



# SC//PLATFORM

Theory of Operations

HyperCore Version 9.1.17+



## About this Document

This document is intended to describe the technology, concepts, and operating theory behind SC//Platform, including the Scale Computing HyperCore™ OS, the Scale Computing Reliable Independent Block Engine (SCRIBE) storage layer, HyperCore Enhanced Automated Tiering (HEAT), Autonomous Infrastructure Management Engine (AIME), and the optional Scale Computing Fleet Manager with zero-touch provisioning (ZTP) and Secure Link.

### Technical Support and Resources

There are many technical support resources available for use. Access this document, and many others, at [www.scalecomputing.com/support](http://www.scalecomputing.com/support).

#### Online Support

You can submit support cases and view account information online through the Portal at [www.scalecomputing.com/support](http://www.scalecomputing.com/support). You can also Live Chat with support through [www.scalecomputing.com](http://www.scalecomputing.com) during standard hours Mon-Fri 8 AM to 6 PM ET.

#### Telephone Support Support

Telephone Support is available for critical issues 24/7 by phone at 1-877-722-5359 in the US and at 0808 234 0699 in Europe. Telephone support is recommended for the fastest response on priority issues.

#### Professional Resources

Scale Computing offers many professional service options for remote and on-site assistance in getting your cluster up and running quickly and knowledgeably. Contact your Scale Computing sales representative today to discuss our service offerings. Find additional information at [www.scalecomputing.com/professional-services-and-programs](http://www.scalecomputing.com/professional-services-and-programs).

#### Document Revision History:

Version 5.0: 04/2023 Addition of AIME, ZTP, and Secure Link

Version 4.0: 11/2022 Addition of SC//Fleet Manager

Version 3.2: 05/2022 Content review and design overhaul for product rebrand

Version 3.1: 05/2021 Minor edits for Versions 8.5+

Version 3.0: 01/2020 Content review and design overhaul for corporate rebrand

## Table of Contents

About this Document.....	2	SC//Fleet Manager—Cloud-based Monitoring and Management.....	12
Design Goals.....	4	Architecture.....	12
SC//Platform Overview.....	4	Data Handling.....	12
SC//HyperCore—The Hypervisor Simplified.....	4	Security.....	13
Software Managed Compute.....	4	Zero-Touch Provisioning (ZTP).....	13
Create a Virtual Machine.....	5	Secure Link.....	13
VM Placement.....	6	SC//Platform in Action.....	14
VM Failover.....	6	Cluster Formation—Resource Aggregation.....	14
VM Level Snapshots.....	6	Networking—High Availability.....	14
VM Snapshot Promote and Revert.....	7	Migration—Move Existing Workloads to SC//HyperCore.....	14
VM Snapshot Based Thin Cloning.....	7	Management—Multi-site Management.....	15
VM Replication and Recovery.....	7	Management—Role-based Access Control.....	15
Individual File Recovery.....	7	Software—Real-Time System Monitoring.....	15
SCRIBE—Mirroring, Tiering, and Wide Striping.....	8	Software—Firmware Upgrades.....	16
Software Defined Storage.....	8	Software—Cluster-to-Cluster Replication.....	16
Mirroring, Tiering, and Wide Striping.....	8	Hardware—Disk Failure Scenario.....	16
Data Placement.....	9	Hardware—Node Failure Scenario.....	17
Storage Pool Verification and Orphan Block Detection.....	9	Hardware—Network Failure Scenario.....	17
HyperCore Enhanced Automated Tiering (HEAT).....	10	Hardware—Expanding the Cluster.....	17
Tiering.....	10	Resources.....	17
Autonomous Infrastructure Management Engine.....	11		

## Design Goals

SC//HyperCore, the Scale Computing Reliable Independent Block Engine (SCRIBE) storage layer, and SC//Fleet Manager were designed to provide highly available, scalable compute and storage services across one or many locations while maintaining operational simplicity through highly intelligent software automation and architecture simplification.

SC//HyperCore adds intelligent automation to the hypervisor and storage layers and was designed to take advantage of low-cost, easily replaceable, and upgradeable “commodity” hardware components, including the virtualization capabilities built into modern CPU architectures. By clustering these components together into a single, unified, and redundant system, the architecture creates a flexible and complete “data center-in-a-box.” SC//HyperCore operates as a redundant and elastic private “cloud” that scales seamlessly with the automatic incorporation of additional nodes and handles hardware failures gracefully with minimal effort or disruption. SC//HyperCore provides flexible configurations (single node systems, replicated pairs, storage-only mode for local tie-breaker, 3+ node configurations, etc.) and various hardware platforms supported. The efficiency of SC//HyperCore allows for running on very small form factor devices, where the efficient use of compute (~4GB of RAM on smaller platforms) minimizes the total cost of ownership. Enterprise-class server hardware running SC//HyperCore also benefits from the efficiency of the stack in the high-performance capability it provides.

SC//Fleet Manager makes it easy to gain the benefits of SC//HyperCore regardless of the scale or geographic distribution of your infrastructure needs. Automation interfaces such as the Red Hat Certified Ansible Collection and the full REST API of SC//HyperCore augment SC//Fleet Manager for full end-to-end configuration and lifecycle management of any size fleet of SC//HyperCore clusters.

## SC//Platform Overview

Run your edge computing applications on a fully integrated platform for the easiest manageability and highest availability. Regardless of your hardware requirements, the same innovative software and simple user interface give you the power to run infrastructure efficiently at the edge. Scale Computing Platform combines everything you need: virtualization, servers, storage, and backup/disaster recovery with powerful fleet management to deliver a single manageable solution at scale for distributed edge locations.

