What's New in VMware vSphere® 5.1 – Platform

VMware vSphere 5.1

TECHNICAL MARKETING DOCUMENTATION
V 2.0 / UPDATED JUNE 2012

vmware®

Table of Contents

Introduction 3
vSphere Platform Enhancements 3
User Access 3
Auditing 4
Monitoring 4
vMotion Enhancements 4
Extended Guest OS and CPU Support 5
Agentless Antivirus and Antimalware 5
Virtual Machine Enhancements 5
New Virtual Machine Features 5
Introducing Virtual Machine Compatibility 7
Auto Deploy 9
Stateless Caching Mode 10
Stateful Install Mode 11
Remote Logging and Dump Collection 11
Improved Scalability 12
Conclusion 12

Introduction

VMware vSphere® 5.1 ("vSphere 5.1") provides many new features and enhancements that further extend the core capabilities of the VMware vSphere® ("vSphere") platform. Included among these are notable improvements in the areas of host security, logging, monitoring, and deployment; new VMware vSphere® vMotion® (vMotion) capabilities that enable virtual machines to be migrated between hosts and clusters with no shared storage; and support for the latest processors and guest operating systems (OS). This paper provides an overview of the new features and capabilities being introduced with vSphere 5.1.

This paper is organized into the following four sections:

- vSphere host features and enhancements
- Enhanced vMotion capabilities
- Virtual machine features and enhancements
- Introduction to VMware vSphere® Auto Deploy (Auto Deploy) stateless caching and stateful install

vSphere Platform Enhancements

A key focus for vSphere 5.1 was on improving host security, auditing, and monitoring. This included eliminating the reliance on a shared "root" user account and adding support for SNMPv3.

User Access

A challenge with earlier vSphere releases was the requirement for administrators to share a common "root" user account when working from the host shell. The reliance on a shared administrative account came with several challenges:

- All shell activity was logged as "root," making it difficult to audit individual user activity on the host.
- The need to coordinate password resets among many administrators often resulted in the "root" user becoming an exception to reoccurring password change policies. As a result, the "root" password was rarely, if ever, changed.
- It was difficult to track the number of administrators with root access.
- It was difficult to remove access when personnel left the company or changed roles.

To address these challenges and improve host security, VMware made several improvements to the ESXi shell in vSphere 5.1:

1. There is no longer a dependency on a shared root account. Local users assigned administrative privileges automatically get full shell access. With full shell access, local users no longer must "su" to root to run privileged commands.
2. Administrative privileges can be removed from the default root account, effectively eliminating a popular target for anyone who might attempt to gain unauthorized access to a vSphere host.
3. The ability to automatically terminate idle shell sessions with a new *UserVars.ESXiShellInteractiveTimeOut* variable. Termination of inactive shell sessions helps prevent security breaches that can occur when administrators fail to properly log out of the shell.

Auditing

In vSphere 5.1, all host activity, from both the shell and the Direct Console User Interface (DCUI), is now logged under the account of the logged-in user. This ensures user accountability, making it easy to monitor and audit activity on the host.

Monitoring

vSphere 5.1 adds support for SNMPv3, which provides many improvements over SNMPv2. These include added security, with SNMP authentication, and added privacy, with SSL encryption. SNMPv3 also provides additional configuration capabilities through SNMP Set. Also in vSphere 5.1, the SNMP agent has been unbundled from the VMkernel and now runs as an independent agent on the host. This makes it easier to incorporate SNMP updates and patches because they are no longer tied to the vSphere kernel.

vMotion Enhancements

Enhancements to vMotion in vSphere 5.1 provide a new level of ease and flexibility for live virtual machine migrations. vSphere 5.1 now enables users to combine vMotion and VMware vSphere® Storage vMotion® (Storage vMotion) into one operation. The combined migration copies both the virtual machine memory and its disk over the network to the destination host. In smaller environments, the ability to simultaneously migrate both memory and storage enables virtual machines to be migrated between hosts that do not have shared storage. In larger environments, this capability enables virtual machines to be migrated between clusters that do not have a common set of datastores.

