SCALE COMPUTING HC3: A QUALITATIVE EVALUATION

HYPERCONVERGED INFRASTRUCTURE

Simplicity is valued in nearly every area of enterprise IT but not often easy to deliver effectively. The past decade has seen the rise of a new kind of IT solution from which compute, storage and networking can be simplified into a single, easy-to-use system, and they can be delivered with enterprise-class scalability and reliability, too.

Hyper-Converged Infrastructure (HCI) consolidates storage, networking and virtualized workloads into an easy-to-deploy, appliance-like experience. Extending beyond these basic, functional blocks, HCI also includes tools for rapid deployment, backup, disaster recovery, simplified management, lowering the total cost of ownership (TCO) and embedded clustering technology to provide in-the-rack, high availability.

Enterprise IT has escaped the confines of the physical datacenter. Today's IT architecture blends compute and storage capabilities on-premise, in the cloud and at the edge. While enterprises have adopted HCI into nearly every facet of IT infrastructure, HCI is well targeted at physical locations where local IT support is at its weakest - this includes retail establishments, banks, small offices, remote/branch offices and other "edge"-like locations where local IT specialists are sparse and ease-of-experience and remote manageability become critical to functionality. The ease of remote manageability directly affects the total lifecycle cost of the solution from deployment through upgrades to long-term maintenance.

At the same time, enterprise workloads where the HCI model simply makes sense are available. HCI is ideal for applications where enterprises wish to deploy discrete workloads with predictable resource requirements. VDI deployments, for example, where a single HCI cluster can serve high numbers of remote desktop users, are a natural fit for an architecture that can be racked, configured and left alone. HCI deploys file and print services, departmental databases and even dev-ops within the enterprise.

The excitement generated by this market has enabled a number of competing solutions, each with positive and negative attributes. This paper delves deep into one such solution, evaluating the HC3 option delivered by HCI pioneer Scale Computing.

SCALE COMPUTING

Scale Computing is one of the earliest innovators in hyperconverged infrastructure. Its founding in 2007 predated the term "HCI" by a number of years. Scale Computing launched its HCI solution, HC3, in 2012 where it quickly found a home in the small and medium business market. It has since found success in the enterprise, where its time-to-value and intrinsic simplicity promise to drive economic benefits well beyond the technology itself.

More than a decade after its founding, Scale Computing continues to deliver innovation into the HCI market. This is a strong indication of Scale Computing's competitiveness and fit within the markets that it serves. But is Scale Computing a truly competitive and differentiated solution? Let's explore what HC3 offers, look at its architecture and then install and use the system to understand its true benefits and competitive differentiation.

SCALE COMPUTING AND HCI

HC3 is Scale Computing's HCI solution. HC3 is an integrated solution, managed through a simple and intuitive web interface, with nodes within a cluster that coordinate with each other to monitor the overall state of the system.

HC3's robust feature set includes the following:

- HC3 to HC3 System VM Replication
- VM Snapshot Scheduling
- Bulk VM actions
- VM Cloning
- Non-disruptive VM live migration between nodes
- VM and System performance monitoring
- Automatic VM failover in a node failure scenario
- Automatic data restriping in the event of a failed disk
- Self-monitoring and self-healing
- Non-disruptive software upgrades
- No single point of failure
- Scale-up system resources with no downtime required

HC3 HyperCore Architecture

The Scale Computing HC3 operating software, HyperCore, takes an innovative (and divergent) path to deliver enterprise-class, failure-tolerant, scalable, clustered HCI.

Software-Defined Computing

HyperCore's virtualization engine is built on the Kernel-based Virtual Machine (KVM) hypervisor, a bare-metal (i.e. "type 1") hypervisor integrated directly into the operating system within each node. KVM is a well-known and widely deployed, open-source hypervisor within the Linux community.

Scale Computing, by marrying KVM with its own internally developed technologies, has delivered a KVM virtualization solution that is both unique and competitively differentiated in the market. Beyond feature-level differentiation, the use of KVM allows Scale Computing to deliver an HCI solution without the common hypervisor "tax" that proprietary HCI solutions often require -- allowing enterprise IT to spend that budget elsewhere.

Two examples of this differentiation follow below:

Capability-based VM Placement: HyperCore's global storage pool (implemented Scale Computing Reliable Independent Block Engine (SCRIBE)), described in the following section) allows for VMs to be placed on any node within an HC3 cluster that contains the capabilities required by that VM (e.g. RAM, vCPUs, etc.).

Performant VM Live Migration: Virtual machines within HyperCore can be migrated without temporary termination and without impact to either the workload within the VM or the clients connected to that virtual machine.

Software-Defined Storage: SCRIBE

The core technology within HyperCore is SCRIBE, Scale Computing's software-defined storage technology. SCRIBE is described by Scale Computing as an enterprise-class, clustered block storage layer that is embedded directly in the KVM hypervisor. The job of SCRIBE is to manage the data both within each node and across the cluster.