## SC//HyperCore—The Hypervisor Simplified

SC//HyperCore is based on a 64-bit, hardened and proven OS kernel and leverages a mixture of patented proprietary and adapted open-source components for a truly hyperconverged product. All components—storage, virtualization, software, and hardware—interface directly through the SC//HyperCore hypervisor and SCRIBE storage layers to create an ideal computing platform deployed anywhere from the data center to the edge of the network, with optional monitoring and management assistance from Scale Computing Fleet Manager.

SC//HyperCore bundles a variety of adapted open-source and proprietary, intelligent software to create a simplified operating system. Custom-built utilizing the KVM architecture to integrate with the SCRIBE storage layer directly, SC//HyperCore makes virtualization and software automation look easy—and it is.

### Software Managed Compute

The SC//HyperCore software layer is a lightweight, type 1 (bare metal) hypervisor that directly integrates into the OS kernel and leverages the virtualization offload capabilities provided by modern CPU architectures. Specifically, SC//HyperCore is based on components of the KVM hypervisor, which has been part of the Linux mainline kernel for many years and has been extensively field-proven in large-scale environments.



SC//HyperCore integrates the [SCRIBE](#) storage pool directly into the KVM hypervisor. This means that virtual machines (VMs) running on SC//HyperCore have direct block-level access to the SCRIBE “virtual storage device” (VSD) virtual disks in the clustered storage pool without the complexity or performance overhead introduced by using remote storage protocols and accessing remote storage over a network (although the backplane network is used to communicate data locally between the cluster nodes).

Unlike other seemingly “converged” architectures in the market, the SCRIBE storage layer does not run inside a virtual machine as a virtual storage appliance (VSA) or controller VM. Instead, it runs parallel to the hypervisor, allowing direct data flows to benefit from zero-copy shared memory performance.

SCRIBE is not a re-purposed file system with the overhead introduced by local file/file system abstractions such as virtual hard disk files that attempt to act like a block storage device. This means performance-killing issues such as disk partition alignment<sup>1</sup> with external RAID arrays become obsolete. Arcane concepts like storing VM snapshots as delta files that later have to be merged through I/O killing brute-force reads and re-writes are also a thing of the past with the SCRIBE design.

## Create a Virtual Machine

Creating a virtual machine is a simple process accomplished through the SC//HyperCore interface in a single dialogue box. The interface is also used to upload Operating System and Virtual Appliance ISO installation images for VM installation as part of the VM creation process (ISOs are virtual DVD / CD images). ISO images can be uploaded by drag-and-drop anywhere on the HyperCore UI, or you can manage and upload ISOs through the media tab in the SC//HyperCore Control Center console. Once uploaded, ISO images are available for use by any future VMs.

Selecting the Create VM option (the + icon in the UI) right from the main VM management screen allows the user to specify required and optional parameters for the virtual machine, including:

- VM name and optional description
- Optional tags that allow logical VM grouping (the VM name, description and tags are searchable and filterable in the main VM management screen)
- Number of virtual CPU cores
- VM RAM
- Number and size of virtual, thin provisioned disks
- A previously uploaded virtual DVD/CD ISO image for installing an operating system

Creating a virtual machine creates not only persistent VM configuration parameters but also creates virtual disks using the SCRIBE distributed storage pool that attaches to the VM when started. SC//HyperCore VMs can access their virtual disks directly as if they are local disks without using any SAN or NAS protocols, regardless of which node the VM is running on at the time.

---

<sup>1</sup> Virtual disk block alignment problems can happen when you try to represent a virtual block device as a file sitting on a file system with its own block size which itself will sit on top of some other physical or logical block storage that may have its own block size and physical data layout.





SC//HyperCore virtual disks are thin provisioned so that virtual disk space is not fully allocated until the virtual machine actually uses it. On systems with tiered storage resources such as flash storage (SSD or NVMe) and spinning disks, SC//HyperCore assigns each virtual disk a default “flash priority” setting designed to provide the ideal balance of performance versus storage cost for every VM. Flash priority and tiered clusters are discussed further in the [SC//HyperCore SCRIBE—Mirroring, Tiering and Wide Striping](#) section.

SC//HyperCore automatically creates a single virtual NIC for each VM, but the administrator can easily add additional virtual NICs or specify VLAN tags to control the VM's access to network resources.

Of course, all of this can be accomplished programmatically through the REST API or via the Red Hat Certified Ansible Collection which allows users to interface with the Infrastructure as Code. Both the [REST API](#) and [Ansible Collection](#) are fully documented complete with example code on how to implement common functionality.

## VM Placement

Because all SC//HyperCore capable nodes have access to the entire pool of storage, virtual machine placement on nodes is determined by the availability of compute resources (RAM and CPU).

For example, in multi-node clusters, if a new VM requires 16 GB RAM, SC//HyperCore selects a node with at least 16 GB RAM available automatically, SC//HyperCore will try to determine the best available node based on the most available capacity for new VM placement. SC//HyperCore allows running VMs to be live migrated to other SC//HyperCore nodes without the VM being shut down and with virtually no noticeable impact.

## VM Failover

If an SC//HyperCore node ran VMs before a node failure in a multi-node cluster, those VM disks and data are still available to the remaining cluster nodes since the cluster uses a single pool of redundant storage. This allows VMs to be automatically restarted within minutes of the node failure on the remaining available nodes, with SC//HyperCore automatically placing VMs on nodes with the correct availability of compute resources or optionally based on pre-defined VM node affinity rules where desired.

