

MASTERING vSPHERE

Best Practices, Optimizing
Configurations & More



ALTARO

CONTENTS

INTRODUCTION	3
Remember: Boring is Good!.....	4
PART 1 - MASTERING ESXI HOST AND VM CONFIGURATION	5
ESXI Host Configuration	5
Automated Deployment - Scripted Installation.....	5
Automated Deployment - Auto Deploy Configuration	7
Setting Up Auto Deploy	9
Virtual Machine Configuration	11
Virtual Machine Template Configuration	12
PART 2 - MASTERING SHARED STORAGE	13
Factors That Affect Performance.....	13
ISCI Best Practices	15
VSAN Best Practices	17
PART 3 - MASTERING NETWORKING	20
Help! My Hosts Won't Communicate With Each Other.....	21
The E1000 Or VMXNET3 Adapter	22
Do I Really Need 10 Gigabits?	22
Scenario Conclusion	24
PART 4 - MASTERING TROUBLESHOOTING AND LOG ANALYSIS BEST PRACTICES	25
Build a Troubleshooting Methodology	25
Know Your Log Files And Locations!	28
Conclusion	30
ABOUT RYAN BIRK	31
ABOUT ALTARO	33

INTRODUCTION

Hi there! If you're here to gather some of the best practices surrounding vSphere, you've come to the right place! In my extensive career as a VMware consultant and teacher (I'm a VMware Certified Instructor) I have worked with people of all competence levels and been asked hundreds - if not thousands - of questions on vSphere. I was approached to write this eBook to put that experience to use to help people currently working with vSphere step up their game and reach that next level. As such, this eBook assumes readers already have a basic understanding of vSphere and will cover the best practices for four key aspects of any vSphere environment.

The best practices covered here will focus largely on management and configuration solutions so should remain relevant for quite some time. However, with that said, things are constantly changing in IT, so I would always recommend obtaining the most up-to-date information from VMware KBs and official documentation especially regarding specific versions of tools and software updates.

This eBook is divided into several sections, outlined below. Although I would advise reading the whole eBook as most elements relate to others, you might want to just focus on a certain area you're having trouble with. If so, jump to the section you want read about here:

- **Part 1: Mastering Host and VM Configuration**
- **Part 2: Mastering Shared Storage**
- **Part 3: Mastering Networking**
- **Part 4: Mastering Troubleshooting and Log Analysis**

REMEMBER: BORING IS GOOD!

Before we begin, I want to note that in a VMware environment, it's always best to try to keep things simple. Far too often I have seen environments be thrown off the tracks by trying to do too much at once. I try to live by the mentality of “keeping your environment boring” – in other words, keeping your host configurations the same, storage configurations the same and network configurations the same. I don't mean duplicate IP addresses, but the hosts need identical port groups, access to the same storage networks, etc.

Consistency is the name of the game and is key to solving unexpected problems down the line. Furthermore, it enables smooth scalability - when you move from a single host configuration to a cluster configuration, having the same configurations will make live migrations and high availability far easier to configure without having to significantly re-work the entire infrastructure.

Now the scene has been set, let's get started!

PART 1 - MASTERING ESXI HOST AND VM CONFIGURATION

ESXI HOST CONFIGURATION

We will start our first section with a few host configuration topics. As mentioned above, your host configuration is very important. Things such as automated deployments, PowerCLI, and Host Profiles, will help prevent human mistakes; make your environments far easier to manage; and keep your environment nice and “boring”. With that said, let’s look at some of what these tools can offer.

AUTOMATED DEPLOYMENT – SCRIPTED INSTALLATION

This can be done with a so-called “kickstart” configuration file, which is supported by VMware. The file contains the configuration for a VMWare ESXi Host to set up settings like IP address, subnet mask, hostname, license key, etc.

The kickstart configuration file can be made available on the following locations:

- FTP
- HTTP/HTTPS
- NFS Share
- USB flash drive
- CD/DVD device

You might be asking yourself: why go through all of this? Automated deployments get more important as your environment grows. For example, if you have a large environment and need to spin up many hosts or maybe you redeploy hosts often and want the same configuration for testing. If you’re currently operating a simpler

environment, it's still well worth mastering automated deployments to maintain consistency down the line.

The HTTP method of deployment is very popular, so we will use that method. We first need to build a configuration file. A Basic script to start with would be called ks.cfg and look something like this:

```
# Accept the VMware End User License Agreement
vmaccepteula

# Set the root password for the DCUI and Tech Support Mode
rootpw mypassword

# The install media is in the CD-ROM drive
install --firstdisk --overwritevmfs

# Set the network to DHCP on the first network adapter
network --bootproto=dhcp --device=vmnic0

# A sample post-install script
%post --interpreter=python --ignorefailure=true
import time
stampFile = open('/finished.stamp', mode='w')
stampFile.write( time.asctime() )
```

Additional configuration options are located out on the VMware Docs [located here](#).

Place the ks.cfg file out on a webserver. Then follow these steps:

Step 1: Boot the ESXi host with a VMware ESXi ISO.

Step 2: Press `shift + o` when the server boots.

Step 3: Enter the following line after runweasel:

- For a HTTP share: `ks=http://%IP_or_FQDN%/kg.cfg`
- For a HTTPS share: `ks=https://%IP_or_FQDN%/kg.cfg`
- For a NFS share: `ks=nfs://%IP_or_FQDN%/ks.cfg`

Step 4: The installation will start and finish the ESXi Host installation based on the ks.cfg configuration file.

Your line will look something similar to this. You'll also notice that you can input your new host IP configuration.

```
<ENTER>: Apply options and boot>  
> runweasel ks=http://192.168.116.1:8080/ks.cfg ip=192.168.116.222 netmask=255.255.255.0 gateway=192.168.116.2| <ESC>: Cancel>
```

The server will boot and install automatically and you now have an automated host deployed!

AUTOMATED DEPLOYMENT – AUTO DEPLOY CONFIGURATION

Before we begin, Auto Deploy requires the following components:

- VMware vCenter Server
 - ♦ The central management server for all servers
- Auto Deploy Service
 - ♦ Serves images and host profiles to the ESXi hosts

- Host and Image Profiles
 - ◆ Defines the set of VIBs (software) to boot the hosts with
- A TFTP Server
 - ◆ Where the Auto Deploy bundle will be stored
- A modifiable/new DHCP scope for the newly deployed hosts
 - ◆ The DHCP scope, will point to the TFTP server above and use the files stored on it

vSphere Auto Deploy lets you provision hundreds of physical hosts with ESXi software. Hosts boot from the network, reading from a central Auto Deploy server which typically will run on your vCenter server. Optionally, hosts are configured with a host profile created from a reference host. The host profile can be configured to prompt the user for input. Once the host is booted up and configuration is complete, the hosts are managed by vCenter Server just like a traditional host would be.

