



What's New in VMware vSphere® Storage Appliance 5.1

TECHNICAL MARKETING DOCUMENTATION
V 1.0/UPDATED JULY 2012

Table of Contents

Introduction	3
Support for Additional Disk Drives	3
Maximum Storage per ESXi Host	4
Increase Storage Capacity Online	4
Remote Office/Branch Office (ROBO) Support.	6
Two-Node Clusters and the VSA Cluster Service in ROBO	7
Improved Security	7
ROBO Install and Upgrade	8
vCenter Server Running on the VSA Cluster	8
vCenter Server Running on the VSA Cluster – Considerations	9
Brownfield Installation of the VSA Cluster	9
Performance Improvements	10
Faster Synchronization Times	10
Synchronization Serialization	10
Reconfigure Network Workflow Changes	10
VMware vSphere 5.1 Enhancements	11
vSphere 5.1 Support	11
Memory Overcommitment Restriction Removed	11
Miscellaneous Enhancements	11
Conclusion	12
About the Author	12

Introduction

The goals of the VMware vSphere® Storage Appliance (VSA) 5.1 version release are twofold: The first is to enhance the current offering to fully address small to midsized business use cases; the second is to support adjacent markets such as remote office/branch office (ROBO).

The primary enhancements in VSA 5.1 are as follows:

- VSA 5.1 has the ability to have more than eight disk drives per VMware® ESXi™ host.
- VSA 5.1 has the ability to select how much local disk space to allocate to the VSA as shared storage, as well as the ability to dynamically increase the storage used by the VSA.
- A single VMware vCenter Server™ instance now can manage multiple VSA instances.
- The vCenter Server instance managing VSA now can route to the ESXi hosts participating in the VSA storage cluster. There no longer is a requirement to have the vCenter Server instance and ESXi hosts on the same subnet.
- VSA 5.1 enables deployment on ESXi hosts that already have been configured and have running virtual machines.

The remainder of this *What's New in VMware vSphere Storage Appliance 5.1* white paper will discuss each of these enhancements in greater detail. It will also review additional, minor enhancements to the product.

Support for Additional Disk Drives

In VSA 1.0, each ESXi host had the capacity for the following number of disks in a RAID 10 configuration (RAID 5 and RAID 6 maximums were slightly different):

- 4, 6 or 8 x 2TB disk drives
- 4 x 3TB disk drives

In VSA 5.1, VMware increases the number of disks per ESXi host to the following:

- 8 x 3TB disk drives

The following points apply to disk sizes that are 2TB or less in size:

- 12 internal disk drives per host
- External expansion chassis with up to 16 drive bays (external expansion chassis is cabled to RAID controller within the host)
- Implies support for up to 28 drives per host

If the user employs disks in both internal drive bays and external expansion chassis, the configuration still will be a single VMware vSphere® VMFS-5 volume using multiple extents. VMware is making a best practice recommendation to have the same RAID configuration on both internal and external disks. This means that the performance doesn't change depending on where the virtual machine is deployed on the VMFS-5 volume.

Maximum Storage per ESXi Host

In VSA 5.1, the maximum storage supported is as follows:

3TB drives

- 8 disks of up to 3TB capacity in a RAID 6 configuration (no hot spare)
- 18TB usable by the VMFS-5 file system per host
- Across three hosts, a total business usable storage of 27TB

2TB drives

- 12 local disks of up to 2TB in a RAID 6 configuration (no hot spare)
- 16 external disks of up to 2TB in a RAID 6 configuration (with hot spare)
- VMware supports maximum VMFS-5 size of 24TB per host in VSA 5.1
- Across three hosts, total business-usable storage of 36TB

The VSA installer automatically adjusts the VMFS heap size on each ESXi host to handle the larger file system sizes.

Increase Storage Capacity Online

In VSA 1.0, users needed to determine how much capacity was required in the cluster at deployment time. There was no supported way to increase the disk capacity of the cluster after it was deployed.

In VSA 5.1, VMware supports the online addition of new disk drives to an ESXi host. This also introduces a user interface mechanism to expand the current VSA datastores to include additional space enabled by the new, larger disks.

