

## **ASA**

### **Regular ONTAP (AFF/FAS)**

When the owning controller fails, the aggregate containing the LUN must give over to the partner via storage failover (SFO).

During this process:

- The active/optimized paths disappear

- The host detects an APD (All Paths Down) event momentarily

- The partner takes over the aggregate and re-advertises the paths as active/optimized

The host ALUA logic must discover and switch to the new paths

This causes a noticeable I/O pause — typically seconds to tens of seconds depending on host multipath settings and ALUA transition timers.

### **ASA (original)**

The LUN still has an owning controller, but because the partner is already serving active/optimized paths, when the owning controller fails:

- The host never sees an APD event — the partner's paths were already in use

- No ALUA transition is needed

- I/O continues on the surviving controller's paths with minimal interruption

The timeout/pause is essentially eliminated because from the host's perspective, nothing changed — it was already load balancing across both controllers.

### **ASA r2**

Same symmetric behavior, but with the shared storage availability zone both nodes already have direct access to all disks, so the failover is even cleaner — there is no aggregate takeover process in the traditional sense.

So the key insight is: on a regular ONTAP cluster the host reacts to the failure.

On an ASA the host doesn't need to because it was never dependent on a single path being optimal.