longer an issue. For details, see http://technet.microsoft.com/en-us/d2cae85b-41ac-497f-8cd1-5fbaa6740ffe(v=ws.10)#bkmk1_planning_to_virtualize.

Finally, when you use Hyper-V snapshots in production, you should implement a policy that requires deletion of a snapshot when it is no longer needed, for example after a successful update of a system.

Example: Broken snapshot tree

One of Patricia's virtual machines is running out of available disk space, so she shuts it down. The virtual machine has some snapshots that a developer created. Patricia uses Hyper-V Manager to expand the original VHD from 127 GB to 160 GB.

When she tries to re-start the virtual machine, she receives the following error message:

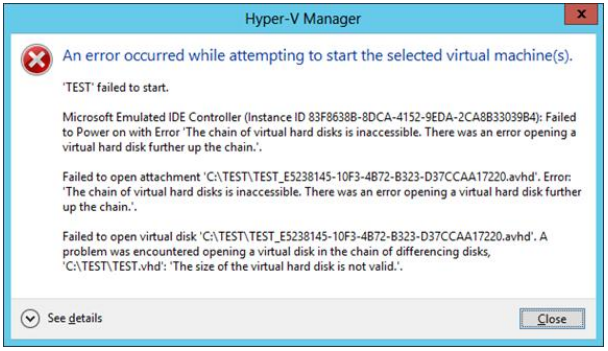

Patricia knows that shrinking the original VHD is not possible with Hyper-V Manager. She has expanded VHDX files with snapshots attached several times without any issue. Patricia decides to use a free utility called *vhdtool* available for download (http://archive.msdn.microsoft.com/vhdtool). This tool repairs a broken snapshot tree by returning a base VHD to its original size. This only works without data loss as long as the content was not changed after the expand operation.

The vhdtool also enables an administrator to create fixed-size virtual hard disks quickly because it doesn't zero out the file like Hyper-V Manager does. Instead it bypasses the file system, and data that reside on the physical disk are available inside the VHD. Keep this in mind and rate the security risk of using the tool for your environment.

—Thomas Roettinger, Program Manager, Partner and Customer Ecosystem Team

Additional resources

Here are a few additional resources concerning this topic:

- About Virtual Machine Snapshots (TechNet Library) at: http://technet.microsoft.com/en-us/library/dd851843.aspx
- Hyper-V Virtual Machine Snapshots: FAQ at: http://technet.microsoft.com/en-us/library/dd560637(WS.10).aspx

File system alignment

File system misalignment will only be a problem if you are using Windows Vista or Windows Server 2008 as an operating system for virtual machines that are stored on a storage attached network (SAN). This is not just a Hyper-V problem; it will also occur for other hypervisors. Thomas Roettinger explains how to identify and deal with this issue.

Identifying file system misalignment

Working with a virtual environment also means working with different layers of storage. Windows Server systems running Hyper-V include the following layers:

- File system inside a VHD/VHDX
- NTFS on the host
- LUN

For the best possible performance of your storage subsystem, you want the file system inside the VHD, the NTFS file system of the host, and the storage array blocks aligned.

Here is a simplified architecture for a properly aligned system:

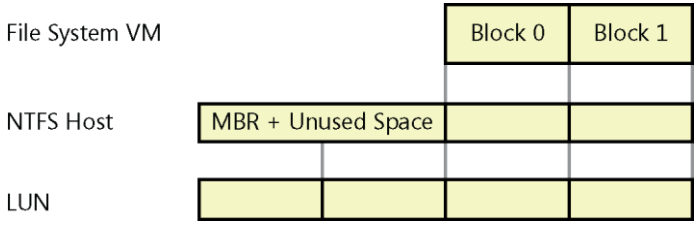

The following block diagram shows a file system misalignment problem what will result in bad performance. Every block that is accessed by the operating system of the virtual machines requires the LUN to access two blocks. This doubles the I/O load on your SAN:

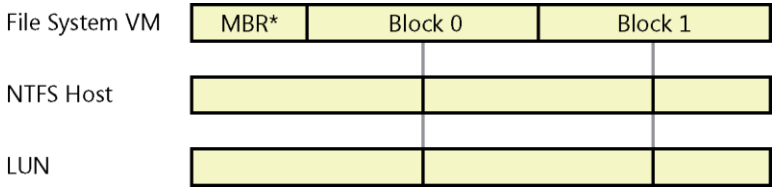

Historically, an operating system like Windows 2000, Windows 2003, and various Linux distributions started the first partition at sector 63. This led to a misaligned file system because the partition did not begin at a sector that is a multiple of 8. Beginning with Windows Vista

and Windows Server 2008, operating systems no longer have this problem because their first partition is at 1,048,457, which is divisible by 4,096.

How do you verify that your virtual machines running an operating system prior to Windows Vista or Windows Server 2008 have a properly aligned file system?

You simply launch msinfo32.exe, like this:

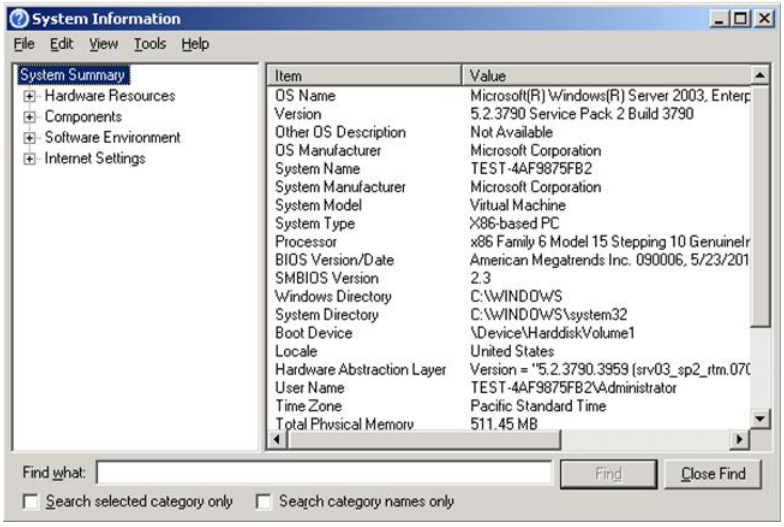

You then expand Components, then Storage, and click Disks. Write down the Partition Starting Offset; you will need this value in the next step:

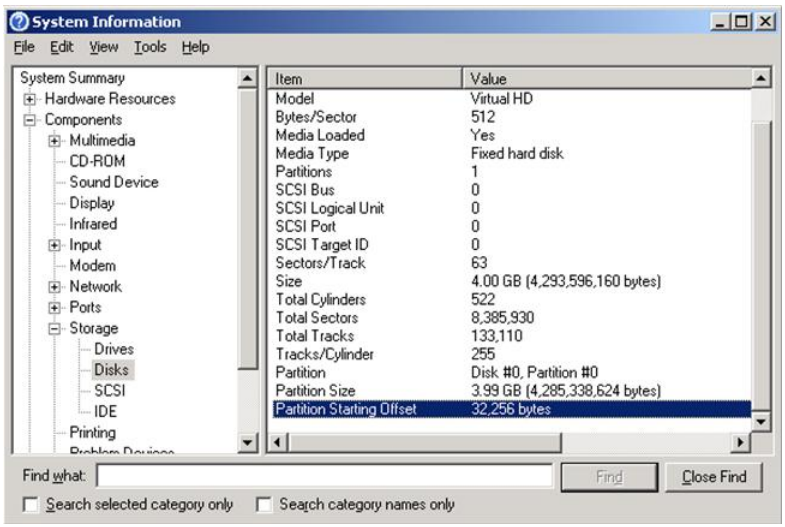

Now open a calculator and divide 32,256 by 4,096. The result is 7.875. Because this value is not even, you know that this virtual machine has a misaligned file system. When the result is an even number, your partition is properly aligned.

If you want to retrieve the information for more virtual machines, you can use Windows PowerShell to query the information via WMI. This command will also calculate to determine if the partition is aligned:

```
gwmi -class win32_diskpartition -Computername VM | select name,startingoffset,
@{label="Aligned";express={if($_.startingoffset%4096 -eq 0){$true} else {$false}}}
```

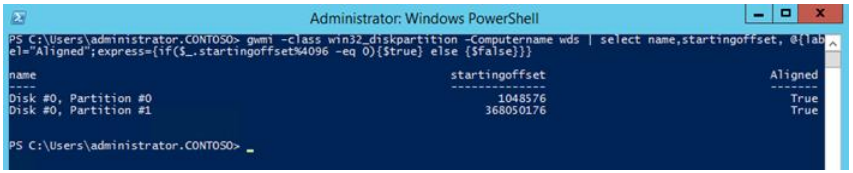

So how can you correct this issue? The sad answer is that there is no easy fix. The only way to get rid of the file system misalignment is to create a backup inside the virtual machine. Create a new VHD with the correct alignment, and restore your backup.

To create a virtual hard disk with the correct alignment, follow these steps:

1. Boot a virtual machine with a Windows PE or use your Windows Installation DVD.
2. Launch Diskpart.exe.
3. Select **Disk**.
4. Create Partition Primary align=32.

Keep in mind when you do a physical-to-virtual (P2V) conversion of a machine running an operating system prior to Windows Vista/Windows Server 2008, the partition does not get properly aligned. So instead of doing a P2V of a machine running an old operating system, you should plan to migrate your applications to a newer version of Windows.

—Thomas Roettinger, Program Manager, Partner and Customer Ecosystem Team

Additional resources

Here is an additional resource concerning this topic:

- Advanced format (4 KB) disk compatibility update at: http://msdn.microsoft.com/en-us/library/windows/desktop/hh848035(v=vs.85).aspx

Virtual disk fragmentation

Fragmentation can cause performance problems on both physical and virtual hard disks. Troubleshooting fragmentation on virtual hard disks can be difficult sometimes. Carlos Mayol Berral demonstrates this with an example.

Large file size records, dynamic disks, differencing disks, and problems managing highly fragmented files

Dynamic and differencing VHDs both use an NTFS file system feature called sparse files.

What is a sparse file?

Sparse files provide a method of saving disk space for files that contain meaningful data, as well as large sections of data composed of zeros. If an NTFS file is marked as sparse, NTFS allocates disk clusters only for the data explicitly specified by the application.

> ***MORE INFO*** **You can read more about how NTFS works at** http://technet.microsoft.com/en-us/library/cc781134(v=ws.10).aspx.

What is the MFT?

When you format a volume with NTFS, Windows creates a *master file table* (MFT) and metadata files on the partition. The MFT is a relational database that consists of rows of file records and columns of file attributes. It contains at least one entry for every file on an NTFS volume, including the MFT itself. In other words, the MFT stores the information required to retrieve files from the NTFS partition.

Why should I care?

NTFS creates a file record for each file and a folder record for each folder created on an NTFS volume. The MFT includes a separate file record for the MFT itself. These file and folder records are 1 KB each and are stored in the MFT.

When a file is very fragmented, NTFS uses more space to save the description of the allocations that is associated with the fragments. The allocation information is stored in one or more file records.

When there is no more space for storing attributes in the file record segment, additional file record segments are allocated and inserted in the first (or base) file record segment in an attribute called the *attribute list*. The number of ATTRIBUTE_LIST_ENTRY structures that the file can have is limited.

MORE INFO **You can read more about MFT at http://msdn.microsoft.com/en-us/library/bb470206(v=vs.85).aspx.**

By default, the size of file record segment (FRS) is 1 KB and is represented as Bytes Per File-Record Segment in the last line of the output of the fsutil.exe command-line tool as shown here:

```
C:\>fsutil fsinfo ntfsinfo c:
NTFS Version   :                 3.1
LFS Version    :                 2.0
Bytes Per Cluster :              4096
Bytes Per FileRecord Segment    : 1024
```

In some situations you can have problems with highly fragmented VHDs and might see errors like this:

```
"The requested operation could not be completed due to a file system limitation"
0xC0000427
STATUS_FILE_SYSTEM_LIMITATION
```

For Windows Server 2008 R2, you can read a Knowledge Base article about this issue at http://support.microsoft.com/kb/967351. The article explains the cause and introduces a new functionality for the format.exe command.

On Windows Server 2012 this new functionality is included by default, and you can now use format.exe to increase the size of file size records, a capability called *Large File Size Records*. The only way to do this is by formatting the volume using format.exe /L, which the Help file explains like this:

```
/L              NTFS Only: Use large size file records.
                By default, the volume will be formatted with small size file records.
```

No other tools such as Disk Manager or diskpart.exe have the ability to do this.

When you use this parameter to format a volume, you can see the result using fsutil:

```
fsutil fsinfo ntfsinfo v:
NTFS Version                   : 3.1
LFS Version                    : 2.0
Bytes Per Cluster              : 4096
Bytes Per FileRecord Segment   : 4096
```

You should consider doing this if you plan to have a very high number of files on the same volume and are using the previously mentioned virtual hard disk format types, which increase the probability to have very large fragmented files.

—Carlos Mayol Berral, Premier Field Engineer

Migrating VHD to VHDX

The new VHDX virtual hard disk format was introduced in Windows Server 2012 to accommodate the growing storage needs of enterprises that use virtual environments. VHDX has numerous benefits over the older VHD format including greater storage capacity and enhanced data protection. In this section Carlos Mayol Berral goes into more detail concerning some of these benefits and how to optimize the conversion of VHD to VHDX using the new block and sector sizes.

I migrated my virtual machines. Now what?

Normally with a Hyper-V migration one is occupied with issues like migrating the virtual machines, the configuration, and the network and storage fabrics. But you also need to be aware of the new VHDX functionalities and what they can offer your environment. Your final migration step should then be to migrate your VHD to VHDX files.

I would say that the key benefits of VHDX are in resiliency and performance, namely:

- VHDX provides a form of transactionality for metadata updates.
- VHDX has better performance than VHD.
- VHDX lets you tune your virtual hard disk files to match the physical disks in your storage fabric by providing additional options for block and sector size.
- VHDX includes embedded protection against data corruption by logging updates to the VHDX metadata structures, which can help prevent corruption due to power failures.

In general, Microsoft has always said that fixed virtual disks provide better performance than dynamically expanding virtual disks for virtual machines running on Hyper-V hosts. This is still valid, but one of the greatest challenges in our industry is to contain storage growth. On this issue, Microsoft has improved VHDX files so they are better aligned with the underlying physical storage, and this is the reason why we can now say that you can use dynamic disks in production environments and expect the same or even better performance compared with fixed disks.

Comparing VHDX and VHD performance

The following two charts show some of the performance improvements provided by the new VHDX format. The first chart shows that VHDX provides about 10 percent improvement in random write performance over VHD when using dynamic disks:

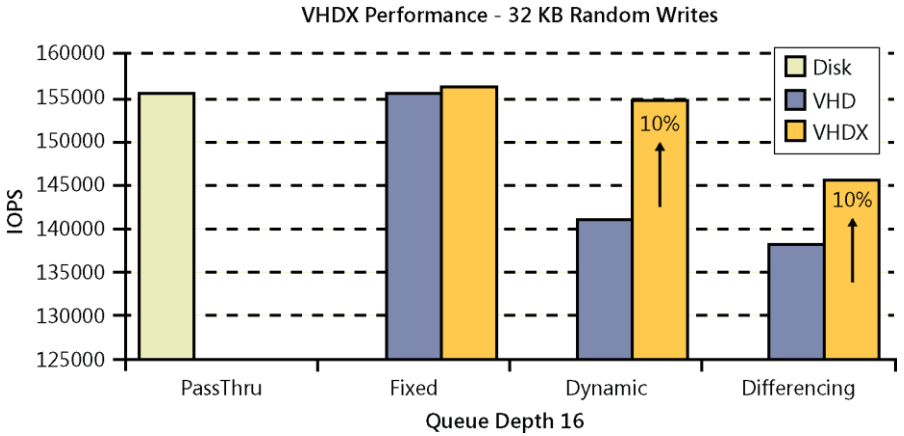

The second chart shows that VHDX has an even greater 25 percent improvement for sequential write performance over VHD when using dynamic disks:

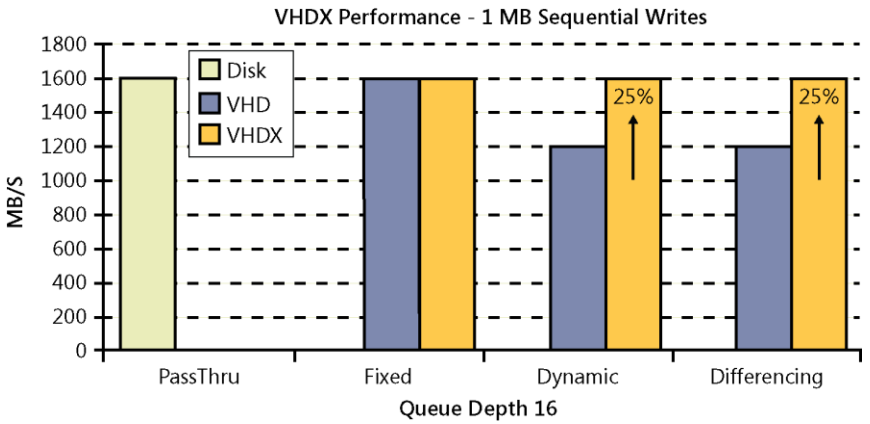

Migrating from VHD to VHDX

Migrating VHD files to the new VHDX format can be accomplished using either the Hyper-V console or with Windows PowerShell. Using the GUI interface you can use the option of creating a new disk and select the Create-from-Source option.

Using Windows PowerShell, you can use this cmdlet:

```
Convert-VHD:
```

The Convert-VHD cmdlet converts a virtual hard disk file by copying the data from a source virtual hard disk file to a new virtual hard disk file of a specified format and version type. Conversion is an offline operation; the virtual hard disk must not be attached when the operation is started. Example:

```
Convert-VHD –Path c:\test\testvhd.vhd –DestinationPath c:\test\testvhdx.vhdx
```

Optimizing block and cluster sector sizes

Another improvement is that VHDX formats can now support new block and cluster sector sizes. Block sizes can be determined during the conversion of the old VHD as well as during creation of a new VHDX dynamic disk. The parameter –BlockSizeBytes determines this value. This value can be 1 MB, 2 MB, 4 MB, 8 MB, 16 MB, 32 MB, 64 MB, 128 MB, or 256 MB.

It is optimal to match the block size to the allocation patterns of the workload using the disk. The default block size for dynamic disk was increased from 2 MB on VHD format to 32 MB on VHDX format.

To determine what block size you have for a specific VHDX file, you can use Windows PowerShell like this:

```
Get-VHD –Path c:\test\testvhdx.vhdx
VhdFormat               : VHDX
VhdType                 : Dynamic
FileSize                : 4194304
Size                    : 1073741824
LogicalSectorSize       : 512
PhysicalSectorSize      : 4096
BlockSize               : 33554432
FragmentationPercentage : 0
Alignment               : 1
```

The other important aspect is the disk sector size, and in the hardware world we have:

- **Logical Sector** The unit that is used for logical block addressing for the media. You can think of this as the smallest unit of write that the storage device can accept. This is the "emulation."
- **Physical Sector** The unit for which read and write operations to the device are completed in a single operation. This is the unit of atomic write.

Nowadays the hard disk industry standard is moving to 4-KB sector hard drives, and Windows 2012 and Windows 8 are the first Microsoft operating systems that support 4-KB native disks. For more information, see http://support.microsoft.com/kb/2510009.

If you are using 4-KB sector hard drives with an earlier version of Windows, you need to use Advance Format or 512e to use it. This will lead to a bad situation, however, because it means a process called Read-Modify-Write will be used which causes poor performance for virtual hard disks.

Windows Server 2012 mitigates some of the performance impact of using 512e disks on the VHD stack by preparing the previously mentioned structures for alignment to 4-KB boundaries in the VHD format. This alignment is another good reason to migrate your old VHD virtual disks to VHDX.

You can use the fsutil tool to determine whether your hard drive is the 512e or 4-KB native type.

```
C:\windows\system32>fsutil fsinfo  ntfsinfo f:
NTFS Version                   :  3.1
LFS Version                    :  2.0
Bytes Per Sector               :  512
Bytes Per Physical Sector      :  512
Bytes Per Cluster              :  4096
```

For VHDX file creation, these values are represented as LogicalSectorSize and PhysicalSectorSize. By default, VHDs are exposed with a LogicalSectorSize of 512 bytes for application compatibility.

Note also that the default sector size of a VHDX is 4 KB. When installing Windows Server 2008 R2, you need to install a hotfix or convert the VHDX to 512-byte sector size.

In this section I have described only a few of the benefits of the new VHDX format. I hope you do not avoid the migration of your virtual hard disk files for your next migration project.

—Carlos Mayol Berral, Premier Field Engineer

Additional resources

Here are a few additional resources concerning this topic:

- Hyper-V Virtual Hard Disk Format Overview (TechNet Library) at: http://technet.microsoft.com/en-us/library/hh831446.aspx
- Convert-VHD (TechNet Library) at: http://technet.microsoft.com/en-us/library/hh848454.aspx
- An update that improves the compatibility of Windows 7 and Windows Server 2008 R2 with Advanced Format Disks is available (Microsoft Support) at: http://support.microsoft.com/kb/982018
- Performance Tuning Guidelines for Windows Server 2012 (MSDN) at: http://msdn.microsoft.com/en-us/library/windows/hardware/jj248719.aspx

Monitoring storage performance

To truly know whether you've managed to optimize Hyper-V hosts and the virtual machines running on them, you need to compare their performance before and after the configuration changes you've made to them. The in-box tool for doing this on the Windows Server platform is Performance Monitor. Thomas Roettinger quickly reviews how to use this tool and summarizes some key performance counters that you might want to consider monitoring.

Using Performance Monitor

The Windows operating system offers performance counters for nearly all different components. You can gather performance data from these counters by using WMI or Performance Monitor. In this section I will show you how to use Performance Monitor to capture performance data for all important storage components and also present thresholds that will help you to understand if there is a potential problem.

To start Performance Monitor, simply type **perfmon** in the modern UI:

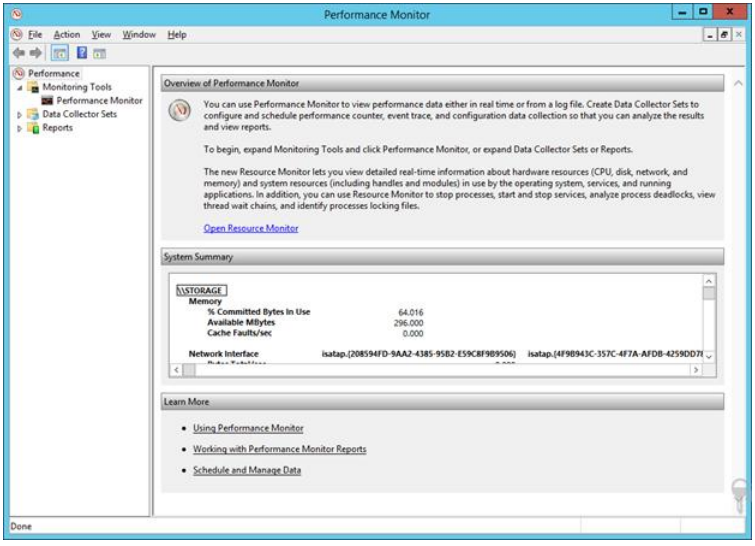

When you click Performance Monitor, you see a real-time view of your system. However, the real-time view does not provide enough data to determine if there is a potential problem. You need to capture performance data over a longer period of time. Microsoft recommends at least 24 hours or even more so that you capture a whole business day or a business week. For example, in an environment that is only used in one given time zone you expect to see load picking up in the morning and going down in the evening. If you host a shared environment for multiple tenants it's even more important to capture data over a longer period of time because usually you have no detailed information about the usage of the virtual machines that belong to the tenants.

For capturing performance data over a longer period of time you need to set up a data collector set. To do so, expand Data Collector Sets, click User Defined, and right-click in the pane on the right to create a new data collector set:

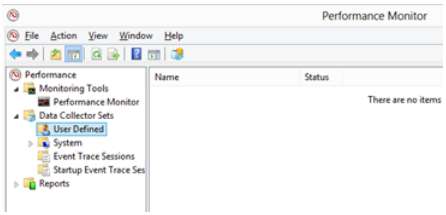

Specify a name for your data collector set and select Create Manually:

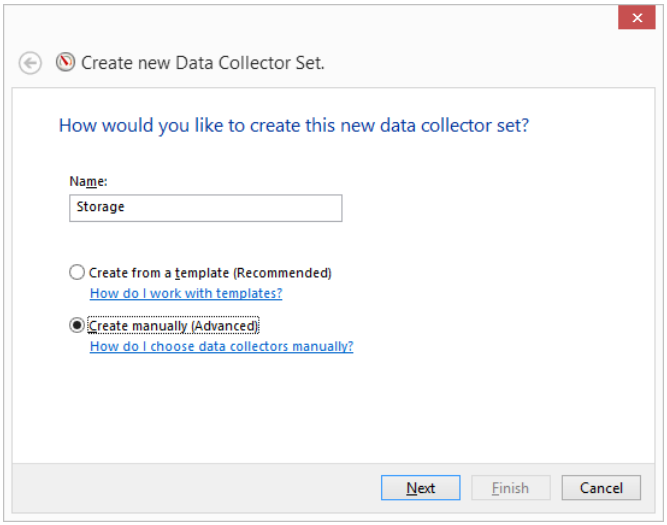

Next, indicate that you want to include performance counters in that collector set by selecting Performance Counter under Create Data Logs:

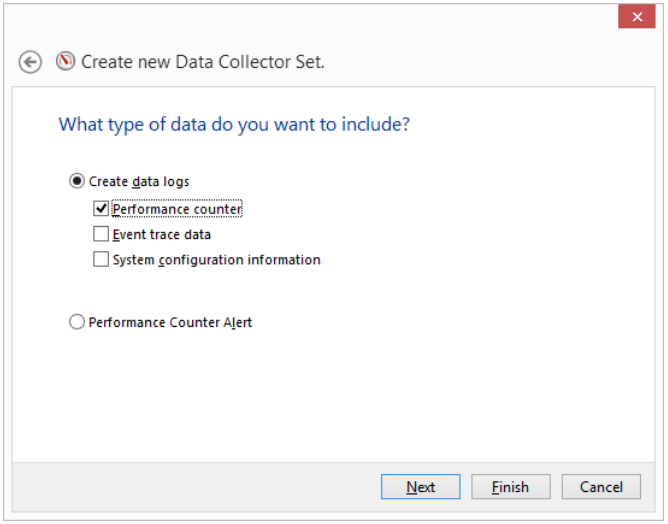

Next, select the appropriate performance counters for storage. Walking through the following examples will help you understand the performance counters so that you can later use them in a data collector set. But first, you should know how to start and stop a data collector set and how to load and analyze data. Notice the green arrow and the stop symbol in the following screenshot:

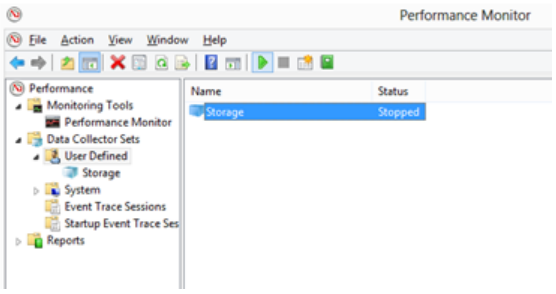

You could also use options in the data collector set properties to schedule the data collector set to run automatically.

To load a data collector set, go to Performance Monitor and right-click the data collector set to open its Properties, and then click the Source tab. There you can specify to load captured data from a log file.

When the file is loaded you also have the option to limit the data that is shown to a specific time window:

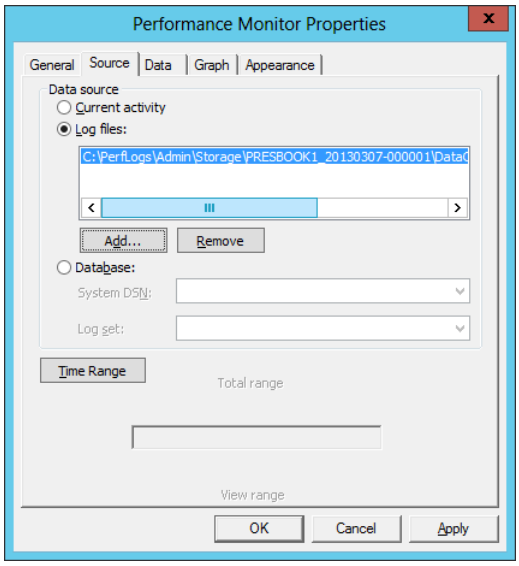

After the file is loaded you can add the counters you captured by clicking the green plus sign control and start investigating your problem.

To make life easier, there is a tool called Performance Analysis of Logs (PAL) available at http://pal.codeplex.com. This tool contains a template with counters and thresholds for various Microsoft Windows Roles, as well as Exchange, SQL, and many others.

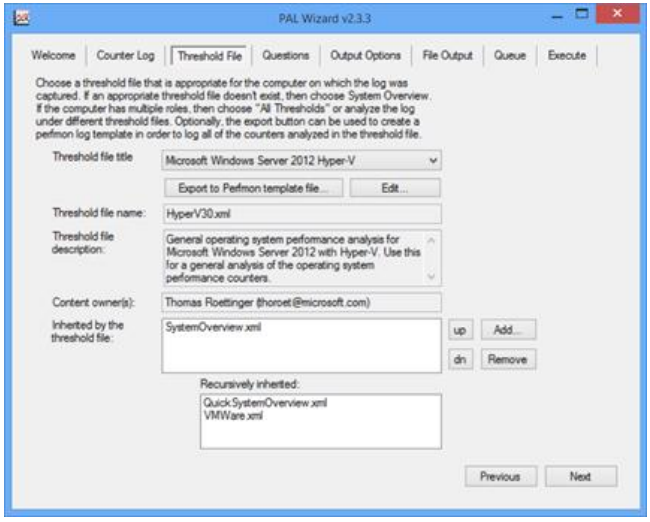

After exporting a template from PAL you can import it to a data collector set. The log file that you then get from the data collector set created from your performance data then needs to be imported into PAL. PAL then analyzes the log file and creates an HTML report with all the findings. Give it a try!

