

Synology PAS Failover Guide

Synology

Table of Contents

Introduction	2
How it works	2
Failover and Giveback	3
Failover	3
Giveback	3
Managing Failover and Giveback	5
During a failover	5
Planning for giveback	5
During a giveback	6
Troubleshooting	7
Identify issues	7
Common issues	7
Conclusion	10

Introduction

In today's data-driven landscape, storage downtime poses a significant operational risk. To ensure continuous service availability, **Synology Parallel Active Station (PAS)** is built on a dual-controller architecture operated by **DSM Enterprise**. Its active-active design lets both controllers share workloads and provide services simultaneously, delivering balanced performance and redundancy.

Within this architecture, **Continuous Availability Manager (CAM)** serves as the central interface for monitoring and managing high-availability operations. It provides real-time visibility into system health, tracks controller interconnections, coordinates network failover groups, and offers recommendations during recovery scenarios.

How it works

PAS maintains service continuity through two key mechanisms: **failover** and **giveback**.

When one controller becomes unavailable, whether due to hardware maintenance or unexpected failure, the system automatically shifts ownership of affected resources to the partner controller. This failover process keeps services online without manual intervention.

Once the affected controller recovers, giveback is triggered, returning resources to rebalance workloads and restore full active-active performance.

Environment-specific configuration

To support seamless failover, complete the following configurations in advance based on your environment:

- **IP network environments:** [Configure failover groups](#) so that service IP addresses can be reassigned to the partner controller. This allows client connections to continue without interruption during a failover event.
- **Fibre Channel environments:** [Enable Multipath I/O \(MPIO\)](#) so the host can automatically switch to an available path if a controller or path becomes unavailable. In addition, make sure that each host has redundant connections to both controllers, and that the multipathing policy supports both failover and load balancing.

Failover and Giveback

Failover

Failover (also referred to as takeover) is the mechanism that maintains service availability by reassigning service components when a controller or part of the system becomes unavailable.

Trigger conditions

PAS automatically initiates a failover during:

- **Planned events:** Scheduled maintenance or hardware replacement
- **Unplanned incidents:** Unexpected or improper controller removal, interconnection errors, or drive connectivity issues

Failover behaviors

During the failover, the remaining controller takes over the storage and network responsibilities of its partner:

- **Storage services:** Ownership of dual pools, volumes, and other storage objects shifts to the remaining controller, ensuring continued data access.
- **Network services:** If network interfaces across controllers are grouped into a [failover group](#), IP addresses will shift to the associated interfaces on the remaining controller to keep services running.

Failover duration varies depending on system load, the number of managed resources, and the scope of services affected. The scope of a failover can differ by situation.

In some scenarios, a partial failover occurs, affecting only specific services instead of the entire controller. For example, a network issue may cause only IP addresses to fail over to the partner controller, while storage remains unchanged. Similarly, a storage issue may trigger storage failover without impacting network services.

Giveback

Giveback is the reverse of failover. When the previously unavailable controller recovers, its resources are returned, restoring the system to its optimal active-active state.

Trigger conditions

PAS automatically initiates a giveback when a controller is ready to resume service after either of these conditions:

- It is physically reinserted after planned maintenance.
- It recovers from a system error through a reboot or repair.

Giveback behaviors

When a giveback begins, ownership of storage and network services is returned to the recovered controller to restore the original configuration.

- **Storage services:** Dual pools, volumes, and other storage resources are reassigned to the recovered controller to rebalance the workloads.
- **Network services:** IP addresses return to their original interfaces on the recovered controller, so traffic follows the intended paths.

The scope of the giveback depends on the resources involved in the preceding failover. If the failover affected only specific services, the giveback is applied only to those services rather than the entire controller. In such cases, only the relevant storage or network resources are returned.

Managing Failover and Giveback

During a failover

A failover means services are temporarily running on a single controller. The goal during this period is to identify the cause and prepare for recovery.

When a failover occurs, take the following actions to ensure system stability:

Step 1: Confirm the failover state

1. Open **Continuous Availability Manager**.
2. During the transition, the overview status appears as **Processing - System transferring**.
3. When complete, the status changes to **Critical**, indicating that services are operating on one controller.

Step 2: Check alerts and system health

1. In **Continuous Availability Manager > Overview**, look for **Warning** or **Critical** notifications and follow the [troubleshooting prompts](#).
2. In **Log Center**, review recent system events to help determine the cause.
3. For suspected storage-related issues, go to **Storage Manager** to check the status and suggestions of your storage and drives.

Planning for giveback

A giveback starts automatically after the failed or removed controller is recovered and ready to use. Because service transitions may temporarily affect system performance, we recommend bringing the controller back online during low-activity periods.