There are two different methods to install your ESXi hosts using Auto Deploy. Stateless caching and stateful installations. Stateless caching is the default option when using Auto Deploy. ESXi configuration and state data is not stored on the host disk, instead a base image profile defines the image that the host is provisioned with, other host attributes and configuration data are applied to the host through host profiles. A host that is set up for stateless caching must connect to the Auto Deploy server (boot media) and the vCenter Server (host profile) every time it boots. Once the ESXi host is running, the installation runs from memory.

In your environment you might want to permanently install ESXi. In that case a stateful install is for you, the host installation image is copied from the Auto Deploy server and stored to a disk that is only accessible to that ESXi host. Subsequent boots are from disk.

Auto Deploy in vSphere 6.5 and 6.7 is easier to implement than in prior versions.

SETTING UP AUTO DEPLOY

On the vCenter Appliance, navigate to this path to start the services.

Home > Administration > System Configuration > Nodes > vCenter Server > Related Objects. Change the Auto Deploy Service and the Image Builder services to Automatic. (Figure 1)

Select the startup type for this service:

☒ Automatic
The service starts automatically when the OS starts.

☐ Manual
You must start the service manually after the OS starts.

☐ Disabled
The service is disabled when the OS starts.

OK Cancel

Figure 1

Then you will need to download the TFTP Boot Zip bundle from vCenter and setup a TFTP server. There are many free options out there. Once you have the TFTP server setup, download the TFTP Boot Zip as shown below. (Figure 2)

Getting Started Summary Monitor **Configu...** Permissions Datacenters Hosts & Clusters VMs

Settings
General
Licensing
Message of the Day
Advanced Settings
Auto Deploy
vCenter HA

Auto Deploy

BIOS DHCP File Name	undionly.kpxe.vmw-hardwired
iPXE Boot URL	https://192.168.1.202:6501/vmw/rbd/tramp
Cache Size	2.00 GiB
Cache Space In-Use	7 MiB

[Download TFTP Boot Zip](#)

Figure 2

On your TFTP server, extract the TFTP bundle you downloaded earlier.

You then will need to setup a new DHCP scope on your DHCP server and set a reservation for the servers that will boot using Auto Deploy. You will need the MAC address of the Management Network adapter to set up the reservation. Setup the 66 and 67 options and point them to the TFTP server and bootfile names. (Figure 3)

Option Name	Vendor	Value
003 Router	Standard	192.168.1.1
006 DNS Servers	Standard	192.168.1.254, 10.107.1.110
015 DNS Domain Name	Standard	lab.local
066 Boot Server Host Name	Standard	192.168.1.2
067 Bootfile Name	Standard	undionly.kpxe.vmw-hardwired

Figure 3

After all the above steps have been carried out, you have now completed the basic setup. Next, we will need to tell our hosts what software to load. This is accomplished by using a Software Depot.

Log back into vCenter and navigate back to the Auto Deploy Section. You will need to download the software depot from the vSphere downloads page and upload it to the software depot. (Figure 4)




Figure 4

Once that is uploaded you will have to select an Image Profile to load. (Figure 5) Followed by creating a new Deploy Rule. A Deploy Rule tells the host to download certain versions of the software, which host profile to use and where the hosts will be put in inventory when configured.

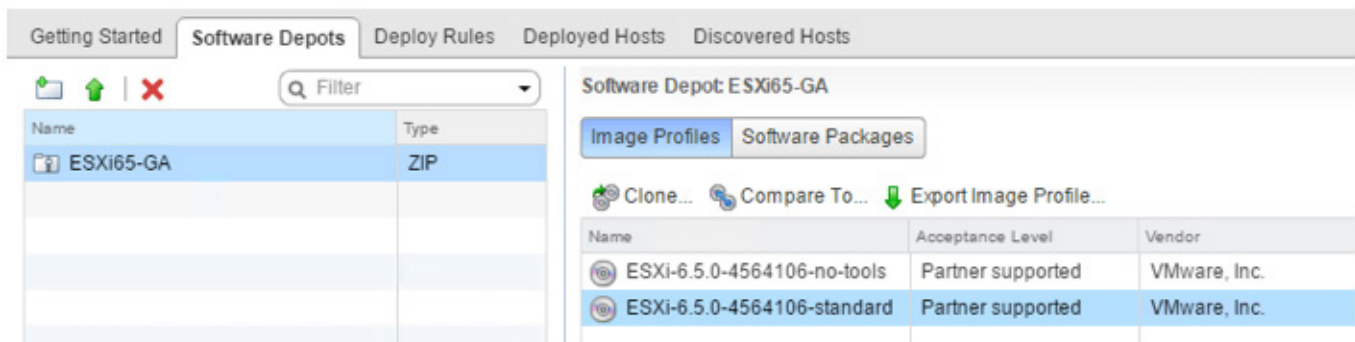


Figure 5

VIRTUAL MACHINE CONFIGURATION

For virtual machine configuration, you should plan your deployment by allocating enough resources for all the virtual machines you will run, as well as those needed by ESXi itself.

Allocate to each virtual machine only as much virtual hardware as that virtual machine requires. Provisioning a virtual machine with more resources than it requires can, in some cases, reduce the performance of that virtual machine as well as other virtual machines sharing the same host.

Disconnect or disable any physical hardware devices that you will not be using. These might include devices such as:

- COM ports
- LPT ports
- USB controllers
- Floppy drives
- Optical drives (that is, CD or DVD drives)
- Network interfaces
- Storage controllers

Disabling hardware devices (typically done in BIOS) can free up resources. Additionally, some devices, such as USB controllers, operate on a polling scheme that consumes extra CPU resources. Some PCI devices also reserve memory, making that memory unavailable to ESXi and other VMs that might need it. If you have no requirements for PCI device passthrough, do not enable it. Between memory reservation overhead and the lack of snapshots, it can become a trouble spot.

Another example and reason to disable CD-ROM devices is that Windows guest operating systems poll optical drives quite frequently. When virtual machines are configured to use a physical drive, and multiple guest operating systems simultaneously try to access that drive, performance could suffer. This can be reduced by configuring the virtual machines to use ISO images instead of physical drives, and can be avoided entirely by disabling optical drives in virtual machines when the devices are not needed.