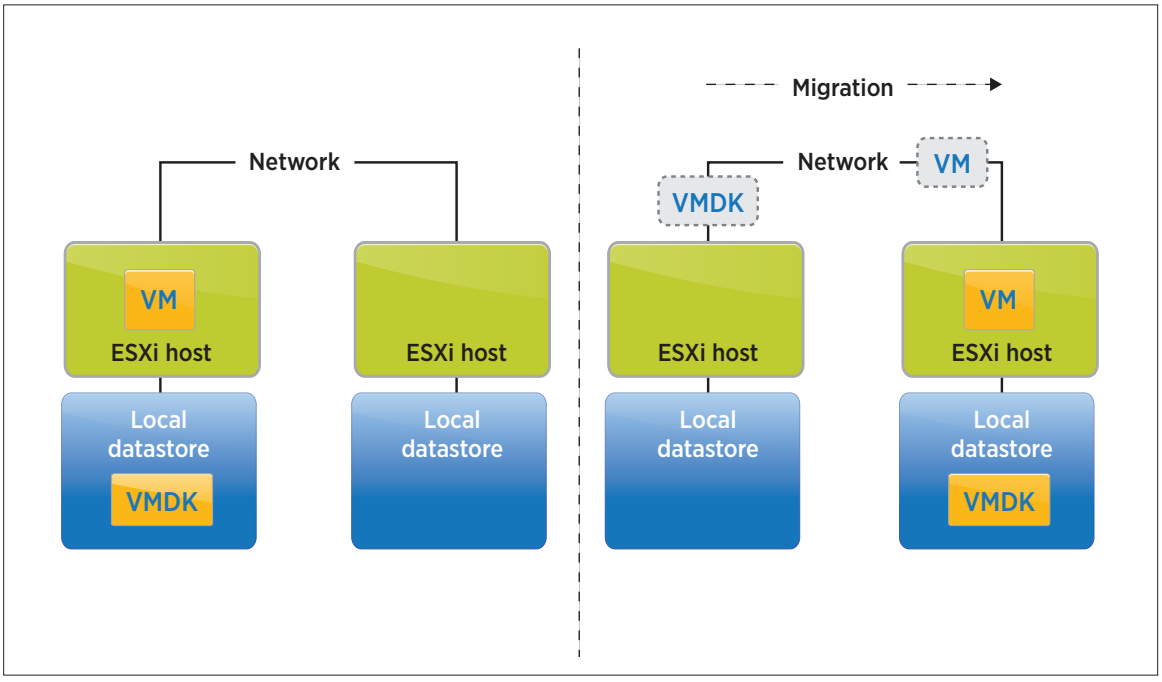

Figure 1.

Extended Guest OS and CPU Support

Along with security, auditing and monitoring updates, and the enhanced vMotion capabilities, vSphere 5.1 adds support for the latest guest operating systems and CPUs. vSphere 5.1 introduces support for Microsoft Windows 8, both server and desktop editions, along with support for the latest AMD and Intel CPUs, including the AMD "Piledriver" series and the Intel "Ivy Bridge" and "Sandy Bridge" series.

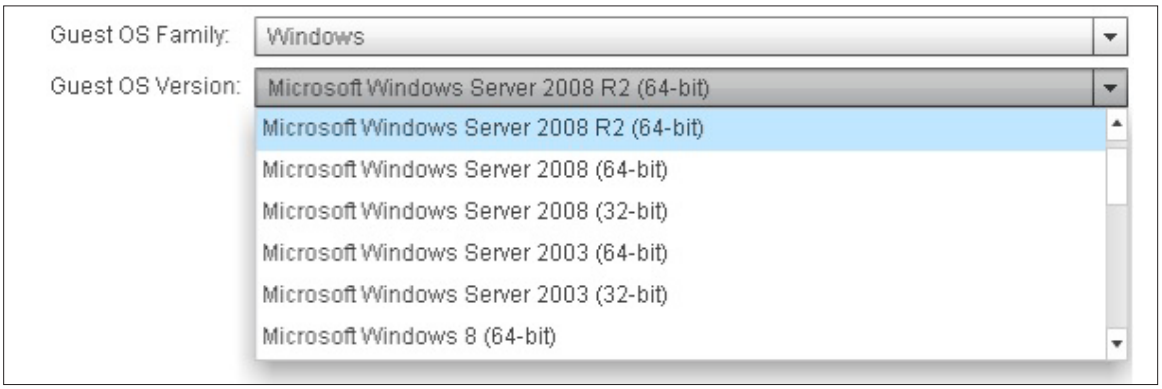

Figure 2.

