



SUPERCHARGING HYPER-V PERFORMANCE

FOR THE TIME-STRAPPED ADMIN

Brought to you by **Altaro Software**,
developers of [Altaro VM Backup](#)

Compiled and written by **Paul Schnackenburg**

Table of contents

INTRODUCTION	3
PART 1: DIAGNOSING AND REMEDIATING PERFORMANCE ISSUES	4
Tool: Windows Performance Monitor 101	5
Tool: PAL is your friend	7
Finding runaway servers through Resource Metering	10
Finding and remediating Storage Performance issues	11
Finding and remediating CPU performance issues	11
Finding and remediating memory performance issues	12
Finding and remediating network performance issues	13
PART 2: PLANNING IS BETTER THAN REMEDIATION	13
Step 1 Planning Hosts to minimize the risk of performance issues	13
Step 2 Host Patching, Host Drivers	16
Step 3 Planning VMs to minimize the risk of performance issues	17
Step 4 Planning Storage to minimize the risk of performance issues	18
Step 6 Planning Networking to minimize the risk of performance issues	20
Step 7 Planning Management to minimize the risk of performance issues	22
TAKE AWAY ACTIONS	22
ABOUT THE AUTHOR: PAUL SCHNACKENBURG	23
ABOUT ALTARO	23

Introduction

The aim of this eBook is to help time-pressed and resource-strapped Hyper-V administrators make the most of their existing infrastructure as well as give guidance as to what to look for when planning for new hardware. Through practical hands-on tips as well as background information you will learn how Hyper-V (and virtualization in general) affects performance and how to find issues in storage, CPU, memory and network components. This is followed by a look at planning hosts, VMs, storage, networking and management for maximum performance.

You can use this eBook in two ways:

- 1** If you're familiar with performance monitoring and management in Hyper-V you can use the Table of Contents to jump directly to the particular page that covers the issue. And if you're in the planning phase for a Hyper-V deployment you should find some good help in the planning chapter.
- 2** Alternatively, if you are relatively new to the idea of monitoring servers for performance it's better to work through the book from start to finish – this should give you a good, basic grounding in the field of performance monitoring and management.

This book should be useful whether you have a small environment with a handful of Hyper-V hosts, whether they are clustered or not, using only Hyper-V Manager and Failover Cluster Manager for operations. It should also serve you well for larger environments managed through System Center Virtual Machine Manager (VMM) and monitored through System Center Operations Manager (SCOM).

PART 1:

Diagnosing and remediating performance issues

It's easy to think that in today's virtualized, cloud world that performance monitoring and management doesn't matter. One approach I hear frequently from fellow IT Pros is to simply size every server VM the same (1 vCPU, 1 GB of virtual memory is common) and only change those values upwards if users complain. The other tactic that I've heard is giving each VM 4 virtual processors and 4 GB of memory, which in most cases should lead to happy users, but wasted hardware resources.

A more practical approach is using common sense and application workload sizing guides to design your VMs for your workloads and then knowing how to track down performance issues when they do appear. **To do this you need to understand how virtualization, specifically Hyper-V, works.** I've heard it said a few times by engineers at Microsoft that **the job of Hyper-V is to lie to Windows.** If you take a normal Windows Server 2012 R2 (and earlier) and you enable the Hyper-V role, what happens is that the small layer of hypervisor code "slips in" under the installed OS, turning it into a guest VM. This "first VM" is special because it's the parent partition which hosts drivers, negating the need for Hyper-V specific drivers. Normal Windows Server certified drivers work fine for Hyper-V. vSphere, in contrast, doesn't use a parent partition and thus needs to have drivers specifically written for it.

Hyper-V "lies" to the parent partition. As an example, on a very large system, not all CPUs are exposed to the root partition (another name for the parent partition), simply because it doesn't need them. And of course Hyper-V is "lying" to each guest VM about what resources they have access to. This means it's important to use a performance monitoring tool that's not being lied to for tracking down the issue.

Say that you have indications (through user complaints, or monitoring tools) that a particular workload or VM(s) are having performance issues. The first step is to ascertain if the issue is ongoing, if it only happens at certain times, or at random times. It's also important that you know your environment. In the old world, each server had its own OS and application(s) and probably its own storage, making the task of tracking down the area to start looking for a performance issue much easier.

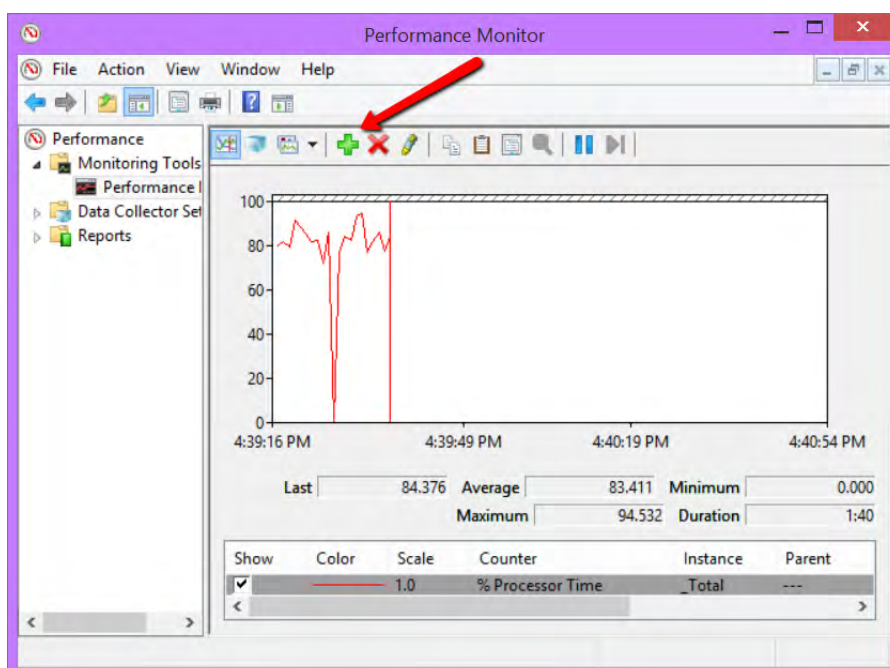
A more practical approach is using common sense and application workload sizing guides to design your VMs for your workloads and then knowing how to track down performance issues when they do appear.

Today most servers are virtualized so the performance issue could be lack of resources assigned to the VM. It could also be a “noisy neighbor” VM on the same host or it could be in the shared storage or somewhere in the complex networking system. So, knowing your environment means understanding how each of your Hyper-V hosts connects to backend storage, how storage is allocated and its performance characteristics, and grasp how the hosts and VMs are connected networking wise. You also need to know the characteristics of your workloads in each VM as well as the processors and memory assignment for the problem VMs and other VM(s) on the same host(s). If you’re taking over an existing environment there’s an [excellent free script](#) which will inventory your stand-alone or clustered hosts and present reports on any issues found.

Task Manager is the first tool most people think of when it comes to performance issues. It is completely useless in a virtual world. In a guest VM it’s being lied to by Hyper-V and if you run it on the host, remember Hyper-V is lying to the parent partition as well. So your first choice should be Performance Monitor – not Task Manager.

