SDC 20

BY Developers FOR Developers

Storage Developer Conference
September 22-23, 2020

Why do customers need media path redundancy in a storage array with simple low-latency hardware paths?

Presenter: Dr. M. K. Jibbe
Technical Director , NetApp ESG
Presenter: Joey Parnell
SW Architect, NetApp ESG

Current Problem

- **Abstract:** Many emerging white box storage systems replace mature, robust drive networks with simple point-to-point drive port topologies

The presentation provides fault tolerance where each storage controller has a dedicated port to a drive, but no hardware path to other port(s) to that drive. The redundant access to media is provided via Non-Transparent Bridging (NTB), Ethernet, or other networks to overcome lack of redundancy in the topology.

Current Problem Cont'd

Background: As commodity white-box storage systems with less custom hardware become prevalent, some parts of the hardware design that provide the fault tolerance previously relied upon to provide high availability in hardware platforms are no longer present. Such hardware designs will enable software-defined storage (SDS) products and will have lower hardware cost and complexity. Additionally, input/output (I/O) access is increasingly expected to be predictable and have extremely low latency, motivating the use of simpler point-to-point peripheral component interconnect express (PCIe)/non-volatile memory express (NVMe) connections in these hardware designs. However, this sacrifices the redundancy provided by more robust traditional storage networks (i.e., cross-connected packet networks or Fabrics) that in the past connected physical media to the storage system.

In a storage system, lack of redundancy in a storage controller's hardware topology to physical media (i.e., HDD, SSD, NVMe, Flash) introduces a single point of failure that is unacceptable to enterprise storage customers.

Goal of the Team (Implements an I/O policy management module to reduce / eliminate non-optimal I/O accesses over time)

1) Provide fault tolerance,
2) Minimize disruptions when faults occur,
3) Provide automatic drive recovery mechanism

