

Synology High Availability (SHA): An Introduction



Table of Contents

Introduction	2
Overview	3
Synology High Availability Architecture	5
Hardware components	5
Network interface	7
Network scenarios	9
Data replication	11
Special conditions	12
Dual Controller Features	14
Performance and data protection	14
Achieving Service Continuity	16
Auto-failover protection	16
Best Practices for Deployment	19
System performance	19
Reliability	20
Performance Benchmark	23
Performance considerations	23
Switchover time-to-completion	25
Summary	25

Introduction

Business challenges

Unexpected downtime can lead to frustration for customers and result in huge losses in revenue. 50% of SMBs worldwide remain unprepared in the case of disaster, and in many cases downtime becomes a potentially fatal problem costing companies as much as over 12,000 USD a day. With the growing demand for uninterrupted availability, businesses seek for solutions that ensure high levels of service continuity.

Synology solutions for service continuity

High availability is a high demand solution for anyone deploying critical services such as databases, company file servers, virtualized storage, and more. All of these services are extremely low tolerance and cannot afford to have services interrupted when unexpected events strike.

High availability is mostly featured as an enterprise-exclusive solution due to its high cost and complicated deployment. However, with **Synology High Availability**, this high-end feature is available on most plus-series and all FS/XS-series, making it a cost-effective solution for protecting important services. Synology High Availability mitigates the impact for IT administrators to fix any system or hardware issues when disaster strikes, while allowing businesses to prevent downtime for mission-critical applications and minimize lost revenue associated with the incident.

Overview

According to Transparency Market Research (TMR), the global market for high availability servers is anticipated to rise at a CAGR of 11.7% between 2017 and 2025. The market is valued at \$4.1bn in 2015 and is expected to reach \$12.3bn by 2025.

A key factor driving the global market demand for high availability solutions is the reliance on data, and the need for data to be more accessible. There has been an increase in demand for a higher level of availability to prevent lost revenue caused by unexpected or undesired incidents that may have a profound impact on data access and an organization's business productivity.

High availability solutions, supporting redundancy and data replication, have become a proven strategy for organizations to retain and manage data while being able to access information in real-time. However, due to the complex and expensive nature of the high availability technology, small and mid-size businesses often lack the resources or budget to implement a high availability solution to protect against data loss. In most cases, only large enterprises with sufficient IT resources can afford to build a business continuity plan, deploy highly available and fault tolerant servers, and store an ever-growing amount of critical data generated on a daily basis to be kept secure and available at all times.

Synology is dedicated to providing a reliable, cost-effective, comprehensive high availability infrastructure to small and mid-range businesses to effectively minimize data loss and downtime.

Synology High Availability (SHA) is introduced as an enterprise-level solution that is affordable especially for smaller organizations with limited IT resources seeking for robust, continuous system availability at a lower budget without worrying about installation and maintenance costs. Contrary to most HA solutions that require expensive, dedicated hardware, SHA is available on most Synology NAS and can be implemented at a lower investment cost.

Key features of Synology High Availability:

- **High availability cluster** provides comprehensive hardware and data redundancy
- **Heartbeat** technology achieves real-time data synchronization
- **Automatic failover** maximizes service uptime and business continuity
- SHA ensures storage for performance-intensive workloads including virtualization solutions such as VMware® vSphere™, Microsoft® Hyper-V®, Citrix® XenServer™, and OpenStack Cinder.
- Offers hassle free high availability configuration, management, and maintenance without requiring additional technical resources through an intuitive set up wizard that comes with built-in, visualized management tools
- Complete protection for file services including SMB, AFP, NFS, FTP, iSCSI, Synology Active Directory, Fibre Channel, and the DSM login page.

- Supports uninterrupted updates of the **Synology High Availability** package.

This white paper aims to provide information on Synology High Availability (SHA) design and architecture, common scenarios, best practices, and performance metrics.

Synology High Availability Architecture

The Synology High Availability solution is a server layout designed to reduce service interruptions caused by system malfunctions. It employs two servers to form a **high-availability cluster** (also called **HA cluster**) consisting of two compatible Synology servers. Once this high-availability cluster is formed, one server assumes the role of the active server, while the other acts as a standby passive server. Full data replication is required to be performed once the cluster is successfully created.

Once the high-availability cluster is formed, data is continuously replicated from the active server to the passive server. All files on the active server will be copied to the passive server. In the event of a critical malfunction, the passive server is ready to take over all services. Equipped with a duplicate image of all data on the active server, the passive server will enable the high-availability cluster to continue functioning as normal, minimizing downtime.

Hardware components

Synology's high availability solution constructs a cluster composed of two individual compute and storage systems: an active and a passive server. Each server comes with attached storage volumes, and the two are linked by a **Heartbeat** connection that monitors server status and

facilitates data replication between the two servers.

