TANEJA GROUP
TECHNOLOGY ANALYSTS

SCALE COMPUTING HC3 AND VMWARE VIRTUAL SAN HYPERCONVERGED SOLUTIONS - HEAD TO HEAD

SEPTEMBER 2014

Scale Computing was an early proponent of hyperconverged appliances and is one of the innovators in this marketplace. Since the release of Scale Computing's first hyperconverged appliance, many others have come to embrace the elegance of having storage and compute functionality combined on a single server. Even the virtualization juggernaut VMware has seen the benefits of abstracting, pooling, and running storage and compute on shared commodity hardware. VMware's current hyperconverged storage initiative, VMware Virtual SAN, seems to be gaining traction in the marketplace. We thought it would be an interesting exercise to compare and contrast Scale Computing's hyperconverged appliance to a hyperconverged solution built around VMware Virtual SAN. Before we delve into this exercise, however, let's go over a little background history on the topic.

Taneja Group defines hyperconvergence as the integration of multiple previously separate IT domains into one system in order to serve up an *entire* IT infrastructure from a single device or system. This means that hyperconverged systems contain all IT infrastructure—networking, compute and storage—while promising to preserve the adaptability of the best traditional IT approaches. Such capability implies an architecture built for seamless and easy scaling over time, in a "grow as needed" fashion.

Scale Computing got its start with scale-out storage appliances and has since morphed these into a hyperconverged appliance—HC3. HC3 was the natural evolution of its well-regarded line of scale-out storage appliances, which includes both a hypervisor and a virtual infrastructure manager. HC3's strong suit is its ease of use and affordability. The product has seen tremendous growth and now has over 900 deployments.

VMware got its start with compute virtualization software and is by far the largest virtualization company in the world. VMware has always been a software company, and takes pride in its hardware agnosticism. VMware's first attempt to combine shared direct-attached storage (DAS) storage and compute on the same server resulted in a product called "VMware vSphere Storage Appliance" (VSA), which was released in June of 2011. VSA had many limitations and didn't seem to gain traction in the marketplace and reached its end of availability (EOA) in June of 2014. VMware's second attempt, VMware Virtual SAN (VSAN), which was announced at VMworld in 2013, shows a lot of promise and seems to be gaining acceptance, with over 300 paying customers using the product. We will be comparing VMware Virtual SAN to Scale Computing's hyperconverged appliance, HC3, in this paper.

Here we have two companies: Scale Computing, which has transformed from an early innovator in scale-out storage to a company that provides a hyperconverged appliance; and VMware, which was an early innovator in compute virtualization and since has transformed into a company that provides the software needed to create build-your-own hyperconverged systems. We looked deeply into both systems (HC3 and VSAN) and walked both through a series of exercises to see how they compare. We

87 Elm Street, Suite 900 · Hopkinton, MA 01748 · T: 508.435.2556 · F: 508.435.2557

aimed this review at what we consider a sweet spot for these products: small to medium-sized enterprises with limited dedicated IT staff and a limited budget. After spending time with these two solutions, and probing various facets of them, we came up with some strong conclusions about their ability to provide an affordable, easy to use, scalable solution for this market.

The observations we have made for both products are based on hands-on testing both in our lab and on-site at Scale Computing's facility in Indianapolis, Indiana. Although we talk about performance in general terms, we do not, and you should not, construe this to be a benchmarking test. We have, in good faith, verified all conclusions made around any timing issues. Moreover, the numbers that we are using are generalities that we believe are widely known and accepted in the virtualization community.

Underlying Technologies

In order to understand how Scale Computing and VMware approached their solutions, we need to have a broad view of the technologies underlying each of these hyperconverged solutions.

Both Scale Computing and VMware rely on their own software stacks to provide a hyperconverged infrastructure. Both have a VM-centric storage model that uses DAS designed only to store virtual machines (VM); neither presents a storage protocol that would be suitable for general data storage. However, because of their heritage, Scale Computing and VMware approach the delivery of their hyperconvergence solution in two different ways. On one hand, Scale Computing delivers their products as an appliance; their hardware is pre-installed with software. On the other hand, VMware has taken a different approach; their software is installed on hardware, and the integration of the hardware and software is done by the customer or a third party. Below is a quick summary of the technology behind these two solutions.