Fabric Port Fail vs. PT to PT Port Fail

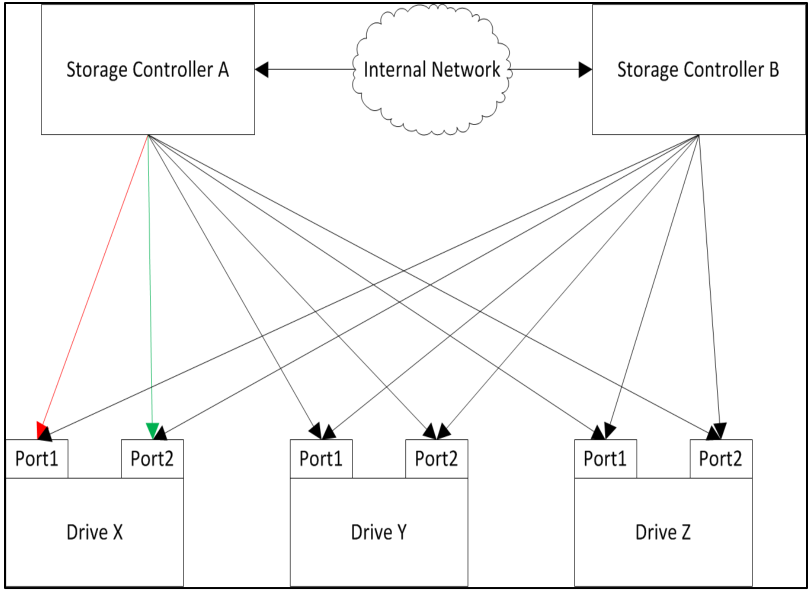

Figure 1: Fabric Port Fail

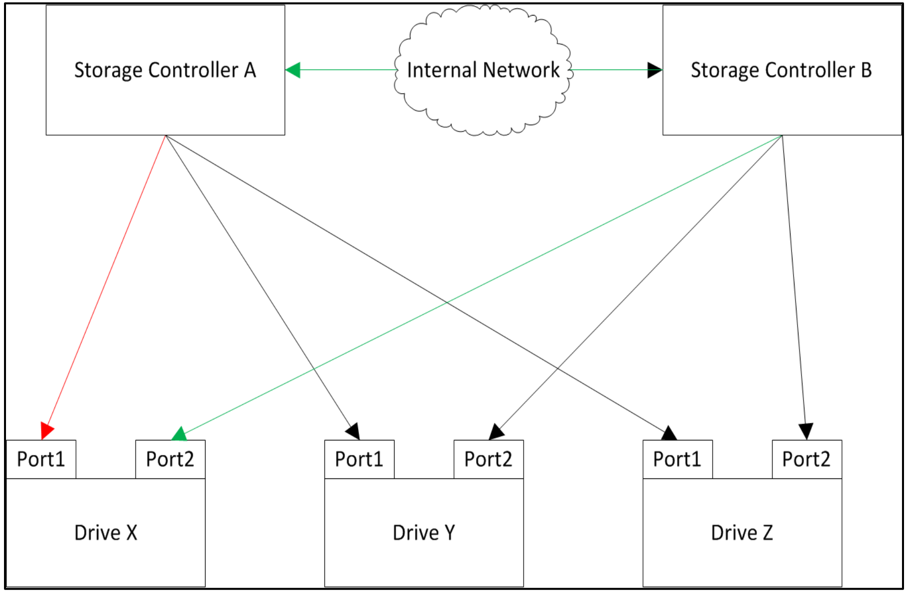

Figure 2: Hardware Path Point-to-Point Port Fail

Solution/Novel Contribution

- The solution described herein provides fault tolerance where each storage controller has a dedicated port to a drive, but no hardware path to other port(s) to that drive.
- The redundant access to media is provided via Non-Transparent Bridging (NTB), Ethernet, or other networks to overcome lack of redundancy in the topology.
- Providing a non-optimal I/O path to physical media introduces a new class of problems because of the potential impact on host applications.
- Any asymmetry in access (i.e., non-optimal accesses) could negatively impact host applications. This problem is mitigated but not eliminated by storage controller caching.

Method/Process

- In the solution described here, the storage controller implements an I/O policy management module to reduce or eliminate such non-optimal I/O accesses over time.

- The redundant storage controllers implement several modules to solve the problems described:
  - I/O Shipping Module to execute I/O operations
  - Multipathing Module to select the path to media
  - I/O Policy Management Module to reduce or eliminate usage of non-optimal path(s)
  - Serviceability Module to aid the user with notifications and actions to take for these features.

- Refer to Figures 3 and 4 which illustrate the paths and functions of the modules.