For a comprehensive list of the supported guest operating systems and server hardware, refer to the *VMware Compatibility Guide*. The guide is updated frequently, so it's a good idea to check it often to get the latest updates.

Agentless Antivirus and Antimalware

VMware® vShield Endpoint™ (vShield Endpoint)—now included in vSphere 5.1—offloads antivirus and antimalware agent processing inside guest virtual machines to a dedicated secure virtual appliance delivered by VMware® partners. Solutions based on vShield Endpoint improve consolidation ratios and performance by eliminating antivirus "storms" due to signature updates, because updating occurs on the security virtual appliance instead of inside each individual virtual machine. These solutions also streamline antivirus and antimalware deployment and monitoring in VMware environments and satisfy compliance and audit requirements through logging of antivirus and antimalware activities.

Virtual Machine Enhancements

vSphere 5.1 includes several virtual machine improvements as well. It now supports up to 64 vCPUs in a single virtual machine, enables improved graphics support with the introduction of hardware-based virtual graphics processing units (vGPUs), introduces guest-level storage reclamation and provides improved CPU virtualization capabilities.

NOTE: The vGPU and guest-level storage reclamation capabilities provided in vSphere 5.1 have a dependency on a future version of VMware® View™ (View). Refer to the View documentation for information on when these features will be available.

New Virtual Machine Features

64-vCPU Support

Virtual machines running on a vSphere 5.1 host can be configured with up to 64 vCPUs, which demonstrates the maturity of the vSphere hypervisor and how it is well suited to take on even the largest of application workloads. With the ability to host virtual machines running up to 64 vCPUs and 1TB of RAM, vSphere 5.1 enables users to virtualize any size workload with confidence.

Hardware-Accelerated 3D Graphics Support

With vSphere 5.1, VMware has partnered with NVIDIA to provide hardware-based vGPU support inside the virtual machine. vGPUs improve the graphics capabilities of a virtual machine by off-loading graphic-intensive workloads to a physical GPU installed on the vSphere host. In vSphere 5.1, the new vGPU support targets View environments that run graphic-intensive workloads such as graphic design and medical imaging.

Hardware-based vGPU support in vSphere 5.1 is limited to View environments running on vSphere hosts with supported NVIDIA GPU cards (refer to the VMware Compatibility Guide for details on supported GPU adapters). In addition, the initial release of vGPU is supported only with desktop virtual machines running Microsoft Windows 7 or 8. Refer to the View documentation for more information on the vGPU capabilities of vSphere 5.1.

NOTE: vGPU support is enabled in vSphere 5.1, but the ability to leverage this feature is dependent on a future release of View. Refer to the View documentation for information on when this feature will be available.

Guest OS Storage Reclamation

Over time, the storage requirements of virtual machines will change. While thin provisioning enables users to dynamically add capacity to a virtual machine as its storage requirements increase, prior to vSphere 5.1 there was not a way to return storage to the storage array free pool when it was no longer being used by the guest OS. With vSphere 5.1, this is now possible using the Guest OS Storage Reclamation feature. With Guest OS Storage Reclamation, when files are removed from inside the guest OS, the size of the VMDK file can be reduced and the deallocated storage space returned to the storage array's free pool. Guest OS Storage Reclamation utilizes a new SE sparse VMDK format available with View.

NOTE: Guest OS Storage Reclamation is enabled with vSphere 5.1, but the ability to leverage this feature is dependent on a future release of View. Refer to the View documentation for information on when this feature will be available.

Improved CPU Virtualization

In vSphere 5.1, the vSphere host is better able to virtualize the physical CPU and expose more information about the CPU architecture to the virtual machine. This improved CPU virtualization, referred to as virtualized hardware virtualization (VHV), provides the guest OS with near-native access to the physical CPU. VHV enables, for example, the running of Microsoft Windows XP Mode from inside a Windows 7 or 8 virtual machine.

