STORAGE DEVELOPER CONFERENCE

SNIA SANTA CLARA, 2017

A large cluster architecture with Efficient Caching Coherency, Intelligent Management, and High Performance for a Low-Cost storage node

Dr. M. K. Jibbe,

Technical Director, NetApp

Dean Lang

DEV Architect, NetApp

Paper Objective, Method, and Results

- Objective
  - **Using a data cache coherency model that employs the concept of logical unit ownership within a cluster of storage nodes allows for optimization of performance for ultra-low latency, even on low-cost storage hardware lacking a high-speed interconnect between nodes**.
- Method:
  - Accomplished the objective by limiting optimal I/O access to any one logical unit to a single storage node with a local high-speed data cache. However, this also implies use of *Asymmetric Logical Unit Access* (ALUA).
- Observations:
  - Providing resilient / fault-tolerant access to shared storage at the logical unit level to initiators in a multi-node compute cluster is difficult when using ALUA.
  - Balancing I/O workload is more complex than with a symmetric active/active model where I/O is routed equally among all storage nodes in the cluster.
  - Handling SAN connectivity faults requiring rerouting of I/O is also complex and the associated ALUA state changes may imply I/O response time spikes.
  - Thrashing of ALUA states may occur on any one logical unit if multiple compute nodes in the cluster disagree as to the optimal paths to access that logical unit.
- Results:
  - Have demonstrated an architecture for a low-cost storage node that uses implicit ALUA and intelligent management of logical unit ownership to implement a highly efficient data cache coherency model and lean I/O path while solving many of the issues plaguing ALUA solutions.
  - Achieve ultra-low I/O latency as well as fault tolerance and stability within the compute and storage cluster at a low price point. This design also provides for intelligent I/O workload balancing across storage nodes in the cluster, optimized ALUA state changes and intelligent detection and reporting of SAN connectivity issues.

2017 Storage Developer Conference. © NetApp, Inc. All Rights Reserved.

Take away

- How to build a large, stable cluster using shared storage in a heterogeneous host OS environment
- How to establish Efficient Caching Coherency in such large cluster (Knobs / Tuning parameters)
- How to sustain an expected high performance in an optimal condition
- How to sustain an expected high performance during and after logical LUNs transition during failover and failback between redundant controllers
- How to establish a low Latency and lean I/O path using ALUA while solving many of the issues plaguing ALUA solutions with fault tolerance and stability

