SCALE COMPUTING HC3

A Look at a Hyperconverged Appliance

SEPTEMBER 2015

(This is an update of a report that we originally released in June of 2012.)

Consolidation and enhanced management enabled by virtualization has revolutionized the practice of IT around the world over the past few years. By abstracting compute from the underlying hardware systems, and enabling oversubscription of physical systems by virtual workloads, IT has been able to pack more systems into the data center than before. Moreover, for the first time in seemingly decades, IT has also taken a serious leap ahead in management, as this same virtual infrastructure has wrapped the virtualized workload with better capabilities than ever before - tools like increased visibility, fast provisioning, enhanced cloning, and better data protection. The net result has been a serious increase in overall IT efficiency.

But not all is love and roses with the virtual infrastructure. In the face of serious benefits and consequent rampant adoption, virtualization continues to advance and bring about more capability. All too often, an increase in capability has come at the cost of complexity. Virtualization now promises to do everything from serving up compute instances, to providing network infrastructure and network security, to enabling private clouds.

Figure 1: Is it storage or compute? With a Scale Computing HC3 appliance it is both.

For certain, much of this complexity exists between the individual physical infrastructures that IT must touch, and the simultaneous duplication that virtualization often brings into the picture. Virtual and physical networks must now be integrated, the relationship between virtual and physical servers must be tracked, and the administrator can barely answer with certainty whether key storage functions, like snapshots, should be managed on physical storage systems or in the virtual infrastructure.

Scale Computing, an early pioneer in HyperConverged solutions, has released multiple versions of HC3 appliances and now includes the 6th generation of Scale's HyperCore Operating System. Scale Computing continues to push the boundary in regards to simplicity, value and availability that many SMB IT departments everywhere have come to rely on. HC3 is an integration of storage and virtualized compute within a scale-out building block architecture that couples all of the elements of a virtual data center together inside a hyperconverged appliance. The result is a system that is simple to use and does away with much of the complexity associated with virtualization in the data

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

center. By virtualizing and intermingling compute and storage inside a system that is designed for scale-out, HC3 does away with the need to manage virtual networks, assemble complex compute clusters, provision and manage storage, and a bevy of other day to day administrative tasks. Provisioning additional resources - any resource - becomes one-click-easy, and adding more physical resources as the business grows is reduced to a simple 2-minute exercise.

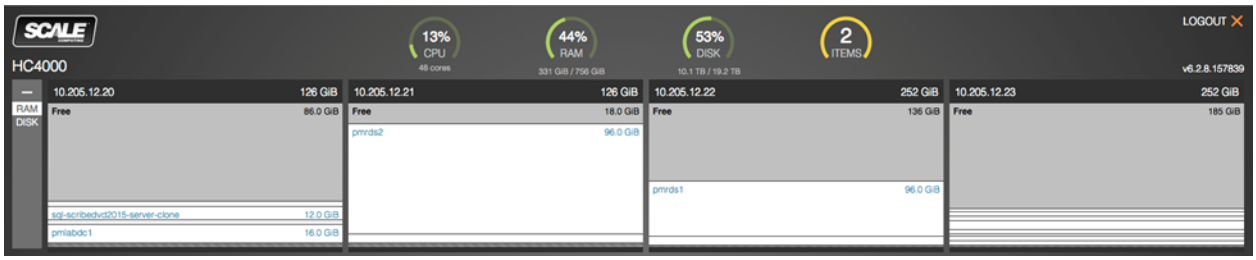

Figure 2 Both storage and compute capacity increased with another node

While this sounds compelling on the surface, Taneja Group recently turned our Technology Validation service - our hands-on lab service - to the task of evaluating whether Scale Computing's HC3 could deliver on these promises in the real world. For this task, we put an HC3 cluster through the paces to see how well it deployed, how it held up under use, and what special features it delivered that might go beyond the features found in traditional integrations of discreet compute and storage systems.

While we did touch upon whether Scale's architecture could scale performance as well as capacity, we focused our testing upon how the seamless integration of storage and compute within HC3 tackles key complexity challenges in the traditional virtual infrastructure.

As it turns out, HC3 is a far different system than the traditional compute and storage systems that we've looked at before. HC3's combination of compute and storage takes place within a scale-out paradigm, where adding more resources is simply a matter of adding additional nodes to a cluster. This immediately brings on more storage and compute resources, and makes adapting and growing the IT infrastructure a no-brainer exercise. On top of this adaptability, virtual machines (VMs) can run on any of the nodes, without any complex external networking. This delivers seamless utilization of all datacenter resources, in a dense and power efficient footprint, while significantly enhancing storage performance.

Time and Dollar Consuming Virtual Infrastructure Challenges	**The HC3 Change**
Deployment/Provisioning	HC3 eliminates provisioning complexity, and greatly reduces the time and effort required to configure a virtual infrastructure In illustration we have that a HC3 3 node cluster in less than 20 minutes. This less time than it takes to set up a single component with a traditional virtual infrastructure.

Capacity Management	With data accessible from all nodes virtual machines can be easily and near instantaneously moved and rebalanced across compute resources. This can drive up utilization and eliminate complex planning to make sure new workloads end up in the right place. This is out-of-the-box; no complex cluster or resource pool configuration is required.
Component Failure and Recovery	An HC3 cluster can tolerate any component failure including an entire node. When the system detects a disk failure it will begin to rebuild that data that was contained on the disk on other disks immediately. In the event of a node failure the VMs that were running on it will be restarted on other nodes immediately.
Growth and Scaling	Expanding storage performance, storage capacity, or compute resources in an HC3 cluster simply takes the addition of more nodes to the cluster. This is distinctly simpler than complex planning required for the expansion of discreet storage pools, data stores, and servers to add more capacity or performance to a traditional datacenter.
Business Continuity	The use of built-in snapshot and remote replication capability on a per-VM granularity enables a complete disaster recovery solution in just a few clicks. This puts leading enterprise-like business continuity in the hand of smaller businesses that previously could not afford such capabilities.
Obsolescence and Migration	HC3 cluster architecture puts an end to obsolescence and migration - for both storage and compute. As hardware ages, new hardware can be introduced to an HC3 cluster, and old nodes can be phased out opportunistically. This rolling upgrade process means no more timely and costly data migrations or availability disruptions.

