SCALE COMPUTING

# Disaster Recovery Strategies

From the Experts at Scale Computing

Technical White Paper

**Understanding how to survive and respond to I.T. threats**

# Table of Contents

# Introduction

Disaster recovery is a concept that asks, "How can an organization survive and respond to various threats, ranging from small hiccups to catastrophic destruction?" The threats to ongoing operations range from human error to malicious attacks to natural disasters. We recognize that in today's 24/7 marketplace, IT infrastructure must be both resilient and highly available to keep organizations operational. Organizations need to prepare in ways that involve both human and technological responses.

At Scale Computing, we keep simplicity and ease of use in mind - built into our disaster recovery (DR) capabilities. These allow our users to recover quickly from disasters that may affect a single file to an entire site. Disaster recovery is often planned for and measured in recovery point objective (RPO) and recovery time objective (RTO). Scale Computing's HyperCore achieves RPO and RTO measured in minutes to minimize downtime and data loss.

Simplicity. Scalability. Availability. These bring significant value to DR, as the prevailing strategies and approaches add complexity and cost. This document will outline the strategies and built-in technologies to protect both data and workloads on SC//HyperCore to get services back online as quickly as possible following even the worst disasters.

# Scale Computing HyperCore Data Protection Suite

## Backup

The concept of backup has evolved over the decades to overlap with more modern snapshot and replication technologies. The days of taking traditional full and incremental backups should be over. With Scale Computing HyperCore, anyone can implement a simple yet robust disaster recovery strategy that combines per VM snapshot scheduling with replication, failover, and recovery. Each appliance has built-in VM snapshots with scheduling capabilities that are flexible enough to implement almost any backup strategy.

Like traditional incremental backups, snapshots only capture data that has changed since the last snapshot, making them highly efficient for storage and enabling flexible scheduling. What's more, SC//HyperCore's snapshots are immutable, meaning they can't be altered or deleted by the guest VM. If a VM's guest OS is compromised, it cannot modify snapshots.

Different workloads will have different requirements for disaster recovery. It is important to know what level of protection each workload in your organization needs. Often, workloads are divided between tiers of priority, recognizing some as more critical to operations than others. The backup strategy should reflect these multiple levels of need.

Examples:

| Snapshot Increments | Minutes | Hours | Days | Weeks | Months |
|---|---|---|---|---|---|
| Strategy 1 Critical | Every 5 mins for 12 hours | Every hour for 36 hours | Every day for 4 weeks | Every week for 6 months | Every month for 2 years |
| Strategy 2 Non-Critical | N/A | Every hour for 1 day | Every day for 2 weeks | Every week for 2 months | N/A |

The backup strategy will impact available disk space, so the number of snapshots per VM may need to be managed based on overall storage availability on an appliance. Snapshot size will vary depending on the data rate of change per VM.

Snapshots alone do not make a backup, even though they are beneficial for local data recovery from many operational disasters. In addition, their immutability means that they can be crucial in recovering from malicious ransomware attacks. Snapshots must be replicated onto another device for a true backup strategy, preferably at another site.

## Replication

SC//HyperCore's native replication capabilities will replicate snapshots on a per-VM basis to another appliance or cluster using efficient network compression and encryption. The replication process begins with an entire replica of the VM and its snapshots. That replica can be used on the target appliance for complete recovery or failover. Replication follows the snapshot schedule assigned to a VM and can replicate snapshots as often as every 5 minutes for a solid RPO.

Replication occurs over standard TCP/IP networks so that it can travel over any distance to any remote site. Replication is not synchronous, so it does not require expensive or restrictive high-speed links. Low bandwidth and latency can affect replication performance. It is important to understand the amount of data changing on a VM that may be replicated over the available bandwidth. WAN acceleration technologies may also help overcome bandwidth and latency challenges between remote sites.

Replication can be directed to another Scale Computing appliance locally or remotely, although remote replication is recommended to protect against site failure. Snapshots are maintained locally on the original appliance based on the snapshot schedule for retention, whether they are replicated or not. Local snapshots can often protect against individual file loss or corruption. Replication is what creates the full backup for protection against appliance or site failure.

