



Disaster Recovery Strategies

From the Experts at
Scale Computing



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



Table of Contents

Introduction.....	3
HC3 Data Protection Suite	3-5
Backup.....	3
Replication.....	4
Failover.....	4
Recovery	5
ROBO and Small Environments	6
Distributed Enterprise.....	6
Small Environments.....	6
Disaster Recovery as a Service	7
On-Prem vs. DRaaS	7
Third-Party Options.....	8
Acronis Backup on Scale Computing HC3	9
Summary.....	10

Introduction

Disaster recovery is a concept that asks the question, “How can an organization survive and respond to a wide variety of threats ranging from small hiccups to catastrophic destruction?” The threats to ongoing operations range from human error to malicious attacks to natural disasters. Organizations need to prepare in ways that involve both human and technological response. At Scale Computing, we recognize that in today’s 24/7 marketplace, IT infrastructure must be both resilient and highly available to keep organizations operational.

In our HC3 architecture, keeping in mind our typical simplicity and ease of use, we have built-in a number of disaster recovery capabilities. These allow our users to recover quickly from a variety of disasters that may affect anywhere from a single file to an entire site. Disaster recovery is often planned for and measured in terms of recovery point objective (RPO) and recovery time objective (RTO). HC3 provides features to achieve both RPO and RTO measured in minutes to minimize both downtime and data loss.

At Scale Computing, we build all of our solutions with three primary considerations: simplicity, scalability, and availability. We approach disaster recovery with these same considerations. The simplicity consideration is of significant value, as the prevailing strategies and approaches add complexity and cost. This document will outline the strategies and built-in technologies that can be used to protect both data and workloads on HC3 to get services back online as quickly as possible following even the worst disasters.

HC3 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 HC3, anyone can implement a simple, yet robust disaster recovery strategy that combines per VM snapshot scheduling with replication, failover, and recovery. Each HC3 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 enable flexible scheduling. 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. 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

A user’s 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 extremely useful for local recovery of data from a number of operational disasters. For a true backup strategy, snapshots must be replicated onto another device, preferably at another site.

Replication


HC3 native replication capabilities will replicate snapshots on a per VM basis to another HC3 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 HC3 appliance for full 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 it can travel over any distance to any remote site. Replication is not synchronous so it does not require high speed links that are expensive and restrict replication within campus or metropolitan areas. 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 can be used with HC3 to help overcome some bandwidth and latency challenges between remote sites.

Replication can be directed to another HC3 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 an HC3 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, consisting of 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. This process can be completed very quickly, in a matter of minutes after failure, providing a very solid RTO.



It is important to understand dependencies 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 of the process, 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 data as it is now in production. 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. If the original VM is intact but now outdated, as is often the case, data can be re-stored quickly with HC3 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 just require a bit more time and effort to restore all the data.





Recovery

While failover is also considered recovery of an entire VM, there are cases where it is neither desirable nor practical to failover 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. In these cases, the data can indeed be recovered either on an individual file basis or as an entire VM.

For file-level recovery, an HC3 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 of 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 powerful point-in-time file recovery that is 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 whole 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 seconds and then you may begin restoring the VM back to the source site. This is a perfect scenario for non-essential VMs that are not needed online during disaster recovery scrambles.


These built-in recovery techniques, along with replication and failover can provide a complete disaster recovery solution for HC3 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. Now, with features that are built-in to HC3, most disaster recovery needs are met without reliance on additional solutions.

ROBO and Small 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. Distributed enterprises with multiple remote offices or branch offices (ROBO) and other small IT environments can be challenged to provide disaster recovery with very limited resources. HC3 addresses these configurations with a single node appliance configuration.

The Distributed Enterprise

In the distributed enterprise, there is typically a healthy IT infrastructure at a central office and then multiple remote or branch offices with minimal IT footprint. In this environment, the HC3 single node appliance 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. For these sites, a full cluster is overkill so the single node appliance makes sense from a cost and resourcing standpoint. 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.




Each single node appliance in the remote offices can be managed remotely and replicated back to an HC3 cluster at the central office for backup, failover, and recovery. When one of these appliances fails, the workloads can be failed over 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 HC3 cluster in production for local high availability. However, in the event of a site disaster they can probably weather the storm running their critical workloads on a single node appliance until the primary site is recovered. 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.

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




A single node appliance can be an ideal configuration for disaster recovery in both distributed enterprises and small environments to meet both 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 HC3 Cloud Unity DRaaS. This disaster recovery as a service offering provides an HC3 DR target running securely in Google Cloud Platform. Workloads can be replicated to the Google cloud for failover or recovery on a per VM basis. Your HC3 on-prem system is already acting as a private cloud so cloud-based disaster recovery fits perfectly.