Meanwhile, within an HC3 cluster, these capabilities are all delivered on top of a uniquely robust system architecture that can tolerate any failure - from a disk to an entire cluster node - and guarantee a level of availability seldom seen by mid-sized customers. Moreover, that uniquely robust, clustered, scale-out architecture can also intermix different generation of nodes in a way that will put an end to painful upgrades by reducing them to simply decommissioning old nodes as new ones are introduced.

HC3's flexibility, ease of deployment, robustness and a management interface is the simplest and easiest to

Technology Validated:

HC3 stands poised to transform the way IT is done in SMB/SME customers, by reducing the complexity and overhead of on-going IT infrastructure management. HC3 delivers outstanding availability protection, and unifies storage and compute into a seamless adaptable pool of resources that can put an end to painful system replacements and the constant purchase and deployment of more hardware. This will change how IT is done.

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

use that we have seen. This makes HC3 a disruptive game changer for SMB and SME businesses. HC3 stands to banish complex IT infrastructure deployment, permanently alter on-going operational costs, and take application availability to a new level. With those capabilities in focus, single bottom-line observations don't do HC3 justice. In our assessment, HC3 may take as little as 1/10th the effort to setup and install as traditional infrastructure, 1/4th the effort to configure and deploy a virtual machine (VM) versus doing so using traditional infrastructure, and can banish the planning, performance troubleshooting, and reconfiguration exercises that can consume as much as 25-50% of an IT administrator's time. HC3 is about delivering on all of these promises simultaneously, and with the additional features we'll discuss, transforming the way SMB/SME IT is done.

FOCUS ON SCALE COMPUTING

Indianapolis-based Scale Computing has its lineage in, and has long been known for, its file and block scale-out storage appliance. Scale Computing had for a number of years sold this uniquely affordable but highly scalable unified storage solution to SMB and SME customers. Behind that storage is a collection of unique data management intellectual property that Scale Computing continues to hone, as detailed in the sidebar "The Scale Computing Cluster Foundation".

This storage foundation pedigree was behind Scale Computing ambitious expansion to further alter data centers for the better. The expansion was a coupling of a Scale storage cluster with virtualized compute and networking running within the same node, and was labeled HC3, an acronym that stands for *hyperconvergence*.

In order to allow Scale Computing to optimize its storage system for virtual machine storage it took the bold decision to forgo SAN or NAS storage and use a proprietary storage layer (SCRIBE) for its HC3 systems. By doing this the storage system is truly VM-centric and is not directly accessible by anything other than the VMs hosted on the HC3 nodes. The benefits of this proprietary storage are many including allowing VMs to directly access the storage without the overhead of a file system and allowed Scale Computing to add new functionality such as VM level snapshots and near instantaneous VM cloning and reduce the storage requirements for VMs.

A Wave of Hyperconvergence

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 scale over time, in a "grow as needed" fashion.

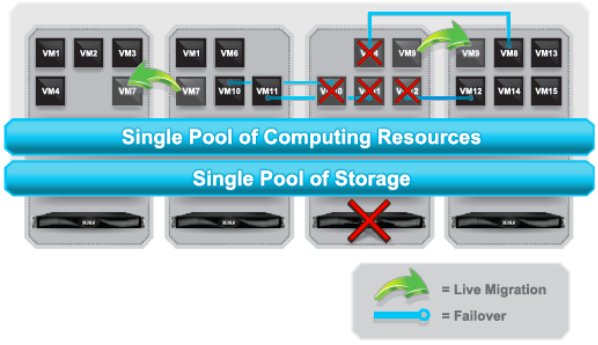

Figure 3: An HC3 cluster combines both compute and storage into a single pool of resources

Whereas convergence has time and again emerged as a focal point of IT innovation, it has most often focused within single IT domains - converging the protocols and functionality of multiple types of networking onto one single fabric for instance. Hyperconvergence is far different,

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

and far more ambitious, aiming to virtualize and integrate nearly all of the technology and functionality required to run a data center.

Sales, Licensing and Support

Scale Computing is dedicated to making computing in the datacenter as simple as practical and this mindset extends to its sales, licensing and support model. Scale Computing sells only hyperconverged appliances. The entire appliance, hardware and software is sold under a single SKU. All of HC3 features are included so there aren't any licensing issues to deal with. Scale has world class support organization to deal with any issues that may arise. Again they keep it is simple as possible, they offer one support plan; 7 x 24 live support to all of their customers.

We did have an issue during our testing and needed to call Scale Computing support. The call was answered on the second ring, the engineer was not only knowledgeable but also affable. She was able to confirm that we were seeing was a known issue that would be taken care of in the next update. She offered a work around and closed the call. From time to dial to resolution was under three minutes.

After our experience with Scale Computing support we were not surprised to find that it has a "Net Promoter" score that places it in the desirable "World Class" category. When I talked to David Demlow – VP of Product Management and Support for Scale Computing - he attributes its world-class support standing to his skilled and highly competent reactive engineering staff. He mentioned that on average each engineer has over 9 years of experience. He also said that 98% of incoming calls are handled live with an average 15 second ring time. Currently the entire support staff is based in Indianapolis, Indiana.

A LOOK AT HC3

HC3 was one of the first of closely coupled compute+storage appliances and was instrumental in advancing hyperconvergence into the mainstream. HC3 is built on top of the Scale Computing storage engine that brings key capabilities to the table. First, since Scale Computing nodes are built using powerful Intel CPUs each node has abundant processing power that can easily be applied to multiple tasks such as virtualization and the management of the storage. The storage functions on the nodes consume little processor power. We observed during our testing that even an intense disk rebuild process did not have an impact on the performance of the VMs hosted on the nodes. Second, its clustered architecture connects individual nodes into a seamless pool of storage and compute this allows VMs to run on any node in the cluster without any additional complexity. Since this pool is constructed by connecting together many individual nodes over a network, the cluster is already network connected and can support the network needs of the VMs in the cluster.

