BROCADE
A Broadcom Inc. Company

VALIDATION TEST REPORT

Brocade Fabric Technology with the Tegile T3100 Hybrid Flash Array

Supporting Fabric OS 7.4.1, 8.0.1, 8.1.0, 8.2.0, and 8.2.1

53-1004546-04
26 February 2019

Contents

Preface

Document History

Date	Part Number	Description
July 21, 2016	53-1004546-01	Supporting Brocade Fabric OS 8.0.1.
May 24, 2017	53-1004546-02	Supporting Brocade Fabric OS 8.1.0.
March 16, 2018	53-1004546-03	Supporting Brocade Fabric OS 8.2.0.
February 26, 2019	53-1004546-04	Supporting Brocade Fabric OS 8.2.1.

Overview

The Storage Fabric Ready (SFR) program is a comprehensive testing and configuration initiative to validate the interoperability of Fibre Channel flash storage with a Brocade FC network infrastructure. This program provides testing of multiple fabrics, heterogeneous servers, and HBAs in a large-port-count Brocade environment. The SFR qualification program helps verify seamless interoperability and optimum performance with solid-state storage systems in Brocade FC storage fabrics.

Purpose of This Document

The goal of this document is to demonstrate the compatibility of the Tegile T3100 FC storage array in a Brocade FC fabric containing Gen 5 (16 Gb) and Gen 6 (32 Gb) FC switches. This document provides a test report on the SFR qualification test plan executed on the Tegile T3100 storage array.

Audience

The target audience for this document includes storage administrators, solution architects, system engineers, and technical development representatives.

Objective

- Test the Tegile T3100 array with the Brocade FC fabric in single and routed configurations for different stress and error recovery scenarios and thereby validate the interoperability and integration of the Tegile T3100 array with Brocade FC fabrics.

- Validate the performance of the Brocade FC fabric in a solid-state storage environment for high-throughput and low-latency applications.

Related Documents

- *Brocade Fabric OS Administration Guide*
- *Brocade Monitoring and Alerting Policy Suite Configuration Guide*
- *Brocade Flow Vision Configuration Guide*
- *Brocade Fabric OS Command Reference Manual*
- *Brocade San Design and Best Practices*
- *Brocade SAN Fabric Resiliency and Administration Best Practices*
- Emulex ExpressLane Configuration

About Broadcom

Broadcom Inc. provides innovative storage networking solutions for data center, campus, and service provider networks, helping to reduce complexity and cost while enabling virtualization and cloud computing to increase business agility. To help ensure a complete solution, Broadcom partners with world-class IT companies and provides comprehensive education, support, and professional services offerings (www.broadcom.com).

About Tegile, a Western Digital® Brand

Tegile, a Western Digital® brand, is pioneering a new generation of all-flash and flash-driven hybrid enterprise storage arrays that balance performance, capacity, features, and price for virtualization, file services, and database applications.

Our hybrid arrays are significantly faster than legacy arrays and significantly less expensive than all solid-state disk-based arrays. Featuring both NAS and SAN connectivity, these virtual data storage systems are easy to use, fully redundant, and highly scalable. Additionally, they come complete with built-in snapshot, replication, near-instant recovery, and virtualization management features.

Tegile's patented IntelliFlash™ technology accelerates performance to solid-state speeds without sacrificing the capacity or cost advantage of hard-disk storage. Additionally it enables on-the-fly de-duplication and compression so usable capacity is far greater than its raw capacity.

Tegile offers an extensive line of hybrid storage solutions as part of the Intelligent Flash Array portfolio. Tegile hybrid arrays leverage the performance of flash, the density of hard disks, and the rich features of the IntelliFlash operating system to deliver a compelling storage platform that accelerates a wide variety of workloads in the enterprise.

Fully redundant with active/active controllers, these arrays are built for enterprise data centers with resilience, data availability, and data protection in mind. With Tegile hybrid arrays, you no longer have to compromise between performance, capacity, and cost for accelerating and protecting your enterprise applications.

Configure DUT and Test Equipment

Task 1: Brocade FC Fabric Configuration

1. Enable MAPS and Fabric Performance Impact (FPI) monitoring on all switches in the fabric.

   MAPS enables health monitoring on switches to detect potential faults and create alerts. MAPS FPI allows fabric monitoring for performance impacts, including timeouts, latency, and throughput. Detailed information on MAPS configuration and setup can be found in the *Brocade Monitoring and Alerting Policy Suite Configuration Guide*.

   - MAPS requires a "Fabric Vision" license to be installed.
   - Enable the desired MAPS policy using any of the available default policies or create a custom policy.
   - FPI monitoring is enabled by default.

   ```
   > mapsconfig --enablemaps -policy dflt_aggressive_policy

   > mapsconfig --actions raslog,email,sw_marginal,sw_critical

   > mapsconfig --show
   Configured Notifications:       RASLOG,EMAIL,SW_CRITICAL,SW_MARGINAL
   Mail Recipient:                 testuser@domain.com
   Paused members :
   ===============
   PORT :
   CIRCUIT :
   SFP :
   ```

   To view the summary of events and rules triggered:

   ```
   > mapsdb --show
   1 Dashboard Information:
   ========================
   DB start time:                  Tue Mar 14 22:44:04 2017
   Active policy:                  dflt_aggressive_policy
   Configured Notifications:       RASLOG,EMAIL,SW_CRITICAL,SW_MARGINAL
   Fenced Ports :                  None
   Decommissioned Ports :          None
   Fenced circuits :               N/A
   Quarantined Ports :             None
   Top Zoned PIDs <pid(it-flows)>: 0x4a1901(30) 0x4a1a01(30) 0x4a1801(30) 0x4a1b01(30) 0x4a0400(25)

   2 Switch Health Report:
   ========================
   Current Switch Policy Status: HEALTHY
   ……
   ```

2. Configure Flow Monitoring with the I/O Insight feature on the Brocade Gen 6 switches in the fabric.

   The I/O Insight feature supported on the Brocade Gen 6 switches allows us to monitor the flow latency statistics at the SCSI I/O exchange level. The monitoring can be configured at an IT (Initiator-Target) flow level on fixed-port switches and at an ITL (Initiator-Target-LUN) flow level on chassis-based switches.

   Requires a “Fabric Vision” license and an “I/O Insight” license.

   a) Create a "Flow Monitor" flow at the source or destination device port on the Brocade G620/G610 switch.

```
> flow --create ios_tegile_1 -fea mon -dstdev e01100 -egrport 17
Monitor feature(s) have been activated.
> flow --create ios_tegile_2 -fea mon -dstdev e01200 -egrport 18
Monitor feature(s) have been activated.

> flow --show ios_tegile_1
================================================================================
Name      : ios_tegile_1        Features: mon(Activated)        noConfig: Off
Definition: EgrPort(17),DstDev(0xe01100)
```

   b) Import the created flows into MAPS.

```
> mapsconfig --import ios_tegile_1
> mapsconfig --import ios_tegile_2

> logicalgroup --show
-----------------------------------------------------------------------------------------------
Group Name                      |Predefined |Type              |Member Count |Members
-----------------------------------------------------------------------------------------------
ios_tegile_1                    |No         |Flow              |1            |Monitored Flow
ios_tegile_2                    |No         |Flow              |1            |Monitored Flow
```

   c) Create MAPS rules to monitor the desired SCSI I/O latency statistics and add it to a custom MAPS policy.

```
> mapspolicy --clone dflt_aggressive_policy -name ios_aggressive_policy

> mapsrule --create ios_tegile1_rd_latency -group ios_tegile_1 -monitor RD_STATUS_TIME_LT_8K -timebase min -op
ge -value 2500 -action email,raslog -policy ios_aggressive_policy

> mapspolicy --show ios_aggressive_policy
Policy Name: ios_aggressive_policy

Rule Name                              |Condition                                        |Actions     |
---------------------------------------------------------------------------------------------------------
ios_tegile1_wr_latency                 |ios_tegile_1(WR_STATUS_TIME_LT_8K/min>=2500)  |email,raslog|
ios_tegile2_wr_latency                 |ios_tegile_2(WR_STATUS_TIME_LT_8K/min>=2500)  |email,raslog|
ios_tegile1_rd_latency                 |ios_tegile_1(RD_STATUS_TIME_LT_8K/min>=2500)  |email,raslog|
ios_tegile2_rd_latency                 |ios_tegile_2(RD_STATUS_TIME_LT_8K/min>=2500)  |email,raslog|
```

   d) Enable the MAPS policy.

```
> mapspolicy --enable ios_aggressive_policy

> mapsdb --show

1 Dashboard Information:
========================
DB start time:                  Thu Mar  9 16:48:07 2017
Active policy:                  ios_aggressive_policy
Configured Notifications:       RASLOG,SW_CRITICAL,SW_MARGINAL
Fenced Ports :                  None
Decommissioned Ports :          None
Fenced circuits :               N/A
Quarantined Ports :             None
Top Zoned PIDs <pid(it-flows)>: 0xe01a01(30) 0xe01b01(30) 0xe01901(30) 0xe01801(30) 0xe02200(22)

2 Switch Health Report:
```

```
 ======================
 Current Switch Policy Status: HEALTHY
```

3. Zoning is configured using the "Peer Zoning" feature in FOS.

   Peer Zoning allows a "principal" device to communicate with the rest of the devices in the zone. The principal device manages a Peer Zone. Other "non-principal" devices in the zone can communicate with the principal device only; they cannot communicate with each other.

```
> zonecreate --peerzone sfr_tegile_1 -principal "21:00:00:0e:1e:2e:d8:00; ...." -members
"10:00:00:90:fa:c7:d2:76; ...."

> zoneshow --peerzone all
zone:  ssr_tegile_1
   Property Member: 00:02:00:00:00:03:00:04
   Created by: User
   Principal Member(s):
                21:00:00:0e:1e:2e:d8:01
                21:00:00:0e:1e:2e:d8:00
                21:00:00:0e:1e:2e:d5:f1
                21:00:00:0e:1e:2e:d5:f0
   Peer Member(s):
                10:00:00:90:fa:c7:d2:76
                10:00:00:90:fa:c7:d2:77
                10:00:00:90:fa:61:92:3b
                10:00:00:90:fa:61:92:3c
                21:00:00:24:ff:7d:3d:ea
                21:00:00:24:ff:7d:3d:eb
                10:00:00:10:86:04:83:e6
                10:00:00:10:86:04:83:e7
                21:00:00:0e:1e:18:99:90
                21:00:00:0e:1e:18:99:91

> cfgactvshow
zone:  ssr_tegile_1
                00:02:00:00:00:03:00:04
                21:00:00:0e:1e:2e:d8:01
                21:00:00:0e:1e:2e:d8:00
                21:00:00:0e:1e:2e:d5:f1
                21:00:00:0e:1e:2e:d5:f0
                10:00:00:90:fa:c7:d2:76
                10:00:00:90:fa:c7:d2:77
                10:00:00:90:fa:61:92:3b
                10:00:00:90:fa:61:92:3c
                21:00:00:24:ff:7d:3d:ea
                21:00:00:24:ff:7d:3d:eb
                10:00:00:10:86:04:83:e6
                10:00:00:10:86:04:83:e7
                21:00:00:0e:1e:18:99:90
                21:00:00:0e:1e:18:99:91
```

4. Configure Fibre Channel Routing (requires an Integrated Routing license).

   The FC-FC routing service provides Fibre Channel routing between two or more fabrics without merging those fabrics. For example, using FC-FC routing, you can share tape drives across multiple fabrics without the administrative problems such as change management, network management, scalability, reliability, availability, and serviceability that might result from merging the fabrics. Detailed information on FCR setup can be found in the *Brocade Fabric OS Administration Guide*.