## VM-Level Snapshots

SC//HyperCore can take multiple point-in-time consistent snapshots of virtual machine storage devices (virtual disks) and the VM configuration—a complete image of the VM at that time for backup purposes. Snapshots occur nearly instantly with virtually no impact on production workloads, even with thousands of snapshots.

SC//HyperCore snapshots use a space-efficient allocate-on-write methodology where no additional storage is used at the time the snapshot is taken, but as blocks are changed, the original content blocks are preserved, and new content written to freshly allocated space on the cluster.

Over time, the space consumed by a particular snapshot represents only the blocks that are unique to that snapshot point-in-time – which given multiple possible snapshots would be a subset of the data that has changed since the snapshot was taken. Blocks that are the same in multiple point-in-time snapshots or the current live image are reference counted so that they will be retained as long as they are referenced by the current image or any point-in-time snapshot to preserve cluster capacity.

This allocate-on-write methodology does not introduce the additional I/O associated with older copy-on-write snapshot methods, which for every change, must first read the original content, write it to a new temporary location, update the snapshot metadata to reflect that new location, then allow the new data to be written and metadata updated.

## VM Snapshot Promote and Revert

SC//HyperCore VM level snapshot can be promoted to a live VM and started nearly instantaneously from the HyperCore UI by leveraging the VM cloning capability. If the goal is to revert the running VM to a previous point in time, for example, to undo a problem update or patch, you can simultaneously clone the previous snapshot to a live VM and start that cloned VM while you power down and optionally delete the original problem VM. Or, retain the problem VM image and start it isolated from the production network to do a postmortem analysis.

When promoting or reverting a VM, the entire VM configuration and its snapshot storage objects are made available, allowing you to edit the VM configuration as you choose and start it up immediately with no data recovery time required.

## VM Snapshot-Based Thin Cloning

It may be obvious from the examples above, but VM Snapshots allow for instant space-efficient cloning or creating “template” VMs. In a single step, you can clone a live, running VM by taking a snapshot and immediately promoting the snapshot to a live image. It is then possible to start the cloned VM in an isolated network and do a trial run of a new software patch, for example.

Golden “master” template VMs can be created with the latest OS, application patches, security and antivirus tools, and customized settings. These template VMs can be used to quickly deploy large numbers of similar server or desktop OS VMs. Combined with the support for Cloud-Init and provisioning VMs using Ansible and the SC//HyperCore REST API, this makes for a powerful IT automation solution.

Each VM created from a snapshot clone originally takes zero additional space and grows based only on the changes made once that clone VM is booted.

## VM Replication and Recovery

In addition to being highly available within an SC//HyperCore cluster, VMs can be replicated between clusters for remote availability. This integrated replication protects against additional disaster scenarios, such as site and regional disasters, which may impact an entire SC//HyperCore cluster, affecting local recovery options.

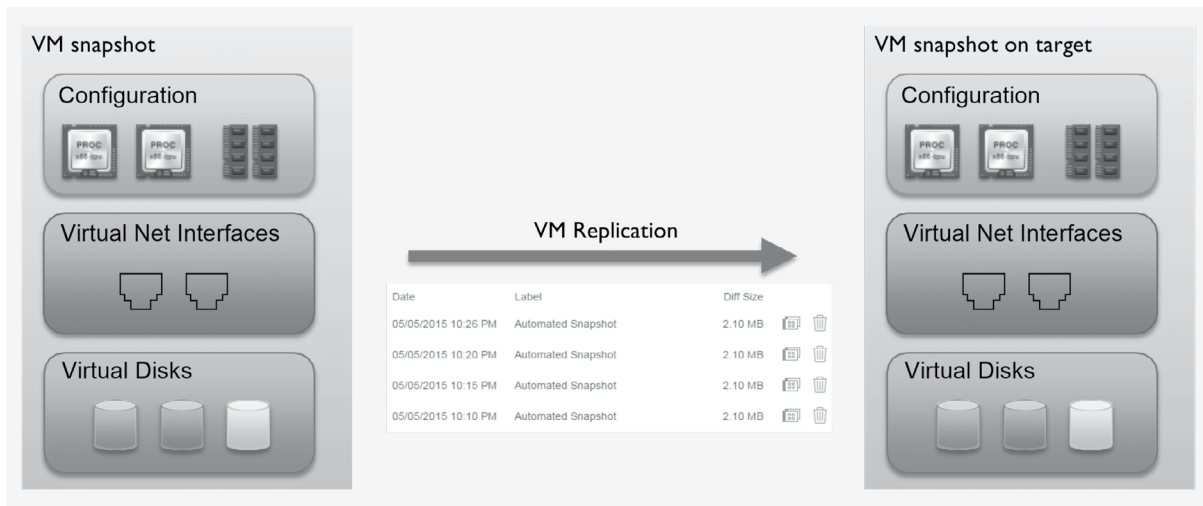
SC//HyperCore replicated VMs and the snapshots used for replication can be manually activated as live VMs with the click of a button. The recovery process from a remote replicated cluster is the same as recovering from a snapshot on the original cluster: promote a clone VM from the VM replication card image on the remote cluster. For larger environments, Ansible and the SC//HyperCore REST API can be used in automating DR testing and actual failover and failback steps.

The entire replication process can later be reversed to replicate and return one or more VMs back to the original cluster to resume operations. As with the original replication, SC//HyperCore can determine the block differences between multiple points in time, which can minimize the amount of data that needs to be sent back to the original cluster.