2017 Storage Developer Conference. © NetApp, Inc. All Rights Reserved.

Typical Cluster Environment

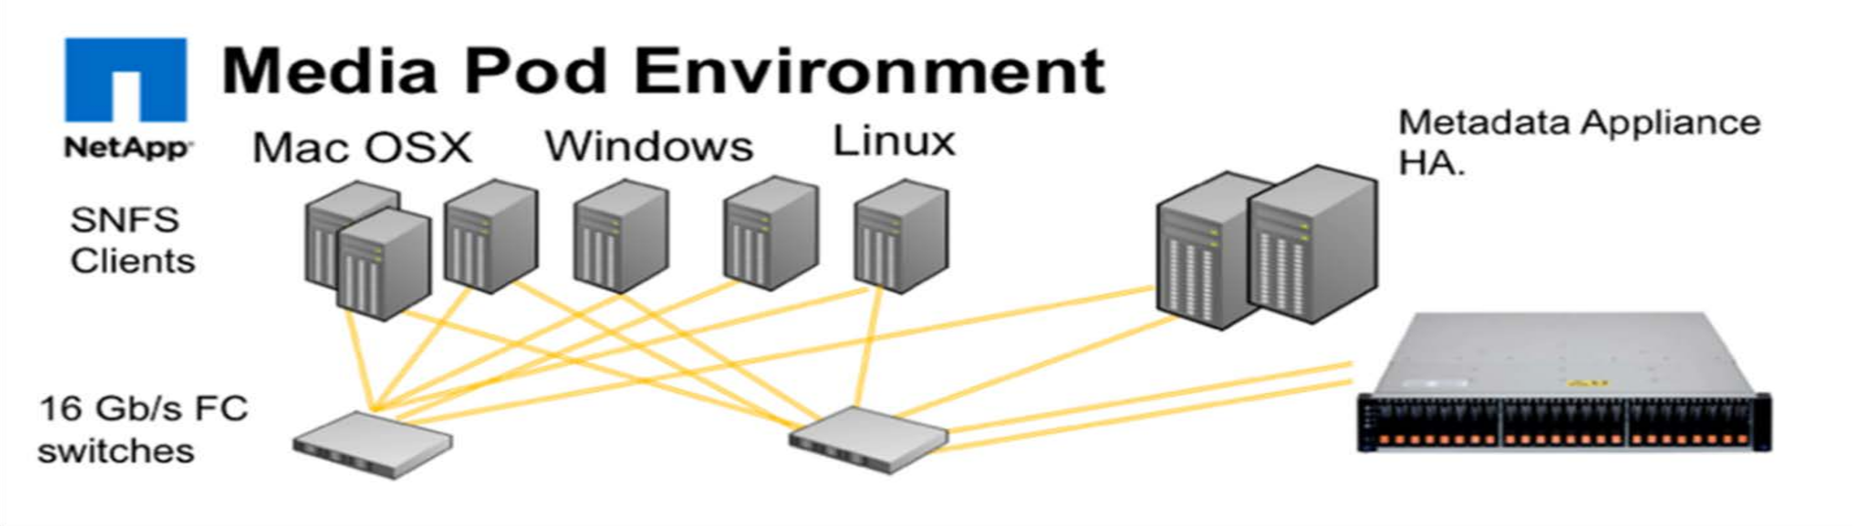

Test Configuration parameters:
Hardware

- 2 x Dell R720 (Linux and Window), and 1 MAC Pro
- Linux CENTOS 6.5 with ATTO FC162E and Windows 2012 R2 with ATTO FC84E, and ATTO FC162 with MAC Pro
- 1 x Brocade E6510 switch and 1 CISO Swicth populated with 16Gb SFPs
- E5624 with 2 8+2 R6 14 TB volumes, and 8 volumes DDP
- 24 NL SAS 7200 HDDs
- RYO Stornext MDCs with 8GB of RAM

Test Configuration parameters:
Software

- Stornext 5.2.0 installed on Server and all hosts
- E-series NVSRAM updated (Implicit Failover)
- I/O tools: Smash and Iometer
- Host HBAs running ATTO driver 1.55MP with config tool 4.25
- Brocade: Fabric OS: v8.0.01, and Cisco: version 6.2(13)

2017 Storage Developer Conference. © NetApp, Inc. All Rights Reserved.

Multi-Node Storage Clusters and ALUA

- Limiting I/O processing to a single storage node for a given logical unit allows for local caching of all data and elimination of the latency implied by multi-storage-node cache state synchronization (Traditional “logical unit ownership” model).
  - This allows for ultra-low-latency on READ operations by eliminating the need for any inter-storage-node communication to process the request.
  - For fault tolerance, an I/O request must be able to be processed by any storage node in the cluster even if that incurs a performance penalty (Asymmetric Logical Unit Access or ALUA).
- Using T10/SPC standards for reporting asymmetric path access to redundant storage nodes, initiator multipath drivers are able to route IO to the storage node(s) capable of most efficiently processing I/O
  - I/O requests are only routed to the non-optimal storage nodes when a SAN connectivity fault occurs.
  - Storage nodes must recognize that I/O is being directed to a non-optimal (non-owning) storage node and adjust ownership to eliminate performance penalty (ALUA state transition).
- Two types of management for ALUA: Implicit and Explicit
  - Explicit ALUA: Initiators may manage ALUA states explicitly by requesting specific storage nodes to be the optimal node for handling I/O to a given logical unit (i.e. initiators may make requests to change which storage node “owns” a logical unit)
  - Implicit ALUA: Initiators only route I/O to the storage nodes that are reported as providing optimal performance. This allows the storage nodes to decide which node is the best owner for a given logical unit.

