



Automate Management of your Hyper-V Environment

**Next level Integration of System Center
and Veeam Management Pack v8**

Pete Zerger
Microsoft MVP

AVAILABILITY
for the Always-On Enterprise™

Contents

Objectives and introduction	3
Audience	3
Environment	3
Sample Use Cases	5
Easing your way into SMA	5
Overview of Veeam Management Pack Hyper-V support	7
Automation Scenarios	11
Scenario #1: Auto-delete aging checkpoints	11
Process flow for auto-deletion of VM checkpoints	12
Variations in the snapshot auto-deletion scenario	15
Scenario #2: Auto Scale-up.....	15
Process flow for auto scale-up	16
Variations in the scale-up scenario	18
Scenario #3: Adding a new VM to Veeam backup job.....	19
Scenario #4: Controlling VM sprawl.....	23
Other scenarios and capabilities enabled by Veeam Management Pack	25
CMDB enrichment	25
Virtualization self-service	25
Application and service modeling	25
Conclusion	25
About the Author	26
About Veeam	26

Objectives and introduction

The focus of this white paper is to provide practical examples of how you can use Veeam® Management Pack™ (MP) v8 for System Center's management of Microsoft Windows Server 2012 Hyper-V (Hyper-V) to augment the native capabilities of Microsoft System Center. We will start with a quick summary of the components and capabilities included in the System Center stack as well as capabilities that are available within the new Veeam MP v8 release. By understanding the capabilities of both, you can better understand how these solutions can work together to provide additional opportunities for automation and operational efficiencies.

Audience

This white paper is for IT professionals with virtualization management experience as well as those who are familiar with System Center and are interested in how to use Veeam MP to augment their environments. This paper assumes basic familiarity with runbook automation concepts with System Center Orchestrator (SCO) and Microsoft Service Management Automation (SMA).

In this white paper, we will focus on leveraging SMA for new automation development while incorporating SCO in automation scenarios in which it can be helpful. This approach will help organizations move into the future, without leaving behind investments in SCO all at once.

Veeam designed the Hyper-V Veeam MP v8 for System Center to integrate seamlessly with System Center Operations Manager (SCOM) in order to provide comprehensive Hyper-V monitoring from the System Center platform. Additionally, through the interaction available between the components in the System Center stack, System Center can utilize alerts raised by Veeam MP to initiate runbooks in SCO and SMA to automate responses to conditions identified by Veeam MP.

Environment

The sample environment consists of an e-commerce web application with a .NET web application front-end: .NET PetShop. Nodes are all virtualized instances of Windows Server 2012 R2 hosted on Hyper-V. As with many web applications, the front-end application components are load-balanced with either network load balancing (NLB) or hardware load balancing. Windows Server Failover Clustering (WSFC) or SQL AlwaysOn implements high Availability in the SQL data tier. In this scenario, I implemented a 2-node SQL Server 2014 cluster in an active-passive configuration (see Figure 1).




Figure 1 - NET PetShop e-commerce application (healthy)

System Center 2012 R2 includes several complementary software components. The key components of this discussion include:

- System Center 2012 R2 Operations Manager (SCOM): Provides a proactive **monitoring** framework for systems, including monitoring for on-premises, hybrid and cloud environments as well as deep-dive diagnostic information for applications through the application performance monitoring (APM) feature
- Orchestrator (SCO) and Service Management Automation (SMA): Provide process **automation** (also called runbook automation) for actions in a variety of systems, including System Center and Azure
 - SCO is very popular with System Center customers today, but is currently considered a legacy platform in which Microsoft is no longer investing
 - SMA is the latest automation platform from Microsoft and is based on PowerShell Workflow

The sample use cases presented in this whitepaper are intended to provide organizations leveraging SCO for runbook automation an opportunity to try SMA through a few meaningful data center automation scenarios.

Note: To successfully test the sample SMA runbooks in this document, the PowerShell modules for Hyper-V 2012 R2, System Center Virtual Machine Manager 2012 R2 (SCVMM) and SMA, as well as the Hyper-V PowerShell module should be installed on your SCO runbook server and SMA runbook worker.

Sample Use Cases

This white paper will present four use cases in which Veeam MP V8 and Veeam Backup & Replication can be utilized, together with System Center to drive more efficient data center operations in Hyper-V environments. These use cases are described briefly here.

- **Scenario #1: Auto-delete Aging Checkpoints**

VM checkpoints provide an easy restore point to return to a specific point-in-time. They offer quick recovery in application upgrade and patch scenarios in which things do not go as planned. However, checkpoints are detrimental to VM and application performance and should not be left in place any longer than necessary, particularly for production environments. In this scenario, you will learn how to delete aging checkpoints detected by the Veeam MP automatically, before they get out of control.

- **Scenario #2: Scaling up in a Hyper-V VM**

Sometimes, bursts in demand affect tiers of an application that do not scale out easily, such as the database tier. In this scenario, you will learn how to add compute resources within a VM automatically when a resourcing issue causes Veeam MP to generate Hyper-V VM resource utilization alerts SCOM.

- **Scenario #3: Add New VM to Backup Job**

When you deploy new VMs in a production environment, ensuring critical application data is recoverable is very important. In this use case, you will learn how to add new Hyper-V VMs to a backup job in Veeam Backup & Replication automatically, protecting critical data from the start, with no manual configuration required.

- **Scenario #4: Controlling VM Sprawl**

While you could utilize runbook automation to identify idle VMs (sometimes called zombie VMs) running in your environment, sometimes the best solution is even simpler. With the ease of deployment that comes with virtualization also comes the increased likelihood of unnecessary resources left unnoticed beyond their useful life. In this scenario, we will explore how the analysis reports in Veeam MP v8 help administrators easily identify opportunities to reclaim resources and prevent VM sprawl.