## Individual File Recovery

To recover individual files that may have been deleted or corrupted accidentally, Virtual Disk Cloning allows individual virtual disks to be cloned and mounted from snapshots so files can be recovered from a previous point in time. A virtual disk clone can be mounted to the same VM it is being cloned from or any other VM for recovery.

The virtual disk is cloned and mounted as a single action, making accessing the files on that disk from a live virtual machine easy. The virtual disk can be cloned from any available snapshot, ensuring files can be recovered as far back as snapshots are retained.



## Scale Computing Reliable Independent Block Engine (SCRIBE)

### Software-Defined Storage

A critical software component of SC//HyperCore is the Scale Computing Reliable Independent Block Engine (SCRIBE). SCRIBE is an enterprise-class, clustered, block storage layer that is purpose-built to be consumed by the KVM-based SC//HyperCore hypervisor directly. SCRIBE discovers all block storage devices—including flash-based solid state disks (SSDs), NVMe storage devices, and conventional spinning disks (SATA or SAS)—and aggregates the block storage devices across all nodes of SC//HyperCore into a single managed pool of storage. All data written to this pool is immediately available for read or write access by any and every node in the storage cluster, allowing for sophisticated data redundancy, load balancing intelligence, and I/O tiered prioritization.

### Mirroring, Tiering, and Wide Striping

SCRIBE is the storage management layer embedded in SC//HyperCore and treats all storage in the cluster as a single logical pool for management and scalability purposes. The real benefits of SCRIBE come from the intelligent distribution of blocks redundantly across the cluster to maximize availability and performance for the SC//HyperCore virtual machines.

Unlike many systems where every node in a cluster requires local flash storage—with a significant dedication of flash for metadata, caching, write buffering, or journaling—SC//HyperCore was designed to operate with or without flash. SC//HyperCore allows mixing of node configurations, which may or may not include flash in addition to spinning disks. Regardless of the configuration, SC//HyperCore uses wide-striping to distribute I/O load and access capacity across the entire SC//HyperCore cluster, achieving levels of performance well beyond other solutions on the market with comparable storage resources and cost.

To best utilize the flash storage tier, SCRIBE assesses the user-configured and natural day-to-day workload priority of data blocks (discussed in more detail in the [Tiering](#) section below) based on recent data history at the VM virtual disk level. Intelligent I/O heat mapping identifies heavily used data blocks for priority on flash storage. SC//HyperCore virtual disks can be configured individually at varying levels of flash prioritization so that an always-changing SQL transaction log disk versus a more static disk in the same VM, for example, will allocate flash appropriately as needed for improved performance.





## Data Placement

The SCRIBE storage layer allocates storage for block I/O requests in chunks ranging from as large as 1 MB to as small as a 512-byte disk sector. SCRIBE ensures that two or more copies of every chunk are written to the storage pool in a manner that not only creates the required level of redundancy (equivalent to a RAID 10 approach) but also aggregates the I/O and throughput capabilities of all the individual disks in the cluster, commonly known as wide striping.

In multi-drive SC//HyperCore clusters, all data is mirrored on the cluster for redundancy. Data writes between the node where the VM issues the I/O request and the nodes where the blocks will be stored occurs over the private backplane network connection to isolate that traffic from incoming user/server traffic.

For example, 10 GB of data is being written to a cluster of three nodes containing three solid-state disks and nine spinning disks. Each node contains four total disks of various available capacities. In this example, each node has one 960 GB SSD and three 4 TB spinning disks, thus contributing 12.96 TB of RAW storage capacity per node or 6.48 TB of effective usable storage space to the cluster after accounting for two copies of all data blocks being stored. That 10 GB of data (or 10240 MB) would require 10240 chunks (1MB each) to store the first copy plus 10240 chunks to store a second copy of each chunk. Not factoring in other I/O that may be occurring simultaneously, you could picture each of the 12 disks doing approximately 1700 x 1MB writes to store that 10 GB redundantly across the SC//HyperCore cluster.

Of course, other I/O is likely going on in the cluster at the same time, further leveraging the I/O capabilities and throughput of all the disks in the cluster. This means that the more nodes you add to the cluster, the more aggregate performance you can achieve. The benefits of wide striping across numerous disks apply to both read and write I/O.

Contrast this methodology to a conventional RAID array controller where you pick a few disks, allocate them to a specific RAID set, and are limited to the performance of the disks in that specific RAID set. If disks in another RAID set happen to be idle, those disks could be sitting there doing nothing while disks in a heavily utilized RAID set have become a bottleneck for that application (note that there are many other common limitations of conventional RAID such as the overhead of parity calculations, rebuild times, dedicated hot spares, etc., that are beyond the scope of this document).

Review the [SC//HyperCore Cluster in Action](#) section below for examples of how the SC//HyperCore system utilizes the SC//HyperCore and SCRIBE features in real-world scenarios.

## Storage Pool Verification and Orphan Block Detection

For temporary failures such as a node power outage, the SC//HyperCore cluster will automatically verify that any data blocks that the node contains were not changed elsewhere on the cluster while the node was offline. If the blocks were changed, SCRIBE would clean up the original blocks in the shared storage pool if they are no longer required (possibly blocks that were held by a snapshot).

As an example of orphan block detection, block 200 is contained on nodes 1 and 2. Node 1 goes offline due to a power outage and block 200 changes when node 1 is offline. Block 200 is then updated on node 2 and a new redundant copy is written elsewhere on the cluster (as node 1 is still offline and all new or changed data always requires a redundant copy to be stored). When node 1 comes back online, SC//HyperCore will detect that the original contents of block 200 are no longer current or needed and correct that discrepancy by releasing them back into the storage pool.

