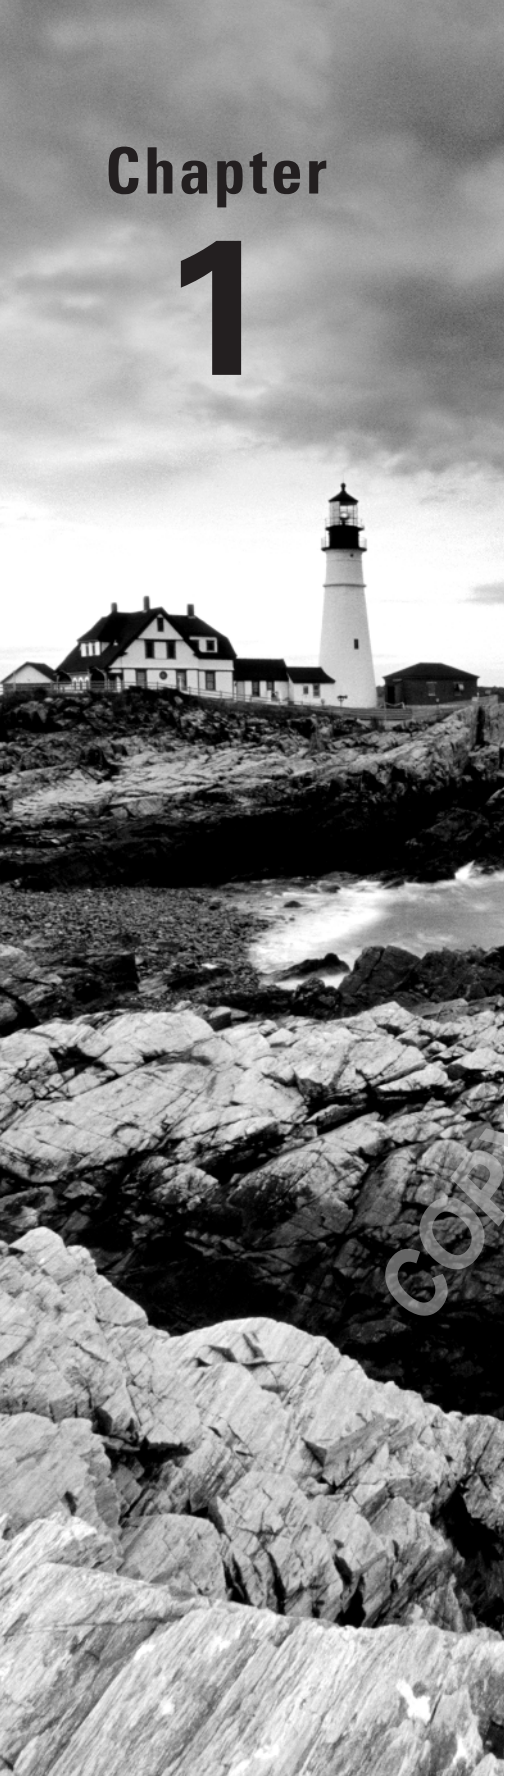

Chapter 1

What's New in vSphere 5

THIS CHAPTER COVERS THE FOLLOWING TOPICS:

- ✓ ESX Retirement
- ✓ VCB Retirement
- ✓ VMI Paravirtualization Retirement
- ✓ VMware GUI Toolbox Retirement
- ✓ Windows 2000 Guest OS Customization Support
- ✓ Newly Created VMCI Sockets Unsupported
- ✓ Requirement of LAHF and SAHF CPU Instruction Sets
- ✓ Intel SMT–Related CPU Scheduler Enhancements
- ✓ Notable Configuration Maximums Changes
- ✓ ESXi Firewall and Management Networks
- ✓ Swap to SSD
- ✓ Support for Hardware and Software FCoE Adapters
- ✓ Host UEFI Boot Support
- ✓ Improved SNMP Support
- ✓ New Command-Line Interface
- ✓ vSphere High Availability Improvements
- ✓ Virtual Machine Enhancements
- ✓ Expanded Support for VMware Tools
- ✓ Mac OS X Server Support