Storage performance counters

The sections below summarize some key performance counters you can track for monitoring the following storage system components:

- Hard disks
- SMB Client
- SMB Server

Hard disk counters

To measure the storage subsystem of either local or SAN-attached disks, use these performance counters:

- **LogicalDisk AVG. Disk sec/Read** This counter is the average time in seconds of a read of data to the disk. Thresholds are greater than 0.015 (15 ms) and greater than 0.025 (25 ms). Spikes above 25 ms are normal.
- **LogicalDisk AVG. Disk sec/Write** This counter is the average time in seconds of a write of data to the disk. Thresholds are greater than 0.015 (15 ms) and greater than 0.025 (25 ms). Spikes above 25 ms are normal.
- **LogicalDisk Disk Transfers/sec** This counter measures the I/O of your disk. Thresholds depend on disk type, e.g., 7,200 rpm SATA can do about 75 to 100 IOPS and 15,000 rpm SAS can do about 175 to 200 IOPS.
- **LogicalDisk Disk Read Bytes/sec** This counter measures how much data was read to the disk in bytes per second. Values depend on the type of disk subsystem.
- **LogicalDisk Disk Write Bytes/sec** This counter measures how much data was written to the disk in bytes per second. Values depend on the type of disk subsystem.
- **Hyper-V Virtual Storage Device Error Count** This counter represents the total number of errors. This should be 0.
- **Hyper-V Virtual Storage Device Write Bytes/sec** This counter measures how much data was written to a virtual device. Performance depends on the physical disk system.
- **Hyper-V Virtual Storage Device Read Bytes/sec** This counter measures how much data was read from a virtual device. Performance depends on the physical disk system.

SMB Client counters

To measure SMB Client performance, use these counters:

- **SMB Client Share Read Bytes/sec** This counter measures how much data was read from a share. Performance depends on various components (disk, network, adapter, RSS).
- **SMB Client Share Write Bytes/sec** This counter measures how much data was written to a share. Performance depends on various components (disk, network, adapter, RSS).
- **SMB Client Share Current Data Queue Length** This counter measures the current queue depth.

SMB Server counters

To measure SMB Server performance, use these counters:

- **SMB Server Share Read Bytes/sec** This counter measures how much data was read from a share. Performance depends on various components (disk, network, adapter, RSS). The value should match what you got from the client.
- **SMB Server Share Write Bytes/sec** This counter measures how much data was written to a share. Performance depends on various components (disk, network, adapter, RSS). The value should match what you got from the client.
- **SMB Server Share Current Data Queue Length** This counter measures the current queue depth.
- **SMB Server Sessions** This counter measures information related to a specific session.
- **SMB Direct Counters** This counter measures information related to monitoring SMB Direct connections. You should measure SMB when it is transported over RDMA.
- **RDMA Activity** This counter measures information related to RDMA, including errors, connections, and performance.

Example: Troubleshooting a storage problem using Performance Monitor

Patricia is an administrator who is notified that several departments are reporting poor performance of their services. She logs on to the Hyper-V host that she knows is running virtual machines from the departments in question and opens Performance Monitor.

She adds the counter for the physical disk first to determine if there is a storage-related problem:

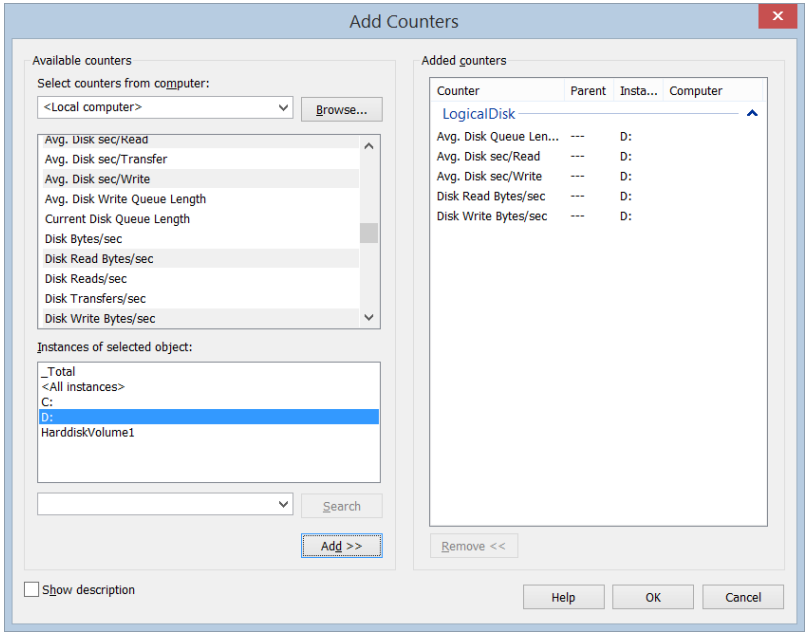

She uses Performance Monitor to monitor the Disk Read Bytes/sec counter for drive D:

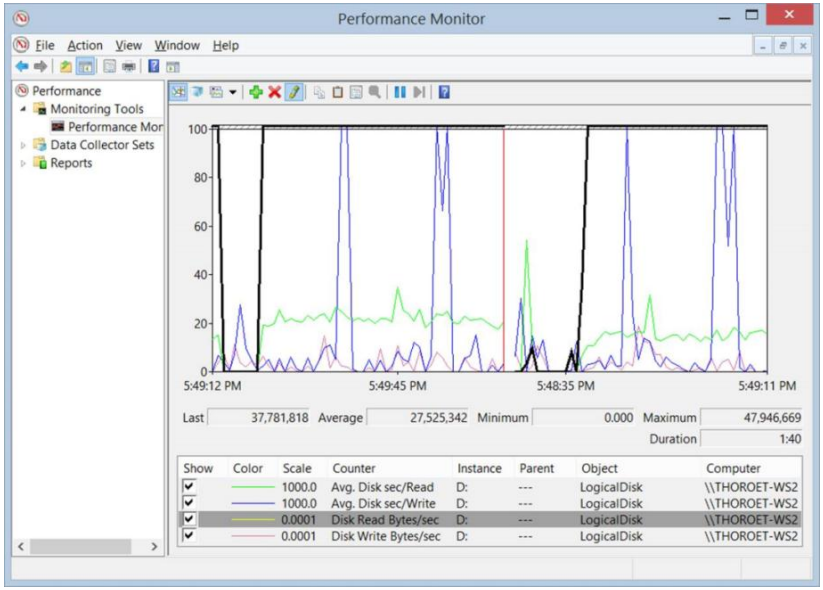

She also examines the Avg. Disk sec/Read counter for drive D:

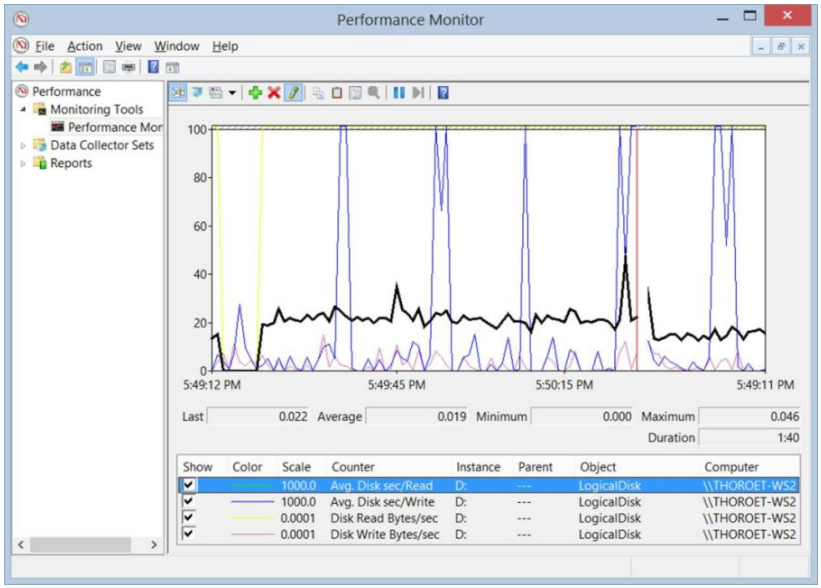

Patricia notices that drive D has an average response time over 0.025 ms. This information indicates that there is a disk problem. She now wants to determine if a specific virtual machine is causing the disk load. So she deletes the previously added counters and adds the counter for virtual hard disks:

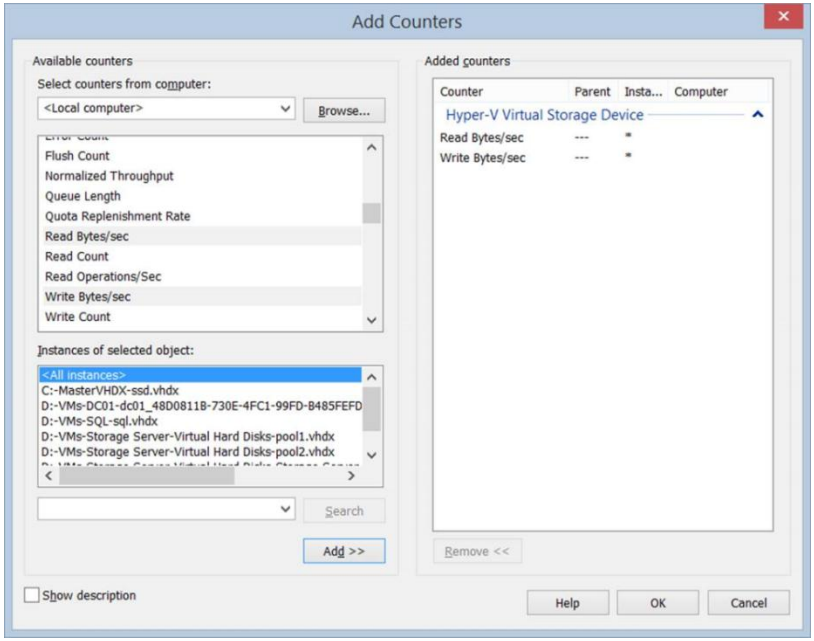

She examines Read Bytes/sec for each of the two virtual hard disks:

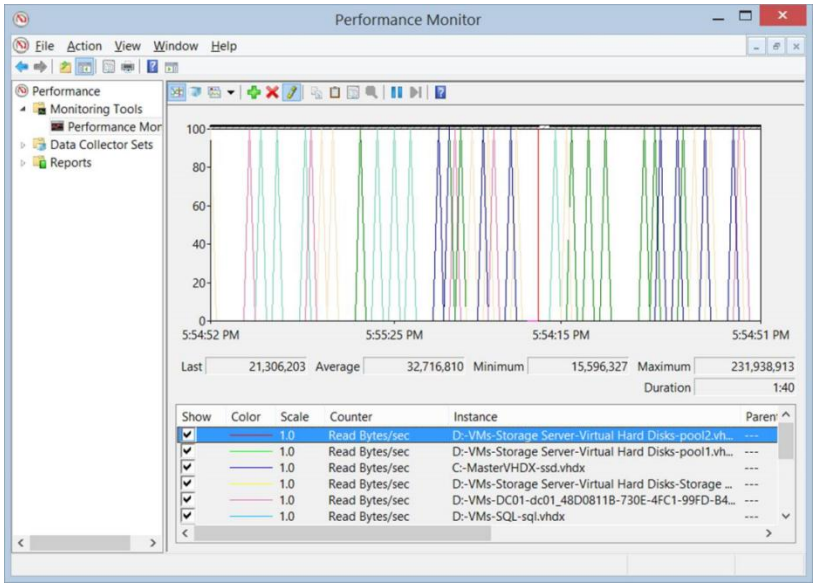

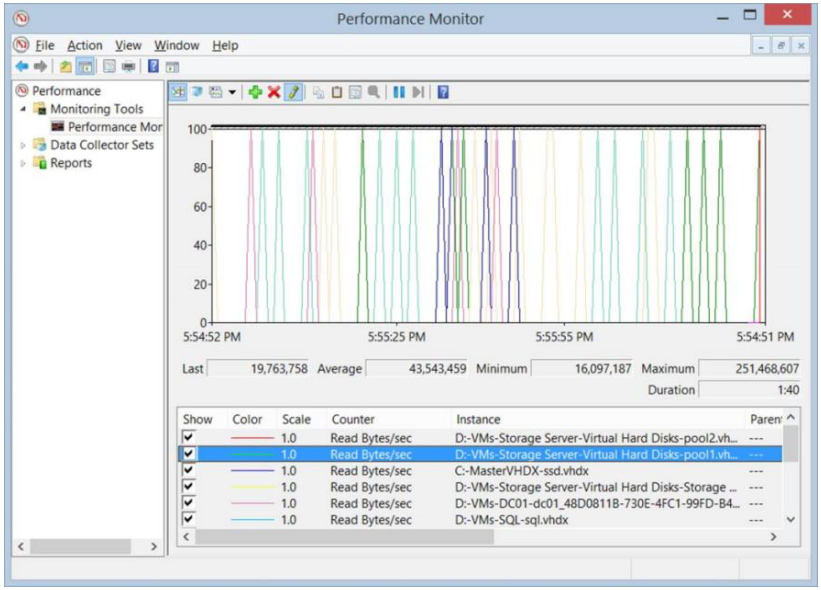

She discovers that two virtual hard disks have values for read bytes per seconds that are nearly matching the physical disk read bytes per second of the physical drive D. Patricia is able to identify the virtual hard disks that are causing the issue, and she knows the virtual machine to which these disks belong.

She can now continue troubleshooting the guest operating system or live migrate the virtual hard disks to another physical disk without interrupting the service.

—Thomas Roettinger, Program Manager, Partner and Customer Ecosystem Team

Additional resources

Here are a few additional resources concerning this topic:

- Windows Performance Monitor (TechNet Library)at: http://technet.microsoft.com/en-us/library/cc749249.aspx
- Performance Tuning Guidelines for Windows Server 2012 at: http://msdn.microsoft.com/en-us/library/windows/hardware/jj248719.aspx

Cluster Shared Volumes

Cluster shared volumes (CSV) is a feature of Failover Clustering that simplifies the configuration and management of clustered virtual machines running on Hyper-V hosts. CSV works by allowing multiple clustered virtual machines to use the same logical unit number (LUN) or disk while still being able to fail over independently of one another.

Windows Server 2012 delivers new capabilities for CSV including:

- CSV support for a scale or file server
- Improved Backup/Restore
- Multi-IP Subnet support
- Support for BitLocker encryption
- Direct I/O for more scenarios
- Block Level I/O redirection
- Support for memory mapped files

Using CSV with clustered virtual machines running on Hyper-V hosts is a big topic, so this section deals with it in two ways. First we have Thomas Roettinger who walks us through several different CSV troubleshooting scenarios. Then we have Subhasish Bhattacharya who explains the CSV data flow and how to optimize the CSV block-level cache that is new in Windows Server 2012.

CSV Redirected Access mode

Cluster shared volumes (CSV) were introduced with Windows Server 2008 R2 and have been enhanced in Windows Server 2012. CSV allows each node that is part of the same Windows failover cluster to access the same disk (LUN) at the same time. CSV allows virtual machines to fail over independently.

Without CSV, a disk (LUN) can be accessed by only one cluster node. CSV is a distributed file system that provides a scalable fault tolerant solution while using the NTFS file system.

The following are requirements for using CSV in a cluster:

- Drive letter for the system disk must be equal on all nodes.
- NLTM authentication must be enabled on all nodes.
- All nodes must use the same IP subnet (Windows Server 2008 R2 only).
- Multi-site clusters require stretched VLANs (Windows Server 2008 R2 only).

CSV is implemented through a mini filter driver that differentiates between redirected I/O and direct I/O. This is very important because even with CSV enabled there is still one node in a cluster that owns the LUN. This node is called the coordinator node. Redirected I/O will be send to the coordinator node over the network, and the coordinator node will send it to the disk via FC, ISCSI, or SAS using the direct storage link.

CSV offers fault tolerance even when a node does lose the direct storage link to a disk (LUN). In such a case, direct I/O is also redirected over the network to the coordinator node. The cluster service assigns each network a metric. The network with the lowest metric will be used for redirected I/O. A cluster node that has lost its direct storage link marks the CSV to be in redirected access.

There are basically four reasons why a CSV might be in redirected access mode.

- The administrator enabled redirected access on purpose (maintenance).
- The direct storage link was lost.
- An incompatible filter driver was installed on a node.
- A backup of a CSV volume is in progress or failed.

Example: Network for redirected I/O

Patricia is an administrator who needs to replace an ISCSI fabric switch. This change will impact one of her Hyper-V cluster nodes (2012N2).

She puts the CSV disk in maintenance mode. The following Event ID is written to the system log.

```
Event ID 5136
Cluster Shared Volume 'Volume1' (Cluster Disk 2) redirected access was turned on. Access to the storage device will be redirected over network from all cluster nodes that are accessing this volume. This may result in degraded performance. Turn off redirected access for this volume resume normal operations.
```

She also transfers the CSV disk ownership to the other node (2012N1):

```
Get-ClusterSharedVolume
```

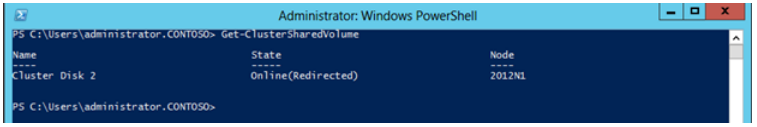


```
Move-ClusterSharedVolume  -name "cluster disk 2" -node 2012N1
```

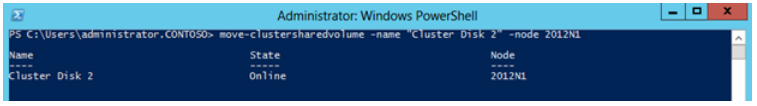

While Hyper-V cluster node 2012N2 is without a direct storage link, Patricia checks which network is used for redirected I/O. She discovers that only one network was being used. Because of the new SMB Multichannel feature in Windows Server 2012, she had expected that two networks would have been used:

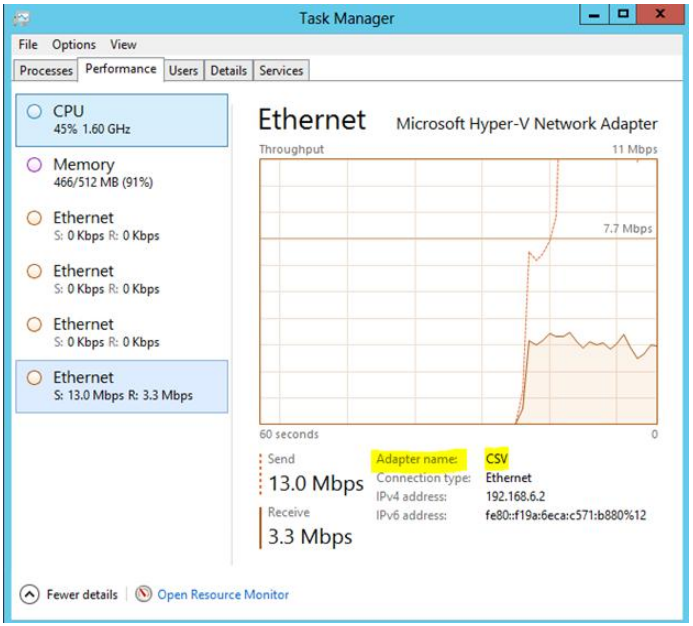

Next Patricia checks the cluster network metrics:

```
Get-clusternetwork | ft name,metric,autometric
```

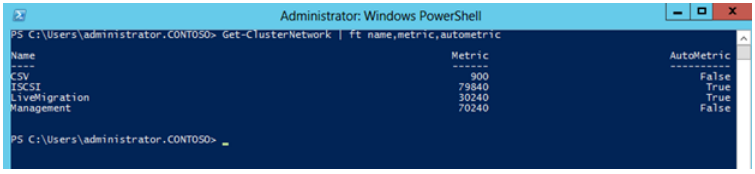

She discovers that AutoMetric is turned off for the network called CSV and also that the lowest metric is assigned to it.

Now, for a Windows Server 2008 R2 cluster, this would be the perfect setting. The cluster is using the network that was implemented for redirect I/O. However, in Windows Server 2012, the algorithm for assigning metrics was adjusted to make use of the new SMB 3.0 Multichannel feature.

Patricia turns on AutoMetric for the CSV network and lists the assigned metrics again:

```
(get-clusternetwork "CSV").AutoMetric=$true
```

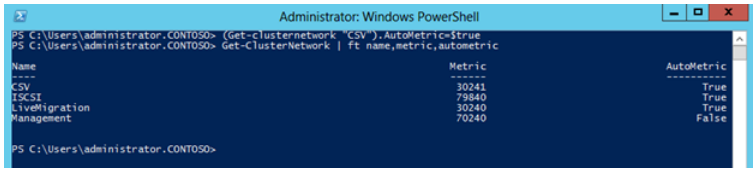

Immediately the cluster is using two networks via SMB Multichannel and as a result the performance of the virtual machines increases while the CSV disk (LUN) is still in redirected access:

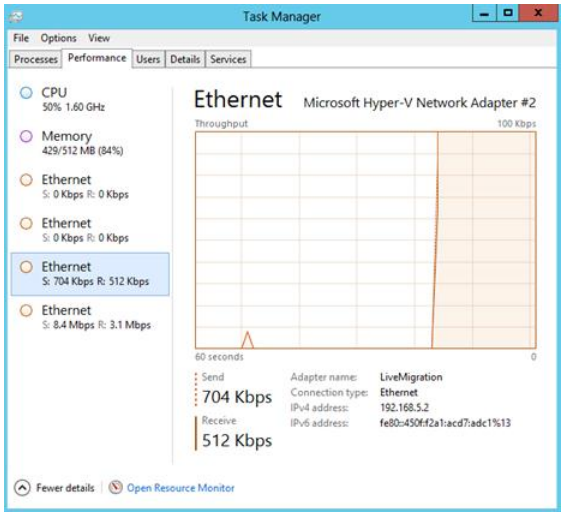

Example: Lost direct storage link

After Patricia has replaced an ISCSI fabric switch, and everything is working as expected. But one hour later, she is notified that redirected access is turned on again.

Patricia opens Event Viewer and discovers the following event:

```
Event ID 5121
Cluster Shared Volume 'Volume1' (Cluster Disk 2)' is no longer directly accessible from this cluster node. I/O access will be redirected to the storage device over the network to the node that owns the volume. If this results in degraded performance please troubleshoot this node's connectivity to the storage device and I/O will resume to a healthy state once connectivity to the storage device is reestablished.
```

She also discovers the following event entry:

```
Event ID 20
Connection to the target was lost. The initiator will attempt to retry the connection.
```

This allows her to understand that the ISCSI link is broken.

Patricia verifies the TCP/IP connection to the ISCSI Target:

```
Ping 192.168.11.1
```

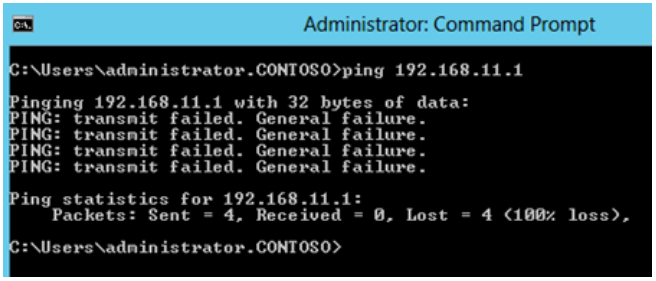

There is no response from the ISCSI Target and the error message from the ping command suggests that the ISCSI Network Adapter is disconnected:

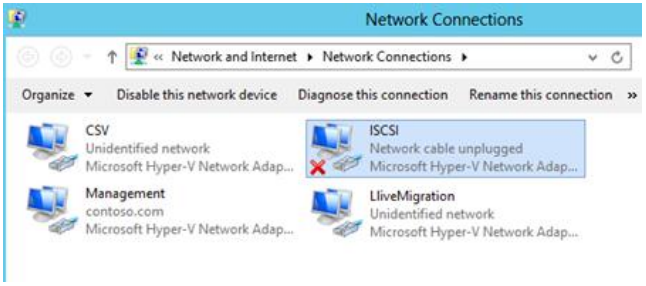

Example: Failed backup

When a backup software does create a backup of a Hyper-V host that stores virtual machines on a CSV disk (LUN) it needs to put the CSV disk (LUN) in redirected access. Depending on what kind of snapshot provider is used, software or hardware it will influence the time the CSV disk (LUN) is in redirected access. When the software snapshot provider is used, the CSV disk (LUN) is in redirected access for the entire time it takes to complete a backup.

If a hardware snapshot provider is installed, the CSV disk (LUN) is only in redirected access as long as it takes to create a VSS snapshot. This behavior was slightly changed in Windows Server 2012 so that the CSV disk (LUN) is in redirected access only for the time it takes to create the VSS snapshot even with the software snapshot provider.

When the backup completes, the backup application must inform the cluster to bring the CSV disk (LUN) out of redirected access.

The administrator Patricia discovers that one of her CSV disks (LUN) is in redirected access and shows "Online (Backup in progress)."

She first opens the backup software to determine if there really is a backup job running, but there is no active job.

Patricia tries to bring the CSV disk (LUN) out from redirected access by simulating a failover of the CSV disk (LUN).

```
Test-ClusterResourceFailure "Cluster Disk 2"
```

Simulating a failover of the CSV disk (LUN) does not solve the problem. If that had worked Patricia would need to inform the vendor of the backup software that they do not clear the state of the CSV by using the cluster API.

Patricia's next step is to determine if there is an existing VSS software snapshot for the CSV disk (LUN) and delete it. This would indicate that the backup job failed after the snapshot was created.

For Windows Server 2012, Patricia uses Diskshadow.exe and run these two commands:

```
List shadows all
Delete shadows all
```

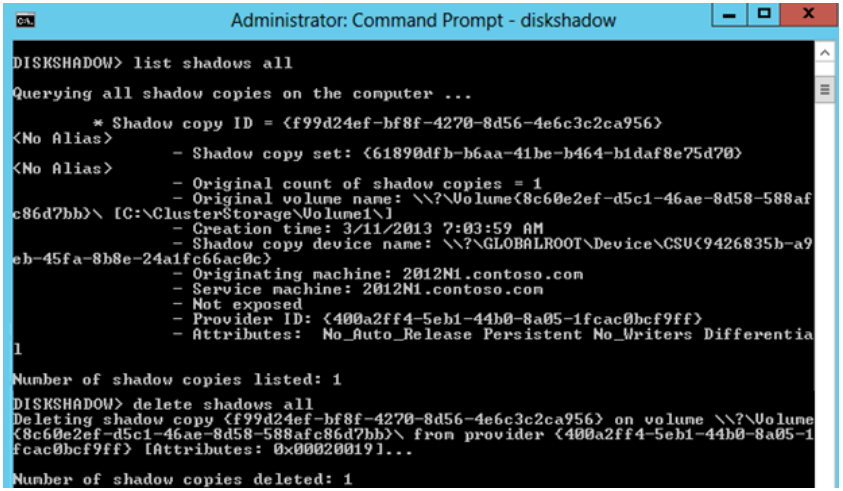

And for Windows Server 2008 R2, she uses Vssadmin.exe like this:

```
Vssadmin list shadows
Vssadmin delete shadows /shadowID
```

After deleting the VSS snapshot, the CSV disk (LUN) comes out of redirected access immediately.

Example: Incompatible filter driver

Cluster shared volumes are implemented through a mini filter driver. If there is an incompatible filter driver such as an antivirus software filter driver, it may end up in redirected access for a CSV disk (LUN).

Patricia discovers that one of her CSV disks (LUN) is in redirected access. After verifying that there is no broken storage link, no backup issue, and that nobody turned on redirected access for maintenance, she checks the event logs.

She discovers the following event entries:

```
Event ID 5125
Cluster Shared Volume '%1' ('%3') has identified one or more active filter drivers on
this device stack that could interfere with CSV operations. I/O access will be
redirected to the storage device over the network through another Cluster node. This may
result in degraded performance. Please contact the filter driver vendor to verify
interoperability with Cluster Shared Volumes.
Event ID 5126
Cluster Shared Volume '%1' ('%3') has identified one or more active volume drivers on
this device stack that could interfere with CSV operations. I/O access will be
redirected to the storage device over the network through another Cluster node. This may
result in degraded performance. Please contact the volume driver vendor to verify
interoperability with Cluster Shared Volumes.
```

Patricia runs the fltmc.exe utility to check for installed filter drivers:

```
Fltmc.exe
```

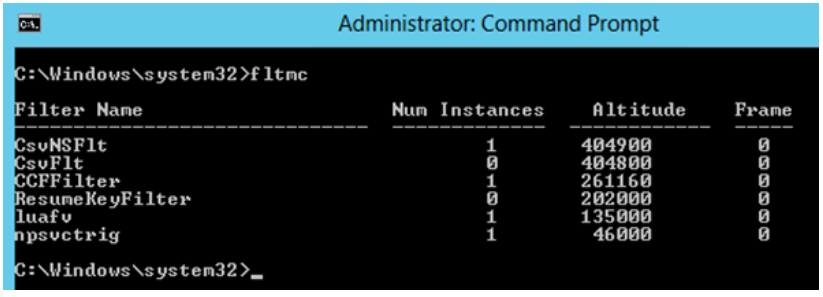

(Note that the above screenshot is from a system running no third-party software.)