## HyperCore Enhanced Automated Tiering (HEAT)

### Tiering

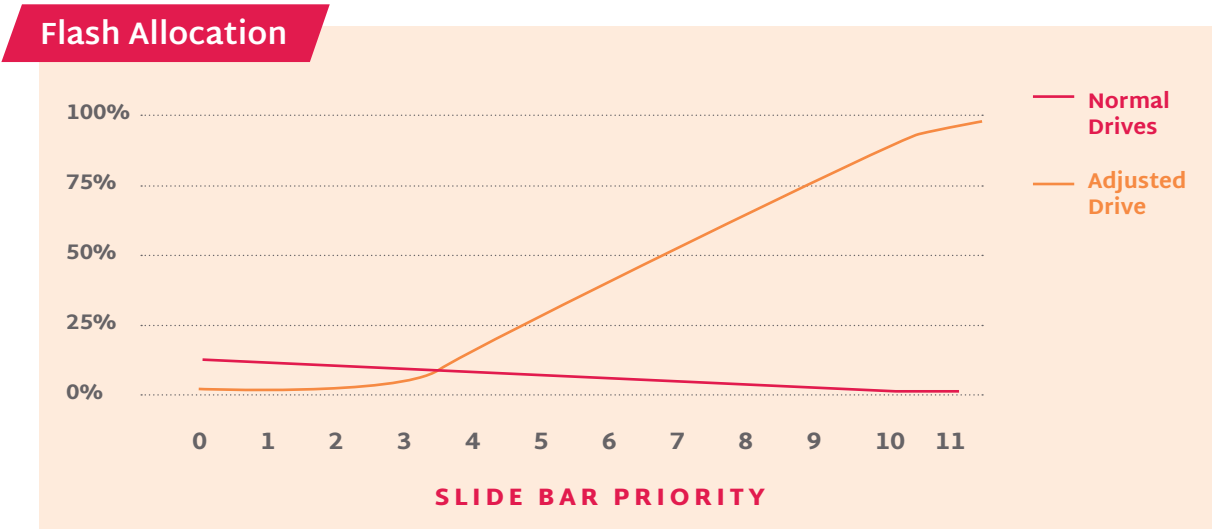
As previously mentioned, the addition of flash storage in SC//HyperCore nodes—known as Tiered SC//HyperCore Nodes—allows more capabilities in the SCRIBE storage architecture. Tiered SC//HyperCore Nodes in conjunction with the version 7.0 and later HyperCore firmware release allow the use of the HyperCore Enhanced Automated Tiering (HEAT) feature.

HEAT includes configurable flash priority allocation at the individual virtual disk level through an easy-to-use slide bar in the HyperCore UI, and intelligent data block priority based on block I/O heat mapping assessed utilizing historical I/O information on each virtual disk.

HEAT-capable clusters will automatically assign all new writes to the flash tier until SCRIBE is able to accurately assess their activity. In theory, an SC//HyperCore cluster with low storage utilization could run completely on flash storage.

By default, a new virtual disk will automatically get assigned a flash priority of 4 on a scale of 0 to 11. The scale, represented by the slide bar in the HyperCore interface, is exponential, meaning that changing the priority of a VM virtual disk from 4 to 5 in the HyperCore UI doubles the priority of that data for flash placement. This scaling means that even one position change on the slide bar can significantly improve performance and requires less experimenting with various settings. The values of 0 and 11 on the slide bar are unique. A setting of 0 eliminates flash storage usage on that virtual disk (convenient for a static FTP drive, for example) and 11 changes the flash priority by order of magnitude, multiplying the priority of the virtual disk data for flash storage placement by 10. When you turn it to 11, you really crank it.

The graph below shows an example of a VM virtual disk being prioritized from 0 to 11, assuming all other VM virtual disks remain at the default setting of 4.





## Autonomous Infrastructure Management Engine (AIME)

Autonomous Infrastructure Management Engine (AIME) is the orchestration and management engine that powers SC//HyperCore. As artificial intelligence for IT operations (AIOps) functionality, AIME drastically reduces the amount of effort required to deploy, secure, manage, and maintain on-premises infrastructure. Think of it as an autonomous system administrator included with every Scale Computing node.

AIME handles day-to-day administrative and maintenance tasks automatically, monitors the system for security, hardware, and software errors, and remediates those errors where possible. It identifies the root cause and minimizes the impact of those issues when it cannot repair them automatically, notifying users with specific problem determination and action, versus just sending a stream of data that must be interpreted. This includes actions to secure the environment. It also maintains current firmware, driver, and OS versions for security and stability purposes.

**State machines.** The core intelligence in the system is a hierarchical finite state machine. AIME's state machine constantly monitors and models the reality that the cluster is currently operating in, and then using that model, AIME can trigger appropriate actions. AIME's model of reality includes the hardware and software the cluster is running, the physical environment (temperature, power, cabling, etc.), and the logical environment (networking and external services). Given its extensive understanding of the environment, AIME can take many actions to maximize performance, security, and uptime for all your workloads.

The state machine model itself encodes the knowledge of over a decade of experience running tens of thousands of clusters in the field and improving that model release after release. It can model an astronomical number of possible states and take appropriate actions from any of those states.