   An example FCR configuration is shown below.

   Enabling FCR on the backbone fabric switches.

   ```
   > fcrconfigure -bbfid 100
   > fosconfig --enable fcr
   ```

   Configuring EX ports connecting to edge fabrics.

   ```
   > portcfgexport [port#] -a1 -m0 -f 10
   > portcfgexport [port#] -a1 -m0 -f 20
   > portcfgexport [port#] -a1 -m5 -f 50
   ```

   LSAN zone is created on both fabrics.

   ```
   zone:  LSAN_ssr067146_tegile
                  21:00:00:0e:1e:2e:d8:01; 21:00:00:0e:1e:2e:d8:00;
                  21:00:00:0e:1e:2e:d5:f1; 21:00:00:0e:1e:2e:d5:f0;
                  10:00:8c:7c:ff:24:6d:00; 10:00:8c:7c:ff:24:6d:01
   ```

   Example output of exported devices

   ```
   > fcrproxydevshow
     Proxy           WWN             Proxy      Device   Physical    State
    Created                           PID       Exists     PID
   in Fabric                                   in Fabric
   ------------------------------------------------------------------------
       20   21:00:00:0e:1e:2e:d5:f0  02f501       10      e01100   Imported
       20   21:00:00:0e:1e:2e:d5:f1  02f401       10      4a0400   Imported
       20   21:00:00:0e:1e:2e:d8:00  02f301       10      e01200   Imported
       20   21:00:00:0e:1e:2e:d8:01  02f201       10      4a0500   Imported
   ```

Task 2: Tegile IntelliFlash T3100 Array Configuration

1. Configure "High Availability" between the array controllers. Create the storage pool in an active/passive configuration and assign it to the active controller in the HA pair.

FIGURE 1 Tegile Data Pool Creation

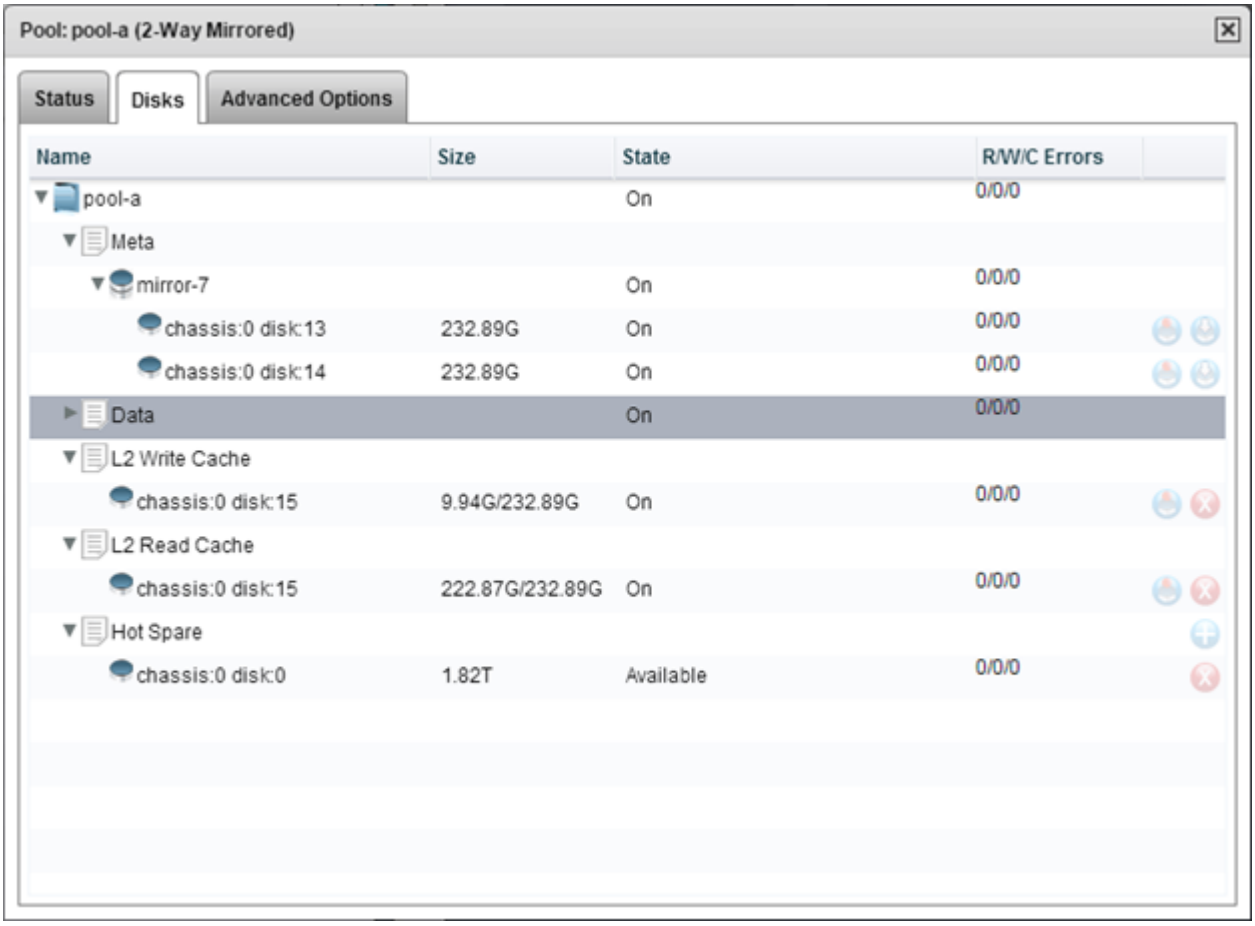

FIGURE 2 Tegile HA Configuration for the Storage Pool

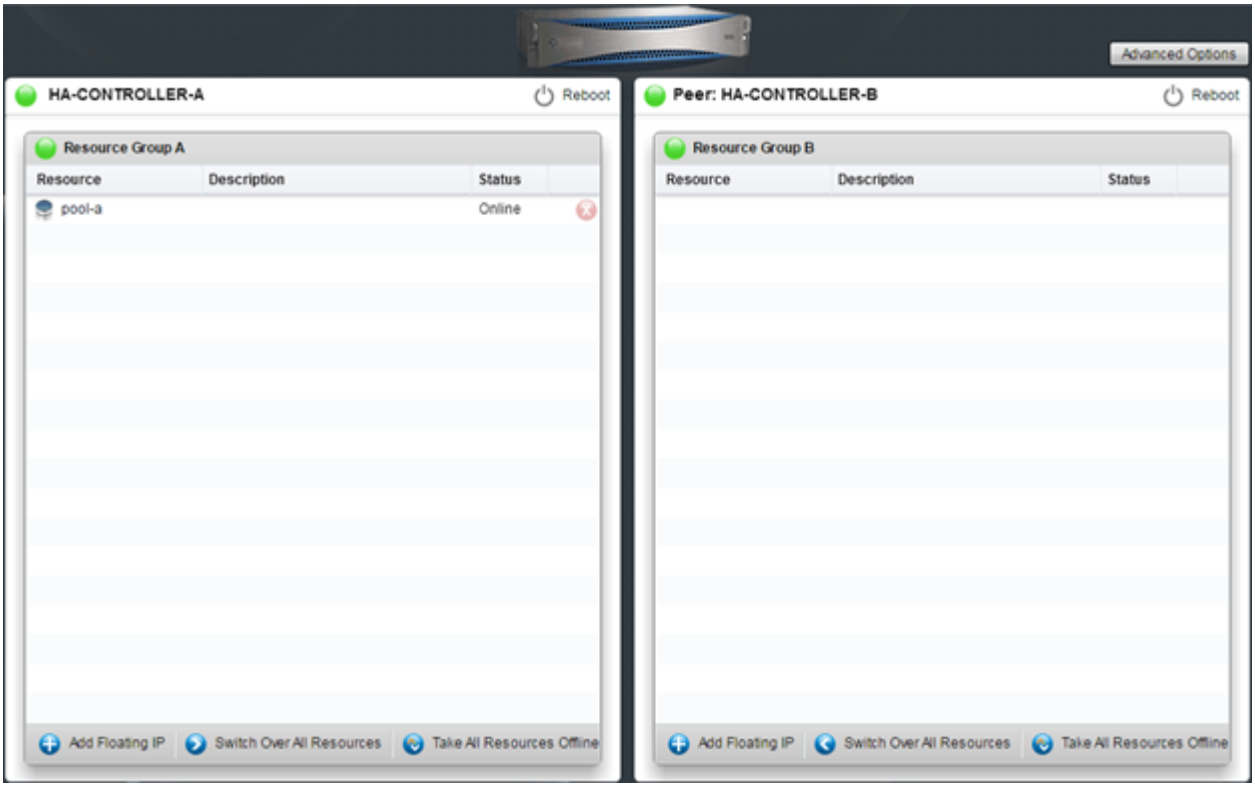

2. The FC target ports from both controllers on the array are connected to the Brocade FC fabric.

FIGURE 3 Tegile FC Target Connections

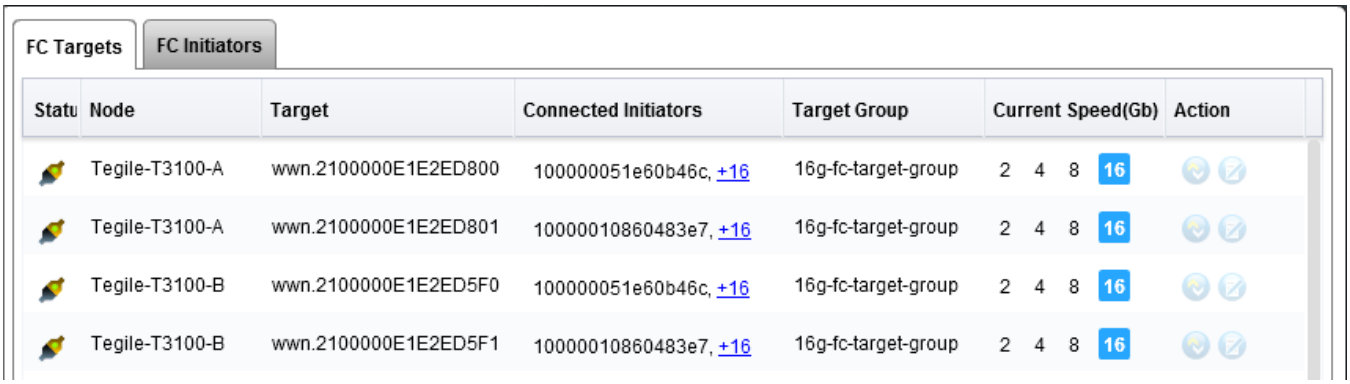

3. Create Initiator groups and volume group "projects" based on host group access. Default values for "Data Compression" and "Deduplication" are accepted. No "Snapshots" or "Replication" is configured.