- **Maintenance scenarios:** If a controller was removed for hardware service or troubleshooting, reinsert it during quieter periods.
- **Recovery scenarios:** If Continuous Availability Manager shows messages such as "Services are now running on Controller X", wait for an off-peak time before following the recommended actions (e.g., recovery or restart).

During a giveback

A giveback returns storage and network ownership to the recovered controller. The focus during this phase is to monitor the transition and verify that workloads are balanced afterward.

Step 1: Monitor the transition

1. Open **Continuous Availability Manager**.
2. During the handoff, the overview status appears as **Processing - Restarting Controller X**.
3. When complete, the status returns to **Normal**, indicating both controllers are active and resources have been redistributed.

Step 2: Confirm workload distribution

1. Open **Resource Monitor** to review CPU, network, and memory usage on each controller.
2. Confirm that resource usage is decreasing on the partner controller and activity is resuming on the recovered controller, indicating an active-active state.

Troubleshooting

If the system encounters issues, use Continuous Availability Manager as the primary diagnostic tool.

Identify issues

1. Sign in to PAS.

If you cannot sign in, the system may be unreachable or experiencing hardware problems. Refer to the following:

- [How do I find my Synology storage system via Web Assistant or Synology Assistant?](#)
- The [product manual](#) for hardware inspection and repair

2. Go to **Continuous Availability Manager > Overview**. If the status shows **Warning** or **Critical**, check the **Suggestion** section for recommended actions.

Common issues

High-availability operations may surface different types of warnings and errors. This section outlines common causes and their troubleshooting methods.

Important: If Continuous Availability Manager instructs you to contact [Synology Technical Support](#), stop any manual recovery steps. Continuing without guidance may lead to data loss or system instability.

Controller-level issues

Issues affecting overall system health or controller interconnection. These may either trigger a controller failover or make it unsafe for the system to proceed with failover.

Error Message	Cause	Recommended Action
• Interconnection issues detected across controllers • Interconnection error	The physical link (i.e., interconnection cable) between controllers may be loose or not seated correctly.	Make sure the interconnection cable is fully inserted.

• Unable to detect Controller X • Controller X is offline	The controller cannot establish physical communication with the partner controller, or the controller did not power on correctly.	Check if the controller is firmly seated and powered on.
• Unable to detect Controller X's system	The operating system on the controller cannot be found.	Follow the on-screen instructions in Continuous Availability Manager to verify or repair the system.
• Inconsistent OS versions across controllers	One controller is running a different OS version, which prevents synchronized operation.	Schedule a maintenance window to update the controller running the older version.
• Inconsistent configurations at risk of loss • Inconsistent configurations on the controllers	Settings between both controllers are out of sync and require reconciliation.	Follow the on-screen instructions in Continuous Availability Manager to resynchronize the settings. Avoid non-essential operations until synchronization is complete.
• Unable to retrieve system information from Controller X • Unknown errors occurred	The system has encountered an unidentified error.	Stop all operations and contact Synology Technical Support immediately.

Storage & network issues

Issues affecting individual storage or network resources (e.g., specific dual pools, volumes, or network interfaces). These may either trigger a partial failover for those resources or make it unsafe for the system to proceed with failover.

Error Message	Cause	Recommended Action
Unable to perform a failover for storage	A system protection mechanism has been activated, temporarily suspending failover.	Follow the Suggestion in Continuous Availability Manager to reboot the system. After the restart, check whether the status returns to Normal . If the issue persists, continue troubleshooting based on the reported errors.
- Severe issues occurred in the storage - Issues occurred in the storage - Storage errors on both controllers - Storage errors on Controller X	One or more storage components are reporting errors that affect availability or data access.	Follow the on-screen instructions in Storage Manager to troubleshoot.
- System drive (Controller X) failed - System drives on both controllers failed	The system drive hosting DSM Enterprise is not functioning properly.	Follow the on-screen instructions in Continuous Availability Manager to troubleshoot.
- Failover group error - Critical failover group error	One or more LAN ports in the failover group are unavailable or misconfigured, blocking IP migration.	Check if the IP addresses and LAN ports are correctly configured in Continuous Availability Manager > Failover Group .

Conclusion

Synology Parallel Active Station (PAS) works with **DSM Enterprise** to deliver a storage system designed for continuous service, even during maintenance or unexpected failures. Automated failover and giveback keep services running, while **Continuous Availability Manager (CAM)** provides visibility for monitoring health and maintaining balanced resource allocation.

By understanding how these mechanisms operate and following the recommended procedures, you can safeguard service availability and ensure business continuity.