SCRIBE's operation is sophisticated as compared to the block storage layer within other HCI solutions. SCRIBE creates an aggregated global storage pool comprised of the available storage within the block storage devices resident on each node within the

cluster. Each node of the cluster has access to SCRIBE's storage pool, allowing for sophisticated data management functions such as data redundancy, deduplication and load balancing to be used by the rest of the system.

The SCRIBE architecture delivers a compelling competitive advantage to Scale Computing HC3. First, SCRIBE deduplication technology provides little-to-no-impact on VM workloads by prioritizing customer workloads over deduplication efforts without sacrificing dedupe performance or effectiveness. The mechanisms allowing this trade-off are integrated into SCRIBE and SCRIBE's integration into KVM allowing for deduplication performance and efficiency that is difficult for competing solutions to match.

Second, SCRIBE does not build "virtual hard disk" files that are resident on local filesystems, removing the intrinsic inefficiencies of that approach.

Third, RAID partition alignment problems common on external RAID systems, and which have dramatic negative performance impacts, do not occur.

Finally, VM-snapshots are complete, leveraging SCRIBE's inherent deduplication and other features in order to avoid implementing VM snapshots as delta files; delta-file snapshots have severe negative performance impacts caused by I/O thrashing during delta-merge operations.

HyperCore uses SCRIBE to integrate block storage objects directly into the hypervisor. This allows for benefits that, again, provide significant architectural and performance advantages to the Scale Computing HC3 solution. Resident VMs have direct access to SCRIBE virtual storage devices, eliminating the complexity, performance and latency impacts found in competing solutions which rely on remote, non-resident storage. Also, native block storage avoids the common "Virtual Storage Appliance" found in nearly all competing solutions, eliminating latency and providing performance equivalent to zero-copy shared memory.

The work that Scale Computing has delivered with SCRIBE provides intrinsic performance and simplification gains. The efficiencies within SCRIBE also enable Scale Computing to deliver similar functionality as competing solutions with required processor and memory, allowing HC3 to run on smaller form-factors (as low as 8GB per node in some cases) to allow edge deployments with lower investments. This adds up to a compelling competitive advantage over similar solutions in the HCI market.

HIGH AVAILABILITY

Downtime in any enterprise IT environment, whether in the datacenter or at the edge, costs money. Delivering business-critical workloads requires robust, high-availability capabilities. Scale Computing's reliability features deliver exactly the kind of functionality that enterprises look for in a solution.

HC3 is designed to maintain operation during multiple failure types: disk failures, node failures and loss of network link connectivity. Let's take a brief look at each of these.

DISK FAILURE TOLERANCE

Disk failures in HyperCore have little effect beyond the temporary loss of the capacity provided by the failing disk. HyperCore, upon detecting a failing disk, will redirect I/O's to a mirrored device.

HyperCore goes a step further and creates a new mirrored copy of the disk from its available capacity. New incoming writes will be mirrored to the new copy, providing a second level of redundancy.

This approach contrasts positively with competing solutions that use RAID technology for data protection. RAID requires expensive and I/O intensive scans of the data to determine which data was lost by the disk failure. HyperCore, in comparison, simply uses the known-good, mirrored copy and creates a new mirror to continue serving highly available data.

NETWORK FAILURE TOLERANCE

All HC3 nodes within a cluster offer redundant network ports for both public LAN and private backplane network for full network redundancy. These ports are in an active-passive bond that allows for port-level failover.

A failed port is replaced by its healthy backup with nearly no disruption to the system node. Once a failed port is detected as restored (these failures are often caused by networking impacts outside the node), traffic is switched back to the restored port.

ENTERPRISE MANAGEABILITY

Delivering the total value of an HCI solution requires a management model and user experience grounded in ease-of-use while maintaining power enough to meet the needs of an expansive and distributed enterprise. Scale Computing HC3 delivers just this combination.

HC3 systems are manageable via a web-based GUI accessible on any cluster node, without the need for a separate management server that many competitive solutions require. Scale Computing provides a Multi-Site Consolidated Management server and UI optimized for monitoring hundreds or thousands of HC3 systems to allow managing HC3 installations across distributed sites. HC3 also provides REST API and has developed plugins for several IT orchestration tools, including Ansible, to support customer automation needs.

Integration with Ansible and Terraform offer the fleet management tools needed in retail, healthcare, manufacturing and other distributed environments where hundreds or thousands of clusters may exist. From staging clusters prior to deploying at the edge, to rolling out new versions of workloads to the various sites (VM life-cycle management), these tools allow customers to script and orchestrate the execution and apply changes to the HC3 environment.

