




# Features Guide: SC//Platform

THE  
COMPETITIVE  
EDGE

## Scale Computing Platform Features

Scale Computing Platform brings together servers, storage, virtualisation, and disaster recovery into a single, feature-rich solution. This guide describes all the features that make SC//Platform the perfect infrastructure solution for any data center or edge deployment.



### SCALE COMPUTING FLEET MANAGER

SC//Fleet Manager is the first cloud-hosted monitoring and management tool built for hyperconverged edge computing infrastructure at scale. It's now easier than ever to securely monitor and manage your entire fleet of clusters running Scale Computing HyperCore. For IT Managers with multiple clusters, Scale Computing's industry-leading low maintenance time will be reduced by 50% or more.

Not just another monitoring tool, SC//Fleet Manager is engineered to leverage SC//HyperCore clusters' unique strengths to create a complete future-proof infrastructure solution.

Log in via any browser on any device to gain real-time monitoring and management of your global fleet of SC//HyperCore clusters. No installation, configuration, or training is required to get your entire Scale Computing infrastructure at your fingertips.

### SCALE COMPUTING HYPERCORE

The foundation of SC//Platform is the Scale Computing 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 (more information below) to automate tiering across our own hybrid flash storage architecture.

### HYPERCORE UI

Access the HyperCore UI from any node in a cluster. The simple interface design allows both storage and compute resource management and virtual machine management from a single 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.

### SCRIBE

Scale Computing Reliable Independent Block Engine (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 and node failure. Wide striping on every drive 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 [VSA]) architecture, which optimises the available resources to the workloads running in the environment.

## HEAT

HyperCore Enhanced Automated Tiering (HEAT) 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 prioritises blocks of data on SSD or HDD based on an I/O heat map that tracks I/O against each virtual disk. Although HEAT is primarily automation happening behind the scenes, it provides users with the ability to tune the relative priority of every virtual disk in the cluster to further optimise the SSD usage where it is needed the most. You can literally turn it to 11.

## RAPID, EASY DEPLOYMENT

A Scale Computing HyperCore cluster can be racked, cabled, powered on, configured in a matter of minutes, and virtual machines (VMs) can be deployed and running in under an hour. Manage from your preferred web browser to the HyperCore UI. There is no extra software to install, simply upload ISOs to deploy VMs with your favorite guest OS.

## REST-APIS COMPLETE WITH CLOUD-INIT

SC//HyperCore includes REST-APIs that enhance the speed and ease with which users can deploy virtual machines (VMs) at scale using cloud-init. Common VM templates may be uniquely configured via scripting at first boot. Avoid the need to manually create and individually customise VMs, but programmatically provide hundreds or thousands of machines with their own settings via script.

## NON-DISSCRIPTIVE SCALE-OUT

When a cluster needs more resources, a new node can be added within minutes without any downtime to the existing nodes or workloads. 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.

## FLEET-WIDE AND CLUSTER- LEVEL USER MANAGEMENT

SC//Platform features multi-user administration to allow each administrator a unique account and login, with role-based access to ensure control and security. Utilise SSO to enhance authentication and ease of management across your entire SC//Platform fleet.

## MIX AND MATCH NODES

With SC//Platform, appliances from different families can be added to and coexist in the same cluster and contribute to the cluster storage pool. Even appliances with dissimilar storage such as all HDD, hybrid HDD + SSD, or all SSDs can be combined as a cluster grows. This allows clusters to grow not just linearly, but also 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. Check with our ScaleCare support engineers for best practices for your particular cluster.

## DEDUPLICATION

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

## LIVE VM MIGRATION

VMs on SC//HyperCore clusters can be non-disruptively migrated between nodes with no downtime. This allows for rebalancing resource allocation across the cluster and allows VMs to be relocated automatically during our rolling update process for the SC//HyperCore OS firmware.

## THIN VM CLONING

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

## NON-DISRUPTIVE ROLLING UPDATES

SC//HyperCore receives software/ firmware updates directly to the HyperCore UI 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 system. The automated update process will automatically relocate VMs between nodes to free up the nodes to be updated, one at a time. When all of the cluster nodes are updated, the VMs are returned to their original configuration across the cluster.

## SELF-HEALING

SC//HyperCore architecture features layers of redundancy, such as dual active/passive network ports, redundant power supplies, and redundant block storage striped across all cluster nodes. Intelligent 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.

## 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 cluster nodes. Failover happens within minutes for minimal disruption. After the node has been restored or replaced, the VMs will automatically be live migrated back to their preferred nodes non-disruptively.

## SNAPSHOT SCHEDULING

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

## CLUSTER-TO-CLUSTER REPLICATION

VMs can be replicated between two SC//HyperCore 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 is required to meet each customer's unique RPO/ RTO requirements. Granular selection of VMs for replication allows a range from one to all the VMs on a cluster. 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.

## VM BACKUP AND QUICK RECOVERY

Full agentless snapshot-based VM backup can protect your VMs to a local or remote site, or to our SC//Platform Cloud Unity Disaster Recovery as a Service (DRaaS) offering. A VM backup can not only be recovered instantly at a DR site, but data can be recovered quickly on the original VM by restoring only differentials. VM snapshots and replication combine with quick file-level restore functionality to provide protection against and recovery from accidental file deletion, corruption, and even ransomware.

## REMOTE SUPPORT ACCESS

SC//Platform offers a remote access point exclusive to ScaleCare support to help diagnose support issues and take corrective actions if necessary. This is only initiated from the customer side for security and does not grant ongoing access to customer clusters. It merely dramatically speeds up support issue resolution by allowing ScaleCare support engineers to see issues firsthand and in real-time.