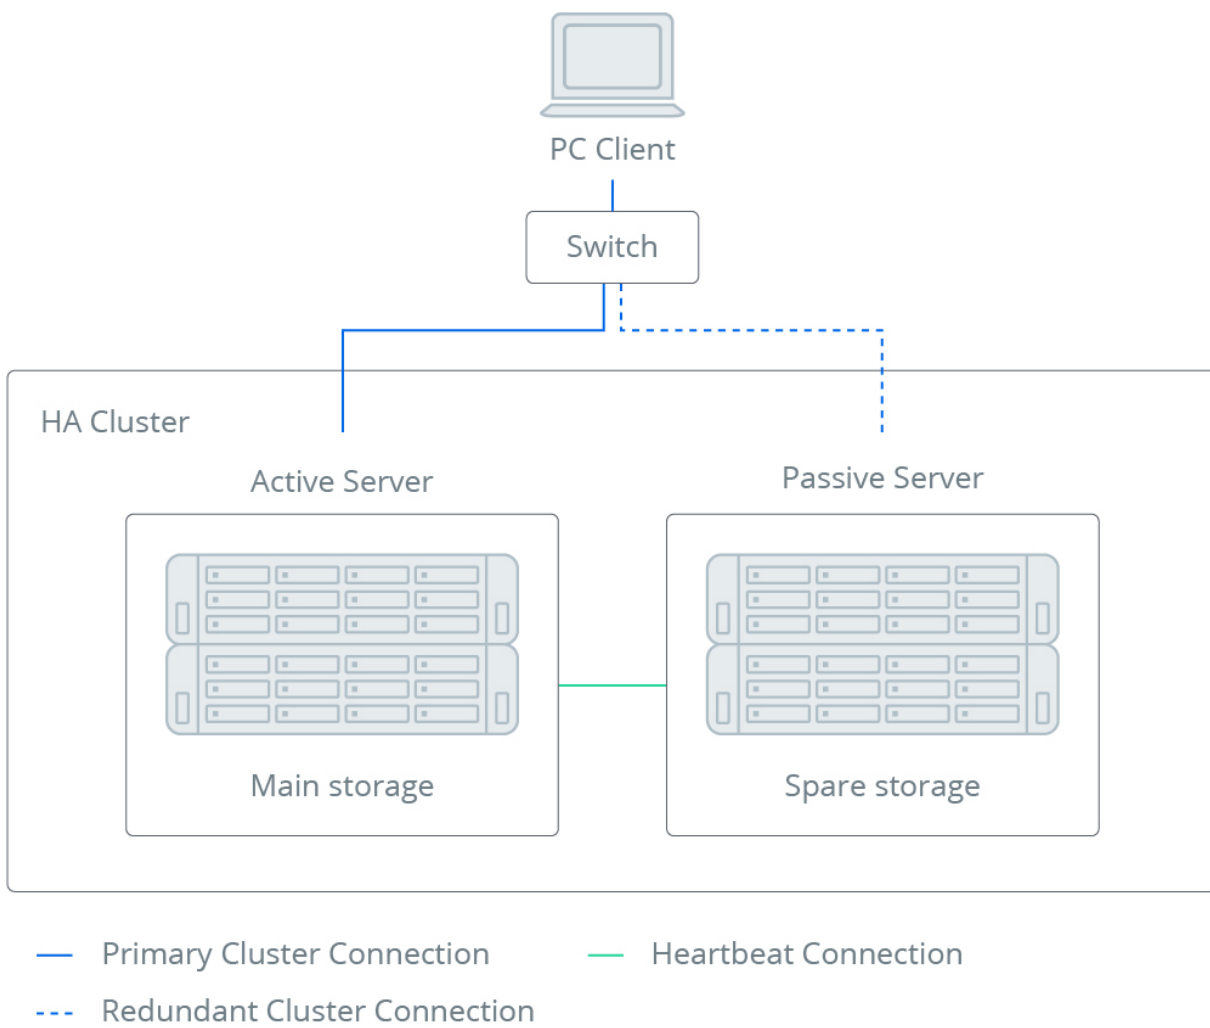


Figure 1: Physical components of a typical Synology High Availability (SHA) deployment

- **Active Server:** Under normal conditions, all services are provided by the active server. In the event of a critical malfunction, the active server will be ready to pass service provisioning to the passive server, thereby circumventing downtime.
- **Passive Server:** Under normal conditions, the passive server remains in standby mode and receives a steady stream of data replicated from the active server.
- **Cluster Connection:** The connection network used for the communication between the clients and high-availability cluster. There is at least one cluster connection for the active server, and one for the passive server, to the client. In order to ensure the communication between both active and passive servers, the cluster connections must go through a switch.
- **Heartbeat Connection:** The active and passive servers of a high-availability cluster are connected by a dedicated, private network connection known as the "Heartbeat" connection. Once the cluster is formed, the Heartbeat facilitates data replication from the active server to the passive server. It also allows the passive server to constantly detect the active server's presence, allowing it to take over in the event of active server failure. The ping response time between the two servers must be less than 1 ms, while the transmission speed should be at least 500 Mbps. The performance of the HA cluster will be affected by the response time and bandwidth of the Heartbeat connection.
- **Main Storage:** The storage volume of the active server.
- **Spare Storage:** The storage volume of the passive server, which continually replicates data received from the main storage via the Heartbeat connection.

Network interface

Cluster interface

When the two servers are combined into a high-availability cluster, a virtual interface (unique server name and IP address) shall be configured. This virtual interface, also called the cluster interface, allows clients to access the cluster resources using a single namespace. Therefore, when a switchover is triggered and the provision of services is moved to the passive server, there will be no need to modify network configurations on hosts in the data network.

Only one cluster interface is configured during cluster creation. Additional cluster interfaces can be added on the **Network** page in Synology High Availability. One of the cluster interfaces is designated as the **Primary Cluster Interface**, which is responsible for communication between the

active and passive servers.