VIRTUAL MACHINE TEMPLATE CONFIGURATION

If possible, simplify your deployments by using virtual machine templates.

A template is a master copy of a virtual machine that can be used to create and provision virtual machines. Templates cannot be powered on or edited, and are more difficult to alter than an ordinary virtual machine. A template offers a more secure way of preserving a virtual machine configuration that you can easily re-deploy.

When you clone a virtual machine, or deploy a virtual machine from a template, the resulting cloned virtual machine is independent of the original virtual machine or template. Changes to the original virtual machine or template are not reflected in the cloned virtual machine, and changes to the cloned virtual machine are not reflected in the original virtual machine or template.

Some ideas for creating clean templates would be to disable any unused devices. Don't just take the default setting. Install VMware Tools, Remove Floppy Drives, enable CPU/Memory HotAdd if you plan to use that feature, set up the machine to automatically check and upgrade VMware Tools and sync the guest time to the ESXi host to eliminate NTP issues later.

PART 2 - MASTERING SHARED STORAGE

Shared storage is key to successful deployments. Most of the issues that come up with migrations and high availability are due to not having your virtual machines on shared storage.

FACTORS THAT AFFECT PERFORMANCE:

HARD DISKS

- Disk subsystem bottlenecks
- Disk speed

PERFORMANCE VERSUS CAPACITY!

- Disk performance does not scale with drive size
- Larger drives generally equate lower performance

IOPS (I/OS PER SECOND) IS CRUCIAL!

- How many IOPS does a certain number of disks provide?
- How many disks are required to achieve a required number of IOPS?
- More spindles generally result in greater performance

Understanding your workload is a crucial consideration when designing for optimal performance.

When looking at storage, consider using VAAI. VAAI is short for vStorage API for Array

Integration. It allows you to offload certain operations to your storage array. The storage array is far more efficient and much faster than ESXi sending the commands for each operation. It's important to note though, that while VAAI helps, if your storage array is already at maximum performance capacity, VAAI will not allow you to do things you otherwise could not do.

Often other bottlenecks can be something altogether different and not performance related. I often see people use Thin Provisioning thinking that they don't have to manage it and that they won't ever run out of space. That is the wrong mentality as Thin Provisioning adds management complexity to your environment.

System administrators or VMware technicians may be asked if the free version of ESXi supports shared storage. Simply put, yes it does. However, it becomes hard to manage because you don't have a central management server like vCenter and you can only manage your VMs individually, host by host. You also lose the possibility to move VMs from one host to another using vMotion technology.


REQUIREMENTS FOR SHARED STORAGE:

- At least one ESXi host with two physical NICs (management and storage traffic)
 - ◆ It is recommended to keep storage traffic on separate physical network infrastructure. When using High Availability, your storage network is used as a heartbeat network if the management network fails.
- Shared storage offering iSCSI, Fiber Channel or NFS
- Network switch between shared storage and the ESXi host
- vSphere client software

In the following section, we will take a deep dive into the best practices for two types of shared storage: iSCSI and vSAN.

ISCSI BEST PRACTICES

- Place only one VMFS datastore on each LUN. Multiple VMFS datastores on one LUN is not recommended and can lead to issues.
- Do not change the path policy the system sets for you unless you understand [the implications](#).
- **Document everything.** Include information about configuration, access control, storage, switch, server and iSCSI HBA configuration, software and firmware versions.



PLAN FOR FAILURE:

- Look for redundancy. Cross off different links, switches, HBAs and other elements to ensure you did not miss a critical failure point. Whiteboard as much of it as you can!
- Ensure that the iSCSI HBAs are installed in the correct slots in the ESXi host, based on slot and bus speed. Balance PCI bus load among the available busses in the server.


- Become familiar with the various monitoring points in your storage network, at all visibility points, including ESXi performance charts, Ethernet switch statistics, and storage performance statistics.
- Be cautious when changing IDs of the LUNs that have VMFS datastores being used by your host. If you change the ID, virtual machines running on the VMFS datastore will fail.

Place each LUN on a RAID group that provides the necessary performance levels. Utilize tools like vCenter performance graphs and [lometer](#) to verify. Pay attention to the activities and resource utilization of other LUNS in the assigned RAID group. A high-performance RAID group that has too many applications doing I/O to it might not meet performance goals required by an application running on the ESXi host.

Provide each server with enough network adapters or iSCSI hardware adapters to allow maximum throughput for all the applications hosted on the server for the peak period. I/O spread across multiple ports provides higher throughput and less latency for each application. Another consideration is the network bandwidth. Network bandwidth is dependent on the ethernet standards used (1Gb or 10Gb). There are other mechanisms such as port aggregation, jumbo frames and bonding links that deliver greater network bandwidth.

When implementing software iSCSI that uses network interface cards rather than dedicated iSCSI adapters, gigabit Ethernet interfaces are required. These interfaces tend to consume a significant amount of CPU resources. One way of overcoming this demand for CPU resources is to use a feature called a TOE (TCP/IP offload engine). TOEs shift TCP packet processing tasks from the server CPU to specialized TCP processors on the network adaptor or storage device. Most enterprise-level networking chip sets today offer TCP offload or checksum offload, which vastly improve CPU overhead.

VSAN BEST PRACTICES



vSAN is a relatively new tool in the vSphere suite. It first came out in vSphere 5.5 and VMware has actively been improving it since. vSAN is a hyper-converged, software-defined storage (SDS) that creates a pool of storage using (DAS) direct-attached storage devices. It is developed by VMware and used in a VMware vSphere cluster to create a distributed, shared data store. The storage requirement is provided by users through storage policy. You can base policies around performance and availability. It ensures that these policies are properly administered and maintained.

vSAN is a piece of the VMware ESXi kernel and runs on industry-standard x86 servers from OEMs, including Cisco, Dell EMC, Hewlett Packard Enterprise, Fujitsu, Lenovo and Supermicro. Since vSAN doesn't require extra software, users can easily enable this feature with a few clicks.

A neat option surrounding vSAN is the ability to enable an iSCSI target that uses the vSAN backed datastore. See the steps required to enable the service in Figure 6 below.