**Conditions.** Below the sophisticated intelligence of the state machines, there is another simpler but essential layer of intelligence, conditions. Conditions are boolean flags that indicate problem areas. Conditions point directly to problems, making it simple for the state machine to understand the problem and enact fixes or an administrator when the solution requires human intervention. A concise list of problems leads to understanding the root cause quickly, and conditions often indicate what actions are required to return the cluster to a healthy state.

**Checked values.** Underpinning AIME is the data layer. As the old adage goes, "Garbage in, garbage out." If we are going to trust AIME to take action to repair problems on the cluster, we have to know that it's going to make good decisions. This can only happen if it can trust the data it uses to make decisions. To ensure the data can be trusted, we use carefully controlled variables we call checked values. Checked values are self-monitoring and recognize when they contain bad or old data. Reading bad data is disallowed, and the conditions and state machines that rely on that data know the data can't be used. This prevents taking actions on the cluster that could be harmful due to bad or stale data.



## Security

When sold as a pre-imaged appliance, SC//HyperCore is imaged at an ISO 9001:2015 and ISO 13485:2016 certified manufacturing facility and drop-shipped to a deployment site. Node detail is harvested during the integration process, including unique identifiers that associate the physical hardware/node with the organization within SC//Fleet Manager. Once the node is plugged in with power and networking, the node authenticates Scale Computing's Fleet Manager via standard HTTPS Public Key Infrastructure. Once the secure connection is established, SC//Fleet Manager delivers a set of keys to the node that allow both authentication and encryption between SC//Fleet Manager and the node going forward. This prevents man-in-the-middle attacks because at no point does the node connect to an unauthenticated service.

SC//Fleet Manager does not store passwords by design. If a user logs in with email and password, a Gcloud Authentication service securely stores the password. Only a token negotiated between the user and that Gcloud service passes to SC//Fleet Manager. If a user logs in with one of the supported SSO providers, even the Gcloud authentication service will not see or store passwords for users. SC//Fleet Manager is invite-only and restricts each user's access by tenant and role.

Scale Computing's in-house IT team controls access to all development and production systems. Access is only granted to users responsible for and knowledgeable of the systems. Access within the systems is scoped as narrowly as possible. Access to all production systems requires a work account login leveraging MFA. All code access also requires a secured VPN and additional login.

## Zero-Touch Provisioning (ZTP)

The zero-touch provisioning (ZTP) feature of SC//Fleet Manager provides cloud-like ease of infrastructure configuration for administrators, allowing them to centrally configure clusters with a few clicks and entry of the bare minimum details: cluster name, LAN and backplane IPs, and LAN gateway. All that is needed on-site is racking, cabling, and powering on, slashing configuration time and potentially removing the need for IT staff to be physically present at installation.

Nodes ordered from Scale Computing ship imaged with SC//HyperCore, which includes all necessary unique identifiers and authentication key requirements out of the box, negating the need for additional node registration or manual authentication steps required by other "zero-touch" solutions. The node simply connects, authenticates, and if SC//Fleet Manager holds a ZTP configuration for that node's unique identifiers, it is securely delivered and executed without any manual setup or action on the physical node, and the provisioned cluster appears in the owner's SC//Fleet Manager cluster list ready for workload provisioning and configuration.

## Secure Link

Secure Link provides IT admins with secure access to the HyperCore UI of any cluster in their fleet. When a user initiates a session, a secure outbound connection is established from the cluster to the Secure Link Server over HTTPS, which is presented in the browser to the authenticated user. All data transmitted between the browser and the cluster is encrypted using TLS 1.2.

During the Secure Link session, only the HTML of the cluster UI is transmitted to the user's browser. No lower-level data is sent or received, ensuring that only the user has access to the UI. Local SC//HyperCore authentication requirements are still in effect when using Secure Link, ensuring cluster-level authorization controls for all users. Once the browser containing the tunnel session is closed, Secure Link automatically closes the outbound connection, eliminating the access route entirely until the next use.

Secure Link eliminates the need for traditional remote access tools like VPNs or remote desktops, which can be costly and time-consuming to set up and maintain. Additionally, Secure Link ensures that there are no misconfigurations or security vulnerabilities associated with the manual setup of remote access tools, making it a more secure and convenient solution for IT admins.



## The SC//Platform in Action

The best way to highlight the benefits of SC//Platform is through real-world scenarios at various levels: cluster formation (installation), networking, migration, software, and hardware. The following sections will describe how SC//HyperCore intelligence simplifies virtualization, how SCRIBE capabilities unify storage management, and how SC//Fleet Manager further simplifies monitoring and managing day-to-day operations for any number of clusters.

### Cluster Formation—Resource Aggregation

Once the SC//HyperCore nodes are racked, cabled, and configured with physical network connectivity, the initialization process takes multiple nodes and logically bonds them together to act as a single coordinated cluster. Initialization can occur via command line, or via [Zero Touch Provisioning \(ZTP\)](#) through SC//Fleet Manager. Once the cluster is formed, it is managed as a single system, and all resources attached to each node are aggregated into a single managed pool of storage and compute resources. This is real hyperconvergence.

SC//HyperCore management architecture eliminates the need for a separate management server (or VM appliance) to install or manage and no single “brain” or “master node” (known as a single point of failure) that controls the overall system. The entire SC//HyperCore cluster is managed locally by pointing a web browser to the public IP address of any node in the cluster for a complete cluster management view utilizing a simple and intuitive web interface. SC//Fleet Manager enables a multi-cluster perspective in the cloud with no additional cluster setup but is not required for the full functionality of any SC//HyperCore cluster.