SCALE COMPUTING HC3 PRODUCT FAMILY

The variety of environments into which HCI is being deployed today is broad. Scale Computing's product strategy is equally broad to best target each segment of the market. Its HC3 offerings come in three series, each with a number of options for performance and capacity:

- **HC3 Edge Series (HE150 and HE500),** designed for edge computing, small remote sites and disaster recovery.
- **HC1000 Series**, designed for data center infrastructure, edge computing, remote site and disaster recovery.
- **HC5000 Series**, designed for large data storage and ideal for data center infrastructure and disaster recovery.

Beyond these fixed configurations, Scale Computing also offers flexible options for edge deployments. The company provides design and consulting services for edge deployments, both directly and through a number of partners. This is a strong offering, as edge deployments, especially across large deployments, often need a little extra attention to ensure a smooth and successful implementation. Scale Computing delivers that extra attention.

EVALUATING SCALE COMPUTING HC3

HCI is about simplicity. This simplicity is inherent in converging functionality into an appliance-like experience. More critically, supplying a stream-lined and intuitive management experience brings simplicity.

HCI is deployed within the datacenter but increases at the edge: in retail environments, small offices, remote offices and similar worlds where a self-contained solution targeted at serving workloads with a single, easy-to-use interface removes the complexity of traditional enterprise IT.

The simplicity and ease-of-use promised by HCI directly yields IT efficiencies in deploying and managing solutions both inside of the traditional datacenter as well as in the field. These IT efficiencies yield reduced operational expenses and minimize time-to-success when managing and deploying HCI.

With that view of simplicity and ease-of-use in mind, we drive the evaluation of the Scale Computing HC3 system. This evaluation looks at the most impactful criteria for deploying HCI solutions into its most common environments:

- Ease of installation, including installation options provided by Scale Computing.
- Ease of use of the management interface.
- VM Operation & Provisioning.
- Data Protection.

We do this by walking through the installation, configuration and workload deployment on a three-node HC3 cluster.

This evaluation does not delve into quantitative benchmarks, as those are workload dependent, and the rich variety of HC3 configurations are as competitive as any solution on the market.

INSTALLING & CONFIGURING HC3

Installation of an HC3 cluster is a straightforward process. The nodes comprising the HC3 cluster are received from Scale Computing with a version of its HyperCore software already loaded (but not configured), along with scripts and instructions that allow an IT engineer to drive the installation. The physical hardware is installed, the network connected and the software configured. At which point, if all goes well, the cluster is ready-to-use.

INSTALLATION SERVICE OPTIONS

Scale Computing recommends that most users let Scale Computing's ScaleCare installation specialist install and configure the software for a new cluster. Moor Insights & Strategy supports this good recommendation. When available, we nearly always recommend a vendor-assisted installation for systems of any complexity or criticality, as it's usually the safest and fastest path to success.

There are two options provided by Scale Computing for assisted installations:

- Remote software configuration, where the customer installs the hardware and configures the network, and the Scale Computing installation specialists do the rest.
- Premium, on-site installation by Scale Computing's professional services team.

Remote installation of the software is most-often the best option for organizations who have an IT staff to assist. Remote installation requires the customer to rack the systems and configure the networking (described in the following section *HC3 Physical Network*), at which point Scale Computing's engineers take over the process.

Customers can also purchase on-site installation, where Scale Computing provides on-site assistance for each step of the process. On-site installation and configuration services are available for the HC3 cluster as well as for network switch configuration, if that's needed. This is a valuable option for smaller organizations who may not have IT staff available or for any organization who is deploying an HC3 cluster in a remote location where technical staff may not be readily available.

Whether one chooses remote installation or on-site installation, Scale Computing allows customers to manage the entire process through a simple and easy-to-use web portal. Once a customer submits a request for installation, a well-defined process begins with a ScaleCare engineer reaching out to schedule a planning call and to guide the new customer through the process.

More adventurous types can choose to do a self-installation, of course, going against both our advice and that of Scale Computing. For this evaluation, Moor Insights & Strategy did just that, with the hope that it leads to the most complete understanding of the Scale Computing HC3 product.

HC3 Physical Network

Configuring the networking for Scale Computing's HC3 is probably the most critical task in installing the system. This isn't to say that the process itself is difficult – it's actually very straightforward. After reading the brief documentation from Scale Computing and understanding the HC3 networking topology, most experienced IT engineers should have no problem quickly cabling the system and communicating network address needs to the networking team.

HC3 requires two logical networks be defined. One network is dedicated to servicing typical network LAN traffic with a gateway to external networks for regular networking traffic. Note the presence of another network, termed the "backplane" network, which is used for what is traditional, out-of-band, cluster-level communication between nodes.