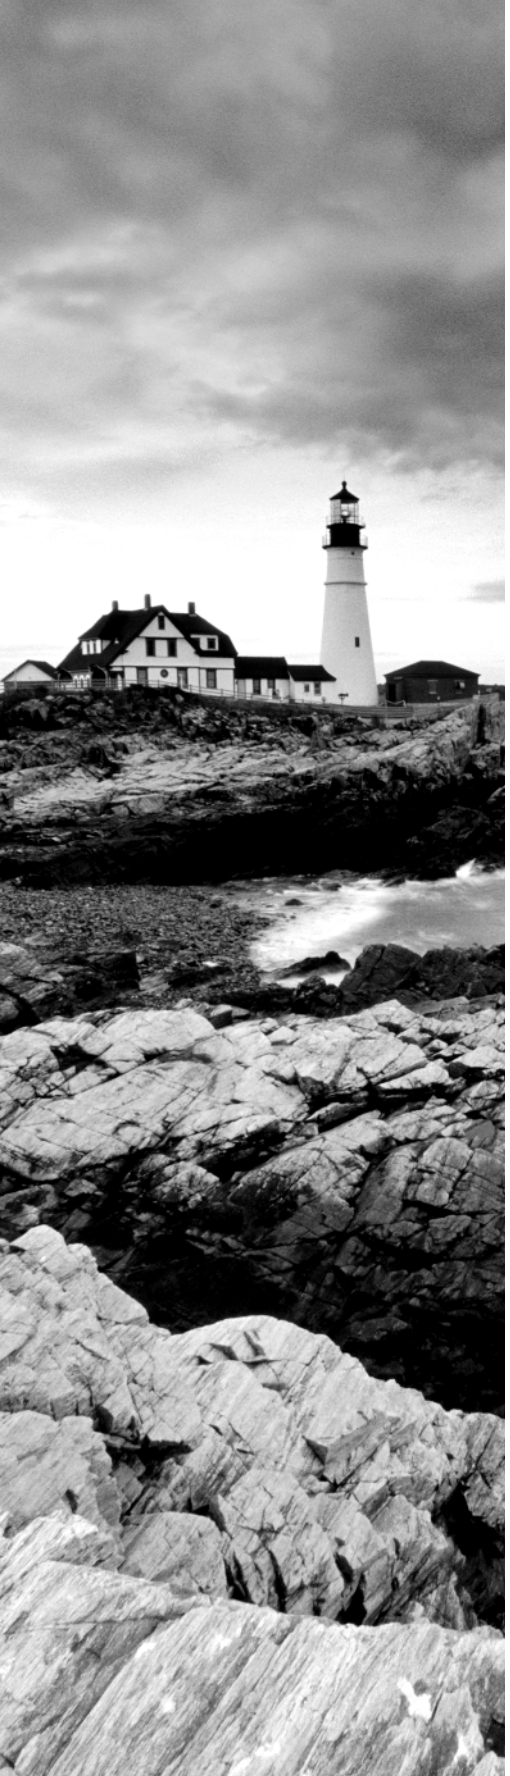

- ✓ vCenter Enhanced Logging
- ✓ VMware vCenter Server Heartbeat Improvements
- ✓ Fault Tolerance (FT) Improvements
- ✓ iSCSI UI Support
- ✓ GUI to Configure Multi-core Virtual CPUs
- ✓ vNetwork Distributed Switch Improvements
- ✓ Network I/O Control Improvements
- ✓ Storage vMotion Improvements
- ✓ VAAI Thin Provisioning Improvements
- ✓ NFS Support Improvements
- ✓ Storage Accelerator
- ✓ VMFS-5
- ✓ vMotion and Metro vMotion
- ✓ vCenter Server Appliance
- ✓ vSphere Storage Appliance (VSA)
- ✓ Storage DRS
- ✓ Profile-Driven Storage
- ✓ Image Builder
- ✓ Auto Deploy
- ✓ VMware Data Recovery 2.0 Improvements
- ✓ Update Manager Improvements
- ✓ Licensing

vSphere 5 continues to build on the rich feature set of vSphere 4.1, which is the industry-leading virtualization platform for building cloud infrastructures. There are new capabilities at all levels of the product, and the introduction of these new features will in many ways be what will differentiate the VCP4 from the VCP5. In this chapter, I will review these new vSphere 5 features. The purpose here is simply to introduce what is new in vSphere 5. Many of these features will be covered in much greater detail in the following chapters.

ESX Retirement

As of vSphere 5, ESX is no longer available. ESXi is the hypervisor architecture that will be used in vSphere 5 and beyond. For those who have been using only ESX and not ESXi, one of the most notable changes is the absence of the Service Console. This represents new operational and administration challenges, but it also improves security, reliability, and management. There is a menu-driven Direct Console User Interface (DCUI) and an ESXi Shell that can each be used in ESXi 5. The ESXi Shell looks similar to the Service Console in ESX.

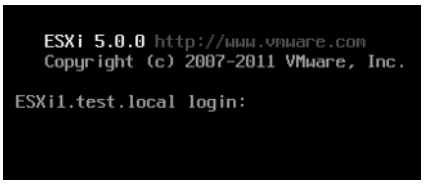

The ultimate idea with removing the Service Console is to move away from the concept of using the Service Console as a direct means to administer your hosts. As mentioned earlier, removing the Service Console improves security, reliability, and host management. The host should almost be thought of as a piece of hot-swappable hardware in your virtual infrastructure. The removal of the Service Console helps enable this way of thinking. If you haven't done so already, take a look at either the VMware vSphere PowerCLI, the vSphere Command-Line Interface (vCLI), or the vSphere Management Assistant (vMA). You will notice that these utilities offer equivalent functionality, but the syntax may be quite different. If you have come to rely on the Service Console for operational or administrative tasks and are not familiar with the vCLI, PowerCLI, or vMA, you will have a bit of a learning curve. Also, you'll need to consider items such as SAN agents, multipathing software, backup agents/software, hardware monitoring, UPS software, and any scripts that were being used in ESX and how these functionalities will work, or be replaced, in ESXi.

VCB Retirement

VMware Consolidated Backup (VCB) is no longer supported in vSphere 5. VMware Data Recovery will be covered in the exam. If you are interested in using third-party backup applications, you should look into whether these vendors support the VMware vStorage APIs for Data Protection.

VMI Paravirtualization Retirement

vSphere 5 does not support the Virtual Machine Interface (VMI) guest operating system paravirtualization interface. This is because innovations made in CPU hardware acceleration technologies from Intel and AMD have allowed these newer processors to outperform VMI. Virtual machines that use paravirtualization can be migrated to vSphere 5, and more information is available in VMware KB 1013842.

VMware GUI Toolbox Retirement

vSphere 5 will be the final vSphere release to support the VMware Tools graphical user interface, VMware Toolbox. The Toolbox command-line interface (CLI) will continue to be updated and supported. The Toolbox command-line interface provides the same functionality that was previously available in the VMware Tools control panel. The Toolbox command-line interface is named differently, depending on the guest operating system it is accessed from:

- In Windows, use `VMwareToolboxCmd.exe`.
- In Mac OS X, use `./vmware-tools-cli`.
- In Linux, FreeBSD, and Solaris, use `vmware-toolbox-cmd`.

For example, to list the available devices for a virtual machine running a Windows OS, you would run the following command from the Windows command line:

```
VMwareToolboxCmd.exe device list
```

Windows 2000 Guest OS Customization Support

vSphere 5 will be the last vSphere release to support guest customization of the Windows 2000 guest operating system. Newer versions of Windows guests will continue to be supported for guest customization.

Newly Created VMCI Sockets Unsupported

Virtual machine–to–virtual machine communications are no longer supported in vSphere 5. This functionality will be removed in the next major vSphere release, but VMware will continue to support host-to-guest communications.

Requirement of LAHF and SAHF CPU Instruction Sets

For ESXi 5 hosts, only CPUs that contain the Load Register AH From Flags (LAHF) and Store AH Register into Flags (SAHF) CPU instruction sets are supported. The ESXi installer will check the compatibility of the host CPU during an installation or upgrade. If your host hardware is not compatible, a purple screen will appear with an incompatibility information message. The install or upgrade cannot continue beyond this point. Always use the VMware HCL (`www.vmware.com/go/hcl`) to verify host compatibility with ESXi 5 before you begin the installation.

Intel SMT–Related CPU Scheduler Enhancements

The Intel simultaneous multithreading (SMT) architecture exposes two hardware contexts from a single processor core. Utilizing two hardware contexts yields a 10 percent to 30 percent improvement in application performance, depending on the workload. The ESXi 5 CPU scheduler's policy is tuned for this type of architecture and ensures high efficiency and performance for mission-critical applications.

Notable Configuration Maximums Changes

While the VCP5 exam is unlikely to quiz you on simple configuration maximum questions, this data is extremely important to know for the exam. It is much more likely that you will see a question where knowing a specific configuration maximum is a prerequisite for knowing the correct answer to the question. The following list is only a

sampling of the configuration maximums and focuses only on those that have changed from vSphere 4.1.