Moreover, Scale's clustered architecture ensures each node has equal access to all resources. Any single node can fail, and while that might temporarily interrupt the VMs running on that node (depending on a business's approach to HA), the VMs can immediately be restarted on any other node in the cluster.

The Scale Computing Cluster Foundation

An HC3 cluster supports a true scale-out architecture. An HC3 cluster starts with a minimum of 3 nodes and can expand with the addition of more nodes up to a maximum of 16 nodes currently. For flexibility Scale Computing offers a range of nodes with different compute and storage profiles – entry level HC1000 nodes, capacity-rich HC2000 nodes, high-performance HC4000 nodes and even

cost effective "storage only" node models. These nodes can be mixed and matched within a cluster allowing customers to tweak their performance and capacity according to their business needs.

Deep within this Scale Computing cluster there is significant intelligence that has been crafted into their VM-centric storage - SCRIBE (Scale Computing Reliable Independent Block Engine). The core of SCRIBE is the technology that aggregates the capacity of all the disks contained on all the nodes in an HC3 cluster into a single pool of storage. Data is protected and performance is insured by laying out the data using a wide striped, mirrored storage scheme. This scalable, distributed, software-based scheme is based on bespoken code written and maintained by Scale Computing engineers. It provides exceptional data protection as well as rapid recovery in the event of a drive failure, by allowing for a "many-to-one" rebuild process where many storage controllers and drives are sending data to the replacement drive without over-taxing the resources of any one drive or controller. Scaling-out a Scale Computing storage cluster only requires adding additional nodes to a cluster.

With utterly simplified management, SMB and SME friendly price points, and optimally sized nodes Scale Computing has been steadily racking up converts. Scale now claims over 1000 customers among a wide range of organizations.

PUTTING SCALE COMPUTING HC3 TO THE TEST

When we visited Scale Computing in the summer of 2014 we examined Scale's HC3 solution in depth to ascertain whether HC3 can deliver unique value and capability to SMB/SME customers. To evaluate and test HC3 we turned to hands-on lab exercises designed to stand up a real world workloads against an HC3 cluster and execute a number of tasks against these workloads. The tasks we performed and questions we asked included:

1. Infrastructure setup, deployment and maintenance, for both the cluster itself and virtual workloads.

Would the HC3 solution be as easy, or easier, to setup and maintain than separate storage and compute systems? Would setting up virtual machines on top of this cluster match or exceed the ease of deploying those same workloads in other infrastructures?

2. Virtual workload management.

Can a hyperconverged solution like HC3 deliver all of the virtual machine management functionality expected from modern hypervisors, in particular those features that are most valuable in a scalable, adaptable infrastructure such as the instantiation of new VMs, live-migration of VMs and VM cloning?

3. Resiliency, failure, and recovery.

Given Scale Computing's roots in building highly resilient storage systems, would an HC3 cluster be more robust and offer better protection from failures and outages for VMs?

4. Scalability and adaptability.

Would Scale Computing's pedigree in scale-out storage system make growing and managing a virtual infrastructure easier?

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

Test Lab Environment

Test Environment		
Equipment	**Description**	**Capacity**
HC3 Cluster A 4 Node Scale Computing HC3 HC4000 cluster (3 Nodes initially, one additional introduced to the cluster during testing)	**Per Node** Dual socket, 12 core, server class Intel processors 128GB RAM 4.2TB Raw Storage Capacity (8x 600GB 10,000 RPM SAS Drives) 2 x Dual-port 10G Ethernet NICs Built-in Storage, Hypervisor and Infrastructure Management	~ 60 Virtual Machines on each node depending on resource requirements of guests
Dell M2024 Network Switches	48 10G Ethernet ports	Up to 24 HC3 nodes
APC HCX 5000KVA UPS	Power failure protection for HC3 nodes and switches	N/A

For this testing we used a 4-node HC4000 cluster interconnected via a Dell M2024 switch. The only hardware used for our testing was the HC3 cluster - no additional servers or storage systems were used.

Deployment and Virtual Infrastructure Setup

Scale Computing HC3 systems are available through a large number of highly qualified VARs and resellers who can carry out initial cluster setup and configuration, but we wanted to set it up ourselves to see how complex it would be to set up compared a more traditional infrastructure.

Deploying an HC3 cluster requires little more than plugging in the power, and attaching network cables from each node to a switch. Since the cluster is highly dependent upon node-to-node network communications many HC3 users choose to use a pair of dedicated switches help provide an extra measure of resiliency. Larger environments may wish to connect the nodes to larger shared switches, and Scale Computing provides a set of VLAN best practices to help isolate, secure, and protect the cluster.

We attached a crash cart to one of the nodes and powered it on. On power up we were presented with an ASCI GUI that guided us through setting up the nodes networking and gave us the option of joining an existing HC3 cluster or setting up a new one, we chose to create a new cluster. The initial setup of the node took less than 2 minutes. We were able to set up and have two additional nodes join our cluster and within 10

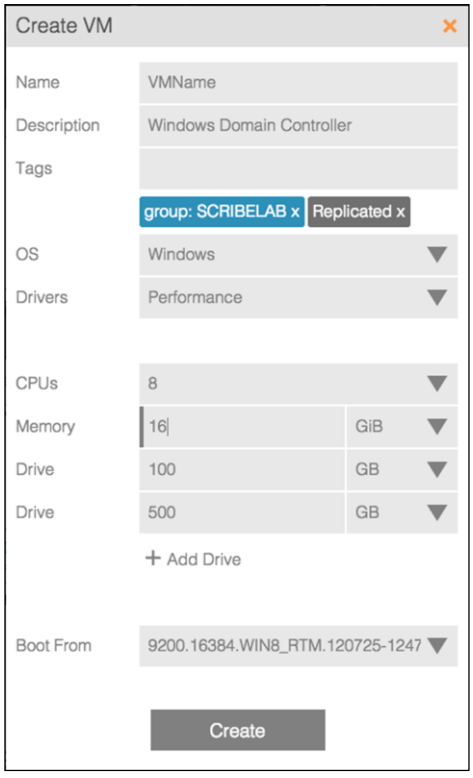

Figure 4: Wizard used to create a VM

minutes a fully functioning HC3 cluster. Later on we added an additional HC3 node in the rack were able to effortlessly add it to the cluster in less than 2 minutes.