Patricia checks each filter driver with the help of the File System Minifilter Allocated Altitudes spreadsheet that she downloaded from the Windows Hardware Developer Central website at http://www.microsoft.com/whdc/driver/filterdrv/alloc-alt.mspx.

After she identifies the responsible filter driver and uninstalls the driver, the CSV disk (LUN) comes out of redirected access.

Using CSV performance counters

When using Performance Monitor to check the performance of a CSV disk (LUN) in Windows Server 2008 R2, you need to use the PhysicalDisk counters. Specifically, use these counters:

- Disk Read Bytes/sec
- Disk Write Bytes/sec
- Disk Transfer/sec (I/O)
- Avg Disk sec/Read
- Avg Disk sec/Write

In Windows Server 2012, however, you now have specific counters for cluster CSV. These counters allow you to monitor your CSV disk (LUN) for more specific information. These counters include:

- Cluster CSV Block Redirection
- Cluster CSV Volume Manager
- Cluster CSV Coordinator
- Cluster CSV File System
- Cluster CSV Volume Cache

For example, the Cluster CSV Coordinator counter will show, for instance, how many create file actions happen or how much metadata is created. The Cluster CSV File System counter can be used to look for redirected bytes, read/write bytes per sec, and queue length. The Cluster CSV Volume Cache counter will only show values if you enabled CSV Cache. These are very powerful in pooled VDI scenarios. For more information, see the topic by Subhasish Bhattacharya in the next section.

The important take away for both operating systems is that using PhysicalDisk or Cluster CSV counters gives you a view from only a single node that is accessing the CSV disk (LUN). You therefore need to aggregate the data from all hosts accessing the same CSV disk (LUN).

—Thomas Roettinger, Program Manager, Partner and Customer Ecosystem Team

Exploring Cluster Shared Volume data flow

Cluster shared volumes (CSV) is a clustered file system in Windows Server 2012. It enables all servers in a failover cluster to access a common NTFS volume by providing a layer of abstraction above NTFS. The NTFS volume is mounted on a cluster node referred to as the coordinator node.

Metadata operations

CSV, like all clustered file systems, needs a mechanism to orchestrate metadata updates. CSV synchronization is done on the server side and therefore avoids I/O interruptions. On non-coordinator nodes, metadata operations are sent over SMB (and therefore the network) to the disk mounted on the coordinator node. Metadata operations are small and infrequent and occur in scenarios such as virtual machine creation, virtual machine power on/off, and backup (snapshot creation):

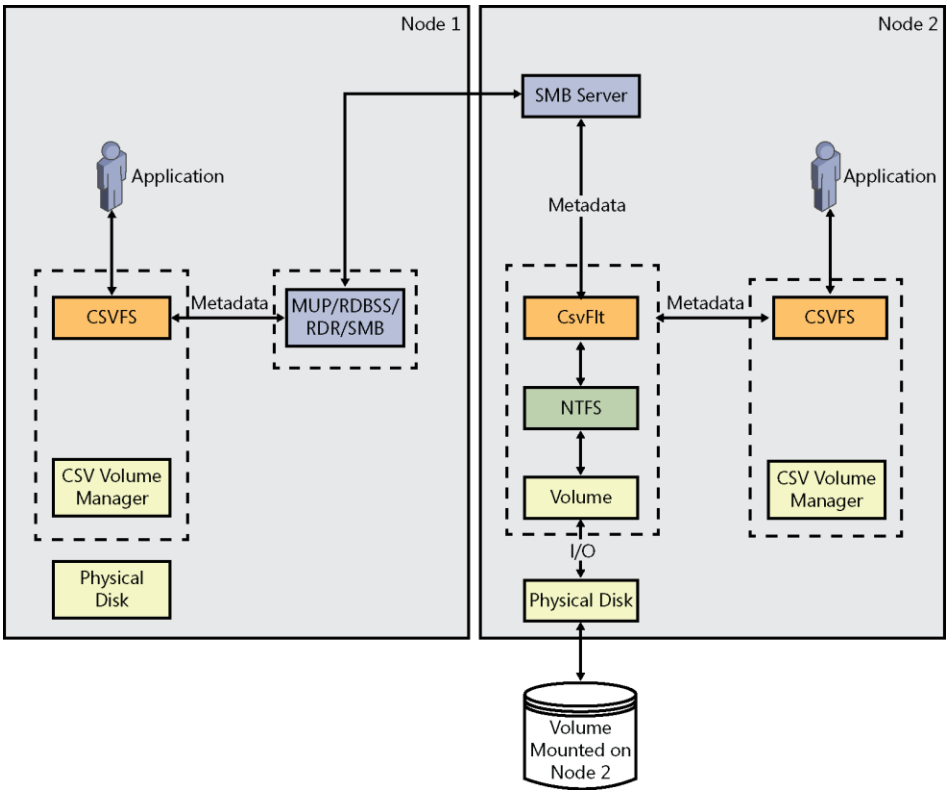

Direct I/O operations

For most CSV operations, when the cluster node has connectivity to the storage, I/O can be sent directly to the storage. It therefore bypasses the NTFS volume stack:

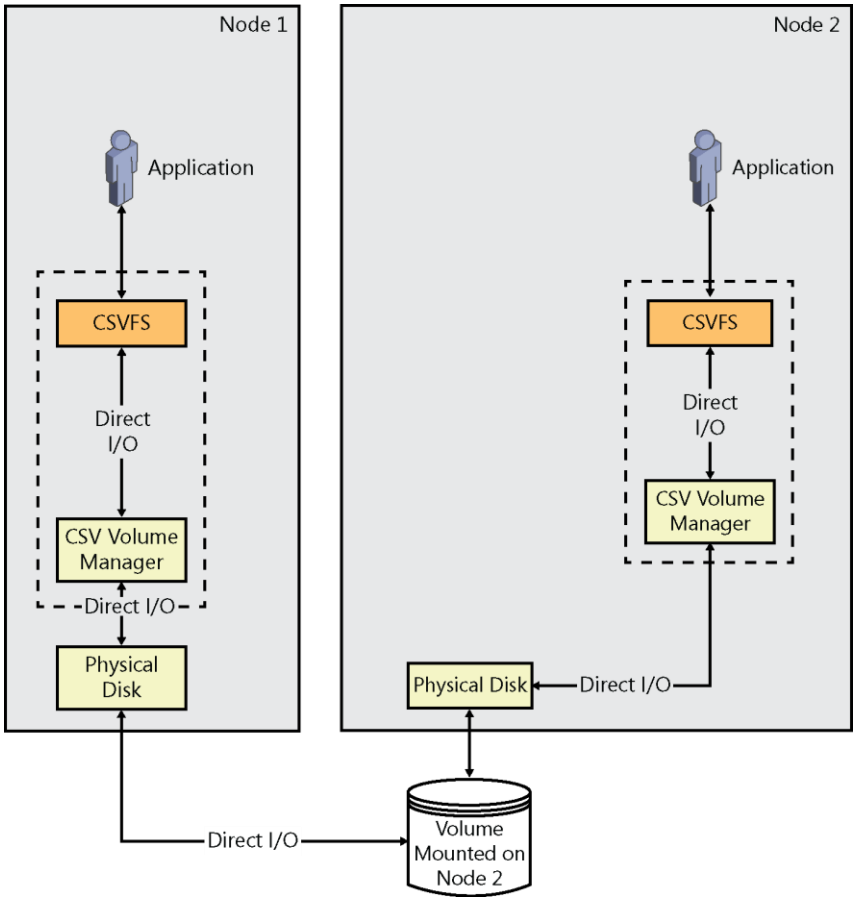

File System level redirection

During File System redirection, I/O on a cluster node is redirected at the top of the CSV Windows pseudo-file system stack over SMB to the disk. This traffic is written to the disk via the NTFS file system stack on the coordinator node. This mode of redirection occurs when the CSV volume is manually put into redirected mode, when BitLocker drive encryption is initiated or when an unsafe file system or volume filter is operating on CSV:

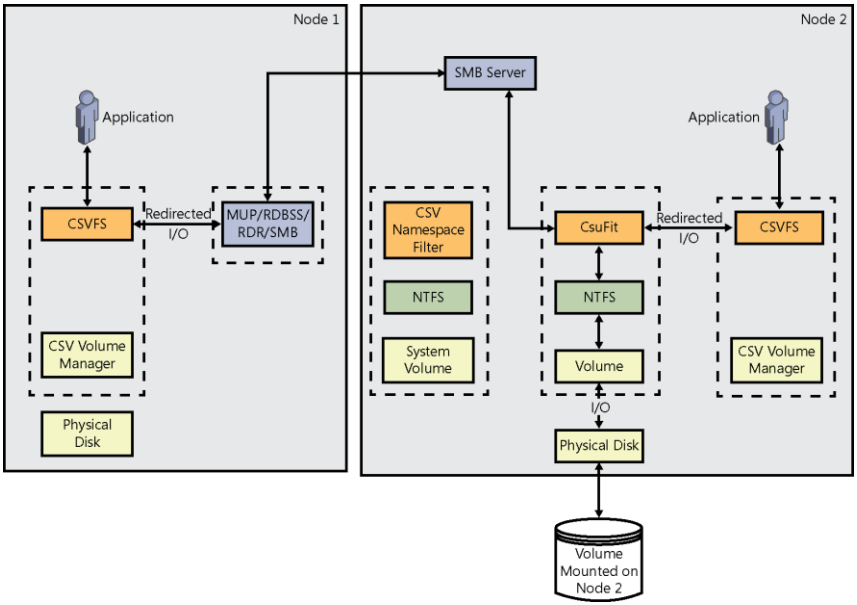

The Cluster CSV File System performance counters enable a deeper exploration of File System redirected traffic on a cluster node. This information will be helpful in planning the network infrastructure between the nodes in the cluster for handling CSV traffic and remedying (if required) any issues resulting in the I/O redirection. Key performance counters to monitor for the Cluster CSV File System performance object include:

- Redirected Read Bytes
- Redirected Read Bytes/sec
- Redirected Reads
- Redirected Reads Avg. Queue Length
- Redirected Reads/sec
- Redirected Write Bytes
- Redirected Write Bytes/sec
- Redirected Writes
- Redirected Writes Avg. Queue Length
- Redirected Writes/sec

Block level redirection

This mode of redirection was introduced in Windows Server 2012 and results in a significant performance improvement over File System redirected mode. In Block level redirected mode, I/O passes through the local proxy file system stack and is written directly to Disk.sys on the coordinator node. As a result it avoids traversing the file system stack twice. Block level redirection occurs in the event of a storage connectivity failure, when CSV is operated in an asymmetric configuration or in Storage Space configurations such as Mirrored spaces:

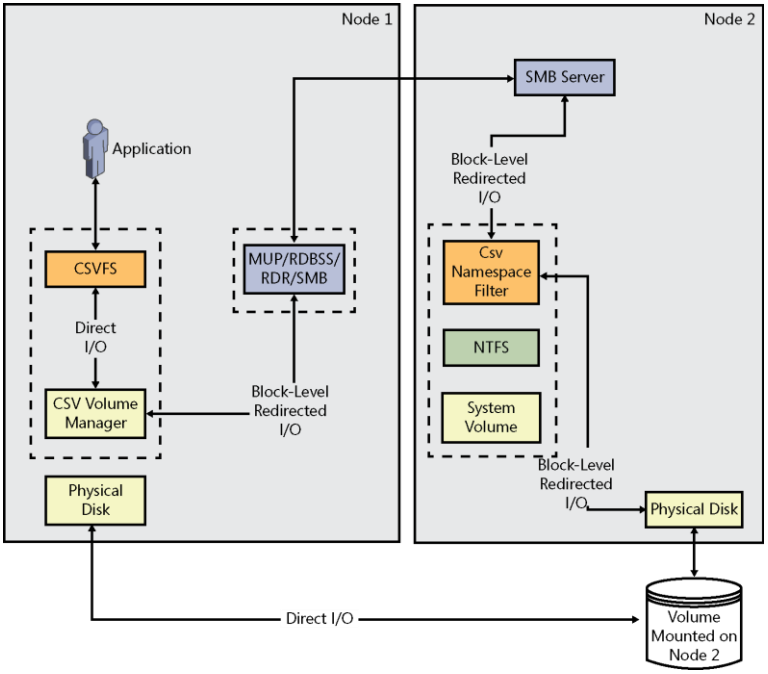

The Cluster CSV Volume Manager performance counters facilitate an examination of Block level redirected traffic on a cluster node. Key performance counters to monitor for the Cluster CSV Volume Manager performance object include:

- Direct IO Failure Redirection
- Direct IO failure Redirection/sec
- IO Read Bytes - Redirected
- IO Read Bytes/sec
- IO Read Bytes/sec - Redirected
- IO Read/sec - Redirected
- IO Reads
- IO Reads - Redirected
- IO Write Bytes - Redirected

- IO Write Bytes/sec
- IO Write Bytes/sec - Redirected
- IO Write/sec - Redirected
- IO Writes
- IO Writes - Redirected

Cluster Shared Volume Cache performance tuning

Windows Server 2012 introduces a new feature, Cluster Shared Volume (CSV) Cache, for caching CSV unbuffered I/O at the block level. CSV Cache allows system memory (RAM) to be allocated as a read-only, write-through cache.

CSV Cache substantially improves performance for applications such as Hyper-V, which conducts unbuffered I/O when accessing a VHD file. Given that CSV Cache caches at the block level, it is able to cache pieces of data being accessed within a VHD file. Additionally, CSV Cache delivers the most value in scenarios where virtual machines are used primarily for read requests, and are less write intensive. These include scenarios such as Pooled VDI virtual machines or also for reducing virtual machine boot storms. Given that the applicability of CSV Cache depends on the workload and the specific deployment, it is disabled by default. The customer feedback on CSV Cache has been overwhelmingly positive, and the general recommendation is to have it turned on for all applicable scenarios, including both Hyper-V Clusters using CSV and Scale-out File Servers using CSV.

Microsoft preliminary testing has found 512 MB to deliver excellent gain at minimal cost, and this is the recommend default value if enabled. Then based on the specific deployment and the I/O characteristics of the workloads in the virtual machines, the amount of memory allocated can be tuned. For a Scale-out File Server deployment, physical memory is typically not a contended resource; therefore it is recommended to allocate a significantly larger CSV cache.

The optimal CSV cache size can be tuned based on monitoring the Cluster CSV Volume Cache performance counters:

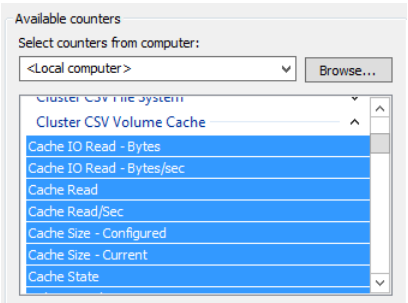

Also important for CSV cache tuning purposes are the following performance counters for the Cluster CSV File System performance object:

- IO Read Avg. Queue Length
- IO Read Bytes
- IO Read Bytes/sec
- IO Read Latency
- IO Read Queue Length
- IO Reads
- IO Reads/sec

— Subhasish Bhattacharya, Program Manager, Clustering and High Availability

Additional resources

Here are a few additional resources concerning this topic:

- Use Cluster Shared Volumes in a Windows Server 2012 Failover Cluster (TechNet Library) at: http://technet.microsoft.com/en-us/library/jj612868.aspx
- Using Cluster Shared Volumes in a Failover Cluster in Windows Server 2008 R2 (TechNet Library) at: http://technet.microsoft.com/en-us/library/ff182346(WS.10).aspx
- How to Enable CSV Cache at: http://blogs.msdn.com/b/clustering/archive/2012/03/22/10286676.aspx

Live Migration

Live Migration enables running virtual machines to migrate between nodes of clustered Hyper-V hosts with no perceived downtime. Live Migration has been enhanced in numerous ways in Windows Server 2012, and an entire book could be written on how to design, implement, optimize, and troubleshoot this feature in all its varied scenarios.

One topic that deserves close treatment is the subject of *Constrained Delegation*, which enables administrators to specify the services from which a computer that is trusted for delegation can request resources over a network. Certain scenarios using Live Migration require Constrained Delegation to be configured, and in this section Manjnath Ajjampur explains when and how you should do this to ensure optimal performance of Live Migration.

Why Constrained Delegation?

Not setting up Constrained Delegation properly could have impact on the ability to do the following:

- Remote Management of a Hyper-V role on a server other than the one you are directly connected to
- Live Migration of a virtual machine when there is no cluster in place—a Shared Nothing Live Migration

The user might get errors indicating authentication failures and messages like "general access denied" and "no credentials are available in the security package."

Constrained Delegation is not required if any of the following scenarios are true:

- RunAs accounts have been configured within Virtual Machine Manager, a component of System Center 2012, the Microsoft software product used for enterprise class management and provisioning of a private cloud.
- Remote Desktop Session is established to the source machine of the Live Migration.
- PowerShell remoting is used together with CredSSP.

NOTE **Constrained Delegation is more secure than CredSSP.**

This discussion will focus on the need for properly setting up Constrained Delegation in Live Migration scenarios.

With Hyper-V in Windows Server 2012, Live Migration has been greatly enhanced. One of the enhancements is the ability to Live Migrate a virtual machine without the need for a cluster. If the virtual machine is stored on SMB 3.0 shares, it can be live migrated from one host to another host, both of which are connected to the SMB 3.0 share.

Some background info

The concept of delegation has been around since Windows 2000. Upon login to an Active Directory domain, the user is issued a security token that in turn is used to access a resource. However, this token is good only for the connection between the user's machine and the first resource the user is attempting a connected to.

If the machine at the first hop wants to connect to resources on yet another machine, via a second hop, there is no mechanism to have the machine at the first hop authenticate to the second machine on the user's behalf. This is where the concept of delegation comes into play.

Delegation allows a service on the first machine to be delegated authentication on the second machine on the user's behalf. Constrained Delegation takes this a step further. Starting with Windows Server 2003, domain administrators can configure service accounts to delegate only to specific sets of service accounts.

The Hyper-V connection

In the context of Shared Nothing Live Migration, each Hyper-V host needs to be trusted for delegation to specific services on other Hyper-V hosts. These services are:

- **CIFS** This allows the SMB 3.0 protocol in Windows Server 2102 to set up, access, and create file shares on the target hosts.
- **Microsoft Virtual System Migration Service** This is the service for the Live Migration of virtual machines.

With this in place, all the Live Migration targets have the right authentication privileges to each other and to the SMB shares to perform the Live Migration.

Setting up Constrained Delegation

Constrained Delegation for Live Migration can be set up using Hyper-V Manager, as shown in the following screenshot:

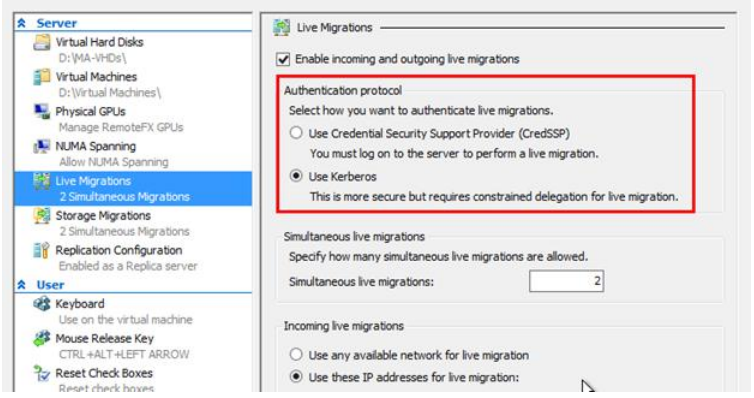

Constrained Delegation also needs to be set up in Active Directory. This can be done either via script or via the use of Active Directory Users and Computers MMC Snap-in. For more information, see the following blog posts:

- http://blogs.msdn.com/b/taylorb/archive/2012/03/20/enabling-hyper-v-remote-management-configuring-constrained-delegation-for-non-clustered-live-migration.aspx
- http://msdnrss.thecoderblogs.com/2012/03/hyper-v-remote-management-with-powershell-2/
- http://blogs.technet.com/b/matthts/archive/2012/06/10/configuring-kerberos-constrained-delegation-for-hyper-v-management.aspx

Prior to a Windows Server 2012 Active Directory infrastructure, Domain Administrator privileges were needed to setup Constrained Delegation. With Windows Server 2012, Resource Based Kerberos Constrained Delegation was introduced. This works across trusts and domains, and it is really easy to configure via PowerShell.

—Manjnath Ajjampur, Principal Datacenter Technologist

Additional resources

Here are a few additional resources concerning this topic:

- Virtual Machine Live Migration Overview (TechNet Library) at: http://technet.microsoft.com/en-us/library/hh831435.aspx
- Configure and Use Live Migration on Non-clustered Virtual Machines (TechNet Library) at: http://technet.microsoft.com/en-us/library/jj134199.aspx

Virtual Fibre Channel

Windows Server 2012 now provides Fibre Channel (FC) ports you can use within the guest operating system running on Hyper-V hosts. This allows you to connect your virtual machines directly to a Fibre Channel SAN. In this section Carlos Mayol Berral explains how you can verify whether your environment can use this feature, some possible errors that can arise from configuration problems, and how to implement redundancy using guest MPIO and Live Migration.

Fibre Channel on the guest

One interesting feature in Hyper-V in Windows Server 2012 is the capability to have Fibre Channel on the guest, which means to have FC LUNs directly attached to your virtual machines.This functionality can provide direct storage path communication, which can be very useful in some scenarios such as application backups using hardware providers, for maximum storage performance of guest virtual machines, or for guest clustering of up to 64 nodes running Windows Server 2012 virtual machines.

Before you implement Fibre Channel on the guest, you need to ensure that your infrastructure meets the necessary prerequisites.

Prerequisites

Begin by reviewing your host bus adapter (HBA) or converged network card to ensure Windows Server compatibility by:

- Checking if it has the Windows Server 2012 Certified Logo
- Checking if you can create a Virtual SAN Switch, that is, whether the physical HBA is compatible with Hyper-V. You can do this with Hyper-V UI or you can do it by running the following Windows PowerShell command:

```
gwmi -n root\virtualization\v2 Msvm_ExternalFcPort | ft -auto Name,
IsHyperVCapable
```

The output of this command might look something like this:

```
Name                                                                IsHyperVCapable
----                                                                ---------------
PCI\VEN_1657&DEV_0013&SUBSYS_00141657&REV_01\4&bc131d5&0&0138_0                True
PCI\VEN_1657&DEV_0013&SUBSYS_00141657&REV_01\4&bc131d5&0&0038_0                True
BDRV\FCOE&PCI_166214E4&SUBSYS_121314E4&REV_01\5&3353d3c6&0&50050500_0         False
EBDRV\FCOE&PCI_166214E4&SUBSYS_121314E4&REV_01\5&25692199&0&50050500_0        False
```

Next, review your HBA or converged network card to ensure you hardware is N_Port ID Virtualization (NPIV) capable. You can do this by:

- Running the Cluster Validation Wizard to check if your HBA has this feature and if it is enabled
- Enabling NPIV using your vendor's driver software

Next you should enable NPIV for your SAN fabric. This means that all of your SAN fabric should be NPIV capable/enabled.

Now create a Virtual FC SAN using the Hyper-V Manager console and select your HBA. Then when you create a FC Virtual Switch, you can add an FC virtual adapter to your virtual machines.

Remember also to configure your storage to present the logical units (LUNs) to the new World Wide Name (WWN) on your SAN Zone/Mask configuration.

Finally, it's important to note that virtual machines that can use this feature must run Windows Server 2008, Windows Server 2008 R2, or Windows Server 2012 as the guest operating system.

NOTE **The vast majority of issues involving virtual Fibre Channel can been resolved by installing the latest HBA drivers from the vendor.**

Virtual machine not starting

Once you have a virtual machine with a Virtual FC adapter installed, try starting the virtual machine. If the virtual machine won't start (i.e., it stops at 10 percent) after adding the Virtual FC adapter, first review your hardware driver and firmware versions and update them with the most recent versions available. This is the most common problem.

Possible errors for misconfigured NPIV or HBAs that aren't NPIV-capable include:

- 15080 Hyper-V-VMMS Error, 'VM' failed to add resources (Virtual machine ID)
- 12004 Microsoft-Windows-Hyper-V-Worker Synthetic FibreChannel Port: Failed to start reserving resources with Error 'Invalid class'
- 32100 Microsoft-Windows-Hyper-V-VMMS N/A NT AUTHORITY\SYSTEM 'Test': NPIV virtual port operation on virtual port () failed with an unknown error
- 32170 Microsoft-Windows-Hyper-V-VMMS N/A NT AUTHORITY\SYSTEM HBA port with instance name ('PCI\VEN_') is not NPIV capable and will not be used for virtual Fibre Channel. A newer version of the HBA driver is needed that supports the new NPIV Methods
- 21502 Microsoft-Windows-Hyper-V-High-Availability for Synthetic FibreChannel Port: Failed to start reserving resources with Error 'Insufficient system resources exist to complete the requested service, virtual port failed with an error: No physical port available to satisfy the request

Additional configuration steps

After your virtual machine is configured with a Virtual FC adapter, you might want to think about adding redundancy and configuring Live Migration. If Multipath I/O (MPIO) has been configured for your SAN, you should consider adding more than one Virtual FC adapter to have a redundant path for your virtual machine. After adding a second adapter you must install MPIO.

NOTE **Host and guest MPIO can coexist.**

For Live Migration purposes you should have two Virtual FC SANs, then connect one host HBA to every Virtual FC SAN and add two Virtual FC adapters to your virtual machine connected to every Virtual FC SAN. During the Live Migration process, Hyper-V automatically alternates between the Set A and Set B WWN addresses:

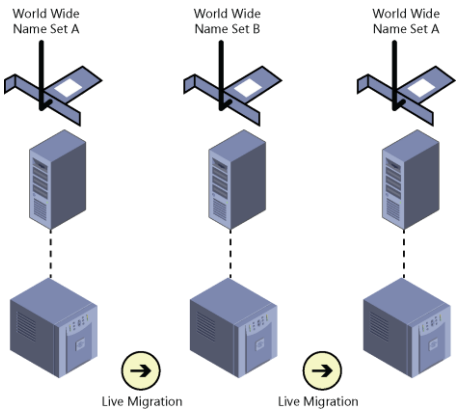

The correct procedure for implementing Guest MPIO is:

1. Have two or more HBA connections to the storage in the host.
2. Create a Virtual SAN in Hyper-V associated with each HBA connection.
3. Configure Virtual FC adapters in each virtual machine connected to each defined Virtual SAN.
4. Configure MPIO in the guest.

—Carlos Mayol Berral, Premier Field Engineer

Additional resources

Here is an additional resource concerning this topic:

- Hyper-V Virtual Fibre Channel Overview (TechNet Library) at: http://technet.microsoft.com/en-us/library/hh831413.aspx

Event logs

The Windows event logs are one of the first places you should look for clues when your systems or applications are not behaving in the way you expect them to behave. In the old days of Windows Server 2003 there were only a few event logs you had to check, with the System, Application, and Security logs being the main ones. Then beginning with Windows Server 2008 the number of different event logs and types of logs jumped into the hundreds, and although you can filter and search these logs using Event Viewer and Windows PowerShell, many long-time server admins still often feel overwhelmed by the sheer amount of information in Windows Server logs and wonder how to discover and identify what might be relevant for the particular scenario they are trying to optimize or troubleshoot. To help get you oriented, Thomas Roettinger explains some basics and also provides two examples below.

Hyper-V storage event logs

Windows Server provides several different event log categories where you can look for Hyper-V related issues. You can open the event log viewer either through the modern UI by using the key combination Windows logo key+X, or by launching the MMC snap-in directly by typing **eventvwr.msc**.

If you expand Application and Services Logs, Microsoft, and Windows, you will find specific application and services logs for Windows components. For example, you will see the Hyper-V logs if the Hyper-V Role is installed.

The storage-related logs are called:

- Hyper-V-SynthFC (Virtual FC Adapter)
- Hyper-V-SynthStor (SCSI Controller)

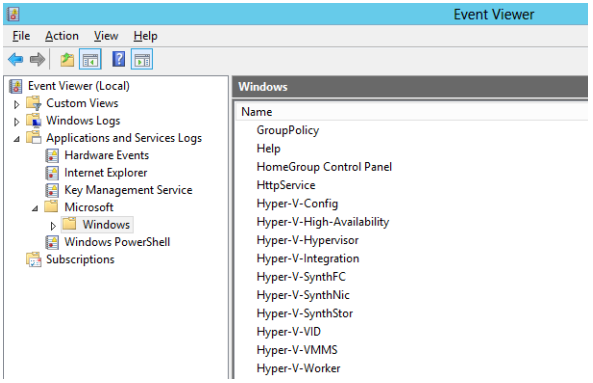

When you start a virtual machine, a virtual machine worker process (VMWP.exe) is launched for each virtual machine. The Hyper-V-Worker log provides all necessary information related to start, stop action, and the run time of a virtual machine. The worker process connects to the virtual machine management services (VMMS). In the following example you will notice that you can look at several Hyper-V logs for troubleshooting and that sometimes it is required to find the root cause of an outage.