Figure 1: High Availability Configuration

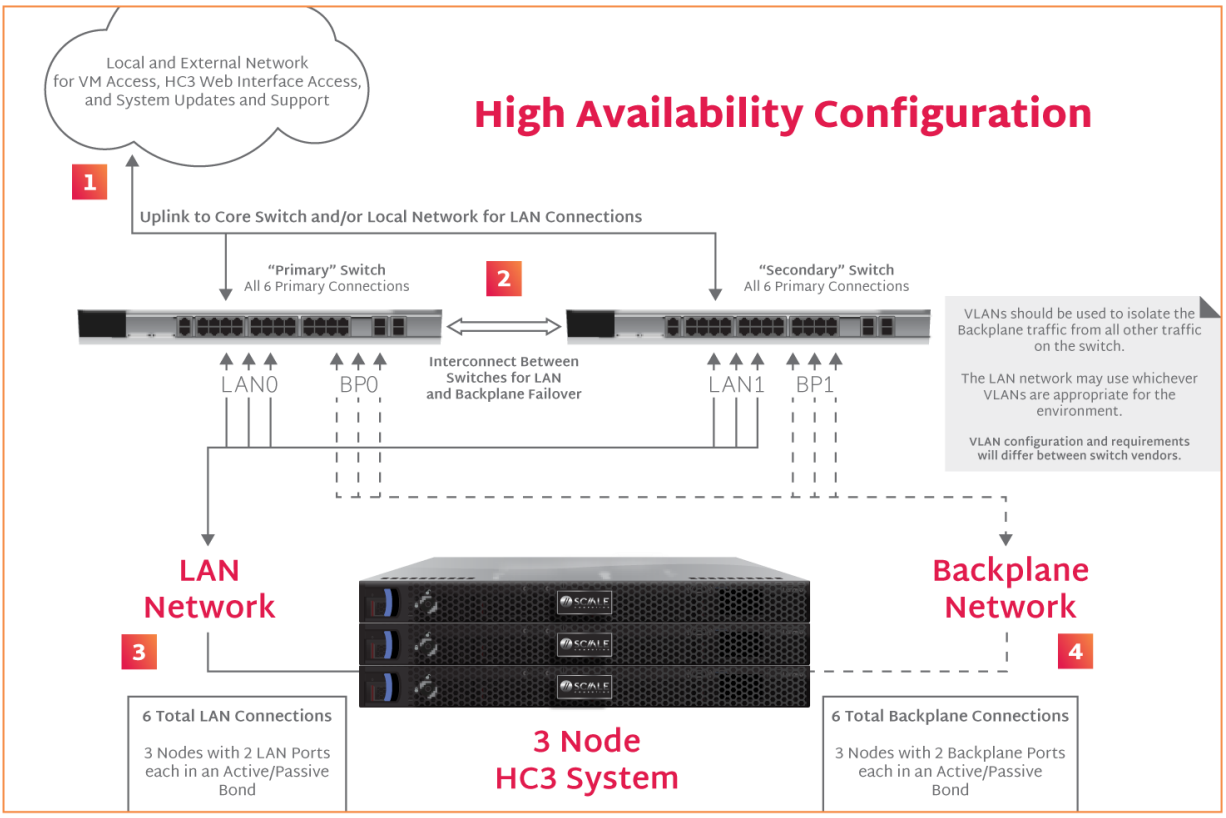

Source: Scale Computing

These two networks can be partitioned flexibly with the use of VLANs, can live within separate subnets on a common set of switches or can coexist within a single network address space[1]. The illustration above from Scale Computing's documentation illustrates a typical configuration[2] and shows what to expect.

The HC3 installation guide provides both required and recommended functionality for the switches to which the HC3 nodes will be connected. The following requirements are easy to meet, because most are typical for modern ethernet switches and routers:

- It must be a managed switch
- It must have support for Spanning Tree Protocol (STP).
- It must have support for 802.3x flow control.
- Optional VLAN support (but which makes it easy to provide robust isolation between LAN and Backplane Networks).
- Optional per-port ability to enable/disable to STP protocol.
- Optional rapid spanning tree protocol (RSTP).
- Optional stacking or dedicated interconnect for High Availability.

HyperCore relies on the networking infrastructure's support for the spanning tree protocol (STP/HSTP) to mitigate any "looping" issues that may occur where multiple network paths exist between nodes. Looping is not an uncommon problem in high-availability and cluster configurations where nodes can be reached through multiple ports or switches.

Enabling STP/HSTP in these environments is a good practice and, in this case, allows HyperCore to utilize backup and other redundant ports without worrying about looping effects. We like that Scale Computing defines these requirements as part of the installation[3].