NOTE: Microsoft Windows XP Mode enables running legacy Windows XP applications that are not compatible with newer Windows versions inside a simulated XP environment on Microsoft Windows 7 or 8.

In addition to the improved CPU virtualization with VHV, vSphere 5.1 adds the ability to expose more low-level CPU counter information to the guest OS, which enables improved debugging, tuning and troubleshooting of the OS and applications running inside the virtual machine. This facilitates the ability to develop, test and optimize application workloads running inside a VMware virtual machine.

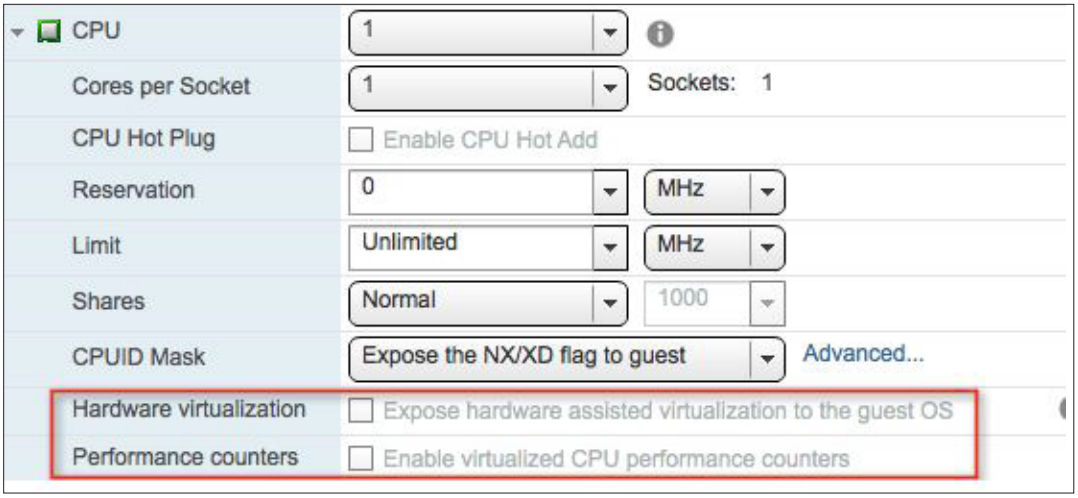

Figure 3.

Introducing Virtual Machine Compatibility

With each new vSphere release, VMware continues to enhance virtual machine features and capabilities. To help track the features and capabilities associated with a given virtual machine, VMware has traditionally used a virtual machine hardware version number. A list of virtual machine hardware version numbers associated with each vSphere release is shown in the following table:

vSPHERE RELEASE	VIRTUAL MACHINE HARDWARE VERSION
Virtual Infrastructure 3.5	Version 4
vSphere 4.0	Version 4
vSphere 4.1	Version 7
vSphere 5.0	Version 8
vSphere 5.1	Version 9

Table 1. Virtual Machine Hardware Version Support Matrix

Using a virtual machine's hardware version number, administrators can readily identify each virtual machine's feature set and capabilities. Historically, it has been considered a good practice that each time an upgrade is made to a vSphere infrastructure (VMware® vCenter Server™ (vCenter Server) and associated vSphere hosts), the virtual machine hardware version for each virtual machine is also upgraded. This would ensure that each virtual machine had the ability to fully benefit from the latest updates and newest features. However, over time several challenges to this approach have been identified:

1. Upgrading a virtual machine's hardware version requires powering off the virtual machine, which requires downtime. With many virtual machines now hosting mission-critical workloads, it can be difficult to schedule downtime.
2. As the size of the virtual infrastructure grows, trying to keep a large number of virtual machines running at a common virtual machine hardware version proves difficult if not impractical.
3. In many cases, a virtual machine's requirements are fully satisfied with the older virtual hardware version, so a virtual machine upgrade might yield little or no benefit.