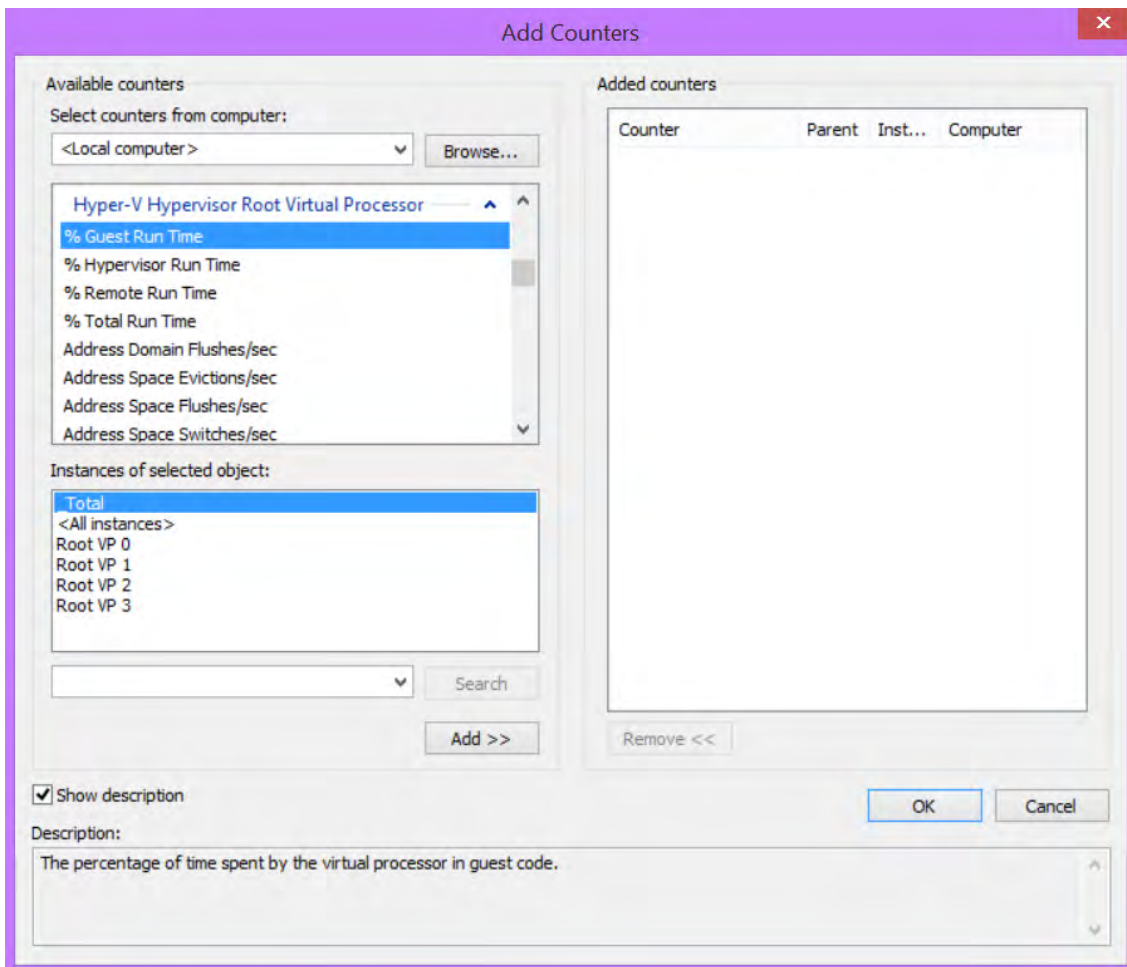
Tool: WINDOWS PERFORMANCE MONITOR 101

This excellent tool has been in Windows since NT 3.1 (for you young ones: that was the hottest OS on the planet just after the dinosaurs died out). And since the first version of Hyper-V, it’s had specific ways of monitoring that are aware of Hyper-V and thus can see beyond the virtualization layer. So if you’re tracking down an issue in a VM you should run Performance Monitor on the Hyper-V host.



Add counters button in Performance Monitor

When you start Performance Monitor (Start – Run – Perfmon), click on the green plus sign. As you can see in the Add Counters dialog box it can run against your local host or against a remote host, in the middle area on the left you pick the object to monitor (Hyper-V Hypervisor Root Virtual Processor in the screenshot). Each object then has multiple counters (% Guest Run Time) and each counter has one or more instances (_Total as well as four Virtual Processors). Once you've picked one or more instances of each counter you need, click **Add** to start monitoring these particular assets.



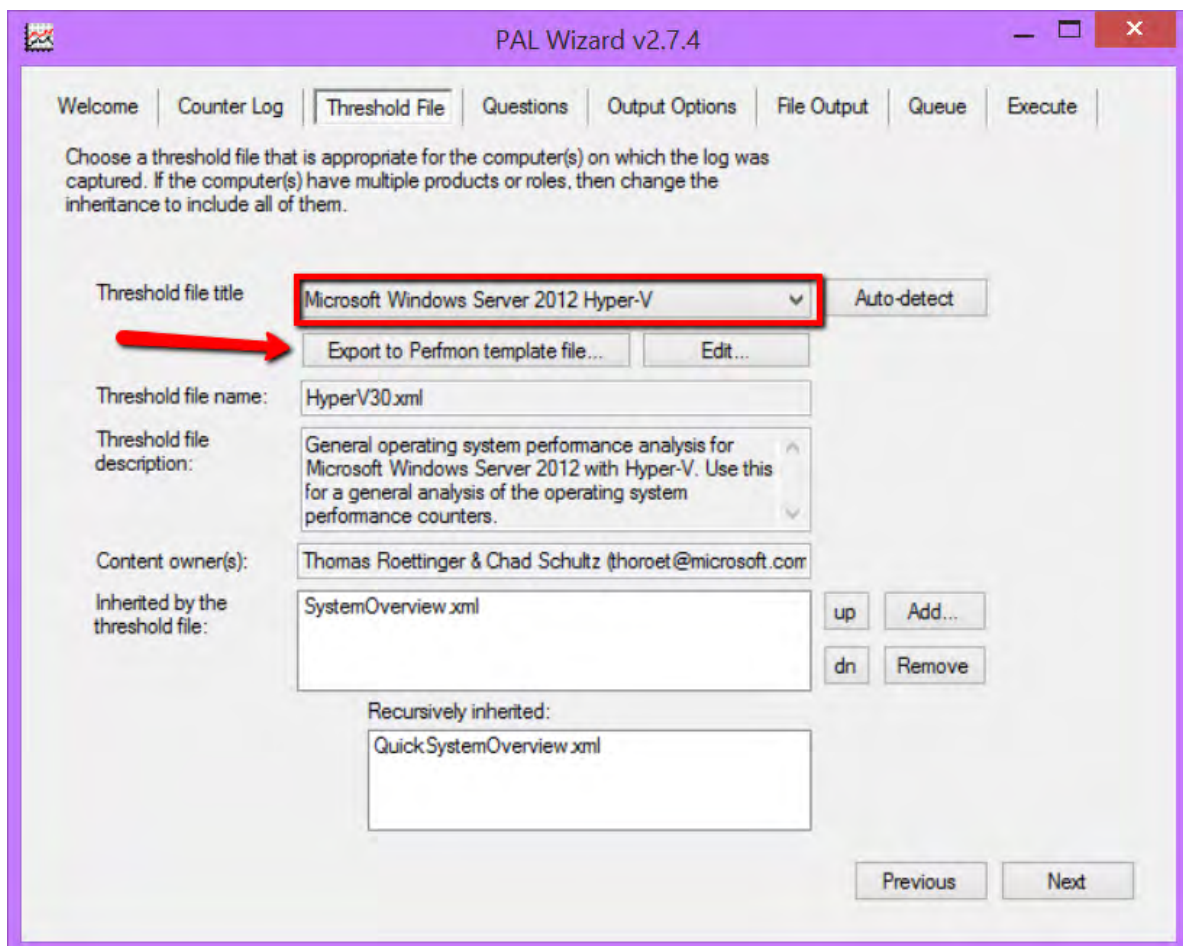
Adding counters in Performance Monitor

This will give you a real time chart of all the counters you have chosen to display, which can be handy if the problem is happening right now. For most problems, you need to monitor for longer periods of time. This is done using Data Collector Sets; these can be set to gather data for hours or even days. To display the collected data you click on the View log data in Performance Monitor. This also allows you to select the time range to show. You can highlight a particular graph line with the highlight button (Ctrl + H).