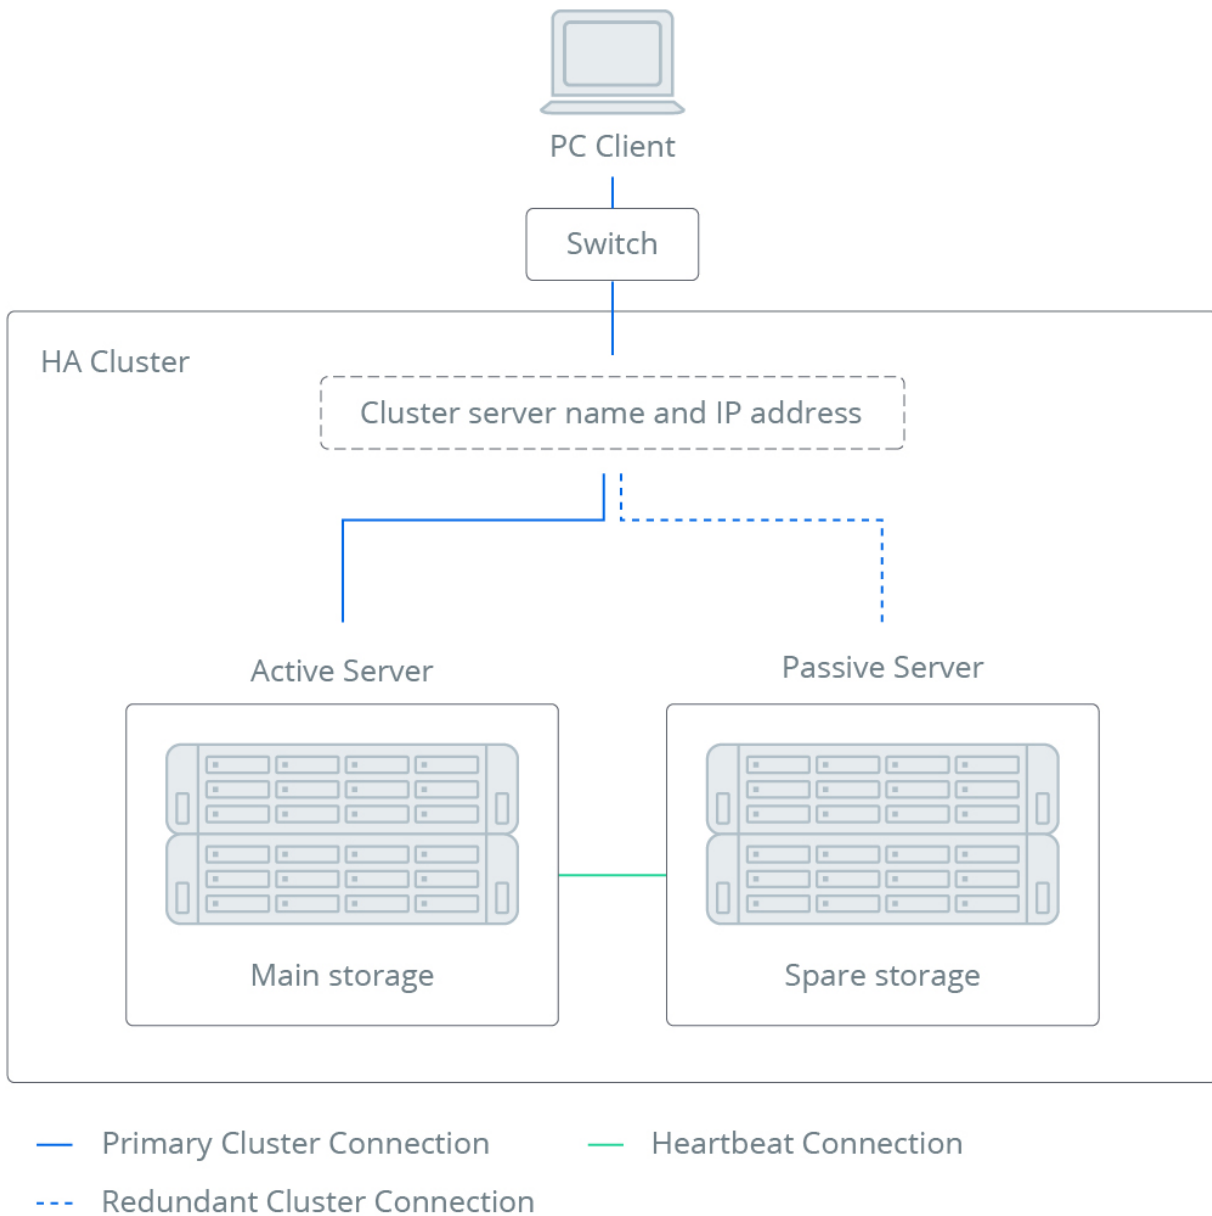


Figure 2: PC clients access a Synology High Availability (SHA) cluster through a single virtual interface

- **Cluster Server Name and IP addresses:** Servers in the cluster will share IP addresses and a server name, which should be used in all instances instead of the original IP addresses and individual server names.

Heartbeat interface

The **Heartbeat** connection is the connection established on the Heartbeat interfaces of the active and passive server and is used for replicating data from active server to passive server, including the differential data, and the real time write operation. All data syncing is performed at block-

level, and it ensures that the active and passive servers contain identical data. As all data is constantly maintained to be up-to-date, switchover can be accomplished seamlessly.

The Heartbeat IP address is automatically selected by system upon the cluster creation.

- The system will randomly select an IP address (169.254.x.x) for the Heartbeat interface.
- **Suggested network configurations:** We suggest you choose the fastest network interface for the Heartbeat connection, or ensure the ability of the network interface is the same as the network interface of the cluster connection. Refer to the following examples:
 1. Both servers have at least two 10GbE network interfaces. One 10GbE is suggested to be the interface of Heartbeat connection, and the other is for the cluster connection.
 2. There is only one 10GbE interface between both servers. It should be used for the Heartbeat connection. If there are more than two 1GbE network interface, it is also suggested to set up link aggregation for the cluster connection.
 3. If there is no 10GbE network interface, make sure the Heartbeat connection and the cluster connection share network interfaces.