Before increasing the VSA shared storage, administrators might first have to introduce new physical storage to the ESXi hosts. VMware supports users' introducing new physical disk space on VSA nodes using the following methodologies:

1. Convert existing RAID 10 configuration to RAID 5/6.

In the initial release of VSA 1.0, VMware required RAID 10 configurations. More recently, this restriction was relaxed to allow both RAID 5 and RAID 6 as well as RAID 10. Users might be interested in changing their RAID 10 configurations to RAID 5 or RAID 6 to provide additional disk space. This method destroys existing data on disk. However, this data will be reconstructed from the surviving replica copies of the datastore.

2. Add disks to existing RAID set.

Users initially can deploy VSA with four physical disks in a host in a RAID 5 configuration. To introduce additional disks, they can add another four physical disks per host, remove the current RAID 5 configuration and create a new RAID 5 configuration using eight disks. This method also destroys existing data on disk.

3. Add a new RAID set, preserving on-disk data.

This final method preserves the data on the existing RAID set. After the new RAID set is created, it is presented to the ESXi host as a new LUN. The existing VMFS-5 volume then can be expanded by adding this new LUN as an extent. This method preserves the existing ESXi install and datastore, and it must be repeated on all ESXi hosts in the VSA cluster.

As noted, both procedures 1 and 2 destroy on-disk data on the ESXi host/VSA node. This entails a reinstall of the ESXi. Because the data is preserved on mirror replicas on other nodes in the VSA cluster, administrators can synchronize it onto the newly created datastores using the *replace node* procedure by implementing it on one node at a time. Steps include 1) removing the node via maintenance mode, 2) adding the new physical storage, and 3) running the *replace node* procedure to put the node back in the cluster. Users then must wait for resynchronization of the data and repeat these steps on additional hosts.

After the physical storage has been added to the hosts, and the VMFS-5 volume on the ESXi hosts has been expanded, a new **Increase Storage** option in VSA Manager can be used to enlarge the shared storage of the VSA, as shown in the following screenshot:

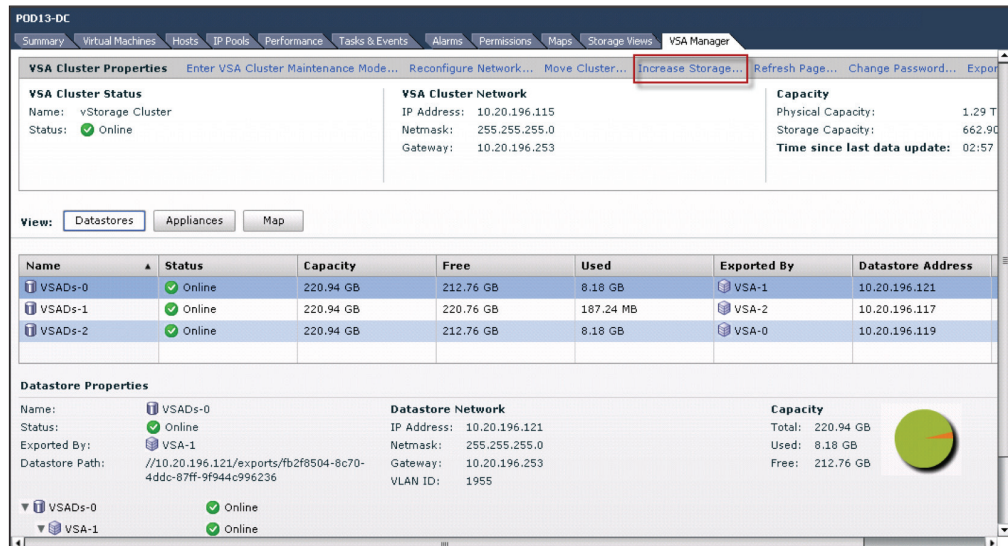


Figure 1. Increase Storage Option

In the following example, it appears that the VMFS-5 volume on each ESXi host is being increased by approximately 150GB and that shared storage is being enlarged by approximately 100GB, leaving approximately 50GB free on each host.