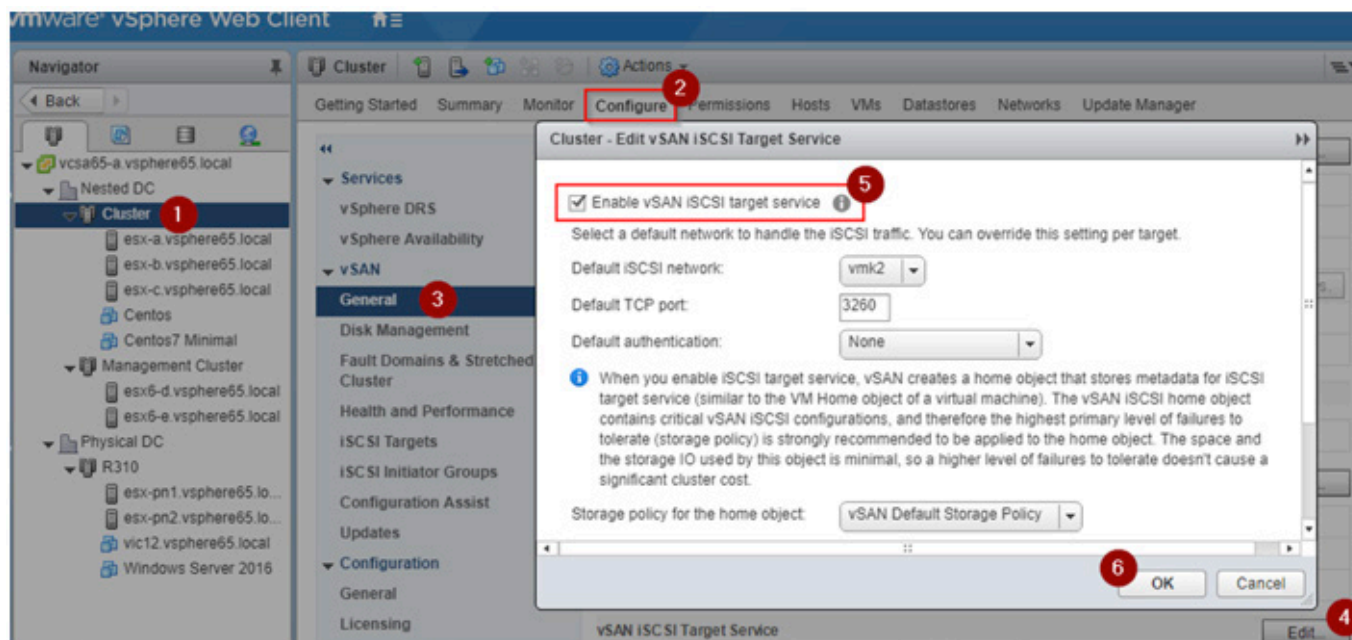


Figure 6

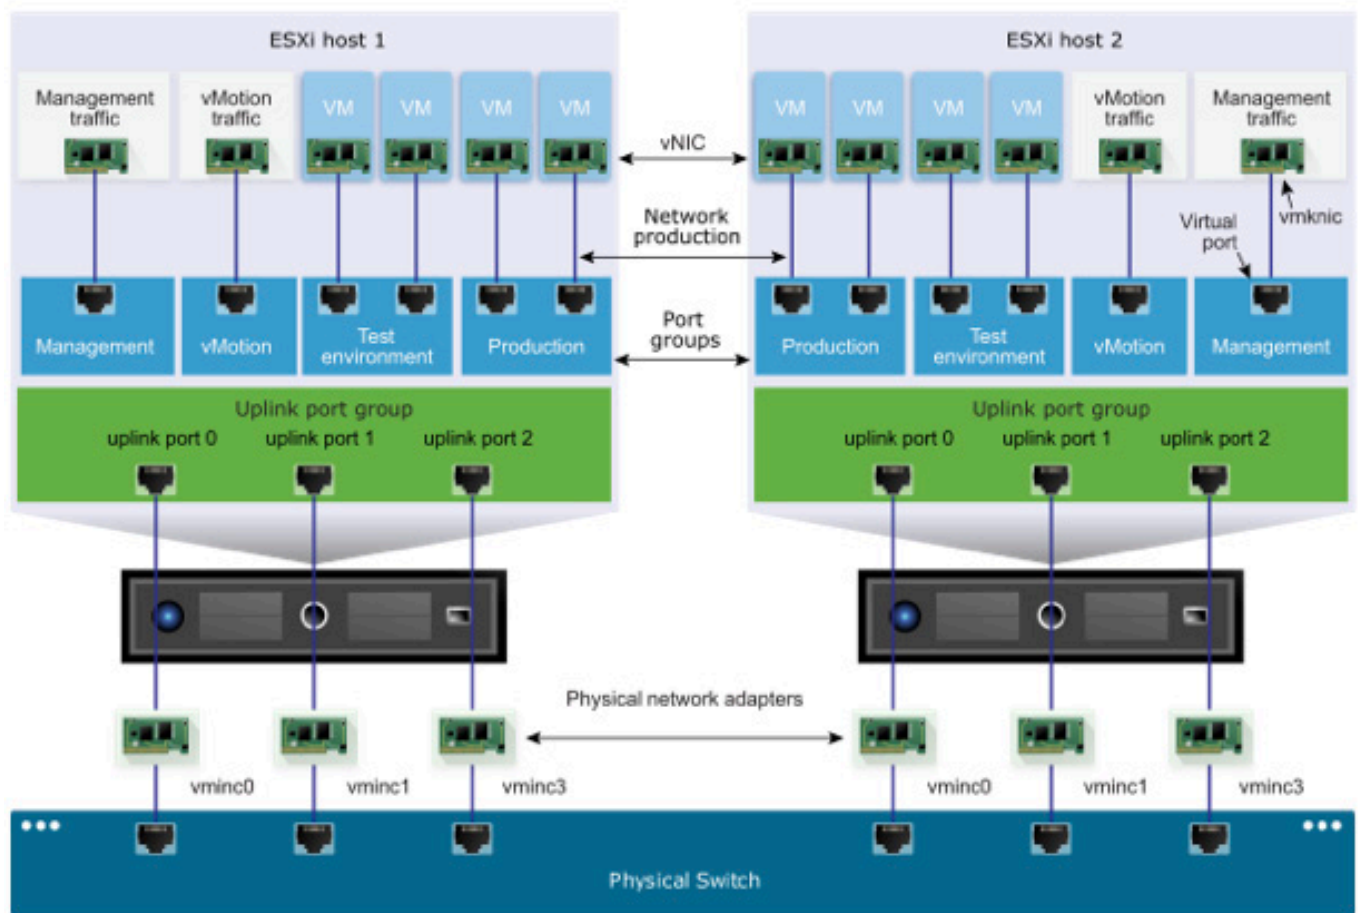
With the above option enabled, you can allow other non-VMware based iSCSI systems to connect to your storage as well. It's an often-overlooked aspect of a vSAN deployment that allows you to extend its use.