Paths and Functions of the modules

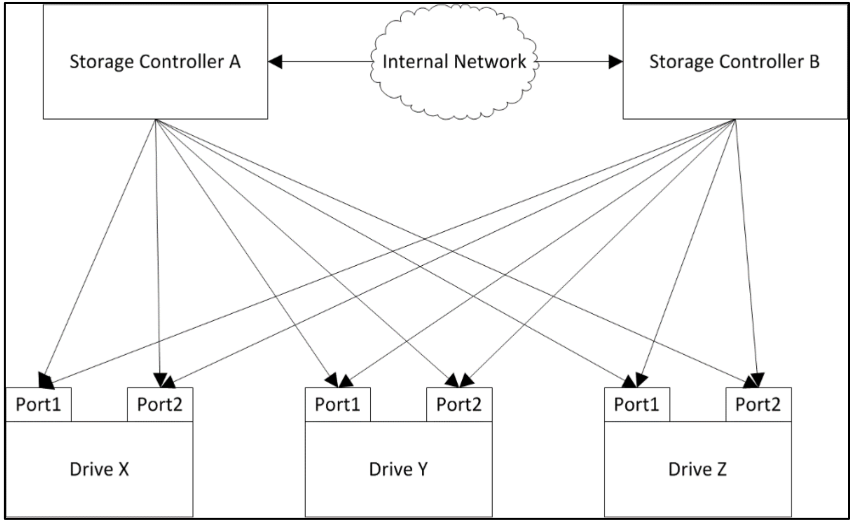

Figure 3: Hardware Path Fabric Port

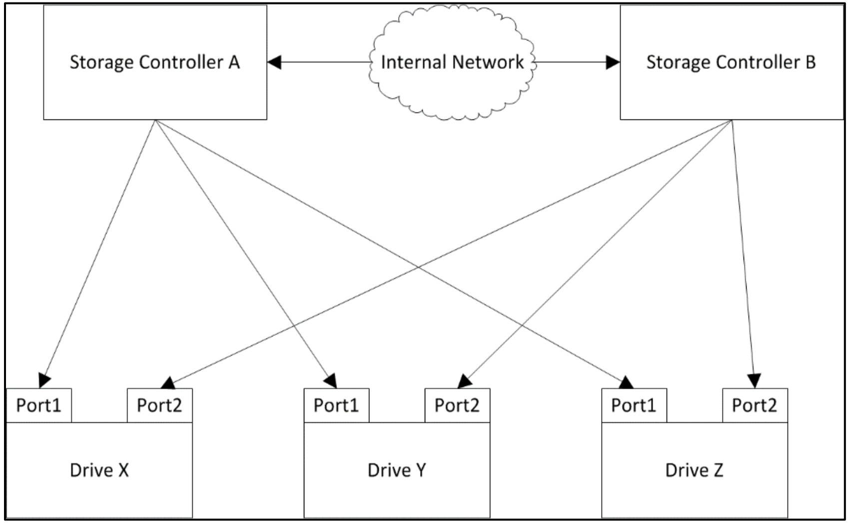

Figure 4: Hardware Path Point-to-Point Port

I/O Shipping Module

- The storage controller implements an I/O Shipping Module.
  - The I/O Shipping Module uses a combination of messages and data transfers to execute I/O operations.
  - Messages for data phases are transmitted between storage controllers via internal network (e.g., Ethernet) or mirrored via PCIe Non-Transparent Bridging (NTB) to known cache addresses
  - I/O data is mirrored between storage controllers via NTB to known cache addresses
- Errors are handled between storage controllers. If any errors are encountered such that an I/O operation cannot complete, then:
  - If storage controllers have inter-connectivity, then those errors are propagated to the Multipathing Module on the originating storage controller of the I/O operation via internal network or NTB
  - If storage controllers do not have inter-connectivity, then the originating storage controller of the I/O operation will observe an I/O timeout

Multipathing Module

- The storage controller implements a Multipathing Multipathing Module. The Multipathing Module selects the path to media.
- The Multipathing Module is notified of media discovery
  - Storage controllers discover drive accessibility changes during Start-of-Day processing and/or when drives are inserted or removed.
  - These drive accessibility changes are communicated between storage controllers via internal network or NTB
- The Multipathing Module maintains a collection of paths to media along with an indication of the expected access latency of each path. Paths that are added from other storage controllers are denoted as non-optimal.

Multipathing Module Cont'd

- The Multipathing Module selects the path to use for each I/O operation to media
- The storage controller selects the optimal path to media for I/O operations under normal conditions
- When a port failure is observed:
  - A non-optimal path is selected based on expected I/O access characteristics
  - The I/O Shipping Module is executed to perform the I/O operation
  - The I/O timeout may be increased, and/or the I/O retry count may be decreased based on the expected latency of the non-optimal path
  - Any errors on the I/O operation are handled and may be retried on the same path or on another path based on the path selection algorithm.

I/O Policy Management Module

- The storage controller implements an I/O Policy Management Module.
- The I/O Policy Management module reduces or eliminates usage of the non-optimal path to media.
- Redundant storage controllers may have asymmetric characteristics for host I/O operations.
  - The storage controller initiates volume ownership transfers to reduce I/O operations going to the non-optimal path to media.
- The storage controller maintains a mapping of physical media to logical blocks on a redundant array of independent discs (RAID) volume.
- Using this mapping, the I/O Policy Management Module can identify volumes that are reporting an optimal path to the host via a storage controller (i.e. providing beyond some threshold of its media to the non-optimal path to the drive.)