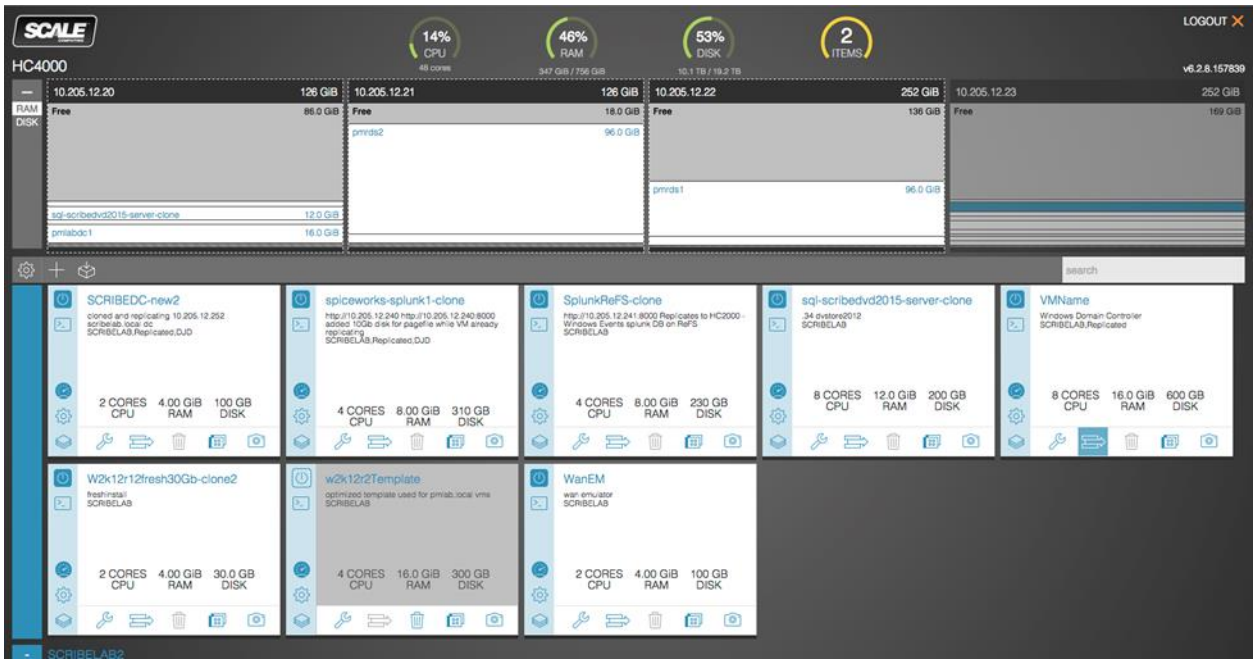

Figure 5: Adding a fourth node to the cluster

From Deployment to Virtual Infrastructure Setup

Moving beyond initial cluster setup, we set our focus on evaluating how much additional configuration would be required before the cluster was ready to use in a production environment. We started at a slight advantage, as we already had a set of virtual machine templates and installation ISO images at our disposal and were able to upload them through the HC3 web management console and had stored them on our HC3 cluster.

Our next step consisted of entering the cluster management web GUI by pointing the web browser on our laptop to the IP address of the cluster. After signing in, we found an intuitive interface, with clear indication of where to go to start deploying VMs. As shown in we created a Windows 2012 VM, attached an ISO image to it and performed our installation. Altogether, we arrived at the first installation screen of the Windows OS within 2 minutes of opening the HC3 management web GUI. Setting up a full virtual infrastructure - storage, networking, and compute - and deploying our first VM in less than 17 minutes is a remarkable achievement.

After initial deployment, configuring VMs on vSphere can be done more easily, but still involves significant complexity in ensuring VM placement, maintaining virtual networks, and walking through all of the configuration permutations of the VMs. Comparatively, we estimate that VM deployment with HC3 takes 1/4 the time and effort, and will be a welcome relief for the many SMB and SME environments who use

Technology Validated:

We measured deployment and configuration of an HC3 cluster as 10 minutes and 2 minutes to deploy the first VM, versus what typically takes hours for traditional hardware configuration, and 16 more minutes to configure a freshly installed hypervisor for the first VM. Bottom line, we give Scale Computing a 10X time and effort advantage in deployment.

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

little of this functionality.

The secret to Scale Computing's performance is simplicity, and that simplicity comes from the close melding of storage with compute, on a scale-out platform where networking is already built-in and handled by default. As a result powering on a Scale Computing cluster means it is ready to go without any additional configuration needed other than providing IP addresses. That is pure simplicity, and given the hundreds of virtual servers that exist in datacenters today, that simplicity has significant pay-off. 8X faster deployment adds up when multiplied across hundreds of virtual machines, an annual cycle of replacement servers and the inevitable goofs or disasters requiring rework or recovery.

Workload Operation and Management

Next we turned to putting HC3 through a deeper set of exercises involving tasks such as protecting, copying, and migrating VMs.

After our initial deployment, we noted that HC3 currently works on firm memory allocations, without any memory over-subscription or dynamic memory management. We suspect this functionality will arrive soon, but this has become less of an issue as real memory has dramatically dropped in price and Scale Computing offers nodes with up to 256GB of RAM.

Scale Computing's Cluster Management web GUI provides visibility into the resource consumption of a cluster in a resource utilization bar that spans the top of the screen. To test how well we could see into utilization, we installed and cloned a variety of machines. HC3 supports thin (AKA linked) clones that will not consume any space until written to and take mere seconds to create.

With our utilization increasing we turned to migrating a VM running a production load. To do this we created a Microsoft SQL Server 2008 VM with DVD Store 2.1 benchmark software installed on it. Then we started DVD Store and proceeded to migrate the VM from one node to another.