Link Aggregation increases the bandwidth and provides traffic failover to maintain network connection in case the connection is down. It is recommended to set up link aggregation for both the **cluster connection** and the **Heartbeat connection**. Note that the link aggregation must be set up before the creation of high-availability cluster. Once the HA cluster is created, the link configuration cannot be modified again. In addition, no matter what type of link aggregation that is chosen for Heartbeat connection, SHA creation will set it to **round-robin** automatically.

Network scenarios

The physical network connections from the data network to the active server and passive server must be configured properly so that all hosts in the data network can seamlessly switch connections to the passive server in the event a switchover is triggered. The following section covers different configurations for various situations and Synology NAS models.

Network implementation for Synology NAS with two LAN ports

In situations where both servers have two network ports only, one network port on each server will be occupied by the Heartbeat connection, so each server will have only one port available for the HA cluster to connect to the data network. Therefore, there will not be sufficient network ports to accommodate redundant paths between the hosts in the data network and HA cluster. However, we still recommend using multiple paths to connect hosts to the data network, as well as more than one switch in your data network to provide redundancy.

Synology High Availability (SHA) provides an option to trigger a switchover when the active server detects network failure. When enabled, if connection failure occurs between the switch connected to the active server or the switch fails, service continuity will be maintained by

switching over to the passive server (assuming the network connection of the passive server is healthy).

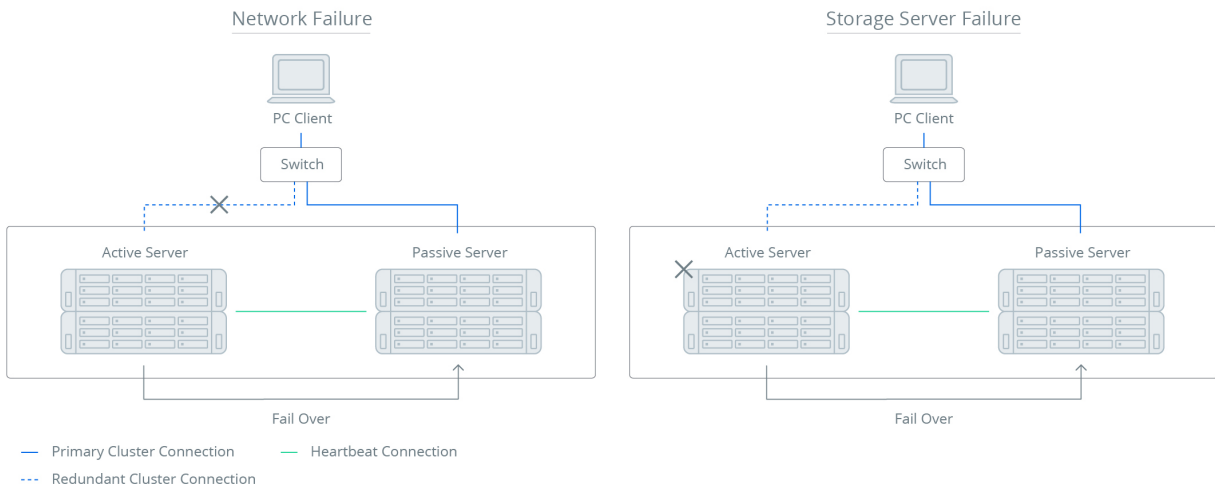


Figure 3: high-availability cluster network configuration on models with two LAN ports

Network implementation for Synology NAS with four or more LAN ports

The best way to create a high availability environment is to use a Synology NAS with at least **four network ports**. In this instance, you can connect multiple paths between the hosts and HA cluster, providing a redundant failover path in case the primary path fails. Moreover, I/O connections between the data network and each clustered server can be connected to more than one port, providing a load balancing capability when all the connections are healthy.

Path redundancy for the Heartbeat connection

For Synology NAS models with four or more network ports, link aggregation may be implemented on the Heartbeat connection to provide failover redundancy and load balancing. This feature does not require a switch between the connections.

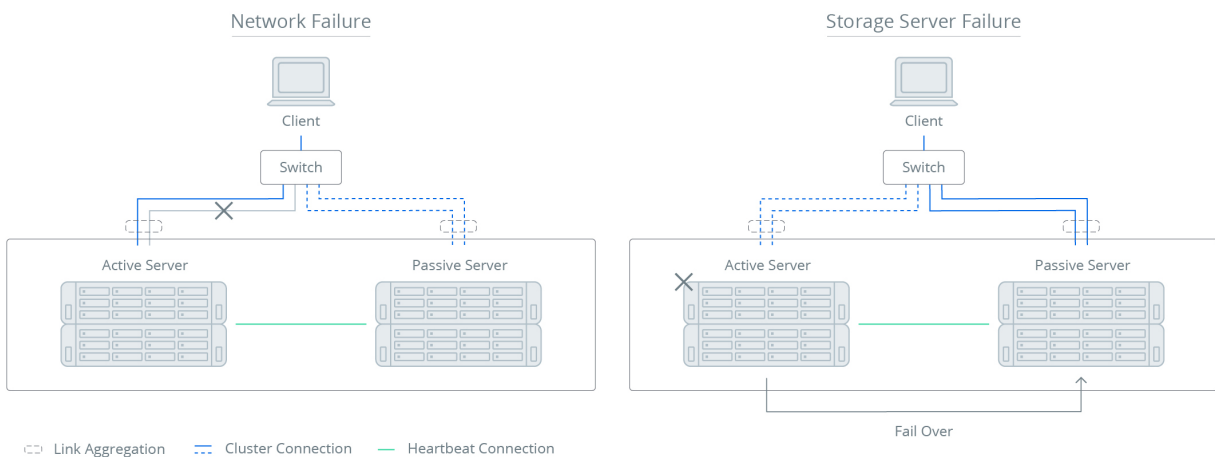


Figure 4: High-availability cluster network configuration on models with four or more LAN ports (NAS)