A solution built-up around VSAN would consist of four parts: a hypervisor, VSAN, a management server, and the hardware. ESXi is VMware's hypervisor, is based on proprietary code, and is well-regarded by the IT community. vCenter server (vCenter) is used to manage the hypervisor and the VMs that are running on it. Plugins and extensions can be used with vCenter in order to expand its capabilities and its reach into the datacenter. VMware can run on a wide variety of commodity hardware; however, in order to insure that the hardware has been tested and certified to work with their software, VMware maintains a hardware compatibility list (HCL). The final piece of VMware's hyperconvergence stack is their VM-centric storage—VSAN.

VSAN was announced at VMworld in August 2013, to much fanfare. The beta was released shortly afterwards and, with over 10,000 downloads, it was one of VMware's most popular beta programs. To create VSAN, VMware took their existing file system (VMFS) and modified and morphed it into a proprietary, VM-centric, distributed storage system. VSAN uses the direct attached storage (DAS) on a server and requires both SSD and HDD drives. VSAN is tightly integrated with vSphere and embedded in the ESX kernel, making it a very VM-centric product, but also limiting it from acting as a general purpose storage provider. A VSAN cluster supports up to 32 nodes but a minimum of three nodes in the VSAN cluster must supply storage. Each of the nodes that is providing storage requires one SSD device (for read/write caching) and at least one HDD (for persistent storage). The HCL for VSAN is a subset of the HCL for vSphere. The nodes that provide storage can also host VMs. VSAN is compatible with, and vSphere does come with vSphere Data Protection (vDP) which can be used to back up the VMs hosted on a VSAN cluster. vDP is a disk based backup product so caution and forethought must be given before using VSAN as the storage for vDP as if you lose the VSAN cluster you will not be able to restore it using vDP.

Underlying Technologies		
Feature	**VMware Virtual SAN**	**Scale Computing HC3**
Maturity of Product	< 1 year	> 2.5 years
Hypervisor	ESXi (proprietary)	KVM (open source)
Hardware	Limited HCL	Appliance based
Memory Overcommitment	Supported	Not supported
Storage Scheme	Many - based on SSD for r/w cache, HDD for persistent data	One - wide striping and mirroring on HDD
Scalability	3-32 nodes	3-16 nodes
Storage Construct	VM-centric	VM-centric
Persistent Storage Media	HDD	HDD
Management Server	Single server – single point of failure	Distributed Architecture – built in
Data Protection	vDP, 3rd party products	Snapshot export, 3rd party products

HC3 is a Scale Computing hyperconverged appliance. The hypervisor used by Scale is KVM, a well-regarded open source hypervisor that was released for general availability in 2008. Scale Computing built and maintains its own storage engine for the VMs. The management of the HC3 virtual infrastructure in an HC3 cluster is built on a distributed architecture that does not rely on a central management server; any node can go offline without affecting the manageability of the cluster. All the software comes pre-installed on 1U servers. The server is built up using commodity hardware spec'd out by Scale Computing. As there are a limited number of hardware configurations, every appliance type is tested with each build of the software, thereby insuring that all software is 100% compatible with all the appliances. An HC3 cluster can be composed of as few as three nodes and can scale up to as many as 16 nodes. Scale Computing offers an export/import feature that creates a snapshot of a running VM which then can be saved to an SMB share. Their upcoming release (Q4 2014) will allow the replication of a snapshots to another HC3 cluster, this will be useful for business continuity. VMs hosted on an HC3 system can use third party product for backup and restoration.

Setup and Configuration

The installation of our HC3 cluster was dead simple: we racked three HC3 series HC4000 nodes; cabled the 10Gb Ethernet into a switch; hooked a keyboard, mouse and monitor to one of the nodes; powered it on and answered a few configuration questions and that was it. There was no software to install, and no HCL to consult to insure compatibility. On the initial boot, the HC3 node that we plugged the monitor into presented an ASCII GUI that guided us through the steps needed to set up its network settings and inter-node cluster communication. Following the initial setup, which took less than two minutes, we powered-on the other two nodes and followed similar steps to have them join the cluster. We were able to have our HC3 cluster set up and configured in less than 15 minutes. HC3 user management is very simple, with only a single user ID being in charge of the cluster. The virtual network backing the VMs on an HC3 cluster is also simple; a virtual bridge is used to attach the VMs to the external network and is set up automatically. Adding a fourth node to the cluster took 2 minutes. Scale Computing offers, at no cost, a one hour "hand holding" session for its new customers to guide them through the installation, configuration and management of an HC3 cluster and its VMs.