- vCPUs per VM changes from 8 to 32
- RAM allocation per VM changes from 255GB to 1TB
- VM swap file size changes from 255GB to 1TB
- VMs per host changes from 320 to 512
- vCPUs per host changes from 512 to 2048
- RAM per host changes from 1TB to 2TB
- NFS mounts per host changes from 64 to 256
- LUN size changes from 2TB minus 512 bytes to 64TB
- Port groups per vSwitch changes from 512 to 256
- Maximum concurrent host HA failover changes from 4 to 32
- Resource pools per cluster changes from 512 to 1,600
- Resource pools per host changes from 4,096 to 1,600
- Number of hosts per datacenter changes from 400 to 500

ESXi Firewall and Management Networks

With the removal of the Service Console in vSphere 5, VMware now provides a new firewall to protect the ESXi host management interface(s). This firewall is provided through a VMkernel network adapter (vmknic)–level firewall module. The following are the key features of this firewall:

- Service-oriented and stateless
- Supports ability to restrict access to services based on IP address and subnet mask
- Can be managed through the vSphere Client or the new `esxcli` command-line interface
- Supports host profiles
- Eliminates the use of iptables

  To provide the ability to create new services and firewall rules, you can also define firewall services through XML description files. You create these XML files manually and activate them by running an `esxcli refresh` command.

Swap to SSD

In vSphere 5, the VMkernel automatically recognizes and tags solid-state drive (SSD) devices. These devices can be located on the local ESXi host or on the network, and the VMkernel scheduler allows ESXi swap to extend to these local or network SSD devices. Swapping to SSD is significantly faster than swapping to a traditional spinning disk, and this feature can help minimize the performance impact of swapping resulting from memory overcommitment.

Support for Hardware and Software FCoE Adapters

Support has been added for many software and hardware Fibre Channel over Ethernet (FCoE) adapters. The following are some of the supported software FCoE adapters:

- 10 Gigabit BR KX4 Dual Port Network Connection
- 10 Gigabit Dual Port Backplane Connection
- 10 Gigabit CX4 Dual Port Network Connection
- 10 Gigabit Network Connection
- 10 Gigabit Dual Port Network Connection
- 10 Gigabit Network Connection
- 10 Gigabit BX Network Connection
- 10 Gigabit AT2 Server Adapter
- 10 Gigabit KX4 Network Connection
- 10 Gigabit KR2 Network Connection
- 10 Gigabit TN Network Connection

These are some of the supported hardware FCoE adapters:

- Emulex OneConnect OCe10100 10GbE, FCoE UCNA
- Emulex OneConnect OCe11100 10GbE, FCoE UCNA

As with any storage or network adapter, always consult the VMware HCL (`www.vmware.com/go/hcl`) to check the compatibility or availability for specific FCoE adapters.

Host UEFI Boot Support

ESXi 5 hosts are supported for booting from the Unified Extensible Firmware Interface (UEFI). UEFI allows the host server to boot from hard disks, CD-ROM drives, or USB media.

Improved SNMP Support

The capability to convert Common Information Model (CIM) indications to Simple Network Management Protocol (SNMP) traps has been added to vSphere 5. vSphere 5 also supports the Host Resources MIB (RFC 2790) and allows for greater control over the types of traps sent from the SNMP agent.

New Command-Line Interface

vSphere 5 introduces a new command-line interface (CLI) called `esxcli`. The introduction of the `esxcli` command marks the beginning of efforts by VMware to standardize on a single command-line interface for both local and remote administration. The `esxcli` command is available on ESXi 5 hosts from the ESXi Shell. It is also available in the vSphere Command-Line Interface or through the vSphere Management Assistant.

vSphere High Availability Improvements

vSphere High Availability (HA) has been completely rewritten for the vSphere 5 release to increase scalability, reliability, and usability. HA now uses the concept of a master-slave relationship between the nodes in a cluster, where before it relied on the concept of primary and secondary nodes. The master host is responsible for the following:

- Monitoring the state of the slave hosts
- Monitoring the power states of all protected VMs
- Managing the lists of cluster hosts and protected VMs
- Serving as the vCenter Server management interface by reporting the cluster health state

The slave hosts run virtual machines locally, monitor their runtime states, and report their state updates to the master host. This new model simplifies planning and design considerations and allows for easier scalability. This is particularly true for vSphere environments running on blade servers and for stretched cluster environments.

HA also no longer has any dependency on DNS servers and can also communicate between the cluster nodes through the storage subsystem in a process called *datastore heartbeating*. Datastore heartbeating allows a master host that has lost communication with a slave host over the management network to communicate over the storage network to determine whether the slave host has failed, is in a network partition, or is network isolated. If a slave host cannot be reached via datastore heartbeating, it is considered failed. Using multiple paths of communication allows for a greater level of redundancy and better identification of the node's actual health.

The HA user interface has been improved to more easily show the role a node plays in the cluster and its current state. There is now a single log file for HA, which should allow for easier troubleshooting of HA events.

Virtual Machine Enhancements

With the release of vSphere 5, VMware introduces a new virtual hardware version for virtual machines. This newest release is virtual hardware version 8. Improvements in this latest version include the following:

- 32-way vCPUs
- Virtual NUMA (vNUMA) support
- 1TB RAM maximum
- Non-hardware-accelerated 3D graphics for Windows Aero support
- USB 3.0 device support
- UEFI virtual BIOS
- Client-connected USB devices
- Smart card reader support

Figure 1.1 shows the new capabilities of virtual hardware version 8.

The CPU and memory capabilities are provided to handle the virtualization of even more Tier 1 applications, while the enhanced graphics capabilities, USB 3.0 support, and smart card reader support are clearly targeted at providing more features to virtual desktop infrastructure (VDI) environments.

Expanded Support for VMware Tools

vSphere 5 supports hosting virtual machines running prior versions of VMware Tools and virtual hardware. This means virtual machines with the 4.*x* version of VMware Tools and older virtual machine hardware versions can run fully supported in a vSphere 5 environment. This compatibility will make the vSphere 5 upgrade process much easier to manage.

Table 1.1 shows the vSphere 4.*x* and vSphere *5* compatibility for VMware Tools and the virtual hardware versions.