Example: Missing virtual hard disk

Patricia the administrator received an error message when she tried, to start one of her virtual machines, but she did not pay attention to it:

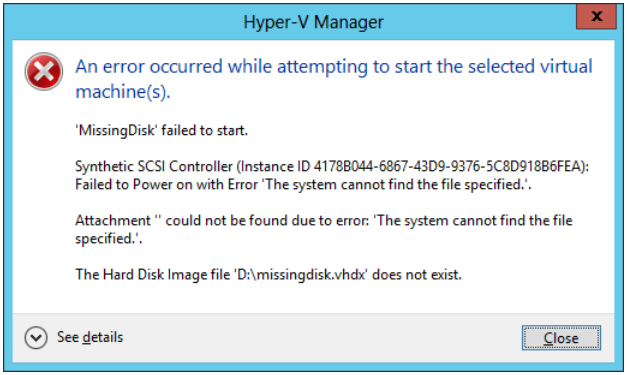

To investigate the problem, Patricia opens Event Viewer and checks the related Hyper-V event logs. She looks at the Hyper-V-VMWP logs first because the issue happened when she tried to power on the virtual machine. She finds the following event entry:

```
Event ID 32902
"Missing Disk" Synthetic SCSI Controller: Failed to Power on with Error 'The system cannot find the file specified.'
```

Patricia next looks at the Hyper-V-SynthStor logs to see if she can find more information. She finds the following:

```
Event ID 12240
'Missing Disk': Attachment" could not be found due to error 'The system cannot find the file specified.'
```

Patricia now needs to identify which file is missing, so she checks the Hyper-V-VMMS logs and finds the following entry:

```
Event ID 32902
The absolute path 'd:\missingdisk.vhdx' is valid for the "Hard Disk Image pool, but
references a file that does not exist.
```

Patricia knows that a virtual hard disk file is missing on drive D. She uses Windows Explorer to find the file. She then discovers that drive D is missing entirely:

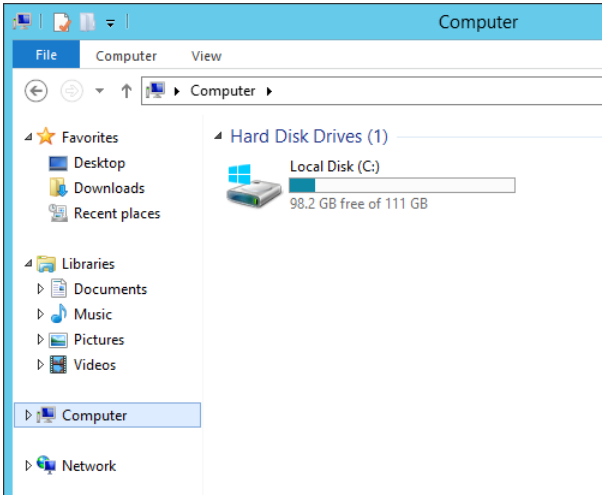

This error could happen if you lose storage connectivity, no matter how it is presented to your host. There are techniques to provide multiple paths to your storage when using SMB, FC, or ISCSI using MPIO or SMB multichannel to provide availability.

In this section you have seen the available Hyper-V event logs for storage. But you also read about their relationship to other important Hyper-V logs such as the Virtual Machine Worker Process and the Virtual Machine Management Service.

Example: Unsupported Fibre Channel adapter

Patricia needs to configure a virtual machine with a virtual Fibre Channel adapter. The Hyper-V host has a Fibre Channel adapter installed, so Patricia opens Virtual SAN Manager to create a new Fibre Channel SAN. She notices that she is not able to select the physical HBA and that the status is "The device or driver does not support virtual Fibre Channel."

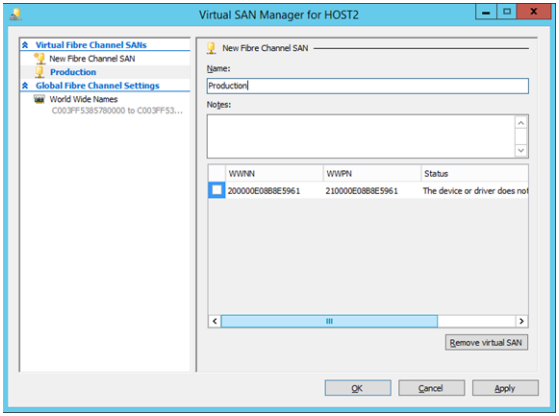

She ignores the message and continues to add a virtual HBA to the virtual machine. When she tries to start the virtual machine, she receives the following error message:

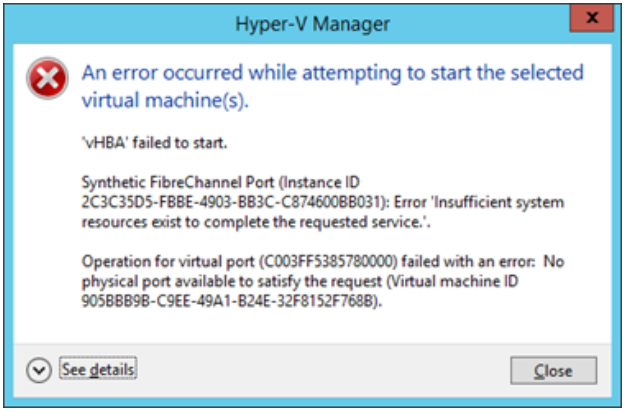

Patricia looks at the Hyper-V-SynthFC event log and discovers the following entry:

```
Event ID: 32161
'vHBA': Operation for virtual port ……failed with an error: No physical port available to
satisfy the request
```

Patricia opens Windows PowerShell to validate the driver. She does a WMI query to see if the driver provides support for Hyper-V. She runs the following command:

```
gwmi –Namespace "root\virtualization\v2" –Class msvm_externalfcport |select *hyper*
```

The resulting output from this command looks like this:

```
IsHyperVCapable
---------------
          False
```

While this output indicates that the driver does not support Hyper-V, Patricia knows that the physical adapter is capable of NPIV and should support Hyper-V. This disparity can happen because some older driver versions do not populate the necessary information correctly.

After Patricia downloads and installs the latest driver from the HBA vendor, everything works as expected.

—Thomas Roettinger, Program Manager, Partner and Customer Ecosystem Team

Additional resources

Here are a few additional resources concerning this topic:

- Event Viewer (TechNet Library) at: http://technet.microsoft.com/en-us/library/cc766042.aspx
- Windows PowerShell Management Cmdlets (TechNet Library) at: http://technet.microsoft.com/en-us/library/hh849827.aspx

SMB storage

A new capability for Hyper-V in Windows Server 2012 is the ability to store virtual machine files on a file share on a network file server. This capability is only possible because of enhancements in version 3.0 of the Server Message Block (SMB) file-sharing protocol introduced in Windows Server 2012. All types of virtual machine files can be stored on SMB 3.0 shares, including:

- Virtual machine configuration files
- Virtual hard disk files
- Virtual machine snapshots

Advantages of this new approach include the ability to leverage your existing converged networking infrastructure and reduced CapEx and OpEx costs due to no longer needing specialized storage hardware and expertise.

In this section Thomas Roettinger demonstrates how to troubleshoot an issue relating to SMB storage of virtual machines using the Windows Internals tool Process Monitor.

SMB share permissions

Windows Server 2012 now allows an administrator to store virtual machines on an SMB file share. The server that presents the file share must be capable of supporting SMB 3.0. When creating a share on a Windows Server 2012 you have three options:

1. SMB Share – Quick
2. SMB Share – Advanced
3. SMB Share - Applications

The main difference between a quick share and a share for applications is a feature called *Continues Availability*. This feature is required for applications like Hyper-V and SQL Server. Continues Availability requires that each write goes directly to the storage subsystem without any interference from the Windows Cache Manager in order to prevent data loss.

The Hyper-V Virtual Machine Management Service (VMMS) runs under the local system. This requires that the Hyper-V host computer account has the right permission to a remote SMB 3.0 share.

Remember the effective permissions are Share and NTFS permissions, and the Hyper-V host computer account needs read and write to the share.

In this section, I show you how to use Process Monitor for troubleshooting an SMB permission issue.

Example: Wrong share permissions

Patricia is creating a virtual machine that is stored on an SMB file share. She creates an SMB file share called "Share" and configures the share permissions with full control for everyone:

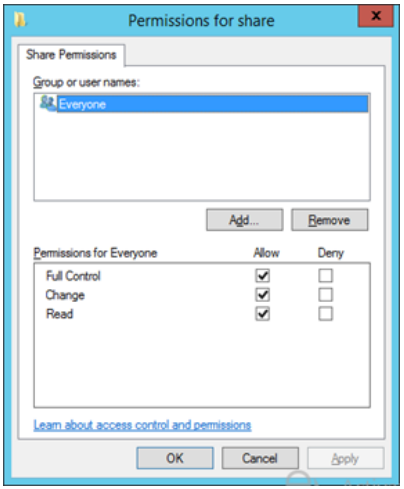

She also verifies that her Administrator account has full NTFS permission:

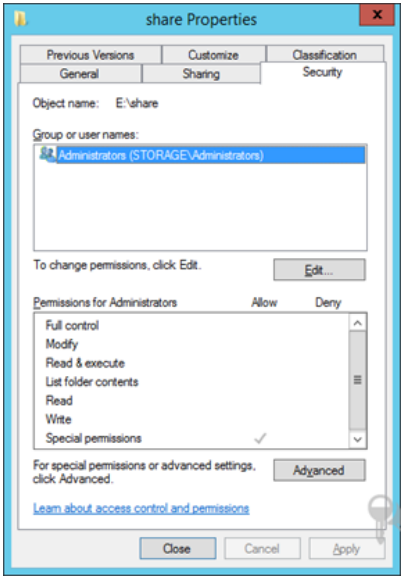

Patricia verifies on the Hyper-V host that she can access the share and has write permission with her Administrator account. She creates a virtual machine called "SMB" and stores it on the share. This works without any issue. But when she tries to power on the virtual machine she receives the following error message:

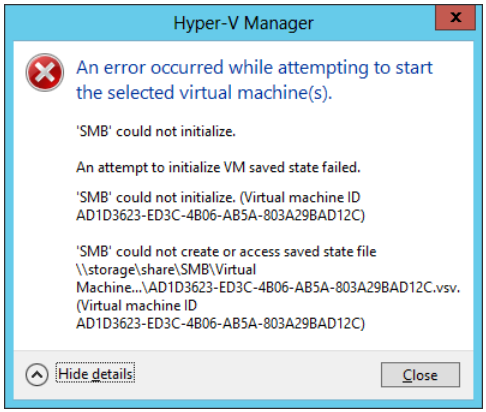

It is unclear to Patricia why this failed, and she starts investigating the issue by using Process Monitor. She starts Process Monitor, configures a filter to show only messages that include \\storage\share. After reproducing the error, she gets one Access Denied message in the Process Monitor log:

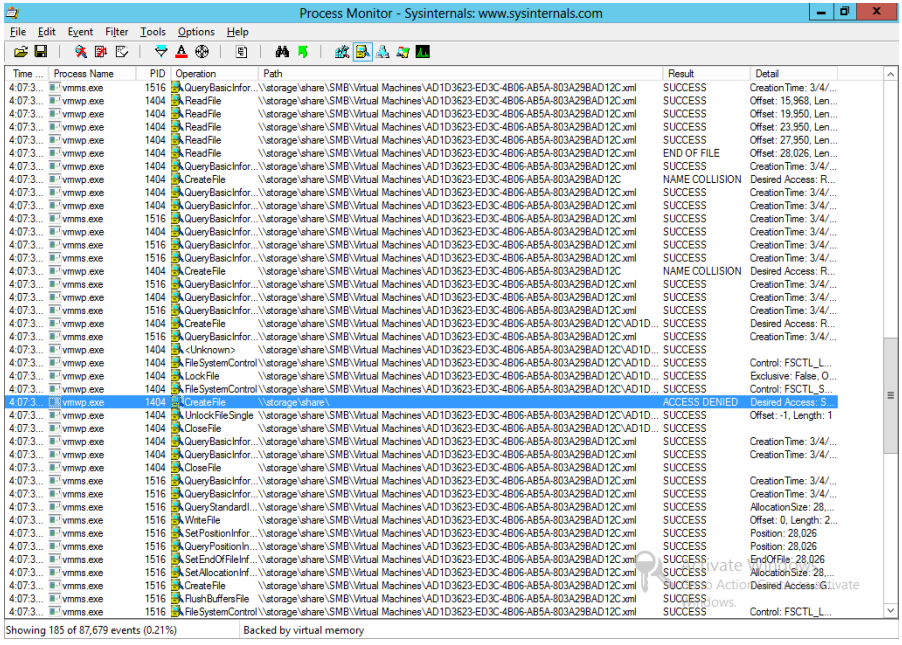

Patricia opens the Access Denied message to see which process is doing the CreateFile operation and getting access denied.

She recognizes that the virtual machine worker process is getting access denied and that the vmwp.exe process also has a parent process. This is the Virtual Machine Management Service (VMMS), which runs under the local system:

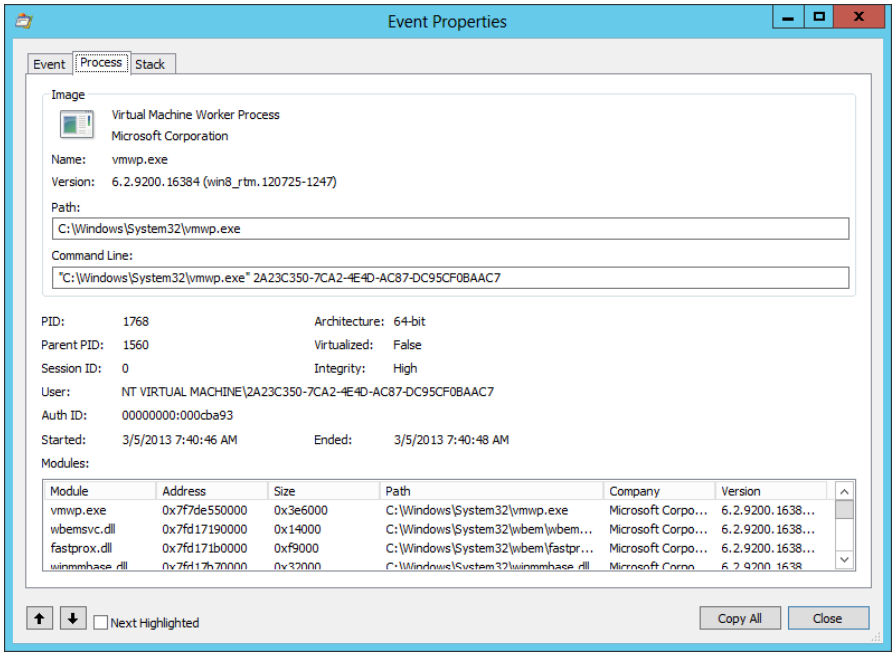

Patricia now adds the Hyper-V host computer account to the NTFS permissions and configures Full Control.

Finally the virtual machine starts and runs as expected. To ensure that she doesn't miss this important step in the future, she creates the following Windows PowerShell script to create the share and assign the proper permissions.

```
# Assign NTFS Permission to folder x:\VMS
ICACLS.EXE X: \VMS -- % /Grant Dom\ HVAdmin:(CI)(OI)F
ICACLS.EXE X: \VMS -- % /Grant Dom\ HV1$:(CI)(OI)F
ICACLS.EXE X: \VMS /Inheritance:R
#Create Share with the right permissions
New - SmbShare -Name VMS - Path X: \VMS - FullAccess Dom\HVAdmin, Dom\ HV1$
```

In this example, you have seen how you can troubleshoot a permission issue for an SMB share. But as I highlighted in the "Hyper-V storage event logs" section earlier, there could be other permission issues. For example, a virtual hard disk could be missing the permission for the NT Virtual Machine account.

—Thomas Roettinger, Program Manager, Partner and Customer Ecosystem Team

Additional resources

Here are a few additional resources concerning this topic:

- Server Message Block overview (TechNet Library) at: http://technet.microsoft.com/en-us/library/hh831795.aspx
- Hyper-V over SMB: Remote File Storage Support in Windows Server 2012 Hyper-V (TechNet Video) at: http://technet.microsoft.com/en-us/video/hyper-v-over-smb-remote-file-storage-support-in-windows-server-2012-hyper-v.aspx

Process Monitor (Windows Sysinternals on TechNet) at:
http://technet.microsoft.com/en-us/sysinternals/bb896645

SMB Multichannel

SMB Multichannel is one of several new features in version 3.0 of the Server Message Block (SMB) protocol introduced in Windows Server 2012. SMB Multichannel allows multiple connections to be used within a single SMB session in order to enhance network performance and ensure greater availability of file shares on Windows servers. In this section Thomas Roettinger discusses Receive-Side-Scaling (RSS), a feature of enterprise network adapters that distributes kernel-mode network processing across multiple processor cores to support higher network traffic loads than a single core can support. Thomas also provides an example of how to troubleshoot an issue when SMB Multichannel doesn't work as expected.

Troubleshooting SMB Multichannel

SMB Multichannel is a new feature of the SMB 3.0 protocol that was introduced in Windows Server 2012. This feature allows several network adapters to connect to a file server. By using SMB Multichannel you increase network performance, but even more importantly you increase the availability of your storage connection by using more than one network path. This feature is turned on by default and automatically discovers and makes use of your network connections.

Verifying Receive-Side-Scaling

SMB Multichannel requires multiple network adapters or multiple network teams. At least one network adapter must be capable of Receive-Side-Scaling (RSS).

To verify that your network adapter is RSS capable, you can run this Windows PowerShell command:

```
Get-SmbServerNetworkInterface
```

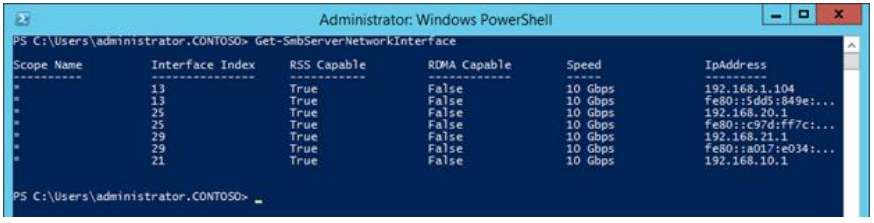

You should also verify that RSS is enabled for your network adapters using this command:

```
Get-NetAdapterRSS
```

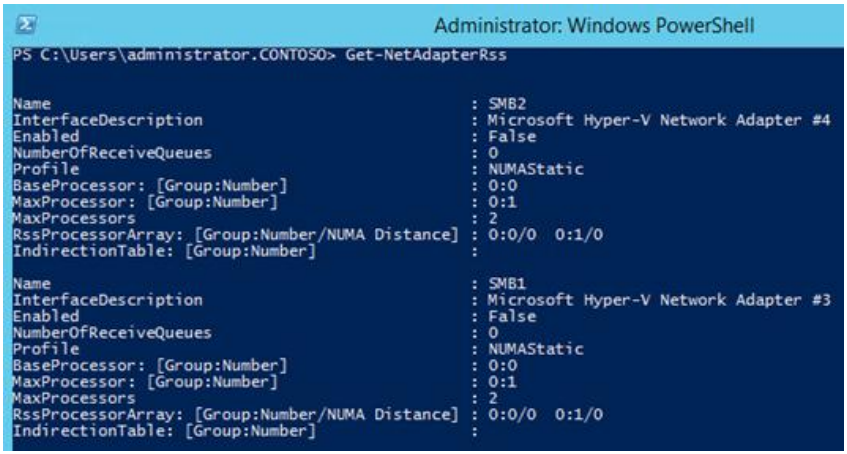

In this example, RSS is not enabled, so the driver settings need to be checked for RSS:

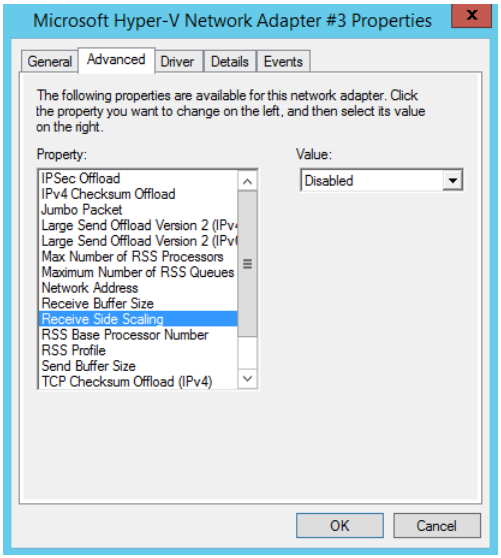

RSS is disabled at the driver level, so enable RSS and verify again with Windows PowerShell:

```
Get-NetAdapterRSS
```

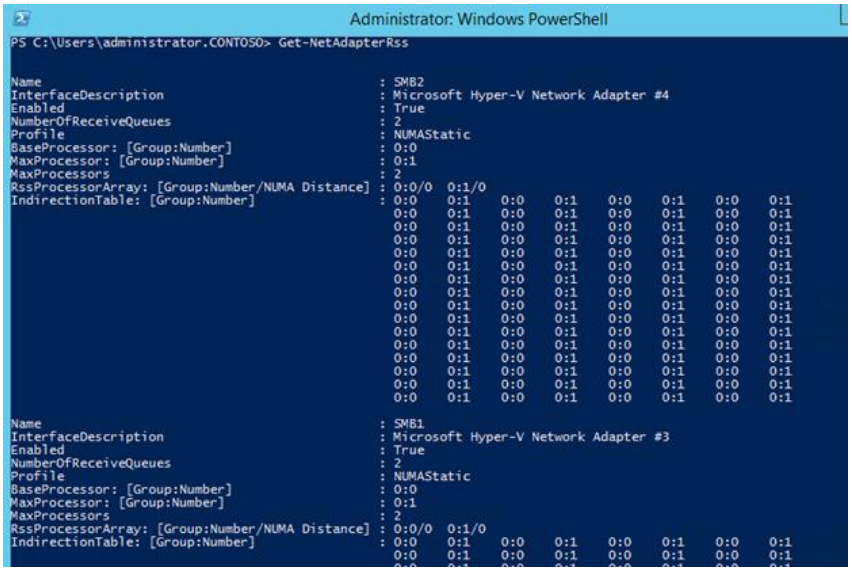

All modern server class network cards support RSS. Using multiple network interfaces to connect to a file server uses multiple TCP/IP connections. RSS ensures that these connections are distributed across multiple CPUs in your system. Otherwise your logical CPU 0 could become a bottleneck when the bandwidth increases and many small I/Os are performed.

It is very important to verify that RSS is enabled on both servers and clients where a Hyper-V host is the client.

Verifying SMB Multichannel

As mentioned earlier, SMB Multichannel is enabled by default, but you can verify it by using these two Windows PowerShell commands depending on whether you are checking a server or client:

```
Get-SmbServerConfiguration
```

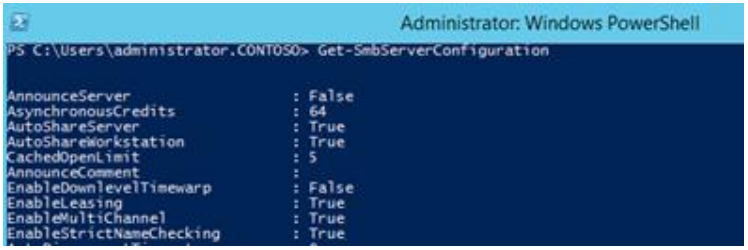

```
Get-SmbClientConfiguration
```

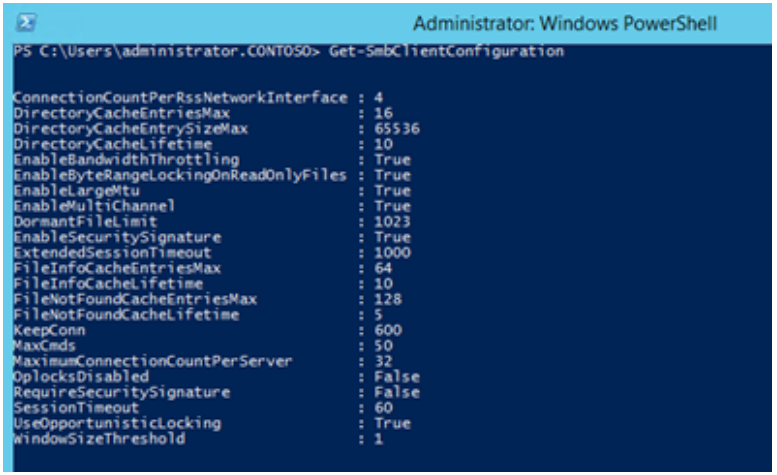

To verify active SMB Multichannel connections, start a long-running copy job from the Hyper-V host to the file server. You can then use this Windows PowerShell command to verify the multichannel connections:

```
Get-SmbMultichannelConnections
```

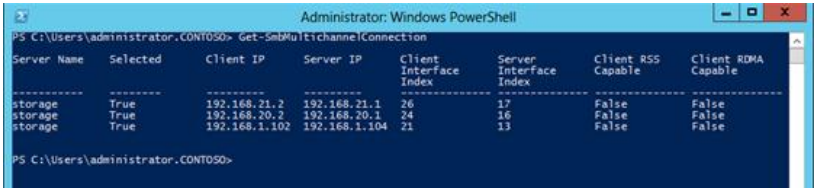

In this example, three active SMB connections are being used: two dedicated connections for SMB and one for the management adapter. The management adapter is used because it has the same link speed as the two dedicated SMB networks.

Excluding a network card

If you want to exclude a specific adapter from being used for SMB Multichannel, you can use the following Windows PowerShell command:

```
New-SmbMultichannelConstraint -ServerName FS1 -InterfaceIndex 21
```

Next, identify the physical network cards that are being used for the SMB connections by running this Windows PowerShell command on the client and server:

```
Get-NetAdapter
```

Here's the client:

And here's the server:

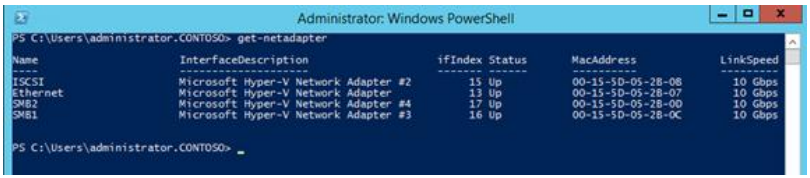

You can now create a table by matching the name from the client and server output to the output from the active SMB connections. This helps you verify the communication path that is being used:

IFINDEX SERVER	NAME		IFINDEX CLIENT	NAME
17	SMB2	< ----- >	26	vEthernet (SMB2)
16	SMB1	< ----- >	24	vEthernet (SMB1)
13	Ethernet	< ----- >	21	vEthernet (Management)

Example: Link down

In a previous section of this book, I covered the Hyper-V storage-related event logs. Now let's take a closer look to the SMB client-related event log.

You can open the event log viewer either through the modern UI by using the key combination Windows logo key+X, or by launching the MMC snap-in directly by typing **eventvwr.msc**.

Expand Application And Services Logs, Microsoft, and Windows to find specific application and services logs for Windows components. The SMB-related log is called SMBClient:

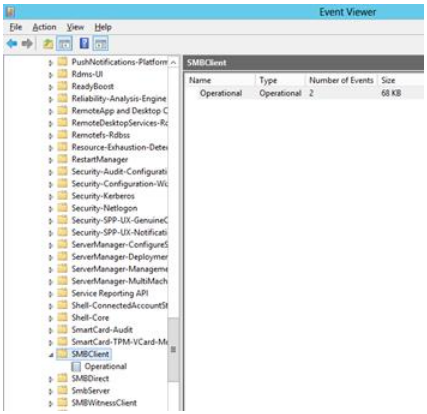

Patricia is an administrator who is notified that some virtual machines are not performing as usual. Patricia runs through the Windows PowerShell commands to verify that SMB Multichannel is working, and she discovers that one SMB connection is lost:

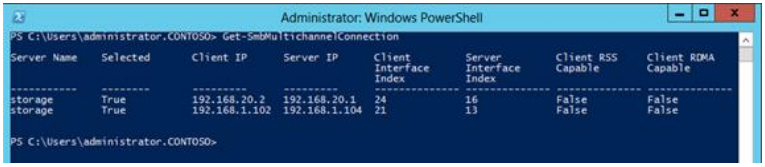

Patricia needs to understand why this happened, so she starts looking into the SMBClient event log on the Hyper-V host. She discovers the following event:

```
Event ID 30620
Connection to server \storage IP Address 192.168.21.1:445 was aborted
```

She checks the connection of the network interface and discovers that one link is down.

—Thomas Roettinger, Program Manager, Partner and Customer Ecosystem Team

Additional resources

Here is an additional resource concerning this topic:

- Network Performance and Availability (TechNet Library at): http://technet.microsoft.com/en-us/library/hh831499.aspx

Online backup

A key aspect of managing storage for your Hyper-V hosts is backing up the virtual machines that run on these hosts. Backups can be performed in various ways depending on the type of storage you're using and the tools you have available.

An important distinction is whether the backup of a virtual machine is performed online or offline. Online backups involve no downtime of the virtual machine, while offline backups involve some downtime.

> ***MORE INFO*** **For more information, see this post from the Virtualization Blog:** http://blogs.technet.com/b/virtualization/archive/2008/08/29/backing-up-hyper-v-virtual-machines.aspx

Of course, even backups can have problems sometimes, and in this section Thomas Roettinger describes how to troubleshoot an issue involving online backup. But first he explains the backup process and how the Volume Shadow Copy Service (VSS) is involved in performing Hyper-V backups.

Hyper-V backups and VSS

VSS can produce consistent shadow copies by coordinating with business applications, file system services, backup applications, fast-recovery solutions, and storage hardware.

VSS consists of several components, including the Hyper-V VSS Writer. The Hyper-V VSS Writer is automatically installed and registered when you install the Hyper-V role.