87 Elm Street, Suite 900 • Hopkinton, MA 01748 • T: 508.435.2556 • F: 508.435.2557

Setting up the VSAN cluster in comparison involved deciding which hardware and software to use, then installing and configuring the different software packages on the server. As you may have guessed, building up a VMware hyperconverged solution based around VSAN took more time. The first thing we needed to do was consult VMware's software interoperability guide to make sure that the various components and the versions of the components would work together. After sorting out the software stack we were going to use, we needed to consult VMware's hardware compatibility list (HCL) to ascertain what hardware is compatible with VSAN. VSAN requires that servers that contribute to the VSAN storage pool have an SSD drive in them, as well as at least one HDD drive and an IO controller that is on the VSAN HCL. We found that the VSAN HCL only supports a subset of the hardware that vSphere supports. It should be noted that we could not find any written statement from VMware about long-term commitment to hardware that would be supported by VSAN. Once the hardware was installed in our rack, we installed vSphere from a DVD to a USB device attached to the server as VSAN consumes the entire HDD drive, preventing it from being partitioned and used as an installation device for ESXi. Installation of vSphere was straightforward, and most of the time for installation was spent waiting for the bits to be transferred from the DVD. It took 21 minutes per node to install and configure each ESXi node. Once ESXi was installed, we installed vCenter Server Virtual Appliance (vCSA) which provides for the clustering and management functions of the ESXi nodes. We found it somewhat annoying that in order to install the vCSA, we needed to download and install the vSphere native client to a Microsoft Windows system. We installed the vCSA on a non-shared disk on one of the nodes to avoid the "chicken and egg" problem as detailed in the sidebar, "Where to install vCenter?" Once vCSA was installed, we needed to create and configure a vSphere cluster. VMware recommends using a 10Gb Ethernet connection and the vSphere Distributed Switch in combination with VMware vSphere Network I/O Control to share the 10GbE interfaces. They also recommend separating the various traffic (management, vMotion, VMs, Virtual SAN, etc.) onto different VLANs and a quality of service mechanism to mitigate contention scenarios. As this was not a performance test we forewent doing this, but expect that it would add time to setting up VSAN.

Setup and Configuration of 4 Node Cluster		
Feature	**VMware Virtual SAN**	**Scale Computing HC3**
Ease of setup	Complex	Simple
Sales SKU	Many hardware and software SKUs are needed	One
Number of Products to Configure	5 (Hardware, vSphere, VSAN, vCenter native client, vCenter)	1
Number of Nodes per Cluster	3-32	3-16
Time To Setup 3 Node Cluster	97 minutes (excluding time to consult documentation or attend 5 day training class)	15 Minutes (excluding one hour training session)
Time to add additional node	23 minutes	5 minutes
Recommended Training	5 day class	One hour
User Management	Role-based access control	Single User
Virtual Networking	Very complex	Simple

87 Elm Street, Suite 900 · Hopkinton, MA 01748 · T: 508.435.2556 · F: 508.435.2557

After the cluster was configured, we then needed to license the vCenter Server and each of the CPU sockets on each node for the hypervisor. It took 18 minutes to install, configure, and license the vCSA. VSAN does not need to be installed but simply enabled and licensed; this took us just two minutes. All the processors on every node in a VSAN cluster (even if they don't provide storage, and all nodes do not need to provide storage) need to have VSAN licenses. Adding a fourth node to the cluster took 23 minutes: 21 minutes to install and 2 minutes to add to the cluster and license. vSphere supports a complex "Role-Based Access Control" scheme to grant permissions, on a very detailed level, to which roles and rights vSphere users have access. Setting up the virtual networking on vSphere involves setting up and managing virtual switches, port groups, vmnics, physical NICs, etc. Virtual networking on vSphere can be very complicated and VMware has a 200 page document to help guide the user through its complexities. VMware recommends that people new to vSphere attend a 5-day Install, Configure and Manage class, and based on our experience, we agree that new vSphere administrators should attend this course. The class can be taken either online or live and costs about $3,900. VMware does not yet offer training on VSAN, so in order to install, configure and utilize VSAN we needed to consult the VSAN documentation on numerous occasions.