## Failover

When a Scale Computing appliance experiences a critical failure event that cannot be recovered quickly, such as a fire, flood, or sustained power outage, replicated VMs can be failed over to the appliance where the replicas reside. The failover process is simple - by creating a clone from the replica, powering on the clone, and in most cases, between sites, redirecting DNS for the IP address and subnet at the remote site. SC//HyperCore also allows for bulk actions, which means groups of replicas can be cloned and powered on all at once. As such, this process can be completed quickly, in minutes after failure, providing a solid recovery time.

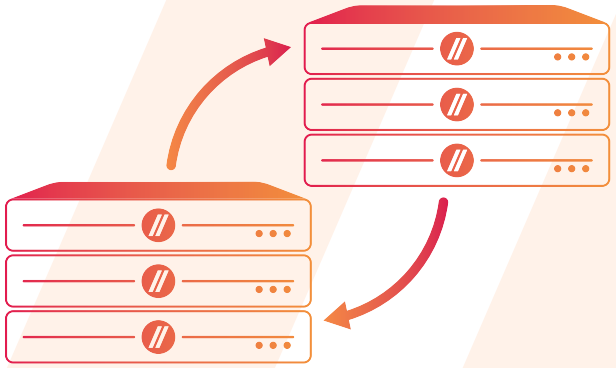

It is important to understand the interdependencies between workloads and create a plan or runbook for failing over. Some workloads may need to start before others to ensure applications can connect to required data and services for startup. Some processes, such as DNS redirection, may also be automated with scripting. The larger the number of workloads, the longer the overall process may take, but with proper planning, it should all be manageable in minutes rather than hours.

After failover, the failover VM will have the most current version of the production data. Although the failover VM could become the new permanent production VM, most likely, it will be relocated back to the original site once the appliance is recovered/replaced. Now suppose the original VM is intact but now outdated, as is often the case. Data restores quickly, with SC//HyperCore replication identifying the last good data point and only needing to restore the changed data. Then the failover process can be reversed, quickly bringing the workload online at the original site. If the original VM no longer exists, it will require more time and effort to restore all the data.