2017 Storage Developer Conference. © NetApp, Inc. All Rights Reserved.

Multi-Node Compute Clusters and ALUA

- Multi-node compute clusters using shared / clustered storage are complex and difficult to manage with ALUA
  - This is especially true when the compute nodes are using different host operating systems and storage driver stacks
- Challenges:
  - **Getting multipath drivers on compute nodes in the cluster to cooperate with each other and behave compatibly** when ALUA state transitions occur can be difficult when those compute nodes are running different operating systems and drivers but sharing access to the same storage.
    - This is a problem regardless of whether implicit or explicit ALUA is being used but is worse with explicit ALUA
  - **Explicit ALUA is susceptible to logical unit ownership / ALUA state "thrash"** if the initiator multipath drivers on the compute nodes / do not cooperate with each other and independently initiate explicit ALUA state change requests.
    - This is often triggered when one compute node in the cluster loses access to one or more storage nodes in the SAN while the remaining compute nodes retain access.
- Solutions:
  - **Our storage architecture uses configuration information provided by the system administrator to understand host operating system and multipath solution on each initiator / compute node** – this allows us to adjust storage node behavior / reporting to help trigger compatible behavior in the initiator multipath driver(s).
  - **Using implicit ALUA where the storage nodes exclusively drive the logical unit ownership assignments** allows our architecture to avoid ownership thrash since the storage node has visibility to all initiators in the cluster and all SAN connectivity within the cluster.
    - This allows the storage node to select the best possible owner of a given logical unit since it is aware of all compute nodes generating I/O to that logical unit and which storage nodes each of those compute nodes can access.
    - This takes advantage of the behavior of the initiator multipath drivers, using T10/SPC standard reporting methods to get the initiators to direct I/O to the owning storage node for a given logical unit even when the storage nodes reassign that ownership.

2017 Storage Developer Conference. © NetApp, Inc. All Rights Reserved.

Automatic Load Balancing

- Challenge: With an ALUA architecture, I/O workload for a given logical unit is not automatically distributed evenly among all storage nodes in a cluster as is the case with a symmetric active/active architecture.
  - This makes it difficult to balance workload and make efficient use of storage node resources.
- Solution: Our architecture uses an intelligent workload analysis method to identify the most optimal storage node to own each logical unit to achieve a balance of aggregate workload across all logical units on all storage nodes.
  - Ownership is periodically rebalanced to handle the dynamic nature of I/O workloads over time.
  - This rebalance results in an implicit ALUA state change – we utilize standard methods for reporting such changes to the initiator multipath drivers so they will "follow" our change and redirect I/O for a given logical unit to the new owning storage node to effectively rebalance workload.
  - We use the connectivity tracking data we maintain to understand when initiators can and can't redirect I/O traffic based on current SAN connectivity, allowing the load balancing process to proactively avoid creating issues when faults occur.
  - In cases of partial SAN connectivity faults, this connectivity data can also be used to retain workload balance even if workload distribution from some initiators is fixed due to lack of connectivity to all storage nodes. Workload can be redistributed to storage nodes from initiators / compute nodes that have retained full SAN connectivity.

2017 Storage Developer Conference. © NetApp, Inc. All Rights Reserved.

Connectivity Tracking

- In order to manage logical unit ownership assignments exclusively from the storage node(s), our architecture has SAN connectivity tracking to understand which initiators have access to each storage node at any one given time.
  - This combines protocol-level connectivity information with an understanding of initiator multipath driver discovery of each SAN connection path by monitoring of incoming commands on each path from each initiator.
  - For example: Receipt of a REPORT TARGET PORT GROUPS request from an initiator multipath driver over a given path is a good indicator that the driver has discovered the presence of that path and is evaluating path state/usability
- We use the connectivity tracking data we maintain to understand when initiators can and can't redirect I/O traffic based on current SAN connectivity, allowing the load balancing process to proactively avoid creating issues when faults occur.
  - In cases of partial SAN connectivity faults, this connectivity data can also be used to retain workload balance even if workload distribution from some initiators is fixed due to lack of connectivity to all storage nodes. Workload can be redistributed to storage nodes from initiators / compute nodes that have retained full SAN connectivity.
- As a side-benefit, maintaining this level of connectivity awareness within the storage nodes also allows our architecture to have a rich reporting mechanism to alert the storage administrator to SAN connectivity faults to expedite recovery
  - Reports condition where a initiator appears to be connected at the protocol-layer to one storage node but not the others.
  - Reports condition where protocol-level connectivity exists at the physical layer but for some reason the initiator multipath driver doesn't seem to be discovering all paths
    - This typically indicates a host multipath configuration issue