As a service, this option allows predictable pricing that can protect anywhere from a single VM to any number of VMs on HC3 clusters. This service uses the built-in snapshot and replication features that can provide RPO and RTO measured in minutes. There is no VPN required for connectivity. HC3 Cloud Unity DRaaS uses L2 networking to provide seamless connectivity between on-prem and remote hosted VMs in the event of failover



HC3 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 both planning and execution. When needed, all protected VMs can be failed over and running in the cloud and then failed back once the on-prem 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 will provide piece of mind for business continuity. More information on HC3 Cloud Unity DRaaS can be found at 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 be used 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 might be appropriate for DR, but not always. Location is important, especially if the remote site is within the same metropolitan or regional area. The further the remote site, the better protected the data because some disasters affect entire regions such as floods, earthquakes, or hurricanes. If there is a suitable remote site available, then on-prem DR may be a more cost-effective solution because there may not be additional hosting fees. If a remote site does not already exist, DRaaS is probably the best choice by default.

Security and Compliance

Your data is valuable and liability for data breaches can be expensive and regulations may dictate specific levels of security. Security of data at the remote DR site is just as important as security at the main office/datacenter. If a remote site is already considered to have the same level of security as the main office, then on-prem 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 datacenters may offer secure and compliant computing environments with the cost of security built into the hosting costs. The time and cost of securing the remote site vs. using secure DRaaS must be considered.

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 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 remote infrastructure is managed by the service. Additionally, with DRaaS, setup of DR protection is included as well as assistance in failing over and recovering data and services in the event of disaster. An organization that has the resources to manage remote DR could save by implementing on-prem remote DR. For other organizations, it may be more cost-effective to employ DRaaS and all of 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. Whether because of specialized work-loads or specific compliance needs, some organizations may choose to deploy agent-based backup agents on individual VM workloads.

HC3 supports any in-guest backup agents that are designed to run on Intel-based virtual machines on our supported OS platforms (Windows, Linux, etc – see **Support Matrix**). 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.

If a disaster recovery strategy involves failing over or recovering VM workloads to a platform other than HC3, a third-party, agent-based solution is required. Examples of third-party, agent-based solutions include Acronis Backup, Symantec Backup Exec, Unitrends Enterprise Backup, and many others.

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.

There are many third-party options that may be used with HC3 and these may all be considered as part of a disaster recovery strategy in addition to the built-in HC3 features. Scale Computing is committed to providing the best, most complete infrastructure solution in the market and that includes complete disaster recovery.

Acronis Backup on Scale Computing HC3

Scale Computing has selected Acronis Backup Advanced to sell directly to our HC3 customers because it provides an excellent assortment of advanced backup and recovery features. Acronis backup uses an agent-based approach to backups on HC3 allowing for granular backup of files, individual VMs, or all the VMs on an HC3 system. Here is a list of the features we expect our customers to take advantage of on their HC3 systems.

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

- Archive locally using an existing SAN/NAS, a low-cost storage array, or to Public Cloud storage. Long-term retention allows compliance to regulation requiring years of data recovery. The ability to direct backup data 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 Cloud AI infrastructure recognize numerous ransomware types and suspicious behavior and prevent actual ransomware from attacking. 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 that's 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 now 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 significantly reduce the recovery time of virtual machines.

Acronis Universal Restore

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



Granular Recovery

- Search for specific documents, emails, and the like, and restore the lost data itself without having to recover full databases or systems. This technology ensures quick, incremental restoration of needed pieces of data and simultaneously reduces recovery times.

Acronis Backup Advanced is ideal for users with requirements beyond disaster recovery and high availability. While nearly any backup solution can be used with HC3, Acronis provides one of the most well-rounded and powerful feature sets for advanced backup and recovery options on the market.

Summary

HC3 provides the ability to create a complete disaster recovery strategy with any combination of easy to use, native HC3 features or third party solutions. The combination of snapshot technology, replication, failover, recovery and DRaaS built-in to HC3 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 HC3 the most hyperconverged infrastructure solution on the market.

Additional Resources

- [HC3 Cloud Unity Datasheet](#)
- [HC3 Replication Setup Video](#)
- [HC3 File Recovery using SnapClone](#)
- [Windows File Recovery ISO Guide](#)
- [Acronis Backup for HC3 Data Sheet](#)