FIGURE 4 Create Host Initiator Group

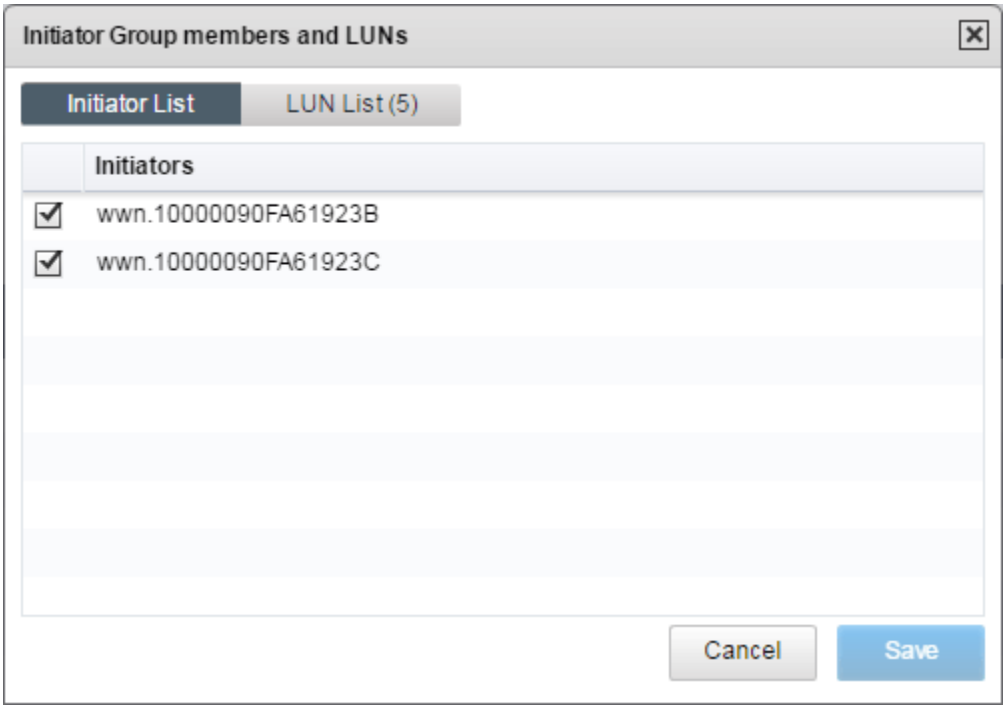

FIGURE 5 Tegile Project Configuration General Settings

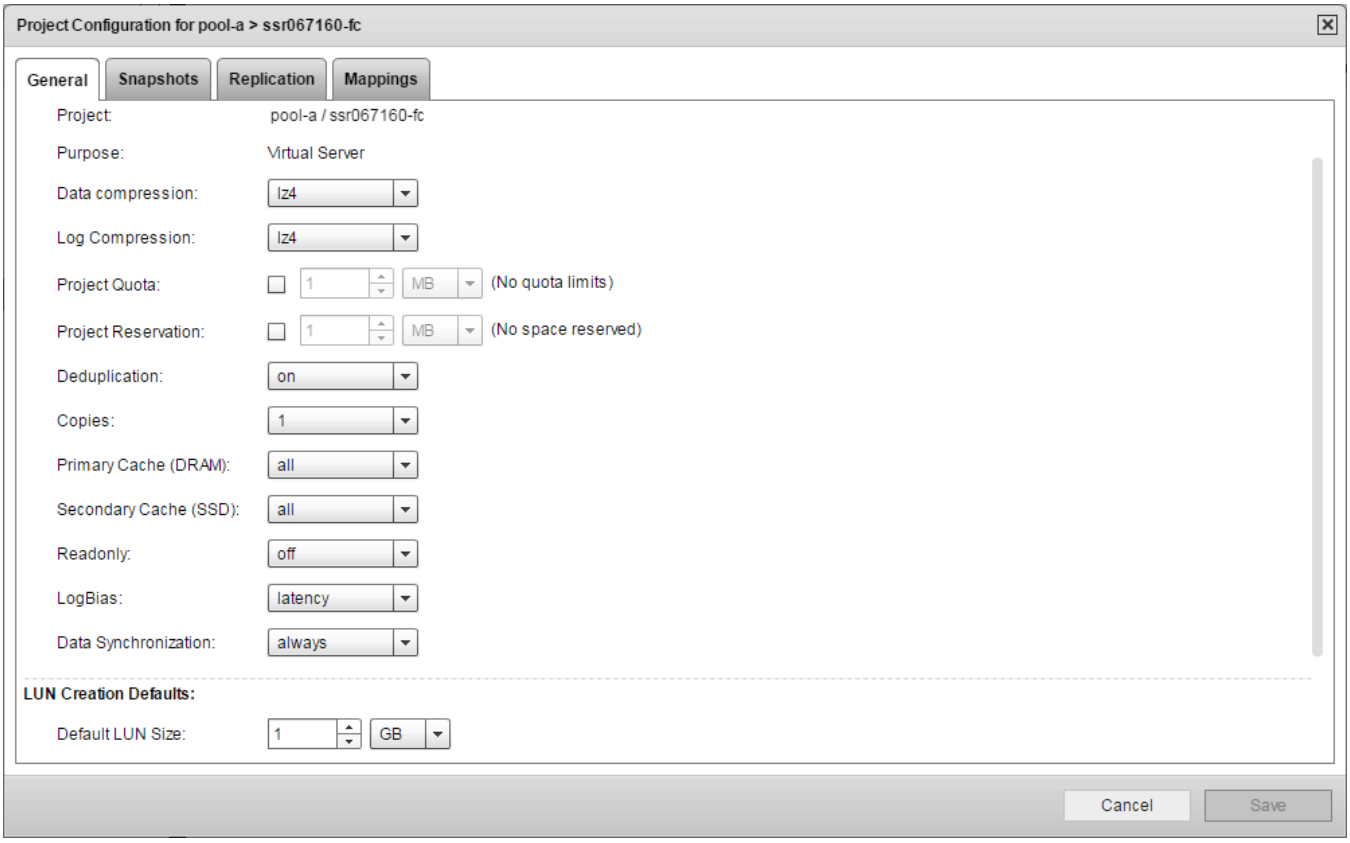

FIGURE 6 Tegile Project Mapping Configuration

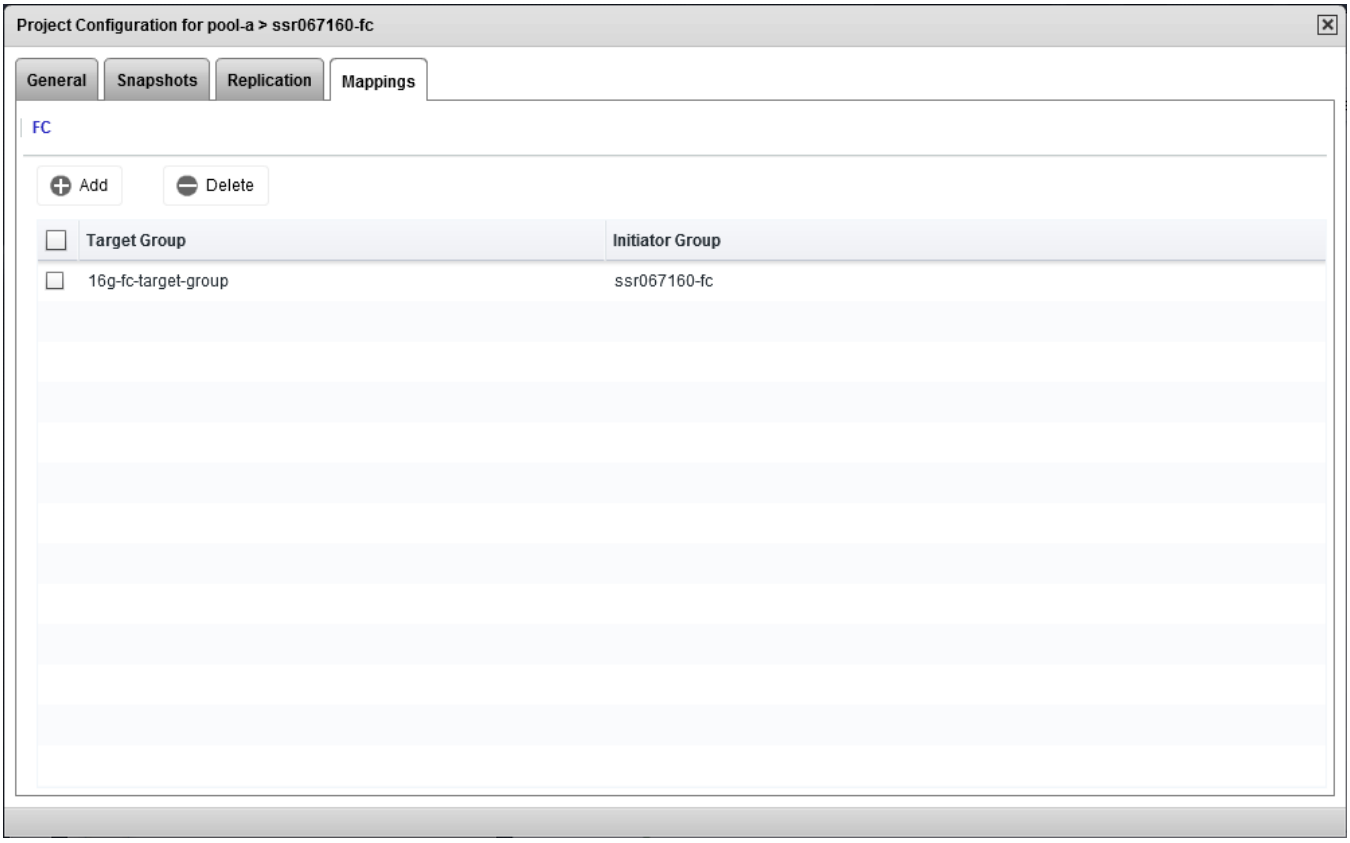

FIGURE 7 Tegile Project Creation

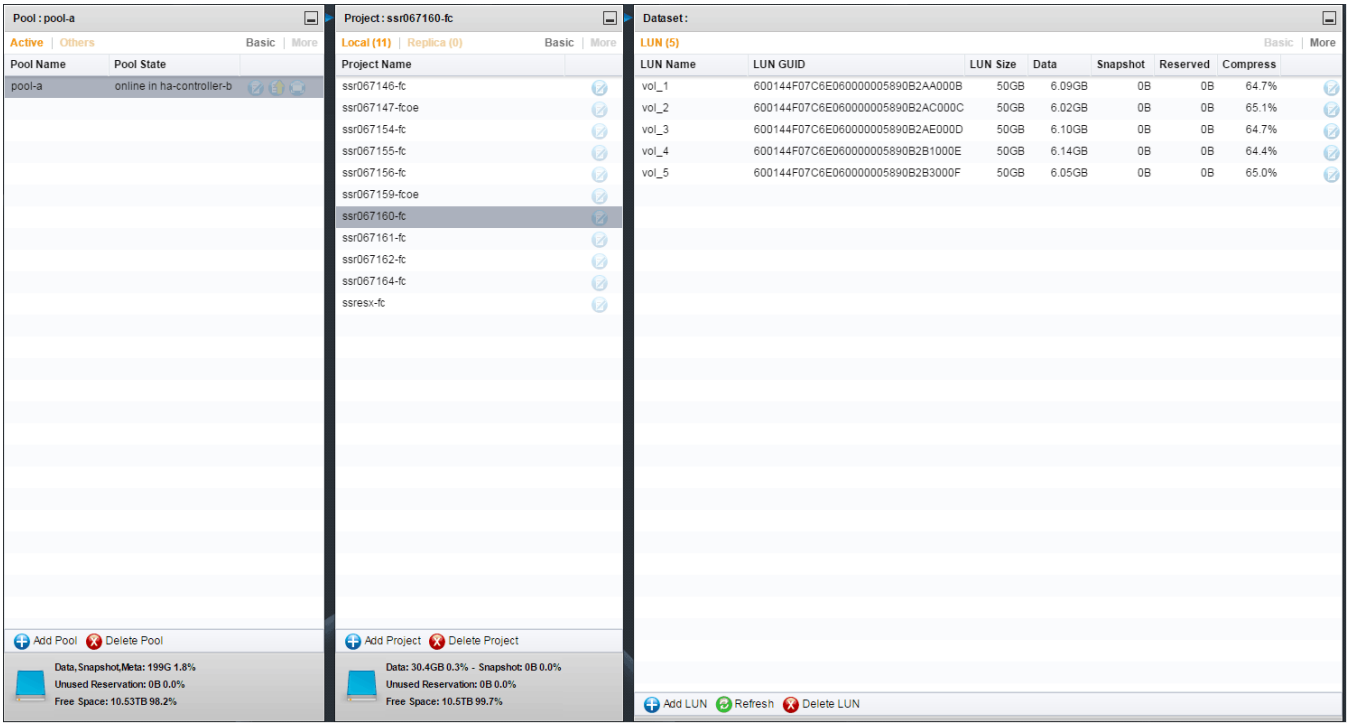

Task 3: Host Setup

1. Provision a minimum of two uplinks from the host to the FC fabric for redundancy and use native multipath tools to manage the available paths and load-balance across them.

2. Configuring the multipath settings allows for proper failover and load balancing across the available links.
   - For Windows, native MPIO is used and the following parameters are tuned in the Windows Registry for optimal path recovery and HA failover. The following Windows MSDSM parameters are required for HA failover to work properly:

```
PathVerificationState     : Enabled
PathVerificationPeriod    : 5
PDORemovePeriod           : 120
RetryCount                : 100
RetryInterval             : 1
```

FIGURE 8 Windows Disk MPIO Properties

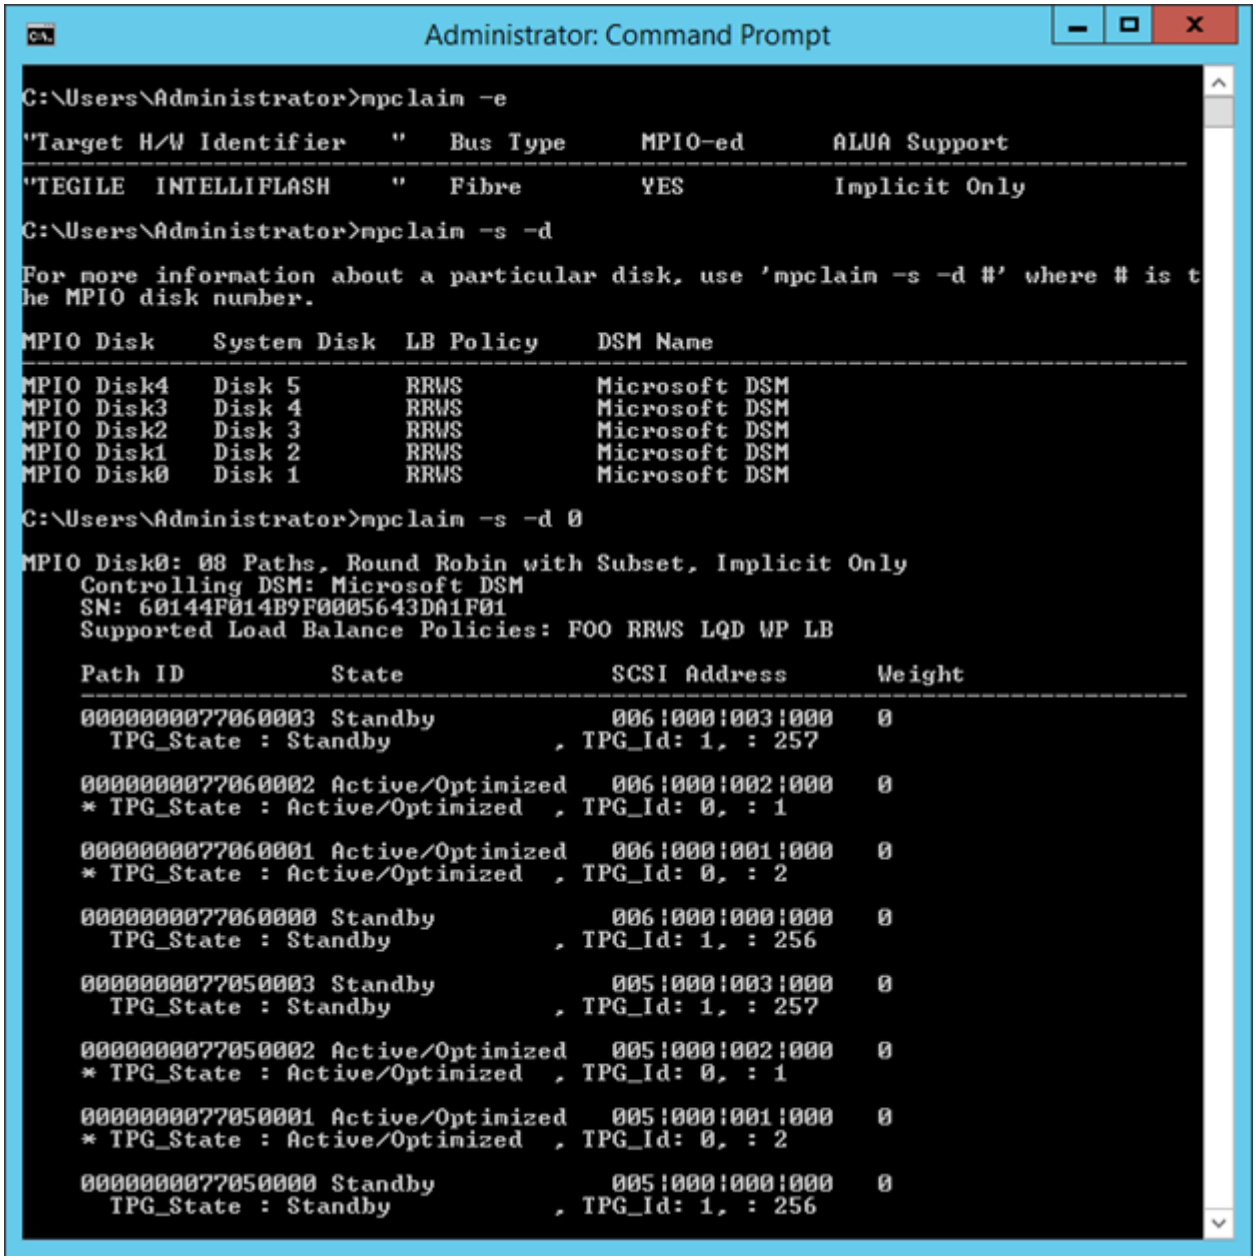

   - For Linux, add the following to **/etc/multipath.conf**:

```
devices {
device {
                vendor                  "TEGILE"
                product                 "INTELLIFLASH"
                hardware_handler        "1 alua"
                path_selector           "round-robin 0"
                path_grouping_policy    "group_by_prio"
                no_path_retry           10
                dev_loss_tmo            50
                path_checker            tur
                prio                    alua
                failback                30
        }
}
```

Sample output:

```
# multipath -ll
mpathaw (3600144f014b90f000000561bfff10006) dm-2 TEGILE,INTELLIFLASH
size=50G features='1 queue_if_no_path' hwhandler='1 alua' wp=rw
|-+- policy='round-robin 0' prio=130 status=active
| |- 0:0:0:5 sda  8:0    active ready running
| |- 1:0:1:5 sdz  65:144 active ready running
| |- 0:0:3:5 sdp  8:240  active ready running
| `- 1:0:3:5 sdaj 66:48  active ready running
`-+- policy='round-robin 0' prio=1 status=enabled
  |- 0:0:2:5 sdk  8:160  active ready running
  |- 1:0:0:5 sdu  65:64  active ready running
  |- 0:0:1:5 sdf  8:80   active ready running
  `- 1:0:2:5 sdae 65:224 active ready running
```

- For VMware, set the Path Selection Policy to "Round Robin" for the discovered device.