Tool: PAL IS YOUR FRIEND

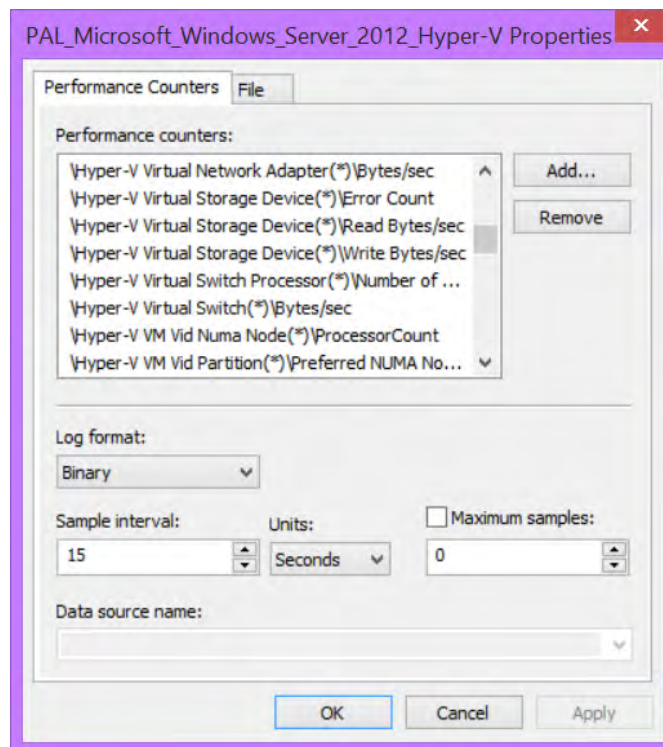
Once you have a feel for Performance Monitor and what it can do, the next step is becoming familiar with [Performance Analysis of Logs \(PAL\)](#), which takes a lot of the guess work out of knowing which counters to use and which values mean there's an issue and which ones are "normal". PAL is free and builds on Performance Monitor (for gathering the logs) but helps tremendously in the analysis phase.

Once you have downloaded and installed PAL itself and the prerequisites (.Net Framework 3.5 and Chart Controls) run PAL. It will show up in All Programs or on your Windows 8 / 10 start screen. A core concept in PAL is the use of threshold files and there are supplied templates for several different Microsoft workloads. In the following screenshot, the Windows Server 2012 Hyper threshold file is loaded:



Choosing a threshold file in PAL

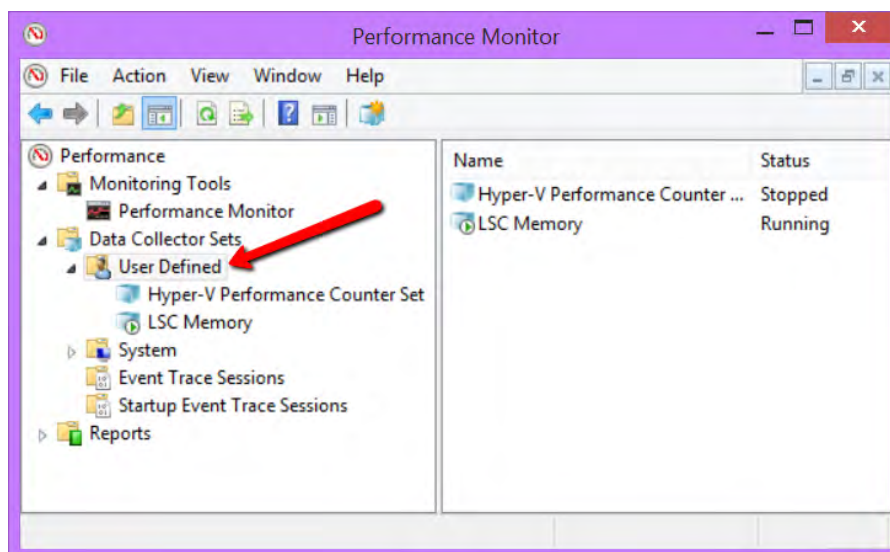
This is a sample of the Performance Counters that are included in the Hyper-V template:



Counters defined in the PAL Hyper-V template

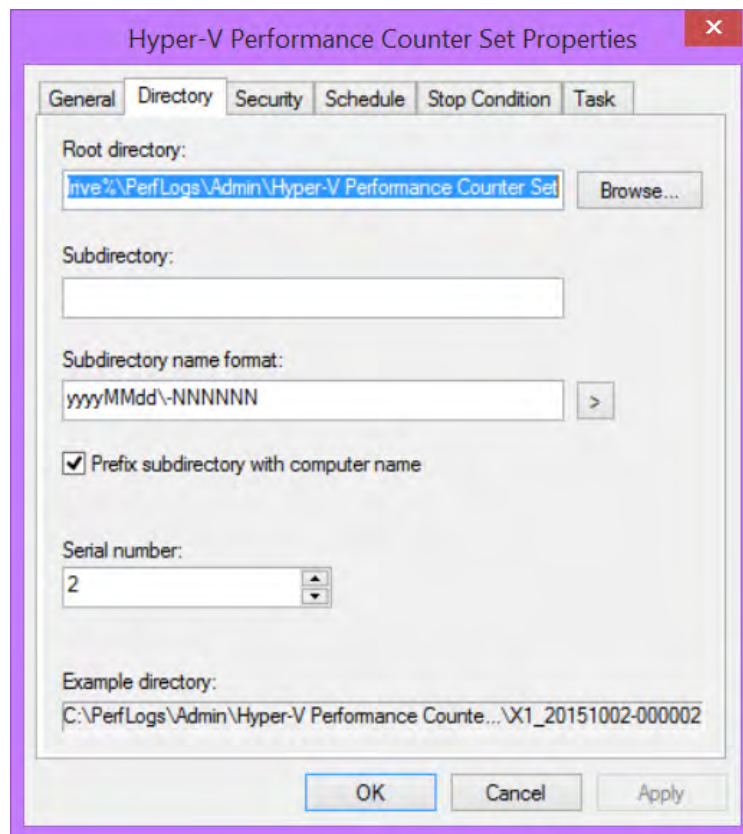
After installation, these are your next steps to start getting some useful information out of PAL:

1) Use the **Export** button to export the threshold file to a Perfmon template file. In Performance Monitor, expand Data Collector Sets – User Defined, right click and select New Data Collector set.



Loading threshold template in Performance Monitor

2) Give the set a name and select Create from a template and click **Next**. On the next page, select **Browse** and pick the threshold file that you saved and click **Finish**. Highlight your new set and click the green **start** button to capture a log during the peak load time on one or more of your systems. By default the file will be saved to C:\Perflogs.



Choosing a folder to save performance log files

Be sure to check that you have ample disk space on your drive – these files can get very big if you capture several hours or days. If you need to save logs on a different drive, right click on your set and go to **Properties** and then to the **Directory** tab. Here you can set another drive to save the log file to.

3) After you have captured the performance data, go back to PAL. On the **Counter Log** tab, load the .blg file created by Performance Monitor. Then, go to the **Questions** tab and enter the relevant information about your host, such as OS and amount of memory. On the **Output Options** tab, select an analysis interval. If you leave it at **Auto**, PAL will divide your analysis into 30 parts so if you captured for 3 hours, each slice will be six minutes long.

