



Lunch Break

S E R I E S

Top 10 Reasons You'll Want to Upgrade to vSphere 5

By Elias Khnaser

Reason 10	1
Reason 9.....	2
Reason 8.....	2
Reason 7.....	3
Reason 6.....	4
Reason 5.....	4
Reason 4.....	5
Reason 3.....	5
Reason 2.....	5
Reason 1.....	6

SPONSORED BY

vmware®

VIRTUALIZATION
REVIEW



Top 10 Reasons You'll Want to Upgrade to vSphere 5

VMware's new vSphere 5 is packed with a great number of new features, but virtualization expert Elias Khnaser knows of 10 features that make this version upgrade-worthy.

VMware vSphere 5 is packed with great new features, many more than we are about to cover in this feature countdown. If you want to see the list from VMware, you can view it at

If we were to look at the whole list—200+—there are quite a large number of features that make this new version upgrade-worthy. It would be a painful read. Instead, I've chosen what I consider the top 10 features that I believe most people will be most

interested in. Of course, my choices are up for debate and you can write to me at editor@virtualizationreview.com to tell me what should also be on this list.

Let's kick off my list with a feature that is a separate product by name only, but really belongs in the list of new features for vSphere 5. By the way, this list is ranked by my order of importance—your list may differ.

10 VMware View Accelerator

At first glance one would think the VMware View Accelerator is a licensing package, since VMware currently has Accelerator packages. But make no mistake: The View

accelerator is one of the coolest features in vSphere 5 and since I am particularly focused and interested in desktop virtualization, I could not help it but start with this feature.

For years now, desktop virtualization technology has been plagued with different types of "storms," from bootup storms to login storms all the way through anti-virus storms and others. vSphere 5 takes a giant leap forward in addressing the issue of boot storms through what is now known as the View Accelerator.

Even though certain aspects of View Accelerator, such as the cache, are configured from the VMware View Composer, View Accelerator is really



more of a hypervisor feature, not a VMware View product feature.

The accelerator works by caching bits of the master image in memory on each ESXi host. When VMs start to boot, it redirects and de-duplicates VDI VMs to boot from those cached bits. This approach significantly reduces if not completely eliminates the boot-up storm issues. The number of VDI VMs to ESXi host is right about 60 VMs, which can very easily be managed by the local ESXi memory cache.


It is worth noting that Citrix XenServer has a similar feature known as IntelliCache. It basically does the same thing by caching the bits of the master image in memory and booting the VDI VMs from this local cache, thereby significantly reducing the IOPS dependency while maintaining centralized control and management.

9 AutoDeploy & Image Builder

Technically, AutoDeploy and Image Builder are two distinct features, but I find that they are very complimentary to one another, so I converged them into the ninth spot on my list.

Let's start with Image builder. Now that ESXi is the only version of the OS that is available (ESX has officially been retired), there is a lot more that goes into installing the hypervisor than was the case with ESX. ESXi has a smaller footprint and is a very thin install. As a result, drivers and software packages for the different hardware that you may be using is not baked into the image.

If you were to install the plain vanilla ESXi 5 image that you download from VMware, odds are some of your hardware won't work, such as network interface cards or Fibre Channel cards, because they don't have a corresponding driver. Image Builder allows you to customize the ESXi install image by



adding the necessary drivers, software packages and any other relevant software bits that are needed for the hardware you are deploying on. In a nutshell, Image Builder streamlines the installation of ESXi.

Now, AutoDeploy streams the installation of ESXi into the memory on the host, so you no longer need to install it on a USB drive, have it burnt on a chip on the motherboard or boot from SAN. Instead, the installation is streamed into memory. The way it works is by PXE booting a stateless host and downloading the assigned copy that is relevant to the hardware you are booting from. When a host's hardware fails, you can very easily swap that hardware for a new set of hardware, PXE boot again and in seconds you have restored that failed host back to productivity. So with AutoDeploy, the personality layer that is traditionally associated with a server is abolished, replaced by a memory-resident install of the OS that can be repeatedly downloaded to new hosts. And because ESXi has such a small footprint, it is easily loadable in memory.

By now you are probably wondering how AutoDeploy and Image Builder interconnect. Think about it this way: If you have IBM, HP and Cisco blades in your environment, you would use

Image Builder to create a customized ESXi install with all the necessary components for each hardware platform. Then you'd use AutoDeploy to automatically deploy the right ESXi install to the right hardware. As a result you have streamlined your host provisioning operation and automated the process significantly, thereby getting you one step closer to true private cloud computing.

The AutoDeploy feature bears a striking resemblance to Citrix Provisioning Services. PVS is an awesome feature and I believe that the VMware community will love AutoDeploy as much as the Citrix community loves PVS.