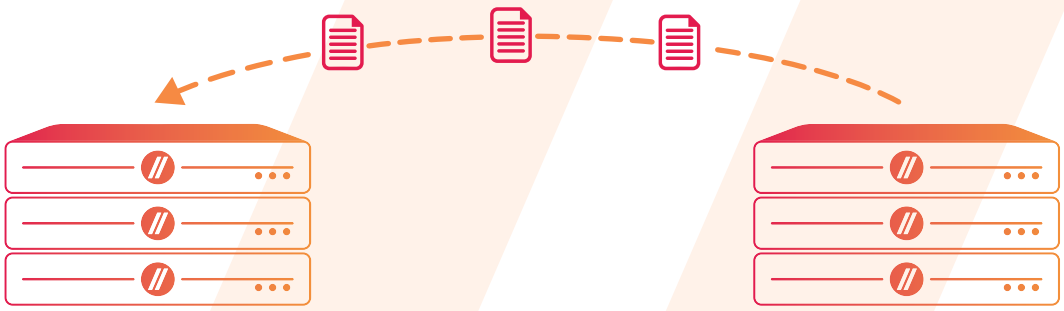

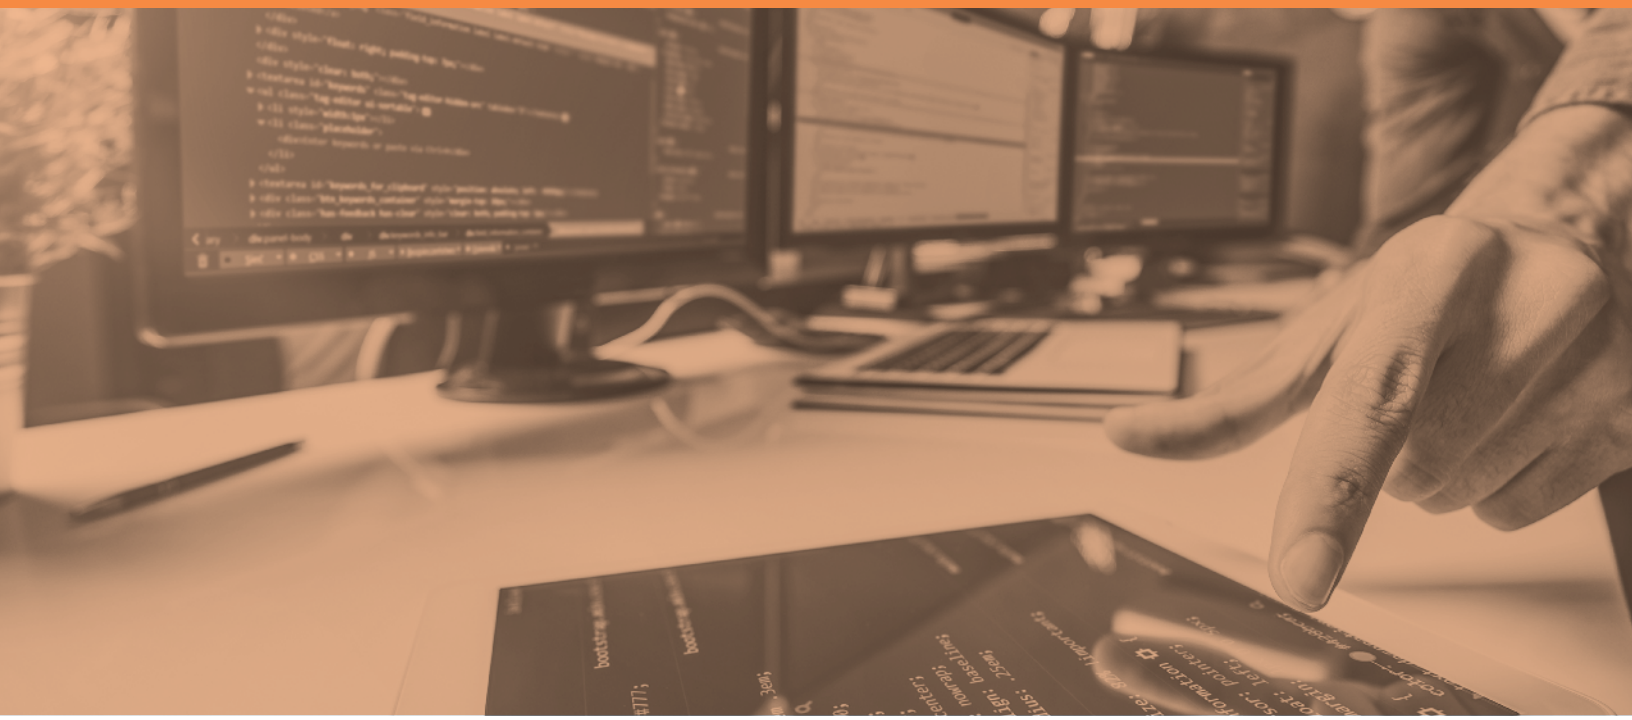

## Recovery

While failover is also considered the recovery of an entire VM, there are cases where it is neither desirable nor practical to fail over a VM. An individual file on a file server may need to be recovered rather than an entire VM. Maybe it is a non-essential VM that is not required to be running during a failover scenario but rather only recovered back to the original site as time permits. The data can be recovered either on an individual file basis or as an entire VM in these cases.

For file-level recovery, an IT administrator may quickly clone a virtual disk and mount it to a live VM to gain access to point-in-time data and recover what is needed. The admin simply chooses a point-in-time snapshot and, with one click, clones and mounts a virtual disk with the desired data to a live VM, and then the admin has full read/write access to all the files on the cloned virtual disk for recovery.

Individual file recovery may also be implemented on individual Windows VMs by enabling the Virtual Shadow Copy Service (VSS) within the Windows guest OS. VSS within the Windows VM can provide robust point-in-time file recovery available to individual users if desired. These snapshots will affect the storage requirements of the VM, but that is simply the tradeoff for quick, easy file recovery for users. The cloned virtual disk may also be used to recover the entire disk if desired.

For full VM recovery, it is just a matter of redirecting replication back to the original site as if you were restoring after a failover, but without ever having failed over. You do need to clone the replica, but this takes only seconds, and then you may begin restoring the VM to the source site. This is a perfect scenario for non-essential VMs not needed online during disaster recovery scrambles.

In the case of a ransomware attack, an administrator can quickly respond by cloning a previous snapshot, creating a new VM from that snapshot, and powering it on. Cloning only takes a few seconds, and SC//HyperCore's snapshots are immutable, meaning they cannot be deleted or altered by the VM guest OS if an attack has already compromised it. Cloning allows administrators to access files that have not been encrypted by a malicious actor and revert to those previous images.