We also like that the installation guide provides a list of tested switches, untested-but-known-good switches and switches that aren't recommended at all for assorted

[1] Moor Insights & Strategy nearly always recommends VLAN isolation for separate networks, especially when one of the networks is for inter-cluster communication.
[2] The configuration tested in this evaluation is nearly identical to the illustrated example from Scale Computing.
[3] Some competing solutions only discuss STP and looping in their troubleshooting documentation – we strongly believe that it's better to configure safety mechanisms from the beginning to avoid triggering any reason to read troubleshooting guides!

reasons. This is much more useful for most IT practitioners to adhere to than simply following a list of "supported" configurations[4].

Once the systems are mounted in the rack with their network connections in place, the HC3 software can be installed and configured.

Software Installation: Networking & Cluster Initialization

The first step of configuring the HC3 software is to assign the allocated network addresses for both the LAN and Backplane networks to each of the nodes within the cluster.

This step must occur whether the cluster is being installed by the customer or the customer is using remote installation.

The process is straightforward. Once a VGA-capable monitor and USB keyboard are attached to the back of a node, the user logs-in to the console and simply follows the prompts to enter the addresses described in the physical network section above. The software takes care of the rest. This happens for each node in the cluster.

At this point, if remote installation is chosen, Scale Computing's ScaleCare installation specialist will take over the operation and complete the configuration.

As we indicated above, our evaluation includes a full manual installation, so we'll continue forward.

Once the network addresses are configured, the monitor and keyboard are reconnected to the node chosen to be the "master" node in the cluster, and a similar process is followed. The user logs into the node which displays a message showing the "run" command.

Typing in that command starts a long string of automated operations that provide baseline configuration of the cluster across each of the nodes. If a problem is encountered, then you will quickly understand why you should have chosen assisted installation, as the terminology in the messages assumes a solid familiarity with the operation of the system.

[4] As with most vendors, network issues that are experienced with switches not on the "recommended" list may require ScaleCare Support to ask you to install one if basic troubleshooting doesn't resolve the issue. On-site support also requires the use of a supported switch.

Assuming no issues were encountered, then the cluster should quickly be ready to manage through the web interface. The web-based management interface can be accessed by navigating to the address of any of the nodes within the cluster.

This process could not have been any easier.

Management Interface & Operation

One configures an HC3 cluster via a web interface accessed by entering the address of any of the nodes within the cluster. The web interface allows for configuration and reporting, resource monitoring, user management, VM configuration and management, as well as a number of support options. Next, we provide an image of the configuration screen:

Figure 2: HC3 Cluster Configuration Screen

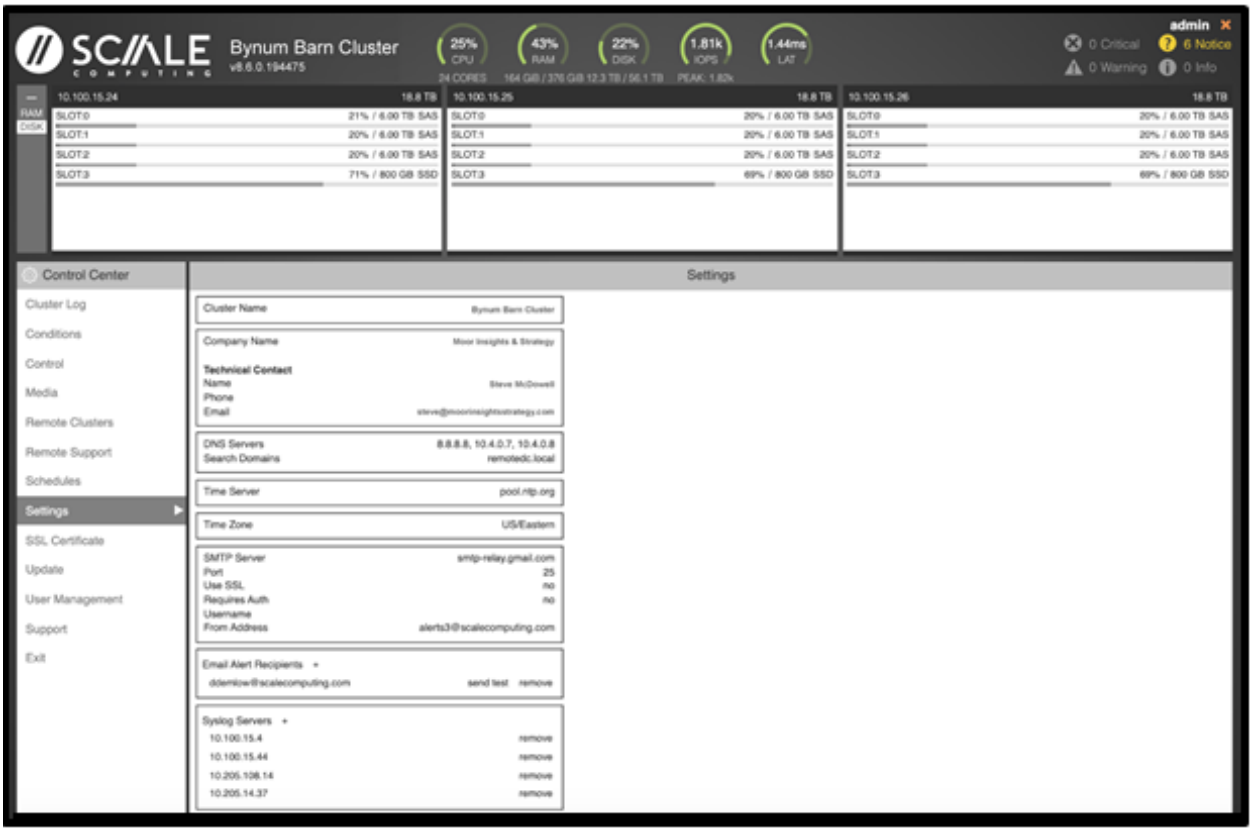

Source: Moor Insights & Strategy

The interface is clean and minimal, with a dashboard + three-panel design that focuses clearly on the items required for configuration and operation. In this example, the resource utilization for each of the three nodes in our evaluation cluster is easily identified across the top of the view (this is consistent regardless of the option selected in the UI), while available operations are detailed along the lower-left pane and operation-dependent content in the larger lower-right pane.

The Scale Computing HC3 user interface is designed with the needs of enterprise IT in mind. The ability for an IT administrator to quickly and simply deploy, manage and troubleshoot an installation are critical factors when choosing an HCI solution. The HC3 user experience will lead to reduced “hands on” time, delivering enterprise IT efficiencies and minimizing precious operational expenses in the process.

We like the HCI interface. It very quickly identifies the basic operations that an administrator requires through easy-to-navigate options. This simplicity and ease-of-use are critical for the user experience in a system like HC3. A cluster should be configured and periodically managed but otherwise left alone to do its job for extended periods of time.

FIGURE 3: OVERALL MANAGEMENT SCREEN TO CONFIGURE VIRTUAL MACHINE

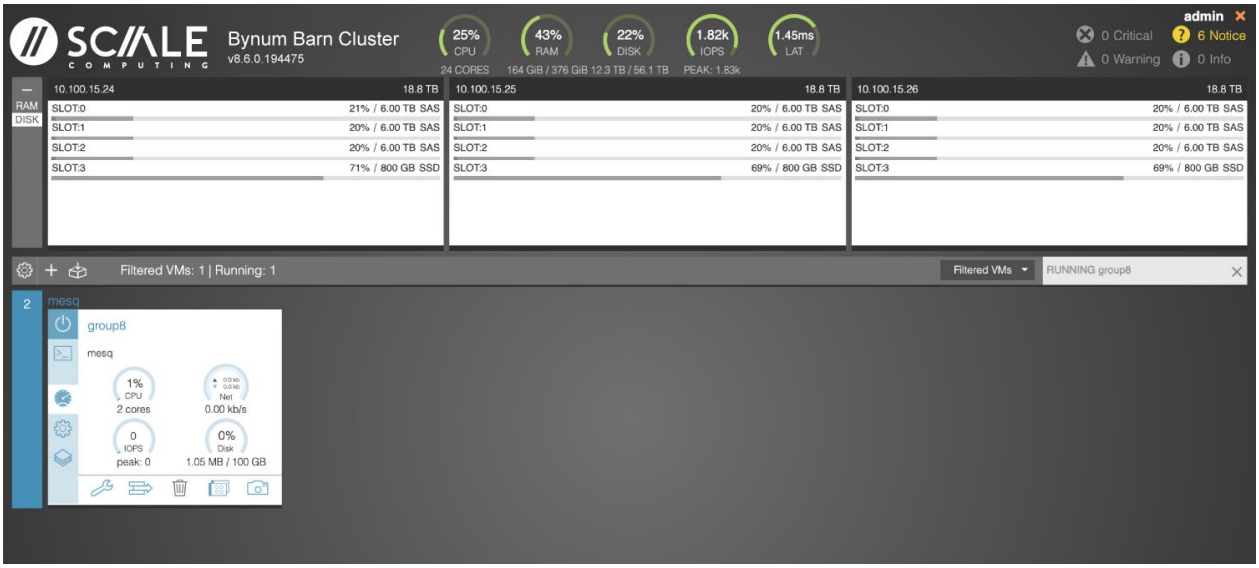

Source: Moor Insights & Strategy

VM MANAGEMENT & PROVISIONING

The cluster is provisioned with a set of intuitive, UI-driven operations. An HC3 cluster has a single pre-defined virtual machine (VM) waiting for the user to configure it. This is where an administrator will spend the majority of their time.