Where to install vCenter?

Installing the vCenter Server Appliance (vCSA) presented an interesting "chicken and egg" dilemma; I needed to have VSAN if I wanted to install it on shared storage and I needed vCenter in order to install VSAN to setup and configure VSAN to gain shared storage! In the end I chose to install and run vCSA on a single disk on a single node of my VSAN cluster, foregoing the benefits of shared storage. I felt more comfortable with the tradeoff of having access to vCenter if VSAN went down as opposed to not having vCenter if the node on which it was running went down. The benefits of running vCSA and other business critical VMs on shared storage cannot be over emphasized but until VMware sorts this dilemma out, their customers will have to make a tough choice. Scale Computing avoids this issue altogether by using a distributed management scheme.

Scale Computing was designed to be simple and simple it is. I would feel comfortable assigning the setup and creation of an HC3 cluster to someone with very limited IT experience. By contrast, setting up a VSAN cluster based on vSphere is considerably more complicated. The one-hour orientation that Scale Computing offers should be enough to bring someone who has a casual acquaintance with computing up to speed on their product.

Creation of Virtual Machines

After setting up our virtual environment, we went about using the systems to create VMs on the platforms.

To create a VM on the HC3 system we logged into the HC3 web portal, clicked on the "Virtualization" tab and answered the four questions that followed, made five more mouse clicks and the VM was created. During this process we were presented with a dialog that asked what OS we were installing. We created a Windows 2008 R2 VM, but HC3 supports most modern Windows and Linux OSes. The size of the VM (number of vCPUs and amount of RAM) is dependent on the OS that is specified, i.e. a Windows 2008 VM will have different needs and will therefore get a different virtual hardware profile than say a Red Hat VM. Scale Computing made the creation of VMs so simple that we feel that anyone with a basic understanding of IT could easily navigate the menus in order to create a VM.

87 Elm Street, Suite 900 · Hopkinton, MA 01748 · T: 508.435.2556 · F: 508.435.2557

Navigating the vCenter management console and creating a VM to use VSAN for storage is much more complicated than it is with HC3. After logging on to the vSphere web client we were presented with a large selection of objects. Someone without experience with vSphere and VSAN would be at a loss to figure out how to go about creating a VM. Luckily we did have this knowledge and we were able to create a VM, but before creating a VM, a VSAN "storage policy" needed to be created. We found the terminology used in the creation of a storage policy somewhat obtuse and we needed to consult the documentation to fully understand it. The storage policy describes the storage attributes, such as level of protection and performance that you would like the VMs using the policy to possess. VSAN does offer more choices than HC3 with regards to the layout of the storage of data as you can specify the number of copies of the data, amount of flash to dedicate to it, how wide to stripe it, etc.. We will not detail all the steps need to create a VM on a VSAN cluster but Figure 1, taken from the "Virtual SAN (beta refresh) POC Guide," shows the workflow for creating a VM using VSAN.

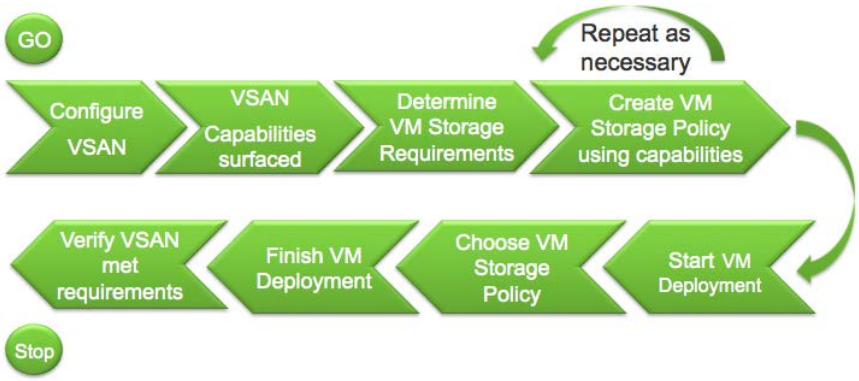

Figure 1 Workflow for creating a VM with VSAN

Creation of Virtual Machines		
Feature	**VMware Virtual SAN**	**Scale Computing HC3**
Management Portal	Very complex	Simple
Creation of VMs	Moderately complex	Simple
Virtual Networking	Complex	Simple

Consumption of Storage

It is common in a datacenter to have multiple VMs running the same OS. Manually installing the OS every time a new VM is needed would be tedious and an unwise use of a system administrator's most prized possession—time. Luckily, both platforms allow the creation of copies of a VM from a parent image through a process known as "cloning." Both systems also support Microsoft "sysprep" which makes the customization for the cloned image extremely simple. The first thing we noticed when cloning a VM was how much faster it was to clone a VM on the HC3 system (seconds) than it was on the VSAN system (minutes). We then noticed how little space was used by HC3 clones, and a deeper look into the technology explains why this was the case.