The migration of a VM was very straightforward. As shown in Figure 6 all it involved was clicking on the move icon button on the selected VM and then the eligible nodes are highlighted that can support the “move to” – simply click on the node you want to move the VM to. As the move (live migration) of the VM took place, several things happen under the covers.

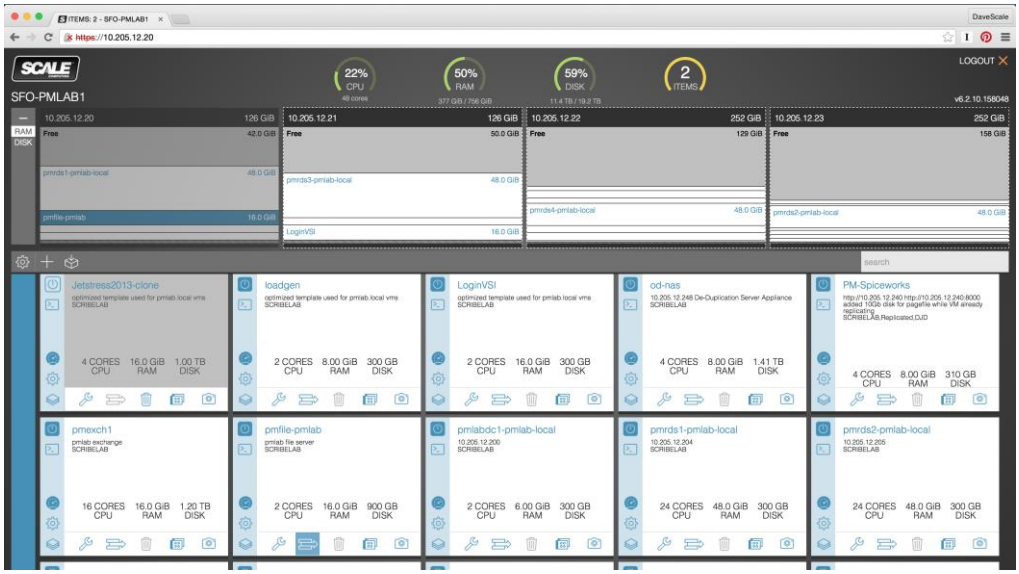

Figure 6: Rebalancing by moving a VM node to node

VMs running on an HC3 cluster always connect to the storage pool from the node they are running on to keep IO traffic to a minimum. Since the cluster runs on top of a clustered, distributed storage engine (SCRIBE) any node in the cluster can access the storage. During a migration, the VM maintains an uninterrupted connection to the underlying storage, but at some point the connection for the VM will need to migrate to the new node. As the migration to the new node takes place, the data in the cache of the node will be flushed, and then repopulated on the other node. We expected to see a disruption in IO during the migration.

What we observed was surprising; there was very little disruption in the IO. In fact, our "Operations per Minute" - the performance measure generated by DVD Store - demonstrated no fluctuation, and the only observable impact was a temporary increase in transaction latency while the cache was being rebuilt.

Our DVD Store benchmark executed an 8 thread test with a heavy product lookup bias and a 0.08 think time on 1GB database.

Technology Validated:

We found HC3 to have class-leading data protection features (even among established vendors) and simplified workload management features that allow typical customers to very easily manage and configure workloads, see resource utilization and optimize the placement of VMs without complex graphs or rules of thumb. Bottom line, this represents consumer class simplicity with enterprise-level features, and we estimate this will reduce virtual workload management time and effort to 25% when compared to traditional virtualization vendors.

Following this migration exercise, we worked with HC3's snapshot functionality. HC3 snapshots are VM-centric and can be used to create a tree with over 5,000 snapshots. These snapshots can be used before a major event on a VM, such as software upgrade, to allow the VM to be reverted back to a known good state if needed. A clone can be created from any snapshot. Creating and working with HC3 snapshots was straightforward, used little I/O and was very quick.

Tuning and optimizing a virtual infrastructure can be a continual exercise in a growing business. Identifying what workloads are using which resources can be riddled with complexity, and taking action to correct bottlenecks or resource consumption issues can be fraught with risk. A vast majority of virtual administrators have horror stories about taking a seemingly sensible course of action to correct a performance issue, only to have a move of a VM to a new node break dependencies and run into entirely different bottlenecks. Scale Computing's HC3 cuts through these issues with a seamless pool of resources that cuts out the typical virtual infrastructure bottleneck, and then radically simplifies visibility into virtual resource utilization. With all this functionality, power and simplicity Scale Computing's HC3 look like a system that can put a business's IT infrastructure on cruise control.

Cluster Upgrade

Sooner or later any system will need to be updated and Scale Computing has made this as non-disruptive as possible by using rolling upgrades. When a system upgrade is available a notification is displayed on the management portal and by clicking the notification the upgrade will be downloaded and the installation of it will begin. In order not to disrupt the running of an enterprise's operations, all the VMs on a 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. It took us 49 minutes to upgrade all four nodes on our cluster, however the time will vary depending on the number of VMs running on the system and the size of the upgrade.

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

The main point here is that the upgrade was totally automated and our running applications did not suffer any downtime during the process.

Resiliency, Failure Tolerance, Recovery

To test what happens if an entire node fails we did just that; we pulled the power on a node. As a consequence, that failed node was dropped from the cluster, and the cluster entered a degraded state. The other nodes were still fully operational and the VMs that were running on the node were automatically restarted on the surviving nodes.

An HC3 clusters can experience two types of failures; either a single component, such as a drive can fail or an entire node can fail. Data protection within an HC3 cluster works by distributing mirrored data blocks across other drives on other nodes in the system ensuring no single node holds two copies of the data. When a single drive fails, the only data that is lost is one copy of mirrored data blocks. The node with the failed drive then enters a degraded state and other nodes begin immediately re-mirroring the data blocks to a fully protected mode. After the data is re-mirrored, the system exits the degraded mode. Node failure is dealt with differently than disk failure as Scale support has found more often than not when a node has gone down the server, not the storage is the issue and an HC3 cluster will not automatically replicate the data unless support has determined that the node is unrecoverable. By doing this they avoid unnecessary and expensive data rebuilds.