Network troubleshooting

- The maximum transmission unit (MTU) and virtual LAN (VLAN) ID between Synology NAS and switch/router must be identical. For example, if the MTU of DS/RS is 9000, do make sure the corresponding switch/router is able to measure up to that size.
- The switch/router should also be able to perform fragmentation and jumbo frame (MTU value: 9000) if the Heartbeat connection goes through the switch/router.
- Ensure the firewall settings do not block the ports for DSM and SHA.
- Ensure that the IP addresses of the active/passive server and of the HA cluster are in the same subnet.
- Unstable Internet (reduced ping rate or slow internet speeds) after binding:
 - Try connecting to another switch/router or connect an independent switch/router to DS/RS and PC/Client for testing.
- The flow control settings on the switch/router would also induce packet loss in the network. Make sure this setting is configured in accordance with DS/RS, which normally will auto-detect and conform to the setting of the corresponding switch/router. Users are advised to manually enable/disable the flow control if any inconsistencies are observed.
- Ensure that Bypass proxy server for local addresses in the proxy setting is enabled.

Data replication

Within the high-availability cluster, all data that has been successfully stored (excluding data that still remains in the memory) on internal drives or expansion units will be replicated. Therefore when services are switched from the active to passive server, no data loss will occur.

While data replication is a continual process, it has two distinct phases spanning from formation to operation of a high-availability cluster:

- **Phase 1:** The initial data replication during cluster creation or the replication of differential data when connection to the passive server is resumed after a period of disconnection (such as when the passive server is switched off for maintenance).
During this phase, the initial sync is not yet complete, and therefore switchover cannot be performed. Data changes made on the active server during this initial replication are also synced.
- **Phase 2:** Real-time data replication after the initial sync has been completed. After the initial sync, all data is replicated in real-time and treated as committed if successfully copied. In this phase, switchover can be performed at any time.

During both phases of data replication, all data syncing is performed at the **block-level**. For example, when writing a 10 GB file, syncing and committing is broken down to block-level operations, and completed piecemeal to ensure that the active and passive servers contain

identical data. As all data is constantly maintained to be up-to-date, switchover can be accomplished seamlessly.

Data and changes to be replicated include:

- **NAS Data Services:** All file services including CIFS/NFS/AFP are covered.
- **iSCSI Data Services:** High-availability clustering supports iSCSI, including iSCSI LUN and iSCSI Target services.
- **DSM and Other Services:** Management applications, including Synology DiskStation Manager (DSM) and its other services and some add-on packages (e.g. Mail Server, Synology Directory Server) are also covered, including all settings and service statuses.

Special conditions

Split-brain error

When a high-availability cluster is functioning normally, only one of the member servers should assume the role of active server. In this case, the passive server detects the presence of the active server via both the Heartbeat connection and primary cluster connection.

If the Heartbeat and primary cluster connections are lost, both servers might attempt to assume the role of the active server. This situation is referred to as a **split-brain** error. In this case, connections to the IP addresses of the high-availability cluster will be redirected to either of the two servers, and inconsistent data might be updated or written on the two servers.

When either of the Heartbeat or primary cluster connections is reconnected, the system will detect the split-brain error and data inconsistency between the two servers, and will enter high-availability **safe mode**.

On the other hand, a **quorum server** helps reduce the split-brain error rate. Users can assign another server to both the active and passive server as the quorum server. For example, a gateway server or DNS server is a good choice because they connect to both servers constantly.

With a quorum server, the following circumstances will be controlled:

- If the passive server cannot connect to both the active and quorum servers, failover will not be performed in order to prevent split brain errors.
- If the active server cannot connect to the quorum server while passive server can, switchover will be triggered in order to achieve better availability.

High-availability safe mode

Instead of performing a complete replication, high-availability safe mode helps users to identify the new active server and re-build the cluster by syncing new data and modified settings from the active server to the passive server.

In high-availability safe mode, both servers and the IP addresses of the high-availability clusters will be unavailable until the split-brain error is resolved. Also, additional information will be shown, including the following:

1. The difference of contents in the shared folders on the two servers.
2. The time log indicating when the server became active.

This information should be found in **Synology High Availability** or the read-only **File Station**. Thus, users would be able to identify the new active server.

When the new active server is selected, all the modified data and settings on the active server will be synced to the passive server. Hence, a new healthy high-availability cluster shall be in place.

Additionally, users can do either of the following:

1. Choose one server to be the new active server and resume high availability services right away (the data on the new passive server will be over-written by that of the active server).
2. Remove one server from the cluster and keep the data of both servers.

To make a complete replication, users should choose one as the active server of the high-availability cluster and unbind the other. Once both servers are restarted, the active server will remain in the high-availability cluster. The unbound server will keep its data and return to Standalone status. Note that a complete replication will entail binding a new passive server onward.

Refer to the **4.7 Split-brain** section in the [Synology High Availability User Guide](#) for detailed information.

Dual Controller Features

Synology High Availability (SHA) offers a reliable solution for ensuring service continuity in the event of unexpected downtime. Those with higher demands for business continuity can utilize Synology's dual-controller models, SA3200D and SA3400D, that come built in with an active-passive high-availability structure.

These models provide comprehensive features and advanced functionalities that allow businesses to enjoy high availability without any compromises.