Exploring the various options on the VM itself yields the following set of potential options:

FIGURE 4: VIRTUAL MACHINE CONFIGURATION OPTIONS

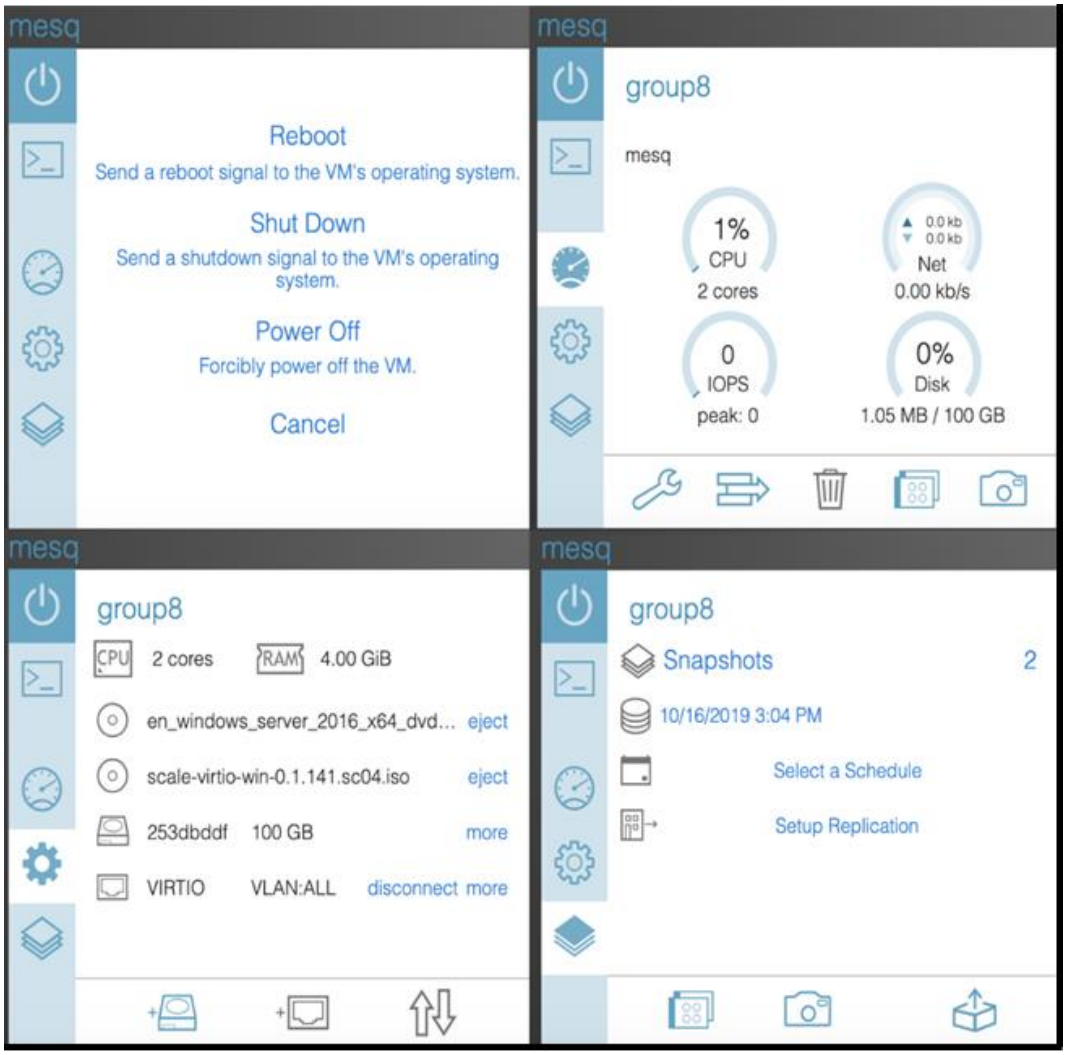

Source: Moor Insights & Strategy

The toolbars along the bottom provide for operations related items selected on the right. These management functions include these items:

- VM Naming and Tagging (tags are searchable through the UI).
- CPU, Memory, Networking and Storage management for the VM.
- VM Migration across nodes.
- VM Cloning.
- VM Snapshots, both on-demand and scheduled.
- VM Web Console.

Every management function and configuration parameter found in a typical VM installation is present. The VM is complete, intuitive and easy-to-use.

DATA PROTECTION

Data protection is implemented in HC3 with snapshots and replication. Snapshots provide the primary means of data protection within HC3. These snapshots are available in multiple options:

- Local one-time VM snapshot.
- Scheduled periodic VM snapshots.
- Snapshot to remote cluster.

Snapshots provide protection for the VM and its data while also allowing for the ability to be restored quickly. Beyond full VM recovery through cloning, HC3 also includes the ability to clone individual virtual disk from snapshots. These virtual disks can then be mounted to running virtual machines, allowing the administrator to quickly navigate the volume to find a file for recovery.