Easing your way into SMA

SCO has been around for several years now, and, with its easy drag-and-drop authoring experience, has been widely deployed by organizations using Microsoft System Center. However, SMA has replaced SCO as Microsoft's preferred engine for runbook automation. Leaving SCO behind all at once is difficult, in part because SCO is better than SMA is for some types of automation tasks, such as monitoring SCOM for new alerts. At the same time, SMA is very well suited to PowerShell-based automation tasks.

Note: Azure Automation is Microsoft's cloud-based automation engine. Like SMA, Azure Automation will execute your PowerShell Workflow runbooks. The SMA runbooks demonstrated in this white paper will also work in Azure Automation with the new [Hybrid Worker](#) role, which enables you to execute Azure Automation runbooks from a runbook worker in your data center!

In the use cases presented in this white paper (scenarios #1 and #2 specifically), you will learn how to use SCO with SMA effectively, enabling your organization to move gradually from SCO to SMA, minimizing unnecessary runbook re-authoring and gaining valuable experience along the way.

Because both the first scenario (auto-deleting aging checkpoints) and second scenario (scaling up in a Hyper-V VM) require monitoring SCOM for new alerts that will trigger an automated response, it is most effective to use a single SCO runbook to monitor for SCOM alerts and then trigger the appropriate SMA runbook from SCO to respond to the situation. This approach minimizes calls to SCOM, as well as runbook complexity and authoring effort, while enabling organizations already using SCO to begin integrating SMA into their runbook automation strategy easily.

The sample runbook (shown in Figure 2) is available as a free download. See the **Download the code** references throughout this document for details.

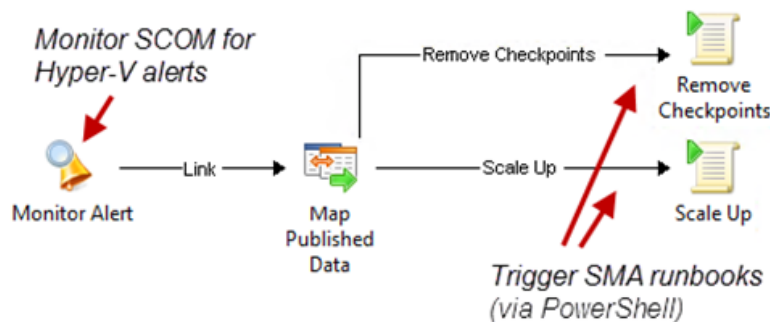


Figure 2 – SCO runbook (to support scenarios #1 and #2)

The **Map Published Data** activity (see Figure 2) allows you to transform the SCOM alert names into values you specify. You can use this activity to convert numeric values to word values, simplify multiple versions of software program names into one name, or perform other string conversion activities within a runbook. In this case, we will use Map Published Data to map alert names from alerts generated by Veeam MP to simple text strings that control the SMA runbook triggered in response to the alert.

Figure 3 shows the configuration details of the Map Published Data. To view the full runbook configuration, simply download the sample runbook and import into a non-production SCO environment.

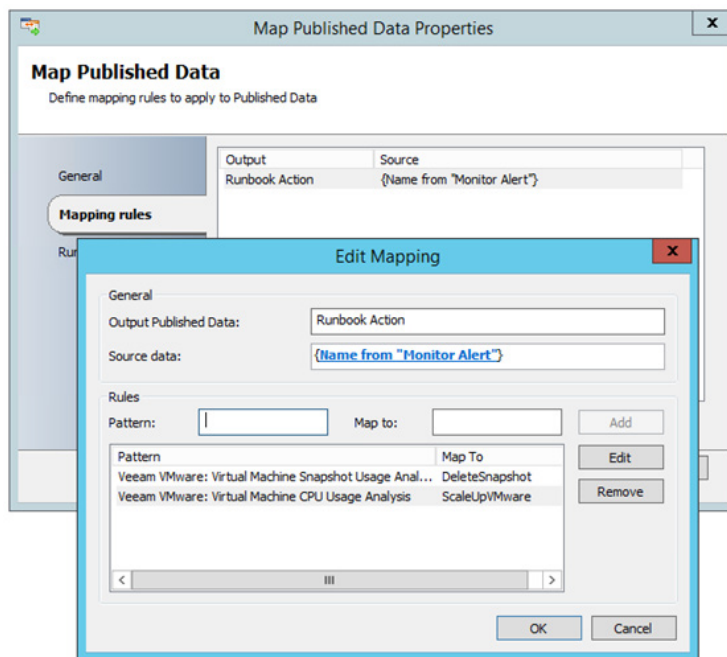


Figure 3 – MAP Published Data activity in SCO runbook (to support scenarios #1 and #2)

Download the Code

Rather than spending a lot of time explaining how to build this SCO runbook step-by-step, it is easier to provide a downloadable sample you can easily import and modify to suit your environment. You can download this SCO runbook from GitHub at <https://github.com/pzgerger/Veeam/tree/master/Automate%20Mgmt%20of%20Hyper-V/SCO%20Runbooks>. Before we delve into the automation scenarios themselves, it will help to examine some of the key capabilities enabled by Veeam MP v8 first.

Overview of Veeam Management Pack Hyper-V support

While Microsoft offers a native management pack for monitoring Hyper-V, it is missing some key capabilities, leaving administrators with "blind spots" when troubleshooting host and guest performance and configuration issues. Veeam MP provides several important features to enable comprehensive monitoring of all supported versions of Hyper-V with SCOM 2012 R2. For the purposes of this white paper, we will focus on four specific Hyper-V items that Veeam MP provides: Enhanced dashboards, smart analysis monitors, advanced reporting and a consistent hypervisor monitoring experience for Hyper-V, very similar to that offered for vSphere in Veeam MP.

Enhanced dashboards:

Veeam MP provides more than 30 views that utilize existing types of views, including the alert view, performance view and state view. In v8, Veeam MP has several prebuilt dashboards that provide an intuitive method to visualize the information that is available for Hyper-V. Veeam MP v8 provides SCOM dashboard views, including Heatmaps, Capacity Forecasting, Top N and Traffic Lights. You can use these views to identify a variety of issues, including large and aging checkpoints that you should delete, VMs and hosts experiencing CPU or memory pressure, and storage errors or network issues. You can see more details about on these new types of dashboards below.