Clones created on HC3 use a "thin cloning" (linked clones) mechanism that will track the changes made to the initial VM, which means storage space will not be consumed until it is actually needed. Not only does this conserve disk space, it makes the creation of a clone almost instantaneous. VSAN uses a more traditional approach to the creation of a clone and copies the entire image from

the parent image. This explained the disparity in the cloning time.

A quick run of the numbers shows how much disk space is initially used by HC3 and VSAN clusters. We determined that on both systems the parent image was 24GB but actually 48GB of disk space was being used as both systems by default create a mirrored copy of the data. Only after clones from the parent image were created did we see the disparity in the amount of storage consumed between the two platforms. The first clone created on the VSAN system consumed the same amount of space as the parent image for a total of 96GB while the clone created on the HC3 system did not require any additional storage, so the total storage consumed on the HC3 system was still 48GB. The first clone saved 48GB of space on the HC3 system over what was used by the VSAN system. Furthermore, each subsequent clone will save the same amount of space, so 39 clones on an HC3 system would still consume only 48GB of space, whereas the same 40 VMs (the parent image and 39 clones) would consume 1,920GB, almost two terabytes of storage.

Consumption of Storage		
Feature	**VMware Virtual SAN (FTT = 1)**	**Scale Computing HC3**
Ease of Clone Creation	Easy	Easy
Support of Sysprep	Yes	Yes
Clone Creation Time	Minutes	Seconds
Clone Deletion Time	Seconds	Seconds
Initial Image	24GB	24GB
Space Consumed by Initial Image and mirror copy	48GB (2 x 24)	48GB (2 x 24)
Space Consumed by Original and Single Clone	96GB (2 x 2 x 24GB)	48GB; will grow as needed
Initial Space Consumed by Original and 39 Clones	1,920GB (2 x 40 x 24GB)	48GB; will grow as needed

We did notice that as applications and data were added to the VMs the clones consumed more storage, but this was to be expected and was approximately the same on both systems. Cloning is not an exotic concept and VMware is currently using a similar technology (linked clones) with its VDI solution and in some of its other product lines.

Snapshot Management

Snapshots capture a point-in-time copy of a VM and is often used before a major event on a VM, such as software upgrade. This allows the VM to be reverted back to a known good state if needed. Most modern virtualization solutions support the creation of VM-centric, as opposed to array based, snapshots; HC3 and VSAN are no exception. With both systems creating a snapshot is simply a matter of clicking on the VM and selecting "snapshot". Although the mechanism for creating VMs is the same we were astonished at how different these two solutions are.

We were surprised to find how quickly HC3 snapshots are created; even snapshots of very large VMs were created in a matter of seconds. We were also surprised to find out that an HC3 VM can have over 5,000 snapshots and each of these snapshots can be used to create a clone. Not only was creating a snapshot quick but the removal of them, even very large ones, was only slightly slower. HC3 snapshots fell short in only two areas: users can only create crash consistent, rather than application

Snapshot Management		
Feature	**VMware Virtual SAN**	**Scale Computing HC3**
Ease of Snapshot Creation and Management	Easy	Easy
Snapshot Creation	Moderately Fast	Very Fast
Snapshot Deletion	Can be very long	Fast
Snapshot Tree Length	32	5,000
Crash consistent Snapshot	Yes	Yes
Quiescing of VM	Yes	No
Memory Snapshots	Yes	No

consistent, snapshots (since quiescing of the VM is not supported, but is on Scale's product roadmap); and snapshots cannot be taken that capture the memory of a running VM. Creating and working with HC3 snapshots was straightforward, used little I/O and was very quick.

