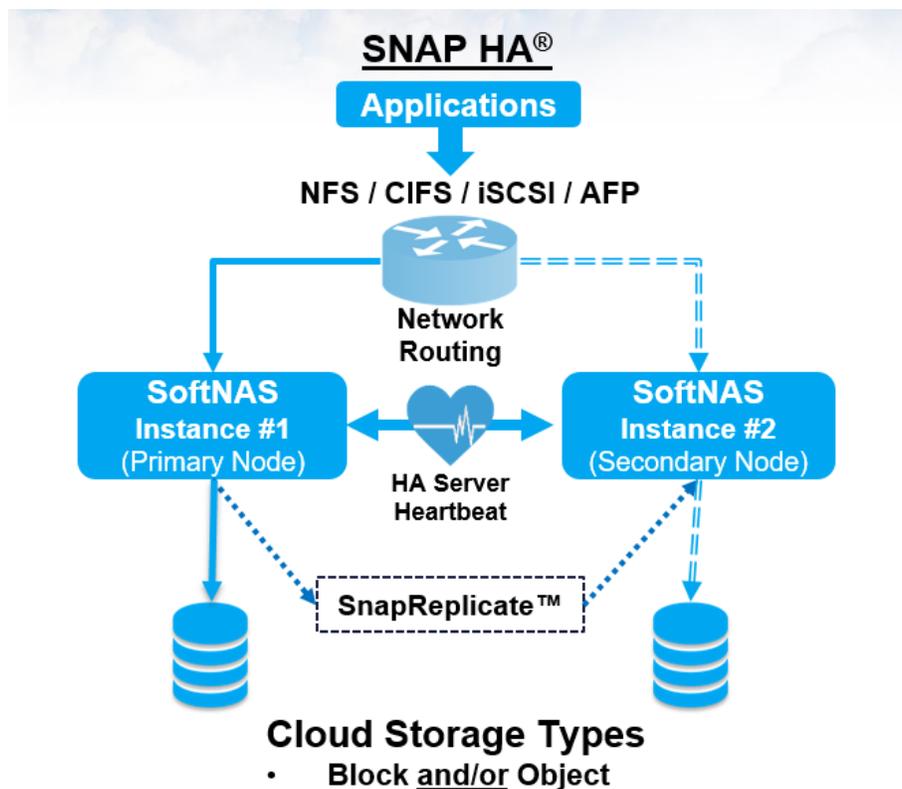


AWS Getting Started: Choosing Your HA Solution

Customers using our product for the first time, or introduced to our newest HA feature, Dual Controller HA, for the first time, may be uncertain of which option best suits their needs. Each high availability option has its own benefits and limitations, which will be covered below, allowing you to make an informed decision. Bear in mind that this does not mean a decision must be made between the two, as both solutions can be used in conjunction with the other. Dual Controller HA is an integrated part of SNAP HA™ but for the purposes of this discussion, SNAP HA will refer to our standard SNAP HA™ replication technology, and Dual Controller HA (or DCHA) will refer to our shared object storage solution.

SNAP HA™



Existing customers will need no introduction to our SNAP HA™ product, which offers a low-cost, low-complexity solution for high-availability clustering that is easy to deploy and manage. A robust set of HA capabilities protect against data center, availability zone, server, network and storage subsystem failures to keep business running without downtime. SNAP HA™ for Amazon Web Services (AWS) includes patent-pending Elastic HA™ technology, providing NAS clients in any availability zone uninterrupted HA access to the storage cluster across availability zones.

SNAP HA™ operates by maintaining a redundant copy of all storage on the primary node, and constantly replicating this data to a secondary node, ensuring that both copies remain synchronized through our patented SnapReplicate™ technology. **SNAP HA™** then monitors all critical storage components, ensuring they remain operational and when there is an unrecoverable failure in a system component, another storage controller detects the problem and automatically takes over, ensuring that no downtime and no business impacts (or very limited business impacts) occur. The data on the second storage controller is identical to the primary node and is presented immediately and seamlessly, ensuring that the end-user never knows an outage occurred.

Key Differences:

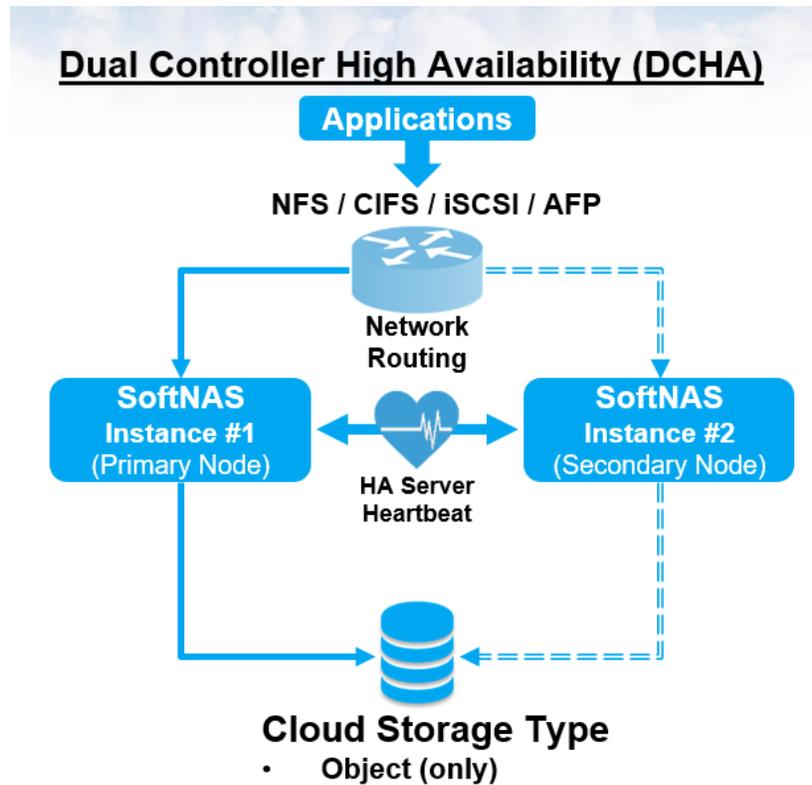
- **Duplication of Storage:** Standard SNAP HA™ replicates data from one node to another. This means that the target node will have copies of the pools and volumes on the primary. This means that each node must maintain an equivalent amount of storage.
- **Storage Type Flexibility:** SNAP HA™ can use any combination of block and object storage. Pools and Volumes can be created from:
 - **Block storage including:**
 - Azure Disk
 - AWS EBS
 - **Object Storage including:**
 - AWS S3
 - AWS S3 Infrequent
 - Azure Cold Blob
 - Azure Hot Blob
 - S3 compliant object storage
- **Separate Locations:** SNAP HA™ enables customers to rely on two completely separate storage systems - where data is completely replicated to 2 different locations ensuring any primary storage use cases remain operational. DCHA is reliant on object storage in a single location - if the cloud location is compromised (for example an AWS region-wide failure), the storage will become unavailable.
- **Faster Recovery:** For most use cases, our standard SNAP HA™ solution offers faster fail-over and recovery times in comparison to DCHA. This is because DCHA requires several ZFS checks behind the scenes to complete the transfer of ownership of the shared storage pool.

Performance and Data Integrity Considerations:

SoftNAS' ZFS based solution provides a great deal of protection to ensure that your data is fully protected, but in a fail-over situation, default settings can potentially result in uncached write bursts not being committed to the target volumes. SoftNAS balances these concerns with performance considerations. If data integrity is of paramount importance, and you want the enhanced protection afforded by hosting your data in two separate locations, we recommend changing the **Sync Mode** when creating your pool to the '**Always**' setting. This setting ensures every file system transaction is written and flushed to stable storage by a system call return, but at a significant performance cost. See [Create a Storage Pool](#) for more information about **Sync Mode** settings.

Alternatively, performance can be maintained while still further protecting against potential data loss by creating a ZIL, or as it is referred to within the SoftNAS UI, a Write Log. This process is covered in [Configuring the Read Cache and Write Log](#).

Dual Controller HA (DCHA)



Dual Controller HA™ (or DCHA) is an extension to our existing SoftNAS Cloud® high availability solution, SNAP HA™. It is designed to provide high availability for a shared pool of object storage. DCHA only applies if a shared pool of object storage, such as AWS S3, or Azure Hot or Cool blob storage, is specified at storage pool creation. After [adding object storage 'disks' via Disk Devices](#), and selecting **Create** in [Storage Pools](#), the following dialog will appear.



If Shared Storage is selected, Dual Controller HA™ will automatically be applied to the shared pool after SNAP HA™ is configured.

Regardless of whether it is a shared pool or dedicated, the customer must first define a SnapReplicate™ relationship between the primary and secondary node, then add the SNAP HA relationship. In other words, there is no change to the [SnapReplicate/SNAP HA process](#). Adding a device to a shared storage pool results in the pool being excluded (skipped) by SnapReplicate; i.e., the data on the underlying device is already shared across nodes, so there is no need to replicate shared storage pools.

This allows SnapReplicate and SNAP HA to function across both types of pools, and to differentiate between them. Existing SNAP HA customer installations continue to operate uninterrupted, and new SoftNAS instances can be paired with both Dual Controller HA shared storage pools and dedicated pools asynchronously replicating via "standard" SNAP HA simultaneously. This also ensures that regardless of which type of pool selected, the customer can confidently set up SNAP HA with the same documentation.

Key Differences

- **No Duplication:** While both DCHA and SNAP HA require the configuration of two nodes, DCHA does not require the duplication of storage. Underlying storage devices, whether they be Azure Blob or AWS S3 object storage, are shared across the two nodes.
- **Object Storage Only:** DCHA applies only to object storage. Object Storage such as AWS S3, or Azure Blob Hot or Cold storage, is typically accessed by network connection, allowing it to be accessed by two or more nodes (only two nodes are supported by DCHA for the moment). Block storage cannot be shared in the same way.
DCHA can be used with the following:
 - AWS S3
 - AWS S3 Infrequent
 - Azure Cold Blob
 - Azure Hot Blob
 - S3 compliant object storage
- **Single Storage Location:** DCHA protects the storage controller and the storage media via object storage; however, it relies on a single data store that may be disrupted if a cloud element experiences an outage.

- **Slower Recovery:** Because standard SNAP HA™ only has to spin up local pre-existing storage, the failover window is much smaller than with Dual Controller HA™, which requires several ZFS data integrity checks to complete.