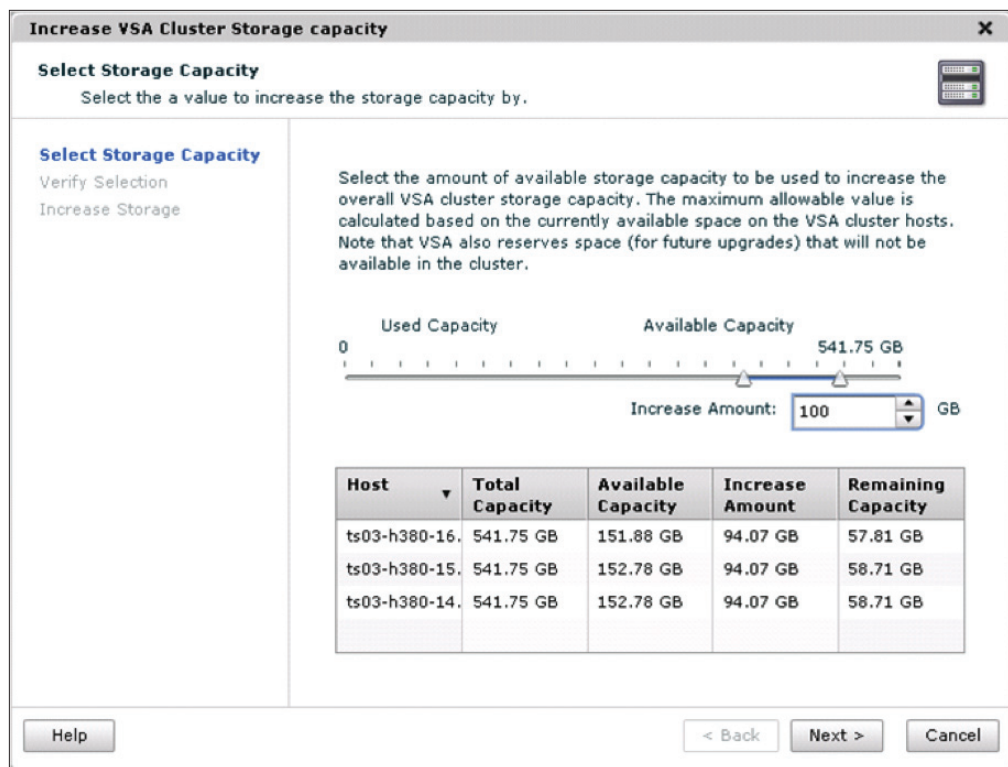


Figure 2. Increase Storage Example

The following is one example of the steps involved in increasing storage capacity online:

1. Each node in the VSA cluster is initially deployed with four physical disks.
2. The administrator now wishes to increase the shared storage capacity, so an additional four physical disks are added to each ESXi host, producing a total of eight per host.
3. The administrator now can either *recreate* the whole RAID set on all eight physical disks, destroying the original RAID set and any datastores on that RAID set, or decide to *preserve* the original RAID set by creating a new RAID set on only the newly introduced disks.
4. Place the VSA and host in maintenance mode. Bring down the host so that the RAID controller configuration menu can be accessed. RAID configuration steps vary from server vendor to server vendor and are beyond the scope of this paper.

4.1. Recreate:

- 4.1.2. Recreate the RAID volume to cover all eight disks.
- 4.1.3. Reinstall the core ESXi packages, which also creates a new VMFS-5 volume.
- 4.1.4. The administrator now can do a *replace host* operation via the VSA Manager user interface. This brings the ESXi host back into the cluster, redeploys the VSA and synchronizes it with its mirror. At this stage, the VSA still has the original shared storage size. It hasn't yet been expanded to use the new storage.