This is a high-level overview of how a backup of a virtual machine works with the Hyper-V VSS Writer.

1. Backup is initiated by a backup program.
3. The backup program sends a request to VSS.
4. VSS talks to Hyper-V VSS Writer to inform Hyper-V that a backup will be taking place.
5. Hyper-V acts as a proxy and forwards this VSS request to running virtual machines with the help of the integration components for VSS running inside the virtual machine. This requires a VSS-aware operating system.
6. The virtual machines then use the integration components to inform the Hyper-V VSS Writer that they are in a consistent state.
7. Hyper-V VSS Writer then informs VSS that it is in a consistent state.
8. VSS creates a VSS snapshot.

You may be asking yourself what happens if you have Linux running as a guest operating system since it does not contain VSS. Here is a high-level overview of the backup process of a virtual machine running Linux or any other operating system that does not include VSS work.

1. Backup is initiated by a backup program.
2. The backup program sends a request to VSS.
3. VSS talks to Hyper-V VSS Writer to inform Hyper-V that a backup will be taking place.
4. Hyper-V acts as a proxy and forwards the VSS request to the virtual machine running Linux but does not receive a response.
5. Hyper-V puts the virtual machine running Linux in a saved state.
6. Hyper-V VSS Writer then informs VSS that it is in a consistent state.
7. VSS creates a VSS snapshot.
8. Hyper-V resumes the virtual machine running Linux.

Example: Online backup issue

Patricia is an administrator who discovers that one of her virtual machines is no longer being backed up online at night. So, she modifies her backup job to also run during the day. She receives a call from the help desk that users are reporting a service interruption.

Patricia starts investigating why the outage happened during backup. She starts by checking the status of the Hyper-V VSS Writer using the following command:

```
Vssadmin.exe list writer
```

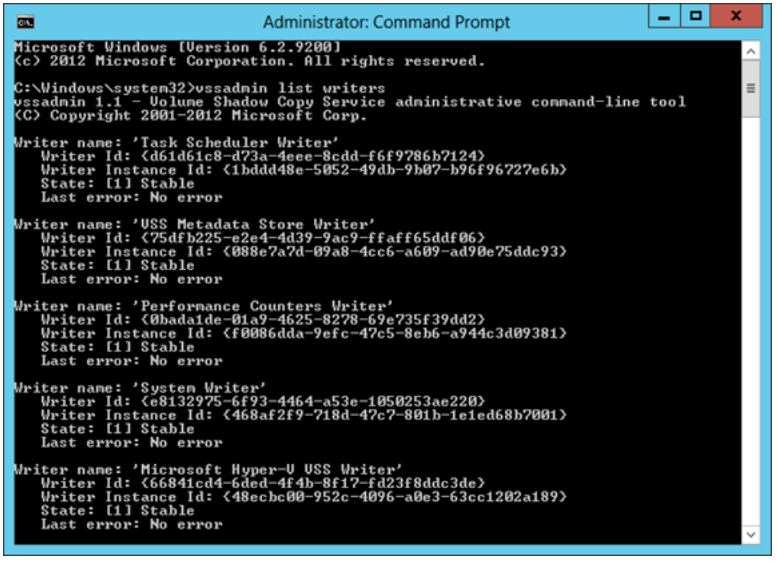

She notes that the Microsoft Hyper-V VSS Writer reports no error.

Next she checks the Application event log of the guest operating system to determine if there are any VSS-related issues. No entries appear, but Patricia knows that there should be at least one event showing that the VSS was started (Event ID 7035, 7036). This makes her think that the VSS request is not reaching the virtual machine.

Patricia opens the event viewer and checks the related Hyper-V event logs. She looks at the Hyper-V-Integration event log because she knows that one of the integration services is responsible for VSS:

```
Event ID 4096
'Online Backup': The Volume Shadow Copy integration service is not enabled.
```

Patricia opens the configuration settings of the respective virtual machine and detects that the Backup (volume snapshot) setting is not enabled:

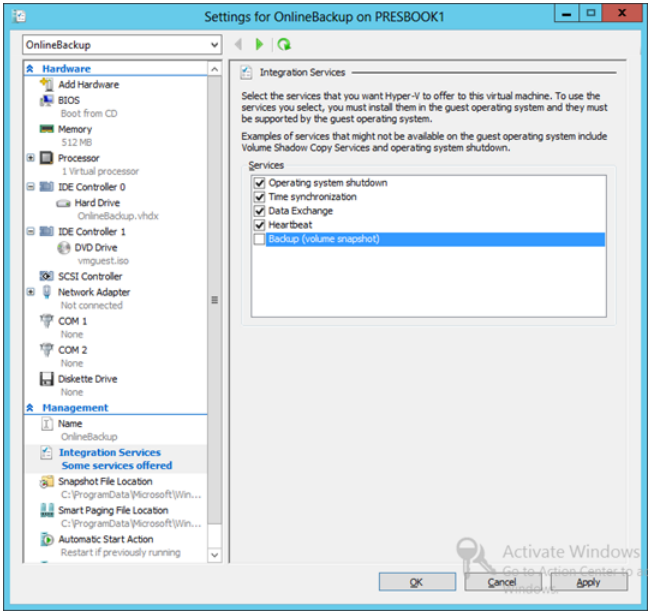

After Patricia turns on the backup feature, the virtual machine no longer goes into saved state while the backup job runs:

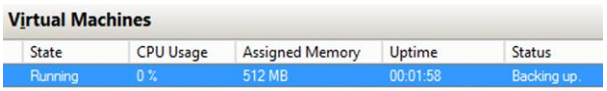

—Thomas Roettinger, Program Manager, Partner and Customer Ecosystem Team

Additional resources

Here is an additional resource concerning this topic:

- Planning for Backup (TechNet Library) at: http://technet.microsoft.com/en-us/library/dd252619(v=WS.10).aspx

Antivirus exclusions

If you plan to use an antivirus product on your Hyper-V host, it's very important that you exclude some important system files to prevent problems. Thomas Roettinger goes into some detail here concerning the importance of this.

Configuring antivirus exclusions

The Hyper-V root partition, also called the parent partition, should not be used to install any software other than agents required to operate a managed system. There are many reasons for this, including memory consumption of processes running in the root partition and also the performance behavior of software.

These are some possible error messages if the necessary exclusions are not configured for Hyper-V:

```
The requested operation cannot be performed on a file with a user-mapped section open. (0x800704C8)
VMName' Microsoft Synthetic Ethernet Port (Instance ID{7E0DA81A-A7B4-4DFD-869F-37002C36D816}): Failed to Power On with Error 'The specified network resource or device is no longer available.' (0x80070037).
The I/O operation has been aborted because of either a thread exit or an application request. (0x800703E3)
```

To make sure you don't run into any of the above errors, you should exclude the following files and folders in your antivirus product real-time scanning component:

- Default virtual machine configuration directory (C:\ProgramData\Microsoft\Windows\Hyper-V)
- Custom virtual machine configuration directories
- Default virtual hard disk drive directory (C:\Users\Public\Documents\Hyper-V\Virtual Hard Disks)
- Custom virtual hard disk drive directories
- Snapshot directories
- Vmms.exe (Note: May have to be configured as process exclusions within the antivirus software)
- Vmwp.exe (Note: May have to be configured as process exclusions within the antivirus software)
- C:\ClusterStorage and all subdirectories (Only when using Cluster Shared Volume)

—Thomas Roettinger, Program Manager, Partner and Customer Ecosystem Team

Additional resources

Here are a few additional resources concerning this topic:

- 6 Best Practices for Physical Servers Hosting Hyper-V Roles (TechNet Magazine) at: http://technet.microsoft.com/en-us/magazine/dd744830.aspx
- Virtual machines are missing, or error 0x800704C8, 0x80070037, or 0x800703E3 occurs when you try to start or create a virtual machine at: http://support.microsoft.com/kb/961804

Windows PowerShell tips

In this section Thomas Roettinger summarizes a number of different Windows PowerShell commands for performing storage-related tasks.

Storage-related tasks and Windows PowerShell

For storage-related tasks, some of the Windows PowerShell commands are not present in the Hyper-V Manager UI, for example changing the physical sector size. These commands can be useful when troubleshooting different kinds of storage-related issues involving Hyper-V hosts and virtual machines.

- Convert a VHD to a VHDX

```
Convert-VHD –Path c:\temp\temp.vhd –destinationpath d:\temp\test.vhdx
```

- Collect VHD/VHDX information, for example Physical Sector Size

```
Get-VHD –Path c:\temp\test.vhdx
```

- Change Physical Sector Size of a VHDX, as required for Windows 2003 for example

```
Set-VHD –Path c:\temp\test.vhdx –PhysicalSectorSizeBytes 512
```

- Expand or shrink a VHDX size (VHD expand only)

```
Expand-VHD –Path c:\temp\test.vhdx –SizeBytes
```

- Test a VHD/VHDX for errors

```
Test-VHD –Path c:\temp\temp.vhdx
```

- Online attach VHD/VHDX to an SCSI Controller

```
Add-VMHardDiskDrive –VMName Test –ControllerType SCSI –Path c:\temp\test.vhdx
```

- Attach a DVD drive to an IDE Controller

```
Add-VMDvdDrive –VMName Test –ControllerNumber 1
```

- Add a Virtual FC HBA to a virtual machine (A Virtual Fibre Channel SAN must exist)

```
Add-VMFibreChannelHba –VMName Test –SanName TestSAN
```

- List all snapshots of a virtual machine

```
Get-VMSnapshot –VMName Test
```

- Run defrag for a CSV disk (LUN). Note: Must be executed on coordinator node

```
Repair-ClusterSharedVolume C:\ClusterStorage\Volume1 -Defrag
```

- Run ChkDsk on a CSV disk (LUN) Note: Must be executed on coordinator node

```
Repair-ClusterSharedVolume C:\ClusterStorage\Volume1 -ChkDsk -Parameters "/F"
```

- Define Size of CSV Cache. Note: This amount will be reserved in host memory.

```
(Get-Cluster). SharedVolumeBlockCacheSizeInMB = 512
```

- Enable CSV Cache for a specific CSV disk (LUN)

```
Get-ClusterSharedVolume "Cluster Disk 1" | Set-ClusterParameter
CsvEnableBlockCache 1
```

—Thomas Roettinger, Program Manager, Partner and Customer Ecosystem Team

Additional resources

Here are a few additional resources concerning this topic:

- Hyper-V Cmdlets in Windows PowerShell (TechNet Library) at: http://technet.microsoft.com/en-us/library/hh848559.aspx
- Failover Clusters Cmdlets in Windows PowerShell (TechNet Library) at: http://technet.microsoft.com/en-us/library/hh847239.aspx

Best Practices Analyzer

Sometimes the best optimizing and troubleshooting tools are the ones right in front of you, staring you in the face. The Best Practices Analyzer (BPA) functionality integrated into Server Manager on Windows Server 2012 is a good example of this.

In this section, Mark Gehazi, a Data Center Specialist with Microsoft U.S. State and Local Government (SLG) team, explains how the BPA functionality for Hyper-V works and why it can be useful.

Troubleshooting with Hyper-V Best Practices Analyzer

In the Windows management and administration world, everyone's accustomed to best practices established by the product manufacturer and expert community to configure a server as defined by them. Many of these best practices are dependent on the environment and usage scenario, but they always can help with establishing a baseline and troubleshooting. For example, it is considered a best practice for Hyper-V virtual machines to use Synthetic Network Adaptor for performance reasons. While best practice violations, even critical ones, are not necessarily problematic, they indicate server configurations that can result in poor performance, poor reliability, unexpected conflicts, increased security risks, or other potential problems.

Windows Server Best Practices Analyzer (BPA) was first introduced in Windows Server 2008 for a few roles and features and was expanded over time through update packages. BPA was well received by the Windows admin community, helping them reduce accidental best practices violations and resulting in less downtime and support costs. Windows also had the admin community utilizing PowerShell integration with BPA and automating the monitoring and reporting of BPA and violations across multiple server farms and various roles.

Hyper-V BPA

In Windows Server 2012, with the introduction of the new Server Manager and its multi-server management capabilities, BPA becomes even more integrated and relevant to administrators who are setting up stand-alone or clustered Hyper-V servers.

Server and virtualization administrators now can utilize Hyper-V BPA while configuring the virtualization fabric or troubleshooting issues.

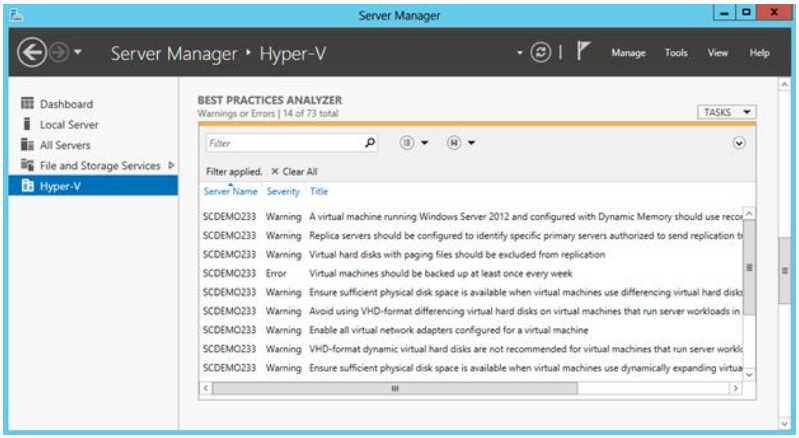

Windows Server 2012 Hyper-V BPA currently includes more than 70 rules and can be expanded via future updates if new best practices are identified. Following are a few examples of these rules:

- The Server Core installation option is recommended for servers running Hyper-V.
- VHDX-format virtual hard disks are recommended for virtual machines that have recovery history enabled in replication settings.
- Run the current version of integration services in all guest operating systems.
- Use RAM that provides error correction.
- Second-level address translation is required when running virtual machines enabled for RemoteFX.
- Use at least SMB protocol version 3.0 for file shares that store files for virtual machines.
- Use at least SMB protocol version 3.0 configured for continuous availability on file shares that store files for virtual machines.
- A virtual machine running Windows Server 2012 and configured with Dynamic Memory should use recommended values for memory settings.
- A Replica server must be configured to accept replication requests.
- To participate in replication, servers in failover clusters must have a Hyper-V Replica Broker configured.
- Configure a policy to throttle the replication traffic on the network.
- Avoid using legacy network adapters when the guest operating system supports network adapters.
- VMQ should be enabled on VMQ-capable physical network adapters bound to an external virtual switch.

Some of these rules pertain to the Hyper-V host configuration and some of them target individual virtual machines for configurations and settings. For instance, if few virtual machines are violating a best practices recommendation, BPA results will report the virtual machine names in the report, making it easier to identify those individual virtual machines from several others running in a Hyper-V farm.

For example, when one or more virtual machines haven't been backed up for over a week, the BPA will report this as an error and the virtual machine name(s) will be reported in the impact section:

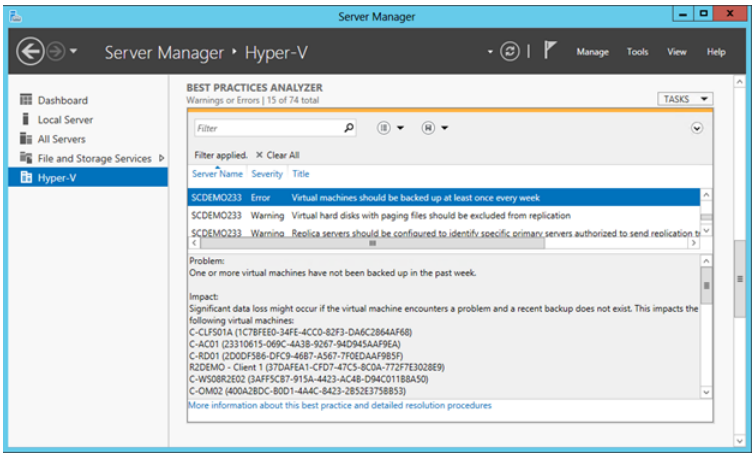

Also, clicking the More Information About This Best Practice And Detailed Resolution Procedure link will take you to the associated TechNet wiki library for that specific item, which is kept current by the Microsoft Hyper-V team.

BPA results are categorized in three severity models: Error, Information, and Warning:

- **Error** Returned when a role does not satisfy the conditions of a best practice rule and functionality problems can be expected.
- **Information** Returned when a role satisfies the conditions of a best practice rule.
- **Warning** Returned if the results of noncompliance can cause problems if changes are not made. The application might be compliant as operating currently, but may not satisfy the conditions of a rule if changes are not made to its configuration or policy settings. For example, a scan of Remote Desktop Services might show a warning result if a license server is unavailable to the role, because even if no remote connections are active at the time of the scan, not having the license server prevents new remote connections from obtaining valid client access licenses.

Server Manager BPA View by default reports only the Warning and Error results; however, clearing the view filter will show all the results, including reports with the Information severity tag.

Individual report items also can be excluded from the report if the administrator is already aware of the item and believes the rule isn't applicable to a specific scenario or environment. To exclude an item, right-click it and select Exclude Result.

Each reported item includes three major severity levels: Problem, Impact, and Resolution. These levels allow administrators to properly document the BPA results and prepare for remediation and a possible change control request process. Here is an example:

Problem:

One or more virtual machines are using differencing virtual hard disks.

Impact:

Differencing virtual hard disks require available space on the hosting volume so that space can be allocated when writes to the virtual hard disks occur. If available space is exhausted, any virtual machine that relies on the physical storage could be affected. This impacts the following virtual machines:

C-CLFS01A (1C7BFEE0-34FE-4CC0-82F3-DA6C2864AF68) - D:\DemoV4.VHD\CLFS01A\Virtual Hard Disks\D-C-CLFS01A.vhdx

C-AC01 (23310615-069C-4A3B-9267-94D945AAF9EA) - D:\DemoV4.VHD\AC01\Virtual Hard Disks\D-C-AC01.vhdx

C-AC01 (23310615-069C-4A3B-9267-94D945AAF9EA) - D:\DemoV4.VHD\AC01\Virtual Hard Disks\D-C-AC01_E1.vhdx

<…>

Resolution:

Monitor available disk space to ensure sufficient space is available for virtual hard disk expansion. Consider merging differencing virtual hard disks into their parent. In Hyper-V Manager, inspect the differencing disk to determine the parent virtual hard disk. If you merge a differencing disk to a parent disk that is shared by other differencing disks, that action will corrupt the relationship between the other differencing disks and the parent disk, making them unusable.

After verifying that the parent virtual hard disk is not shared, you can use the Edit Disk Wizard to merge the differencing disk to the parent virtual hard disk.

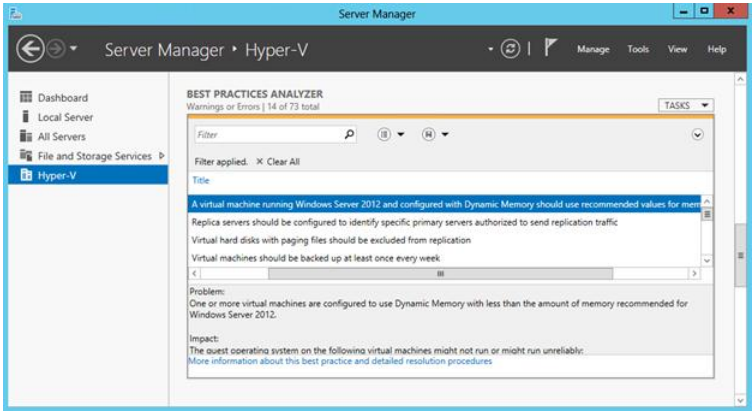

In a multi-server environment, when adding all the remote servers to Server Manager, administrators will be able to view all servers that have the Hyper-V role installed under the Hyper-V section. They can then click Start BPA Scan on the Tasks menu to choose Hyper-V servers (remote or local) for analysis:

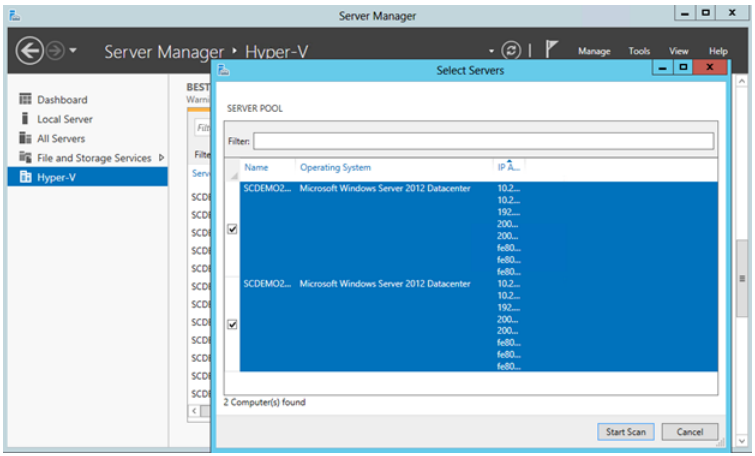

The result will include reports from all those servers, which can be sorted or grouped by Server Name, FQDN, Severity, Title, and Category. Keep in mind that the result will be populated for the Hyper-V servers that are selected in the top section of the page under Servers:

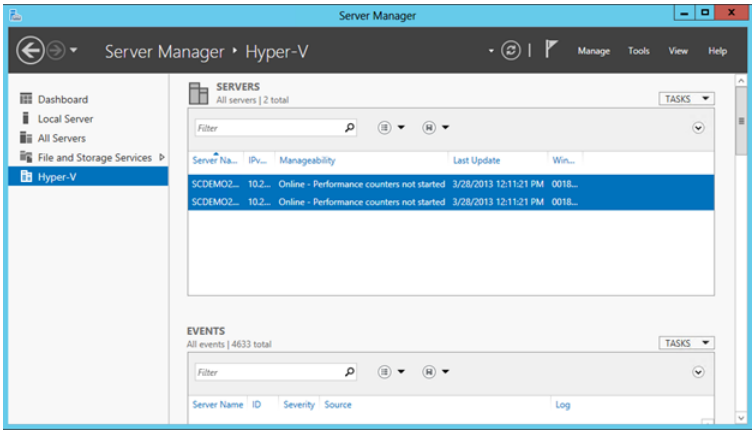

PowerShell and automation

Windows Server 2012 Best Practices PowerShell module contains four simple PowerShell cmdlets that empower administrators to invoke multiple BPA analyses across several servers and capture the results in different formats:

```
PS C:\> Get-Command -Module BestPractices

CommandType  Name               ModuleName
-----------  ----               ----------
Cmdlet       Get-BpaModel       BestPractices
Cmdlet       Get-BpaResult      BestPractices
Cmdlet       Invoke-BpaModel    BestPractices
Cmdlet       Set-BpaResult      BestPractices
```

Here is some more information about each cmdlet:

- **Get-BpaModel** Retrieves and displays the list of BPA models installed on the system.
- **Get-BpaResult** Retrieves and displays the results of the most recent BPA scan for a specific model.
- **Invoke-BpaModel** Starts a BPA scan for a specific model that is installed on a computer.
- **Set-BpaResult** Excludes or includes existing results of a BPA scan to display only the specified scan results.

NOTE **To make sure you have the latest PowerShell help files for Windows Server 2012, use the Update-Help cmdlet.**

The BPA model name is normally needed to work with that BPA. In the case of Hyper-V the BPA model is Microsoft/Windows/Hyper-V.

Thus, invoking the Hyper-V BPA model will be possible by running the following PowerShell command:

```
PS C:\> Invoke-BpaModel -ModelId Microsoft/Windows/Hyper-V
```

The following command will show the results:

```
PS C:\> Get-BpaResult -ModelId Microsoft/Windows/Hyper-V
```

Administrators have the option to use PowerShell and output the results to text, HTLM, CSV or other formats as they wish.

MORE INFO **Cristian Edwards has a blog post on his TechNet blog that demonstrates how someone can generate an HTML output of Hyper-V BPA Results from PowerShell at http://blogs.technet.com/b/cedward/archive/2011/01/11/hyper-v-bpa-html-report.aspx.**

If you'd like to have a little bit of fun, run the following PowerShell command after you've already executed Hyper-V BPA using the invoke-BpaModel cmdlet and see what happens:

```
PS C:\> Get-BpaResult -ModelId Microsoft/Windows/Hyper-V | Out-GridView
```

Cool, isn't it?

Hyper-V BPA also integrates with Microsoft System Center 2012 Operations Manager through Windows Server Operating System Management Pack and can enable an automated method of raising alerts and notification when any violations are detected. The monitoring pack now collects BPA results from monitored servers and returns the BPA state to Operations Manager. Customers do not want BPA data to be collected on all systems by default, so it needs to be enabled after importing the Management Pack.

You can download the System Center Management Pack for Windows Server Operating System here: http://www.microsoft.com/en-us/download/details.aspx?id=9296.

Failover clustering

Of course, most enterprise customers are taking advantage of highly available (HA) Hyper-V virtual machines powered by the Microsoft Windows Server 2012 Failover Clustering feature. All the new enhanced functionality and superb scalability numbers (support for 64 Hyper-V cluster nodes running up to 8,000 virtual machines along with Cluster Shared Volumes 2.0, Scale-Out File Servers for Hyper-V virtual machines over SMB, Cluster-Aware Updating and VM Application monitoring) of failover clustering are music to the ears of virtualization administrators and IT operations efficiency enthusiasts.

Windows Server 2012 Failover Cluster Validation test now supports Cluster Shared Volumes (CSVs), Hyper-V and virtual machines (when the Hyper-V role is installed), along with all other aspects of clustering including networking and storage and host configuration. Plus it runs significantly faster than its predecessor.

As part of troubleshooting a Hyper-V environment when running in an HA environment, administrators should run the Microsoft Failover Cluster validation test and review the report for any errors or warnings. This should help reduce time to resolution for configuration issues or storage/networking problems that occur as a result of changes in the infrastructure.

Similar to Hyper-V BPA, the Microsoft Failover Cluster validation test has a PowerShell cmdlet (Test-Cluster) and allows administrators to invoke and generate reports for one or more clusters. For more information and examples, please see the Test-Cluster page on Microsoft TechNet library at http://technet.microsoft.com/en-US/library/hh847274.

Summary

When troubleshooting Hyper-V issues, the best course of action is to start by using the most advanced tool available to administrators on Windows Server 2012: Server Manager. If all the Hyper-V servers in question (remote or local) have been added to the Server Manager Console, an administrator can review all the Hyper-V event channel logs, services and their status, BPA results, and real-time key performance metrics (CPU and memory) from that same pane. In an HA Hyper-V environment, adding a fresh report from a Microsoft Failover Cluster Validation test would be very useful as well. Hyper-V BPA can help administrators to validate their configuration before adding a Hyper-V server to production and on an ongoing basis. Integration of Hyper-V BPA with System Center Operations Manager along with automation using PowerShell adds advanced monitoring and alerting as soon as any deviation from the best practices is detected.

While these aren't the only troubleshooting tasks one should or might perform, in many cases, they are valid and time saving, and they can lead to higher uptime and better Service Level Agreements (SLAs).

—Mark Ghazai, Data Center Specialist with Microsoft US State and Local Government (SLG) team

Additional resources

Here are a few additional resources concerning this topic:

- Run Best Practices Analyzer Scans and Manage Scan Results at: http://technet.microsoft.com/en-us/library/hh831400.aspx
- Best Practices Analyzer Cmdlets in Windows PowerShell at: http://technet.microsoft.com/en-us/library/hh868084.aspx
- Best Practices Analyzer for Hyper-V: Configuration at: http://technet.microsoft.com/en-us/library/ee941122(v=WS.10).aspx
- Best Practices Analyzer at: http://technet.microsoft.com/en-us/library/dd392255(v=ws.10).aspx

Storage Spaces

The storage needs of companies continues to grow rapidly these days, and while hard drive costs have fallen as capacity has increased, the cost of enterprise storage systems such as Storage Area Networks (SANs) remains high and can be a drain on shrinking IT budgets. The need for innovative new storage solutions to address these issues has led Microsoft to introduce a new technology in Windows Server 2012 called *Storage Spaces* that allows you to virtualize storage by grouping commodity disks into pools, from which virtual disks (also called storage spaces) can be easily provisioned as your storage needs evolve.

To make optimum use of Storage Spaces in a Hyper-V environment, you need to understand it thoroughly and configure it properly. In this section Satya Ramachandran explains the technology underlying storage spaces and provides some advice on how to plan, deploy, maintain, and troubleshoot this feature.

What is Storage Spaces?

Storage infrastructure in today's enterprise is a complex world. Gone are the days where each server had a few local disks and the only layer of abstraction was hardware or software RAID. As we have progressed in the journey toward achieving fast and reliable storage, we have lost flexibility, transparency, and simplicity while provisioning storage.

Today, if an application team needs storage capacity, they reach out to the storage team, and once they have carved out some storage, the application team continues with their deployment. These transactions cost our businesses a precious commodity—time.

In this section we are going to talk about Storage Spaces, which is a new feature of Windows Server 2012. It not only addresses some of the issues described above but many more practical challenges an administrator faces in the real world.

Storage Spaces is a new storage virtualization technology introduced in Windows Server 2012 which gives enterprises access to enterprise-class storage features on just about any hardware…really, *any* hardware. Imagine having options to pick different types of resiliency, thin provisioning, high availability, and scalability all from commodity hardware. Because of its simple requirements, it drastically cuts down storage costs associated with hosting virtualized workloads.

Before we get into technical implementation details, let's take a moment to understand what Storage Spaces brings to the table; it will definitely help you engage your peer in a nice technical conversation over lunch or help you save your business a lot of cash during your next storage upgrade.