Moving on, vSAN requires a dedicated network and the following best practices should be observed:

- For hybrid configurations, dedicate *at least* 1-GbE physical network adapter. Place vSAN traffic on a dedicated or shared 10-GbE physical adapter for the best networking performance.
- For all-flash configurations, use a dedicated or shared 10-GbE physical network adapter.
- Provision one additional physical NIC as a failover NIC.
- If you use a shared 10-GbE network adapter, place the vSAN traffic on a distributed switch and configure Network I/O Control to guarantee bandwidth to vSAN.
- Always make sure that the hardware you are using is listed in the [VMware Compatibility Guide](#). I have seen people run labs on non-supported hardware just fine, but it is certainly a no-go in a production environment.

- In addition to hardware, also make sure that you run the supported software, driver and firmware versions in your cluster.
- Keep your environment updated and patched.
- Keep your hosts similar in terms of configuration and keep ESXi hosts similar in size. This is done to ensure that there is an even balance of virtual machine storage components across all the hosts and their disks.
- Having minimum 4 nodes in your vSAN cluster would give you higher availability when compared to running a 3-node cluster. Although even a 2-node configuration is supported with a quorum, I try to stick to 4-node configurations. VMware has a great whiteboard session [located here](#), if you're interested in 2-node deployments.
- When you size your vSAN cluster to support the number of VMs that you want to deploy, also make sure to size the ESXi hosts appropriately. This makes sure you have enough CPU/memory available for your VMs and not a large stock pile of disk.
- Always enable vSphere HA. Keep in mind that to enable vSAN, you will have to disable HA. Remember to turn it back on.
- Avoid inconsistent performance (remember: boring is good!) Make sure not to use hybrid and all-flash disk groups as part of the same cluster.
- Ensure that there are enough hosts in the cluster to accommodate the desired number of failures to tolerate. Use the formula $2n + 1$, where n is the number of failures to tolerate. If you want to tolerate 2 failures, you should at least 5 hosts in your cluster.
- Multiple smaller disk groups are recommended vs single larger disk group as multiple smaller groups allow for a higher cache to capacity ratio, thus leading to an accelerated performance of virtual machines.
- The cache tier should be sized to be at least 10% of the capacity consumed by virtual machines. You should revisit this number as your deployment grows to maintain the 10% recommendation.

PART 3 - MASTERING NETWORKING



First, we need to cover some differences between standard and distributed switches.

In a VMware environment, switches bring the physical network to virtual machines, while standard virtual switches and distributed virtual switches enable a network topology between VMs, hosts and host clusters.

A network switch directs network traffic. Similarly, a virtual switch carries VMs' traffic to the physical network and to other VMs. Without a switch, you have no connectivity.

Distributed vSwitches, which are also known as VMware vDS, enable more features than standard vSwitches, sometimes called VMware vSS. A standard vSwitch works

within one ESX/ESXi host only. Distributed vSwitches allow different hosts to use the switch, if they exist within the same host cluster. A distributed vSwitch extends its ports and management across all the servers in a cluster, supporting up to 500 hosts per distributed switch.

Instead of making virtual networks more complicated with its additional options, the distributed vSwitch simplifies operations and helps catch configuration errors and increase network visibility.

HELP! MY HOSTS WON'T COMMUNICATE WITH EACH OTHER!

When it comes to VMware networking, this is the number one problem. Generally, start with checking the following list:

- Does the ESXi host network configuration appear correct? IP, subnet mask, gateway?
 - ♦ `esxcli network nic list`
- Is the uplink plugged in? (Yes, that had to be said!)
 - ♦ `esxcli network vswitch standard portgroup list`
- If using VLANs, does the VLAD ID of the port group look correct?
 - ♦ `esxcli network vswitch standard portgroup list`
- Check the trunk port configuration on the switch. Have there been any recent changes?
- Does the physical uplink adapter have all settings configured properly? (speed, duplex, etc.)
 - ♦ `vicfg-nics -d duplex -s speed vmnic#`
- If using NIC teaming, is it setup and configured properly?
- Are you using supported hardware? Any driver issues?
- If all the above test ok, check that you don't have a physical adapter failure.

THE E1000 OR VMXNET3 ADAPTER?

If your default network adapter is E1000, you should consider changing the NIC to a vmxnet3 adapter.

The best practice from VMware is to use the VMXNET3 Virtual NIC unless there is a specific driver or compatibility reason where it cannot be used. In many cases, however, the E1000 has been installed, since it is the default.

The E1000 is a long existing, commonly available Intel-based device and most operating systems include built-in support. Because of that, there is no special driver required or any exceptional effort required to make it operate in a virtual environment. The problem is that the virtual device is just as described, a piece of software acting as if it was hardware. That can lead to performance issues.

The VMXNET3 virtual NIC is a completely virtualized 10 GB NIC. With this device the device drivers and network processing are integrated with the ESXi hypervisor. That means there is no additional processing required to emulate a hardware device and network performance is much better. There is no native VMXNET device driver in some operating systems such as Windows 2008 R2 and RedHat/CentOS 5 so VMware Tools is required to obtain the driver. VMware Tools is, however, highly recommended in any case so normally that is a non-issue.

DO I REALLY NEED 10 GIGABITS?

A 10-gigabit NIC is not required for basic connectivity but is slowly become more of the norm as time passes. More and more deployments are utilizing dual 10 GbE adapters. The key benefits include better utilization of I/O resources, simplified management, and reduced CAPEX and OPEX. While this deployment provides these benefits, there are some challenges when it comes to the traffic management aspects. Specially, in highly consolidated virtualized environments where more traffic types are carried over fewer 10-Gigabit Ethernet network adapters, and it becomes critical to prioritize traffic types that are important and provide the required SLA guarantees. Businesses often want it all and this can be difficult in terms of management. The NIOC (Network I/O Control) feature available on the VDS helps in this traffic management activity.