4.2. Preserve:

- 4.2.1. If the original RAID set has been preserved and a new RAID set has been added (using the four new disks), the VMFS-5 volume must be expanded to consume this new storage via the addition of an extent. This procedure is discussed in the *vSphere Administration Guide*. When complete, this host now has a larger VMFS-5 volume than before.
5. Take the host out of *maintenance mode* and allow the mirrors to synchronize.
6. When all of the hosts have new, larger VMFS-5 volumes, the VSA appliances can expand using the free space on the VMFS-5 volume. This is done via a new **Increase Storage** wizard in the VSA Manager UI. It involves automatically enlarging the virtual machine disk files on the appliances, modifying volume sizes on the appliances and resyncing. Appliances reboot as part of this process.
7. After the appliances are rebooted, the NFS exports from appliances are updated to present the larger size, and the ESXi hosts mounting these datastores automatically pick up the new datastore size.

Remote Office/Branch Office (ROBO) Support

VSA 5.1 now enables multiple VSA clusters to be managed by a single remote vCenter Server instance. The vCenter Server instance can also reside on a different subnet from the VSA cluster in this new release. Both of these were limiting factors in VSA 1.0. VSA 5.1 now becomes an attractive shared storage product for ROBO.

Each VSA storage cluster is located in its own unique datacenter object in the vCenter Server inventory. The following is a screenshot of a two-node cluster and a three-node cluster, both managed by a single VSA Manager on a single vCenter Server instance.

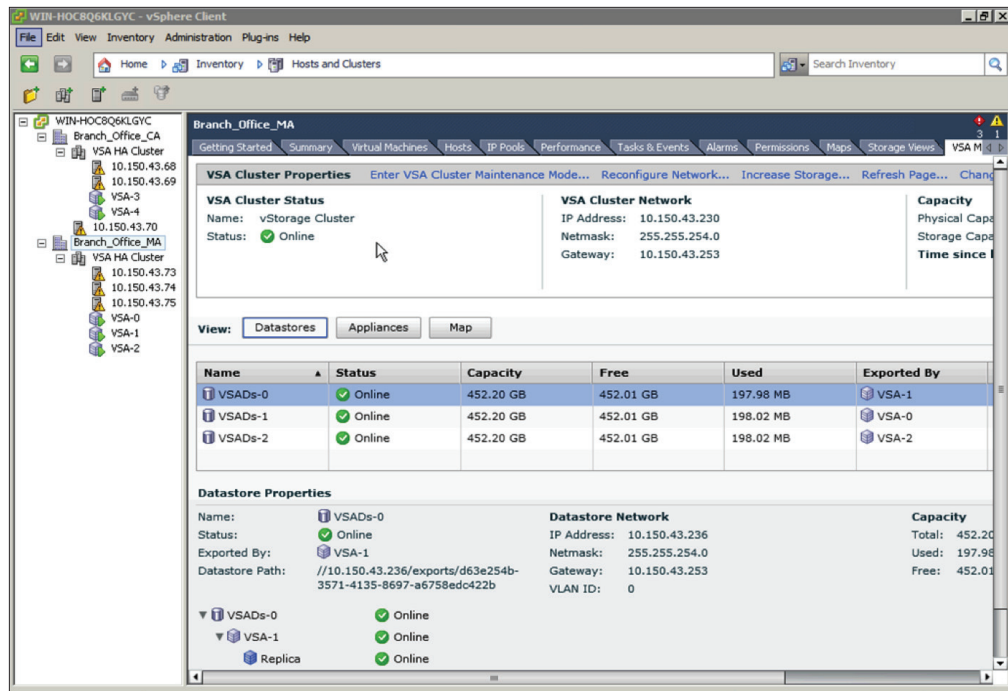


Figure 3. ROBO Support

VMware targets the management of at least 150 VSA 5.1 storage clusters by a single vCenter Server instance.

Two-Node Clusters and the VSA Cluster Service in ROBO

The two-node VSA storage configuration uses a special “VSA Cluster Service,” which typically runs on the vCenter Server instance. It performs like a cluster member and is used as a tiebreaker to ensure that there still would be a majority of members in the cluster if one ESXi host/VSA member were to fail. The new designed tiebreaker for two-node configurations still can run as a vCenter Service in VSA 5.1. However, in a two-node ROBO deployment, the vCenter Server instance managing VSA is now remote. This means that the VSA tiebreaker code must be *locally located at the branch office* and not on the central vCenter Server instance.