FIGURE 9 VMware Path Selection Policy for Tegile

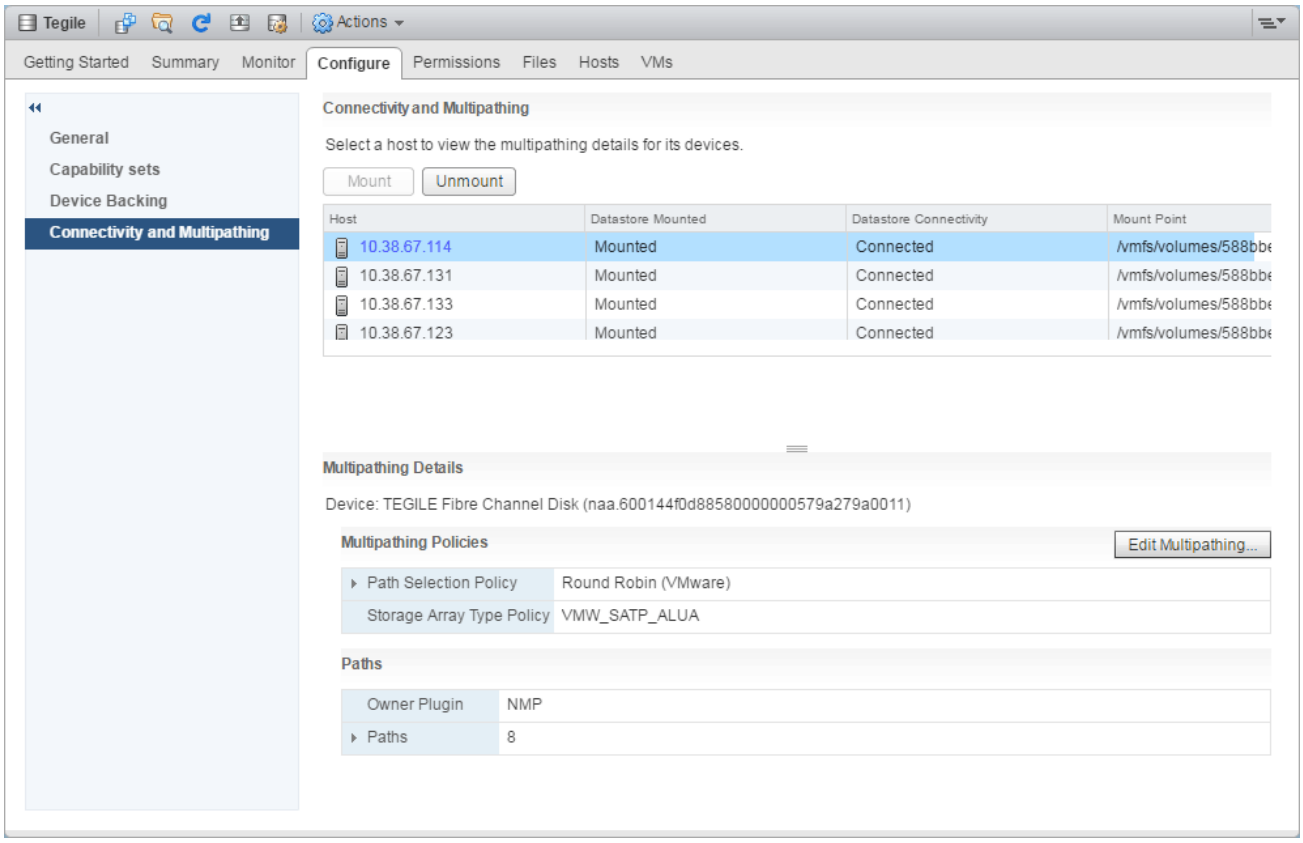

FIGURE 10 Tegile Devices Discovered on VMware Host

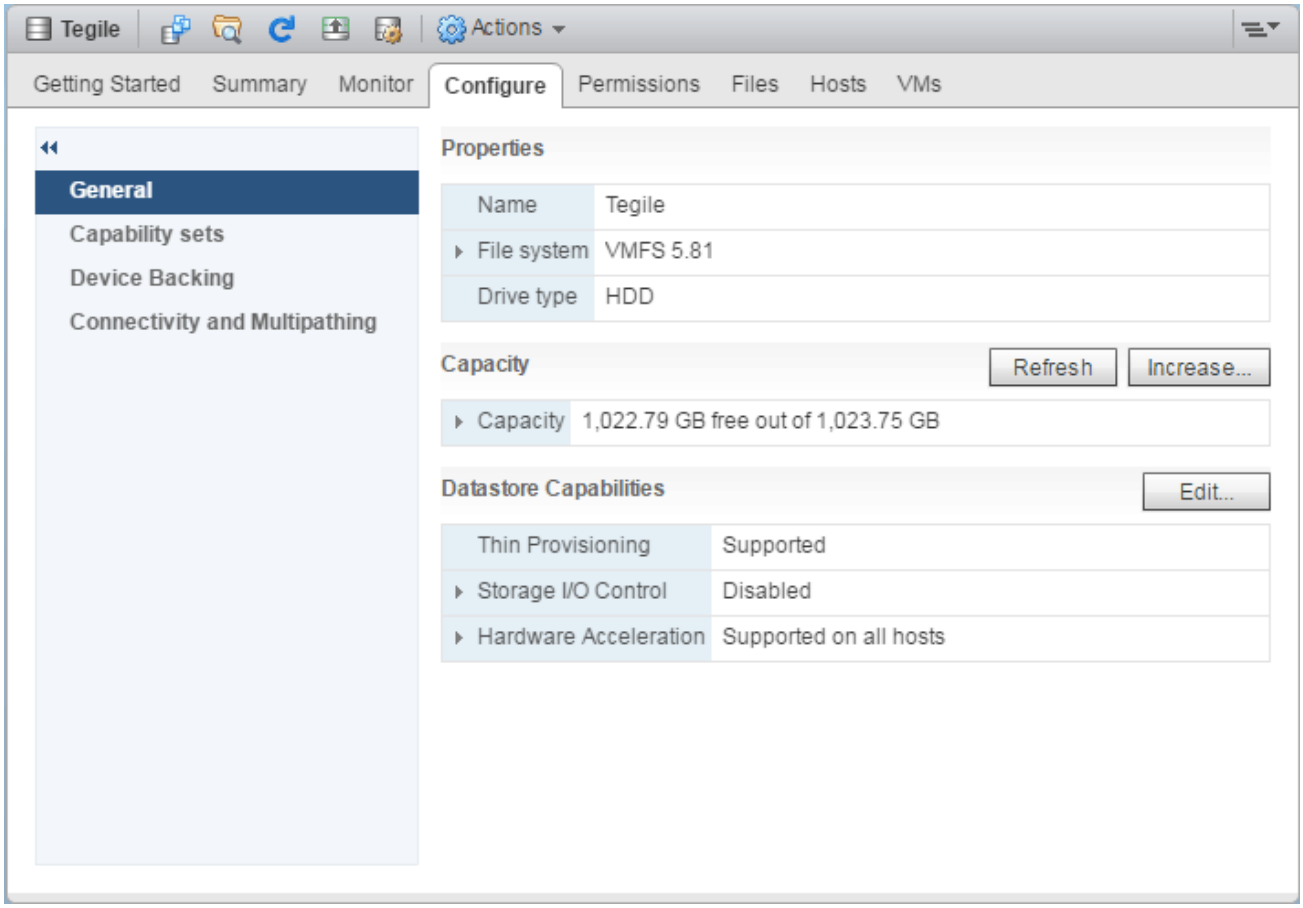

Tegile IntelliFlash T3100 Test Report

What's New in This Report

- The array under test is the Tegile IntelliFlash T3100 running firmware version 3.7.1.2.181123(GA).
- The Brocade Fabric OS (FOS) version under test is 8.2.1.
- All HBAs have updated firmware and drivers.

Test History

Storage Model	Storage Firmware	Brocade FOS Version	Date
Tegile T3100 (8Gb)	3.0.1.150929(HF1)	FOS 7.4.1	February 2016
Tegile T3100 (8Gb)	3.0.2.160404(GA)	FOS 8.0.1	July 2016
Tegile T3100 (16Gb)	3.5.2.3.170510(GA)	FOS 8.1.0	March 2017
Tegile T3100 (16Gb)	3.7.0.2.171005(GA)	FOS 8.2.0	January 2018
Tegile T3100 (16Gb)	3.7.1.2.181123(GA)	FOS 8.2.1	February 2019

Test Plan

The Tegile IntelliFlash T3100 array is connected to a Brocade FC fabric with two 16Gb FC target ports from the "active" and "standby" controller balanced across the Brocade FC fabric.

Scope

Testing focusses on interoperability of the Tegile storage array and determining an optimal configuration for performance and availability.

Testing covers various I/O stress and error handling scenarios. Performance is observed within the context of best practice fabric configuration; however absolute maximum benchmark reporting of storage performance is beyond the scope of this test.

Details of the test steps are covered under the "Test Cases" section. Standard test-bed setup includes IBM/HP/Dell servers with QLogic/Emulex HBAs with two uplinks from every host to the Brocade FC fabric. I/O generator tools used include Medusa Labs Test Tools and VMware I/O Analyzer.

Test Configuration

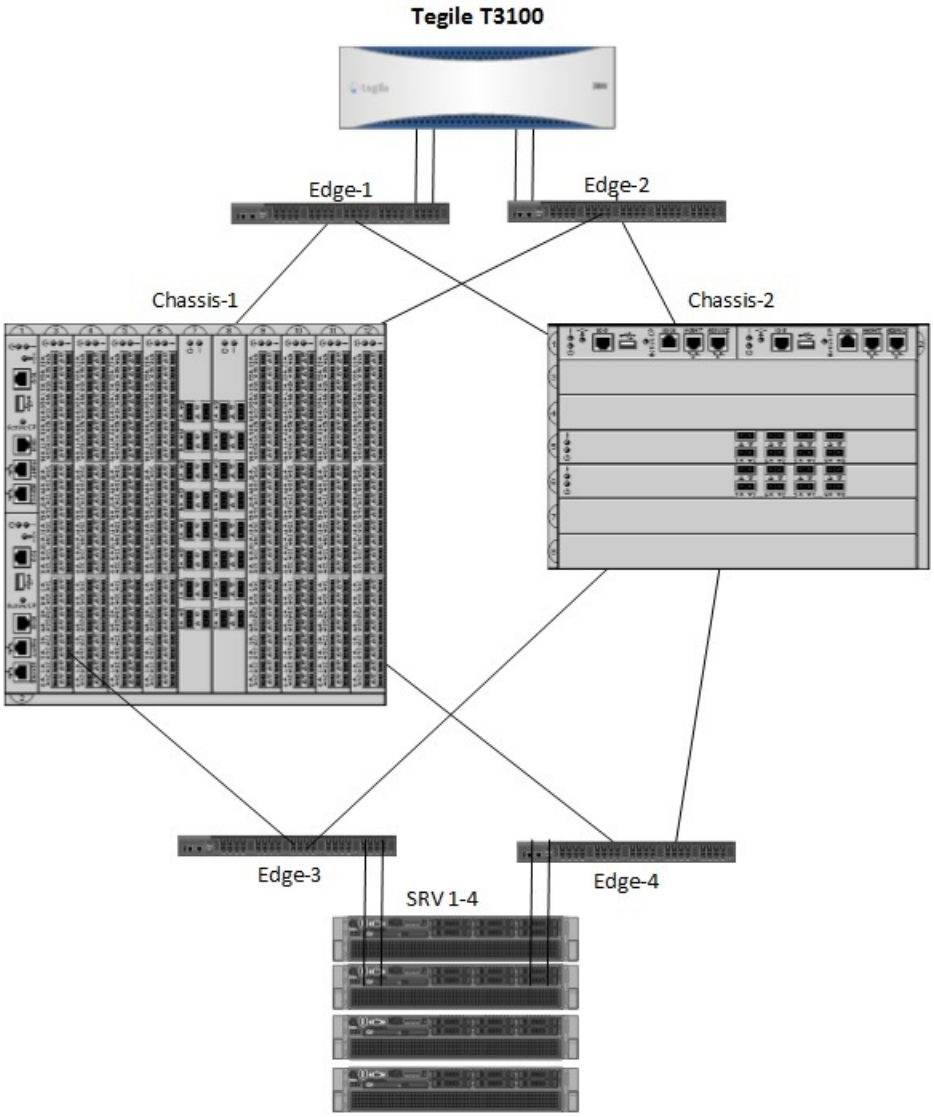

DUT Descriptions

The following tables provide details about the devices under test (DUTs).

TABLE 1 Storage Array

DUT ID	Model	Vendor	Description
Tegile T3100	T3100	Tegile	The T3100 array is a hybrid-flash array setup with dual controllers in an HA resource cluster. The controllers are in an active/standby configuration with 2x16Gb FC and 2x10GbE ports on each controller. The array supports FC, iSCSI, NFS, CIFS, and SMB protocols.

TABLE 2 Switches

DUT ID	Model	Vendor	Description
Edge-1	Brocade 6510	Brocade	48-port Gen 5 16Gb FC switch
Edge-2	Brocade G620	Brocade	64-port Gen 6 32Gb FC switch
Edge-3	Brocade G610	Brocade	28-port Gen 6 32Gb FC switch
Edge-4	Brocade G620	Brocade	64-port Gen 6 32Gb FC switch
Chassis-1	Brocade X6-8	Brocade	8-slot Gen 6 32Gb director
Chassis-2	Brocade X6-4	Brocade	4-slot Gen 6 32Gb director

Test Cases

These test cases are designed to verify basic and advanced functionality features between the Brocade FC fabric and the Tegile T3100 hybrid flash array and host devices; to stress all devices; and to confirm successful error recovery.