These built-in recovery techniques, along with replication and failover, provide a complete disaster recovery solution that should meet the needs of many organizations. In the past, it was a common requirement to purchase and license a third-party solution for every IT infrastructure environment. With features built into SC//HyperCore's, most disaster recovery needs are met without reliance on additional solutions.

# ROBO and Small Environments

Providing disaster recovery with limited resources challenges distributed enterprises with multiple remote offices or branch offices (ROBO) and other small IT environments. Not all organizations look the same nor have the same disaster recovery requirements. It is important to call out some specific types of environments with unique needs.

## The Distributed Enterprise

There is typically a healthy IT infrastructure in the distributed enterprise at a central office and multiple remote or branch offices with minimal IT footprint. In this environment, the SC//HyperCore 3-node cluster configuration provides a simple, easy to manage virtualization platform without the high availability of a cluster but with replication and failover capabilities for disaster recovery. Many of the critical workloads in these environments are already protected by high availability and replication within the central office, so the remote site does not need the same level of availability. SC//HyperCore lets you run the most applications on the smallest hardware with the most reliability and least amount of effort.

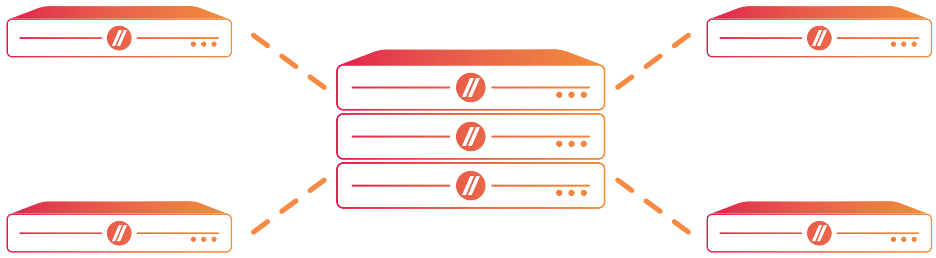

HyperCore UI makes monitoring the health of your entire infrastructure deployment easier than ever. Each single node appliance or cluster in remote offices can be managed remotely and replicated back to an SC//HyperCore cluster at the central office for backup, failover, and recovery. When one of these appliances fails, the workloads can failover to the central office. The failed-over workloads can then be accessed remotely until the node can be recovered or replaced.

It provides a right-sized configuration and price while providing built-in disaster recovery.

## Small Environments

Smaller environments benefit from a smaller SC//HyperCore cluster in production for local high availability. For these smaller organizations, it may not make sense financially or practically to deploy a whole second cluster for disaster recovery, even when they have a second site to host it. However, they can probably weather the storm by running their critical workloads on a single node appliance until the primary site is recovered in a site disaster.

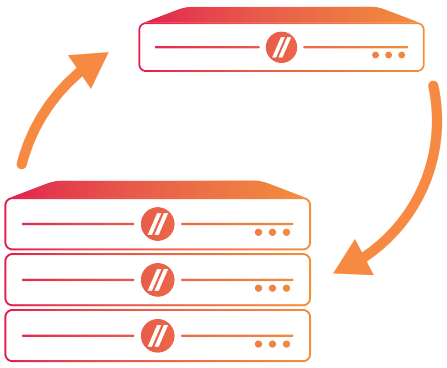

By only protecting the critical workloads that will allow their business continuity after a disaster, these smaller organizations can avoid downtime by failing over to a single node. Non-essential workloads may not need to be backed up on these systems, or simply backed up to storage without failover. As soon as they recover or replace their primary cluster, they can fail back and return to normal operation.

A single node appliance can be ideal for both distributed enterprise and small environments to meet disaster recovery needs and costs.

## Disaster Recovery as a Service (DRaaS)

For organizations that either lack a remote site for disaster recovery or would simply prefer not to deploy their own disaster recovery site, Scale Computing offers SC//Platform Cloud Unity DRaaS. This disaster recovery as a service offering provides an SC//HyperCore DR target running securely in the Google Cloud Platform. Workloads can be replicated to the Google Cloud for failover or recovery per VM. Your on-premise system already acts as a private cloud, so cloud-based disaster recovery fits perfectly.