Performance and data protection

In addition to providing high availability through failover and cluster technologies, Synology's dual-controller solutions also offer features that improve performance and data protection devices.

RAID bitmap

RAID bitmap, also known as write-intent bitmap, is a feature available on certain dual-controller models. RAID bitmap continuously records and updates areas in the RAID that may not be maintaining data consistency during write operations. By enabling bitmap, RAID scrubbing can detect areas of inconsistent data and perform data maintenance only on those areas, which helps to speed up the synchronization process.

Btrfs cache protection

Btrfs cache protection is another feature that enhances data protection on dual-controller models. By default, file system write operations are asynchronous and are only reported successful when they are written to the memory cache. In the event of a system failure, any data in the cache that has not been written to the drive will be lost. On a dual-controller device, Btrfs cache protection ensures that new data is written to the second controller's memory cache. This means that if the primary controller fails during a write operation, the secondary controller can take over and access the data from its own memory cache, thus reducing the risk of data loss.

Non-Transparent Bridge (NTB)

With both RAID Bitmap and Btrfs cache protection, there may be some concerns in regard to the potential performance impact. On Synology dual-controller storage systems, **PCIe Non-Transparent Bridge (NTB)** technology is implemented help address these issues.

By enabling direct memory access between the storage controllers and memory, NTB technology can improve the speed and reliability of data transfers. This can help to reduce the impact of bitmap management on system performance, while also ensuring the integrity of the bitmap in case of a system failure. With Btrfs cache protection, NTB technology is utilized when writing data to the second controller to improve efficiency without significantly impacting system performance. Overall, NTB technology helps optimize the performance and reliability of dual-controller storage systems, especially in high-demand environments.

Achieving Service Continuity

Auto-failover protection

To ensure continuous availability, Synology High Availability allows switching from the active server to the passive server in a normally functioning high-availability cluster at any time.

Switchover can be manually triggered for system maintenance, or automatically initiated in the event of the active server malfunctioning, which is known as **failover**. After the servers exchange roles, the original active server assumes the role of the passive server and enters standby mode. As resources within the cluster are accessed using a single cluster interface, switchover does not affect the means of access.

- **Switchover:** The active and passive server can be manually triggered to exchange roles without interruption to service for occasions such as system maintenance.
- **Failover:** In the event of critical malfunction, the cluster will automatically initiate switchover to maintain service availability.

Auto-failover mechanism

Auto-failover can be triggered by either the active or passive server depending on the situation.

- **Triggered by the active server:** This happens when the active server is aware of system abnormality and attempts to smoothly transfer services from the active server to the passive server. The active server continuously monitors itself to ensure services are functional. When detecting failed management services (e.g. storage space crashed, service error, network disconnection) the active server will halt services in the beginning and then verify that the data on the storage space and system configuration are synced with the passive server.
After this process, the passive server starts to boot up all services. As a result of the transferring process, users will be unable to manage the cluster and services will stop functioning for a brief period of time (depending on the number of services and storage space).
- **Triggered by the passive server:** This happens when the active server is in a state that it is unable to respond to any requests (e.g. accidental shut-down). The passive server tracks the status of the active server through the cluster connection and the Heartbeat connection. Takeover by the passive server will be prompted when network connections with the active server are dropped.

Synology High Availability is designed to protect the system when errors occur under normal status. Services cannot be guaranteed when more than one critical error occurs concurrently. Therefore, to achieve high availability, issues should be resolved immediately each time a failover is performed to allow the cluster to return to normal status.

Failover events

The following situations are common triggers of system failover:

Notes:

- Manual switchover is not possible when the storage space on the passive server is busy with a **Storage Manager** related process (e.g., creating or deleting a volume), however, auto failover is still allowed.

- **Crashed storage space:** If a storage space (e.g., volume, Disk Group, RAID Group, SSD Cache, etc.) on the active server has crashed and the corresponding storage space on the passive server is functioning normally, failover will be triggered unless there are no volumes or iSCSI LUNs (block-level) on the crashed storage space. Storage spaces are monitored every 10 seconds. Therefore, in the worst case, switchover will be triggered 10 to 15 seconds after a crash occurs.
- **Service error:** If an error occurs on a monitored service, failover will be triggered. Services that can be monitored include SMB, NFS, AFP, FTP, and iSCSI. Services are monitored every 30 seconds. Therefore, in the worst case, switchover will be triggered 30 seconds after an error occurs.
- **Power interruption:** If the active server is shut down or restarted, both power units on the active server fail, or power is lost, failover will be triggered. Power status is monitored every 15 seconds. Therefore, in the worst case, switchover will be triggered 15 seconds after power interruption occurs.
However, depending on the client's protocol behavior (e.g., SMB), the client may not be aware of the fact that data was still in the active server's cache during power interruption. If this is the case, the data that has not been flushed into the storage might not be re-sent by the client after the power interruption, resulting in partial data loss.
- **Cluster connection lost:** If an error occurs on the cluster connection and the passive server has more healthy cluster connections, failover will be triggered. For example, if the active server has two cluster connections and one of them is down, the active server will check whether the passive server has two or more available connections. If it does, failover will be triggered in 10 to 15 seconds. Note that for connections joined with **link aggregation**, each joined connection group is considered one connection.

After switchover has occurred, the faulty server may need to be replaced or repaired. If the unit is repaired, restarting the unit will bring the cluster back online and data-synchronization will automatically take place. If the unit is replaced, the cluster will need to be re-bound in order to recreate a functioning cluster. Any USB/eSATA devices attached to the active server will have to be manually attached onto the passive server once switchover is complete.