1.	**FABRIC INITIALIZATION—BASE FUNCTIONALITY**
1.1	Storage Device—Physical and Logical Login with Speed Negotiation
1.2	Zoning and LUN Mapping
1.3	Storage Device Fabric I/O Integrity
1.4	Storage Device Multipath Configuration—Path Integrity
2.	**FABRIC—ADVANCED FUNCTIONALITY**
2.1	Storage Device Bottleneck Detection Using MAPS FPI—With Congested Host
2.2	Bottleneck Detection Using MAPS FPI—With Congested Fabric
2.3	Flow Monitoring with I/O Insight and MAPS
2.4	QoS Integrity with QoS Zone-Based Traffic Prioritization
2.5	QoS Integrity with CS_CTL-Based Frame Prioritization
2.6	Storage Device—FC Protocol Jammer Test Suite
2.7	Clear Link Diagnostics (D_Port) Test
3.	**STRESS & ERROR RECOVERY WITH DEVICE MULTIPATH**
3.1	Storage Device Fabric I/O Integrity—Congested Fabric
3.2	Storage Device Nameserver Integrity—Device Recovery with Port Toggle
3.3	Storage Device Integrity—Device Recovery with Device Relocation
3.4	Storage Device Stress—Device Recovery with Device Port Toggle—Extended Run
3.5	Storage Device Recovery—ISL Port Toggle (Sequential)
3.6	Storage Device Recovery—ISL Port Toggle (Entire Switch)
3.7	Storage Device Recovery—Director Blade Maintenance
3.8	Storage Device Recovery—Switch Offline
3.9	Storage Device Recovery—Switch Firmware Download
4.	**STORAGE DEVICE—FIBRE CHANNEL ROUTING (FCR) INTERNETWORKING TESTS**
4.1	Storage Device Internetworking Validation with FC Host
4.2	Storage Device Internetworking Validation with FCoE Using VDX FlexPort
4.3	Storage Device Edge Recovery After FCR Disruptions
4.4	Storage Device Backbone Recovery After FCR Disruptions
5	**OPTIONAL/ADDITIONAL TESTS**
5.1	Non-disruptive Firmware Upgrade on Storage Device
5.2	Workload Simulation Test Suite—Medusa
5.3	Workload Simulation Test Suite—VMware

1. Fabric Initialization—Base Functionality

1.1 Storage Device—Physical and Logical Login with Speed Negotiation

Test Objective

Verify device login to switch and name-server with all supported speed settings.

Test Configuration

Set switch ports to 4/8/16/Auto_Negotiate speed settings.

```
portcfgspeed <port> [4/8/16/0]
```

Result Validation

1. Validate link states on the array and verify speed negotiation and device login at different speeds.
2. Check switch port status and verify the “actual” and “configured” link speed. Check name server for device login.

```
# nscamshow
# portshow X
```

Test Results

PASS. Speed negotiation, device login, and connectivity verified.

1.2 Zoning and LUN Mapping

Test Objective

Verify host-to-LUN access exists with valid zoning.

Test Configuration

1. Create peer zones on the fabric with the target WWNs as the principal members and the initiator WWNs as the peer members.
2. Create host groups and LUNs on the array with access to the initiator WWN.

Result Validation

Verify that LUNs are discovered on the hosts with host-specific tools.

1. Linux: Check output of **lsscsi**
2. Windows: Check output of **Computer Management > Storage > Disk Management**
3. VMware: Check output at **Configuration > Storage > Devices**

Test Results

PASS. Host has read/write access to the presented LUNs.

1.3 Storage Device Fabric I/O Integrity

Test Objective

Validate single path host-to-LUN I/O with write/read/verify testing. Include short device cable pulls/port-toggle to validate device recovery.

Test Configuration

1. Set up read/write I/O to LUN using Medusa.
2. Perform link disruptions by port toggles, cable pulls.

Result Validation

Check the Medusa I/O logs and verify that I/O resumes after a short downtime. Medusa I/O may pause, but should recover without error.

Test Results

PASS. I/O resumes after the disruption.

1.4 Storage Device Multipath Configuration—Path Integrity

Test Objective

Verify multi-path configures successfully. Each Adapter and Storage port to reside in different switches. For all device paths, consecutively isolate individual paths and validate I/O integrity and path recovery.

Test Configuration

1. Set up the host with at least two initiator ports zoned with two target ports on the array.
2. Set up multipath on the host, and start I/O.
3. Perform sequential port toggles across initiator and target switch ports to isolate paths.

Result Validation

1. Check host multipath properties to verify that the toggled path recovers.
    a. Windows: **mpclaim –s –d**
    b. Linux: **multipath –ll**
    c. VMware: Check the paths at **Configuration > Storage > Devices > Manage Paths**
2. Check the host and storage logs for failures.

3. Check the switch error logs and the switch port status after toggling.

```
# errdumpall
# portstatsshow X
# portshow X
```

4. Check I/O logs, and verify that I/O continues without error.

Test Results

PASS. I/O fails over to the remaining active paths and recovers when the disrupted path is restored.

2. Fabric—Advanced Functionality

2.1 Storage Device Bottleneck Detection Using MAPS FPI—With Congested Host

Test Objective

Verify that congestion on host ports is detected. Verify storage device and fabric behavior during congestion.

Test Configuration

1. Enable MAPS monitoring and MAPS FPI on all switches. A Fabric Vision license is required.
2. Start I/O from single host initiator to multiple targets.
3. Monitor switch logs for Congestion and Latency (IO_PERF_IMPACT/IO_FRAME_LOSS) warnings.

Result Validation

Check the switch error logs and MAPS dashboard for bottleneck warnings.

```
# errdumpall | grep IO_
# mapsdb --show all
```

Test Results

PASS. The bottlenecked ports are displayed on the MAPS dashboard, and a raslog warning is also created.

2.2 Bottleneck Detection Using MAPS FPI—With Congested Fabric

Test Objective

Create congestion on switch ISL port. Verify that congestion in the fabric is detected. Verify storage device and fabric behavior during congestion.

Test Configuration

1. Enable MAPS monitoring and MAPS FPI on all switches. A Fabric Vision license is required.
2. Isolate single ISL in the fabric.
3. Start I/O from multiple host initiator to multiple targets.
4. Monitor switch logs for Congestion and Latency (IO_PERF_IMPACT/IO_FRAME_LOSS) warnings.

Result Validation

Check the switch error logs and MAPS dashboard for bottleneck warnings.

```
# errdumpall | grep IO_
# mapsdb --show all
```

Test Results

PASS. The bottlenecked ports are displayed on the MAPS dashboard, and a raslog warning is also created.

2.3 Flow Monitoring with I/O Insight and MAPS

Test Objective

Monitor I/O latency statistics on the target ports. Verify that stats are reported accurately and that alerts are generated when thresholds are hit.

Test Configuration

1. Baseline a target LUN's latency running 4k reads from a workload generator. Here we are seeing about a 2–3 ms range. Hence we set the rules to a 2500-µs or 2.5-ms threshold.

   ```
   Avg Completion Time   0.002669
   ```

2. Set up flows and monitoring as per "Step 2 - Configure Flow Monitoring with I/O Insight" under the "Task 1. Brocade FC Fabric Configuration" section.
3. Start I/O, and adjust the traffic pattern to cause a rise in latency above the configured monitoring threshold; and confirm that RASLog, MAPS dashboard, and email notifications are generated.

Result Validation

1. Check flow statistics and MAPS alerts to verify that the metrics are reported correctly and that alerts are generated when thresholds are crossed.

```
# flow --show

RASLog:
> errdumpall
2016/05/25-15:19:28:481554, [MAPS-1003], 167953/165659, FID 128, WARNING, G620_066_223, Flow (ios_tegile_1),
Condition=ios_tegile_1(RD_STATUS_TIME_LT_8K/min>2500), Current Value:[ RD_STATUS_TIME_LT_8K,2839 Microseconds],
RuleName= ios_tegile1_4k_rd_status, Dashboard Category=Traffic Performance., raslogAction.c, line: 96, comp:md,
ltime:2016/05/25-15:19:28:481336

Email:
```

```
Switch Time:        May 25 15:19:28
Affected Entity:    Flow (ios_tegile_1)
Monitor:        Read completion time (RD_STATUS_TIME_LT_8K)
Rule Name:        ios_tegile1_4k_rd_status
Group:            ios_tegile_1
Condition:        ios_tegile_1(RD_STATUS_TIME_LT_8K/min>2500)
Current Value:        2839 Microseconds
Dashboard Category:    Traffic Performance
Switch Name:        G620_066_223
Switch WWN:        10:00:c4:f5:7c:2a:8b:c8
Switch IP:        10.38.66.223
Fabric Name:        SSR
VFID:            128
```

MAPS Dashboard:

```
> mapsdb –show

3.2 Rules Affecting Health:
===========================
Category(Rule Count)|RepeatCount|Rule Name               |Execution Time   |Object             |Triggered
                                                                                               |Value(Units) |
-------------------------------------------------------------------------------------------------------------
Traffic Performance(|16         |ios_tegile1_4k_rd_status|05/25/16 15:19:28|Flow (ios_tegile_1)|2839 Microsec|
```

Test Results

PASS. Verified that I/O stats are reported and that alerts are generated when thresholds are exceeded.

2.4 QoS Integrity with QoS Zone-Based Traffic Prioritization

Test Objective

Verify storage device behavior and validate traffic characteristics with different QoS zones.

Test Configuration

1. Set up initiator-target pairs with Low/Medium/High QoS zones in the fabric.

```
zone:  QOSH_sfr_tegile_1
                00:02:00:00:00:03:00:04; 21:00:00:24:ff:48:b9:6a;
                21:00:00:24:ff:48:b9:6b; 50:01:43:80:06:2d:08:2a;
                50:01:43:80:06:2d:08:28; 10:00:00:05:1e:60:b4:6c;
                10:00:00:05:1e:60:b4:6b; 10:00:8c:7c:ff:14:e0:01;
                10:00:8c:7c:ff:14:e0:00; 10:00:8c:7c:ff:03:bc:00;
                10:00:8c:7c:ff:03:bc:01; 10:00:8c:7c:ff:03:9b:00;
                10:00:8c:7c:ff:03:9b:01
```

2. Start I/O from all hosts, and verify I/O statistics.

Result Validation

1. Check the I/O logs, and verify that I/O continues without error.
2. Check the switch error logs and switch port status for errors.

```
# errdumpall
# porterrshow
```

Test Results

PASS. I/O from hosts in all QoS zones completed successfully without error.

2.5 QoS Integrity with CS_CTL-Based Frame Prioritization

Test Objective

Verify CS_CTL I/O prioritization using the Emulex ExpressLane feature.