The HyperCore replication within HC3 provides cluster-to-cluster replication for VM availability during disaster recovery scenarios. The replication within HC3 is handling at the per-VM level, allowing for traditional source-to-target replication, as well as cross-replication, in which a target replicates back to the source system.

In this age of ransomware and other dangerous service disruptions, the ability for local and remote data protection is absolutely required. Scale Computing's implementation of this functionality is full-featured and competitive.

OTHER MANAGEMENT FUNCTIONS

Beyond what has already been described, other management functions within the management interface, typically present in an HCI system, include the following:

- ACL-based user management, including Read-Only, Back-Up Only and other task-driven user personalities to be defined.
- Remote system management, for managing multiple clusters in a disaster-recovery grouping.
- Alerts and logging, for intelligent monitoring and reporting.
- System maintenance functions, including software update, rolling updates across nodes and remote support capabilities.

SUMMARY

The Scale Computing HC3 solution is strong, capable and extremely competitive. Much of this competitiveness derives from a simplified, yet complete, user experience designed with enterprise IT in mind. Scale Computing's assisted installation, whether remote or on-site, removes the pain of deployment for the company, and the simplicity of the HyperCore user interface allows for intuitive operation and quick installation.

HC3's intuitive user-experience drives overall, faster time-to-deployment for new installation while allowing for seamless manageability and troubleshooting for existing deployments. These attributes should lead to lower operational expense and increased IT efficiency and productivity.

The intrinsic functionality of the HC3 system leads to differentiation with Scale Computing and HC3, beyond the user experience:

- HC3 scalability and availability deliver a favorable and compelling TCO model.
- HC3 is available across a broad spectrum of configurations, leading to deployment targets ranging from the edge to an enterprise's core datacenter.
- HC3's scalable capabilities lead to natural workload targets in VDI, file and print, and the sorts of general-purpose workloads expected in edge and remote deployments.
- At the same time, the HC5000 series is ideal for serving enterprise applications across the performance spectrum.

Architectural distinctions are what provide Scale Computing long-lasting and difficult to disrupt competitive differentiation. The smart and strategic use of the most widely adopted open source hypervisor, KVM, allow Scale Computing's solution to slide seamlessly into the enterprise, all without the burden of the hypervisor "tax" that accompanies many of the competing solutions.

The HyperCore architecture itself provides a set of intrinsic set of capabilities that provide for unmatched, cluster-wide scalability and performance that cannot be matched by a competing system without significant reengineering:

- SCRIBE, Scale Computing's innovative and unique block storage architecture, directly integrated into KVM.
- The global memory pool with cluster-wide shared state.

Modern enterprise IT architecture increasingly blends traditional, on-premises capabilities with public cloud and edge capabilities. Time-to-deployment and ease-of-manageability, each of which drive cost efficiencies, are now core criteria in choosing solutions to meet the needs of enterprise IT.

Scale Computing HC3 delivers all the essential elements to meet the demands of enterprise IT. The Scale Computing portfolio of solutions deliver the consistent and scalable experience required to any enterprise HCI need.

IMPORTANT INFORMATION ABOUT THIS PAPER

CONTRIBUTOR

Steve McDowell, Senior Analyst at Moor Insights & Strategy

PUBLISHER

Patrick Moorhead, Founder, President, & Principal Analyst at Moor Insights & Strategy

INQUIRIES

Contact us if you would like to discuss this report, and Moor Insights & Strategy will respond promptly.

CITATIONS

This paper can be cited by accredited press and analysts but must be cited in-context, displaying author's name, author's title, and "Moor Insights & Strategy". Non-press and non-analysts must receive prior written permission by Moor Insights & Strategy for any citations.

LICENSING

DISCLOSURES

This paper was commissioned by Scale Computing. Moor Insights & Strategy provides research, analysis, advising, and consulting to many high-tech companies mentioned in this paper. No employees at the firm hold any equity positions with any companies cited in this document.

DISCLAIMER

The information presented in this document is for informational purposes only and may contain technical inaccuracies, omissions, and typographical errors. Moor Insights & Strategy disclaims all warranties as to the accuracy, completeness, or adequacy of such information and shall have no liability for errors, omissions, or inadequacies in such information. This document consists of the opinions of Moor Insights & Strategy and should not be construed as statements of fact. The opinions expressed herein are subject to change without notice.

Moor Insights & Strategy provides forecasts and forward-looking statements as directional indicators and not as precise predictions of future events. While our forecasts and forward-looking statements represent our current judgment on what the future holds, they are subject to risks and uncertainties that could cause actual results to differ materially. You are cautioned not to place undue reliance on these forecasts and forward-looking statements, which reflect our opinions only as of the date of publication for this document. Please keep in mind that we are not obligating ourselves to revise or publicly release the results of any revision to these forecasts and forward-looking statements in light of new information or future events.