8 Profile-Driven Storage

Profile-Drive Storage is a new feature of vSphere 5 that allows you to easily profile the capabilities of your datastore in order to deploy your VMs on the most appropriate storage.

Today, most virtualization administrators gather a lot of information about what the VMs function will be and what kind of load it could potentially generate against the datastore. Is this a database server? Is this a Microsoft Exchange server? And so on. After that, we scramble to find



the most appropriate datastore. We later also try and follow up to make sure that those virtual disks have not moved from this datastore.

Profile-Driven Storage aims to address this process by simplifying it and automating it to some extent. You can now profile your datastore by manually adding some kind of an identifier, like a “tag” which would tag a datastore as a RAID-10 or RAID-5 datastore or some other characteristic. Once you do that, you can link VMs to this profile, thereby ensuring that the VM linked always exists on the right type of datastore. You can also link the VM to the right datastore during provisioning as well.

VMware has also introduced a new set of APIs with vSphere 5, known as vSphere Storage APIs for Array Awareness (VASA), which expose the capabilities of the storage to vCenter. In other words, you can see what type of datastore this is, what RAID level, etc. These APIs make it easier to profile the different types of storage.

PDS significantly reduces a lot of the manual work that one had to go through when provisioning VMs by automating it and allowing you to adhere to different levels of SLAs depending on the application profile of the VM. Not only are you assigning

the VM to the right storage, but you are also validating that the VM is where it is supposed to be.

Depending on the environment you are in and the level of process, SLs and automation you are required to have, the PDS feature may or may not be useful to you. In larger enterprises, I can see how it can be of significant help.

7 VM Capabilities

vSphere 5 is a platform upgrade and naturally there are feature enhancements to its VM capabilities. After all, this entire environment is built to host VMs. So what’s new and why do I put vSphere’s virtual machine capabilities at the seventh spot? Plenty, take a look:

First VM Hardware version 8 unlocks the capabilities of VMs running on the vSphere 5 platform. Upgrading your existing VMs’ virtual hardware gives you access to the following capabilities:

- ▶ Up to 32 vCPUs. Do you think you can tap all this vCPU power? Any excuse that applications need CPU processing power?
- ▶ Up to 1 million IOPS. That’s a lot of IOPS for a VM to handle; are there any apps that can’t run with that?
- ▶ Up to 1TB of virtual memory. I still don’t see physical servers with 1TB

of memory, but if there’s ever a need, now VMs can handle that load as well.

▶ Support for client connected USB. While I welcome this feature, it is limited to USB devices that are connected to the machine on which you are using the vSphere client or web client. So, you still cannot connect a USB in the ESXi server and pass it through.

▶ Support for USB 3.0.

▶ Non-hardware accelerated 3D capabilities. It’s a great enhancement for Windows Aero and VDI in general.

▶ Smart Card Readers. Similar to client connected USB devices, I welcome the addition of support for Smart Card readers. Again, it is limited to the host you are using the vSphere client and web client from.

And here are some other useful features:

▶ GUI for multicore vCPU configuration. It’s very similar to what you see in VMware Workstation where you can assign the number of vCPUs and vCPU cores.

▶ Apple Mac OS X Server guest operating system support. vSphere 5 now supports Mac OS X 10.6 Snow Leopard as a guest OS in a VM, albeit it is only supported on Apple Xserve3.1. Nonetheless, this is a very welcome step in the right direction.

▶ VMware Tools versions support matrix. New version of ESXi means a new version of VMware Tools. The cool thing here is that VMware Tools from version 4.x are supported on 5.x and vice versa, so you can have VMware Tools running version 5.x on ESX/ESXi hosts running vSphere 4.x

It is quite an impressive collection of capabilities that you can now assign to a virtual machine. Looking back, I would never have thought we would come such a long way with empowering VMs in such a short period of time. The excuse against virtualizing



certain types of applications and servers is becoming very slim, one would truly have to fight hard to make the case for a physical server after deploying vSphere 5.

6 ESXi Firewall, vSphere Web Client and Storage vMotion

Six is a party, with three features I put at a tie at this spot: ESXi Firewall, vSphere Web Client and Storage vMotion.

Let's start with Storage vMotion. Up until VMware vSphere 5 you were unable to use Storage vMotion if the VM had a snapshot attached to it. With vSphere 5, that limitation is lifted thereby allowing a storage vMotion operation to successfully complete. This makes a lot of sense as an enhancement in vSphere 5. Without it, vSphere 5 wouldn't be able to implement one of the most anticipated features: Storage DRS.