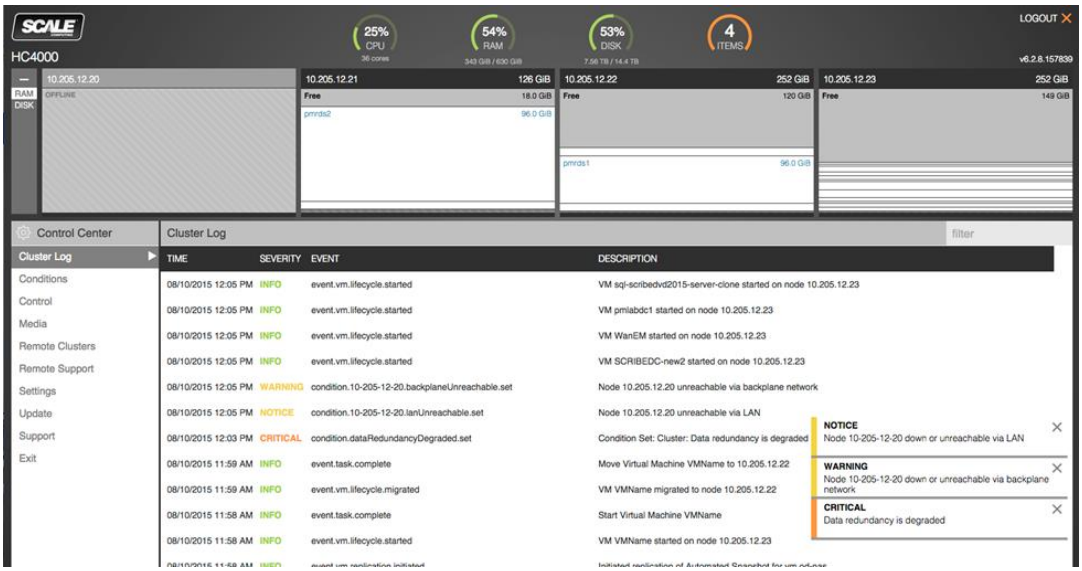

Figure 7: Cluster in "DEGRADED" mode demonstrating an integrated log with severity indications

During a storage rebuild, which happens in degraded mode, Scale Computing claims there is only a small impact on performance. We wanted to test this so we ran our degraded cluster under load. To do this we ran our Microsoft Exchange Server 2008 with LoadGen test on a node while the rebuild was going on. This VM was provisioned with 4 vCPUs and 8GB of memory and we ran the LoadGen test for 500 Outlook 2007 online users for 1 hour. The cluster we were using contained about 1 TB of VM data. We and we observed the CPU utilization and the disk IO happening on our HC3 node while the data was being redistributed using the "top" shell command. Top indicated a peak of 29% CPU utilization on the node during the rebuild process. As the node was still running benchmarking workload that was consuming 19% of the CPU we calculated that the rebuild itself used less than 10% of the node's CPU resources. We found that even in a cluster with a missing node, our Exchange LoadGen test ran flawlessly.

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

In addition to this test, we also examined what happens when a single drive fails, by pulling an active, in-use drive during an IOMeter test on the cluster. Our IO test file fit within the cache memory of the Scale Computing HC3 node - each node is equipped with cache - and as a consequence we saw no change in performance with our on-going IOMeter test. More importantly, as drive failure is not an uncommon event in the real world, we were keenly interested in how long a recovery from the loss of a drive would take. It took 37 minutes to re-mirror the data on a drive that was more than 80% occupied with data. This marks a class-leading recovery time compared to other systems that we have seen that use a traditional RAID approach to protect storage and must rebuild the entire disk volume - a process that can take hours or even days.

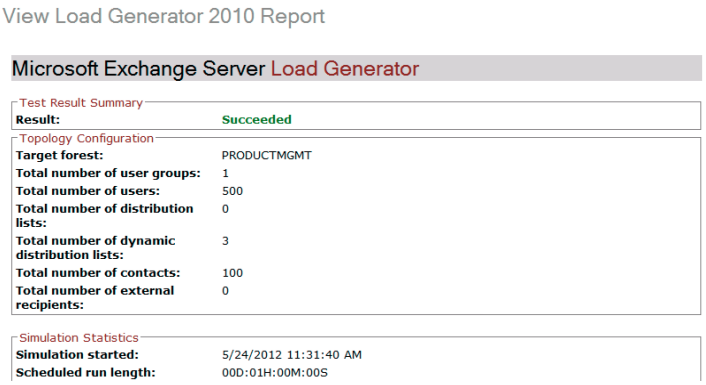
View Load Generator 2010 Report

Microsoft Exchange Server Load Generator

Test Result Summary

Result:	**Succeeded**

Topology Configuration

Target forest:	PRODUCTMGMT
Total number of user groups:	1
Total number of users:	500
Total number of distribution lists:	0
Total number of dynamic distribution lists:	3
Total number of contacts:	100
Total number of external recipients:	0

Simulation Statistics

Simulation started:	5/24/2012 11:31:40 AM
Scheduled run length:	00D:01H:00M:00S

Figure 8: Results from our LoadGen test on a degraded HC3 cluster that was reconstructing data after the failure of one node.

Disaster Recovery Using Built In Remote Replication

What happens when an entire site has an outage in the event of a disaster? To test how easy HC3 could handle this event we evaluated HC3's built in remote replication feature. Remote replication is as easy to configure and setup as the previously discussed snapshot feature.

Figure 9 demonstrates how easy it to set up a remote cluster to the primary site cluster. Once the remote cluster is recognized by HC3 any VM can now be replicated to that remote site. To invoke a replication relationship between two sites the administrator only needs to hit the VM snapshot icon and now a new option is available for setting up the remote replication target.