2017 Storage Developer Conference. © NetApp, Inc. All Rights Reserved.

I/O Latency Spikes / ALUA State Changes

- Challenge: Changes in logical unit ownership / ALUA state transitions can result in I/O response time spikes if the ownership change process is not optimally designed.
- Solution: Our architecture utilizes an optimized ownership transfer process to minimize this impact and in many cases eliminate or greatly reduce that response time spike.
  - Logical unit ownership is transferred in batches grouped by logical unit performance characteristics such as media type and drive-group distribution. This allows for stop time during the ownership transfer process to be commensurate with the ongoing average response time for that group of logical units.
  - Unwritten data in cache is flushed to media while continuing to process incoming I/O requests.
  - Initiator multipath drivers are immediately notified of the state transition using a UNIT ATTENTION, minimizing the time taken by the multipath driver to redirect I/O to the new owning storage node.
- Solution: When logical unit ownership is transitioned between storage nodes for non-critical purposes such as workload rebalancing, the ownership distribution is designed to minimize the number of transitions necessary to achieve the workload balance.
  - A Return on Investment (ROI) computation is also performed to determine if the potential temporary performance impact to rebalance workload is justified by the improvement in workload balance.
  - Workload rebalancing is avoided altogether if the total workload on the storage nodes is insufficient to stress resources such that a rebalance would be beneficial

2017 Storage Developer Conference. © NetApp, Inc. All Rights Reserved.

Automatic Recovery / Fail-Back

- When a SAN connectivity fault is resolved, logical unit ownership may need to be redistributed if previously limited by the connectivity issues (fail-back).
  - This fail-back process takes advantage of the connectivity tracking data in our architecture to automatically recover and redistribute workload for optimal performance following recovery from a fault.
  - This also allows us to not fail-back if conditions occur such that a partial recovery occurs – this avoids limiting bandwidth.
- By having the storage nodes drive logical unit ownership changes for fail-back in addition to changes for load-balancing and fault tolerance (fail-over), all logical unit ownership can be driven exclusively by the storage nodes.
  - **This means our architecture uses implicit ALUA exclusively** – in no case are any of the initiators / compute nodes in the cluster making logical unit ownership decisions
  - **This avoids cluster ownership "thrash"** and creates a stable cluster environment that is tolerant of SAN connectivity faults yet retains the ability to balance aggregate workload on the storage nodes.
- Managing logical unit ownership entirely from the storage nodes in an automatic/intelligent fashion eliminates the need for the storage administrator to understand / manage this ownership
- This greatly simplifies storage administration while retaining ALUA and logical unit ownership in our architecture to support local data caching and ultra-low I/O latency with cost-effective hardware

2017 Storage Developer Conference. © NetApp, Inc. All Rights Reserved.

Reference Solution

- We applied these techniques to build a stable multi-node compute cluster using shared storage
- Example application used a high-bandwidth media workload
- Tests were performed to both show I/O throughput of the cluster and resiliency during a SAN fault and recovery triggering fail-over and fail-back
  - Specific tests demonstrated reduction in I/O latency during ALUA state transitions using implicit ALUA with state transition optimizations