Also tied at number six is the new and enhanced vSphere Web client. In older vSphere versions, we had access to a Web-based administration interface to vCenter and the ESX/ESXi host. That interface, though, had limited functionality. vSphere 5 extends the reach of the vSphere Web client, allowing you to perform the vast majority of tasks that you could

perform from the regular vSphere client today. My only disappointment here is that the vSphere 5 Web client is not HTML5-compliant.

The third in this trinity-featured spot goes to the new ESXi Firewall. It's a new service-oriented, stateless firewall used to protect the management interface of ESXi. It also eliminates the reliance on IP tables and rule sets to define port rules for services. The firewall can be configured either from a GUI through the vSphere client or via the command-line using `esxcli`.

5 Swap to SSD, and VAAI

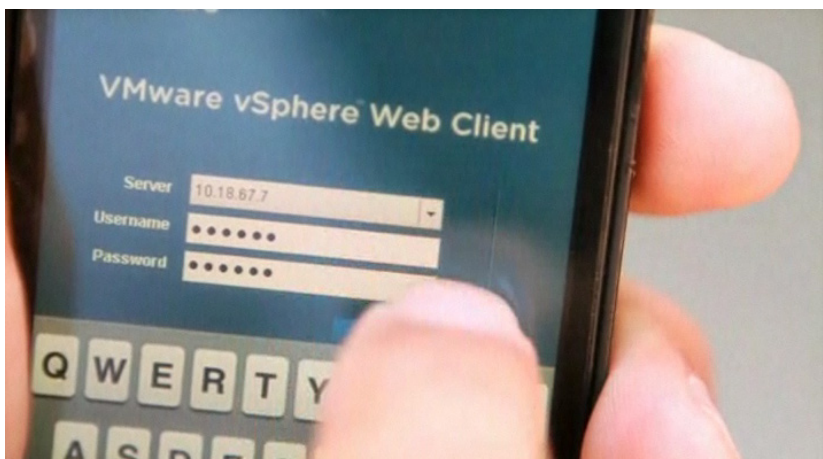
These two features are a tie at 5 on my list.

The Swap to SSD feature is of particular interest to me because of its overcommitment potential, especially as it relates to desktop virtualization. This cool new feature allows the VMkernel to automatically detect and tag any SSD storage that the ESXi host has access to, whether it is local storage or network storage. Once detected and tagged, the VMkernel will then instruct the scheduler to take advantage of this storage by allowing the ESXi swap to extend to it.

While I am not suggesting by any means that you use SSD as a replacement for memory (and it still is best practice to add physical RAM if you are swapping), this cool new feature does let you take advantage of SSDs in the amplification of memory overcommitment.

As for vSphere Storage APIs for Array Integration, or VAAI, it has always been one of my favorite storage features. With vSphere 5, it got better. One new capability I am particularly excited about is the automated free disk reclamation for thin provisioned disks.

Prior to vSphere 5, the process of reclaiming free disk space as a result of using thin provisioning was a relatively manual process. I am sure some of us have automated and scripted it, but I still consider this a manual process—





you still have to maintain the script. Another capability of VAAI I am excited about is that it now extends support for NFS; it was previously just supported on Fibre Channel.

And yet another cool capability of VAAI is known as NFS Space Reservation, which gives you the capability of reserving space on NFS. With NFS, one of the selling points is that it is thin-on-thin by design—the datastore grows and shrinks automatically. There are times I wished I could reserve all the underlying space and not necessarily have to do thin-on-thin. Now you can accomplish that with the new VAAI.

The last feature of the new VAAI I want to share with you is also around NFS and known as Full Copy, which can offload certain copy operations like “clone” or “deploy from template” from the host to the array. Now, I understand that certain NAS vendors had vCenter plugins that would allow you to do this from within vCenter. The enhancement with vSphere 5 is, now you can do it natively using the traditional vCenter commands, provided that your storage vendor can support it.

4 vCloud Director 1.5

Even though vCloud Director 1.5 is not necessarily a feature of vSphere 5, vCD relies heavily (not to mention that VMware made vCD a part of the general vSphere 5 announcement back in July).

vCD has three features that I like a lot. First off is, of course, Linked Clones. Virtualization introduced us to fast provisioning of VMs, but it was at the expense of cloning the entire disk, resulting in the unnecessary duplication of hundreds of Windows operating systems. Linked Clones, as with VMware Workstation and VMware View, allow you to provision VMs faster by not necessarily performing a full clone, but rather what’s called a “fast

clone.” This means you still have reliance on the master image, but you commit your changes to a difference file. This saves significantly on storage costs, de-duplication software and more.

Something tells me that in the next update of vSphere 5, maybe update 1, we will see Linked Clones in vSphere as well, as vSphere seems to be the only tool that lacks this Linked Clones functionality.