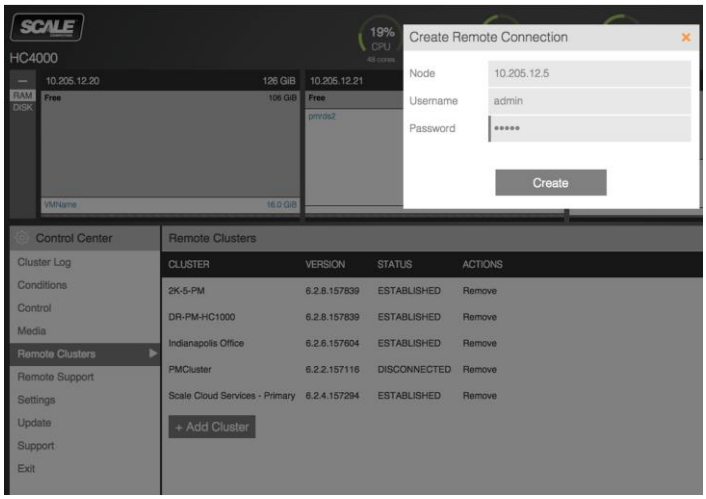

Figure 9: One Click to add a Remote Replication Cluster

Once the relationship is setup HC3 takes it from there by continually taking local snapshots at ~5 minute intervals and then replicating those to the remote site using a space and bandwidth efficient algorithm. HC3 will replicate in a way that automatically maintains the same snapshot versioning on the target HC3 as is on the local machine automatically. The target HC3 can be across campus or across country and gives a SMB significant flexibility in costs as the target machine can be sized to just provide infrastructure forth most critical VMs.

When the inevitable site outage hits, Scale's HC3 has made the failover and failback procedures almost child's play. Once the critical VMs are replicating to the remote site, failing over an application takes only two steps. First clone the latest replicated snapshot to prepare it to be an active VM and then boot the VM, which will enable the downed application at the primary site to

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

now be running at the remote location. It is that simple and this simplicity enables a business to easily test out their DR procedures that previously were too painful and costly to exercise. Figure 10 shows how easy it is to setup both the replication of a VM to remote site and also the cloning step needed to failover the VM at the remote site.

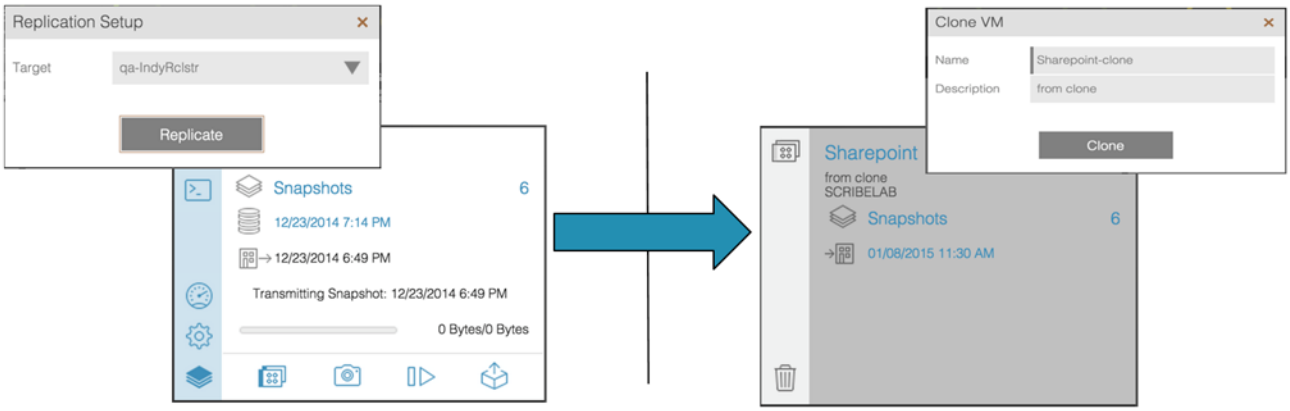

Figure 10: Illustration of replication setup then failing over to remote site after cloning

Typically, when a business plans for a disaster they tend to think about floods, tornados, fires and other acts of God. In reality, for most small businesses an outage is just as likely to come from a backhoe digging across your internet feed or an elongated power outage. With this in mind Scale's HC3 has implemented failback in very novel and unique way that radically reduces the RTO of the primary site. HC3 continuously keeps track of the snapshot relationship between both the primary and secondary site even after the initial failover event. Now when a business is ready to failback to the primary site HC3 can replay back to the primary site only those changes that have happened since the failover which in turn reduces the RTO significantly for the primary site. What this means is there is no penalty to initiate a failover more frequently, rather than waiting out the typical power or internet outage.

SCALE AND ADAPTATION

For the final element of our testing, we were looking forward to evaluating whether Scale Computing's reputation for scalable storage would bring benefits to an HC3 cluster. Traditional storage require either replacing entire storage silos or dealing with multiplying isolated storage silos - Scale Computing's HC3 allow customers to easily add additional nodes to bring on more capacity and performance. HC3 cluster even allow the intermixing of nodes with different performance characteristics (RAM levels, CPU types, and disk capacity) in order to tweak capacity and performance as business demands change

We then turned to the task of adding a new HC3 node to the cluster. We once again walk through the ASCII setup wizard, and configure the IP information so that the node can communicate with the existing cluster. This operation took us about two minutes and then the node appeared in the cluster and was ready for use.

In the past, scale-out storage products have solved some challenging problems and did away with painful customer experiences like storage migrations. This was great, but it still left a lot of complexity unaddressed - namely the scaling of compute. While storage scaling ended storage obsolescence, server sprawl continued on unabated, and each new server or hypervisor addition required a convoluted exercise around configuring the server's connection with storage. Scale-out

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

storage alone couldn't solve this. Now with HC3 the problems with coupling of storage+compute has been solved it. Scale-out simplicity remains, and the issue with server deployment is gone.

A big benefit of HC3's scale-out architecture is obsolescence avoidance and the constant requirement for disruptive migration events that are riddled with business impacts and risks. Scale Computing's clustering architecture is designed to allow the intermixing of different nodes, this applies to different generations of nodes as well. In turn, customers are not restricted to buying the same type of node over and over again. When a new node comes out, a customer can add this new node to their existing cluster, even if it is twice as powerful and has twice as much capacity. Over time, as the customer gradually adds nodes, they may choose to gradually retire their older nodes. It is easy to evacuate the data and VMs off of those nodes and it will have zero impact to availability. With an HC3 cluster disruptive hardware migrations and obsolescence will be a thing of the past.