The **File Output** tab lets you define where to save files, the **Queue** tab lets you see the PowerShell script that PAL will run, and the **Execute** tab allows you to add more jobs to the queue or start the process. The end result will be an HTML report with green (OK), yellow and red color coding for problem areas.

Alerts

An alert is generated if any of the thresholds were broken during one of the time ranges analyzed. The background of each of the values represents the highest priority threshold that the value broke. See each of the counter's respective analysis section for more details about what the threshold means.

Time Range	Condition	Counter	Min	Avg	Max
2015.08.23 22:37:01 2015.08.23 22:37:16	Greater than 1 error per second on this virtual storage device	\\X1\Hyper-V Virtual Storage Device(C:\Program Files\Microsoft Learning\20411-Drives-20411D-1-ON-SVR1-Virtual Hard Disks\20411D-1-ON-SVR1-Encrypted_7B3D492E-953E-444D-80CA-F46A32CD7A7.avhdx)\Error Count	9	9	9
	Greater than 1 error per second on this virtual storage device	\\X1\Hyper-V Virtual Storage Device(E:\Win2016TP3\HH-Win2016TP3-Virtual Hard Disks\Win2016TP3.vhdx)\Error Count	9	9	9
	Greater than 1 error per second on this virtual storage device	\\X1\Hyper-V Virtual Storage Device(C:\Users\Paul-Downloads-10514.0.150008-1579.T1D_RTFASF_SERVER_ODMRTT_X64\FRT_CIN-US-150)\Error Count	22	22	22
	Value should be below 20,000 per logical processor	\\X1\Hyper-V Hypervisor Logical Processor(_Total)\Context Switches/sec	73,641	73,641	73,641
	Value should be below 20,000 per logical processor	\\X1\Hyper-V Hypervisor Logical Processor(Hv LP 2)\Context Switches/sec	22,108	22,108	22,108
	Value should be below 20,000 per logical processor	\\X1\Hyper-V Hypervisor Logical Processor(Hv LP 0)\Context Switches/sec	21,738	21,738	21,738
	More than 20% privileged (kernel) mode CPU usage	\\X1\Processor(_Total)\% Privileged Time	21	21	21
	More than 20% privileged (kernel) mode CPU usage	\\X1\Processor(3)\% Privileged Time	21	21	21
	More than 20% privileged (kernel) mode CPU usage	\\X1\Processor(1)\% Privileged Time	26	26	26
	More than 20% privileged (kernel) mode CPU usage	\\X1\Processor(0)\% Privileged Time	21	21	21
	Less than 10% free disk space	\\X1\LogicalDisk(Harddisk\Volume5)\% Free Space	7	7	7
	Less than 5% free disk space	\\X1\LogicalDisk(G:)\% Free Space	2	2	2
	Less than 10% free disk space	\\X1\LogicalDisk(C:)\% Free Space	5	5	5
	Greater than 25 ms physical disk READ response times	\\X1\PhysicalDisk(3 E:)\Avg. Disk sec/Read	.054	.054	.054
	Greater than 25 ms physical disk WRITE response times	\\X1\PhysicalDisk(0 C:)\Avg. Disk sec/Write	.032	.032	.032
	Greater than 25 ms logical disk READ response times	\\X1\LogicalDisk(E:)\Avg. Disk sec/Read	.054	.054	.054
	Greater than 25 ms logical disk WRITE response times	\\X1\LogicalDisk(C:)\Avg. Disk sec/Write	.032	.032	.032
	More than 1000 hard page faults per second - this only indicates 4 MB/s or more of disk activity which may or may not be related to usage (disk)	\\X1\Memory\Pages/sec	2,746	2,746	2,746
	Greater than or equal to 54 MB IO/sec. Generally, the larger the IO size, the more data can be transferred per second, but the response times are longer.	\\X1\LogicalDisk(C:)\Avg. Disk Bytes/Write	482,478	482,478	482,478
	Less than 4 GB of free disk space	\\X1\LogicalDisk(Harddisk\Volume6)\Free Megabytes	3,305	3,305	3,305
	Less than 500 MB of free disk space	\\X1\LogicalDisk(Harddisk\Volume5)\Free Megabytes	26	26	26
	Less than 500 MB of free disk space	\\X1\LogicalDisk(G:)\Free Megabytes	87	87	87

An example PAL report

PAL will also help you with baselining, which is a critical part of your performance analysis. When things are going well, users are happy and the fabric is providing a good level of performance; that's the time to create a baseline. If you later encounter performance related issues, gather some logs and compare against the good baseline, this should quickly demonstrate where the problem lies. In other words if you don't know what normal looks like, it's hard to spot abnormal.

Finding runaway servers through Resource Metering

Another helpful addition to Hyper-V in Windows Server 2012 is Resource Metering. It's intended for service providers to track resource usage of each VM (the data follows the VM, even if it is Live Migrated to another host) but can be a life saver if you're trying to track down which VM is using more than its fair share of storage, networking and compute. These are the steps to follow:

1) To enable it open PowerShell as Administrator and run:

```
Get-VM | Enable-VMResourceMetering
```

This will enable metering for all VMs on that host. Alternatively you can define metering for a particular VM with **Get-VM VMName**.

2) Once the data has been collected use:

Get-VM VMName | Measure-VM

This displays the gathered information which covers CPU, memory, disk and network usage. If you suspect that a particular VM on a host (or in a cluster) is a “runaway”, using metering for 10 minutes or so and then displaying the results usually pins down the offender.

```
PS C:\WINDOWS\system32>
PS C:\WINDOWS\system32> get-vm Win2016TP3 | Enable-VMResourceMetering
PS C:\WINDOWS\system32>
PS C:\WINDOWS\system32> Get-VM Win2016TP3 | Measure-VM
VMName      AvgCPU(MHz) AvgRAM(M) MaxRAM(M) MinRAM(M) TotalDisk(M) NetworkInbound(M) NetworkOutbound(M)
-----
Win2016TP3  130         681       1024      512       130048      0              2
PS C:\WINDOWS\system32>
```

Resource Metering

Finding and remediating Storage Performance issues

It's vital that your backend storage can deliver enough performance to each VM, whether that storage is a single host with internal disks, a [Scale out File Server](#) or a SAN. If you suspect that storage is the culprit (and honestly, if you start here you'll be right 90% of the time) there are good Performance Monitor counters to employ.

Use **LogicalDisk Disk Transfers/sec** for the backend storage disk. Note that this will work for local storage as well as block storage that shows up as local drives (iSCSI or FC SAN). If you have virtual disks on SMB file shares you'll need to use Performance Monitor on the physical file server that hosts the shares. If a particular VM is hogging storage IO you should use storage QoS to rein it in (see Planning Storage).