With these challenges in mind, and to help simplify virtual infrastructure upgrades, VMware formally extended the virtual machine compatibility matrix in vSphere 5.0. The extended support matrix enables virtual machines running older virtual hardware versions to run fully supported on newer versions of vSphere. The following table shows the extended virtual machine hardware version support matrix:

vSPHERE RELEASE	SUPPORTED VIRTUAL MACHINE HARDWARE VERSIONS
Virtual Infrastructure 3.5	Version 4
vSphere 4.0	Version 4
vSphere 4.1	Version 4, 7
vSphere 5.0	Version 4, 7, 8
vSphere 5.1	Version 4, 7, 8, 9

Table 2. Extended Virtual Machine Hardware Version Support Matrix

A notable benefit of the extended support matrix is the ability to upgrade a vSphere infrastructure (vCenter Server and vSphere hosts) with no virtual machine downtime. Following the vSphere infrastructure upgrade, any virtual machines that do not require any of the newer features or capabilities provided with the latest version can continue to run with the older virtual hardware, completely unaffected by the underlying infrastructure upgrade.

Although the extended support matrix relaxes the requirement to upgrade a virtual machine's hardware version, many administrators still felt compelled to upgrade simply to "keep up" with an ever-increasing virtual machine hardware version number. In an effort to help eliminate this latent pressure to keep up, VMware has made a critical change to the virtual machine hardware nomenclature in vSphere 5.1 with the introduction of *virtual machine compatibility.*

Virtual machine compatibility provides the means to readily identify the features and capabilities associated with a virtual machine by indicating the corresponding vSphere release on which the virtual machine is based. Unlike the virtual hardware version number, virtual machine compatibility is tied to a specific vSphere release. For example, a virtual machine created on vSphere 4.1 will show as "compatibility level 4.x"; a virtual machine created on vSphere 5.1 will show as "compatibility level 5.1." The following table shows the virtual machine compatibility levels in vSphere 5.1:

vSPHERE RELEASE	VIRTUAL MACHINE HARDWARE VERSION	vSPHERE 5.1 COMPATIBILITY
Virtual Infrastructure 3.5	Version 4	VMware ESX 3.x and later
vSphere 4.0	Version 4	VMware ESX 3.x and later
vSphere 4.1	Version 7	VMware ESX 4.x and later
vSphere 5.0	Version 8	VMware ESX 5.0 and later
vSphere 5.1	Version 9	VMware ESX 5.1 and later

Table 3. Virtual Machine Compatibility Levels In vSphere 5.1

Virtual machine compatibility removes ambiguity regarding virtual machine upgrades by eliminating the need to associate, for example, a virtual machine running virtual machine hardware version 7 with vSphere 4.1 or one running virtual machine hardware version 9 with vSphere 5.1.

Figure 4.

Changing Virtual Machine Compatibility Levels

Similar to upgrading a virtual machine's hardware version, users can also upgrade from one compatibility level to a higher one. However, a virtual machine cannot be downgraded to a lower level. Upgrading a virtual machine's compatibility level does require downtime, because the virtual machine must be powered off.

Setting a Default Virtual Machine Compatibility Level

To help with managing virtual machine compatibility across clusters, vSphere 5.1 enables users to define a default compatibility at both the host and cluster levels. This enables users to ensure that new virtual machines created on a host/cluster are done so using the preferred compatibility level. This prevents a new virtual machine from automatically defaulting to the highest compatibility level. It is especially useful during rolling upgrades, when users might be running a mix of vSphere host versions together in a cluster.

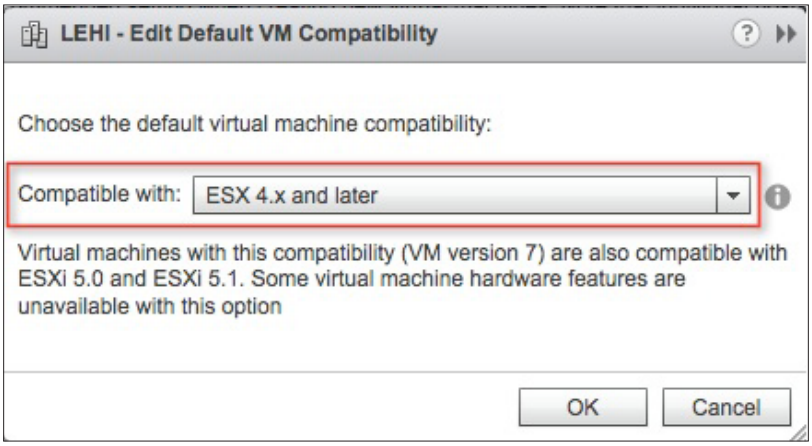

Figure 5.