Test Configuration

1. Configure all switches in the fabric to be in "Auto" CS_CTL QoS mode.

```
root> configurechassis
Configure...
cfgload attributes (yes, y, no, n): [no]
Custom attributes (yes, y, no, n): [no]
  system attributes (yes, y, no, n): [no]
  fos attributes (yes, y, no, n): [no] y

   Reboot needed to effect new CSCTL Mode
        CSCTL QoS Mode (0 = default; 1 = auto mode): (0..1) [0] 1

root> configshow -all | grep csctlMode
fos.csctlMode:1
```

2. Enable CS_CTL mode on initiator and target switch ports.

```
root> portcfgqos –enable [slot/]port csctl_mode

root> portcfgshow 22
……
CSCTL mode:               ON
……
```

3. Set up initiator-target zones in the fabric and discover the LUNs on the host.
4. Enable the ExpressLane feature on the host Emulex ports, and set the ExpressLane Priority (CS_CTL value) to high (3).
5. Enable ExpressLane on any of the discovered LUNs, and start write I/O from the host.
6. Verify I/O statistics, and verify the CS_CTL prioritization in the fabric.

Result Validation

1. Check I/O logs and verify I/O continues without any errors for all LUNs.
2. Verify I/O performance is improved on ExpressLane enabled LUN.
3. Verify CS_CTL prioritization in fabric by monitoring the VC (VC2-5=Medium; VC8-9=Low; VC10-14=High) buffer credits on the ISLs. Look for buffer credit values changing in the bold locations during I/O.

```
root> portregshow 0 | grep -E "_trc |bbc_mbc"
0x88982800: bbc_trc     4    0    2    2    2    2    1    1
0x88982820: bbc_trc     2    2    2    2    2    2    2    0
0x88982840: bbc_trc     0    0    0    0    0    0    0    0
0x88982860: bbc_trc     0    0    0    0    0    0    0    0
0x88982880: bbc_trc     0    0    0    0    0    0    0    0
```

4. Check switch error logs and switch port status for any errors.

```
# errdumpall
# porterrshow
```

Test Results

PASS. I/O completed successfully for all LUNs. ExpressLane-enabled LUN performance is improved, and I/O to the LUN is prioritized throughout the FC fabric.

2.6 Storage Device—FC Protocol Jammer Test Suite

Test Objective

Perform FC Jammer tests including areas such as: CRC corruption, packet corruption, missing frame, host error recovery, and target error recovery.

Test Configuration

1. Insert Jammer device in the I/O path on the storage end.
2. Execute the following Jammer scenarios:
   a. Delete one frame
   b. Delete R_RDY
   c. Replace CRC of data frame
   d. Replace EOF of data frame
   e. Replace "good status" with "check condition"
   f. Replace IDLE with LR
   g. Truncate frame
   h. Create S_ID/D_ID error of data frame

Result Validation

1. Check the host and storage logs for any errors.
2. Check the switch logs and interface stats for any errors.

```
# errdumpall
# porterrshow
```

3. Verify jammer operations and recovery with Analyzer.

Test Results

PASS. Host and target were able to recover from the errors and continue I/O operations.

2.7 Clear Link Diagnostics (D_Port) Test

Test Objective

Execute the Clear Link Diagnostics feature on all supported HBAs, and check for any port and link problems.

Tested with FOS 7.4.1.

Test Configuration

1. The Clear Link Diagnostic test is performed with Emulex LPe16202, QLogic 2672, and QLogic 1860 HBAs.
2. The Brocade FC switch can be configured in static, dynamic, and on-demand D_Port modes. The switch supports dynamic D_Port by default at the chassis level.

```
root> configure
Configure...
  Fabric parameters (yes, y, no, n): [no]
  D-Port Parameters (yes, y, no, n): [no] y
  Dynamic D-Port (on, off): [on]
               On Demand D-Port (on, off): [off]
```

3. With the Emulex LPe16202 adapter, the switch port is in dynamic mode, and the test is manually initiated on the HBA port from the Emulex OneCommand Manager.

```
root> switchshow
0   0   030000   id    N16      In_Sync     FC  D-Port 10:00:00:90:fa:61:92:3b Dynamic
```

4. With the QLogic 2672 adapter, the HBA port is in dynamic D_Port mode, and the test is initiated by configuring the switch port as a static D_Port.

```
root> portdisable 19
root> portcfgdport --enable 19
Caution: D_Port functionality is only available on 16Gb-capable platforms with 16Gb FC SFPs, 10Gb FC
SFPs, 8Gb LWL/ELWL FC SFPs, QSFPs or QSFP+.
root> portenable 19

root> portcfgshow 19
……
D-Port mode:              ON

root> switchshow
19  19   581300   id    N16      Online      FC  D-Port Loopback->Port 19
```

5. With the QLogic 1860 adapter, the HBA port can be in static or dynamic D_Port mode.

```
Enabling static D_Port on the HBA:

# bcu port --disable 2/0
port disabled

# bcu diag --dportenable 2/0
D-port mode for port 2/0 enabled.
```

6. Perform the D_Port test, and verify that all tests pass and that no port and link problems are reported.

Result Validation

Check the D_Port test results on the switch ports and host HBA diagnostic utilities.

```
Mode = Manual -> Static D_Port
Mode = Automatic -> Dynamic D_Port

root> portdporttest --show X
```

Test Results

PASS. D_Port tests passed with all tested HBAs.

3. Stress and Error Recovery with Device Multipath

3.1 Storage Device Fabric I/O Integrity—Congested Fabric

Test Objective

1. From all available initiators, start a mixture of READ/WRITE/VERIFY traffic with random data patterns continuously to all their targets overnight.
2. Verify that no host application failover or unexpected change in I/O throughput occurs.
3. Configure the fabric and devices for maximum link and device saturation.

Test Configuration

1. Start FC I/O to the storage array from multiple hosts.
2. Set up a mix of READ/WRITE traffic.

Result Validation

1. Check the host and storage logs for any errors.
2. Verify the link congestion and check the switch logs for errors.

```
# errdumpall
# portperfshow
# porterrshow
```

3. Check the I/O generator tool logs to verify that I/O runs without error.

Test Results

PASS. All I/O completed without error. All validation checks passed.

3.2 Storage Device Nameserver Integrity—Device Recovery with Port Toggle

Test Objective

1. With I/O running, perform a quick port toggle on every storage device and adapter port.
2. Verify that host I/O recovers.
3. Sequentially performed for each storage device and adapter port.

Test Configuration

1. Set up multipath on host, and start I/O.
2. Perform multiple iterations of sequential port toggles across initiator and target switch ports.

Result Validation

1. Check switch port status after toggle and for any errors in the switch error logs.

```
# errdumpall
# portstatsshow X
# portshow X
```

2. Check host multipath status on hosts to verify the toggled path recovers.
   a. Windows: **mpclaim -s -d**
   b. Linux: **multipath -ll**
   c. VMware: Check the paths at **Configuration > Storage > Devices > Manage Paths**
3. Check host and storage error logs, and verify that I/O continues without error.

Test Results

PASS. I/O failed over and recovered successfully. All validation checks passed.

3.3 Storage Device Integrity—Device Recovery from Device Relocation

Test Objective

1. With I/O running, manually disconnect and reconnect port to different switch in same fabric.
2. Verify host I/O will failover to alternate path and toggled path will recover.
3. Sequentially performed for each storage device and adapter port.
4. Repeat test for all switch types.

Test Configuration

1. Set up multipath on host, and start I/O.
2. Move storage target ports to different switch ports in the fabric.

Result Validation

1. Check for any errors in the switch error logs and the switch port status at the new switch port.

```
# errdumpall
# portstatsshow X
# portshow X
```

2. Check host multipath status on hosts to verify the toggled path recovers.
   a. Windows: **mpclaim -s -d**
   b. Linux: **multipath -ll**
   c. VMware: Check the paths at **Configuration > Storage > Devices > Manage Paths**
3. Check host and storage error logs, and verify that I/O continues without error.

Test Results

PASS. I/O failed over and recovered successfully. All validation checks passed.

3.4 Storage Device Stress—Device Recovery from Device Port Toggle—Extended Run

Test Objective

1. Sequentially toggle each initiator and target port in the fabric.
2. Verify that host I/O recovers to an alternate path and that the toggled path recovers.
3. Run the test for 24 hours.

Test Configuration

1. Set up multipath on the host, and start I/O.
2. Perform multiple iterations of sequential port toggles across initiator and target switch ports.

Result Validation

1. Check the switch port status after the toggle, and check for errors in the switch error logs.

```
# errdumpall
# portstatsshow X
# portshow X
```

2. Check host multipath properties for iSCSI hosts to verify that the toggled path recovers.
   a. Windows: **mpclaim -s -d**
   b. Linux: **multipath -ll**
   c. VMware: Check the paths at **Configuration > Storage > Devices > Manage Paths**
3. Check host and storage error logs, and verify that I/O continues without error.

Test Results

PASS. I/O failed over and recovered successfully. All validation checks passed.

3.5 Storage Device Recovery—ISL Port Toggle (Sequential)

Test Objective

1. Sequentially toggle each ISL path on all switches. Host I/O may pause but should recover.
2. Verify fabric ISL path redundancy between hosts and storage devices.
3. Verify host I/O throughout test.

Test Configuration

1. Set up host multipath with links on different switches in the FC fabric, and start I/O.
2. Ensure ISL redundancy by provisioning multiple ISLs connected to different switches to provide multiple paths through the fabric.

```
# islshow
```

3. Perform multiple iterations of sequential ISL toggles across the fabric.

Result Validation

1. Check FC fabric status after ISL toggle. Verify all nodes are online.

```
# fabricshow
```

2. Check the switch logs for any errors and verify I/O failed over to alternate ISL path in the fabric.

```
# errdumpall
# portperfshow
# porterrshow
```

3. Check host and storage error logs, and verify that I/O continues without error.

Test Results

PASS. I/O re-routes to available paths in the fabric and recovers when the link is restored. All validations checks passed.

3.6 Storage Device Recovery—ISL Port Toggle (Entire Switch)

Test Objective

1. Sequentially, and for all switches, disable all ISLs on the switch under test.
2. Verify fabric switch path redundancy between hosts and storage devices.
3. Verify switch can merge back in to the fabric.
4. Verify host I/O path throughout test.

Test Configuration

1. Set up host multipath with links on different switches in the FC fabric, and start I/O.
2. Ensure ISL redundancy by provisioning multiple ISLs connected to different switches to provide multiple paths through the fabric.

```
# islshow
```

3. Perform multiple iterations of sequentially disabling all ISLs on a switch in the fabric.

Result Validation

1. Check the FC fabric status after ISL toggling. Verify that all nodes are online.

```
# fabricshow
```

2. Check the switch logs for errors, and verify that I/O failed over to an alternate ISL path in the fabric.

```
# errdumpall
# portperfshow
# porterrshow
```

3. Check host and storage error logs, and verify that I/O continues without error.

Test Results

PASS. I/O failed over to alternate path and recovered once the switch merged back in the fabric. All validations checks passed.

3.7 Storage Device Recovery—Director Blade Maintenance

Test Objective

1. Toggle each blade on the director in sequential order.
2. Include blade enable/disable, power on/off, and reboot testing.

Tested with FOS 7.4.1.

Test Configuration

1. Uplink edge switch ISLs to different blades on the directors.
2. Set up host multipath with links on different switches in the FC fabric, and start I/O.
3. Perform multiple iterations of sequential disable/enable, power on/off, and reboot of all blades on the 8510 directors.