FIGURE 1.1 Virtual Machine Properties editor

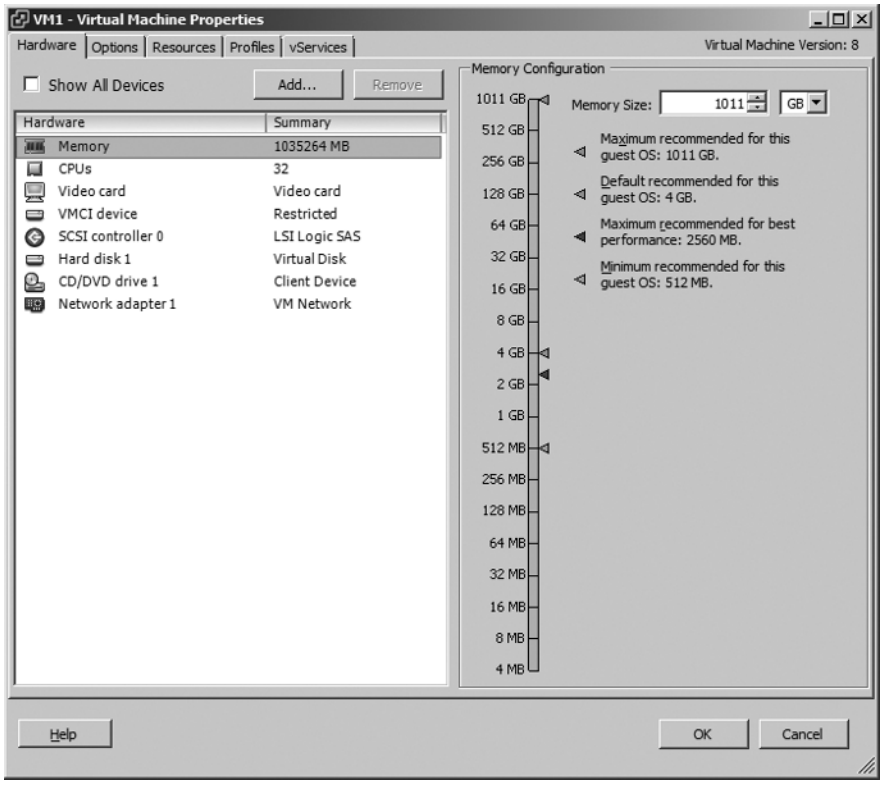

TABLE 1.1 vSphere 4.*x* and vSphere 5 VMware Tools compatibility

Version	vSphere 4.*x*	vSphere 5
VMware Tools 4.*x*	Yes	Yes
VMware Tools 5	Yes	Yes
Virtual hardware	3, 4, 7	4, 7, 8

Mac OS X Server Support

vSphere 5 supports OS X Server 10.6 (Snow Leopard) as a guest operating system when ESXi is installed on Apple Xserve servers. It is important to note that Apple hardware is an absolute requirement for this guest OS support, and the ESXi host will verify whether Apple hardware is being used when provisioning an OS X Server 10.6 guest.

The use of vMotion to migrate Mac OS X Server virtual machines requires both the source and destination ESXi 5 hosts to be on Apple Xserve hardware. Mac OS X 10.6 running in a virtual machine will also require the virtual hardware version 8 EFI BIOS option, which vSphere 4.*x* does not support. It is also important to note that VMware Fault Tolerance (FT) is not supported with Mac OS X Server virtual machines.

vCenter Server Enhanced Logging

vCenter Server system message logging has several enhancements. All log messages are now syslog-generated, and these messages can be logged locally or on one or more remote syslog servers. There is also a bundled VMware Syslog Collector on the vCenter Server installation media that can be used to collect these syslog messages. If you plan on using the vCenter Server Appliance, know that it can use the native Linux syslog-ng facility for syslog messages. You can configure this message logging with either the vSphere Client or the `esxcli` command.

VMware vCenter Server Heartbeat Improvements

When designing your virtual infrastructure, you must consider the availability of vCenter Server. This availability is covered in objective 1.1 of the VCP5 exam, so it is important to understand how to leverage VMware vCenter Server Heartbeat. vCenter Server Heartbeat is used to provide a highly available vCenter Server. The latest version of VMware vCenter Server Heartbeat, version 6.4, provides improvements in the areas of manageability, usability, and application support. Figure 1.2 shows the basic design of vCenter Server Heartbeat.

FIGURE 1.2 vCenter Server Heartbeat architecture

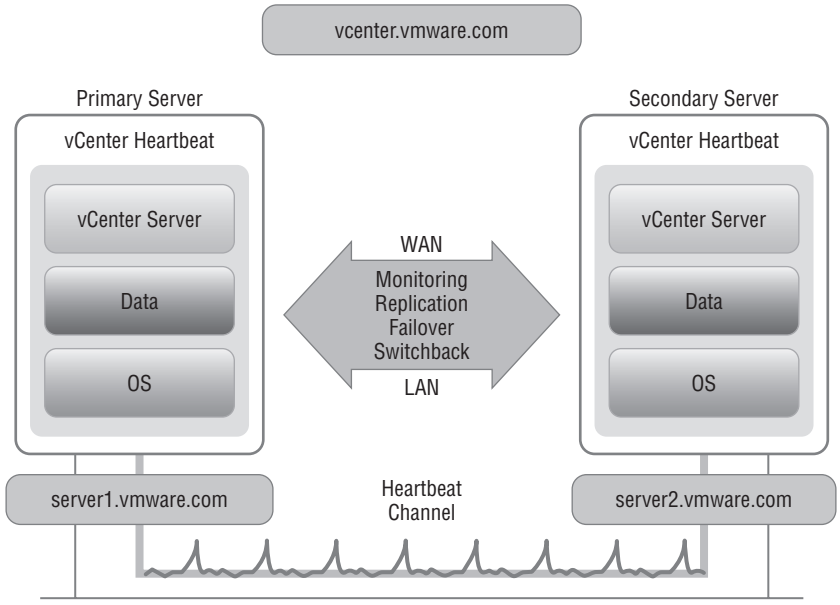

vCenter Server Heartbeat provides a high availability solution that is easy to manage. While other clustering solutions can be used to provide a highly available vCenter Server, the complexity can also be high with these solutions. If a solution becomes too complicated, errors could inadvertently be introduced and decrease the overall effectiveness of the solution.

Like other clustering solutions, VMware vCenter Server Heartbeat 6.4 works with two servers. vCenter Server Heartbeat allows both the active and passive servers to be represented as unique entities within Microsoft Active Directory. Each vCenter Server is assigned a unique IP address, and each vCenter Server is accessible on the network at all times.

Each vCenter Server instance can be associated with a virtual IP address. The virtual IP address follows the active VMware vCenter Server instance, if and when a failover is required. This means that users need to know only a single vCenter address to connect to and that administrators can perform maintenance actions in a transparent manner.

vCenter Server Heartbeat leverages vCenter Server and the vSphere Client to allow a single pane of glass for VMware vCenter Server Heartbeat operations. Administrators now have the ability to monitor and perform management functions from within the vSphere Client via a new vSphere Client plug-in. In addition, vCenter Server Heartbeat operational tasks and alarms are registered with the vCenter Server.

vCenter Server Heartbeat now also provides availability for VMware View Composer and Microsoft SQL Server 2008 R2. This support provides high availability for virtual desktops and protection for the vCenter Server database, even if it is installed on a separate server.

It is important to note that VMware vCenter Server Heartbeat is a separately licensed product available from VMware, and it is not included with any edition of vSphere 5. Although it is not an included component, a separate clustering solution used to provide high availability for vCenter Server would likely also include additional costs. There is more information on these scenarios in the next chapter, including supported configurations for achieving high availability for vCenter Server.