2017 Storage Developer Conference. © NetApp, Inc. All Rights Reserved.

Test configuration

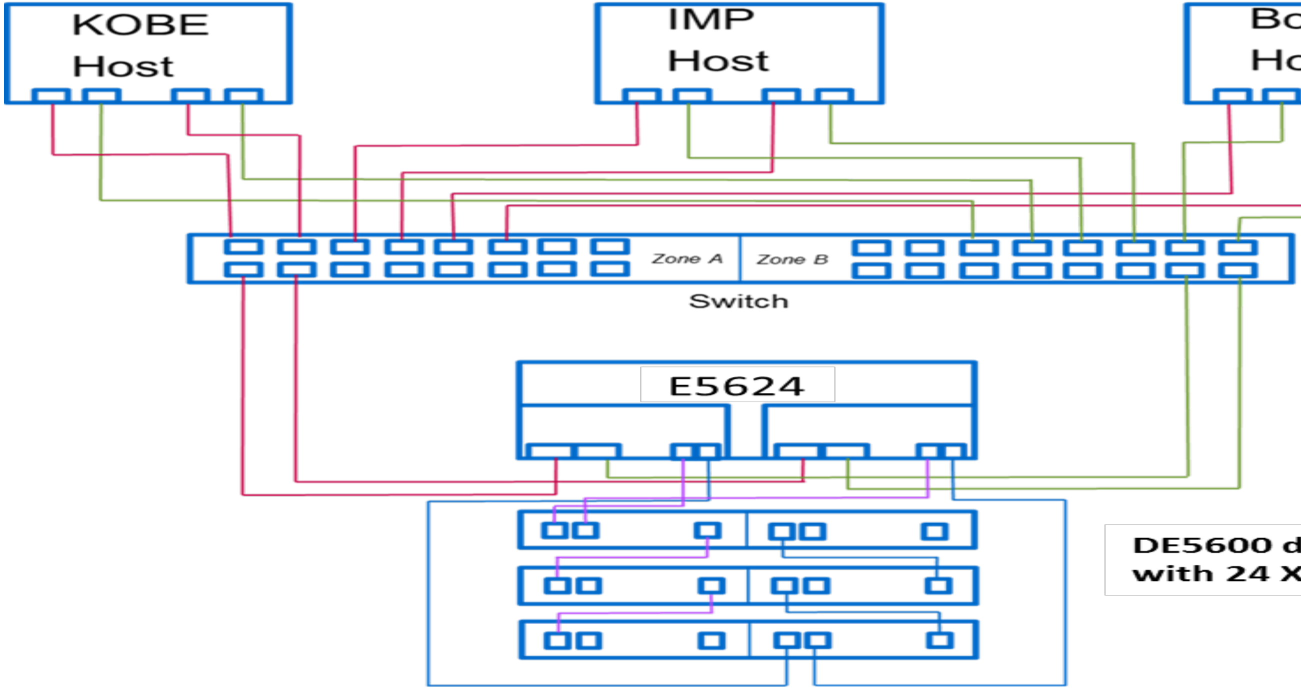

DE5600 drive shelves each with 24 X 10K SAS HDDs

2017 Storage Developer Conference. © NetApp, Inc. All Rights Reserved.

Brief Description of Test Configurations

- Established IP addresses for both controllers in the RBOD and specified the storage array name (fn1014fp-flame-5624_03 was used)
- Connect the Fibre Channel (FC) HIC of the controllers in the RBOD to the two FC switches using fiber cables. The switches were connected to three hosts, namely IMP, Bootleg, and Kobe, via FC cables. Each host had two FC ATTO HBAs installed, and the hosts had the following characteristics:
  - HP z840 40cpu 3.1ghz
  - 64G Ram
  - M6000 12G (USED K6000) server video card
- Installed Santricity Storage Manager (Version: 11.25.0G00.0024) on a local management server,
- Used Santricity Storage Manager to download and upgrade to the following software package on the storage array system: (CFW: 08.25.06.00, NVSRAM Version:N5600-825834-D03 , ESMs: 039E
- Defined the zoning on Switch A and Switch B: The zones are defined such that any HBA in each host can access one target port from one array controller from each switch
- Defined host groups using the Array Management Window
- Assign volumes to the host group (cluster)
- Configure the StorNext MDCs HA cluster by following the StorNext installation guide

2017 Storage Developer Conference. © NetApp, Inc. All Rights Reserved.

Test suite and Results

- **Ran Self-certification with Autodesk Test Suites**
  - Play All frames Playback
  - Real Time Playback
  - Rendering Project
  - Test results: Exceeds expectation
- **Ran Performance Stress Testing of a StorNext Cluster with Video Applications from Autodesk**
- **Test specifications:**
  - One client running 4 streams of unthrottled 4K
  - One client running 2 streams of HD
  - One client running 1 stream of unthrottled HD
  - Pixel/Depth of 8 bits.
- **Test Results:**
  - E-Series:
  - Avg. IOPs of 12k
  - Avg. throughput of 6.2 Gbps (max: 3 clients reading 2000 frames 4K → Max Throughput: 9.1 Gbps)

2017 Storage Developer Conference. © NetApp, Inc. All Rights Reserved.

(8 hosts accessing 8 volumes) Current I/O Latency in the Cluster measured by Management Software: Failover and Failback with explicit ALUA vs Implicit