- *Heatmaps:* Heatmaps provide a way to quickly visualize the health of multiple objects by their size and their color. Prebuilt heatmaps are available for Host Compute, VM Storage and Power, VM Checkpoint Finder (see Figure 4).




Figure 4 – Heatmap dashboard for VM Storage and Checkpoints

- *Capacity Forecasting:* The capacity planning widget is used in the Cluster Capacity Forecast dashboard. An enterprise plus license is required to enable the capacity planning features of the Veeam MP.
- *Traffic lights:* The traffic light widget is used in several dashboards, including the top host and top virtual machines. The traffic light dashboards provide a view of top performance counters where you can define a warning and critical level for the counter so they are easy to identify depending on their value. As an example, see the Top Virtual Machines dashboard shown in Figure 5 below.

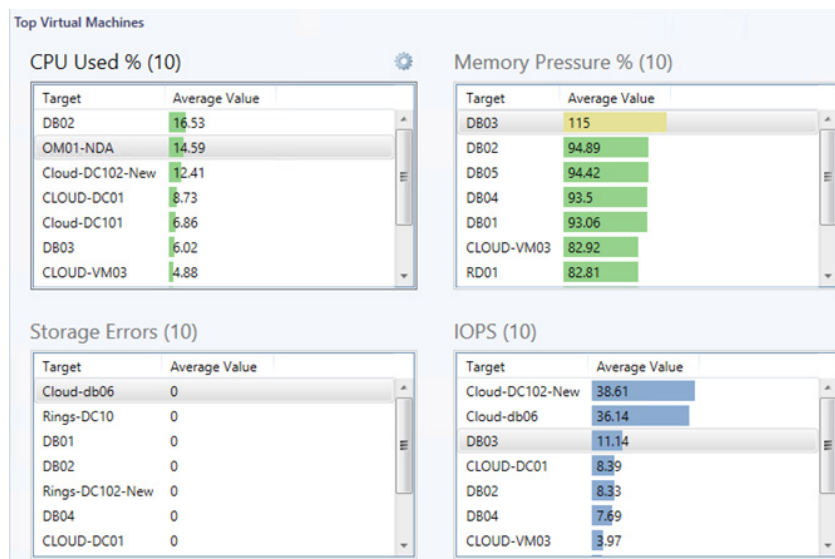


Figure 5 – The traffic light widget in use in the Top Virtual Machines dashboard

The capacity widget, traffic light widget and heatmap widget, all of which Veeam MP uses, are available so you can author your own dashboards using these widgets. You can also combine these widgets with other widgets to provide dashboards for any information available within Operations Manager (including, but not limited, to Hyper-V).

Smart analysis monitors:

Veeam MP gathers information from Hyper-V and adds the information to the Operations Manager framework. This information includes the health of the Hyper-V environment, issues occurring within the environment and performance of components within the environment. Gathering this type of information into the Operations Manager framework provides a single pane of glass, which can effectively show the health, alerts and performance of the hypervisor to provide a more comprehensive view of the environment’s health. Veeam MP also goes beyond gathering performance counters to provide its own counters using smart analysis monitors, such as the VM memory pressure counter.

VM memory pressure calculates the amount of memory the VM requires compared to the amount of memory currently allocated to the VM. The memory pressure is visible in the performance view or dashboards in Operations Manager; for example, see Memory Pressure % shown in in the top right corner of Figure 5.

Advanced reporting:

Veeam MP includes a variety of reports that provide in-depth information about Hyper-V. These reports fall into three major categories:

- *Configuration Tracking and Alert Correlation:* Analyzes virtual infrastructure configuration changes and the number of alerts triggered to see correlation between the configuration changes and triggered alerts

- *Storage Performance Analysis*: Analyzes IOPS and latency statistics for datastores, shows correlation between the IOPS and latency metrics and helps identify datastores with performance issues
- *Virtual Machines*: Reports are available for idle VMs as well as oversized and undersized VMs. Figure 6 illustrates right-sizing, showing VMs that require additional memory and CPU.

▲ Oversized VMs: Memory

RAM analysis	
Total VMs processed:	2
Oversized VMs:	2
VMs that cannot be analyzed (*):	0
Total vRAM that can be reclaimed (**):	1.9

(*) Analysis may fail and VM will be excluded from report if:

- There is missing historical configuration or performance data in the data warehouse for that VM;
- VM has no current configuration.

(**) All recommendations are based on historical performance data. Please verify current VM configuration and workload before applying recommended changes.

Top 2 VMs are shown. VMs are ranked by reclaim value and name.

Virtual Machine	Memory Consumed (GB)		Current vRAM Allocation (GB)	Recommended vRAM Allocation (GB)	Reclaim (GB)	Memory Pressure Avg (%)
	Avg	Max				
SRVR01	1.83	2.00	4.00	3.00	1.00	24.73
SRVR02	1.79	2.00	4.00	3.15	0.85	26.33

▲ Oversized VMs: CPU

CPU analysis	
Total VMs processed:	2
Oversized VMs:	0
VMs that cannot be analyzed (*):	0
Total vCPU resources that can be reclaimed (**):	0

Figure 6 – Right-sizing: VMs undersized for memory and CPU report

Consistent hypervisor monitoring experience:

While this white paper focuses on Veeam MP v8 in Hyper-V environments, you will find that Veeam MP’s monitoring, data visualization and reports for vSphere include much of the same functionality, providing a very consistent monitoring experience for mixed environments in which Hyper-V and vSphere are both present. The addition of Hyper-V support to Veeam MP v8 for System Center provides an experience that is similar to the one available for vSphere servers. These similarities exist in the views that are available, types of performance information gathered and level of depth provided by the diagram views, also called Veeam Topology Views. Figure 7 shows a Hyper-V diagram view on the left and a vSphere diagram view on the right.