Fault Tolerance Improvements

With vSphere 5, VMware Fault Tolerance (FT) gains support for new CPU architectures and guest operating systems. Refer to VMware KB 1008027 for the most up-to-date information on supported CPUs and guest operating systems that can be used with VMware FT. Also, check out my regularly updated blog entry "VMware Fault Tolerance Requirements and Limitations" for much more information on VMware Fault Tolerance:

```
http://communities.vmware.com/people/vmroyale/blog/2009/05/18/vmware-fault-tolerance-requirements-and-limitations
```

iSCSI UI Support

For anyone who has ever had to configure iSCSI in previous versions of vSphere, the iSCSI GUI will be a welcome addition in vSphere 5. The new iSCSI GUI allows you to configure dependent hardware iSCSI and software iSCSI adapters along with their network configurations and port bindings in a single dialog box from within the vSphere Client. Figure 1.3 shows the new Network Configuration tab.

FIGURE 1.3 iSCSI configuration GUI

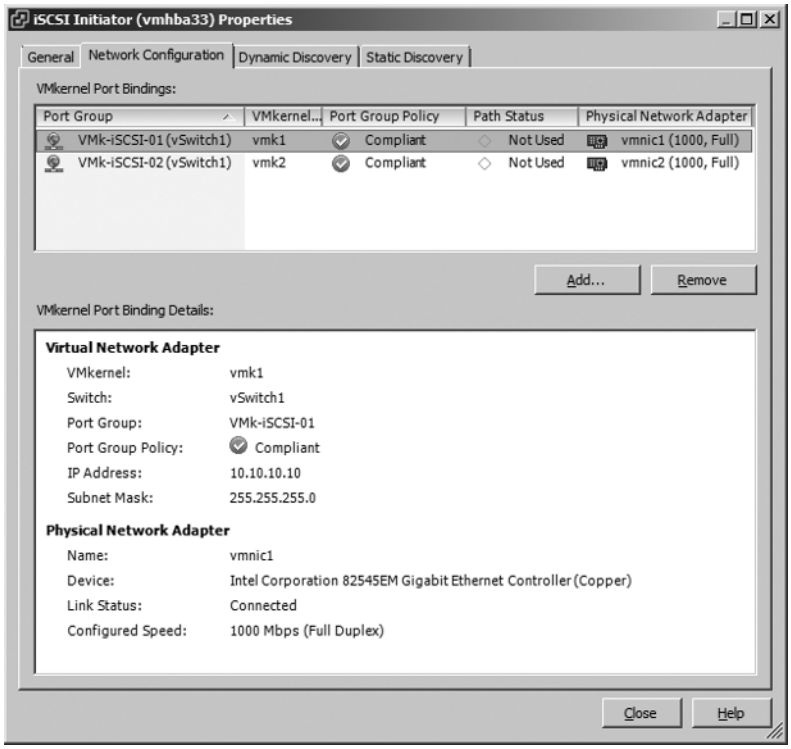

GUI to Configure Multi-core Virtual CPUs

The new GUI on the Virtual Machine Properties dialog box allows you to configure multi-core vCPUs. You can now modify the number of virtual CPU cores per socket without having to use the Advanced Settings options. Figure 1.4 shows the new CPU/multi-core configuration settings.

vNetwork Distributed Switch Improvements

vSphere 5 introduces two new features in the Distributed vSwitch that provide more visibility into traffic that is flowing in the virtual infrastructure. The first of these features is

NetFlow; vSphere 5 supports NetFlow 5. The NetFlow capability provided in vSphere 5 provides the following visibility:

- Intrahost virtual machine traffic (virtual machine–to–virtual machine traffic on the same host)
- Interhost virtual machine traffic (virtual machine–to–virtual machine traffic on different hosts)
- Virtual machine–to–physical infrastructure traffic

FIGURE 1.4 New CPU/multi-core configuration settings

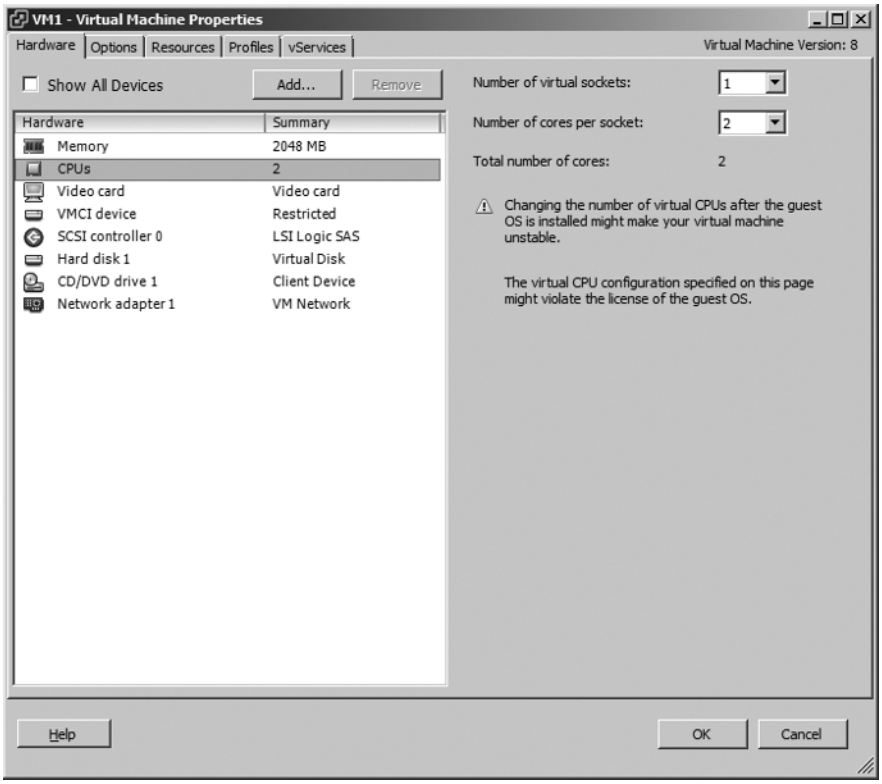

NetFlow can be enabled on Distributed vSwitches at the port group level, at the individual port level, or at the uplink level. Enabling flow monitoring on a Distributed vSwitch allows the monitoring of application flows. This NetFlow data helps in capacity planning and in ensuring that I/O resources are properly utilized.