### Networking—High Availability

While each node in the cluster has its own physical network identity and interfaces, the ability to move and place virtual machines on specific nodes lets SC//HyperCore intelligently balance compute and I/O load among nodes and network interfaces in the cluster, maximizing overall system performance.

SC//HyperCore nodes have two distinct physical networks in which they participate—a public LAN network and a private backplane network. SC//HyperCore nodes offer redundant network ports for both the LAN and backplane network connections to allow for full network redundancy. The LAN network provides a path to allow network access to the management interface, virtual machines and end users accessing data from VMs running on the cluster. The private backplane connection is for intra-cluster communication, such as the creation of redundant copies of data being written on one node or disk to a mirrored copy on another node.

The backplane connection is for intra-cluster communication only and has additional security measures to the internal firewall already in use. The backplane bond is prevented from being assigned an IP in the same subnet range as those of the LAN IPs (to ensure the backplane IPs are isolated on the network as an additional security measure). In SC//HyperCore, for nodes with four network ports, two are reserved to physically separate the backplane network from the public LAN network. In nodes with one or two network ports, the backplane network utilizes a tagged VLAN for logical isolation of backplane traffic.

### Migration—Move Existing Workloads to SC//HyperCore

If the customer has existing workloads that they want to move to SC//HyperCore, there are a variety of Scale Computing or third-party tools available that can migrate existing physical server (Physical-to-virtual— P2V) and virtual machine (Virtual-to-virtual—V2V) images to SC//HyperCore VMs. Tools like Acronis Cyber Cloud even support “convert to VM” functionality, allowing you to back up a VMware or Hyper-V-based VM and convert it to a VM running on SC//HyperCore directly.

Scale Computing Move Powered by Carbonite is available from Scale Computing and offers near-zero downtime migration of Windows servers and applications onto SC//HyperCore. Several other tools have been outlined in Application Notes that can be found on the Scale Computing Customer Portal. In addition, Scale Computing and its partners offer full or quick-start migration services to make it easy for customers to begin using SC//HyperCore.



## Management – Multi-Site Management

SC//Fleet Manager enables monitoring and managing thousands of clusters from a single cloud-based interface. Other monitoring and management solutions need detailed setup and configuration of the nodes, the monitoring tool, or both before clusters can be synced. Some even require a local application with its own dedicated resources and maintenance to act as the relay or middle-man for monitoring, which dramatically increases costs and complexity.

SC//HyperCore nodes auto-join SC//Fleet Manager within minutes of cluster initialization. Once clusters have joined, you not only see clusters in real-time, but can also monitor your fleet from an entirely different perspective with cross-cluster searchable and sortable views of all nodes, VMs, and even individual conditions. With this flexibility and scale, monitoring a multi-thousand cluster deployment is not just possible, but simple.

## Management – Role-based Access Control

Administrator roles can be managed per cluster within SC//HyperCore by assigning specific sets of access controls to individual administrators. As a new administrator account is created, access to specific administrative functions can be granted or restricted to provide only the level of access needed. Access controls are logically grouped to include role-based functions such as VM creation, VM deletion, backup/disaster recovery, and more. SC//Fleet Manager roles are independent of SC//HyperCore cluster roles but also use role-based access to ensure only required permissions are granted to each user.

HyperCore Role-Based Access Control (RBAC) allows specific user accounts to be restricted to performing very specific tasks and actions. For example - certain accounts can be “read-only” allowing view/monitoring access to the HyperCore UI (and REST API) only - with no ability to create new VMs, or delete or modify existing VMs or cluster settings. Other accounts might be established strictly for backup purposes - for example in conjunction with third-party backup products such as Acronis that integrate with SC//HyperCore via API.

OpenID Connect (OIDC) integration extends HyperCore RBAC to utilize third-party identity providers (IDP) such as Microsoft Azure AD or ADFS with other integrations planned for identity providers leveraging the OIDC standard. SC//HyperCore roles can be managed locally or in a centralized directory based on group membership or other policies. Further, access policies including the use of multifactor authentication (MFA) and many others can be enforced by the OIDC identity provider.

## Software – Real-Time System Monitoring

Utilizing a management state machine at the hypervisor layer, all nodes in the SC//HyperCore cluster work in tandem for real-time health and state monitoring—essentially creating a self-healing cluster. All nodes, hardware components, software components, and guest VMs are monitored in real-time. The cluster will automatically take action on any identified conditions and notify configured technical contacts through email or syslog of any issues, again in real-time. Contacts can then take action or contact Scale Computing Support if needed—often, the SC//HyperCore system is intelligent enough to resolve issues without manual intervention ever being required.

Additionally, SC//Fleet Manager and the HyperCore UI will display active cluster conditions in real-time. Conditions are persistent across all nodes in a cluster, so it is not required to access the identified node LAN IP in order to view or take action on an issue.

## Software—Firmware Upgrades

SC//HyperCore offers automatic, non-disruptive firmware upgrades across the entire SC//HyperCore cluster. SC//Fleet Manager and SC//HyperCore detect and display available updates that can be initiated by the administrator from either interface. Minor upgrades that do not impact the SCRIBE storage layer are “in-place” upgrades with zero effect on running VMs. For larger rolling upgrades, SC//HyperCore live-migrates running VMs from a node, updates the node firmware, returns VMs (through live migration) to their original node location, and repeats this process for each node until the entire cluster is updated.