NOTE: In vSphere 5.1, virtual machine compatibility is available only from the new Web UI. The VMware vSphere® Client™ (vSphere Client), VMware vSphere® Command-Line Interface (vSphere vCLI), VMware vSphere® PowerCLI (vSphere PowerCLI) and VMware vSphere® Storage APIs continue to use the virtual hardware version number.

Auto Deploy

vSphere 5.0 introduced VMware vSphere Auto Deploy, a new way to rapidly deploy new vSphere hosts. With Auto Deploy, the vSphere host PXE boots over the network and is connected to an Auto Deploy server, where the vSphere host software is provisioned directly into the host's memory. After the vSphere software has been installed on the host, the Auto Deploy server configures the host, using a VMware® vCenter™ (vCenter) host profile. Then the host is connected to the vCenter Server.

Auto Deploy significantly reduces the amount of time required to deploy new vSphere hosts. And because an Auto Deploy host runs directly from memory, there is no requirement for a dedicated boot disk. This not only provides cost savings, because there is no need to allocate boot storage for each host, but it also can simplify the SAN configuration, because there is no need to provision and zone LUNs each time a new host is deployed. In addition, because the host configuration comes from a host profile there is no need to create and maintain custom pre- and postinstall scripts.

Along with the rapid deployment, cost savings and simplified configuration, Auto Deploy provides the following benefits:

- Each host reboot is comparable to a fresh install. This eliminates configuration drift.
- With no configuration drift between vSphere hosts, less time will be spent troubleshooting and diagnosing configuration issues.
- Simplified patching and upgrading. Applying updates is as easy as creating a new image profile, updating the corresponding rule on the Auto Deploy server and rebooting the hosts. In the unlikely event you must remove an update, reverting back to the previous image profile is also easy: 1) Reupdate the rule to assign the original image profile and 2) do another reboot.

NOTE: Because an Auto Deploy host runs directly from memory, it often is referred to as being "stateless." This is because the host state (i.e., configuration) that is normally stored on a boot disk comes from the vCenter host profile.

In vSphere 5.0, Auto Deploy supported only one operational mode, which was referred to as "stateless" (also known as "diskless"). vSphere 5.1 extends Auto Deploy with the addition of two new operational modes, "stateless caching" and "stateful installs."

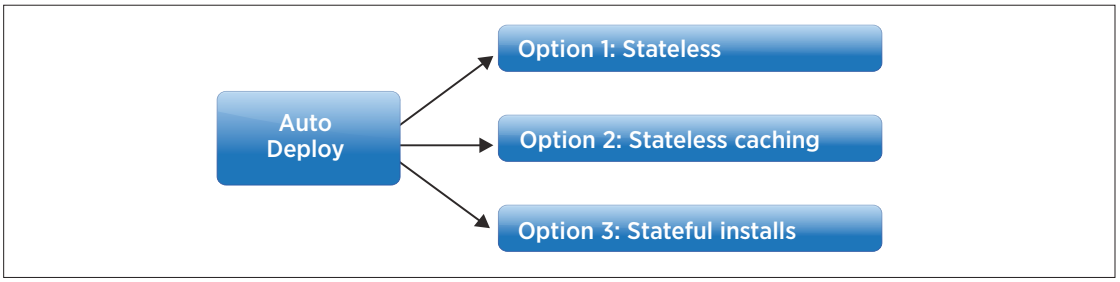

Figure 6.

Stateless Caching Mode

The Auto Deploy stateless caching mode was implemented to help address availability concerns with the PXE boot infrastructure and Auto Deploy server. With stateless hosts, if the PXE boot failed, or if the Auto Deploy server was unavailable, the host would not be able to boot until the outage was corrected. However, with stateless caching, if a host cannot boot due to a problem with the PXE environment or Auto Deploy server, it is able to fall back to booting off a cached image saved to a dedicated boot device. After booting from the cached image, the administrator is able to use the host to help troubleshoot and identify why the PXE boot might have failed.