I/O Policy Management Module cont'd

As part of the volume ownership transfer, the storage controller marks the volume as ineligible for non-critical volume transfers to disallow automatic load balancing actions from thrashing ownership against these transfers. The storage controller:

1) Initiates drive power cycle operations, reducing disruptions where drive problems can be corrected with a power cycle

A. The I/O Policy Management Module manages the list of drives that are power-cycled, avoiding power-cycling too many drives at once and maintaining data availability and integrity

B. Drive power cycle may initiate a rapid reconstruction of data, which is dependent on eligibility based on the RAID level of the volume group to which the drive belongs. The drive power cycle:
   - treats I/O as degraded
   - intercepts and logs write I/O operations to memory
   - performs reconstruction when the drive becomes optimal

2) Initiates drive failure operations. If the above recovery actions do not address the problem after some duration, then the storage controller may fail the drive so they can provide optimal I/O performance and expected low latency to host applications. The duration of time that can pass before taking this action may depend on the redundancy level of the volume group to which the drive belongs and the number of drives in that same group with port failures.

3) (Optional) Initiates drive evacuation for drives that have failed

A. Migrates data to unused or spare capacity on either a hot spare or available disk pool space

B. Automatic drive data migration may be enabled or disabled by user configuration

The storage controller provides an optional user configuration to allow the customer to specify whether optimal I/O performance is valued above data integrity for their use case. If data integrity is preferred, then the non-optimal path to media may be used in perpetuity. Some or all of the options listed above might not automatically occur via the storage controller but, optionally, may be manually initiated by the user.

Serviceability Module

For the Serviceability Module, the storage controller performs the following:

1. Reports drive port failure conditions to the customer via the management interface. Methods can include, but are not limited to:
   - A. Turning on/off a drive fault LED to indicate a failure or problem with the drive, etc.
   - B. Turning off a drive fault LED if a port has failed to indicate a problem with the drive
   - C. Turning off a controller fault LED for another storage controller that is servicing I/O operations via the non-optimal I/O path

2. Implements an isolation diagnostics method to distinguish between port failure(s) on some number of media and controller PCI bus failures to take the corrective action that matches the conditions and to report the cause to the user. Storage controller actions can include, but are not limited to:
   - A. Implementing certain proactive operations while in drive port(s) failure conditions to protect the storage system configuration
   - B. Disallowing online firmware upgrades (e.g., controller firmware, drive firmware, etc.) if a drive port has failed
   - C. Migrating internal storage system metadata to drives without a port failure.

- In an alternative embodiment, the method can attach the drives via NVMe over Fabrics to the redundant storage controllers in a storage system. This would result in significantly higher (worse) latency, but the implementation would not need to include the non-optimal path handling described in this disclosure.
- Another approach would be to extend the addressable memory map between the redundant storage controllers and submit I/O operations to the NVMe submission queue

What is the solution?

1. Provide redundancy for point-to-point connections - the resiliency of a mature drive network topology with inexpensive white box hardware
2. Solve the problems that come from providing the redundancy in software such as:
   a) Redistributing workload to minimize the performance penalty
   b) Automatically recovering misbehaving drives
   c) Restoring the customer workload to optimal configuration afterward
   d) Notifying the customer about faulty drive components

Key Takeaways and Closing Thoughts

- SAN customers increasingly expect enterprise quality at low cost, which is being driven down by hardware convergence/commoditization
  - Some enterprise features are lacking from converged HW
- We need to provide the robustness of a drive network topology in white box server hardware
- We can provide drive path redundancy with software and utilize ethernet and NTB, but it comes with its own set of problems to solve
  - These problems can be solved by redistributing workload to drives via asymmetric logical volumes and automatically recovering misbehaving drives

Questions

Thank You