Concepts and terms

One of the challenges associated with adopting any new technology is the learning curve associated with it. The terms used in Storage Spaces are mostly the same ones storage administrators have been using with traditional storage vendors. Let's go over a few of the important ones:

- **Primordial** The dictionary meaning of primordial is constituting a beginning, and that's the role it plays in Storage Spaces. All the unallocated disks that meet the requirements of Storage Spaces show up here.
- **Storage pools** A storage pool is an equivalent of a pool or an array in the storage world, out of which we can carve different types of virtual disks. It's a unit of aggregation, administration, and isolation. A single pool can consist of heterogeneous physical disks, which can be different sizes or connected via different interfaces/interconnects. We can specify the role of disks as data or hot spare.
- **Virtual disk, or storage space** This is the layer at which we specify the virtual disk size and resiliency. After this we can create a volume and define:
  - File system (can be NTFS or REFS)
  - Cluster size
  - Deduplication settings
  - Drive letter

  After the volume is created, the storage can be utilized for any of the Windows Server 2012 capabilities to host applications and services.

Here is a diagram that illustrates the connection between all of these concepts:

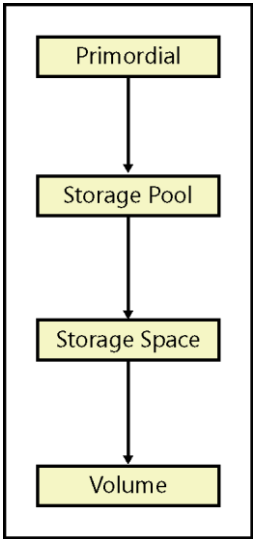

This next diagram shows in more detail how you can create storage pools and spaces in Windows Server 2012:

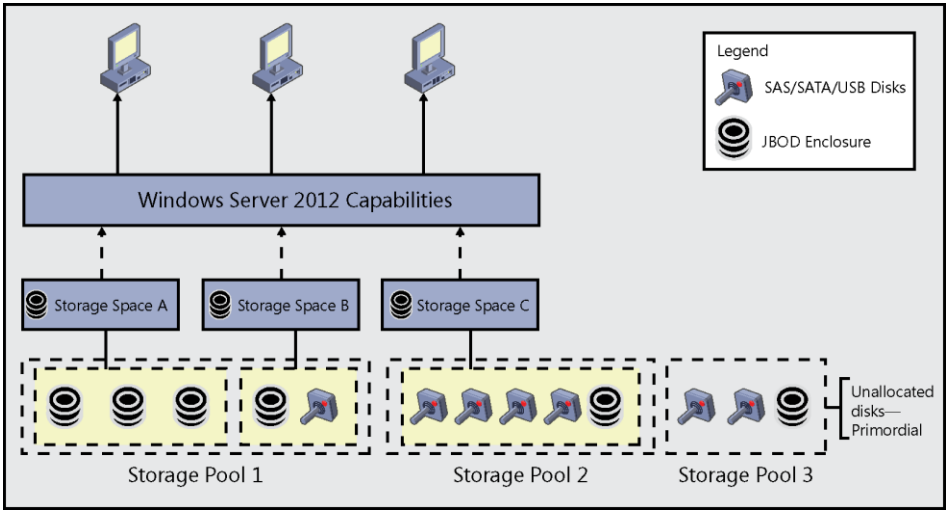

Deployment modes

Depending on the storage space deployment model, we can offer block level or file level access. So as an enterprise we do not need to make a huge investment to purchase an expensive storage solution for either of the two deployment modes.

The following diagram illustrates these two different usage scenarios that Storage Spaces enables:

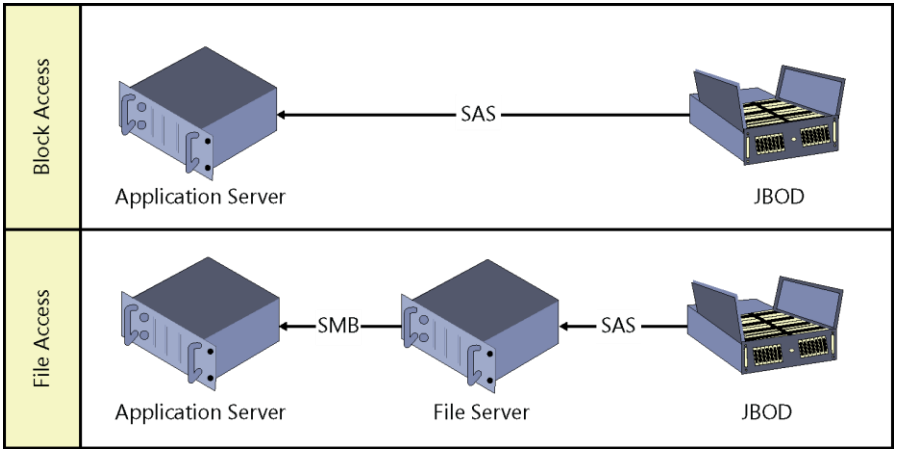

Benefits of Storage Spaces to enterprises

Now that we know what Storage Spaces is, let's take a few minutes to discuss its benefits to an enterprise. Though Storage Spaces can be used for many scenarios, since this book is about Hyper-V storage we will talk about benefits of Storage Spaces when used in Hyper-V deployments.

Cost effective platform for business critical storage

Storage consumes a large chunk of the virtualization hardware budget. With the help of storage spaces we can utilize a couple of JBOD enclosures presented via a few file servers and have a high performance, highly available and scalable storage solution:

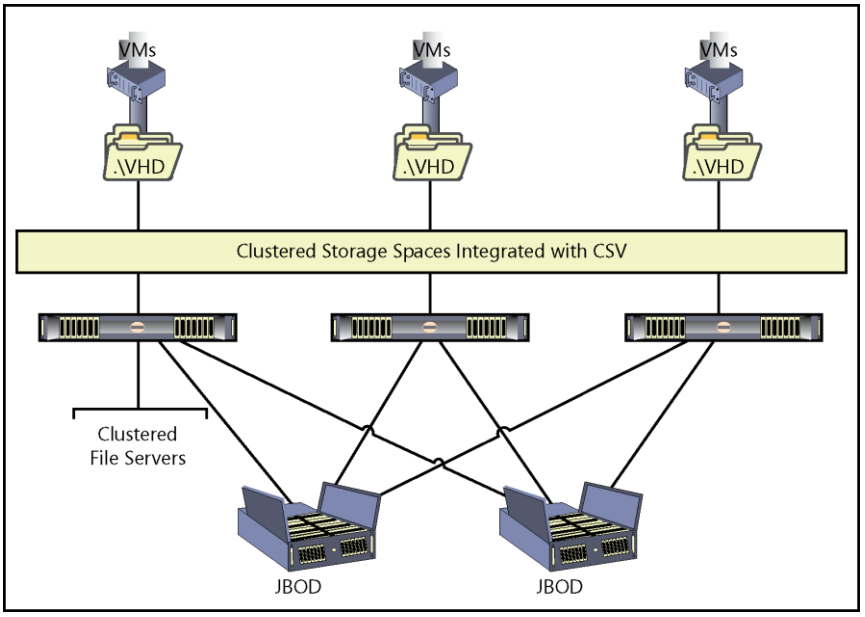

With SSDs in enclosures we would be looking at millions of IOPS from the solution with only the server CPU and storage controllers being the bottleneck. Once we saturate the solution, we can scale out by adding more servers or JBODs.

Flexibility and elasticity

Administrators today are tasked with sizing a solution that is expected to scale to infinity, but with limited resources. It's every administrator's dream to have a solution that can be scaled out without having to worry about application reconfiguration and massive outage windows. With Storage Spaces, as the business requirement grows we can increase the size of the pools by adding more disks; the spaces created on top can be thinly provisioned and can be of much larger size than the currently available storage in the pool.

For example, we can create a storage pool with 20 terabytes worth of disks and then create a space of 380 terabytes on top. As the usage increases, we would be alerted once we are at the seams of the real storage and can add more disks to the pool with no impact to application or service.

The other benefit is storage optimization. Although the application teams may ask for a large chunk of storage up front, they may not utilize the capacity immediately. With monitoring tools we can determine when we are running out of real capacity and add more disks as we go. This also helps us stagger capital expenditure on the infrastructure, at the same time application teams, or customers in the case of a hosted model, would be charged for real storage used.

NOTE **Only fixed storage spaces can be created when integrated with cluster shared volume.**

Resiliency and data integrity

With storage spaces we can create a space with four different types of protection:

- **Simple** No protection from disk failure; data is striped across all disks. It should only be used when we have easily replaceable data. This should not be used to host business-critical data.
- **Mirror** We maintain multiple copies of data; ideal for hosting business critical data.
  - **Two-way mirror** Two copies are maintained and we can tolerate one disk failure.
  - **Three-way mirror** Three copies are maintained and we can tolerate two disk failures.
- **Parity** Data is striped across all disks along with parity information to regenerate data in case of a disk failure. We can tolerate a single disk failure with this model.

Based on the protection type, the minimum number of disks required varies.

Storage Spaces, when integrated with cluster shared volume, can unify storage access and simplify management. All cluster nodes can access the storage concurrently, irrespective of the number of nodes and the number of JBOD enclosures. In case of a server failure, the workloads will transparently failover to an alternative server without any outage. This also makes it possible to take servers offline for maintenance without impacting service.

Multi-tenancy

As an organization, we may have requirements to segment storage for charge back, compliance, or political reasons, and as a multiple-customer hosting provider, we need to offer isolation at all layers. With Storage Spaces we can delegate management of storage pools to individual customers or application teams so that they can provision virtual disks and volumes as per their business requirements.

Ease of management

Storage Spaces can be managed from Server Manager UI or via PowerShell. With the new Server Manager an administrator can manage and monitor storage fabric from a single user interface. From creating a pool to creating a volume, all of it can be accomplished using the same management console. To make it even simpler, the wizards have a chaining feature by which the administrator has an option of triggering the next logical action. For example, after creating a pool, the next logical step is to create a virtual disk. All the administrator needs to do is select the check box for launching the next wizard on the final screen and as soon as the action completes the next wizard is presented. Storage management can't get simpler than this:

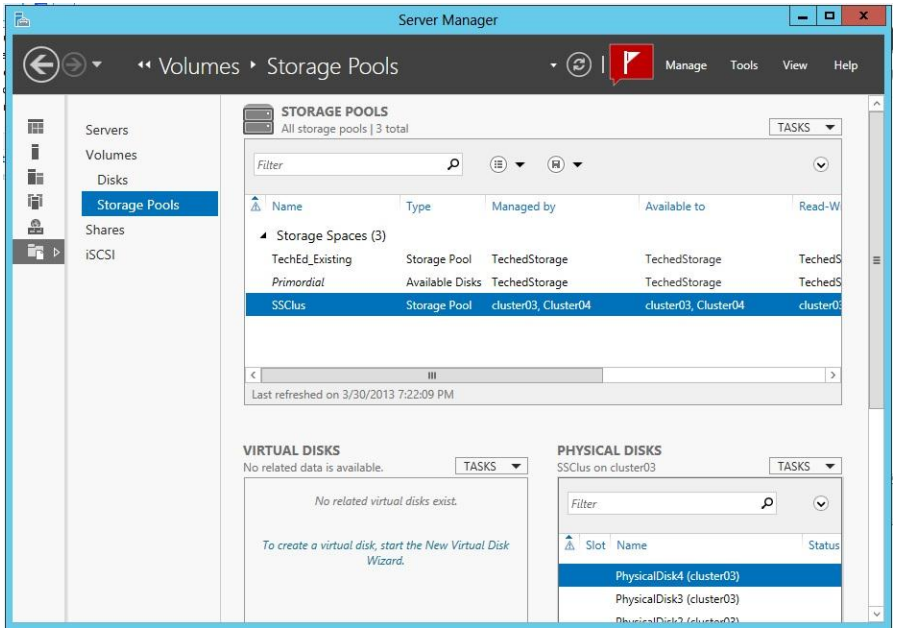

For the command line ninjas, there is PowerShell support for Storage Spaces:

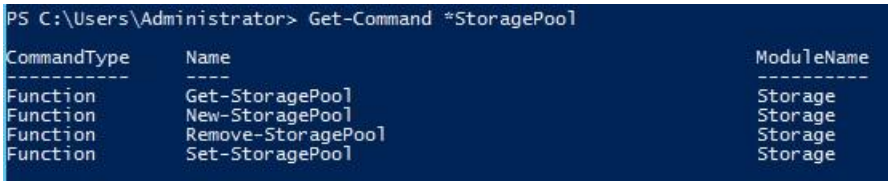

This might get you thinking: how about a storage self-provisioning portal for your dev/test environment with a Web UI that calls the PowerShell cmdlets and completes the task?

Before we start

I am sure by now you are very excited to deploy your first storage space and get a feel for the solution, so let's go over the prerequisites of Storage Spaces:

- **Server operating system** You need Windows Server 2012 Standard or Data Center Edition. Although the Windows 8 client also has this capability, we will not spend any time on client-side implementation of Storage Spaces.
- **Interconnects** Storage Spaces supports:
  - SATA
  - SAS
  - USB
  - SCSI

  It's important to call out that there is no support for FC or iSCSI.
- **Storage with no magic at any layer** All the storage should be directly exposed to the operating system with no abstraction. So we need to turn off all the fancy features of the storage controller and let the disk be completely managed by the operating system.

- **Disks** The disk should be a minimum of 10 GB, unpartitioned, and uninitialized. Depending on the resiliency options, the number of disks required will vary. For simple space, you can start with one disk, two for two-way mirror, three for parity, and five for three-way mirror. The disk requirements may change depending on the number of columns you plan to have (to be discussed later).
- **Hardware** Although this works with any JBOD enclosure, there are some new Storage Spaces certified enclosures that support a feature called SCSI Enclosure Services, or SES. With SES an administrator can turn on the drive light to make it easier to identify a disk that needs to be replaced:

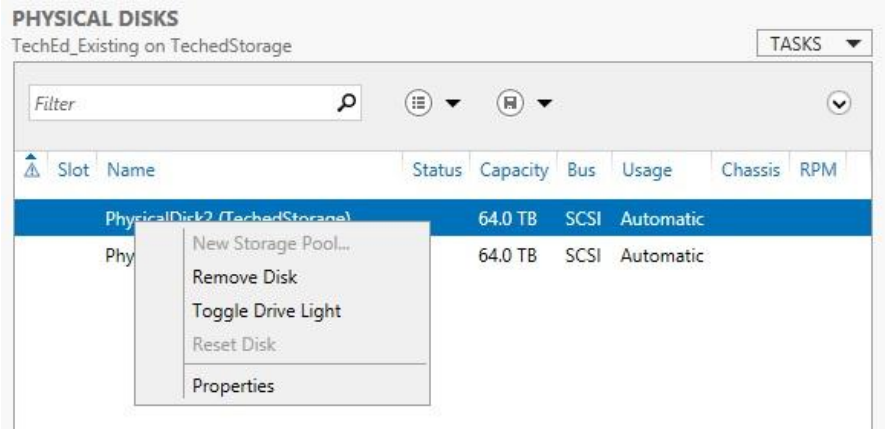

Deploying your first storage space

To start your journey of creating the first storage space, I would recommend a system with Hyper-V enabled and a virtual machine with Windows Server 2012 Standard Edition. (Yes, the evaluation edition works too.) While authoring this section, I set the whole lab on my laptop running Windows 8 Enterprise with Hyper-V enabled. I created a bunch of virtual machines with Windows Server 2012 and I was set! But if you are one of those fortunate souls who has access to a lab with all kinds of hardware, you can definitely try the same on real hardware and experience the performance.

For this walkthrough I am using my member server, which hosts all the storage requirements of my lab enterprise.

1. Launch Server Manager and click File And Storage Services:

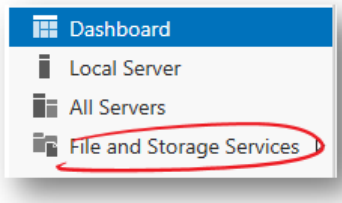

2. Click Storage Pools:

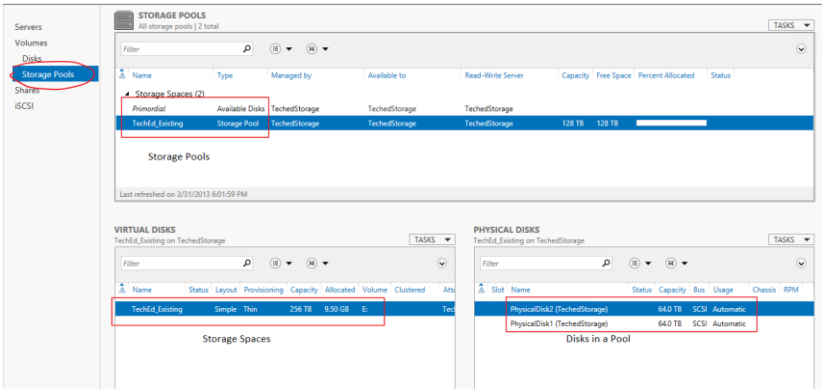

Let's take a moment to go over the user interface. The top box is for Storage Pools. We start with one—Primordial—and all the available disks that meet the requirement for Storage Spaces are listed below it. In the lower-right pane are the physical disks in a pool; information about disk size, the interconnect, and its usage (Automatic or Hot Spare) is displayed. Virtual disks or storage spaces are in the lower-left pane; here we get information about resiliency, provisioning type (Thin or Fixed), and the size of the spaces.

The previous illustration shows two 64-terabyte disks in the pool adding up to 128 terabytes of storage pool. In this storage pool we have created a thinly provisioned simple virtual disk of 256 terabytes.

From PowerShell, we can get the same information via Get-StoragePool and Get-PhysicalDisk. There is more detail available, such as allocated size, physical sector size, and so on, which is available only via user interface.

Hence, I encourage you to explore the PowerShell cmdlets:

```
PS C:\Users\Administrator.VIRTUAL> Get-StoragePool

FriendlyName          OperationalStatus       HealthStatus          IsPrimordial          IsReadOnly
------------          -----------------       ------------          ------------          ----------
TechEd_Existing       OK                      Healthy               False                 False
Primordial            OK                      Healthy               True                  False


PS C:\Users\Administrator.VIRTUAL> Get-PhysicalDisk

FriendlyName        CanPool           OperationalStatus   HealthStatus        Usage                          Size
------------        -------           -----------------   ------------        -----                          ----
PhysicalDisk3       True              OK                  Healthy             Auto-Select                    30 GB
PhysicalDisk0       False             OK                  Healthy             Auto-Select                    30 GB
PhysicalDisk1       False             OK                  Healthy             Auto-Select                    64 TB
PhysicalDisk2       False             OK                  Healthy             Auto-Select                    64 TB
```

3. Select the Primordial storage pool, and then on the Tasks drop-down menu select New Storage Pool:

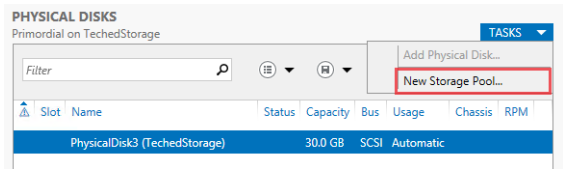

This starts the New Storage Pool Wizard.

4. Click Next on the first page, and on the second page enter a name for your first storage pool:

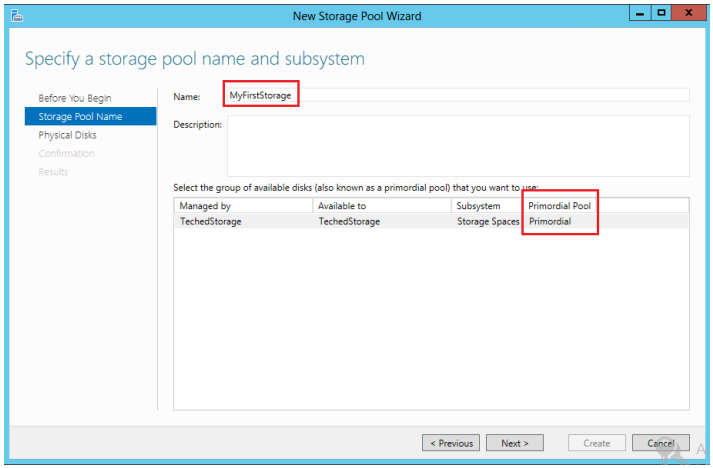

Primordial Pool is listed in the lower pane. Click Next.

5. On the Physical Disks page, select the disks you would like to add to the pool; you can also select the allocation. For production setups you definitely want a few disks marked as hot spare so that Storage Spaces can replace the dysfunctional disks in the pool automatically:

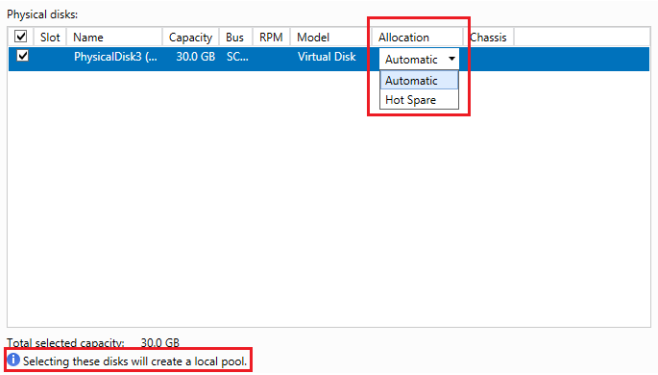

The other useful user interface element is the text indication of whether the created pool will be a clustered or local pool. Click Next.

NOTE **There has been some debate about when the hot spare gets activated. Some people pulled out a disk to simulate a problem and then observed that the hot spare did not immediately register any I/O. From my experience, the hot spare gets activated when there is a bad disk and substantial write I/O is done to the space.**

6. Review the summary and click Create.
7. On the Completion Summary page, another interesting user interface element is the option to chain wizards.

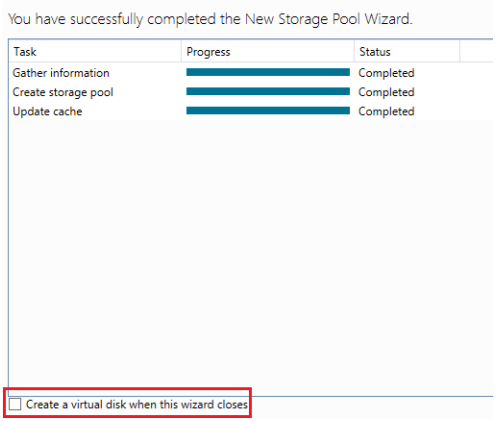

The option guides a first-time user to the next logical step. Select the check box and click Finish. This starts the New Virtual Disk Wizard.

8. Before we go over the New Virtual Disk Wizard, let's review the Server Manager user interface after the storage pool has been created:

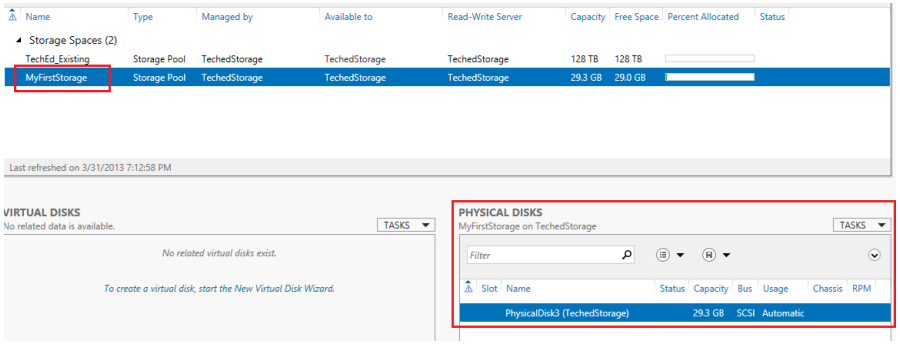

You can see the new storage pool. Click it to show the disks mapped to it in the Physical Disks section. As expected, at this point there are no virtual disks hosted on this pool.

9. Returning to the New Virtual Disk Wizard, click Next on the Before You Begin page.
10. Select the newly created storage pool and click Next:

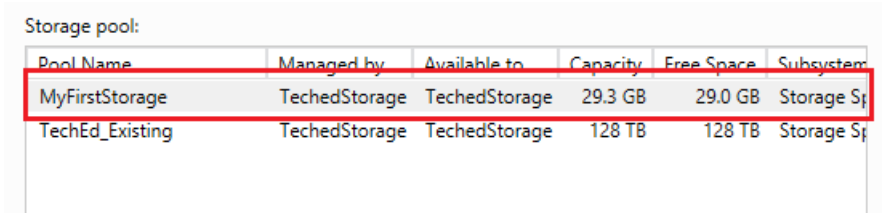

11. Provide a virtual disk name and click Next.
12. On the Storage Layout page, click Mirror and then click Next:

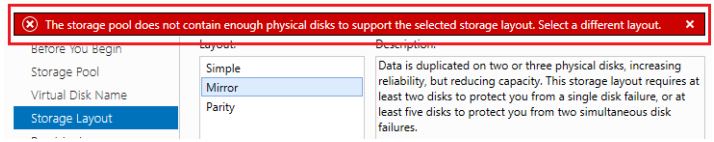

Since the storage pool I created had only one disk, the wizard blocks the creation of mirrored storage space. Click Simple and then click Next.

13. Select Thin as the provisioning type and then click Next:

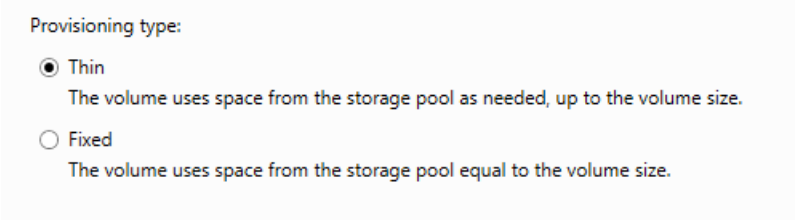

14. On the Size page, set the size upward of 1 GB and click Next, review the summary screen, and then click Create. Once the virtual disk creation is complete, the wizard presents an option to launch the next logical action: Create a volume:

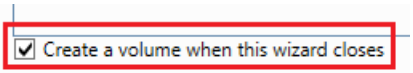

15. Click Next on the Before You Begin page, select the newly created virtual disk, and then click Next:

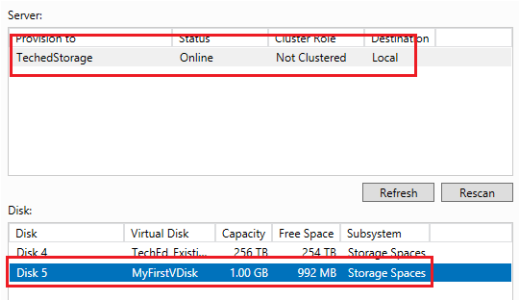

16. Leave the default value on the Size page and click Next.
17. Assign a drive letter and click Next.
18. Select either REFS or NTFS as the file system, provide a volume label, and click Next.

NOTE **REFS is a new file system introduced in Windows Server 2012. It works very well with mirrored spaces where, when the metadata or data fails the integrity check or checksum, REFS restores a good copy from a working mirror.**

19. Storage spaces supports deduplication under two conditions:
    - The file system selected is NTFS.
    - Storage Spaces is not integrated with Failover Cluster.

    Enable deduplication, taking a moment to familiarize yourself with the scheduling options available, then click Next.
20. Click Create.

You now have your first fully functional storage space ready to host data.

A little bit of theory

Now that we have our first storage space up and ready, the next step is to understand the implementation details, or as we like to call it at Microsoft— to go under the hood.

First, let's understand the components of the storage stack (see the following diagram). I have limited the scope of discussion to where Storage Spaces plugs in. Those who are interested in learning more can consider reading an excellent reference book from Microsoft Press called *Windows Internals, Fifth Edition* by Mark E. Russinovich, David A. Solomon, and Alex Ionescu. (You can find it here: http://shop.oreilly.com/product/9780735625303.do.)

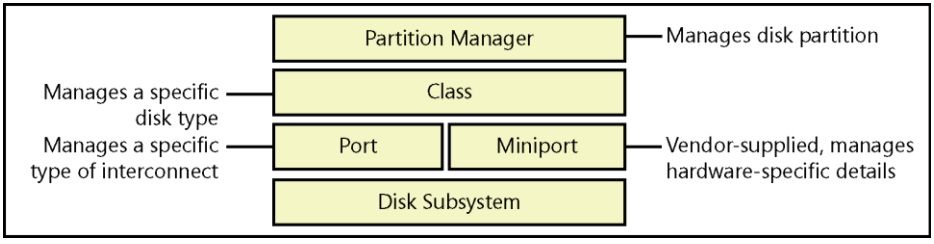

- **Disk class, port, and miniport drivers** During system start, Windows I/O manager starts the disk storage drivers. Storage drivers in Windows follow a class/port/miniport architecture where Microsoft provides a storage class driver that implements functionality common to all storage devices and a port driver that implements functionality common to a particular bus.
- **Partition manager** The partition manager (partmgr.sys) is responsible for discovering, creating, deleting, and managing partitions.