Network I/O Control will greatly reduce the issues that might arise. A typical 10 GbE scenario would involve the following:

REQUIREMENTS:

- Ensure high availability for all network traffic
- Provide high performance and redundant access to the IP Storage (if required)
- Prevent significant impact to storage performance by vMotion/Fault Tolerance and Virtual machines traffic
- Ensure ESXi hosts could be evacuated in a timely manner for maintenance

CONSTRAINTS:

- Two (2) x 10GB NICs

SOLUTION:

Use one dvSwitch to support all VMKernel and virtual machine network traffic and use “Route based on Physical NIC Load”.

Use Network I/O control to ensure in the event of contention that all traffic get appropriate network resources.

Configure the following Network Share Values:

- IP Storage traffic: 100
- ESXi Management: 25
- vMotion: 25
- Fault Tolerance: 25
- Virtual Machine traffic: 50

Configure two (2) VMKernels for IP Storage and set each on a different VLAN and Subnet.

Configure a VMKernel port for each of the following: vMotion (or Multi-NIC vMotion), ESXi Management and Fault Tolerance (if required) and set to active on both 10GB interfaces (default configuration).

All Distributed Port Groups for virtual machine traffic (typically in user defined vLANS) will be active on both interfaces.

The above utilizes LBT to load balance network traffic which will dynamically move workload between the two 10GB NICs once one or both network adapters reach $\geq 75\%$ utilization.

SCENARIO CONCLUSION

Even when your ESXi hosts only have two 10Gb interfaces, VMware provides enterprise-grade features to ensure all traffic, including IP storage, can get access to sufficient bandwidth to continue serving production workloads until the contention subsides.

This design ensures that in the event a host needs to be evacuated, even during production hours, it will complete in the fastest possible time with minimal or no impact to production. The faster your vMotion activity completes, the sooner DRS can get your cluster running as smoothly as possible. In the event you are patching, this means the sooner your maintenance can be completed and the hosts being patched are returned to the cluster to serve your VMs.

PART 4 - MASTERING TROUBLESHOOTING AND LOG ANALYSIS BEST PRACTICES

BUILD A TROUBLESHOOTING METHODOLOGY

ESXi and vSphere can problems arise from various origins, but they generally fall into one of the following categories:


- Hardware issues
- Resource contention
- Network attacks
- Software bugs
- Configuration problems

A typical troubleshooting process contains several tasks:

1. Define the problem and gather information
2. Identify what is causing the problem
3. Fix the problem, implement a fix

One of the first things you should do when experiencing a problem with a host is to try to reproduce the issue. If you can find a way to reproduce it, you have an ideal way to validate that the issue is resolved when you do fix it. It can be helpful as well to take a benchmark of your systems before they are implemented into a production environment. If you know HOW they should be running, it's easier to pinpoint a problem.

Top Down



You should decide if it's best to work from a "top down" or "bottom up" approach to determine the root cause. Guest OS Level issues typically cause a large amount of problems. Let's face it, some of the applications we use are not perfect. They get the job done but they utilize a lot of memory or other resources doing it.

In terms of virtual machine level issues, is it possible that you could have a limit or share value that's misconfigured?

At the ESXi Host Level, you could need additional resources. It's hard to believe sometimes, but you might need another host to help with the load!

Once you have identified the root cause, you should assess the impact of the problem on your day-to-day operations. When and what type of fix should you implement? A short-term one or a long-term solution? Assess the impact of your solution on daily operations.

- **Short-term solution:** Implement a quick workaround
- **Long-term solution:** Reconfiguration of a virtual machine or host

Know the various tools that VMware has available to make troubleshooting easier.

- VMware PowerCLI
 - ♦ VMware PowerCLI provides an easy-to-use Windows PowerShell interface for command-line access to administration tasks or for creating executable scripts.
- `esxcli` commands (`esxcli network`, `esxcli storage`, `esxcli vm`, etc.)

- A set of **esxcfg-*** commands
 - ♦ The **esxcfg** commands are deprecated but you will likely still see some older documentation with them. The recommendation today is to use **esxcli**.
- The host shell can be accessed a couple of ways, either by using the local DCUI (Direct Console User Interface) or via SSH.
 - ♦ Local access by using the Direct Console User Interface (DCUI):
 - Enable the vSphere ESXi Shell service, either in the DCUI or vSphere Web Client. *Typically, this is running by default.*
 - Access the ESXi Shell from the DCUI by pressing Alt-F1 after logging in.
 - When finished, disable the ESXi Shell service when not using it.
 - ♦ Remote access by using PuTTY or an SSH client.
 1. Enable the SSH service on your ESXi host, either in the DCUI or through the vSphere Web Client.
 2. Use PuTTY or your preferred SSH Client to access the ESXi host.
 3. Disable the SSH Service when finished.

One of the best ways to be able to help troubleshooting an issue is invoking cmdlets in PowerCLI. In this example below, you can collect debugging information by invoking the 'Get-ErrorReport' cmdlet to reproduce the problem and generate relevant information to assist you.

To reproduce the problem, the cmdlet runs the script block which is required to pass to the 'ProblemScript' parameter, in this example we will establish a connection to the vCenter Server and assume that invoking the cmdlet 'Get-VM' is generating an issue.

```
$Script = { Connect-VIServer vcenter.domain.local Get-VM }
```

Now that we have created the script text to which is generating an issue we need to invoke the 'Get-ErrorReport' cmdlet to generate the debugging information.

```
Get-ErrorReport -ProblemScriptTimeoutSeconds 60  
-ProblemDescription "Get-VM is not returning data" -Destination  
'D:\Output' -ProblemScript $Script
```

In the above example, we will be running the script to reproduce the issue, specifying a period of time to wait until considering the problem script is not responding, creating a description of the problem and a destination to collect the debugging information as a compressed file.

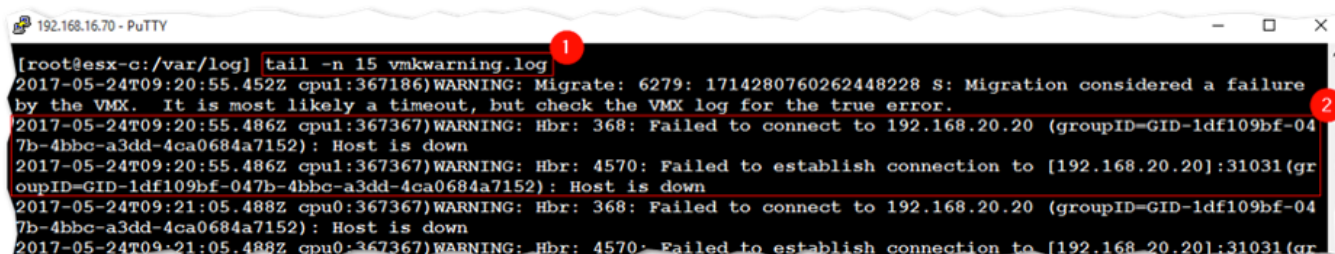
KNOW YOUR LOG FILES AND LOCATIONS!