On a VSAN cluster, due to the underlying snapshot technology, it can be time consuming to create and delete snapshots. We have seen the creation of snapshots take minutes and we have seen reports of the deletion of snapshots taking hours during which it will be consuming resources from the nodes. On a VSAN a VM is limited to 32 snapshots. We did like the ability to create quiesced snapshots using MS VSS, rather than the more common crash consistent snapshots. We also liked the ability to create a snapshot that captured the RAM as well as the underlying storage of the VM.

Upgrading System

All systems will eventually need to have software patches and upgrades. Scale Computing has made this as non-disruptive as possible by enabling one-click rolling upgrades. When a system upgrade is available, a notification is displayed on the management portal; by clicking the notification, the upgrade will be downloaded and its installation will begin. In order to not disrupt the running of an enterprise's operations, all of the VMs on one node are live migrated to the other nodes in the cluster while the node is being upgraded. After the node has been upgraded, the process is repeated for the other nodes in a cluster. The time to perform an upgrade will vary depending on the number of VMs running on the system and the size of the upgrade. The main point here is that an upgrade to an HC3 cluster is totally automated and running applications will not suffer any downtime during the process.

System Upgrade		
Feature	**VMware Virtual SAN**	**Scale Computing HC3**
Ease of Upgrade	Unknown	Very easy
Upgrade Granularity	Product by product	Holistically
Automated Rolling Upgrades	Requires Enterprise Edition or better	Standard feature

As VSAN has yet to have its second release, we are unsure how it will be upgraded. Our best guess is that it will use "vSphere Update Manager" (VUM), which is used to upgrade other vSphere products. VUM is a separate feature that needs to be installed, and consists of both a server component and a client component. VUM updates VMware products independently of each other, not together in a holistic fashion. The

87 Elm Street, Suite 900 · Hopkinton, MA 01748 · T: 508.435.2556 · F: 508.435.2557

Update Manager server can either be installed on the same system as vCenter Server or on a different system; however, it can only be installed on a 64-bit Windows system, negating the ability for it to be installed on the vCSA. Performing rolling upgrades requires "Distributed Resource Scheduler" (DRS), which is only included with the vSphere Enterprise or Enterprise Plus editions of vSphere, not the less expensive Standard edition.

Pricing and Licensing

HC3 systems can be purchased directly from Scale Computing or from one of their partners. In keeping with Scale Computing's overriding theme of simplicity, they keep their purchase process as simple as possible; a single SKU is all that is needed to purchase an HC3 cluster. This SKU covers the software, hardware, and the first year of support. For this comparison, we decided to compare entry level systems. Scale offers a 3-node cluster built with HC1000 nodes. Each HC1000 node is comprised of a 4 core Intel processor, 32GB of RAM, 4 x 500GB 7.2K drive, and 2 x 1GbE NICs. The price for the entire configuration is $25,499.

Pricing and Licensing		
Feature	**VMware Virtual SAN**	**Scale Computing HC3**
Sales SKU	Multiple HW and SW SKUs	One
Support	Multiple support contracts with multiple vendors	Single support contract for SW and HW
Support Contracts	4 tiers Worldwide	1 Tier - 7 x 24 US based
Support Rating	World class	World class

Calculating the price for a VSAN cluster is a more complex process; the hardware along with all of the components of the software stack must be considered. For the hardware, we priced it out to the most comparable Dell Server we could find: a PowerEdge R320 comprised of a 4 core Intel processor with 32GB of RAM, 4 x 500GB 7.2K drive, 2 x 1GbE NICs, and a 300GB SSD drive. We were not able to configure the R320 to meet the exact requirements required by the VSAN HCL, as items such as a compatible IO controller were not available but we were looking for a general comparison and felt comfortable with the configuration. The list price for each server hardware was $5,024. vSphere Enterprise (Enterprise is needed to get DRS to enable rolling upgrades) with one year production support adds an additional $3,594 per node. To manage the cluster, an instance of vCenter Server is also required, which with one year of production support comes out to $6,244. VSAN is priced at $2,495 per host CPU socket for a total on all three nodes of $39,583. The cost of the VSAN and vSphere licenses would have doubled if we had gone with more dual processor mother boards as all processor sockets on every node in a VSAN cluster need to be licensed.