Notes:

- When a switchover occurs, all existing sessions are terminated. A graceful shutdown of the sessions is not possible, and some data loss may occur. However, re-transmission attempts should be handled at a higher level to avoid loss. Note that if the file system created on an iSCSI LUN by your application cannot handle unexpected session terminations, the application might not be able to mount the iSCSI LUN after a failover occurs.

Switchover limitations

Switchover cannot be initiated in the following situations:

- **Incomplete data replication:** When servers are initially combined to form a cluster, a period of time is required to replicate existing data from the active to passive server. Prior to the completion of this process, switchover may fail.
- **Passive server storage space crash:** Switchover may fail if a storage space (e.g., volume, Disk Group, RAID Group, SSD Cache, etc.) on the passive server is crashed.
- **Power interruption:** Switchover may fail if the passive server is shut down or restarted, if both power units on the passive server malfunction, or if power is lost for any other reason.
- **DSM update:** When installing DSM updates, all services will be stopped and then come online after DSM update installation is completed.

Best Practices for Deployment

Different customers of Synology NAS products may attach importance to various aspects of practices according to their needs and purposes. Here, we provide the best practices regarding **system performances** and **reliability** respectively. These two types of configuration are not mutually exclusive. You may apply both practices to optimize the overall system environment.

Before the configuration, make sure that your system environment complies with the basic requirements for Synology High Availability. Both servers must be of the same model, and with the identical configuration including the **capacity**, **quantity**, and **slot order** of drives and RAM modules.

Aside from the drive and RAM module configuration, both servers must have the same number of attached network interface cards and LAN ports.

System performance

To meet the needs of operating performance-intensive services, and of processing a massive amount of connections or frequent data access for a long time, you can optimize the performances with the following best practices:

Implement SSD cache

SSD cache brings a significant rise in data reading and writing speeds of the system, especially under the circumstance where the data storage consists of hard disk drives (HDD). Since solid-state drives (SSD) are specifically designed for high performance usage, by promoting frequently accessed data into SSDs, one can fully utilize the system's random I/O access to effectively reduce the latency and extra data seek time as on HDDs.

SSD cache must be configured identically on each server and each SSD must be inserted in the same disk slots on both the active server and passive server. The size of system memory must also be identical on the active and passive servers, since some of the memory needs to be allocated for operating SSD cache and different memory sizes may result in unavailability of system failover.

The failover mechanism also applies to SSD cache. This means that if SSD cache on the active server fails, a system failover to the passive server will be triggered.

Use a fast Heartbeat connection

When data is transferred to the HA cluster, a copy of the data will be transferred to the passive server through the Heartbeat connection at the same time. The writing process is only complete

once both transfers are completed. In this case, if a Gigabit network environment is applied, the writing speed will be limited to 1Gbps by the network environment.

Most plus-series and all FS/XS-series are equipped with the capability of adding additional external network cards for additional high-speed network interfaces (e.g. 10Gbps).

The most basic principle of network settings for a high availability cluster is that the Heartbeat connection bandwidth must be greater than or equal to that of the cluster network interface. The Heartbeat connection is one of the fastest network interfaces, including link aggregation or 10G/40G network interface.

Synology offers optional 10GbE external network interface cards to be used with high availability clusters. When working with multiple external network interface cards, link aggregation must be set up interchangeably to increase fault tolerance. See the next section for an example.

Set up with single-volume storage pool

When using a single-volume storage pool, the system can avoid the impact on system performance from LVM (Logical Volume Management) that generally happens with a multiple-volume storage pool. We highly recommend setting up with a single-volume storage pool for the peer-to-peer storage architecture of Synology High Availability.

Reliability

Synology High Availability is dedicated to the hardware and software protection of Synology NAS servers. However, aside from the NAS storage itself, the normal operation of the whole services also depends on a couple of other factors, including stable network connections and power supply.

Direct Heartbeat connection between two servers

When the Heartbeat connection between the active passive servers passes through a switch, it can become more challenging to manage network failures caused by the switch or its server connections.

When taking system reliability into consideration, we recommend you to make a direct Heartbeat connection between the active and passive servers.

Heartbeat and HA cluster connections with link aggregation

A link aggregation is formed by combining two or more network interfaces. This not only increases transfer rates, but enhances the availability of such network connections. When one of the interfaces or connections in a link aggregation fails, the others can still maintain the connection.

Consider the following scenario for example:

- Two RS4017xs+ models are each equipped with two E10G17-F2 external network interface cards.
- There are four 1GbE network interfaces (local network 1, 2, 3, 4) and four 10GbE network interfaces (local network 5, 6, 7, 8) on each NAS.
- Local network 5 and 6 are provided by external network interface card 1, and local network 7 and 8 are provided by external network interface card 2.

With this scenario, we recommend the following environment setup:

- **Heartbeat interface:** Link Aggregation of the 10GbE network interface. One interface from external network interface card 1 and one interface from external network interface card 2.
- **Cluster interface:** Link Aggregation of the 10GbE network interface. One interface from external network interface card 1 and one interface from external network interface card 2.

These configurations ensure the performance of both the cluster and Heartbeat connections are maximized while maintaining redundancy. The network service provided by the cluster will not be affected by the Heartbeat connection, thus increasing fault tolerance for the external network interface card. In the case that a problem occurs with external network interface card 1, all services can still be provided through external network interface card 2.

To prevent slow write speeds due to the process of replicating data to the passive server when data is written, we recommend that the Heartbeat connection have the same or higher bandwidth as the service. The Heartbeat connection must be configured on the fastest network interface. For instance, if the servers are equipped with 10GbE add-on network cards, the Heartbeat connection must be configured by using 10GbE cards.