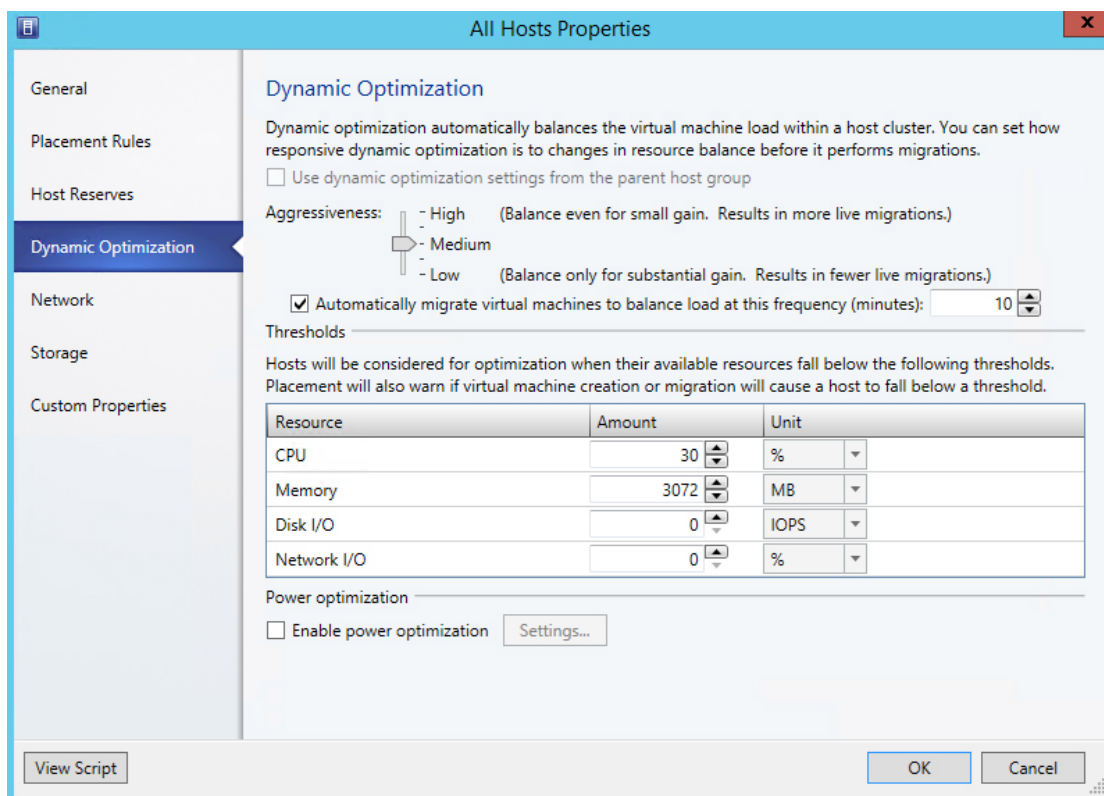
Finding and remediating CPU performance issues

It's rare that processor related performance issues are the root cause of slow VMs. The exception is misbehaving applications; occasionally you will have a hung process or a poorly written program that hammers the virtual CPU(s).

A good counter to use here is **Hyper-V Hypervisor Logical Processor\% Total Run Time** which is the CPU load from all guests and the hypervisor itself for a given host. You can gather this on a per LP basis or as an aggregate across all LPs. Use **Hyper-V Hypervisor Virtual Processor\% Guest Run Time** to track the CPU load on a per VM basis to find the culprit.

Finding and remediating memory performance issues

There are a number of aspects to memory-related performance problems. The power of [dynamic memory](#) (introduced in Hyper-V 2008 R2) takes a lot of the guess work out of establishing how much memory a particular workload requires. Note that some applications, particularly those that do their own memory management, don't play well with dynamic memory. An example is [Exchange Server](#). **Memory\Available Mbytes** is your first stop to see if the host is running out of memory. To see if dynamic memory has to work hard to balance the load amongst the VMs monitor **Hyper-V Dynamic Memory Balancer\Average Pressure**. It should stay under 80 on a healthy host.



Setting Dynamic Optimization in Virtual Machine Manager.png

In a clustered environment that's managed by VMM you should definitely enable [Dynamic Optimization \(DO\)](#). Right click on **All Hosts** in the **Fabric** pane. Select **Properties** and then the **Dynamic Optimization** area. Configure the appropriate frequency and aggressiveness (how often VMM checks to see if VMs should be moved to balance the hosts and how serious it is about making sure all cluster hosts are evenly balanced).

Finding and remediating network performance issues

Networking performance problems are more commonly encountered because of incorrect configuration. There are a lot of options for configuring your host network connectivity. Teamed or single NICs? Dedicated network links for each type of traffic or converged networks using Quality of Service (QoS) or Datacenter Bridging (DCB) for carving up the bandwidth for each flavor? All of these different types of network setups make it difficult to give general performance sleuthing tips.

Nevertheless, aiming for **Network Interface (*)\OutputQueue Length** less than 1 is a good rule of thumb. Comparing **Network Interface (*)\Bytes Sent /sec and Received/sec** against the **Network Interface (*)\Current Bandwidth** is another good place to start to see if network links are saturated. Remember that RDMA-enabled NICs won't show up in normal performance monitoring and neither will SR-IOV network links.

PART 2: Planning is better than remediation

The first half of this book has covered several ways to track down performance issues. But it is better to build or expand your environment with the right gear in the first place to minimize performance problems.

Here we're going to look at how to plan your hosts, VMs, storage, networking and management to end up with a smoothly running, reliable and highly performant fabric on which your business can depend.

For a comprehensive checklist for performance configuration in Hyper-V, check out [this list](#) on [Altaro's Hyper-V Hub](#).

Step 1 | Planning Hosts to minimize the risk of performance issues

A few words on planning the setup of your hosts before we dive into performance details. You can run Hyper-V on a single host which can fit some scenarios such as very small businesses or branch offices. If you don't have the budget for the shared storage required by clustering an option can be Hyper-V Replica.

This fantastic technology is built into Windows Server 2012+ and lets you replicate running VMs to another Hyper-V host, either across the datacenter or, for true Disaster Recovery purposes, another, distant datacenter. So in a small environment you can still provide decent availability by replicating VMs from one host to another in the same datacenter.

It isn't clustering and failover isn't automatic if the primary host fails for some reason but it doesn't require shared storage.

In larger scenarios you'll want to plan for clustering your Hyper-V hosts for High Availability (HA). For this there needs to be some form of shared storage to house the VMs virtual disks and configuration files. This can be Fibre Channel or iSCSI SAN or SMB file shares (see Planning Storage, below).

To design a new (or expand) a cluster you need to know the number of VMs and their requirements for memory, CPU and storage IO. You also need to take into account the resiliency of the cluster itself. Think of this as "cluster overhead"; in a two node cluster you can really only use 50% of the overall capacity of the hardware for workloads. Any less and you won't be able to keep all your VMs running during patching. If there's an outage (planned or unplanned) the remaining node has to be able to run the VMs from the downed node as well as its original workload.

To manage this you can really only plan to use half of the memory, processing and storage IO on each node. In a three node cluster on the other hand, the load can be spread across the remaining two nodes, resulting in 66% available resources being useable. Four nodes have 25% overhead and so forth. By the time you start reaching eight node clusters you may have a requirement to survive two nodes being down simultaneously, again leading to 25% overhead. The largest cluster size in Windows is 64 nodes. **Given cluster overhead I recommend more nodes rather than a few really powerful ones**, although you need to factor licensing costs into your calculations.

As long as any third party software you use works with Windows Server 2012 R2 that is the recommended version to today. Windows Server 2016, due out next year, brings many new features to Hyper-V but as always, it's best to wait a while after RTM and let others find the teething problems. If you use the Standard SKU you only get two virtual Windows VM licenses, whereas Datacenter gives you unlimited Windows VM licenses on that host. For more information on licensing for virtual environments check out [this free Altaro eBook](#).

The option of using Server Core (or Nano Server in the forthcoming Windows Server 2016) for your Hyper-V hosts is recommended by Microsoft but you really need to inventory your own or your team's skills with PowerShell and command line troubleshooting before selecting that installation option. [This survey](#) by Aidan Finn clearly shows that most administrators choose the full UI option for their Hyper-V hosts. Larger environments with very tight standardization of server hardware might be good candidates for Server Core. Select AMD or Intel processors that support Second Level Address Translation (Intel calls it Extended Page Tables and AMD calls it Rapid Virtualization Indexing).