With hyperconvergence, HC3 not only delivers scale-out storage, but by integrating a hypervisor and an easy to use management tools, it makes server configuration problems go away. Compute power becomes a seamless part of a single system that can scale on demand, with no need to ever again configure compute settings, interfaces, or connectivity. That's a massive simplification of the infrastructure.

OUR FINDINGS: A NEW GENERATION OF IT INFRASTRUCTURE

From our hands-on testing, HC3's capabilities look poised to transform a number of human-driven processes that currently take place in the traditional IT infrastructure. All of these processes come with significant financial costs. Below is a brief recap of our findings.

We Investigated	Our Findings
Time to setup versus traditional	10X faster
Time to configure post-setup to first VM deployment	8X faster
Time to deploy additional VMs in a growing environment with multiple hypervisors	4X faster
Workload protection	HA will restart VMs in the case of node failure
Management and ease of use	Radical consumer-class simplicity
Resilience	Industry leading, with an architecture that reduces points of failure and builds-in enterprise class failure tolerance
Disaster Recover with remote replication	Built in remote replication makes Disaster Recovery easy to implement and validate giving the ultimate peace of mind.
Scalability and adaptability	Easily scalable to adapt to changes, with a support for cross-generation scaling and retirement of nodes that puts an end to infrastructure system replacement and migration
Update system software	No application downtown was experienced during the upgrade process. System will alert you when an upgrade is available and the upgrade is initiated with a single click.

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

TANEJA GROUP OPINION

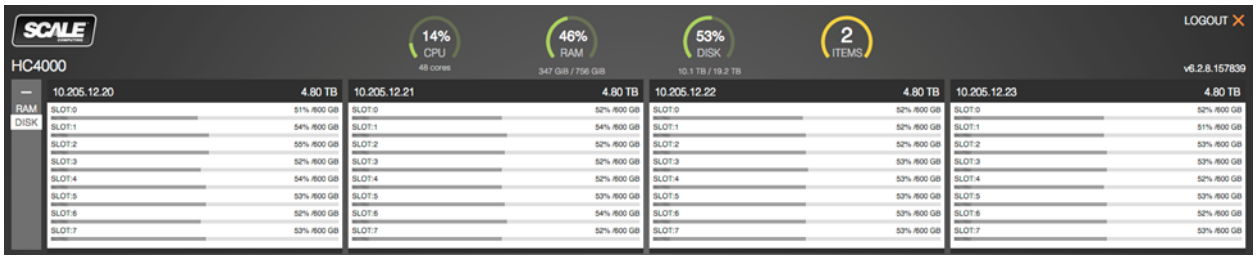

Figure 11: HC3 cluster interface showing top level dashboard

Seldom do we use screenshots to sum up our final thoughts, but a notable summary screenshot seems too tempting to pass up. Figure demonstrates the power of HC3 that is easy to overlook in the summary of its hyperconvergence features. The power of HC3 comes from the power of the unified, clustered platform that underlies it. Hyperconvergence makes it possible to supply both storage and compute on the same node. A Scale-out architecture makes it possible to add capacity to a cluster by adding more nodes. HC3s innovative management console make it intuitive and easy to manage.

This makes it possible to store all VMs in a single storage pool where they are accessible by all nodes of a cluster. When the nodes of a cluster are simultaneously capable of providing storage capacity and running virtual machines, the infrastructure is suddenly much more versatile and easy to use. This is the power that overcomes the complexity inherent in traditional virtualization, and makes HC3 into a truly hyperconverged platform that can provide storage and also serve up compute too and without any additional complexity.

Clearly an HC3 cluster can have a big impact on the IT infrastructure practices of its users, and we believe it will have a transformational impact on their cost as well. Towards the end of our testing we were running 18 VMs on four HC3 nodes. These VMs were ranging between 4 and 16GB of RAM and were running production workloads that ranged from simple desktop to heavily loaded Exchange servers. In a traditional non-virtualized infrastructure using best practices each one of these workloads would be serviced by a single server. By combining these 18 servers on a single four node HC3 cluster the physical servers and their accompanying infrastructure is not needed. In total, this could represent a cost savings of $15,000 to $40,000. An HC3 cluster is by design built to be more affordable than a traditional non-virtualized datacenter and many competitive hyperconvergence systems on the market, and may represent a cost savings of as much as 50% versus systems with similar capabilities, especially when scale-out features and obsolescence protection are considered. Given those numbers an HC3 cluster starts to rapidly look like a solution that can deliver an entire infrastructure for perhaps as little as 1/3rd what it might traditionally cost. This is without considering operational cost advantages, such as less management time and effort (the high spots of which we detailed in Table 3), or less floor space, power and cooling.

What Scale Computing has done with HC3 is nothing short of a radical simplification of the IT infrastructure. In many cases, when vendors try to radically simplify a complex system, they find there's a reason for the complexity, and as a consequence they fail because they've given up functionality for the sake of simplicity. In the case of a Scale Computing HC3 cluster, there appears to be no compromise in any of the IT infrastructure capabilities that SMB and SME customers need. As we looked at HC3, we kept a careful eye toward how the typical business uses virtualization capabilities like workload monitoring, rapid provisioning and reuse, and data protection. For mainstream users, Scale Computing delivers a comprehensive set of these features.

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

Once again, this is largely because Scale Computing invested tremendous time and effort in perfecting a radically simplified scale-out storage solution. The polish of that storage layer supported the introduction of HC3 as a hyperconvergence appliance, and is allowing Scale Computing to deliver a solution that is more capable than most infrastructures that could be pieced together from separate servers and storage from any major vendor. In some ways, this is not surprising, as storage shortcomings have long been the downfall of the datacenter, irrespective of whether the datacenter is large or small. For Scale Computing's customers, we're left with no doubt that HC3 will fundamentally alter how IT is done, for the better.

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