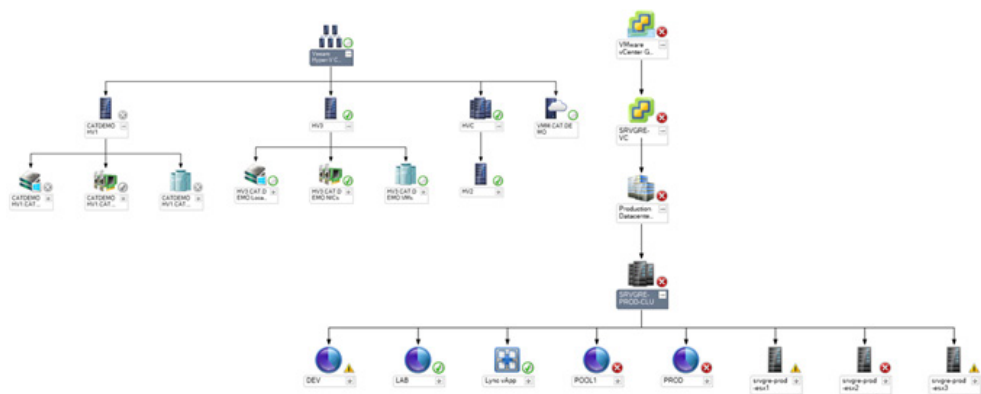


Figure 7 – Hyper-V and vSphere diagram views

Automation Scenarios

The four scenarios in this white paper are just a few examples of efficiencies and automation that you can drive with data from Veeam MP. Veeam MP for System Center enables an additional automation and integration scenario beyond hypervisor and VM guest monitoring for Hyper-V environments. Additionally, the third use case covers the potential for automation in data backup and recovery with Veeam Backup & Replication™.

These scenarios are simply examples of some common use cases suitable for testing and training purposes.

Scenario #1: Auto-delete aging checkpoints

VM checkpoints, designed to provide an easy restore point to return to a specific point in time, provide quick recovery in application upgrade and patch scenarios when things do not go as planned. However, checkpoints are detrimental to VM and application performance and should not be left in place any longer than necessary, particularly for production environments.

The Veeam MP includes monitoring for both checkpoint age and size a monitor named **Veeam Hyper-V: VM Checkpoint Analysis**, which generates an alert when a checkpoint age reaches 2 days or checkpoint size exceeds 2 GB. Because checkpoints generally have no place in a production environment (for extended periods), you could use alerts from this monitor as a trigger to remove forgotten checkpoints on production VMs automatically.

Process flow for auto-deletion of VM checkpoints

Like most automation scenarios, you can adjust the process for auto-deleting checkpoints to suit the specific application scenario. In this example, the auto-deletion process flows as follows (shown in Figure 8):

- An alert is raised in SCOM for aging checkpoints on a VM
- A SCO runbook monitoring for these alerts is triggered to evaluate alerts raised in the last poll of SCOM
- When a checkpoint-related alert is identified, the checkpoint deletion process is initiated (using the Remove Checkpoints branch in the SCO runbook shown in Figure 8)
- SCO triggers a SMA runbook (shown in Figure 9), which — in this example — deletes all checkpoints (parent and all children) on the VM named in the alert

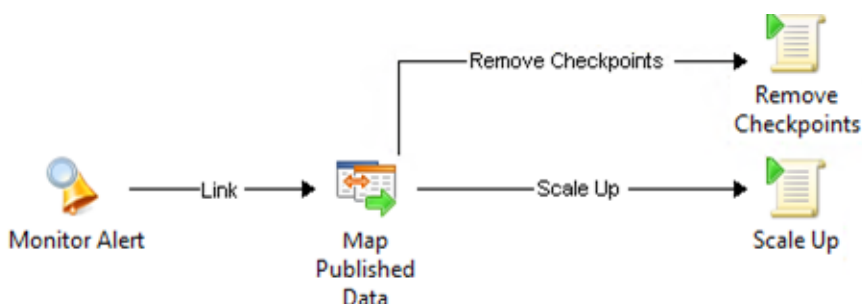


Figure 8 – SCO runbook for auto-deletion of VM checkpoints

The SMA runbook for deleting aging checkpoints on a Hyper-V VM is shown in Figure 9. You can also download it at the link provided in this white paper. If you require guidance on how to import a runbook into SMA, see **Creating or Importing a Runbook** on the Microsoft website at <https://technet.microsoft.com/en-us/library/dn919921.aspx>.

```
workflow Remove-HypervCheckpoints
{
    param(
        $VmmMgmtServer,
        $VMName
    )
    InlineScript
    {
        # load variables for use in the InlineScript block
        $VmmMgmtServer = $Using:VmmMgmtServer
        $VMName = $Using:VMName

        # Import the VMM and Hyper-V modules
        Import-Module VirtualMachineManager
        Import-Module Hyper-V

        # Connect to VMM Server to get the Hyper-V server currently hosting the VM
        Get-SCVMMServer -ComputerName $VmmMgmtServer
        $HypervHost = (Get-SCVirtualMachine -Name $VMName).VMHost.Name

        <#
        Connect to the Hyper-V server and delete all snapshots on the VM.
        We are using the Hyper-V "Remove-VMSnapshot" cmdlet for this step rather
        than the SCVMM "Remove-SCVMCheckpoint" cmdlet because the entire snapshot
        tree can be deleted in one pass using the Hyper-V cmdlet and
        the -IncludeAllChildSnapshots parameter.

        Get the VM properties, select the first (root) snapshot
        (which will be the oldest),
```

```

delete the snapshot and all child snapshots. This has the practical effect
of
deleting ALL snapshots on the VM.
#>
Hyper-V\Get-VM -Name $VMName -ComputerName $HypervHost | `
Hyper-V\Get-VMSnapshot | `
Select-Object -First 1 | Hyper-V\Remove-VMSnapshot -IncludeAllChildSnapshots
-Confirm:$false
}
}

```

Figure 9 – SMA runbook for auto-deletion of checkpoints

A comment on the runbook sample Figure 9:

You may notice Hyper-V\ in front of some of the PowerShell cmdlets. This is necessary because the Hyper-V PowerShell module and the VMware PowerCLI snap-in (loaded in this environment) share some cmdlets with the same names. Pre-pending the cmdlet with the PowerShell module or snap-in (a practice known as pathing) ensures the cmdlet from the correct module or snap-in is always called. This is common in mixed Hyper-V/vSphere environments utilizing runbook automation.