- failover and failback with explicit ALUA (Latency=2.6 sec)

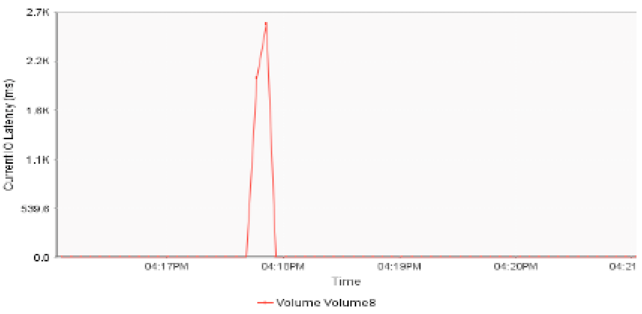

- With Implicit Failover in the Cluster (Latency = 0.669)

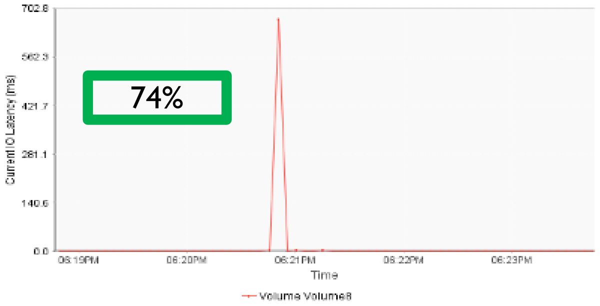

- failover and failback with explicit ALUA (Latency=3.6sec)

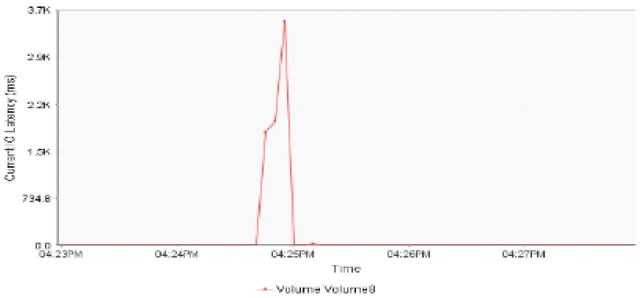

- With Implicit Failover in the Cluster (Latency = 0.753)