SC//Platform Cloud Unity DRaaS uses L2 networking to provide seamless connectivity between on-premise and remote-hosted VMs in the event of a failover. This service uses the built-in snapshot and replication features that can provide RPO and RTO measured in minutes - without a VPN required for connectivity. This option allows predictable pricing that can protect anywhere from a single VM to any number of VMs on SC//HyperCore clusters. When needed, all protected VMs can be failed over and running in the cloud and then failed back once the on-premises resources are restored.

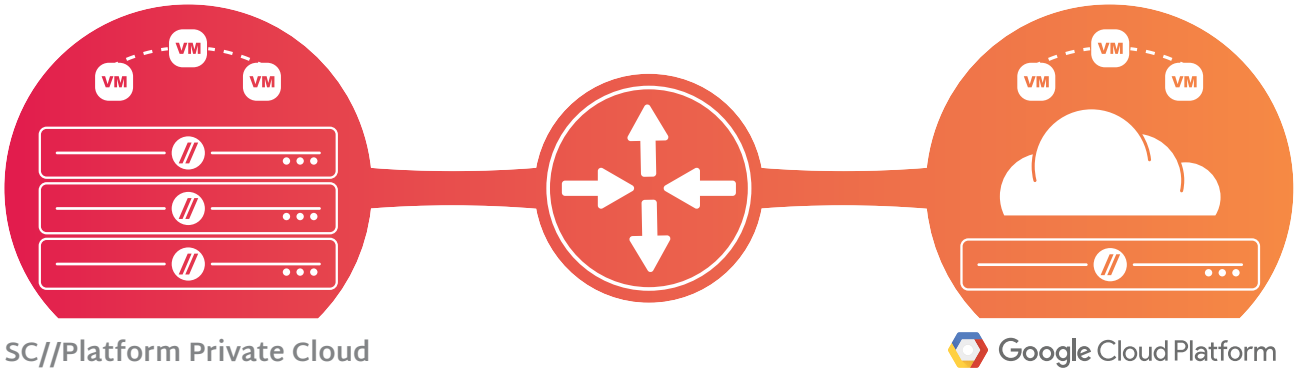


Whether an organization does not have a second site or would rather not manage one, DRaaS is a perfect fit. The combination of predictable cost and reliable recovery provides peace of mind for business continuity.

SC//Platform Cloud Unity DRaaS includes award-winning ScaleCare support at every stage to assist in setup, testing, failover, and recovery. The service also comes with a runbook to assist with planning and execution. For more information on SC//Platform Cloud Unity DRaaS please visit **www.scalecomputing.com/google**.

## On-Prem vs. DRaaS

Choosing between hosting a DR site or using DRaaS is not always easy. There are many factors to consider, and these are some of the biggest:

### Remote Site

Many organizations have more than one site, and one of these sites may function as a DR site. Even organizations without their own remote site may already be leasing space in a remote hosting facility for some of their IT needs. Either of these options could be appropriate for DR, but not always. Location is important, especially if the remote site is within the same metropolitan or regional area. The farther the remote site, the better protected the data becomes because some natural disasters affect entire regions. If a suitable remote site is available, then on-premise DR may be more cost-effective. If a remote site does not already exist, DRaaS is probably the best choice by default.

### Security and Compliance

Your data is valuable. Liability for data breaches can be expensive, and regulations may dictate specific levels of security.

Data security at a remote DR site is just as important as security at the main office/data center. If a remote site has the same level of security as the main office, then on-premises DR may be the best choice. It is important to determine if a remote site complies with industry-specific or general data protection regulations. For example, if your remote site is across an international border, it may not comply with regulations for your primary site data. If the remote site is not suitable or not in compliance, DRaaS data centers may offer secure and compliant computing environments with the cost of security built into the hosting costs. An organization must consider the time and cost of securing the remote site versus using secure DRaaS.

### Management

Hosting a DR site requires additional management of IT resources at the remote site. There may not be IT staff at the remote site, and the site will at least require an initial on-site setup of hardware and routine remote management and maintenance.

DR sites should be a significant distance from primary sites, and the further the distance, the greater the cost may be for managing the site. With DRaaS, the service manages the remote infrastructure. Additionally, DRaaS includes the setup of DR protection, assistance in failing over, and recovering data and services in the event of a disaster.