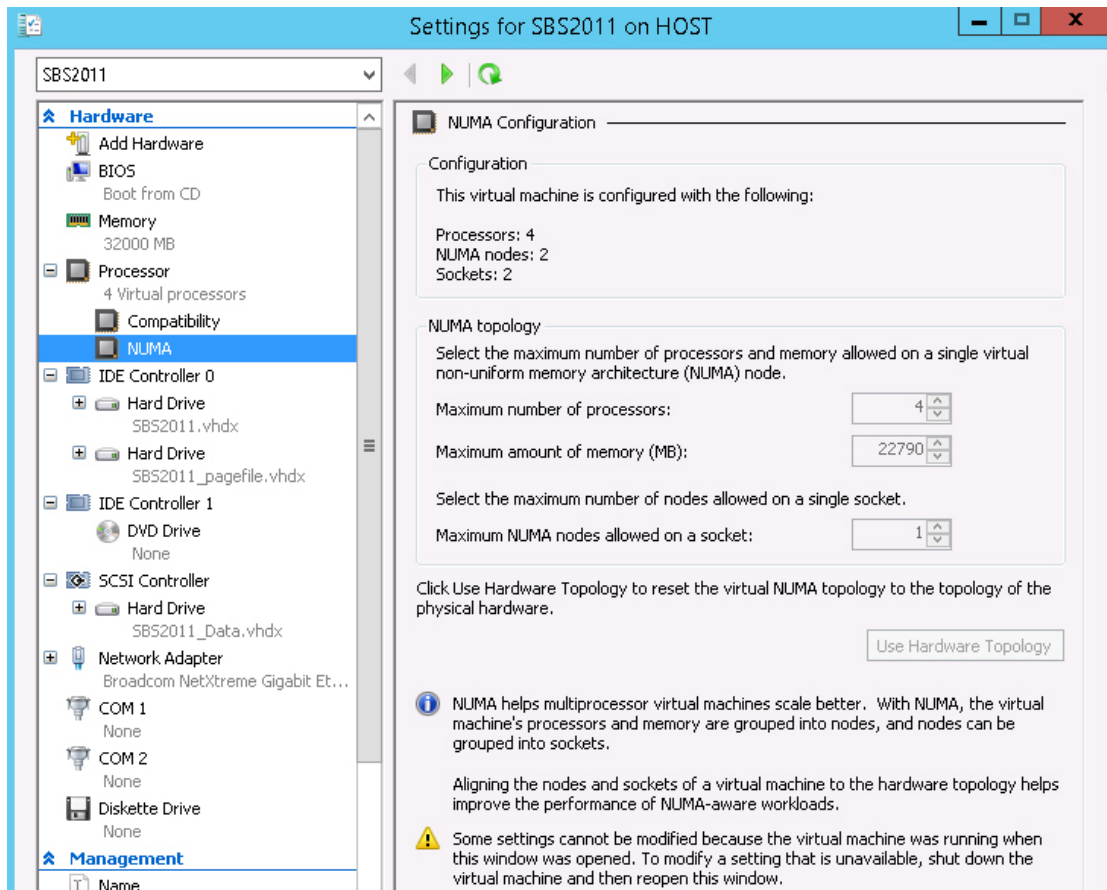
The hosts should only run the Hyper-V role and the only other software that should run on the hosts are management and backup agents. The debate on whether anti-virus software should be run on hosts or only in VMs isn't settled. If you do decide to put it on your hosts please follow [this list](#) to exclude the relevant files, folders and processes.

For a comprehensive checklist with details on how to design a Windows Server 2012 / 2012 R2 Hyper-V environment, or fix one that has been neglected, there's an excellent blog post [here](#). Hosts should have RDP printer mapping disabled.

The single most important tip for hosts is to make sure YOUR DRIVERS ARE UP TO DATE! In many cases simply updating drivers leads to significant performance and stability improvements. That said, it pays to check the forums of your hardware vendor to see if anyone has had issues with a particular driver version.

Another performance related issue is understanding Non Uniform Memory Access (NUMA) for your particular hardware. A standard virtualization host today might come with two physical processors with eight cores each and 128 GB memory. All of those cores do not have the same high speed access to all the memory; the server is instead divided into NUMA nodes. This book's example system has four cores and 32 GB of memory in each node. If you deploy a VM that's larger than a node, the performance of that VM could be adversely impacted. Hyper-V has always had the option to allow a VM to span NUMA nodes and since Window Server 2012, the underlying NUMA topology is projected into VMs so that OSs and applications inside of the VM can configure themselves optimally based on the underlying NUMA topology. One other caveat here is if you have Hyper-V hosts of different memory configuration in the same cluster and the NUMA node sizes don't match between hosts you need to configure your VMs to fit into the smallest NUMA node in the cluster. Otherwise you might get great performance on one host but when the VM is Live Migrated to another host performance might be adversely impacted.

Larger environments with very tight standardization of server hardware might be good candidates for Server Core



NUMA settings for a VM

Step 2 | Host Patching

Keeping your hosts up to date with patches is another important step in performance management. Just turning on Windows Update or relying on an internal Windows Server Update Services (WSUS) or System Center Configuration Manager (SCCM) is not enough. There are several updates specific to Hyper-V that aren't rated as critical so with default settings they'll never be applied to your hosts. Find the list for Windows Server 2008 [here](#), 2008 R2 [here](#), 2012 [here](#) and 2012 R2 [here](#).

Providing a highly performant fabric for your business VMs to run on and providing good uptime for those VMs should never compromise security

Providing a highly performant fabric for your business VMs to run on and providing good uptime for those VMs should never compromise security. Mid July 2015 we learnt of the first “break out of Guest VM” vulnerability in Hyper-V. Read the [bulletin](#) and obtain [the patch](#) if your hosts aren’t up to date. Make sure the location where you keep virtual disks as well as checkpoints are secure with appropriate permissions.

If you have a Hyper-V cluster the easiest way to keep it up to date is using [Cluster Aware Updating \(CAU\)](#). This automates the process of draining VMs from one host and then installing the patches, restarting the host if required and then repeating the process on the next host. If VMM is deployed it has [built in functionality](#) for creating baselines and applying updates across all fabric servers. If you have stand-alone hosts another option is the free [Patch Solution](#) which uses PowerShell to schedule patching.

A word of caution; over the last few years the quality of Windows Server and System Center updates has been declining to the point where **it’s prudent to wait for a few weeks for most updates** (critical security updates not included) to see if other organizations have issues. Depending on when you read this, hopefully this situation has improved.

Host Drivers

It was said above but it bears repeating: it’s very important to update drivers in your hosts for optimal performance. This is a step that often is overlooked, after all “it is working now, don’t change anything” is a well-worn mantra. But the fact is that if you check on your server manufacturers support pages you’ll likely find many issues fixed in later updates for drivers and firmware. **Make sure to check software for network cards, RAID controllers, mainboards, HBAs and other devices on your hosts.**

Step 3 | Planning VMs to minimize the risk of performance issues

The most critical point here is ensuring that your Integration Services (sometimes called Integration Components) are up to date. In Windows Server 2012 R2 you do this from the host (Action-Insert Integration Services Setup Disk). In Windows Server 2016 Integration Services will be distributed through Windows Update (or WSUS / Configuration Manager) instead.

Use the latest versions of Windows and Linux in your VMs. Successive generations of OSs are better at running as VMs.