Result Validation

1. Check the FC fabric status after the blade toggle. Verify that all nodes are present in the fabric.

```
# fabricshow
```

2. Check the switch logs for errors, and verify that I/O failed over to an alternate ISL path in the fabric.

```
# errdumpall
# portperfshow
# porterrshow
```

3. Check host and storage error logs, and verify that I/O continues without error.

Test Results

PASS. I/O failed over to an alternate path and recovered once the blade recovered from the disruption.

3.8 Storage Device Recovery—Switch Offline

Test Objective

1. Toggle each switch in sequential order.
2. Include switch enable/disable, power on/off, and reboot testing.

Test Configuration

1. Set up host multipath with links on different switches in the FC fabric, and start I/O.
2. Perform multiple iterations of sequential disable/enable, power on/off, and reboot of all switches in the fabric.

Result Validation

1. Check FC fabric status after the switch toggle. Verify all nodes are present in the fabric

```
# fabricshow
```

2. Check the switch logs for any errors and verify the toggled switch has recovered.

```
# errdumpall
# switchshow
```

3. Check host and storage error logs, and verify that I/O continues without error.

Test Results

PASS. I/O failed over to an alternate path and recovered once the switch merged back in the fabric.

3.9 Storage Device Recovery—Switch Firmware Download

Test Objective

1. Sequentially perform firmware maintenance procedure on all device connected switches under test.
2. Verify host I/O will continue (with minimal disruption) through the “firmware download” and device pathing will remain consistent.

Test Configuration

1. Set up host multipath with links on different switches in the FC fabric, and start I/O.
2. Sequentially perform firmware upgrades on all switches in the fabric.

Result Validation

1. Verify firmware upgrade completes successfully on each switch node and they merge back in the FC fabric.

```
# version
# fabricshow
```

2. Check I/O generator tool logs to verify I/O runs without errors throughout the firmware upgrade.
3. Check the switch logs for any errors and verify I/O resumes on the node after the firmware upgrade is complete.

```
# errdumpall
# portperfshow
```

Test Results

PASS. I/O operations completed without error. I/O failed over to an alternate path during the switch reload after the firmware upgrade and resumed after the switch was online. All validation checks passed.

4. Storage Device—Fibre Channel Routing (FCR) Internetworking Tests

(Covered in FOS 8.1.0 and previous releases.)

4.1 Storage Device Internetworking Validation with FC Host

Test Objective

1. Configure two FC fabrics with FCR.
2. Verify that edge devices are imported into adjacent name servers and that hosts have access to their routed targets after FC routers are configured.

Test Configuration

1. Set up FCR in an Edge-Backbone-Edge configuration.
2. Set up LSAN zoning, verify host access to target LUNs, and start I/O.

Result Validation

1. Verify the name server and FCR fabric state.

```
# fcrfabricshow
# fcrproxydevshow
```

2. Verify that I/O runs successfully without error.

Test Results

PASS. I/O completed successfully. Both edge fabrics have the corresponding proxy name server entries for the host and target ports.

4.2 Storage Device Internetworking Validation with FCoE Using VDX FlexPort

Test Objective

1. Configure an FC fabric with FCR connected to an FCoE fabric.
2. Verify that edge devices are imported into adjacent name servers and that hosts have access to their routed targets after FC routers are configured.

Tested with FOS 7.4.1.

Test Configuration

1. Add the FCoE VCS fabric to the FCR setup.
2. Set up LSAN zoning, verify host access to target LUNs, and start I/O.

Result Validation

1. Verify name server and FCR fabric state.

```
# fcrfabricshow
# fcrproxydevshow
```

2. Verify that I/O runs successfully without error.

Test Results

PASS. I/O completed successfully. Both edge fabrics have the corresponding proxy name server entries for the host and target ports.

4.3 Storage Device Edge Recovery After FCR Disruptions

Test Objective

1. Configure FCR for an Edge-Backbone-Edge configuration.
2. With I/O running, validate device access and pathing.
3. Perform reboots, switch disables, and port toggles on edge connections to disrupt device pathing and I/O.
4. Verify path and I/O recovery once switches and ports recover.

Test Configuration

1. Set up FCR in an Edge-Backbone-Edge configuration.
2. Set up LSAN zoning, verify host access to target LUNs, and start I/O.
3. Perform sequential reboots, switch disables, and ISL port toggles on the switches in the edge fabric.

Result Validation

1. Verify the FCR fabric state throughout the disruptions.

```
# fcrfabricshow
# fcrproxydevshow
```

2. Check the switch logs for errors.

```
# errdumpall
# portperfshow
```

3. Check host and storage logs, and verify that I/O runs without error.

Test Results

PASS. I/O fails over to available switch path and recovers when the disrupted switch is restored.

4.4 Storage Device Backbone Recovery After FCR Disruptions

Test Objective

1. Configure FCR for a Backbone-Edge configuration.
2. With I/O running, validate device access and pathing.
3. Perform reboots, switch disables, and port-toggles on backbone connections to disrupt device pathing and I/O.
4. Verify path and I/O recovery once the switches and ports recover.

Test Configuration

1. Connect array target ports to the backbone fabric in an Edge-Backbone configuration.
2. Set up LSAN zoning, verify host access to target LUNs, and start I/O.
3. Perform sequential reboots, switch disables, and ISL port toggles on the switches in the backbone fabric.

Result Validation

1. Verify FCR fabric state throughout the disruptions.

```
# fcrfabricshow
# fcrproxydevshow
```

2. Check the switch logs for errors.

```
# errdumpall
# portperfshow
```

3. Check the host and storage logs, and verify that I/O runs without error.

Test Results

PASS. I/O fails over to the available switch path and recovers when the disrupted switch is restored.

5. Optional/Additional Tests

5.1 Non-disruptive Firmware Upgrade on Storage Device

Test Objective

1. Perform firmware maintenance procedure on the storage device.
2. Verify that host I/O will continue (with minimal disruption) through the “firmware download” and that device pathing will remain consistent.

Test Configuration

1. Set up host multipath with links on different switches in the FC fabric, and start I/O.
2. Perform firmware update on all nodes of the storage array.

Result Validation

1. Check the I/O generator tools logs to verify I/O completes without any error.
2. Check the host and storage logs for any errors throughout the I/O operations.
3. Check the switch error logs and port stats for any errors or I/O drops.

```
# errdumpall
# porterrshow
```

Test Results

PASS. I/O completed successfully throughout the firmware upgrade process.

5.2 Workload Simulation Test Suite—Medusa

Test Objective

1. Validate Storage/Fabric behavior while running a workload simulation test suite.
2. Areas of focus may include random and sequential data patterns of various block sizes and database simulation.

Test Configuration

1. Set up 4 standalone hosts with 2 multipathed initiator ports for I/O generation.
2. Use the Medusa I/O tool for generating I/O and simulating workloads.
3. Run random and sequential I/O in a loop at block transfer sizes of 512, 4k, 8k, 16k, 32k, 64k, 128k, 256k, 512k, and 1m. Include a nested loop of 100% read, 100% write, and 50% read/write.
4. Run the Medusa Application I/O workload suite, which includes OLTP, Decision Support System (DSS), Exchange Email, File Servers, Media Streaming, OS Drive, OS Paging, SQL, Video on Demand, VDI, and Web Server profiles.

Result Validation

1. Check the I/O generator tools logs to verify I/O completes without any errors.
2. Check the host and storage logs for any errors throughout the I/O operations.
3. Check the switch error logs and port stats for any errors or I/O drops.

```
# errdumpall
# porterrshow
```

Test Results

PASS. All workload runs were monitored at the host, storage, and fabric and completed without I/O error or fault.

5.3 Workload Simulation Test Suite—VMware

Test Objective

1. Validate Storage/Fabric behavior while running a virtual workload simulation test suite.
2. Areas of focus include VM environments running de-duplication/compression data patterns, and database simulation.

Test Configuration

1. Set up an ESX cluster of two hosts with four worker VMs per host.
2. Use VMware IOAnalyzer tool for generating I/O and simulating workloads.
   a. Run random and sequential IO at large and small block transfer sizes.
   b. Run SQL Server simulation workload.
   c. Run OLTP simulation workload.
   d. Run Web Server simulation workload.
   e. Run Video on Demand simulation workload.
   f. Run Workstation simulation workload.
   g. Run Exchange server simulation workload.
3. Set up VMs with Medusa I/O tool and run random and sequential I/O in a loop at block transfer sizes of 512, 4k, 8k, 16k, 32k, 64k, 128k, 256k, 512k, and 1m. Include a nested loop of 100% read, 100% write, and 50% read/write.
4. Run the Medusa Application I/O workload suite, which includes OLTP, Decision Support System (DSS), Exchange Email, File Servers, Media Streaming, OS Drive, OS Paging, SQL, Video on Demand, VDI, and Web Server profiles.

Result Validation

1. Check the I/O generator tools logs to verify I/O completes without any errors.
2. Check the host and storage logs for any errors throughout the I/O operations.
3. Check the switch error logs and port stats for any errors or I/O drops.

```
# errdumpall
# porterrshow
```

Test Results

PASS. All workload runs were monitored at the host, storage, and fabric and completed without I/O error or fault.

Test Conclusions

1. Achieved 100% pass rate on all the test cases in the SFR qualification test plan. The network and the storage were able to handle the various stress and error recovery scenarios without any issues.
2. Different I/O workload scenarios were simulated using Medusa and VMware I/OAnalyzer tools and sustained performance levels were achieved across all workload types. The Tegile T3100 array and the Brocade FC fabric handled both the low latency and high throughput I/O workloads with equal efficiency without any I/O errors or packet drops.
3. The results confirm that the Tegile T3100 array interoperates seamlessly with Brocade FC fabrics, and demonstrate high availability and sustained performance.
4. The Brocade Gen 5 (16Gb) and Gen 6 (32Gb) FC switches were able to handle the sustained throughput and latency performance requirements efficiently with fewer ISL trunks. Multiple ISLs to different switches in the fabric should be set up for providing path redundancy through the fabric.
5. It is recommended to enable the Monitoring and Alerting Policy Suite (MAPS) heath monitor on all switches in the FC fabric to report fabric-wide events and traffic performance metrics. The MAPS feature of Fabric Performance Impact monitoring is enabled by default and should be used to detect fabric bottlenecks in the form of timeouts and latency.
6. It is recommended to implement the I/O Insight feature with MAPS alerting to closely monitor the critical flows in the fabric at the SCSI level.
7. Utilizing Peer Zoning helps reduce the zone database size and the zoning complexity, while providing the RSCN and hardware resource efficiencies of Single-Initiator Zoning.
8. QoS Zoning should be used to classify host-target traffic into high, medium, or low priority zones; to provide traffic prioritization through the FC fabric for the desired host-target pair by allocating more resources to the traffic in the higher priority zone.
9. Enabling Emulex ExpressLane on a LUN provides prioritized queuing on the HBA for traffic to that LUN and also sets the CS_CTL tag on the frame which allows the traffic to be prioritized through the FC fabric based on the value of the CS_CTL tag and the corresponding priority level.
10. Host multipath should be configured for optimal availability and performance. Multipath configuration details for the hosts are provided in the host setup section of the “Configure DUT and Test Equipment” section.