Note: You may also notice this runbook uses both the SCVMM and Hyper-V cmdlets, which was intentional as these two PowerShell modules are quite simply better together. Because the SCVMM cmdlets make host name retrieval quite simple, while the Hyper-V cmdlets provide a means to delete a checkpoint tree (a set of multiple parent and child checkpoints), using the two sets of cmdlets in combination enabled a noticeable reduction in the complexity of the final solution!

Figures 10 and eleven provide before and after screenshots of VM checkpoint properties in Hyper-V Manager to show you how to view the before and after state of your VM to validate success.

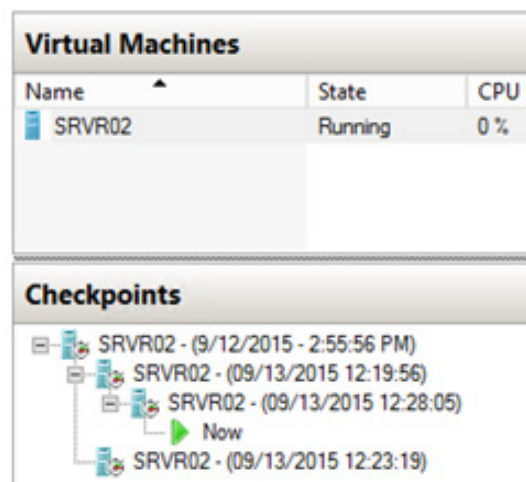
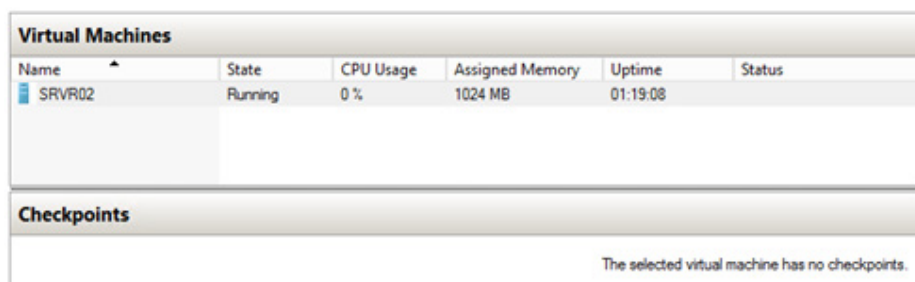


Figure 10 – Checkpoint “Before” in VM properties in Hyper-V Manager



Virtual Machines					
Name	State	CPU Usage	Assigned Memory	Uptime	Status
SRVR02	Running	0 %	1024 MB	01:19:08	

Checkpoints	
The selected virtual machine has no checkpoints.	

Figure 11 – Checkpoint “after” in VM properties in Hyper-V Manager

Download the Code

You can download the SMA runbook for auto-delete aging checkpoints from <https://github.com/pzenger/Veeam/blob/master/Automate%20Mgmt%20of%20VMware/SMA%20Runbooks/Remove-vSphereSnapshots.ps1>.

Variations in the snapshot auto-deletion scenario

The circumstances of the checkpoint scenario may vary based on your application-specific scenarios. You can adjust the scenario above in a number of minor ways to accommodate differing circumstances, including:

- Sending an email to your ITSM tool of choice (if your organization does not use SCSSM) to open an incident or change
- Modifying the SMA runbook to delete only checkpoints older than a specific number of hours or days

The bottom line is that the appropriate level of automation will vary by environment. This example may not be appropriate for all applications in your environment, so feel free to adjust as necessary.

Scenario #2: Auto Scale-up

In other infrastructure components, such as back-end SQL servers or VMs hosting specialized middle-tier business logic, scaling out on-demand may be impractical or not possible from a process perspective. In these scenarios, providing additional resources to VMs in a machine tier may be a suitable alternative. While you could perform this change automatically in some high Availability configurations (as with servers behind physical or virtual load balancers), the complexity of some configurations and external application dependencies may require manual intervention to review the environment more closely before approving the change.




Figure 12 – Resource contention in database tier

Process flow for auto scale-up

Scaling up effectively based on an alert for high CPU or memory utilization can be challenging. In some cases, high utilization of one subsystem can cause a spike in the other. For example, high memory utilization can result in paging to disk, potentially causing a spike in CPU utilization and disk. It is easy to create a single PowerShell script to increase CPU, memory or both. In some cases — such as a stopgap to relieve pressure until an engineer has time to investigate — it may be easiest to scale both. The scale-up process will flow roughly as follows:

- An alert is raised in SCOM for resource pressure (processor and/or memory) for VM in SQL database tier
- A SCO runbook monitoring for these alerts is triggered to begin the scale-up process
- A filter in the runbook ensures the alert(s) raised correspond to one of the monitors approved to trigger the scale-up workflow
- SCO triggers a SMA runbook to shut down the VM
- Then, CPU and memory are increased by the values specified to the runbook (you can specify the number of CPU cores and memory — in GB — to add to the VM)
- The VM is then started when the update is complete

Note: Changing a VM's CPU and memory settings in Hyper-V requires powering down the VM before making the changes in Windows Server 2012 R2. This limitation is resolved in Windows Server 2016.

