











HC3[®] vs. VMware's VSAN

VMware's Virtual SAN (VSAN) was released into public beta in vSphere 5.5 and will presumably be made generally available in early 2014. VMware's VSAN utilizes the local storage attached to each host to create a shared datastore presented to all participating hosts, thus "converging" the servers and storage in the infrastructure. The new functionality brings with it a spotlight on the Hyper-Convergence space where HC3 has seen such wide adoption. This document is intended to highlight the differences between HC3 and VMware's beta version of VSAN.

The Purchasing and Setup Process: *HC3 is easier to buy and implement*

With VMware's VSAN, customers are required to select their own servers. In each server, there is a minimum of one internal Solid State Disk (suggested to be at least 10% of the size of the HDDs) and one Hard Drive from VMware's limited Hardware Compatibility List. With the hardware in hand, it is left to the user to install the ESX hypervisor, install vCenter (and a SQL Database), activate the appropriate features, create a vCenter cluster and activate VSAN as a datastore.

HC3 is sold as a fully integrated system. This means that all software licenses are included, with balanced compute and storage hardware, ready to use out of the box, and designed for scale out expansion. HC3 customers are able to go from unboxing to running highly available Virtual Machines in a matter of hours with no prior training on the hypervisor or a SAN/NAS!

Feature	HC3	VSAN
High Availability, Live Migration without the use of Storage Protocols (iSCSI, NFS)		
Resiliency Against Multiple Hardware Failures		
Seamless Storage Integration with Hypervisor		
Scalability of Performance and Capacity (Scale Out)		
High Performance without SSD Caching		
Independence from Licensing Fees associated with Hypervisor		

Infrastructure Management: *HC3 is easier manage*

VMware requires a separate vCenter Server or virtual appliance to manage the Virtual Machine hosts and the VSAN. This functionality has to be setup by the user and can use the VSAN storage as its datastore, but requires a very specific set of steps to deploy in this manner. Once the vCenter Server Appliance has been configured, users can finally login to the vSphere Web Client to manage the environment.

HC3 fully integrates cluster-wide management and storage redundancy, which is accessible using any browser supporting HTML5. Users simply point their browser to any of the IP addresses associated with the cluster and they are able administer their environment from within the Scale GUI.

Expansion: *HC3 is easier expand*

Expanding a VSAN environment in the vSphere 5.5 version requires that the new host have only components from the VMware vSphere 5.5 HCL. It is then on the user to manually install the hypervisor and add the hosts to the cluster before their local storage can be contributed to an existing VSAN datastore.

With HC3, users have the ability to mix and match node types, which can include various RAM footprints, CPU types and drive types depending on the model selected. Initialize a new node in the environment and Scale's ICOS automatically aggregates that node's resources into the cluster, which in turn gives existing workloads immediate storage performance by increasing the number of spindles available to write data. Adding additional VSAN nodes does not improve the performance of existing VMs.

Support: One throat to choke

While VSAN eliminates the need for a separate SAN or NAS in the environment, it does not eliminate the need for multiple vendors. Users are required to select their own hardware from the VMware HCL to cobble together their converged infrastructure. In the event of an issue, it is on the user to determine who to call for support. Even then, without strict adherence to the HCL as hardware and software components are updated independently, customers could be told that they are unsupported.

HC3 includes 1 year of ScaleCare support in the price of the nodes, which covers not only the software, but also the HC3 hardware as well. As new firmware versions are released, HC3 users covered under ScaleCare are able to upgrade the hardware and software components in an integrated fashion without any disruption to their running VMs. Scale's converged model of support gives HC3 users the assurance that they are fully covered with no potential of finger pointing that can arise with multiple vendors. HC3 was built with supportability at its core and includes a convenient remote support "call home" feature that allows Scale's support engineers access to the console as if they were sitting in the room next to the user.

Data Placement

VMware's VSAN stores copies of each VM storage object using very static data placement (by default a single copy, but can be increased to as many as 12 replicas). Each capacity-contributing node is required to have at least one SSD, which is used as a cache and is not included in the overall capacity of the cluster. All I/O must first be written to the SSD as a read/write cache before being re-written to spinning disk media for long-term storage.

HC3 utilizes of a single wide-striped pool of SAS and/or SATA disks for RAID-10 like redundancy and performance while making use of RAM for caching. This protects against a single point of failure similar to the setup of VMware's VSAN using a replica factor of 2+, but requires no end user effort to ensure protection.

Storage Architecture

VSAN still retains the model of Virtual Machine disks being stored as "files" (VMDK) and using file system type operations. For example, taking a snapshot of a VM essentially freezes writes to the VMDK files for the VM so that changes to the VM are written into a new "delta" virtual disk file. Additional snapshots create chains of new delta disks that are inefficient to grow or update and have very high overhead to remove. Reads have to traverse the "chain" of snapshots to determine where the latest copy of the required block lives. Removing even a single snapshot requires reading all the changes from the delta snapshot virtual disk file and writing those changes back into the original VMDK disk file.

With the latest version of HC3, the added overhead of a filesystem is completely eliminated from the environment. Instead, HC3 utilizes an enterprise class clustered block storage layer that has been purpose-built to be consumed by the HC3 embedded KVM based hypervisor directly.

	Scale Computing	VMware Virtual SAN
Entry Level HA System	3 Nodes -fully pre-integrated hardware/software	Installable software only; Minimum 3 certified servers with at least 1 SSD and 1 disk per host
Entry Cost MSRP	\$25,499	VMware software only - \$23,412 (expected vSphere + vCenter + vSAN)
Storage Architecture	Single wide-striped pool of SAS or SATA disks, RAID 10 redundancy and performance	Static copies of each VM storage object (typically primary and secondary), all I/O goes through SSD layer first as read/write cache, SSD does not count towards capacity
Infrastructure Management	Web based management built in to cluster	Separate vCenter Server, SQL Database, Windows Client Component
Cache / Metadata	RAM and wide-striped disk pool	SSD required as a read cache and write buffer, not used for capacity, recommend 10% of disk size