HC3®

Features Guide

What is HC3?

The HC3 virtualization platform brings together servers, storage, virtualization, and disaster recovery into a single, feature rich solution. This guide describes all of the features that make HC3 the perfect infrastructure solution for any datacenter.

HYPERCORE

The heart of the HC3 platform is the HyperCore operating system which includes a fully integrated KVM-based hypervisor with our patented block access, direct attached storage system called SCRIBE. SCRIBE includes HEAT to automate tiering across our own hybrid flash storage architecture.

SCRIBE

SCRIBE is a wide-striped storage architecture that combines all disks in the cluster into a single storage pool that is tiered between flash SSD and spinning HDD storage. Blocks are striped across all nodes in a cluster redundantly to protect both against individual drive failure and node failure. The use of every drive through

wide striping gives a performance advantage to every VM on the cluster. Performance is also enhanced through direct block access because of the direct integration between the storage and the hypervisor resulting in Hypervisor Embedded Storage. There are no inefficient storage protocols that would normally be found in SAN or NAS-based storage solutions including Virtual Storage Appliance architectures.

HEAT

HyperCore Enhanced Automated Tiering is the part of SCRIBE that manages data efficiently between the flash SSD tier and the spinning HDD tier in the cluster storage pool. It prioritizes blocks of data on SSD or HDD based on a I/O heat map that
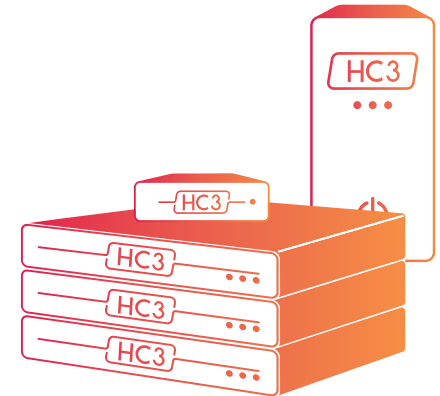
tracks I/O against each virtual disk. Although HEAT is primarily automation happening behind the scenes, it provides the ability for the user to tune the relative priority of every virtual disk in the cluster to further optimize the SSD usage where it is needed the most. You can literally turn it to 11.

Click here for a video demo.

HC3 WEB INTERFACE

Multiple The entire HC3 solution including hardware and software is managed from a single web interface that can be accessed from any node in a cluster. The simple interface design allows both storage and compute resource management and virtual machine management from a single

Simplicity Engineered // scalecomputing.com // +1 317-856-9959

SCALE COMPUTING

pane of glass. Resource statistics are displayed in real-time along with IOPS, both per VM and cluster wide. VM consoles are also immediately available from the web interface.

REMOTE CLUSTER MANAGEMENT

Multiple clusters or single nodes across multiple sites may be managed within the HC3 Web Interface. The remote systems can be monitored for errors, warnings, or other issues at a glance. Manage dozens of sites from a single interface for simple, scalable management.

MULTI-USER MANAGEMENT

For clusters and sites with multiple administrators, HC3 features multi-user administration to allow each administrator a unique account and login. Multiple administrators can be managed with unique logins and activity logging to help troubleshoot issues based on administrator activity.

VM HIGH AVAILABILITY

If a node fails within the cluster, all VMs running on that node are automatically failed over to one of the remaining nodes. Failover happens within minutes for minimal disruption. After the node has been restored, VMs are non-disruptively live migrated back if desired.

Click here for a video demo.

DEDUPLICATION

SCRIBE includes data deduplication that eliminates duplicate blocks on virtual disks, significantly increasing disk capacity. Deduplication is done with minimal impact to running VMs. Reducing the storage footprint of data to increase your data storage capacity.

LIVE VM MIGRATION

VMs on HC3 clusters can be non-disruptively migrated between nodes with no downtime. This not only allows for rebalancing resource allocation across the cluster but also allows VMs to be relocated automatically during our rolling update process for the HyperCore OS firmware.