Figure 13 shows the sample SMA runbook (PowerShell Workflow) to drive this automated process. Rather than trying to copy and paste this runbook into your SMA environment, simply download the copy from the **Download the code** link provided in this white paper.

```
workflow ScaleUp-HypervVM
{
    param(
        $VmmMgmtServer,
        $VMName,
        [int]$RamToAddMB,
        [int]$vCpuToAdd
    )

    InlineScript
    {
        # load variables for use in the InlineScript block
        $VmmMgmtServer = $Using:VmmMgmtServer
        $VMName = $Using:VMName
        $RamToAddMB = $Using:RamToAddMB
        $vCpuToAdd = $Using:vCpuToAdd

        # Import the VMM module
        Import-Module VirtualMachineManager

        # Connect to the VMM Server
        Get-SCVMMServer -ComputerName $VmmMgmtServer

        # Get the properties of the VM to be scaled up
        $VM = Get-SCVirtualMachine -Name $VMName

        # If the VM is running, shut it down
        IF ($VM.Status -ne "PowerOff")
        {
            Stop-SCVirtualMachine -VM $VM
        }
    }
}
```



```
# Calculate the new RAM and CPU counts
# this script assumes the Hyper-V VM is configured with Dynamic Memory
[int]$NewRAMCount = $VM.DynamicMemoryMaximumMB + $RamToAddMB
[int]$NewCpuCount = $VM.CPUCount + $vCpuToAdd

# Update the VM configuration with the new settings
Set-SCVirtualMachine -VM $VM -DynamicMemoryMaximumMB $NewRAMCount `
-CPUCount $NewCpuCount

# Start the VM
Start-SCVirtualMachine -VM $VM
}
}
```

Figure 13 – SMA runbook for VM scale-up

Below, you can see some of the monitors in Veeam MP that can serve to trigger the automation in Hyper-V environments in this scenario. You can extend the sample SCO runbook to monitor for any — or all — of these in only a couple of minutes easily.

- **Veeam HyperV: VM Memory Pressure Analysis:** This monitor tracks threshold breaches for the Hyper-V Dynamic Memory VM \ % Memory Pressure performance counter, which measures average pressure in the VM.
- **Veeam HyperV: VM CPU Scheduling Analysis:** This monitor tracks threshold breaches for the Hyper-V VM Virtual Processor \ % Avg Dispatch Time performance counter, which measures the average time (in percentage) that CPU operations of the VM were ready, but could not get scheduled to run on the physical CPU.
- **Veeam HyperV: VM CPU Usage:** This monitor tracks threshold breaches for the Hyper-V Virtual Processor \ % Total Run Time performance counter, which measures the amount of CPU usage for the guest OS.

Download the code

You can download the SMA runbook for the auto scale-up scenario from <https://github.com/pzenger/Veeam/blob/master/Automate%20Mgmt%20of%20VMware/SMA%20Runbooks/ScaleUp-vSphereVM.ps1>.

Variations in the scale-up scenario

The circumstances of the scale-up scenario may vary based on your application-specific scenarios. You can adjust the scenario above in a number of ways to accommodate differing circumstances or process requirements in an organization, including:

- Enhancing automated scale-up by integrating SCOM maintenance mode before powering down and updating VM resources.
- Automating request for management approval of the auto-scale process through an IT service management platform like System Center Service Manager is achievable with SCO and SMA
- Automating repeating the process (in the reverse direction) when the spike in resource demand subsides to reduce CPU and memory resources assigned to the node

As your comfort level with the solution increases, you may even replace the manual approval process with an automatically approved (pre-approved) process to scale up (and back down) without human intervention to expedite the assignment of compute resources when they are needed most. This provides additional time for administrators to review the situation and make any necessary more permanent adjustments.

Scenario #3: Adding a new VM to Veeam backup job

When you provision new VMs in production environments, it is important that the VM system and application data are recoverable from data backup. Veeam Backup & Replication v8 provides an enterprise solution for effective server and application backup for Hyper-V VMs. We will assume a new VM has just been provisioned in a production environment and needs to be added to an existing backup job in Veeam Backup & Replication.

Unlike the first two scenarios, this is not a monitoring scenario. Because you can call this runbook when needed, the third scenario is a pure SMA automation scenario.

In this example, the process will flow as follows:

- After a new VM is provisioned through any means of automation, this SMA runbook can be called and passed the following information:
 - The VM name
 - Backup job name in Veeam Backup & Replication (Figure 14 highlights the job name used in this sample scenario)

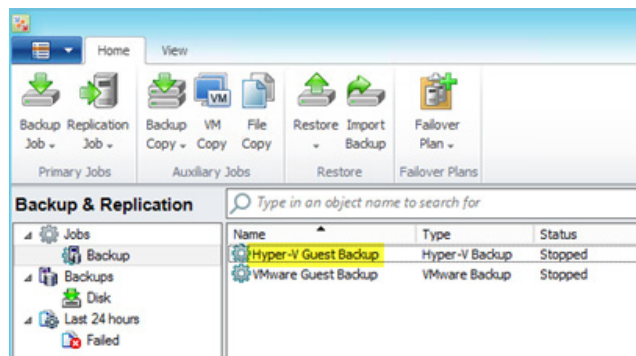


Figure 14 – Backup job name in Veeam Backup & Replication

- The SMA runbook will connect to the Hyper-V server to retrieve the VM details (see the runbook in Figure 17)
- Then, it will connect to the Veeam Backup & Replication Server to retrieve the backup job
- The VM will then be added to the backup job (you can see before and after images of job in Figures 15 and 16)