The stateless caching mode is very similar to the stateless mode, in that during normal operation the host PXE boots from the Auto Deploy server. However, the difference is that with stateless caching, an additional step is taken where the software image running in memory is cached (saved) to a dedicated boot device (local disk, SAN, USB). This cached image then acts as a backup from which the host can boot in the event there is a problem with the PXE boot or Auto Deploy infrastructure.

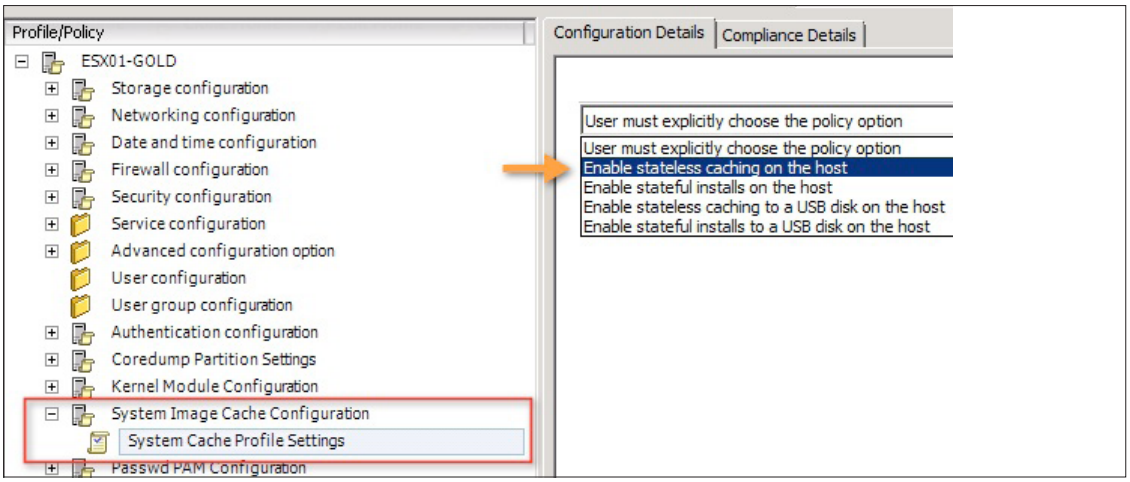

Figure 7.

Unlike the stateless mode, stateless caching requires a dedicated boot device for each vSphere host. In addition, users must configure the host's BIOS settings to first attempt to boot over the network and, if that fails, to fall back to booting from disk.

With stateless caching, if there is a localized outage that affects the PXE boot infrastructure (DHCP or TFTP server) or the Auto Deploy server but does not affect the vCenter Server instance, by using the cached image the host will be able to boot and the administrator able to manually reconnect to the vCenter Server.

NOTE: Stateless caching does not protect against a vCenter Server failure. Always protect a vCenter Server by running it in a vSphere cluster protected by VMware vSphere® High Availability (VMware HA) or VMware vCenter Server Heartbeat™ (vCenter Server Heartbeat).

Stateful Install Mode

Auto Deploy stateful install mode enables administrators to leverage the Auto Deploy infrastructure to provision new vSphere hosts. With stateful install, users perform a one-time PXE boot of a new host from the Auto Deploy server. Following the one-time PXE boot, all subsequent reboots will take place from the dedicated boot disk.

Setting up stateful installs is similar to configuring stateless caching. The difference is the BIOS boot order configured on the server. Where stateless caching is set to boot from the network first and fall back to the local disk only when the PXE boot fails, with stateful installs the host is configured to always try to boot from the local disk first and boot from the network only when no boot image can be found on the disk. With Auto Deploy stateful install mode, a new host will perform an initial one-time PXE boot using the Auto Deploy infrastructure to configure the host. After the initial boot, all subsequent reboots take place using the boot device.

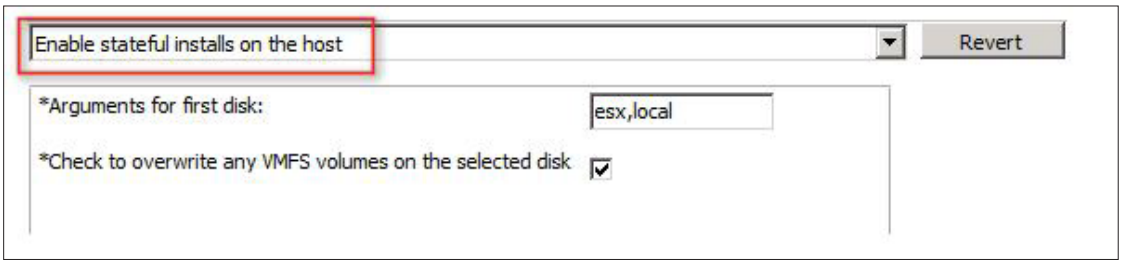

Figure 8.