VMWare EVO:RAIL

As we were finishing this Technology Validation VMware announced EVO:RAIL, which is VMware's hyperconvergent offering. EVO:RAIL uses VSAN as the storage layer, ESXi as the hypervisor and it has a new, greatly simplified management GUI. We see this as a validation of Scale's HC3 philosophy and methodology of creating a simple easy to use and procure virtualization solution. We will be performing a more detailed Technology Validation on EVO:RAIL, as more details come out about it and as EVO:RAIL based solutions become available in the marketplace.

A Scale Computing support contract covers both the hardware and the software. We like that Scale Computing takes ownership of the entire package, as it eliminates any finger-pointing if a problem does arise. Scale offers only one tier of support: 7 x 24. All support is based in the US. Scale Computing support has a "Net Promoter" score that places it in the "World Class" category.

87 Elm Street, Suite 900 · Hopkinton, MA 01748 · T: 508.435.2556 · F: 508.435.2557

VMware offers four levels of support from basic to mission critical. Support is delivered by a worldwide team of support engineers. VMware support also has "World Class" customer satisfaction ratings.

As with most products, a true apples-to-apples comparison was not possible, as each product has its own nuances. For example, the vSphere system comes with features that the HC3 system does not, such as "Fault Tolerance" and "Data Protection." It would also require roughly 20% less memory due to its real memory deduplication, but would require more memory to run its virtual infrastructure software stack (i.e. vCenter Servers, VSAN, VUM, etc.) The HC3 cluster system would require less spinning disks due to its cloning technology and would not require an SSD drive. We used list prices for all pricing and expect that discounts could be obtained from vendors. The point of this pricing exercise was to get a broad overview of the CAPEX costs to deploy these two hyperconverged products. As shown below, the software cost alone of the VMware system is comparable to the total cost of an HC3 system.

Cost to Deploy Entry level 3 Node Cluster **Specification Per Node** **Quad Core i7 Intel System** **32GB RAM** **2 1GbE NIC** **4 x 600GB SAS Drives**				
	Product(s)	**Software Cost**	**Hardware Cost**	**Total Cost**
Scale Computing HC3 1000 Cluster	3 Node HC3 HC1000 system	Included	Included	$25,499
VSAN Deployment **(Total HW + SW Cost)**				$39,583
vSphere Hardware Cost	3 x Dell ProLiant R320 Servers - **(3 x $5,024)** **NOTE: Does not include IO control**		$15,072	
vSphere Software Cost	3 x VMware Virtual SAN licenses **(3 x $2,495)** 3 x vSphere ESXi Enterprise **(3 x $3,594)** vCenter Server Standard **($6,244)**	$24,511		

87 Elm Street, Suite 900 · Hopkinton, MA 01748 · T: 508.435.2556 · F: 508.435.2557

Taneja Group Opinion

Scale Computing has done an outstanding job simplifying their virtualization solution: from the acquisition of the cluster (one SKU), to the setup (under 20 minutes), to adding capacity (2 minutes per additional node). However, this simplicity does not come at the expense of the ability to virtualize and consolidate a company's IT infrastructure. Companies that want the benefits of abstracting, pooling and allocating their compute, networking and storage resources without unnecessary complexity will find that Scale Computing's HC3 hyperconvergence cluster will be able to deliver the benefits of consolidation and business continuity without any unnecessary complexity. To be sure, VMware does offer capabilities that make it suitable for large-scale enterprise deployments that can support its additional cost and complexity. For the majority of small to medium-sized companies that just want to get work done in the most efficient and non-complex manner possible, an HC3 cluster will fit their bill.

Looking back over this paper, you may have realized that a lot more type was given to VSAN than to HC3; this was not by design, but by necessity. We simply needed more words to explain how to do tasks on a VSAN-based solution than on a comparable HC3 solution. For example, to set up an HC3 cluster, you simply plug it in, answer five questions and start to provision VMs; it can take less time to set up an HC3 cluster and begin to deploy VMs on it than it does to check the HCL and product compatibility list to begin to plan the purchase of a VSAN-based solution. It takes both time and effort to make something simple, and it is obvious that Scale Computing has invested a lot of both into HC3; they have done a fine job of making their product as easy to use as possible.

For companies that are looking for a simple to use, yet powerful virtualization solution that is attractively priced, Scale Computing is an easy choice. We were impressed with VSAN's capabilities, but found that its price simply didn't justify the additional cost and training that would be needed in order to take advantage of them, especially for smaller companies with limited resources and a non-dedicated IT staff.