When you enable Storage Spaces, you utilize a driver Spaceport.sys to plug into the storage stack above the partition manager. It's responsible for sending notifications about arrival, removal, and modification of a storage pool, drive, or spaces:

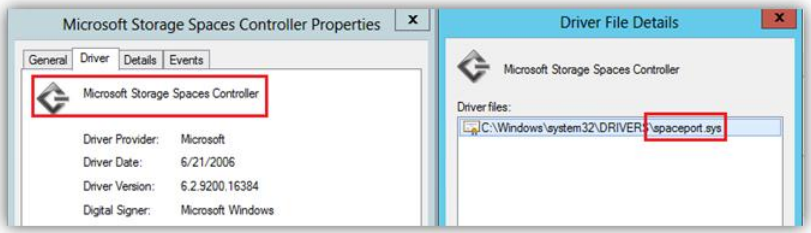

On each disk hosting a storage pool, we create a special partition called Spaces Protective Partition. It contains the metadata and the hosted data in a space in that pool:

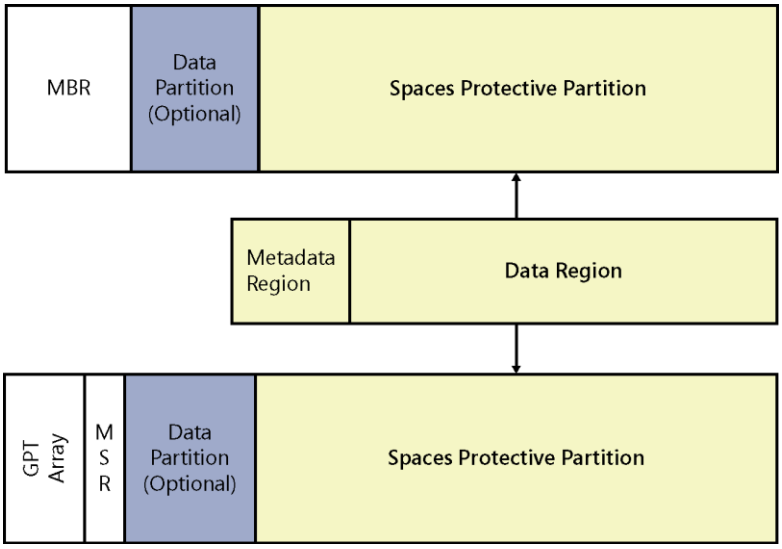

Spaceport.sys interfaces with an inbox hardware provider (smspace.dll) to manage the Storage Spaces subsystem:

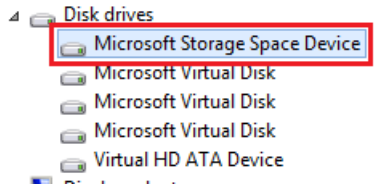

Now let's take a look at how a spaces pool shows up in Device Manager and Disk Management.

In Device Manager, under Disk Drives, the virtual disk is listed as Microsoft Storage Space Device. Disk Management lists the disk as a basic disk; however, the device type is Storage Spaces:

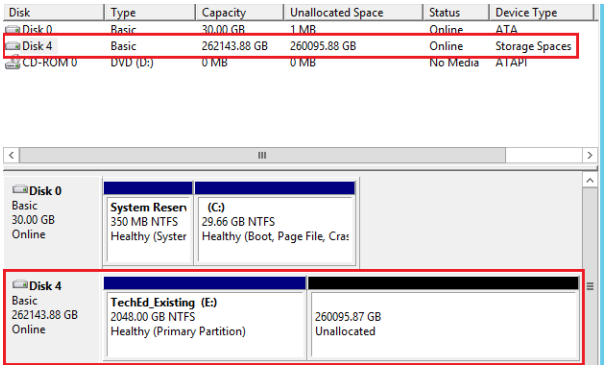

Planning your storage space

Before you start hosting business-critical data with Storage Spaces, a good amount of time should be spent on planning and making the right choices. I am sure no administrator likes to sit on a multi-terabyte mistake that requires many hours to correct over the weekend.

Resiliency and performance tuning

Before we click any option on the drop-down menu, we need to consider what workloads are best suited for each protection. Each protection type has pros and cons. For example, Simple offers no protection in case of disk failure, but it is the fastest in terms of performance. Also the usable storage varies for each protection type.

For extracting the best performance out of storage spaces, we need to fine tune the column and interleave values. Let's first understand what a column is and what an *interleave* is.

The following diagram shows the data is striped across three disks; each disk is a column, and the stripe is the equivalent of a row in a table spanning the three disks. Interleave would be each stripe unit:

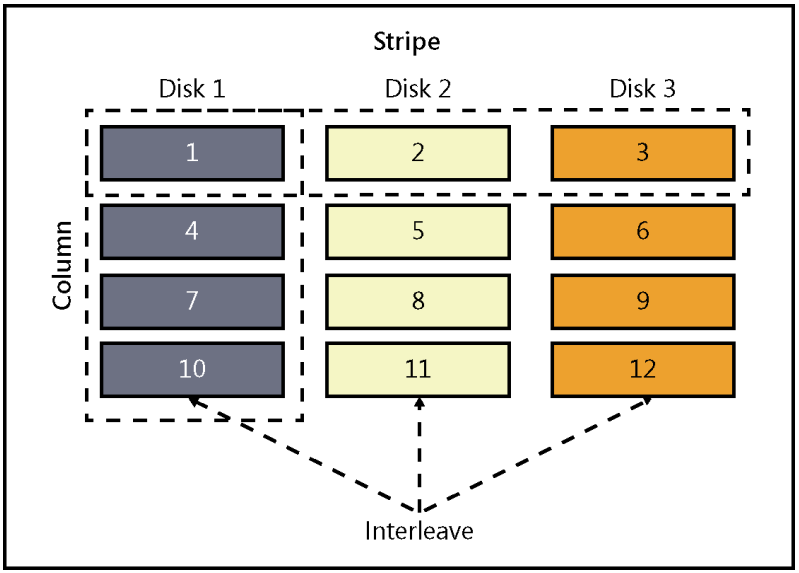

When creating a storage space, we can specify how many columns we want or the size of an interleave.

We can define a column value equal to or less than the number of disks. This is important because the number of columns defines how many disks we concurrently access. Some points to keep in mind when creating storage spaces:

- The maximum number of columns GUI uses is eight.
- PowerShell allows you to specify a parameter—NumberofColumns greater than eight.
- Parity spaces cannot have more than eight columns.
- Mirrored spaces, when created using the UI, allow for up to four columns; you can specify a larger number when creating the space from PowerShell.

Interleave is the second most important parameter because it specifies the amount of data that will be stored in each column. To determine the interleave size, we need to know the typical I/O size for the data being streamed to the space. If the I/O size is larger than the interleave size, data will be scattered across columns, and that translates into a single write operation resulting in multiple writes.

For mirrored spaces, another parameter to consider is NumberOfDataCopies. We have the option of creating two or three copies of data, but this comes at the cost of usable storage space. For a two-way mirror the effective usable space is 50 percent, and for a three-way mirror the effective usable space is 33 percent.

Even when customers are placing non-critical data, I recommend they go with a model that offers at least a single disk failure tolerance so that they can avoid starting from scratch in the case of a disk failure or they can upgrade disks in enclosure without having to backup, delete the existing virtual disk, and create a new virtual disk.

Another deployment model to consider is a mix in the number of columns and mirrors. For example, with four disks, you can specify the number of columns as two and the number of copies as two. If you have more disks, you can consider increasing the column count, thereby increasing the number of disks concurrently servicing the I/O requests.

Thin provisioning

Beyond resiliency and performance tuning, another sought-after feature is thin provisioning.

The thick provisioning method provisions or allocates the same amount of resources reported to the storage subsystem above. The thin provisioning method provisions or allocates fewer resources than what is reported to the storage subsystem above.

The immediate business benefit of thin provisioning is efficiency in storage utilization because it does not block off storage when it will not be utilized immediately. Instead storage is allocated when files expand.

Thin provisioning is supported for standalone spaces deployment, but it's not available for spaces integrated with clustering. Only fixed or thick provisioning is available when deploying on a cluster.

So there is no one answer to which model is best suited for Hyper-V; it depends on the resiliency and performance goals plus I/O size.

Maintaining storage spaces

Here we will examine three aspects of maintaining storage spaces:

- Extending a virtual disk
- Removing a disk from a pool
- Rebuilding a server that hosts storage spaces

Extending a virtual disk

Increasing the size of an existing storage space is simple but requires careful planning. When we create a storage space via UI or PowerShell, we either explicitly specify the number of columns or the system assigns the number. When increasing the size, we cannot alter the number of columns; the same striping model needs to be followed.

For example, if you tried to add a single disk to a two-disk simple space with two columns, the operation would fail since it would break the striping model. Disks must be added in multiples of two to allow the system to maintain the current striping model.

So a simple formula to follow is

*number of data copies * number of columns*

This can be determined by using the following PowerShell command:

```
Get-VirtualDisk | ft FriendlyName, ResiliencySettingName, NumberOfColumns,
NumberOfDataCopies
```

Removing a disk from a pool

The other activity that requires careful planning is removing a disk from a pool. Once a disk is in use, you may have to remove it for one of two reasons:

- Disk has failed or is in an unhealthy state
- Disk upgrades are required

While the replacement of an unhealthy disk is a straightforward operation, replacement of a healthy disk is a bit tricky since it holds data. Here we need to explore Hyper-V storage migration since the only graceful method of removing an actively used disk is to delete the virtual disk hosted on it and then remove the disk from the pool. So an approach to consider is to migrate the virtualized workloads to the new virtual disk and then delete the old virtual disk.

Rebuilding a server that hosts storage spaces

Sometimes a server rebuild may be necessary. In those cases it's important to know what post-rebuild actions are required to recover storage spaces.

After a server is rebuilt, the storage pool defaults to read-only mode and all the spaces hosted on it are detached. To bring the storage pool and space to normal operational state, use the following PowerShell cmdlets:

```
Get-StoragePool | Where-Object {$_.IsReadonly –eq  $True } | Set-StoragePool –IsReadonly
$False
Get-VirtualDisk | Where-Object {$_.IsManualAttach –eq $True} | Set-VirtualDisk –
IsManualAttach $False
```

The next section on troubleshooting storage spaces will go over other maintenance scenarios, such as addressing a storage pool that's not visible or a corrupt pool.

Troubleshooting storage spaces

Before we discuss scenarios, I would like to share some common guidelines for avoiding long cycles of troubleshooting:

1. **Read the prerequisites** A majority of the questions about unsuccessful storage spaces deployment are due to prerequisites not being met. For example, many people miss the interconnect requirement; Storage Spaces supports only SAS, SATA, SCSI, and USB. FC and iSCSI are unsupported. So before you start your deployment take a moment to make sure all the prerequisites are met.
2. **Check the health and operational status of the disks** Before looking around for help, run the following commands to review the health of the disks, pools, and spaces. These PowerShell cmdlets will save a lot of time when you see unexpected behavior:
   - List unhealthy spaces:

   ```
   Get-VirtualDisk | Where-Object {$_.HealthStatus -ne "Healthy"}
   ```

   - List unhealthy storage pools:

   ```
   Get-StoragePool | Where-Object {$_HealthStatus -ne "Healthy"}
   ```

   - List unhealthy physical disks:

   ```
   Get-PhysicalDisk | Where-object {$_HealthStatus -ne "Healthy"}
   ```

3. **Take a look at event logs** Event logs are great friends of an administrator. In the event logs, you will find the errors encountered along with potential remedies. Good starting points are:
   - Applications and Services Logs\Microsoft\Windows\Spaceport\Analytic
   - Applications and Services Logs\Microsoft\Windows\Spaceport\Operational

   Sometimes you may experience issues because of a broken underlying component. Good starting points for identifying these issues are:
   - Applications and Services Logs\Microsoft\Windows\Disk
   - Applications and Services Logs\Microsoft\Windows\StorDiag
   - Applications and Services Logs\Microsoft\Windows\StorPort

We will now examine three troubleshooting scenarios:

- Creating a storage space fails
- Deleting a storage space fails
- Expanding a storage space fails

To make it easier, I have created easy to follow flowcharts for these three scenarios. I hope you find these useful when troubleshooting your implementation of storage spaces.

Creating a storage space fails

The following flowchart can be used to troubleshoot a scenario where creating a storage space fails:

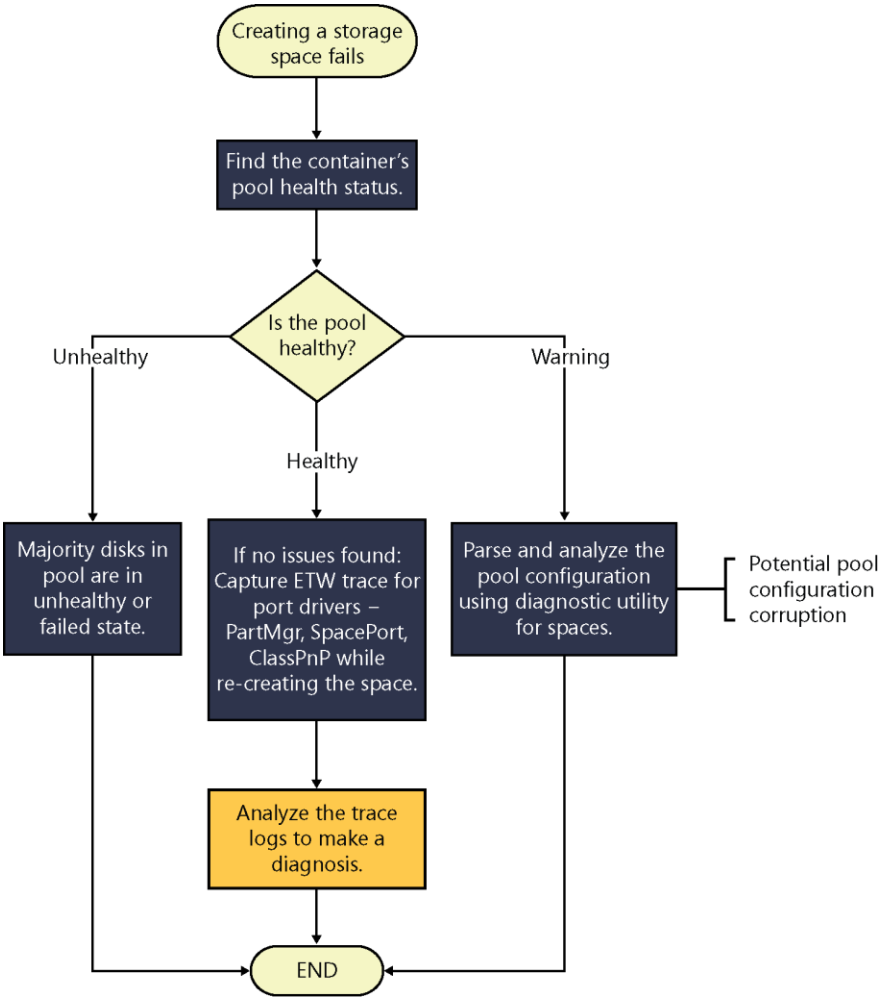

Deleting a storage space fails

The following flowchart can be used to troubleshoot a scenario where deleting a storage space fails:

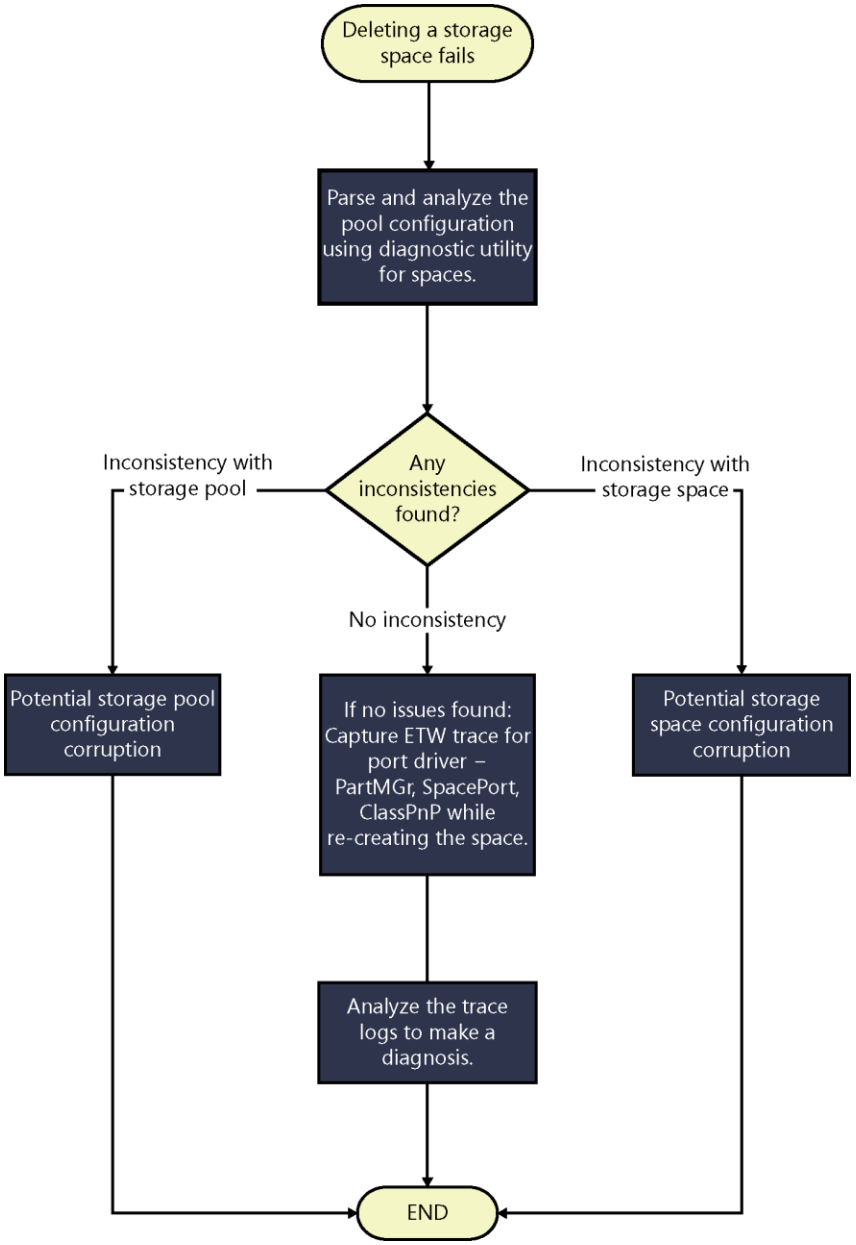

Expanding a storage space fails

The following flowchart can be used to troubleshoot a scenario where expanding a storage space fails:

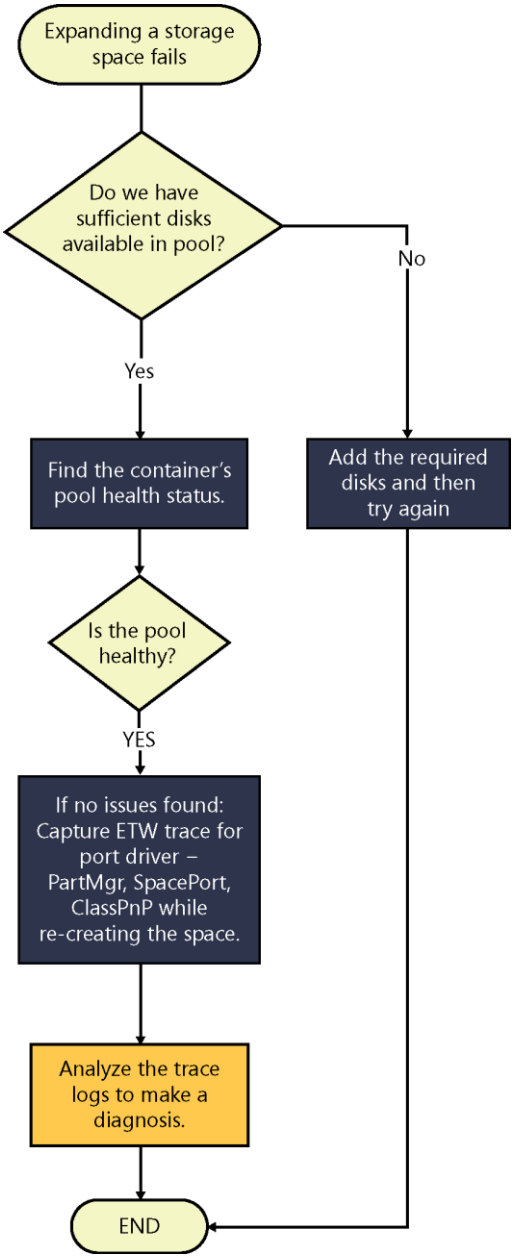

—Satya Ramachandran, Premier Field Engineer

Additional resources

Here are a few additional resources concerning this topic:

- Storage Spaces Overview (TechNet Library) at: http://technet.microsoft.com/en-us/library/hh831739.aspx
- Storage Spaces Frequently Asked Questions (FAQ) (TechNet Wiki) at: http ://social.technet.microsoft.com/wiki/contents/articles/11382.storage-spaces-frequently-asked-questions-faq.aspx

Building a demo environment

Sometimes the best way to troubleshoot something is to avoid problems in the first place. Complex Hyper-V configurations can be time consuming to set up manually and can lead to issues that are hard to troubleshoot. Windows PowerShell is useful because you can use it to automate the configuration process to ensure it's free of errors. Of course that means your Windows PowerShell commands and scripts must also be free of errors.

A good way to learn how to create error-free configuration scripts is to study and then customize scripts created by experts. In this section Jose Barreto provides an in-depth look at the scripts he uses to set up a demo environment for a fault-tolerant Hyper-V over SMB 3.0 infrastructure based on Windows Server 2012. Jose also concludes his section with some links to blog posts where he has covered a number of different SMB 3.0 tips and tricks.

TIP **You can download a zip file containing all the Windows PowerShell scripts in this demo from http://aka.ms/TroubleshootHyper-VStorage/files.**

Hyper-V over SMB: Step-by-step installation using Windows PowerShell

This section describes the steps I used to create a Windows Server 2012 File Server test environment that I used for some of my Hyper-V over SMB demonstrations. The subsections below are as follows:

- Overview
- Environment details
- Script #1: Configuring FST2-DC1 (DNS, Domain Controller, iSCSI Target)
- Script #2: Configuring FST2-FS1 (File Server 1)
- Script #3: Configuring FST2-FS2 (File Server 2)
- Script #4: Configuring FST2-HV1 (Hyper-V host 1)
- Script #5: Configuring FST2-HV2 (Hyper-V host 2)
- Script #6: Configuring the Cluster FST2-FSC (run from FST2-FS1)

- Script #7: Configuring the Classic File Server Cluster FST2-FS (run from FST2-FS1)
- Script #8: Configuring the Scale-Out File Server Cluster FST2-SO (run from FST2-FS1)
- Script #9: Configuring the virtual machines in FST2-HV1
- Script #10: Configuring the virtual machines in FST2-HV2
- Script #11: Creating a Hyper-V Cluster using file share storage
- Script #12: Optional steps to create a nonclustered file share on FST2-FS1
- Additional resources

Overview

The goal is to share some of configuration details and the exact Windows PowerShell scripts I used to configure the environment. (If you look carefully, you might be able to spot a few Windows PowerShell tricks and tips.)

This setup uses five physical machines, since the scenario involves deploying Hyper-V hosts and you can't virtualize Hyper-V itself. I also use RDMA interfaces on the setup with SMB Direct, and those also can't be virtualized. The demo setup includes one domain controller (which also doubles as an iSCSI target), two file servers, and two Hyper-V hosts.

This is probably the most basic fault-tolerant Hyper-V over SMB setup you can create that covers the entire spectrum of new SMB 3.0 capabilities (including SMB Transparent Failover, SMB Scale-Out, SMB Direct, and SMB Multichannel).

Please keep in mind that this is not a production-ready configuration. I built it entirely using five-year-old desktop class machines. To improved disk performance, I did add three SSDs to one of the machines to use as storage for my cluster, which I configured using Storage Spaces and the Microsoft iSCSI Software target included in Windows Server 2012. However, since I only had three small SSDs, I used a simple space, which cannot tolerate disk failures. In production, you should use mirrored spaces. Also keep in mind that the FST2-DC1 machine itself is a single point of failure, so you're really only tolerant to the failure of one of the two Hyper-V hosts or one of the File Server nodes. In summary, this is a test-only configuration!

Environment details

The environment is deployed as five physical machines, all using the FST2.TEST domain. Here's a diagram of the setup so you can better understand it:

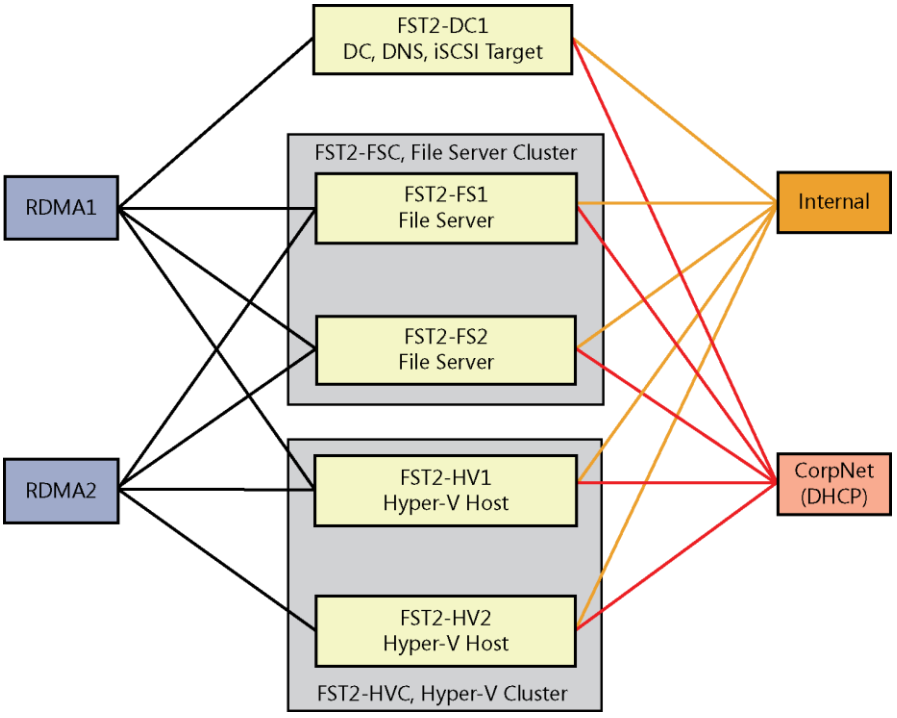

Here are the details about the names, roles, and IP addresses for each of the computers involved, including the cluster objects and virtual machines:

Computer name: FST2-DC1

- Roles: DNS, Domain Controller, iSCSI Target
- External network: DHCP
- Internal network: 192.168.100.10/24
- RDMA 1: 192.168.101.10/24
- RDMA 2: N/A

Computer name: FST2-FS1

- Roles: File Server 1
- External network: DHCP
- Internal network: 192.168.100.11/24
- RDMA 1: 192.168.101.11/24
- RDMA 2: 192.168.102.11/24

Computer name: FST2-FS2

- Roles: File Server 2
- External network: DHCP
- Internal network: 192.168.100.12/24
- RDMA 1: 192.168.101.12/24
- RDMA 2: 192.168.102.12/24

Computer name: FST2-HV1

- Roles: Hyper-V Server 1
- External network: DHCP
- Internal network: 192.168.100.13/24
- RDMA 1: 192.168.101.13/24
- RDMA 2: 192.168.102.13/24

Computer name: FST2-HV2

- Roles: Hyper-V Server 2
- External network: DHCP
- Internal network: 192.168.100.14/24
- RDMA 1: N/A
- RDMA 2: 192.168.102.14/24

Computer name: FST2-FSC

- Roles: File Server Cluster Name Object
- External network: DHCP
- Internal network: N/A
- RDMA 1: N/A
- RDMA 2: N/A

Computer name: FST2-FS

- Roles: Classic File Server Cluster
- External network: N/A
- Internal network: 192.168.100.22/24
- RDMA 1: 192.168.101.22/24
- RDMA 2: 192.168.102.22/24

Computer name: FST2-SO

- Roles: Scale-Out File Server Cluster
- External network: N/A
- Internal network: N/A
- RDMA 1: N/A
- RDMA 2: N/A

Computer name: FST2-HVC

- Roles: Hyper-V Cluster Name Object
- External network: DHCP
- Internal network: N/A
- RDMA 1: N/A
- RDMA 2: N/A

Computer name: FST2-VM*

- Roles: Virtual Machine
- External network: DHCP
- Internal network: N/A
- RDMA 1: N/A
- RDMA 2: N/A

Last but not least, here's a picture of the setup, so you can get a sense of what it looks like:

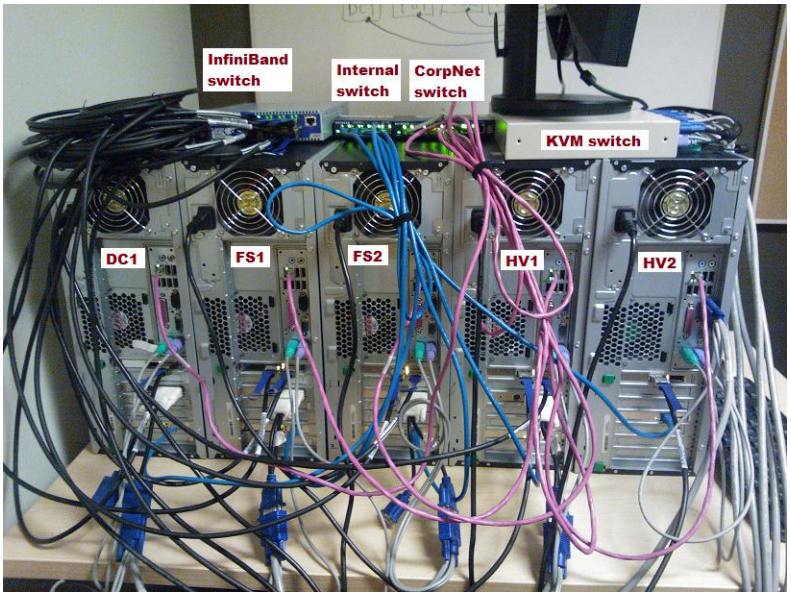

Script #1: Configuring FST2-DC1 (DNS, Domain Controller, iSCSI Target)

```
# Note 1: This assumes you already installed Windows Server 2012 and configured the
computer name.
# Note 2: This setup uses InfiniBand RDMA interfaces.

#
# Set power profile
#
POWERCFG.EXE /S SCHEME_MIN

#
# Configure all 4 interfaces (1 DHCP, 3 static)

#
# Rename External, no further action required, since this is DHCP
#
Get-NetAdapter -InterfaceDescription "*Intel*"   | Rename-NetAdapter -NewName "External"

#
# Rename Internal, set to manual IP address, configure IP Address, DNS
#
Get-NetAdapter -InterfaceDescription "*Realtek*" | Rename-NetAdapter -NewName "Internal"
```

```
Set-NetIPInterface -InterfaceAlias Internal -DHCP Disabled
Remove-NetIPAddress -InterfaceAlias Internal -Confirm:$false
New-NetIPAddress -InterfaceAlias Internal -IPAddress 192.168.100.10 -PrefixLength 24
Set-DnsClientServerAddress -InterfaceAlias Internal -ServerAddresses 192.168.100.10

#
# Rename RDMA1, set to manual IP address, configure IP Address, DNS
#
Get-NetAdapter -InterfaceDescription "*IPoIB*" | Select -Last 1 | Rename-NetAdapter -
NewName RDMA1
Set-NetIPInterface -InterfaceAlias RDMA1 -DHCP Disabled
Remove-NetIPAddress -InterfaceAlias RDMA1 -Confirm:$false
New-NetIPAddress -InterfaceAlias RDMA1 -IPAddress 192.168.101.10 -PrefixLength 24
Set-DnsClientServerAddress -InterfaceAlias RDMA1 -ServerAddresses 192.168.100.10
Get-NetAdapter -InterfaceDescription "*IPoIB*" | ? {$_.Name -ne "RDMA1"} | Rename-
NetAdapter -NewName RDMA2

#
# Disable RDMA2, since this system only uses one RDMA interface
#
Disable-NetAdapter -InterfaceAlias RDMA2 -Confirm:$false
#
# Configure Storage Spaces, create pool with 3 disks, single simple space
#
$s = Get-StorageSubSystem -FriendlyName *Spaces*
New-StoragePool -FriendlyName Pool1 -StorageSubSystemFriendlyName $s.FriendlyName -
PhysicalDisks (Get-PhysicalDisk -CanPool $true)
Set-ResiliencySetting -Name Simple -NumberofColumnsDefault 3 -StoragePool (Get-
StoragePool -FriendlyName Pool1)
#
# Create Space (virtual disk)
#
New-VirtualDisk -FriendlyName Space1 -StoragePoolFriendlyName Pool1 -
ResiliencySettingName Simple -UseMaximumSize
#
# Initialize Space, partition, create volume, format as X:
#
$c = Get-VirtualDisk -FriendlyName Space1 | Get-Disk
Set-Disk -Number $c.Number -IsReadOnly 0
Set-Disk -Number $c.Number -IsOffline 0
Initialize-Disk -Number $c.Number -PartitionStyle GPT
New-Partition -DiskNumber $c.Number -DriveLetter X -UseMaximumSize
Initialize-Volume -DriveLetter X -FileSystem NTFS -Confirm:$false
#
# Install iSCSI Software Target
#
```

```
Install-WindowsFeature FS-iSCSITarget-Server

#
# Create iSCSI target for two initiators (configured by IP address) with 5 LUNs (1GB for
witness disks, four 100GB for data disks)
#
New-IscsiServerTarget -TargetName FSTarget -InitiatorID IPAddress:192.168.101.11,
IPAddress:192.168.101.12
New-IscsiVirtualDisk -DevicePath X:\LUN0.VHD -size 1GB
1..4 | % {New-IscsiVirtualDisk -DevicePath X:\LUN$_.VHD -size 100GB}
Add-iSCSIVirtualDiskTargetMapping -TargetName FSTarget -DevicePath X:\LUN0.VHD
1..4 | % {Add-iSCSIVirtualDiskTargetMapping -TargetName FSTarget -DevicePath
X:\LUN$_.VHD}
#
# Install Active Directory
#
Install-WindowsFeature AD-Domain-Services
#
# Create AD forest, reboots at the end
#
Install-ADDSForest `
-CreateDNSDelegation:$false `
-DatabasePath "C:\Windows\NTDS" `
-DomainMode "Win2008R2" `
-DomainName "FST2.TEST" `
-DomainNetBIOSName "FST2" `
-ForestMode "Win2008R2" `
-InstallDNS:$true `
-LogPath "C:\Windows\NTDS" `
-SafeModeAdministratorPassword (Read-Host -AsSecureString -Prompt "Enter Password") `
-SYSVOLPath "C:\Windows\SYSVOL"
```

Script #2: Configuring FST2-FS1 (File Server 1)

```
#
# Set service power profile
#
POWERCFG.EXE /S SCHEME_MIN

#
# Configure all 4 interfaces (1 DHCP, 3 static)
#

#
# Rename External, no further action required, since this is DHCP
#
```

```
Get-NetAdapter -InterfaceDescription "*Intel*" | Rename-NetAdapter -NewName "External"

#
# Rename Internal, set to manual IP address, configure IP Address, DNS
#
Get-NetAdapter -InterfaceDescription "*Realtek*" | Rename-NetAdapter -NewName "Internal"
Set-NetIPInterface -InterfaceAlias Internal -DHCP Disabled
Remove-NetIPAddress -InterfaceAlias Internal -Confirm:$false
New-NetIPAddress -InterfaceAlias Internal -IPAddress 192.168.100.11 -PrefixLength 24
Set-DnsClientServerAddress -InterfaceAlias Internal -ServerAddresses 192.168.100.10

#
# Rename RDMA1, set to manual IP address, configure IP Address, DNS
#
Get-NetAdapter -InterfaceDescription "*IPoIB*" | Select -Last 1 | Rename-NetAdapter -
NewName RDMA1
Set-NetIPInterface -InterfaceAlias RDMA1 -DHCP Disabled
Remove-NetIPAddress -InterfaceAlias RDMA1 -Confirm:$false
New-NetIPAddress -InterfaceAlias RDMA1 -IPAddress 192.168.101.11 -PrefixLength 24
Set-DnsClientServerAddress -InterfaceAlias RDMA1 -ServerAddresses 192.168.100.10
#
# Rename RDMA2, set to manual IP address, configure IP Address, DNS
#
Get-NetAdapter -InterfaceDescription "*IPoIB*" | ? {$_.Name -ne "RDMA1"} | Rename-
NetAdapter -NewName RDMA2
Set-NetIPInterface -InterfaceAlias RDMA2 -DHCP Disabled
Remove-NetIPAddress -InterfaceAlias RDMA2 -Confirm:$false
New-NetIPAddress -InterfaceAlias RDMA2 -IPAddress 192.168.102.11 -PrefixLength 24
Set-DnsClientServerAddress -InterfaceAlias RDMA2 -ServerAddresses 192.168.100.10

#
# Join Domain, restart the machine
#
Add-Computer -DomainName FST2.TEST -Credential (Get-Credential) -Restart

#
# Install File Server
#
Install-WindowsFeature File-Services, FS-FileServer, Failover-Clustering
Install-WindowsFeature RSAT-Clustering -IncludeAllSubFeature

#
# Start iSCSI Software Initiator
#
Set-Service MSiSCSI -StartupType automatic
Start-Service MSiSCSI
```

```
#
# Configure iSCSI Software Initiator
#

New-iSCSITargetPortal -TargetPortalAddress 192.168.101.10
Get-iSCSITarget | Connect-iSCSITarget
Get-iSCSISession | Register-iSCSISession

#
# Configure the five iSCSI LUNs (initialize, create partition, volume, format as drives
J: to N:
#
1..5 | % {
    $Letter ="JKLMN"[($_-1)]
    Set-Disk -Number $_ -IsReadOnly 0
    Set-Disk -Number $_ -IsOffline 0
    Initialize-Disk -Number $_ -PartitionStyle MBR
    New-Partition -DiskNumber $_ -DriveLetter $Letter -UseMaximumSize
    Initialize-Volume -DriveLetter $Letter -FileSystem NTFS -Confirm:$false
}
```

Script #3: Configuring FST2-FS2 (File Server 2)

```
#
# Set service power profile
#
POWERCFG.EXE /S SCHEME_MIN

#
# Configure all 4 interfaces (1 DHCP, 3 static)
#

#
# Rename External, no further action required, since this is DHCP
#
Get-NetAdapter -InterfaceDescription "*Intel*" | Rename-NetAdapter -NewName "External"

#
# Rename Internal, set to manual IP address, configure IP Address, DNS
#
Get-NetAdapter -InterfaceDescription "*Realtek*" | Rename-NetAdapter -NewName "Internal"
Set-NetIPInterface -InterfaceAlias Internal -DHCP Disabled
Remove-NetIPAddress -InterfaceAlias Internal -Confirm:$false
New-NetIPAddress -InterfaceAlias Internal -IPAddress 192.168.100.12 -PrefixLength 24
Set-DnsClientServerAddress -InterfaceAlias Internal -ServerAddresses 192.168.100.10
```

```
#
# Rename RDMA1, set to manual IP address, configure IP Address, DNS
#
Get-NetAdapter -InterfaceDescription "*IPoIB*" | Select -Last 1 | Rename-NetAdapter -
NewName RDMA1
Set-NetIPInterface -InterfaceAlias RDMA1 -DHCP Disabled
Remove-NetIPAddress -InterfaceAlias RDMA1 -Confirm:$false
New-NetIPAddress -InterfaceAlias RDMA1 -IPAddress 192.168.101.12 -PrefixLength 24
Set-DnsClientServerAddress -InterfaceAlias RDMA1 -ServerAddresses 192.168.100.10

#
# Rename RDMA2, set to manual IP address, configure IP Address, DNS
#
Get-NetAdapter -InterfaceDescription "*IPoIB*" | ? {$_.Name -ne "RDMA1"} | Rename-
NetAdapter -NewName RDMA2
Set-NetIPInterface -InterfaceAlias RDMA2 -DHCP Disabled
Remove-NetIPAddress -InterfaceAlias RDMA2 -Confirm:$false
New-NetIPAddress -InterfaceAlias RDMA2 -IPAddress 192.168.102.12 -PrefixLength 24
Set-DnsClientServerAddress -InterfaceAlias RDMA2 -ServerAddresses 192.168.100.10

#
# Join Domain
#
Add-Computer -DomainName FST2.TEST -Credential (Get-Credential) -Restart

#
# Install File Server
#
Install-WindowsFeature File-Services, FS-FileServer, Failover-Clustering
Install-WindowsFeature RSAT-Clustering -IncludeAllSubFeature

#
# Start iSCSI Software Initiator
#
Set-Service MSiSCSI -StartupType automatic
Start-Service MSiSCSI

#
# Configure iSCSI Software Initiator
#
New-iSCSITargetPortal -TargetPortalAddress 192.168.101.10
Get-iSCSITarget | Connect-iSCSITarget
Get-iSCSISession | Register-iSCSISession

#
```

```
# No need to configure LUNs here. In a cluster, this is done only from one of the nodes.
We did it in FS1.
#
```

Script #4: Configuring FST2-HV1 (Hyper-V host 1)

```
#
# Set service power profile
#
POWERCFG.EXE /S SCHEME_MIN

#
# Configure all 4 interfaces (1 DHCP, 3 static)
#

#
# Rename External, no further action required, since this is DHCP
#
Get-NetAdapter -InterfaceDescription "*82566DM*" | Rename-NetAdapter -NewName "External"

#
# Rename Internal, set to manual IP address, configure IP Address, DNS
#
Get-NetAdapter -InterfaceDescription "*PRO/100*" | Rename-NetAdapter -NewName "Internal"
Set-NetIPInterface -InterfaceAlias Internal -DHCP Disabled
Remove-NetIPAddress -InterfaceAlias Internal -Confirm:$false
New-NetIPAddress -InterfaceAlias Internal -IPAddress 192.168.100.13 -PrefixLength 24
Set-DnsClientServerAddress -InterfaceAlias Internal -ServerAddresses 192.168.100.10
#
# Rename RDMA1, set to manual IP address, configure IP Address, DNS
#
Get-NetAdapter -InterfaceDescription "*IPoIB*" | Select -Last 1 | Rename-NetAdapter -
NewName RDMA1
Set-NetIPInterface -InterfaceAlias RDMA1 -DHCP Disabled
Remove-NetIPAddress -InterfaceAlias RDMA1 -Confirm:$false
New-NetIPAddress -InterfaceAlias RDMA1 -IPAddress 192.168.101.13 -PrefixLength 24
Set-DnsClientServerAddress -InterfaceAlias RDMA1 -ServerAddresses 192.168.100.10

#
# Rename RDMA2, set to manual IP address, configure IP Address, DNS
#
Get-NetAdapter -InterfaceDescription "*IPoIB*" | ? {$_.Name -ne "RDMA1"} | Rename-
NetAdapter -NewName RDMA2
Set-NetIPInterface -InterfaceAlias RDMA2 -DHCP Disabled
Remove-NetIPAddress -InterfaceAlias RDMA2 -Confirm:$false
New-NetIPAddress -InterfaceAlias RDMA2 -IPAddress 192.168.102.13 -PrefixLength 24
```

```
Set-DnsClientServerAddress -InterfaceAlias RDMA2 -ServerAddresses 192.168.100.10

#
# Install Hyper-V
#
Install-WindowsFeature Hyper-V, Hyper-V-Windows PowerShell, Hyper-V-Tools, Failover-
Clustering
Install-WindowsFeature RSAT-Clustering -IncludeAllSubFeature

#
# Join Domain, restart
#
Add-Computer -DomainName FST2.Test -Credential (Get-Credential) -Restart
```

Script #5: Configuring FST2-HV2 (Hyper-V host 2)

```
#
# Set service power profile
#
POWERCFG.EXE /S SCHEME_MIN

#
# Configure all 4 interfaces (1 DHCP, 3 static)
#

#
# Rename External, no further action required, since this is DHCP
#
Get-NetAdapter -InterfaceDescription "*82566DM*" | Rename-NetAdapter -NewName "External"
#
# Rename Internal, set to manual IP address, configure IP Address, DNS
#
Get-NetAdapter -InterfaceDescription "*PRO/100*" | Rename-NetAdapter -NewName "Internal"
Set-NetIPInterface -InterfaceAlias Internal -DHCP Disabled
Remove-NetIPAddress -InterfaceAlias Internal -Confirm:$false
New-NetIPAddress -InterfaceAlias Internal -IPAddress 192.168.100.14 -PrefixLength 24
Set-DnsClientServerAddress -InterfaceAlias Internal -ServerAddresses 192.168.100.10

#
# Rename RDMA1, set to manual IP address, configure IP Address, DNS
#
Get-NetAdapter -InterfaceDescription "*IPoIB*" | Select -Last 1 | Rename-NetAdapter -
NewName RDMA1
Set-NetIPInterface -InterfaceAlias RDMA1 -DHCP Disabled
Remove-NetIPAddress -InterfaceAlias RDMA1 -Confirm:$false
New-NetIPAddress -InterfaceAlias RDMA1 -IPAddress 192.168.102.14 -PrefixLength 24
```

```
Set-DnsClientServerAddress -InterfaceAlias RDMA1 -ServerAddresses 192.168.100.10

#
# Disable RDMA2, since this system only uses one RDMA interface
#
Get-NetAdapter -InterfaceDescription "*IPoIB*" | ? {$_.Name -ne "RDMA1"} | Rename-
NetAdapter -NewName RDMA2
Disable-NetAdapter -InterfaceAlias RDMA2 -Confirm:$false

#
# Install Hyper-V
#
Install-WindowsFeature Hyper-V, Hyper-V-Windows PowerShell, Hyper-V-Tools, Failover-
Clustering
Install-WindowsFeature RSAT-Clustering -IncludeAllSubFeature

#
# Join Domain, restart
#
Add-Computer -DomainName FST2.Test -Credential (Get-Credential) -Restart
```

Script #6: Configuring the Cluster FST2-FSC (run from FST2-FS1)

```
#
# Run Failover Cluster Validation
#
Test-Cluster -Node FST2-FS1, FST2-FS2
#
# Create cluster
#
New-Cluster –Name FST2-FSC -Node FST2-FS1, FST2-FS2

#
# Rename Networks
#
(Get-ClusterNetwork | ? {$_.Address -like "192.168.100.*" }).Name = "Internal"
(Get-ClusterNetwork | ? {$_.Address -like "192.168.101.*" }).Name = "RDMA1"
(Get-ClusterNetwork | ? {$_.Address -like "192.168.102.*" }).Name = "RDMA2"
(Get-ClusterNetwork | ? {$_.Address -like "172.*" }).Name = "External"
#
# Configure Cluster Network Roles (0=Not used, 1=Cluster only, 3=Cluster+Clients)
#
(Get-ClusterNetwork Internal).Role = 3
(Get-ClusterNetwork RDMA1).Role = 3
(Get-ClusterNetwork RDMA2).Role = 3
```

```
(Get-ClusterNetwork External).Role = 1

#
# Rename Witness Disk
#
$w = Get-ClusterResource | ? { $_.OwnerGroup -eq "Cluster Group" -and $_.ResourceType -
eq "Physical Disk"}
$w.Name = "WitnessDisk"
```

Script #7: Configuring the Classic File Server Cluster FST2-FS (run from FST2-FS1)

```
#
# Move all disks to node one, rename Cluster Disks
#
Get-ClusterGroup | Move-ClusterGroup -Node FST2-FS1
(Get-Volume -DriveLetter I | Get-Partition | Get-Disk | Get-ClusterResource).Name =
"FSDisk1"
(Get-Volume -DriveLetter J | Get-Partition | Get-Disk | Get-ClusterResource).Name =
"FSDisk2"

#
# Create a classic file server resource group
#
Add-ClusterFileServerRole -Name FST2-FS -Storage FSDisk1, FSDisk2 -StaticAddress
192.168.100.22/24, 192.168.101.22/24, 192.168.102.22/24
#
# Create Folders
#
Move-ClusterGroup -Name FST2-FS -Node FST2-FS1
md I:\VMS
md J:\VMS

#
# Create File Shares
#
New-SmbShare -Name VMS1 -Path I:\VMS -FullAccess FST2.Test\Administrator,
FST2.Test\FST2-HV1$, FST2.Test\FST2-HV2$
New-SmbShare -Name VMS2 -Path J:\VMS -FullAccess FST2.Test\Administrator,
FST2.Test\FST2-HV1$, FST2.Test\FST2-HV2$
#
# Set NTFS permissions
#
(Get-SmbShare VMS1).PresetPathAcl | Set-Acl
(Get-SmbShare VMS2).PresetPathAcl | Set-Acl
```

Script #8: Configuring the Scale-Out File Server Cluster FST2-SO (run from FST2-FS1)

```
#
# Add two remaining disks to Cluster Shared Volumes
#
Get-ClusterResource | ? OwnerGroup -eq "Available Storage" | Add-ClusterSharedVolume
#
# Create a scale out file server resource group
#
Add-ClusterScaleOutFileServerRole -Name FST2-SO

#
# Create Folders
#
MD C:\ClusterStorage\Volume1\VMS
MD C:\ClusterStorage\Volume2\VMS

#
# Create File Shares
#
New-SmbShare -Name VMS3 -Path C:\ClusterStorage\Volume1\VMS -FullAccess
FST2.Test\Administrator, FST2.Test\FST2-HV1$, FST2.Test\FST2-HV2$
New-SmbShare -Name VMS4 -Path C:\ClusterStorage\Volume2\VMS -FullAccess
FST2.Test\Administrator, FST2.Test\FST2-HV1$, FST2.Test\FST2-HV2$
#
# Set NTFS permissions
#
(Get-SmbShare VMS3).PresetPathAcl | Set-Acl
(Get-SmbShare VMS4).PresetPathAcl | Set-Acl
```

Script #9: Configuring the virtual machines in FST2-HV1

```
#
# Create VM Switch (if doing this remotely, you will need to reconnect)
#
New-VMSwitch -NetAdapterName External -Name External
Get-NetAdapter -InterfaceDescription Hyper* | Rename-NetAdapter -NewName ExternalVirtual

#
# Create VHD files for two VMs
#
New-VHD -Path \\FST2-FS\VMS1\VM1.VHDX -Fixed -SizeBytes 20GB
New-VHD -Path \\FST2-SO\VMS3\VM3.VHDX -Fixed -SizeBytes 20GB

#
```

```
# Create two VMs
#
New-VM -Path \\FST2-FS\VMS1 -Name VM1 -VHDPath \\FST2-FS\VMS1\VM1.VHDX -SwitchName
External -Memory 1GB
New-VM -Path \\FST2-SO\VMS3 -Name VM3 -VHDPath \\FST2-SO\VMS3\VM3.VHDX -SwitchName
External -Memory 1GB
Set-VMDvdDrive -VMName VM1 -Path D:\WindowsServer2012.iso
Set-VMDvdDrive -VMName VM3 -Path D:\WindowsServer2012.iso
Start-VM VM1, VM3
```

Script #10: Configuring the virtual machines in FST2-HV2

```
#
# Create VM Switch (if doing this remotely, you will need to reconnect)
#
New-VMSwitch -NetAdapterName External -Name External
Get-NetAdapter -InterfaceDescription Hyper* | Rename-NetAdapter -NewName ExternalVirtual
#
# Create VHD files for two VMs
#
New-VHD -Path \\FST2-FS\VMS2\VM2.VHDX -Fixed -SizeBytes 20GB
New-VHD -Path \\FST2-SO\VMS4\VM4.VHDX -Fixed -SizeBytes 20GB

#
# Create and start two VMs
#
New-VM -Path \\FST2-FS\VMS2 -Name VM2 -VHDPath \\FST2-FS\VMS2\VM2.VHDX -SwitchName
External -Memory 1GB
New-VM -Path \\FST2-SO\VMS4 -Name VM4 -VHDPath \\FST2-SO\VMS4\VM4.VHDX -SwitchName
External -Memory 1GB
Set-VMDvdDrive -VMName VM2 -Path D:\WindowsServer2012.iso
Set-VMDvdDrive -VMName VM4 -Path D:\WindowsServer2012.iso
Start-VM VM2, VM4
```

Script #11: Creating a Hyper-V Cluster using file share storage

```
#
# on FST2-HV1
#
#
# Create Hyper-V Cluster called FST2-HVC
#
New-Cluster –Name FST2-HVC –Node FST2-HV1, FST2-HV2
#
# on FST2-FS1
#
```

```
#
# Create Folder and File Share for File Share Witness
#
MD C:\ClusterStorage\Volume1\Witness
New-SmbShare -Name Witness -Path C:\ClusterStorage\Volume1\Witness -FullAccess
FST2.Test\Administrator, FST2.Test\FST2-HVC$
(Get-SmbShare Witness).PresetPathAcl | Set-Acl
#
# on FST2-HV1
#
#
# Configure FST2-HVC Cluster with a File Share Witness
#
Set-ClusterQuorum -NodeAndFileShareMajority \\FST2-SO\Witness
#
# Make VMs in FST2-HV1 Highly available
#
Add-VMToCluster VM1
Add-VMToCluster VM3
#
# on FST2-HV2
#
#
# Make VMs in FST2-HV2 Highly available
#
Add-VMToCluster VM2
Add-VMToCluster VM4
```

Script #12: Optional steps to create a nonclustered file share on FST2-FS1

```
#
# on FST2-FS1
#
MD D:\VMS
New-SmbShare -Name VMS5 -Path D:\VMS -FullAccess FST2.Test\Administrator,
FST2.Test\FST2-HV1$, FST2.Test\FST2-HV2$
(Get-SmbShare VMS5).PresetPathAcl | Set-Acl
#
# on FST2-HV1
#
New-VHD -Path \\FST2-FS1\VMS5\VM5.VHDX -Fixed -SizeBytes 20GB
New-VM -Path \\FST2-FS1\VMS5 -Name VM5 -VHDPath \\FST2-SO\VMS3\VM3.VHDX -SwitchName
External -Memory 1GB
Set-VMDvdDrive -VMName VM5 -Path D:\WindowsServer2012.iso
Start-VM VM5
```

Conclusion

As you see, we speak more Windows PowerShell than English around here! I hope you enjoy the scripting samples and try at least some of it in your configurations.

To conclude, here is a set of SMB tips and tricks I have collected over time:

- Switch to the High Performance power profile: http://blogs.technet.com/b/josebda/archive/2012/11/10/windows-server-2012-file-server-tip-switch-to-the-high-performance-power-profile.aspx
- Make sure your network interfaces are RSS capable: http://blogs.technet.com/b/josebda/archive/2012/11/10/windows-server-2012-file-server-tip-make-sure-your-network-interfaces-are-rss-capable.aspx
- Use multiple subnets when deploying SMB Multichannel in a cluster: http://blogs.technet.com/b/josebda/archive/2012/11/12/windows-server-2012-file-server-tip-use-multiple-subnets-when-deploying-smb-multichannel-in-a-cluster.aspx
- Disable 8.3 Naming (and strip those short names too): http://blogs.technet.com/b/josebda/archive/2012/11/13/windows-server-2012-file-server-tip-disable-8-3-naming-and-strip-those-short-names-too.aspx
- Continuous Availability does not work with volumes using 8.3 naming or NTFS compression: http://blogs.technet.com/b/josebda/archive/2012/11/13/windows-server-2012-file-server-tip-continuous-availability-does-not-work-with-volumes-using-8-3-naming-or-ntfs-compression.aspx
- Enable CSV Caching on Scale-Out File Server Clusters: http://blogs.technet.com/b/josebda/archive/2012/11/14/windows-server-2012-file-server-tip-enable-csv-caching-on-scale-out-file-server-clusters.aspx
- Avoid loopback configurations for Hyper-V over SMB: http://blogs.technet.com/b/josebda/archive/2012/11/14/windows-server-2012-file-server-tip-avoid-loopback-configurations-for-hyper-v-over-smb.aspx
- Run the File Services Best Practices Analyzer (BPA): http://blogs.technet.com/b/josebda/archive/2012/11/15/windows-server-2012-file-server-tip-run-the-file-services-best-practices-analyzer-bpa.aspx
- Use Windows PowerShell to find the free space on the volume behind an SMB file share: http://blogs.technet.com/b/josebda/archive/2012/11/19/windows-server-2012-file-server-tip-use-Windows PowerShell-to-find-the-free-space-on-the-volume-behind-an-smb-file-share.aspx
- New per-share SMB client performance counters provide great insight: http://blogs.technet.com/b/josebda/archive/2012/11/19/windows-server-2012-file-server-tip-new-per-share-smb-client-performance-counters-provide-great-insight.aspx

- Minimum version of Mellanox firmware is required for running SMB Direct in Windows Server 2012: http://blogs.technet.com/b/josebda/archive/2013/01/18/minimum-version-of-mellanox-firmware-required-for-running-smb-direct-in-windows-server-2012.aspx
- How much traffic needs to pass between the SMB Client and Server before Multichannel actually starts? http://blogs.technet.com/b/josebda/archive/2013/01/18/how-much-traffic-needs-to-pass-between-the-smb-client-and-server-before-multichannel-actually-starts.aspx

—Jose Barreto, Principal Program Manager, File Server and Clustering Team

Additional resources

Here are a few additional resources concerning this topic:

- Deploy Hyper-V over SMB (TechNet Library) at: http://technet.microsoft.com/en-us/library/jj134187.aspx
- Hyper-V Cmdlets in Windows PowerShell (TechNet Library) at: http://technet.microsoft.com/en-us/library/hh848559.aspx

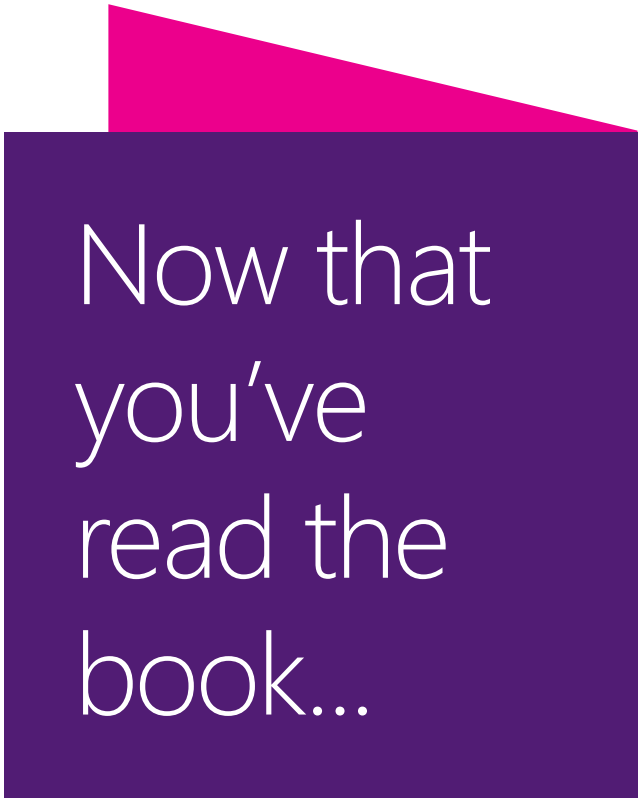

Tell us what you think!

Was it useful?
Did it teach you what you wanted to learn?
Was there room for improvement?

Let us know at http://aka.ms/tellpress

Your feedback goes directly to the staff at Microsoft Press,
and we read every one of your responses. Thanks in advance!

Microsoft