In VSA 5.1, the tiebreaker code can be installed at the branch office. The administrator provides details on the location of tiebreaker code during VSA deployment, and the VSA installer validates these settings. The tiebreaker is simply a set of Java Archive (JAR) files. The JAR is a way of distributing a Java program, along with all its libraries.

VMware provides installers to run the JAR file on a user-supplied platform, either Windows or Linux. The administrator is responsible for configuring the platform, including installing the base operating system (OS). The installer then introduces and sets up the tiebreaker code. VMware provides installation documentation for all the platforms that support running the tiebreaker code.

Improved Security

Because the VSA 5.1 now supports management from a remote vCenter Server instance, security has been improved in a number of places to prevent spoofing and so on. User-visible changes include a certificate exchange at cluster creation time. Nonvisible changes include using secure communication protocols such as HTTPS and SSL.

ROBO Install and Upgrade

Supporting ROBO environments complicates the existing installation process. VSA 5.1 now supports three new installation/upgrade use cases in addition to those that existed in VSA 1.0.

1. Unattended remote installation: A VSA 1.0 cluster at a remote office is upgraded from the central office over the network. This might not be viable if there is limited bandwidth between the central office and the remote offices. This requires no intervention at the remote office.
2. Attended local installation: An administrator located at the remote office does the installation or upgrade of the VSA, using installation media at the remote office.
3. Offline installation: In this option, the VSA 5.1 cluster is configured at a central office, boxed up and shipped to a remote office where it is added back into the central office vCenter Server instance via the **Reconfigure Network** workflow.

vCenter Server Running on the VSA Cluster

In VSA 1.0, VMware did not support the installation of the vCenter Server instance used by the VSA Manager, either on the ESXi hosts participating in the cluster or on the NFS datastore presented by the VSA cluster.

In VSA 5.1, one supported option is to install a vCenter Server instance in a virtual machine on a local datastore on one of the nodes in a VSA storage cluster. The vCenter Server instance then can be used to install VSA by allocating a subset of local storage, excluding the amount allocated for vCenter Server (on all hosts) for VSA.

To do this, we must ensure that the VSA appliances do not consume all of the local VMFS-5 storage. This is achieved by modifying the installation wizard so it now calculates how much capacity is available. The new installation wizard now allows a certain amount of local VMFS-5 space to be given to virtual machines already running on the ESXi hosts, including vCenter Server, as illustrated in the following figure:

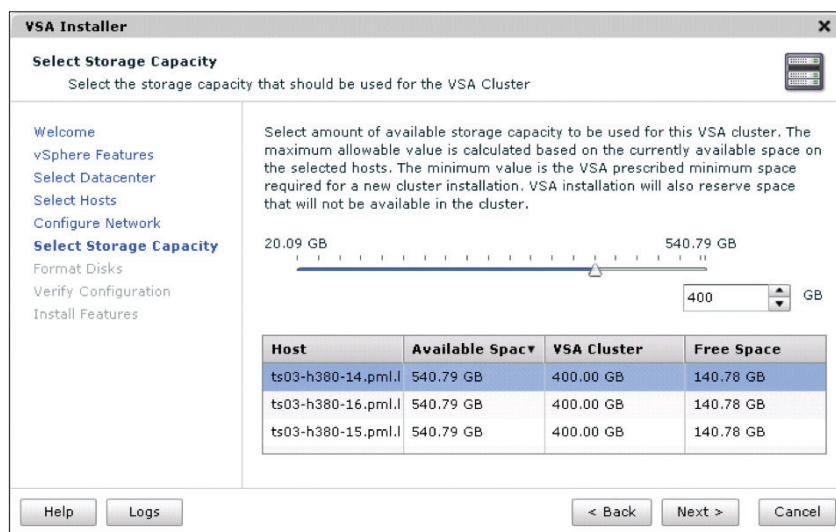


Figure 4. Select Storage Capacity

When shared storage is created, VSA 5.1 also supports the migration of vCenter Server from local storage to shared storage. Shared storage now can be resized with the reclaimed local storage passed on to the VSA. As discussed previously, this can be done via the user interface.

vCenter Server Running on the VSA Cluster – Considerations

The following points should be considered regarding support of a vCenter Server instance running in a VSA cluster.

1. A vCenter Server instance running in a virtual machine is supported on a three-node cluster configuration because it is not involved in tiebreaker situations.
2. A vCenter Server instance running in a virtual machine is supported on a two-node cluster configuration when the VSA Cluster Service is running *outside* of the VSA storage cluster (on an external system or GuruPlug).
3. A vCenter Server instance running in a virtual machine is not supported on a two-node cluster configuration when the VSA Cluster Service is running on the vCenter Server instance itself.
 - In fact, the tiebreaker cannot run on either of the VSA server's ESXi hosts, whether using local storage or shared storage.
4. If a vCenter Server instance running in a virtual machine is also running the VSA Cluster Service, it must not run on the VSA shared storage, even if it is running on an ESXi server that is not a VSA cluster node.

Brownfield Installation of the VSA Cluster

This is a feature very similar to a vCenter Server instance running on local storage of one of the ESXi hosts in the cluster. VSA 5.1 enables vCenter Server to be installed in a virtual machine running on the local datastores of one of the nodes. In VSA 1.0, VMware required that a baseline/new version of ESXi 5.0 be installed on the two or three nodes (what was called a greenfield installation). Consider the case where there are already virtual machines running on local storage in a user environment. This enhancement enables VSA 5.1 to be installed on two or three of these ESXi hosts, even when virtual machines already are deployed on the local storage of the hosts. This is what is referred to as a brownfield installation.

For brownfield installations, the user might already have created vSwitches. If this is the case, the wizard audits the configuration and fails it if it is not set up correctly. The user is responsible for fixing the configuration. For greenfield installations, or brownfield deployments that do not have preconfigured vSwitches, the wizard configures the vSwitches appropriately.

With brownfield installations, ESXi hosts must have their heap size increased to handle the larger VMFS-5 volume sizes now supported in VSA 5.1. This change is automatically implemented by the VSA installer but requires a reboot of the ESXi hosts. This is performed in a rolling manner so the cluster remains online. However, it still is strongly recommended that it be performed within a maintenance window.

The VSA 5.1 installer uses the free space remaining on the local VMFS-5 file system and hands this space off to the VSA appliances. After the shared storage has been configured, the virtual machines running on local storage can be migrated to the shared storage, using either cold migration (virtual machines powered off) or VMware vSphere® Storage vMotion® if the user has a license for this feature.

Using the **Increase Storage** feature discussed previously, the disk space that was occupied by the virtual machines on local storage now can be reclaimed by the VSA.

Performance Improvements

Faster Synchronization Times

The multiple device (MD) RAID 1 synchronization algorithm in the VSA appliances was changed to provide fairer and more controlled resource sharing between the background RAID 1 synchronization occurring across appliances and the foreground application-initiated I/O. Background I/O is given a larger time window to do sequential reads without interruption, provided that this does not violate the desired foreground I/O level. This effects a threefold improvement of sync speed in the presence of foreground I/O. As a result, VSA datastores remain degraded for a shorter period of time and there is no need to pause foreground I/O for a full MD synchronization. Without foreground I/O, sync speed is approximately the same.

Synchronization Serialization

Synchronization operations on two mirrors (exported to different volumes) on a single node utilize the same set of disk drives and thereby compete for resources on read/write operations. In this scenario, the operation is serialized, eliminating contention. The RAID mirror with less data to synchronize is granted permission first; the synchronization speed on the other is throttled down, and this second mirror must wait until the first one completes synchronization or is preempted if it is already syncing. When mirror one has finished syncing, mirror two is scheduled to synchronize.

Reconfigure Network Workflow Changes

Several changes have been made to the "Reconfigure the VSA Cluster Network" workflow to improve the user experience.

In VSA 1.0, during a **Reconfigure Network** operation, the cluster was put into maintenance mode. Then the user was asked to enter the new IP addresses. In VSA 5.1, first the user is asked to enter the new IP addresses. Then, before the cluster is put into maintenance mode, a check for any conflicts is conducted.