NON-DISRUPTIVE ROLLING UPDATES

HC3 receives software/firmware updates directly to the HC3 Web Interface where they can be applied automatically with no downtime. A single click updates the hypervisor, the storage system, the firmware, and any other part of the complete HC3 system. The automated update process will automatically relocate VMs between nodes to free up the nodes, one at a time, to be updated. When all of the cluster nodes are updated, the VMs are returned to their original configuration across the cluster.

Click here for a video demo.

SELF HEALING

The HC3 architecture is built with layers of redundancy, such as dual active/passive network ports, redundant power supplies, and redundant block storage striped across all cluster nodes. Intelligent state machine automation handles drive failures and node failures, redistributing data across remaining drives and VMs across remaining nodes and automatically absorbing replacement drives and replacement nodes into the resource pools.

CLUSTER TO CLUSTER REPLICATION

VMs can be replicated between two HC3 clusters with native, built-in replication. Replication can be local or remote across any distance and can be configured to replicate changes as often as every 5 minutes. Granular selection of VMs for replication allows a range of between one and all of the VMs on a cluster to be replicated. Replica VMs on the target cluster can be booted up within minutes in the event of a failure of the VM on the primary cluster.

Click here for a video demo.

SC/\LE
COMPUTING

SNAPSHOT SCHEDULING

VM snapshots can be scheduled intuitively and flexibly from the HC3 Web Interface. Schedules can include multiple rules spanning schedule intervals based on minutes to months. Schedules are created as templates that can be assigned to one or more VMs across a cluster.

Click here for a demo video.

VM BACKUP AND QUICK RECOVERY

Full agentless snapshot-based VM backup can protect your VMs to a local or remote site or to our ScaleCare Remote Recovery hosted datacenter. Not only can VM backups be recovered instantly at the DR site, but the data can be quickly restored to the original VM by restoring only differential data.

Click here for more Disaster Recovery Strategies

FILE-LEVEL AND DISK-LEVEL RECOVERY

With one click, HC3 administrators can clone a virtual disk from a snapshot and mount it to a live VM. A cloned and mounted virtual disk makes it easy to access data from the desired point in time to recover files or other data that was deleted, corrupted, or otherwise lost. These cloned, mounted virtual disks offer full read/write access and can be used for file-level recovery or full disk recovery.

Click here for a demo video.

THIN VM CLONING

HC3 uses a unique thin cloning technique that allows cloned VMs to share the same data blocks as their parent VM for storage optimization, but with no dependencies. If the parent is deleted, the clone is not affected and continues operating without disruption.

Click here for a video demo.

RAPID, EASY DEPLOYMENT

An HC3 cluster can be racked, cabled, powered on, configured in a matter of minutes, and VMs can be deployed and running in under an hour. Manage from your preferred web browser to the HC3 web interface. There is no extra software to install, simply upload ISOs to deploy VMs with your favorite guest OS.

ROLE-BASED ACCESS CONTROL

Administrators can manage users by assigning varying customized levels of access starting with full administrator access or controlled down to specific functions or read-only access. Role-based access control combined with user logging gives administrators powerful control over user activity.

NON-DISRUPTIVE SCALE-OUT

When a cluster needs more resources, a new node can be added within minutes without any downtime to the existing nodes or VMs. After racking, cabling, powering up, and assigning an IP address, the new node and its resources are seamlessly absorbed into the cluster, including the storage capacity being immediately added to the storage pool. Adding nodes results in an immediate performance increase to the running VMs as the resource pool is expanded. Click here for a demo.

MIX AND MATCH NODES

With HC3, any two cluster nodes across all HC3 families (1000, 2000, and 4000) can coexist in the same cluster and contribute to the cluster storage pool. This allows clusters to grow not just linearly but in whatever way is needed or desired. This support for different size cluster nodes can facilitate non-disruptive upgrades to the infrastructure by adding newer, bigger nodes and then retiring or repurposing older small nodes if desired, all while keeping the VMs running.

› **Learn more about HC3 - Request a Demo**
www.scalecomputing.com

Simplicity Engineered // scalecomputing.com // +1 317-856-9959

SCALE COMPUTING