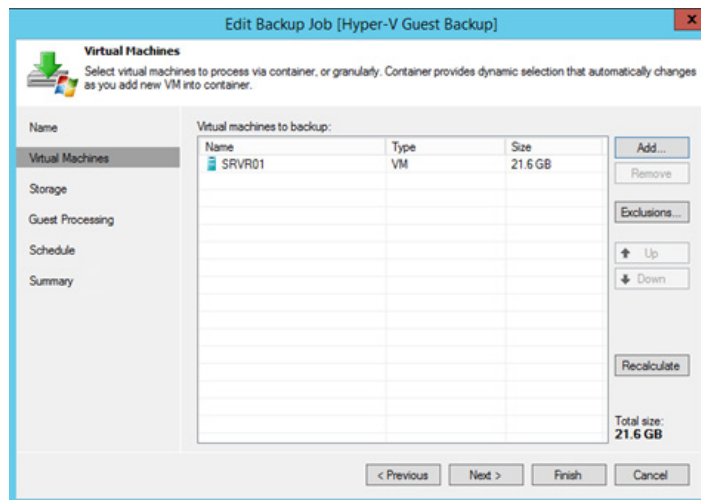


Figure 15 – Backup job “before” in Veeam Backup & Replication

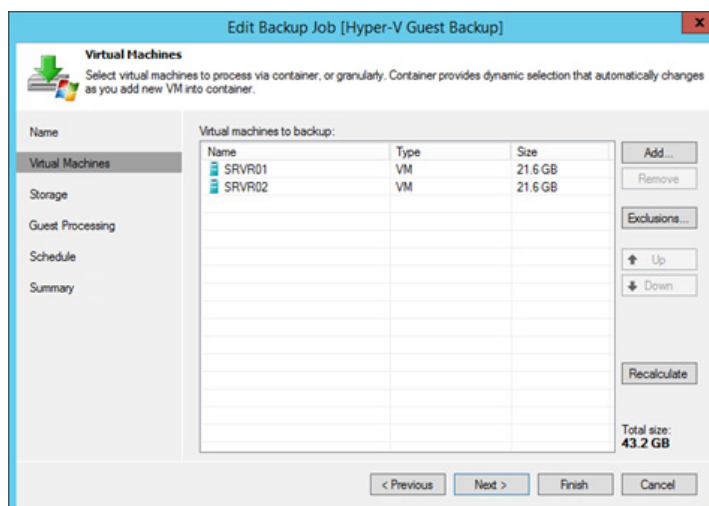


Figure 16 – Backup job “after” in Veeam Backup & Replication

```
workflow Add-HypervVmToVbrBackup
{
    param (
        $PSCredName,
        $VBRMgmtServer,
        $BackupJobName,
        $HvServer,
        $VM
    )

    # Retrieve the previously created SMA/AA Credential object
    "VbrMgmtServerCred"

    # for use when invoking the InlineScript block on the remote VBR Mgmt Server
    $PSUserCred = Get-AutomationPSCredential -Name $PSCredName

    $Result = InlineScript
    {
        # load variables for use in the InlineScript block
        $BackupJobName = $Using:BackupJobName
        $HvServer = $Using:HvServer
        $VM = $Using:VM

        # Loads Veeam Powershell Snapin
        Add-PSSnapin -Name VeeamPSSnapIn -ErrorAction SilentlyContinue

        # Get an existing backup job
        $BackupJob = Get-VBRJob | Where {$_.Name -eq $BackupJobName}

        # Get the VBR managed virtualization host
        $VbrManagedServer = Get-VBRServer | where {$_.Name -eq $HvServer}

        # Find the VM that needs to be added to the backup
        # job on the VBR managed virtualization host.
        $VbrVmToBackup = Find-VBRHvEntity -Server $VbrManagedServer `
        -HostsAndVMs -Name $VM

        # Add the VM to the backup job
        Add-VBRHvJobObject -Job $BackupJob -Entities $VbrVmToBackup
    }
}
```

```
#Check if VM was successfully added to the backup job
$JobCheck = $VM -match (Get-VBRJob -Name $BackupJob | `
Get-VBRJobObject | where {$_.name -eq $VM}).Name

IF ($JobCheck -eq $false)
{
    # Job did not succeed, return error code
    [int]$ReturnCode = 1
}
ELSE
{
    #job succeeded, return success code
    [int]$ReturnCode = 0
}

# Create a custom property to store the ReturnCode value
# for return to the parent runbook.
New-Object PSObject -Property @{
    ReturnCode = $ReturnCode
}

# Invoke the InlineScript block on the remote computer
} -PSComputerName $VBRMgmtServer -psCredential $PSUserCred

# Display the Return Code
$Result.ReturnCode

}
```

Figure 17 – SMA runbook to add a VM to a Veeam Backup & Replication backup job

Download the code

You can download the SMA runbook for adding a VM to a Veeam backup job from <https://github.com/pzenger/Veeam/blob/master/Automate%20Mgmt%20of%20VMware/SMA%20Runbooks/Add-vSphereVmToVbrBackup.ps1>.

Notes: There are a couple of important points to consider in order to implement this scenario successfully:

- The Veeam Backup & Replication PowerShell snap-in is not loaded by default (instructions for installing this snap-in are available in **How to Install the Veeam PowerShell Snap-in** on the Veeam website at <https://www.veeam.com/kb1489>)
- Because the Veeam Backup & Replication PowerShell snap-in requires Veeam Backup & Replication to be installed on the system as well, the SMA runbook for this scenario implements PowerShell remoting to connect to the Veeam Backup & Replication Server from the SMA runbook worker

Backing up on-demand

In the above scenario, the newly added VM will be backed up the next time the job is scheduled to run. If you want to start the backup job immediately to take an initial backup, you can use the Get-VBRJob and Start-VBRJob PowerShell cmdlets as shown in this example:

```
Get-VBRJob -Name "WebApplications Server Backup", "Fileserver Copy Job" | Start-VBRJob
```

This command starts the jobs named **WebApplications Server Backup** and **Fileserver Copy Job**. The job is obtained with Get-VBRJob and passed through the pipeline to Start-VBRJob.

Scenario #4: Controlling VM sprawl

As the ease of deployment automation drives provisioning activities in a private cloud environment, an imbalance in VM counts and resourcing levels can occur as resources are provided, used for a time and then forgotten. To eliminate unnecessary resource consumption, the focus is on automatically identifying likely candidates for de-provisioning, based on resource utilization to prevent or eliminate VM sprawl.