I also really like vCD Messages. In all of my deployments, executive IT always wants some type of visibility into the environment. They want to know by simply looking at a console or a monitor that something is going on in the system, and they want to know when it is fixed. vCD Messages is a perfect companion in this scenario. It will tie into your existing CMDB or ticketing system and can display notifications and messages. vCD comes bundled with up to 100 tasks out of the box—pretty cool!

And finally, I also like that vCD has good integration with vShield. It’s a major enhancement and now supports:

- ▶ Five-Tuple Firewall services, which means you can now build your firewalls based on destination IP, destination port, protocol, source IP and source port

- ▶ IPSec VPN Services allow the user to establish site-to-site VPN tunnels. This is pretty powerful and unique to vCD in that as a customer or user of vCD you can create this site-to-site yourself. It’s completely self-service and pretty cool.

To wrap it all up, vCD 1.5 now has support for Microsoft SQL Server, which means you don’t need an Oracle license anymore to deploy it. SQL Server support is a major enhancement, as most of my customers were hesitant about vCD because they did not want to introduce Oracle and Oracle licensing into their environments.

3 VMFS-5 at #3

I’m a big fan of VMFS. So, it’s no surprise that VMFS-5 is third on my list.

VMFS is a purpose-built file system optimized to run virtual machines. One of the big drawbacks of VMFS-3 was its size limitation of 2TB. With the introduction of the vSphere APIs for Array Integration, it was simply a matter of time before VMware lifted that limitation. VAAI has several useful features, but one in particular is more relevant to our current discussion: VAAI offloads SCSI reservation from the hypervisor onto the storage array, thereby significantly enhancing the performance of the VM in general.

VMFS-5 now supports datastores of up to 64TB in size, thereby stripping NFS from its only strategic advantage as far as I am concerned. The other feature enhancement that is now enabled within VMFS-5 is the ability to leverage VAAI with thin provisioning and deliver automatic free space reclamation. This is huge. Prior to this version of VMFS, many did not know that when you use thin provisioning and as you start deleting data, you’d need to go back and reclaim that free space. Well, the bigger problem was that it had to be done manually. And so for those who did know about the free space that could be reclaimed probably had some manual methodologies or scripts that reclaimed the free space.

Now, you will be happy to know that VMFS-5 automates the task, thus removing yet another annoying task from your to-do list.

2 High Availability

High Availability has been completely re-architected for the better—cosmetically HA looks and feels the same to the virtualization admin, but



there have been significant changes under the hood.

First, HA no longer requires DNS to properly function. In earlier versions, this was a gotcha that many neglected and it indeed caused issues. That being said, my favorite new enhancement to HA has got to be the fact that it no longer relies on just the management network to determine if a host has failed. The new HA in vSphere 5 will check the management network. If it detects a failure, it will then use the storage connection as a secondary way of verifying if the host really failed. If it cannot contact the host using either method, only then will it deem the host as failed and HA will kick in.

There are other changes to the way HA works, like the fact there no longer is a primary/secondary relationship between the hosts in the cluster. Instead, there's a single master server which is determined via an election process. All other hosts in the cluster are slave hosts. Elections occur at different intervals, but it occurs primarily when a new host is added or if a host is rebooted. There are other things that can also trigger an election but these two events are in the forefront.

1 Storage DRS

Storage DRS is a fantastic new feature—and number 1 in my book—that load balances virtual machines across the different datastores to optimize performance.

Up until vSphere 5, DRS was a cluster function limited to load balancing compute resources (CPU and memory), and that is a great feature because as VMs start to consume a lot of physical resources on a particular ESXi host, DRS would instruct vMotion to migrate VMs to a more suitable host that can better satisfy their resource needs. In the beginning this was

enough—we were just virtualizing some servers and some of us were testing and using it in development environments.

Now that virtualization has gone mainstream in production and the consolidation ratios have increased, it has become increasingly more difficult to load balance VMs across datastores. Moreover, those of us that were successful in an initial deployment of properly load balancing VMs across datastores recognized quickly that maintaining this load balancing was not easy.

Storage DRS now increases the functionality of our clusters by extending VM load balancing across the different datastore leveraging Storage vMotion. As a result of this immensely helpful feature, it's number 1 on my vSphere 5 list. [VR](#)

Elias Khnaser owns professional services firm The Elias Khnaser Company and is an expert on VMware ESX, Microsoft Hyper-V and Citrix XenServer. Elias has authored many books and published a number of training videos on the subjects of virtualization and server based computing, primarily Citrix and Terminal Services. He has blogged for Forbes and eWeek.