In VSA 1.0, all IP addresses related to the VSA node had to be changed. In VSA 5.1, the user can select a subset of IP addresses to change, and only those IP addresses will be changed.

In VSA 1.0, the **Reconfigure Network** wizard window to enter IP addresses was blank. After the **Cluster Management** IP address was populated, the remainder of the IP addresses were autopopulated. In VSA 5.1, the same window is prepopulated with the current IP addresses. This enables the user to change only the relevant IP addresses.

There also were significant changes on the back end to improve logging and the flow of the **Reconfigure Network** process. The user-facing aspect of the back-end changes provides better support for running the **Reconfigure Network** operation multiple times in the event of a transient failure.

The **Reconfigure Network** wizard can be run repeatedly as IP address conflicts or other network configuration issues are resolved. The overall robustness of the workflow has increased significantly.

VMware vSphere 5.1 Enhancements

There are a number of enhancements to VSA 5.1 that are related specifically to the VMware vSphere® 5.1 release.

vSphere 5.1 Support

VSA 5.1 is supported on vSphere 5.0 *and* vSphere 5.1. VSA 5.1 also supports the single sign-on (SSO) functionality found in vSphere 5.1.

Memory Overcommitment Restriction Removed

VSA 1.0 had a restriction that did not allow the memory overcommitment of virtual machines running on VSA datastores. There were recommendations in the administration guide that stated that an administrator should 1) reserve all memory for virtual machines using VSA datastores and 2) disable VMX swapping for all virtual machines using VSA datastores. Memory overcommitment now is supported for VSA v5.1 on ESXi 5.1. There are no new restrictions beyond those that already exist for ESXi 5.1 memory overcommitment.

Miscellaneous Enhancements

1. Node replacement for a three-member cluster no longer requires a temporary vCenter Server standard license.
2. When new hosts are added to an existing VSA datacenter, the access control lists (ACLs) are automatically updated and the NFS shares are automatically mounted on those hosts. When an ESXi server is removed from the datacenter, the ACLs are removed, but the datastores are not automatically unmounted because the vCenter Server instance no longer can access the ESXi server to do so.
3. VSA now can be installed on a vCenter Server instance that has IPv6 enabled.
4. MemberOffline, StorageOffline and ClusterOffline events now generate alarms that are sent to vCenter Server.











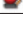
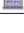
Triggered Alarms				
Object	Status	Name	Defined In	Triggered
 VSADs-0	 Alert	 VSA Storage Entity VSADs-0 Offline	 VSADs-0	3/30/2012 12:19:42 PM
 VSA-0	 Alert	 VSA Member VSA-0 Offline	 VSA-0	3/30/2012 12:19:40 PM
 st	 Alert	 VSA Cluster Service for VSA Cluster in Datacenter: st Offline	 st	3/30/2012 12:22:21 PM

Figure 5. VSA Events Surfacing as vCenter Alarms

Conclusion

A number of significant features have been added to the VMware vSphere Storage Appliance 5.1 version release, many as a direct result of user feedback. VMware has removed restrictions that were in place in VSA 1.0 by providing several enhancements. These include increasing the storage capacity after VSA deployment, offering more disks per host and enabling the vCenter Service instance managing VSA to run on the VSA cluster. The ability for a single vCenter Server instance to manage multiple VSA deployments, as well as the fact that the vCenter Server instance no longer must be on the same subnet as the ESXi hosts participating as VSA cluster nodes, makes the VSA an excellent fit for the ROBO market.

About the Author

Cormac Hogan is a senior technical marketing architect in the Cloud Infrastructure Product Marketing group at VMware. His focus is on virtual storage in general as well as core VMware vSphere storage technologies in particular, including the VMware vSphere Storage Appliance. He has been in technical marketing since 2011.

- Follow Cormac's blogs at <http://blogs.vmware.com/vsphere/storage>.
- Follow Cormac on Twitter: [@VMwareStorage](https://twitter.com/VMwareStorage).