Component	Location	Purpose
Authentication	/var/log/auth.log	Contains all events related to authentication for the local system.
ESXi host agent log	/var/log/hostd.log	Contains information about the agent that manages and configures the ESXi host and its virtual machines.
Shell log	/var/log/shell.log	Contains a record of all commands typed into the ESXi Shell as well as shell events.
System messages	/var/log/syslog.log	Contains all general log messages and can be used for troubleshooting.
vCenter agent log	/var/log/vpxa.log	Contains information about the agent that communicates with vCenter Server.
VMkernel	/var/log/vmkernel.log	Records activities related to virtual machines and ESXi.
VMkernel summary	/var/log/vmksummary.log	Used to determine uptime and availability statistics for ESXi.
VMkernel warnings	/var/log/vmkwarning.log	Records activities related to virtual machines.

For the most up-to-date list of log locations, visit [this page](#). Log locations change often between versions.

An effortless way to view logfiles on VMware ESXi is to SSH to the host and use old-fashioned Linux commands such as [cat](#), [more](#), [less](#), [tail](#) and [head](#) and use [grep](#) to filter for things you might want. A few examples are outlined below.

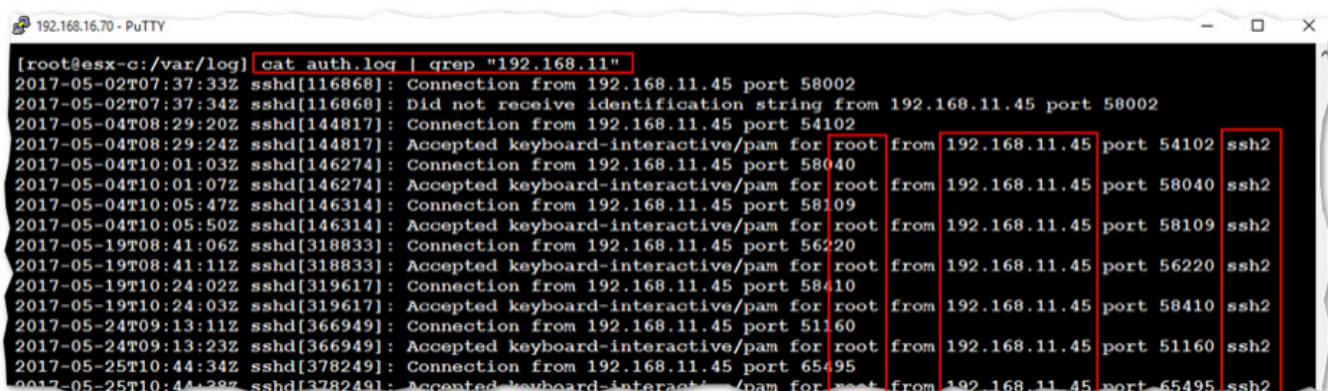
In this first example (Figure 7), the last 15 lines from the `vmkwarning.log` using the command `tail -n 15 <filename>`. This is marked (1) in the next screen screenshot. The text labeled (2), tells us that the host is failing to connect to host 192.168.20.20. As it turns out, 192.168.20.20 happens to be the IP address of a replication server I had set up on a second vCenter instance, which was powered off at the time.



```
[root@esx-c:/var/log] tail -n 15 vmkwarning.log
2017-05-24T09:20:55.452Z cpu1:367186)WARNING: Migrate: 6279: 1714280760262448228 S: Migration considered a failure by the VMX. It is most likely a timeout, but check the VMX log for the true error.
2017-05-24T09:20:55.486Z cpu1:367367)WARNING: Hbr: 368: Failed to connect to 192.168.20.20 (groupID=GID-1df109bf-047b-4bbc-a3dd-4ca0684a7152): Host is down
2017-05-24T09:20:55.486Z cpu1:367367)WARNING: Hbr: 4570: Failed to establish connection to [192.168.20.20]:31031(groupID=GID-1df109bf-047b-4bbc-a3dd-4ca0684a7152): Host is down
2017-05-24T09:21:05.488Z cpu0:367367)WARNING: Hbr: 368: Failed to connect to 192.168.20.20 (groupID=GID-1df109bf-047b-4bbc-a3dd-4ca0684a7152): Host is down
2017-05-24T09:21:05.488Z cpu0:367367)WARNING: Hbr: 4570: Failed to establish connection to [192.168.20.20]:31031(gr
```

Figure 7

In this next example (Figure 8), the `auth.log` log file is used to determine if connections are being established from subnet 192.168.11.0 and by whom. To do this, `cat auth.log` and pipe it into `grep` filtering by the string `192.168.11` as shown. The output shows many successfully established SSH connections via the root account from the 192.168.11.45.



```
[root@esx-c:/var/log] cat auth.log | grep "192.168.11"
2017-05-02T07:37:33Z sshd[116868]: Connection from 192.168.11.45 port 58002
2017-05-02T07:37:34Z sshd[116868]: Did not receive identification string from 192.168.11.45 port 58002
2017-05-04T08:29:20Z sshd[144817]: Connection from 192.168.11.45 port 54102
2017-05-04T08:29:24Z sshd[144817]: Accepted keyboard-interactive/pam for root from 192.168.11.45 port 54102 ssh2
2017-05-04T10:01:03Z sshd[146274]: Connection from 192.168.11.45 port 58040
2017-05-04T10:01:07Z sshd[146274]: Accepted keyboard-interactive/pam for root from 192.168.11.45 port 58040 ssh2
2017-05-04T10:05:47Z sshd[146314]: Connection from 192.168.11.45 port 58109
2017-05-04T10:05:50Z sshd[146314]: Accepted keyboard-interactive/pam for root from 192.168.11.45 port 58109 ssh2
2017-05-19T08:41:06Z sshd[318833]: Connection from 192.168.11.45 port 56220
2017-05-19T08:41:11Z sshd[318833]: Accepted keyboard-interactive/pam for root from 192.168.11.45 port 56220 ssh2
2017-05-19T10:24:02Z sshd[319617]: Connection from 192.168.11.45 port 58410
2017-05-19T10:24:03Z sshd[319617]: Accepted keyboard-interactive/pam for root from 192.168.11.45 port 58410 ssh2
2017-05-24T09:13:11Z sshd[366949]: Connection from 192.168.11.45 port 51160
2017-05-24T09:13:23Z sshd[366949]: Accepted keyboard-interactive/pam for root from 192.168.11.45 port 51160 ssh2
2017-05-25T10:44:34Z sshd[378249]: Connection from 192.168.11.45 port 65495
2017-05-25T10:44:38Z sshd[378249]: Accepted keyboard-interactive/pam for root from 192.168.11.45 port 65495 ssh2
```

Figure 8

When you cannot troubleshoot the issue yourself, you will need to contact VMware for support. It is best to generate a log file bundle. The support team will ask you for it and likely won't provide support until you upload it.