In addition, we strongly recommend you to build a direct connection (without switches) between the two servers, the distance between which is usually shorter than 10 meters. If an HA cluster requires two servers with a larger distance, the Heartbeat connection between two servers must have no other device in the same broadcast domain. This configuration can be achieved by configuring a separate VLAN on the Ethernet switch to isolate the traffic from other network devices.

Make sure that the cluster and Heartbeat connections are in different loops, in case of sudden interruption when functioning.

Separate switches for cluster connections

A network switch is required for the HA cluster connection between the active/passive server and the client computer. To avoid network switch malfunctions, we recommend connecting each of the two servers to a different switch.

For example, you can set up link aggregation with two connections between one switch and the active server, and set up another one between another switch and the passive server. Then, the client computer will be configured with two connections, each of which is linked to one of the two switches.

Separate power supplies for servers and switches

Aside from the reliability of network connections among the servers, switches, and clients, the stable power supply is also an important consideration for the system. Most FS/XS-series and plus-series with RP are equipped with redundant power supplies, allowing you to allocate different electric power sources to the servers.

For NAS models without redundant power supplies, we recommend allocating a power supply to one server and its connected switch, and another one to the other server and switch. This helps mitigate the risk of system failure due to the power outage of both servers/switches.

Split-brain prevention with a quorum server

In an actual implementation of Synology High Availability, there are a certain number of possibilities that, even upon network abnormalities, the passive server takes over the workload while the failover from active server is not triggered. Both servers may assume the services at the same time and, in this case, the split-brain occurs. To avoid split-brain, you can configure a quorum server to detect the connections between itself and the active/passive servers respectively.

Performance Benchmark

Performance considerations

Synology High Availability employs the synchronous commit approach by acknowledging write operations after data has been written on both the active and passive servers at the cost of performance.

To enhance the random IOPS performance and reduce latency, we recommend enabling the SSD cache for volumes that require high performance for random I/O workloads. For a better understanding of the benchmark for different levels of hardware, refer to the following table containing the performance benchmark for DS920+ and SA3600.

	SHA ¹	Standalone
Model	SA3600	
Testing Version	DSM 7.0.1	
Cluster Connection Bandwidth	10GbE *1	N/A
Heartbeat Connection Bandwidth	10GbE *1	N/A
File System	Btrfs	
Disks and RAID Type	12 SSD (960 GB) with RAID 5	
SMB - 64KB Sequential Throughput - Read (MBps)	1180.3	1180.5
SMB - 64KB Sequential Throughput - Write (MBps)	1179.6	1180.4
SMB - 64KB Encrypted Sequential Throughput - Read (MBps)	1180.4	1180.4
SMB - 64KB Encrypted Sequential Throughput - Write (MBps)	1179.6	1180.1

SMB - 4KB Random IOPS - Read	295488.6	295516.4
SMB - 4KB Random IOPS - Write	140422.68	176020.7
iSCSI - 4KB Random IOPS - Read	298221.71	298286.85
iSCSI - 4KB Random IOPS - Write	148824.43	176399.07

¹ The configuration of each server in the SHA cluster.

	SHA ¹	Standalone
Model	DS920+	
Testing Version	DSM 7.0.1	
Cluster Connection Bandwidth	1GbE *1	N/A
Heartbeat Connection Bandwidth	1GbE *1	N/A
File System	Btrfs	
Disks and RAID Type	4 HDD (6 TB) with RAID 5	
SMB - 64KB Sequential Throughput - Read (MBps)	113.1	113.1
SMB - 64KB Sequential Throughput - Write (MBps)	112.4	112.9
SMB - 64KB Encrypted Sequential Throughput - Read (MBps)	113	113
SMB - 64KB Encrypted Sequential Throughput - Write (MBps)	112.1	112.9

¹ The configuration of each server in the SHA cluster.

Switchover time-to-completion

When a switchover is triggered, the active server becomes the passive server and the original passive server will take over. During the exchange, there will be a brief period where both servers are passive and services are paused.

The time-to-completion varies depending on a number of factors:

- The number and size of volumes or iSCSI LUNs (block-level)
- The number and size of files on volumes
- The allocated percentage of volumes
- The number of running packages
- The number and total loading of services on the cluster

The following table provides estimated time-to-completion:

Settings	DS920+	SA3600
Switchover	64	30
Failover	62	31

- Unit: Seconds
- Switchover is triggered manually in the Synology High Availability package.
- Failover is triggered by unplugging the power cord of the active server.

Summary

Synology is committed to delivering an exceptional user experience for our customers. Synology High Availability (SHA) comes with an intuitive user interface with built-in graphical management tools to simplify complex high availability configurations, streamline workflows, reduce implementation time and installation expenses, and prevent costly operational efficiency problems. SHA is a solution especially suitable for small and mid-range businesses with limited IT resources.

With availability on most Synology NAS models, new and existing customers can easily set up a comprehensive business continuity plan when acquiring a Synology NAS. Synology High Availability (SHA) provides a cost-effective and reliable means of minimizing service downtime while offering enterprise-grade features that are designed to meet common scenarios and fulfill data protection and high-availability requirements.

This white paper has outlined the basic principles and benefits of Synology High Availability (SHA) and provided recommendations, best practices, and strategies for companies of any size. For

more information and customized consultations for your NAS deployment, contact Synology at www.synology.com.

More resources

- [DSM Live Demo](#): Take advantage of our free 30-minute DSM 7.0 trial session. Experience our technology yourself before making your purchase!
- [Tutorials and FAQs](#): Learn more about SHA and get the information you need with step-by-step tutorials and frequently asked questions.
- [Where to Buy](#): Connect with our partners and find out how to choose the best product and solution to suit your business needs.