The second new feature in the Distributed vSwitch is port mirroring. *Port mirroring* is when a network switch sends a copy of network packets seen on a switch port to a network monitoring device connected to another switch port. This is also known as Switch Port Analyzer (SPAN) on Cisco switches. Port mirror configuration can be done at the

distributed switch level and is used for network monitoring or troubleshooting. The port mirroring capability in vSphere 5 provides the following visibility:

- Intrahost virtual machine traffic (virtual machine–to–virtual machine traffic on the same host)
- Interhost virtual machine traffic (virtual machine–to–virtual machine traffic on different hosts)

FIGURE 1.5 Configuring Discovery Protocol settings

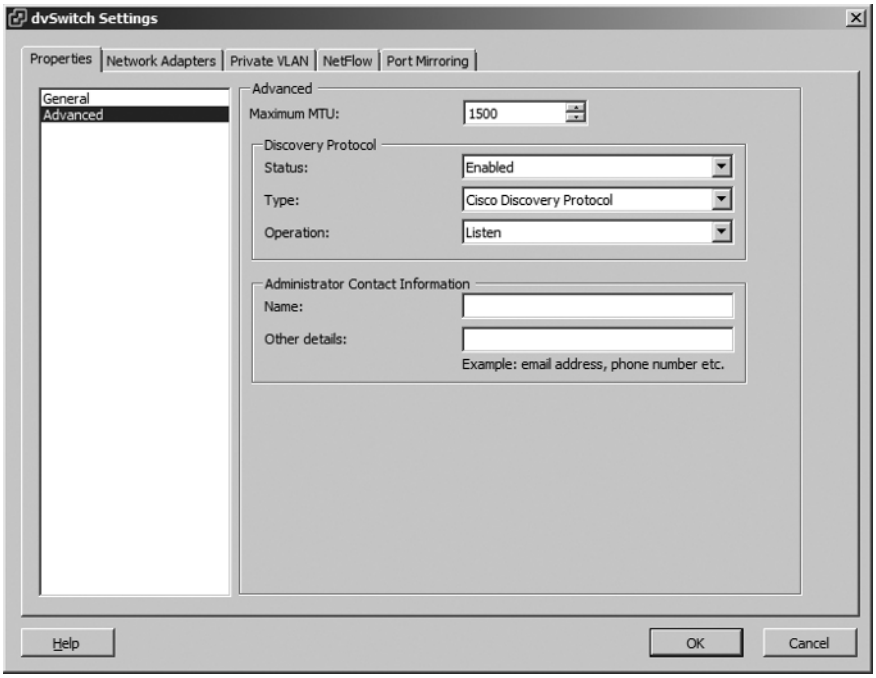

If you have been using vSphere 4.*x*, then you have likely used Cisco Discovery Protocol (CDP) data before and know the value it brings to identifying network devices and their associated configurations. vSphere 5 now also supports IEEE 802.1AB standards-based Link Layer Discovery Protocol (LLDP). Like CDP, LLDP is used to discover information about network devices, but it is a vendor-neutral discovery protocol. LLDP can be enabled at the distributed switch level by selecting either the CDP or LLDP discovery protocol type. You can also configure the operation mode for the discovery protocol from the distributed switch settings. The following are the three available options:

`Listen` Only listen for upstream network information.

`Advertise` Only advertise information about this distributed switch.

`Both` Listen for upstream network information and advertise information about this distributed switch.

You can configure the Discovery Protocol settings in the distributed switch settings, under Advanced, as shown in Figure 1.5.

Network I/O Control Improvements

Network I/O Control (NIOC) allows the creation of resource pools containing network bandwidth. Users can create new resource pools to associate with port groups and specify 802.1p tags, allowing different virtual machines to be in different resource pools. This allows a subset of virtual machines to be given a higher or lower share of bandwidth than the others.

In vSphere 5, NIOC supports traffic management capabilities for the following traffic types:

- Virtual machine
- Management
- iSCSI
- NFS
- Fault Tolerance logging
- vMotion
- vSphere replication
- User-defined

Storage vMotion Improvements

Storage vMotion now supports virtual machines in snapshot mode and the migration of linked clones. In addition, there have been improvements made in the copying mechanisms used by Storage vMotion, which will result in greater efficiency and migration time predictability.

VAAI Thin Provisioning Improvements

VAAI enhancements have been introduced in vSphere 5 to optimize thin provisioning capabilities. The challenges around using thin provisioning have typically been the reclamation of dead space and monitoring the storage to ensure that an out-of-space condition never occurs. VAAI Thin Provisioning introduces Dead Space Reclamation and Out-of-Space Conditions to address these problems.

Dead Space Reclamation does exactly what its name implies. When a virtual disk is deleted or migrated off of the datastore, the storage device will be made aware that these blocks are no longer in use, and the blocks may then be reclaimed. Out-of-Space Conditions is a set of advanced warnings and errors that may be used to notify

administrators when important thresholds are reached for thin-provisioned datastores. Also introduced are mechanisms that can temporarily pause a virtual machine, when disk space is depleted. This enables the issue to be addressed by a virtual infrastructure administrator, without resulting in the failure of a virtual machine.

NFS Support Improvements

vSphere 5 introduces NFS support for Storage I/O Control. Also introduced is hardware acceleration for NAS, which will allow faster provisioning and the use of thick virtual disks on NFS volumes. This is accomplished through two newly introduced VAAI primitives:

- Full File Clone: Enables virtual disks to be cloned
- Reserve Space: Enables creation of thick virtual disk files

It is important to point out that the NAS VAAI plug-ins do not come with vSphere 5. This means your NAS vendor will have to develop and distribute this technology.

Storage Accelerator

A storage accelerator has been added, to be specifically used with VMware View (VDI) workloads. When this setting is configured in ESXi, a read cache is constructed in memory. This read cache is optimized for recognizing, handling, and deduplicating VDI client images and can result in a significant reduction in IOPs from ESXi hosts to the storage.

VMFS-5

VMFS-5 is arguably one of the best new features in vSphere 5. VMFS-5 is the latest version of VMware's Virtual Machine File System (VMFS). VMFS volumes can now be up to 64TB in size using a single extent, but it should be noted that a single VMDK is still limited to 2TB minus 512 bytes. If this VMDK size limit is considered bad news, then the good news is that pass-through (physical) RDMs larger than 2TB are supported.

Also absent in VMFS-5 are the different VMFS block size options, because a unified block size of 1MB is now used for newly created VMFS-5 volumes. Upgraded

VMFS-3 volumes will retain their given block sizes, but these volumes may be upgraded while VMs are running on them. Coupling VMFS-5 capacity with 32-way vCPU and 1TB memory support could lead to some pretty massive workloads running in vSphere 5!

vMotion and Metro vMotion

vMotion is certainly not new, but there are some new capabilities of vMotion in vSphere 5 that are worth mentioning. Performance enhancements made to vMotion now allow it to effectively saturate a 10 GbE network adapter's bandwidth. vSphere 5 also introduces the ability for vMotion to use multiple network adapters. Even a single vMotion will see the VMkernel transparently load-balance the traffic over all vMotion-enabled vmknics. These enhancements speed up vMotion and are especially beneficial for VMs with large memory footprints. There seems to be a theme here of vSphere 5 having the ability to support large Tier 1 workloads.

vSphere 5 also introduces the Metro vMotion feature. Metro vMotion is a latency-aware feature that allows vMotion to work with higher network latencies than were supported in previous versions of vSphere.

vCenter Server Appliance

The vCenter Server Appliance (vCSA) is a preconfigured Linux-based VM that has been optimized to run only vCenter Server and its associated services. It is deployed as an Open Virtualization Format (OVF)–formatted VM and is supported only on ESX/ESXi 4.0/4.1 and ESXi 5.