You can generate it two ways, via the ESXi host client, under Monitor -> Logs tab. Or you can generate a support bundle via the command line by running `/usr/bin/vm-support` from within an SSH session while logged as root. Once the bundle file is generated, you can copy it using `scp`.


The easiest method by far is just generating it via a web browser. Point your web browser to `http://<ESXi IP address>/cgi-bin/vm-support.cgi`. You are then asked for host credentials. Then the `vm-support` script is executed on the ESXi host. The generated bundle is then downloaded as a compressed tar file (TGZ). The process may take a while depending on the size of the logs, host utilization, uptime, etc.

CONCLUSION

To wrap up, there is no perfect way to become a Master of vSphere but by following best practices and taking the time to understand the underlying theory, you will become more at ease navigating your way through the platform and you will see your productivity rise. I highly recommend building a home lab and building out a practice environment. That's the best way to learn. Altaro has great resources on their [VMware blog](#) to get you headed in the right direction.

Thanks for reading!

ABOUT RYAN BIRK



Ryan has been working in Information Technology for 15+ years in a number of different roles. He has been a Virtualization Consultant, Engineer and, most recently, as a VMware Certified Technical Instructor. Since 2012, he has been a proud VMware vExpert and runs a blog at ryanbirk.com, which focuses on VMware home labs, and providing tips to those working with VMware on a daily basis.

[@ryanbirk](https://twitter.com/ryanbirk)

VMware vExpert 12-18

VMware Certified Instructor

Altaro VM Backup - Trusted by over 30,000 SMBs

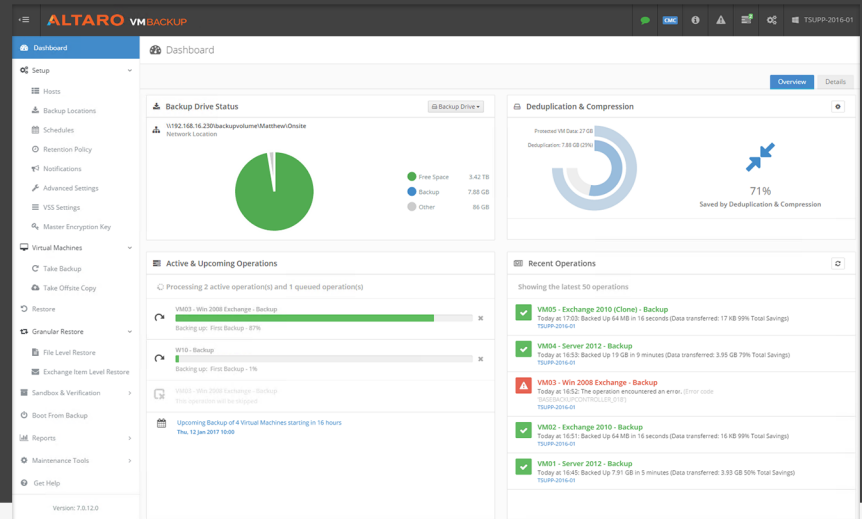
New v7.6! Altaro VM Backup for VMware & Hyper-V. Hassle-free and affordable VM backup software. Grab your free copy for 2 VMs now!

- ✓ Hassle-free and effective
- ✓ Unbeatable Value
- ✓ Outstanding Support

Free for 2 VMs, forever.

Back up unlimited VMs for 30 -days. After 30-days you get 2 VMs for free, forever. Download now!

**Download your
30-day trial**



Up and running quickly, without the need for complex configurations!

With Altaro VM Backup, you can install and run your first virtual machine (VM) backup in less than 15 minutes. Get up and running quickly, without the need for complex configurations or software dependencies.

Altaro VM Backup is designed to give you the power you need, without the hassle and steep learning curve.

- **Easy to use, intuitive UI** - making it easy to implement a rock solid backup strategy
- **Managing and configuring backup/restore jobs across multiple hosts has never been simpler**
- **Full control & scalability** - Monitor and manage all your Hyper-V and VMware hosts from a single console



Virtual machine backup software packed with powerful features for **VMware** and **Hyper-V**.

View features

ABOUT ALTARO

[Altaro Software](#) is a fast-growing developer of easy-to-use backup solutions which backs up and restores both Hyper-V and VMware-based virtual machines, built specifically for MSPs and SMBs customers with up to 50 host servers. Altaro take pride in their software and their excellent level of personal customer service and support, and it shows. Founded in 2009, Altaro already services over 40,000 satisfied customers worldwide and are a Gold Microsoft Partner for Application Development and Technology Alliance VMware Partner.

FOLLOW ALTARO

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



[Subscribe to our VMware blog](#) and receive best practices, tips, optimization guides and more here:

<https://www.altaro.com/vmware>

Follow Altaro at:



SHARE THIS RESOURCE!

Liked the eBook? Share it now on:



PUBLISHED BY
Altaro Software
<http://www.altaro.com>

Copyright © 2018 by Altaro Software

All rights reserved. No part of this book may be reproduced or transmitted in any form or by any means without the prior written permission of the publisher or authors.

WARNING AND DISCLAIMER

Every effort has been made to make this book as complete and as accurate as possible, but no warranty or fitness is implied. The information provided is on an “as is” basis. The authors and the publisher shall have neither liability nor responsibility to any person or entity with respect to any loss or damages arising from the information contained in this book.

FEEDBACK INFORMATION

We’d like to hear from you! If you have any comments about how we could improve the quality of this book, please don’t hesitate to contact us by visiting www.altaro.com or sending an email to our Customer Service representative Sam Perry: sam@altarosoftware.com