With stateful installs, the Auto Deploy infrastructure is being leveraged as a provisioning tool similarly to how scripted installations or kickstart might be used. The advantage to Auto Deploy stateful install is that users are able to rapidly deploy hosts without having the need to create and maintain custom scripts. The software to be installed is determined using the Auto Deploy rules engine, and the host is configured using the vCenter host profiles and therefore doesn't rely on external scripts.

With stateful installs, users leverage the Auto Deploy infrastructure to provision new hosts but forgo most of the benefits of stateless or stateless caching because after the vSphere hosts have been deployed, they must be maintained, patched and updated individually.

Remote Logging and Dump Collection

Because Auto Deploy hosts running in the stateless or stateless caching mode run directly from memory, it's important to enable remote logging and set up remote dump collection. Without remote logging or dump collection, log files and dumps will be stored on a RAM disk in memory and will be lost each time the host is rebooted.

To help with remote logging and dump collection, VMware provides two Auto Deploy add-on components: the Syslog Collector and the ESXi Dump Collector. Both of these add-on components are free of charge and are provided as part of the vCenter 5.1 installation media.

NOTE: The VMware® vCenter™ Server Appliance™ (VMware VCSA) comes with the Auto Deploy Server, Syslog Collector and ESXi Dump Collector preinstalled. It is not necessary to install these components when using the VCSA. Users simply must enable them with the VCSA configuration interface.

Improved Scalability

In vSphere 5.1, optimizations have been made to the way the host image is passed from the Auto Deploy server to the host, enabling up to 80 concurrent host reboots from a singleAuto Deploy server.

NOTE: To scale beyond 80 concurrent reboots per Auto Deploy host, use reverse Web proxies.

Conclusion

VMware vSphere 5.1 introduces many new features and enhancements that further extend the core capabilities of the vSphere platform. The core platform enhancements in vSphere 5.1 include the following:

- Improved security and auditing by removing reliance on a shared "root" account
- Improved SNMP monitoring with SNMPv3 authentication and SSL encryption
- Enhanced vMotion support, enabling virtual machines to be migrated between hosts and clusters that do not have shared storage
- Support for the latest Intel and AMD CPUs along with support for Microsoft Windows 8
- Bundling of VMware® vShield Endpoint™ as part of the core vSphere product

Along with the core platform improvements, vSphere 5.1 provides the following several virtual machine–related enhancements:

- Support for up to 64 vCPUs per virtual machine, doubling the number of support vCPUs from vSphere 5.0
- Enhanced CPU virtualization, enabling the passing of low-level CPU counters and other physical CPU attributes directly to the virtual machine, where they can be accessed by the guest OS
- Introduction of virtual machine compatibility, making it easier to identify and track virtual machine capabilities—removes ambiguity regarding virtual machine upgrades and helps eliminate pressure to keep up with an ever-increasing virtual hardware version

For VMware View environments, vSphere 5.1 puts in place the following critical enabling technologies that will be used with future releases of View:

- Enhanced hardware-accelerated graphics support with virtual GPUs (vGPUs)
- Guest OS Reclamation, providing the ability to return deallocated storage space inside the guest OS to the storage array

vSphere 5.1 also provides several critical enhancements to VMware vSphere Auto Deploy:

- Improved scalability of the Auto Deploy server, with support for up to 80 concurrent (re)boots, twice the number of concurrent reboots with vSphere 5.0
- A new stateless caching mode that enables the Auto Deploy host to reboot in the event of an outage affecting the PXE infrastructure or Auto Deploy server
- A new stateful install mode that enables the Auto Deploy infrastructure to be used as a provisioning tool to deploy new vSphere hosts

VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001 www.vmware.com