Limitations of the vCenter Server Appliance include no IPv6 support, no Linked Mode capability, and no use of Microsoft's SQL Server. The vCenter Server Appliance can be an appealing option to non-Windows environments but can also be used for smaller deployments because it includes a bundled database.

vSphere Storage Appliance

The vSphere Storage Appliance (VSA) is aimed at small-to-medium business (SMB) customers and provides a shared storage solution, without requiring these customers to purchase a SAN or NAS. The VSA can be deployed in a two-node or three-node configuration and

uses the local disks from the ESXi hosts to create a mirrored NFS share, as shown in Figure 1.6. The VSA allows features like High Availability (HA), Distributed Resource Scheduler (DRS), and vMotion to be used.

FIGURE 1.6 VSA cluster with two members

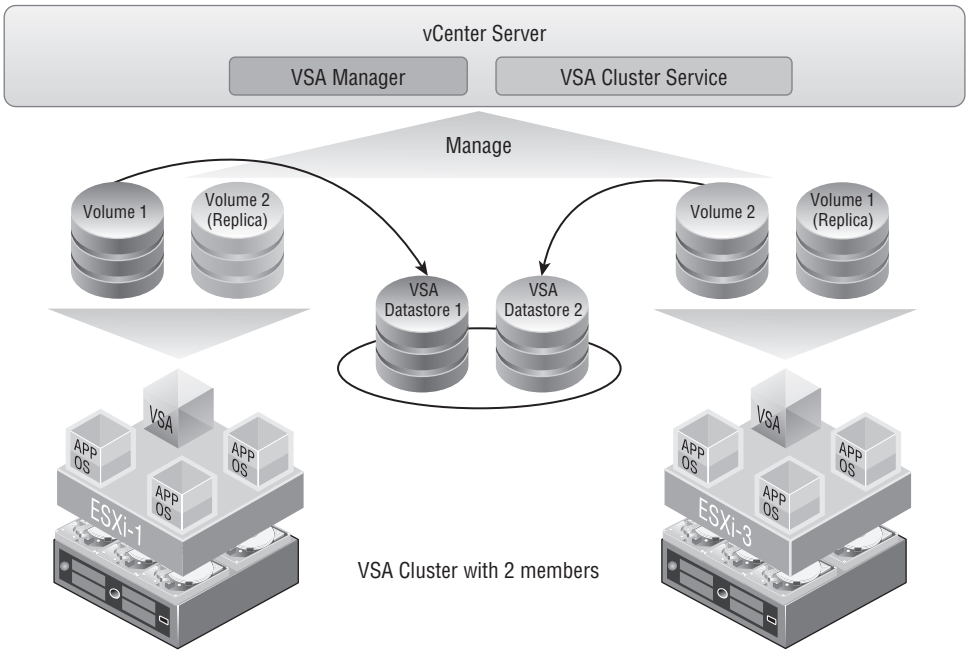

Storage DRS

This is essentially VMware's Distributed Resource Scheduler (DRS) for storage. Storage DRS, much like DRS, provides virtual machine placement and load-balancing mechanisms based on disk I/O and capacity. Storage DRS allows you to group similar datastores into a load-balanced storage cluster.

A new object called the *datastore cluster* is the basis of Storage DRS. A datastore cluster is simply an aggregated group of VMFS or NFS datastores (see Figure 1.7). A datastore cluster can contain either all VMFS or all NFS datastores. Mixing and matching VMFS and NFS is not allowed in the same datastore cluster.

FIGURE 1.7 Datastore cluster

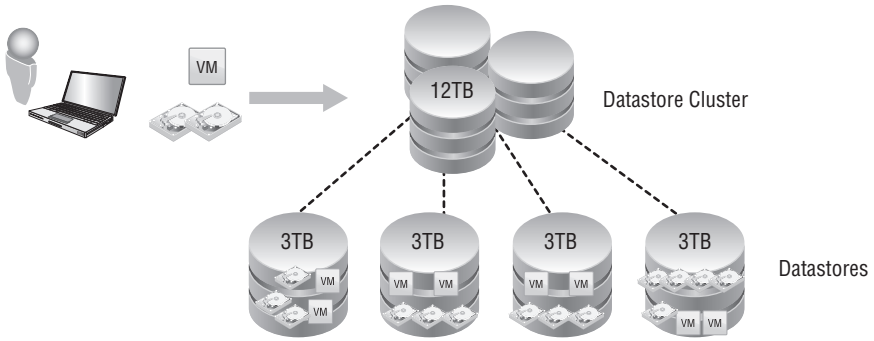

Profile-Driven Storage

Referred to as VM Storage Profiles in vCenter, profile-driven storage decreases the amount of administrator interaction required to deploy VMs. Instead of the VMware administrator having to know (or learn) a VM's particular disk requirements, a profile is created, and the VM is placed accordingly on the appropriate storage.

Image Builder

Image Builder is a new feature of vSphere 5; it is a PowerShell CLI command set that is used to manage software depots, image profiles, and VMware Installation Bundles (VIBs). These components, shown in Figure 1.8, are used to create and maintain ESXi images to be used during ESXi installations and upgrades.

FIGURE 1.8 Image Builder components

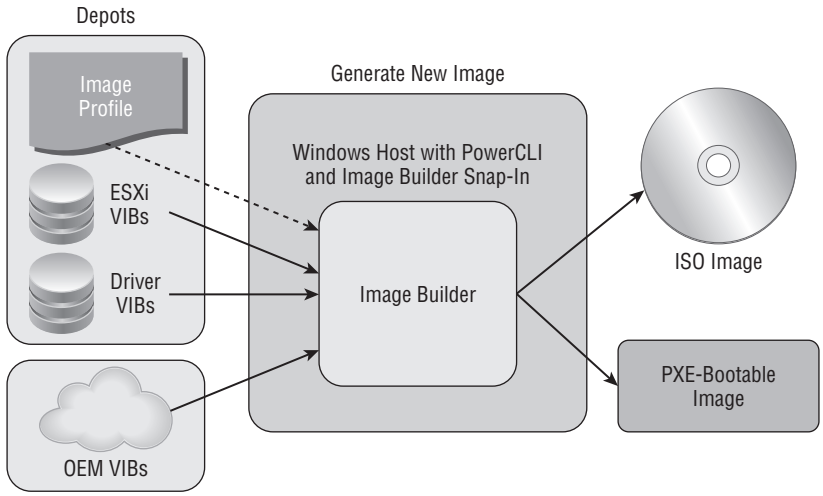

Auto Deploy

Auto Deploy is used to provision ESXi hosts, using a vCenter Server (see Figure 1.9). With Auto Deploy, the ESXi host will use the Preboot Execution Environment (PXE) at boot, and then vCenter Server will load the ESXi image directly into the host's memory. Auto Deploy stores no ESXi state on the host disk and allows diskless ESXi host configurations.