SC//HyperCore updates require sufficient compute RAM available to allow the VMs running on each node to be migrated across the remaining nodes in the cluster (essentially, the ability for all VMs to run with one node’s RAM resources unavailable in the cluster). This is also a requirement to allow for proper VM failover and SC//HyperCore continuously monitors used versus available RAM and generates an alert condition through the state machine if there is insufficient memory to allow VMs from any node to failover to remaining nodes in the cluster.

## Software—Cluster-to-Cluster Replication

SC//HyperCore cluster-to-cluster replication provides the ability to replicate VMs—at a per VM level—on one SC//HyperCore cluster another, often at a remote location. Cluster-to-cluster replication is designed to run continuously and to transmit changes to a secondary cluster as quickly as possible. SC//HyperCore VM replication leverages the SCRIBE snapshot capabilities plus additional compression functionality to efficiently determine the data changes between the most recent data snapshot and the snapshot that represents the last completed replication point (often the last completed point is just minutes old) and transmit the changes to the remote cluster.

Leveraging change tracking limits the impact of replication on the cluster by reducing the I/O required to read and transmit changes. This also eliminates the need to “brute force” read and compare data to determine which blocks have changed since the last replication cycle.

For VMs that were created from snapshots or via cloning that share common data blocks (such as multiple VMs created from a “template” and cloned), SC//HyperCore is intelligent enough to transmit those common blocks across the network to the remote cluster a single time which can greatly reduce the bandwidth required for the initial replication of a new VM created from an already replicated template.

## Hardware—Disk Failure Scenario

SC//HyperCore was designed with the expectation that hardware eventually fails, and such failures should not be critical to the system. Based on this design, disk failures are considered trivial and have little effect on the cluster beyond a temporary loss of the capacity and I/O resources of the affected disk. When a disk is determined to have failed, in addition to raising alerts in the management interface and through email, SC//HyperCore will continue servicing I/O transparently using the mirrored data copies. Additionally, SC//HyperCore will automatically use the remaining available space to re-generate a second copy of any chunk that was previously located on the failed disk. Any new writes are mirrored across the remaining disks to ensure immediate protection of new and changed data and reduce future data movement once the failed disk is replaced.

Unlike many RAID systems, disk failure and replacement do NOT require time-consuming and I/O-intensive parity calculations to determine what data was lost. SC//HyperCore simply reads the mirror blocks and writes a new copy using any available space on the cluster. Note that both read and write operations here leverage the same wide striping performance benefits as normal I/O.

Not only does SC//HyperCore not require a dedicated “hot spare” disk that sits idle most of the time, it doesn’t require a failed disk to be replaced or any other manual steps before it automatically re-creates the lost mirror copies to regain full data redundancy. This ensures there is no critical need to immediately replace a failed disk in the system for ease of system management.

## Hardware—Node Failure Scenario

SC//HyperCore was designed not only to be resilient against disk failures but to sustain the temporary loss of an entire node while continuing to process data and running VMs using redundant resources from the remaining cluster nodes. A single node failure (loss of access to all node disks—flash or spinning disks— and compute resources) does not impact data availability for VMs, unlike some systems on the market where even a single flash storage device or spinning disk failure may result in data loss due to caching or RAID arrays.

## Hardware—Network Failure Scenario

Nearly all SC//HyperCore nodes offer redundant network ports for both the public LAN and private backplane network connections for a total of 4 network ports per node to allow for full network redundancy. These ports are in an active/passive bond for failover—2 ports bonded for LAN access and 2 ports bonded for backplane access. If access to the primary port is lost on either the public or private node network, the secondary bonded port will activate with little to no disruption to the cluster node. Once the primary port becomes available again, the network will fail back to that port. All onboard network cards are used as the primary port for the SC//HyperCore nodes.

## Hardware—Expanding the Cluster

After the initial cluster installation and formation, SC//HyperCore allows additional nodes to be added to the cluster at any time without downtime to the system. Joining a new node to the cluster automatically adds the new node's capabilities to the overall available storage pool and compute capacity. New nodes can be of the same or different type or capacity as the existing nodes in the cluster. SC//HyperCore allows newer, faster nodes with different CPU, RAM and disk configurations to be added to existing SC//HyperCore clusters to increase capacity and ultimately replace and retire older legacy nodes as needed.

When adding nodes with storage resources to an existing cluster, existing storage may be optionally rebalanced to immediately utilize the new nodes and free up resources on older nodes. This process will relocate blocks from the older nodes onto the newer nodes to achieve an equal distribution of data across cluster nodes and disks. Even without using the option to rebalance data, SCRIBE will prioritize the node with the freest storage capacity for any new data written to the cluster, so the new node storage resources will immediately become an active part of the cluster. With or without rebalance, the new resources will begin providing more distributed I/O and improved system performance across the cluster.

Running SC//HyperCore VMs may be immediately live-migrated to the newly added nodes as needed to take advantage of the new compute resources with no downtime or disruption to VMs. New CPUs and RAM, in addition to storage resources, can be utilized by both VMs that are live migrated or new VMs that are created within the SC//HyperCore cluster.

## Resources

Additional informational and technical resources are available from Scale Computing at [www.scalecomputing.com/resources](http://www.scalecomputing.com/resources).

### Disclaimer

Any information listed here is supplement to the SC//HyperCore User Guide, product information, and Knowledge base. Scale Computing is not responsible for any issues or damages arising out of the use of this document.

Corporate Headquarters  
525 S. Meridian Street - 3E  
Indianapolis, IN 46225  
P. +1 317-856-9959  
[scalecomputing.com](http://scalecomputing.com)

EMEA B.V.  
Europalaan 28-D  
5232BC Den Bosch  
The Netherlands  
+1 877-722-5359