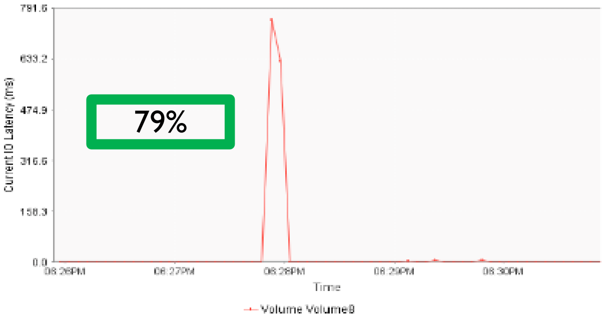

2017 Storage Developer Conference. © NetApp, Inc. All Rights Reserved.

Example Reporting: Host Redundancy Lost

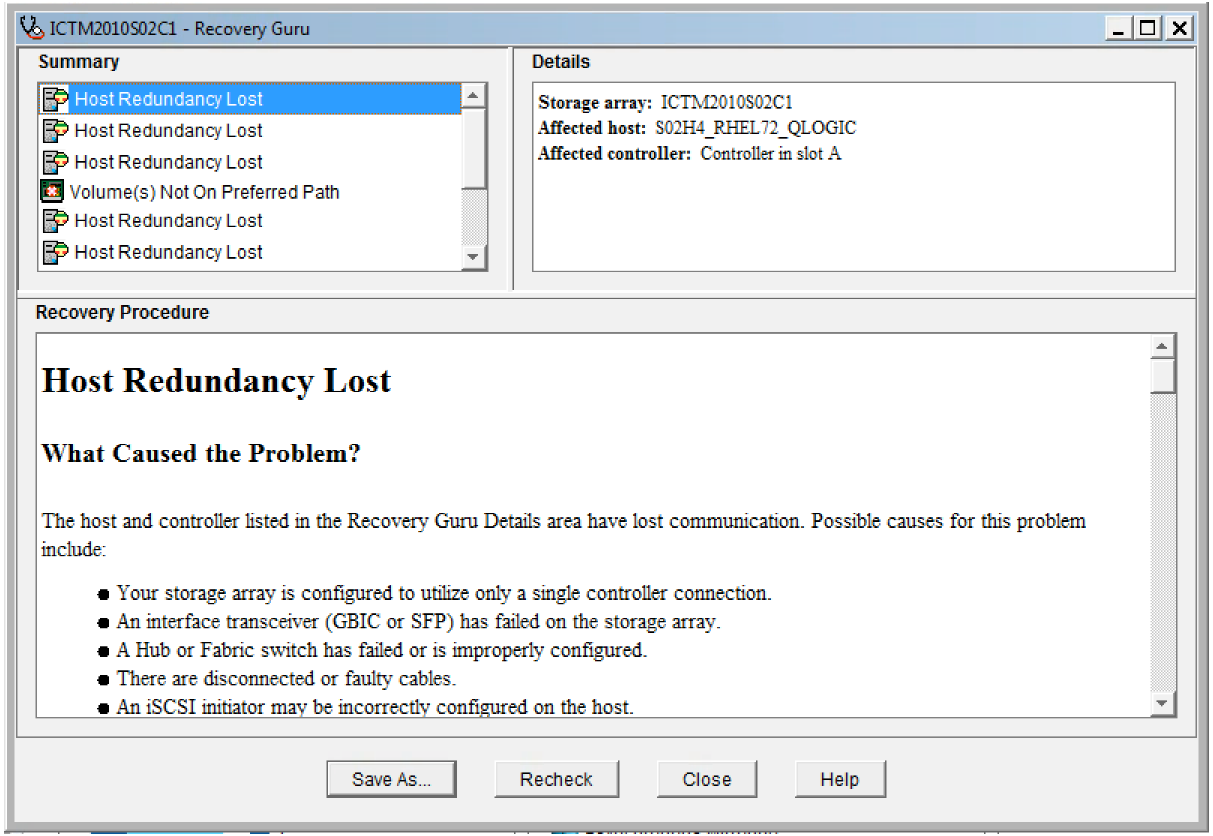

2017 Storage Developer Conference. © NetApp, Inc. All Rights Reserved.

Challenges and Solutions Summary

- Challenge: Logical unit ownership thrashing in a multi-node heterogenous operating system compute node cluster
  - Solution: Move to implicit ALUA with intelligent distribution of logical unit ownership among the storage nodes
- Challenge: I/O response time spikes during logical unit ownership / ALUA state transitions:
  - Solution: Optimize ownership transfer process to intelligently batch transitions and minimize I/O stop time during transition.
  - Demonstrated 79% improvement in I/O latency during fail-back and 74% during fail-over as compared to non-optimized explicit ALUA solution
- Challenge: Balancing workload across storage nodes is difficult with ALUA
- Challenge: Manual management of logical unit ownership by end-user to achieve workload balance and maintain stability is complex and burdensome
  - Solution: Track SAN connectivity intelligently to provide auto fail-over and fail-back capability
  - Solution: Use connectivity data and workload analysis to automatically assign logical units to storage nodes.
  - Solution: Use connectivity data to provide enhanced SAN reporting to the end-user

2017 Storage Developer Conference. © NetApp, Inc. All Rights Reserved.

Conclusions

- NetApp engineering continues to improve host/array interoperability and compatibility with host multipath solutions to better support clustered host configurations with shared access to the same set of storage volumes
    - Minimizing risk of volume ownership thrash while retaining failover/redundancy capabilities
- Testing of media-specific configurations has demonstrated real-world performance aspects of solutions using these new capabilities
    - Firmware releases passed the Autodesk self-certification suites with much better than expected stream results
        - The Quantum-Autodesk Stress tests mimicking the FLAME application also delivered superior performance throughput (Sustained 9.1 Gbps end to end) without any frame drop
- E-Series design also continues to evolve in a direction that is more cluster-friendly, emphasizing implicit ALUA to support failover in shared volume access environments today and identifying potential opportunities to continue to expand these capabilities going forward.

2017 Storage Developer Conference. © NetApp, Inc. All Rights Reserved.

Thank you

2017 Storage Developer Conference. © NetApp, Inc. All Rights Reserved.