An organization with the resources to manage remote DR could save by implementing on-premise remote DR. For other organizations, it may be more cost-effective to employ DRaaS and all the included management resources.

| Primary Deciding Factors | | |
|---|---|---|
| Does the organization have a suitable remote site? | Yes | No |
| Is the remote site secure? | Yes | No |
| Can the organization manage the remote DR site? | Yes | No |

## Third-Party Options

Built-in disaster recovery features provide a complete and effective backup, failover, and recovery strategy for most organizations, but some organizations require more specialized disaster recovery. Some organizations may choose to deploy agent-based backup agents on individual VM workloads due to specialized workloads or specific compliance needs.

Scale Computing supports any in-guest backup agents designed to run on Intel-based virtual machines on our supported OS platforms (Windows, Linux, etc). These backup agents generally interact with the application directly.

These solutions create specialized backups for various recovery scenarios or are more specialized in creating system-state backups for recovery between hardware and virtualization platforms.

A third-party, agent-based solution is required if a disaster recovery strategy involves failing over or recovering VM workloads to a solution other than SC//HyperCore. Examples of third-party, agent-based solutions include Acronis, Symantec, Unitrends, etc.

For workloads that require a more aggressive level of RPO and RTO down to seconds rather than minutes, there are third-party agent-based replication solutions like Double-Take Availability. These solutions provide automatic failover for maximum RPO and RTO over any geographical distance. This option is for highly critical workloads where nearly any data loss or downtime, even minutes' worth, is unacceptable.

Many third-party options may be used with SC//HyperCore. All may be considered as part of a disaster recovery strategy in addition to the built-in SC//HyperCore features. Scale Computing is committed to providing the best, most complete infrastructure solution in the market, including complete disaster recovery.

## Acronis Backup on Scale Computing HyperCore

Our strategic partnership with Acronis offers many benefits to our customers. Acronis provides an excellent assortment of advanced backup and recovery features. Acronis uses an agent-less-based approach to backups on SC//HyperCore, allowing for granular backup of files, individual VMs, or all the VMs on an SC//HyperCore system.

### Low-cost, Long-term Retention Options for Regulated Industries

- Archive locally using an existing SAN/NAS, a low-cost storage array, or public cloud storage. Long-term retention allows regulatory compliance requiring years of data recovery. The ability to back up data directly to other storage platforms provides multiple options for storing data at the most reasonable cost.

### Proactive ML-based Ransomware Protection

- Secure all data, including network shares and removable devices, from ransomware attacks. The machine learning models trained in the Acronis Cyber Infrastructure recognize numerous ransomware types and suspicious behavior to prevent an actual attack. Acronis Active Protection detects zero-day attacks, stops questionable behavior, and automatically recovers damaged files.

### Automated, Remote, and Bare-metal Recovery

- Accelerate your recovery up to two times with a complete system image ready for reinstallation and smart technology that automatically detects boot requirements. Restoring a full system on a computer or virtual machine with an empty “bare-metal” disk drive is a snap.

### vmFlashback

- Recover after a failure up to 100 times faster with minimum network impact. This technology tracks and saves only changed blocks of information in the backup to reduce virtual machines’ recovery time significantly.

### Acronis Universal Restore

- Restore an entire system to new, dissimilar hardware with a few clicks. This technology ensures quick, easy system migration between physical and virtual platforms by overcoming compatibility issues. You’ll be up and running on a new virtual or physical machine in minutes.

### Convert to VM- Any VM You Need

- Utilizing the convert to VM functionality in Acronis Cyber Protect Cloud, users can back up VMware, Hyper-V or even physical workloads and instantly power them on as VMs on SC//HyperCore at a moment’s notice. This allows SC//HyperCore to function as a target.

## Conclusion

Scale Computing HyperCore provides the ability to create a complete disaster recovery strategy with any combination of easy to use, native features or third party solutions. The combination of snapshot technology, replication, failover, recovery and DRaaS is one of the reasons Scale Computing is leading the market with innovation and ease of use. Scale Computing's goal of eliminating IT complexity from infrastructure isn't just about hardware. Including virtualization and disaster recovery makes SC//HyperCore the most hyperconverged infrastructure solution on the market.

## Additional Resources

- SC//Platform Cloud Unity Theory of Operations
- SC//HyperCore Replication Setup Video
- SC//HyperCore File Recovery using SnapClone
- Acronis Backup for SC//HyperCore Data Sheet

Corporate Headquarters
525 S. Meridian Street - 3E
Indianapolis, IN 46225
P. +1 317-856-9959
scalecomputing.com

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

SCALE COMPUTING