There is no direct relationship between an LP and a VP. Hyper-V doesn’t have the issue around gang scheduling that VMWare has and it’s safe to assign many VPs to a VM. If a VM doesn’t use the CPU resources, other VMs can use them.

When assigning processor resources to guests, you will see the terms “logical processor” (LP) and “virtual processor” (VP). A LP is the operating system’s way of showing a single pipeline for processing which only operates on one thread at a time. On AMD CPUs, and on Intel CPUs with Hyper-threading disabled, each LP is a physical core. With Hyper-threading enabled each core counts as two LPs. Note that the processing on this second LP only provides about 25% more performance because execution of instructions are tightly bound to the primary core. For VMs that are going to be Active Directory Domain Controllers, turn off the Hyper-V Time Synchronization guest service and make sure that the PDC FSMO role in the first domain in your forest is synchronizing its clock with an authoritative NTP server.

Windows Server 2012 R2 Hyper-V introduced [generation 2 VMs](#). These VMs have no emulated hardware and install and boot faster (but do not run faster). If you’re using recent Linux distros or Windows Server 2012 or later as the OS in your VMs, consider using generation 2 VMs. Windows VMs benefit from virtual Secure Boot in generation 2 and in Windows Server 2016 Hyper-V, so can Linux VMs.

Another benefit in 2012 R2 is Enhanced Session mode, giving you remote access to VMs even before they have an OS installed; great copy and paste functionality between host and VMs, as well as USB device redirection from a remote desktop into a VM.

Step 4 | Planning Storage to minimize the risk of performance issues

As stated before, by far the most common mistake in Hyper-V deployments is not providing enough IO performance. The scariest is where capacity is mistaken for IOPS. For instance “each server takes up 50 GB, so that means we should be able to fit in 70 servers on this new 4 TB hard drive”. The single most important thing you need to consider when planning a Hyper-V deployment or expansion is to **account for IOPS required, not capacity**.

Before virtualization, when each OS lived on its own physical server, how did we plan storage? For example – “a couple of SAS disks in a mirror for the OS, better make them 15,000 RPM disks”, and, “For the data - a RAID 5 set with six disks to get the required performance”.

Nothing in that equation has changed when you move to a virtualized server. It still needs a certain minimum number of IOPS, even if it’s now delivered from a virtual disk on a SAN array or a SOFS file share. Also, plan for future growth. Many systems were probably planned with adequate capacity a few years ago but since then more hosts have been hooked up to the same storage with more VMs, resulting in IOPS starvation.

As stated before, by far the most common mistake in Hyper-V deployments is not providing enough IO performance

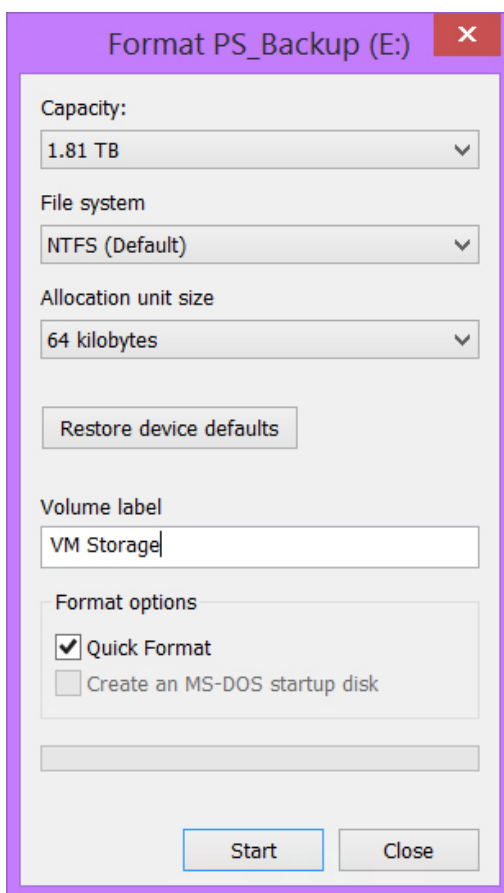
One solution here is SSD or flash based storage. Because many server workloads rely more on reading data than writing it, they're eminently suited to SSD performance. Make sure your storage uses tiering to combine the best performance characteristics of flash memory with the cost effectiveness of HDD storage.

Profiling the IOPS requirements of your workloads is an important part of planning for new deployments. Use Performance Monitor to profile them to see what IOPS requirements they have. You can also use [this script](#) which tests your storage to give you performance information. It relies on the slightly older SQLIO tool; there's also the newer [DiskSPD](#) which can be used to give your storage a good workout. For a bit more real world testing you can use [PDT](#) which is a set of scripts to create new VMs for the entire System Center suite, as well as install all components. It'll create a great deal of load on a single or multiple Hyper-V hosts. The type of shared storage to use for your Hyper-V clusters is also a planning consideration with three main options. If you already have an iSCSI SAN, it'll work well with Hyper-V, as long as it has the needed IO (not just storage capacity) capability for your VMs. Fibre Channel SANs will also work well with Hyper-V, and you can use virtual Fibre Channel SAN to connect VMs directly to shared storage. The final option is Scale out File Server (SOFS) where two, three, or four Windows Server 2012 R2 storage servers connect to backend SAS enclosures.

Hyper-V hosts connect to shared storage using ordinary SMB 3.0 file shares for easy management while the SOFS servers do the "smarts" that a SAN would do, including moving hot blocks to the SSD tier and less frequently accessed blocks to the HDD tier. SOFS can be more cost effective than SANs in many scenarios as well as being easier to manage. Here's a [great spreadsheet](#) to help you plan a SOFS implementation. With Windows Server 2016 we will get another option in the form of [Storage Spaces Direct \(S2D\)](#) where internal storage in each host is pooled together and shared.

Starting in Windows Server 2012 R2 you can assign a minimum and maximum storage IOPS (counted in 8KB increments) to a virtual disk to ensure that a single, hungry VM doesn't starve other VMs of storage performance. Note that the maximum will always be enforced, but if the SAN / SOFS can't deliver the minimum storage IO you have specified to all VMs, there's nothing the hosts can do about it. In Windows Server 2016 there will be a centralized storage controller where policies can be set across all hosts and groups of VMs to control storage IO.

Format your disks with the recommended 64 KB allocation unit size.



Changing to 64 KB allocation units

Step 5 | Planning Networking to minimize the risk of performance issues

Not so long ago your aim for networking was to have as many 1 Gbps network connections in each host as you could, and then dedicate them to different types of traffic such as Live Migration, backup, storage, client connections, cluster heartbeat etc.

Today, most typical Hyper-V hosts for medium to large enterprises come with two (or more) 10 Gbps interfaces. This often calls for a converged infrastructure where software QoS divides up the bandwidth for different traffic types.

Since Windows Server 2012, a consideration in larger implementations is Remote Direct Memory Access (RDMA) NICs for high speed networking. It comes in three flavors, iWarp, Infiniband and RoCE. RDMA gives you high speed (40, 56 and 100 Gbps) networking, with zero CPU overhead. If your shared storage for your cluster is SOFS (see Planning for Storage) it's recommended to use RDMA to connect the Hyper-V hosts to the SOFS hosts when possible. In Windows Server 2012 R2 you can also use RDMA for Live Migration traffic (see below) but these need to be on two separate networks, increasing cost. Windows Server 2016 allows you to merge both types of traffic on the same physical RDMA NICs.