There are several benefits of using Auto Deploy, including the decoupling of the ESXi 5 host from the physical server. For example, if an ESXi 5 host fails and needs to be replaced, there is no manual ESXi 5 installation required when using Auto Deploy. Think about the manual steps required to install ESXi 5 and the configuration that has to happen. Using Auto Deploy, the ESXi host configurations will be consistent and significantly faster to provision. Another benefit of using Auto Deploy with diskless servers would be no RAID monitoring! Because Auto Deploy relies on vCenter Server, you should give serious consideration to providing high availability for vCenter Server when using Auto Deploy.

FIGURE 1.9 Auto Deploy architecture

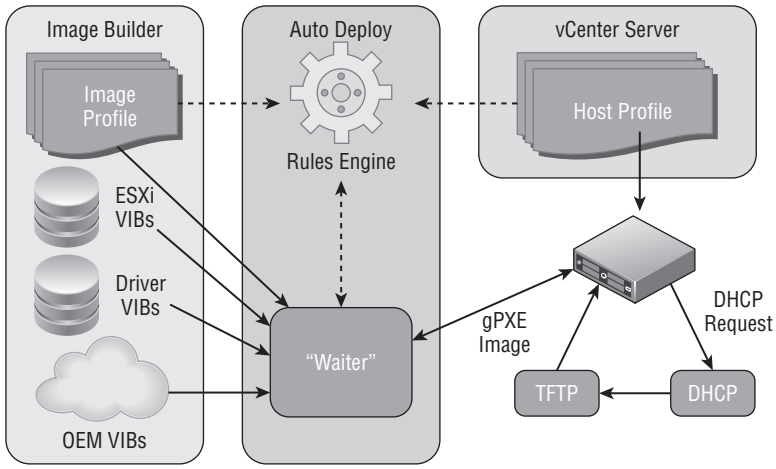

VMware Data Recovery 2.0 Improvements

VMware Data Recovery 2.0 has new capabilities that make it faster, more reliable, more easily managed, and better integrated with vCenter Server. These improvements include the following:

- You can generate and email scheduled reports to as many as 10 e-mail addresses. Reports include an application summary, job summary, destination summary, virtual machines that failed backup, virtual machines not backed up, virtual machines successfully backed up, and nonbackup warnings and errors.
- It includes user-specified Destination Maintenance Windows in which backups will not run.
- You can suspend individual jobs by right-clicking the backup job and selecting Suspend Future Tasks.

- You can change the number of days between automated integrity checks.
- You can specify the day of the week on which the automated integrity check is run.
- Deduplication improvements include a new compression algorithm to improve the commit speed of compressed data, utilization of the I/O path, and a more efficient layout of data on disk that reduces time to restore from backup.
- Improvements to the integrity check process include periodic checkpoints during integrity checking of the datastore and bulk processing of similar jobs.
- VMware Data Recovery 2.0 is built on CentOS 5.5 64-bit, allowing for better scalability and stability for the appliance.
- Linux swap partitions and Windows paging files are no longer included in backups.
- If an integrity check is stopped, the check can now be resumed without having to restart the entire process.
- Data Recovery is now more resilient against transient network failures.

Update Manager Improvements

vCenter Update Manager 5 has also been updated for vSphere 5. Improvements include the following:

- Support for ESX/ESXi 4.*x* to ESXi 5 upgrades
- Improved integration with vSphere clusters
- Improved VMware Tools upgrade
- Enhanced Update Manager download service
- Update Manager UI improvements

The Update Manager 5 release also includes the VMware vSphere Update Manager Utility. This utility helps users reconfigure the setup of Update Manager, change the database password and proxy authentication, re-register Update Manager with vCenter Server, and replace the SSL certificates used by Update Manager.

Licensing

There have been changes to the licensing in vSphere 5. Gone are the restrictions on physical core counts for processors and the amount of physical RAM installed in an ESXi 5 host. While the physical RAM limitations are gone, there is a new concept introduced called vRAM. vSphere 5 is still licensed per processor but will now include a vRAM entitlement for each processor license purchased. vRAM is the amount of memory configured to powered-on virtual machines, and it can be pooled across an environment.

Another change in vSphere 5 licensing is in the number of versions being offered. There are now five editions, or "kits," being offered:

- Essentials Kit
- Essentials Plus Kit
- Standard Acceleration Kit
- Enterprise Acceleration Kit
- Enterprise Plus Acceleration Kit

The Advanced Edition of vSphere 4 is gone, and customers who were licensed for this edition of vSphere 4 will be entitled to the vSphere 5 Enterprise Edition.

It is highly unlikely that the specifics of licensing will be covered in the VCP 5 exam, because the exam blueprint specifies only that users know how to license ESXi and vCenter Server. With that being said, understanding VMware licensing is a very beneficial thing for a VMware administrator to understand.

Summary

There have been many changes and new features introduced in vSphere 5. Some of these changes are around retired features like the ESX hypervisor architecture, VCB, and paravirtualization. These technologies served their purpose at one point in the VMware road map but have been replaced by improved technologies.

Many of the new features in vSphere 5 are targeted at being able to support large Tier 1 workloads in VMs. You can now have a 32-vCPU virtual machine with 1TB of RAM and a single disk of up to 64TB. That is a huge improvement over vSphere 4.1, where you could have an 8-vCPU virtual machine with 255GB of RAM and a single disk of up to 2TB minus 512 bytes. There also have been improvements to vMotion, Storage I/O Control, and Network I/O Control to accommodate these large workloads.

There are also many improvements in management included in vSphere 5. New features like Storage DRS, Profile-Driven Storage, and Auto Deploy will go a long way in making the VMware administrator's day-to-day tasks simpler. Improvements made to vSphere HA, Update Manager, and VMware Data Recovery will all help simplify the management of the virtual infrastructure.

There are also several exciting firsts for vSphere 5, like Mac OS X support, a vCenter Server Appliance that runs on Linux, and a cross-platform vSphere Web Client. Of course, with all these changes in feature sets, there also comes a new licensing model that you will need to get familiar with.

Although there are many changes to vSphere 5, many of which we will explore in the coming chapters of this book, there will be features that are not covered by the VCP5 exam blueprint. Thinking that these features omitted on the blueprint are not likely to show up on the exam would be a mistake. I encourage you to learn all of the vSphere 5 products and explore the feature sets as you move through this book. Learn how the products work and where they might effectively be used in your virtual infrastructure, but also be mindful of the "why" of these products and feature sets.