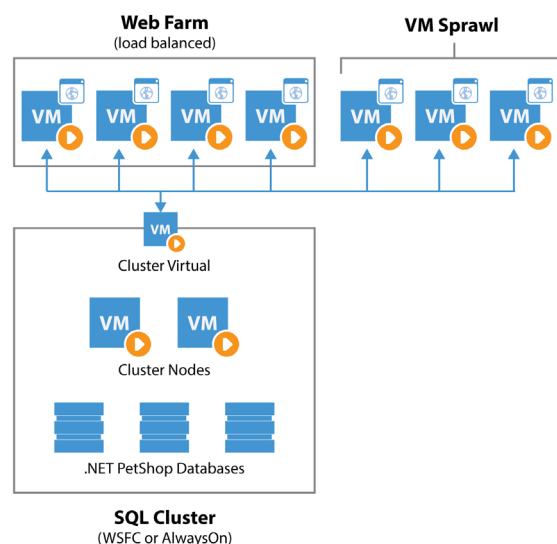


Figure 18 – VM sprawl in N-tier example

While you could use SCO or SMA automation to identify VMs that should be de-provisioned, the lowest effort solution to support this scenario is built right into Veeam MP v8. You can identify potential sprawl with the following Veeam MP v8 Analysis Reports:

- **Virtual Machines. Right-sizing - VMs Oversized for Memory and CPU** – The report (shown in Figure 19) helps you to detect VMs that have more allocated RAM or CPU resources than they require
- **Virtual Machines. Idle VMs** – This report (shown in Figure 20) provides a list of idle VMs in terms of CPU, memory, disk and network utilization

Top 2 VMs are shown. VMs are ranked by reclaim value and name.

Virtual Machine	Memory Consumed (GB)		Current vRAM Allocation (GB)	Recommended vRAM Allocation (GB)	Reclaim (GB)	Memory Pressure Avg (%)
	Avg	Max				
SRVR01	2.00	2.00	4.00	2.05	1.95	17.08
SRVR02	2.00	2.00	4.00	2.70	1.30	29.87

Top 2 VMs are shown. VMs are ranked by reclaim value and name.

Virtual Machine	CPU Usage (Pct)		Current vCPU Allocation (Count)	Recommended vCPU Allocation (Count)	Reclaim (vCPUs)
	Avg	Max			
SRVR01	0.71	7.70	2	1	1
SRVR02	0.64	15.45	2	1	1

Figure 19 – A subsection of the Veeam MP Hyper-V report for right-sizing oversized VMs



Figure 20 – Veeam MP Hyper-V report for idle VMs

Running these reports on a frequent schedule (daily) and delivering them to the Hyper-V and application administrators will enable the teams responsible for cloud and application capacity to make good decisions about which VMs to de-provision in order to reduce spending.

Other scenarios and capabilities enabled by Veeam Management Pack

You can implement or enhance a number of additional by implementing Veeam MP. We'll look at a few of these scenarios along with high-level guidance about implementing in your data center.

CMDB enrichment

By allowing SCSM to import Veeam Hyper-V classes into the CMDB via the SCOM CI Connector, you can populate a richer set of host properties for the Hyper-V servers in the CMDB. You can enable this scenario through a process known as whitelisting. Once these CIs are available within SCSM, you can use them in change control processes and easily identify through views that can be created for these specific CIs.

Virtualization self-service

You can leverage the richer selection of CIs in the CMDB in query-based lists in the SCSM Self-Service Portal to create a more intuitive user experience for IT operations and application development teams deploying VMs through the portal. By leveraging the richer VM classes in Veeam MP v8, you can present VM template options in SCSM's Self-Service Portal to improve the self-service VM management and provisioning experience.

Application and service modeling

You can also leverage the monitored object classes in SCOM added through Veeam MP to diagram richer application models in SCOM. You can then import these models through the SCOM CI Connector in SCSM where they will be added to the CMDB as business services. In fact, by simply configuring the SCOM CI Connector in SCSM, you can populate discovered instances of Hyper-V hosts and host clusters represented as distributed applications in SCOM.

Conclusion

System Center provides a framework for effective monitoring, automation and service desk functionality. The integration of Veeam MP provides many benefits, including the building blocks to automate auto-scale, environment cleanup, VM data protection and a host of other scenarios. By using Veeam MP, you can make intelligent decisions about appropriate automated response in a wide variety of situations.

About the Author



Pete Zerger

Consultant, Author, Speaker and 10-time Microsoft MVP focusing on MS System Center, process automation, virtualization and cloud computing. Pete is a founder of systemcentercentral.com, founding member of mscloudcommunity.com and a frequent presenter at technical conferences around the world, including MS TechEd and System Center Universe. Pete is an author or contributor on several books in the popular Unleashed series from Sams Publishing, as well as a book on Microsoft Azure coming in late 2015.

About Veeam

Veeam® recognizes the new challenges companies across the globe face in enabling the Always-On Business™, a business that must operate 24/7/365. To address this, Veeam has pioneered a new market for Availability for the Always-On Enterprise™. Unlike “legacy backup” solutions, which provide recovery time (RTO) and recovery point (RPO) objectives of hours or days, Veeam helps organizations meet recovery time and point objectives (RTPO™) of less than 15 minutes for all applications and data. This is achieved through a fundamentally new kind of solution that delivers high-speed recovery, data loss avoidance, verified protection, leveraged data and complete visibility. [Veeam Availability Suite™](#), which includes [Veeam Backup & Replication™](#), leverages virtualization, storage and cloud technologies that enable the modern data center to help organizations save time, mitigate risks and dramatically reduce capital and operational costs.

Founded in 2006, Veeam currently has 34,500 ProPartners and more than 168,000 customers worldwide. Veeam’s global headquarters are located in Baar, Switzerland, and the company has offices throughout the world. To learn more, visit <https://www.veeam.com>.


COMING SOON



AVAILABILITY
for the Always-On Enterprise™

NEW Veeam® Availability Suite™ v9

RTPO™ <15 minutes for
ALL applications and data



Learn more and preview
the upcoming v9 release

vee.am/v9