For Live Migration traffic in particular Windows Server 2012 R2 and 2016 offer more choices with the default being compression. In this mode Hyper-V trades some CPU cycles (and most Hyper-V hosts that I've seen have lots of spare processor capacity) for decreased network load by compressing the data stream on the source server and then decompressing on the destination server. Compression is also pretty smart and it continually monitors the hosts, if CPU load increases, **compression will back off. Compression can lead to 50 % reduction in times for Live Migrations** (your mileage may vary). Another option is RDMA / SMB Direct if your network cards offers the capability, this will make your Live Migrations lightning fast.

For applications in your VMs where clients need to have very low latency access an option to consider in your planning is Single Root I/O Virtualization (SR-IOV) NICs. As long as your server motherboards and BIOS supports it SR-IOV will project virtual copies of itself into VMs, providing close to bare-metal performance. Using SR-IOV NICs doesn't stop your VM from being Live Migrated, if the destination host doesn't have SR-IOV NICs (or the ones that are there have no more virtual functions to spare) the Hyper-V virtual switch should be able to simply transition over to a standard synthetic adapter with minimal impact.

On top of all this you can use network teaming, either switch dependent or switch independent flavors, to provide increased bandwidth as well as failover in case of link failure. Teaming can also be done both at the host level and inside of individual VMs.

Because there are so many different business situations and requirements it's difficult to provide general network recommendations beyond the above. It is, however, important to [familiarize yourself](#) with the alphabet soup of different [network-enhancing technologies](#) and their different applications.

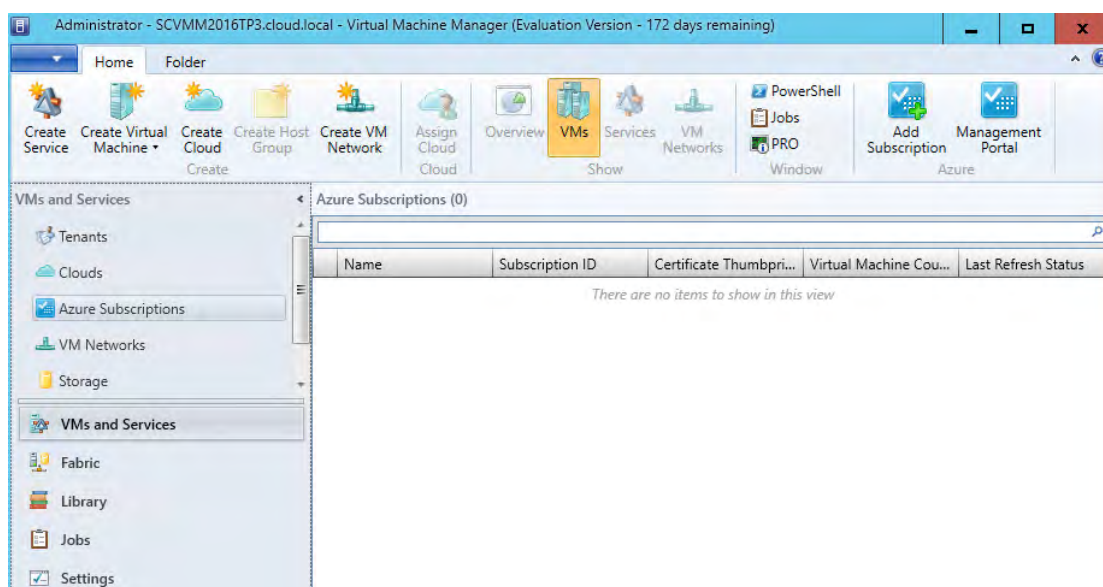
The virtual switches on each Hyper-V host in the cluster need to have the same name and settings for Live Migration to work. In a larger environment you should use the Logical switch in Virtual Machine Manager to push out a centrally controlled switch to all cluster hosts. Centrally defining all your virtual switch settings negates the need to configure the switch on each host individually.

It's also worth mentioning that Hyper-V comes with built in network virtualization, part of Microsoft's overall Software Defined Networking (SDN) story. Unless you are planning a very large Hyper-V infrastructure or you have specific isolation requirements for different parts of your business, it'll likely be overkill. Hyper-V works well with traditional VLAN-based network isolation.

6 | Planning Management to minimize the risk of performance issues

In small environments you can get away with just using Hyper-V Manager and Failover Cluster Manager, along with PowerShell to keep things humming along. For anything larger, System Center Virtual Machine Manager (VMM) is your friend. It does a lot more than just manage VMs; deploys the OS and Hyper-V / SOFS roles to brand new, bare-metal servers and interfaces with both your SANs and SOFS storage for automatic provisioning.

It manages both Hyper-V and VMware environments, manages Top of Rack switches through OMI and can create private clouds which abstracts the underlying fabric and presents quota-based VM resources to application owners. VMM uses templates to deploy individual VMs as well as multi-tier applications with potentially many VMs as services which can be managed as a unit.



Virtual Machine Manager

To keep track of performance, the right System Center family member is Operations Manager (OM). Extensible through Management Packs, OM will collate events and performance data from your host hardware, storage, networking fabric as well as Hyper-V itself. Microsoft offers a Management Pack (MP) for Hyper-V.

Take away action items:

- 1) Monitor the performance of your hosts – use PAL.
- 2) Keep a close eye on the performance of your fabric, storage in particular.
- 3) Use performance data to understand your environment and VMs.
- 4) Plan your fabric environment to make sure you minimize the risk of performance problems.
- 5) Keep your hosts firmware, drivers and OS patches up to date.



About the Author: Paul Schnackenburg

Paul Schnackenburg started in IT when DOS and 286 processors were the cutting edge. He works part time as an IT teacher at a Microsoft IT Academy. He also runs Expert IT Solutions, a small business IT consultancy on the Sunshine Coast, Australia. Paul writes in-depth technical articles, focused on Hyper-V, System Center, private and hybrid cloud and more recently Office 365 and Azure public cloud technologies. He has MCSE, MCSA, MCT, MCTS and MCITP certifications. He can be reached at paul@expertitsolutions.com.au, follow his blog at TellITasITis, <http://tellitasitis.com.au/>.

About Altaro

Altaro Software (www.altaro.com) is a fast growing developer of easy to use backup solutions used by over 30,000 customers to back up and restore both Hyper-V and VMware-based virtual machines, built specifically for SMBs with up to 50 host servers. Altaro take pride in their software and their high level of personal customer service and support, and it shows; Founded in 2009, Altaro already service over 30,000 satisfied customers worldwide and are a Gold Microsoft Partner for Application Development and Technology Alliance VMware Partner.

About Altaro VM Backup

Altaro VM Backup (formerly known as Altaro Hyper-V Backup) is an easy to use backup software solution used by over 30,000 SMB customers to back up and restore both Hyper-V and VMware-based virtual machines. Eliminate hassle and headaches with an easy-to-use interface, straightforward setup and a backup solution that gets the job done every time.

Altaro VM Backup is intuitive, feature-rich and you get outstanding support as part of the package. Demonstrating Altaro's dedication to Hyper-V, they were the first backup provider for Hyper-V to support Windows Server 2012 and 2012 R2 and also continues support Windows Server 2008 R2.

For more information on features and pricing, please visit:
<http://www.altaro.com/vm-backup/>

Don't take our word for it – Take it for a spin!

[DOWNLOAD YOUR FREE COPY OF ALTARO VM BACKUP](#)

and enjoy unlimited functionality for 30 days. After your 30-day trial expires you can continue using the product for up to 2 VMs for free, forever. No catch!

Follow Altaro

Like our eBook? **There's more!**

Subscribe to our Hyper-V blog <http://www.altaro.com/hyper-v/> and receive best practices, tips, free Hyper-V PowerShell scripts and more here: <http://www.altaro.com/hyper-v/sign-up/>

Follow Altaro at:



Share this resource!

Share now on:

