BROCADE

VALIDATION TEST REPORT

Brocade Fabric Technology with the Kaminario K2 All-Flash Array Validation Test Report

Supporting the Kaminario K2 All-Flash Array with Fabric OS 8.0.1

53-1004539-01
29 July 2016

Contents

Preface

Document History

Date	Part Number	Description
July 29, 2016	53-1004539-01	Initial release.

Overview

The Storage Fabric Ready (SFR) program is a comprehensive testing and configuration initiative to provide Fibre Channel (FC) SAN and IP interoperability with flash storage. This program provides testing of multiple fabrics, heterogeneous servers, NICs, and HBAs in large-port-count Brocade environments. Testing covers future Brocade OS versions and the vendor's software releases.

The SFR qualification program helps verify seamless interoperability and optimum performance with solid-state storage, software-defined storage, and hyper-converged systems in Brocade storage fabrics.

Purpose of This Document

This document provides the validation of Brocade fabric technology with the Kaminario K2 Fibre Channel storage array, using multiple switch platforms, HBAs, and server operating systems. This validation shows that the K2 interoperates properly within Brocade Gen 5 and Gen 6 Fibre Channel fabrics, while supporting the performance and low latency associated with solid-state storage.

Audience

This document is written for a technical audience, including solution architects, system engineers, and technical development representatives.

Objectives

- Test the Kaminario array with Brocade FC fabrics, in single and routed configurations, for different stress and error recovery scenarios and validate the interoperability and integration of the array with Brocade FC fabrics.
- Validate the performance of the FC fabric in a solid-state storage environment for high-throughput and low-latency applications.

Related Documents

- Brocade Fabric OS Administration Guide, 8.0.1
- Brocade SAN Design and Best Practices
- Brocade Monitoring and Alerting Policy Suite Configuration Guide, 8.0.1
- Brocade Network OS Layer 2 Switching Configuration Guide, 7.0.0
- Brocade Flow Vision Configuration Guide, 8.0.1
- Emulex OneCommand Manager User Manual
- QLogic QLE2672 Adapter User Guide
- QLogic BR-1860 Adapter Administrator's Guide

About Brocade

Brocade® (NASDAQ: BRCD) networking solutions help the world's leading organizations transition smoothly to a world where applications and information reside anywhere. This vision is designed to deliver key business benefits such as unmatched simplicity, non-stop networking, application optimization, and investment protection.

Innovative Ethernet and storage networking solutions for data center, campus, and service provider networks help reduce complexity and cost while enabling virtualization and cloud computing to increase business agility.

To help ensure a complete solution, Brocade partners with world-class IT companies and provides comprehensive education, support, and professional services offerings (www.brocade.com).

About Kaminario

Kaminario, the leading all-flash storage company, is redefining the future of modern data centers. Its unique solution enables organizations to succeed in today's on-demand world and prepares them to seamlessly handle tomorrow's innovations. Only the Kaminario K2 all-flash array delivers the agility, scalability, performance, and economics that a data center requires to deal with today's cloud-first, dynamic world, and it provides real-time data access—anywhere, anytime. Hundreds of customers rely on Kaminario K2 to power their mission-critical applications and safeguard their digital ecosystem.

Headquartered in Needham, MA, Kaminario works with an extensive network of resellers and distributors, globally. For more information, please visit www.kaminario.com.

Configure DUT and Test Equipment

Task 1. Brocade FC Fabric Configuration

Configuration settings of the Brocade switches in the test bed are covered here.

1. Zoning is configured using the Peer Zoning feature. Peer zoning allows a "principal" device to communicate with the other 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.

```
root> zonecreate -peerzone kaminario_peer -principal "50:02:4f:40:55:82:01:00; 50:02:4f:
40:55:82:01:01;
50:02:4f:40:55:82:02:00; 50:02:4f:40:55:82:02:01" -members "10:00:8c:7c:ff:24:4c:00; 10:00:8c:7c:ff:
24:4c:01 …"

root> zoneshow --peerzone all
zone:  kaminario_peer
   Property Member: 00:02:00:00:00:03:00:08
   Created by: User
   Principal Member(s):
                50:02:4f:40:55:82:01:00; 50:02:4f:40:55:82:01:01;
50:02:4f:40:55:82:02:00; 50:02:4f:40:55:82:02:01   Peer Member(s):
                10:00:8c:7c:ff:24:4c:00; 10:00:8c:7c:ff:24:4c:01
                10:00:00:90:fa:61:92:3b; 10:00:00:90:fa:61:92:3c;
                10:00:8c:7c:ff:05:60:01; 10:00:8c:7c:ff:05:60:00;
                10:00:8c:7c:ff:14:e0:01; 10:00:8c:7c:ff:14:e0:00;
                10:00:8c:7c:ff:03:bc:01; 10:00:8c:7c:ff:03:bc:00;
                10:00:8c:7c:ff:03:9b:00; 10:00:8c:7c:ff:03:9b:01;
                10:00:8c:7c:ff:05:72:02; 10:00:8c:7c:ff:05:72:03
```

2. Configure MAPS on switches (requires a Fabric Vision license). This will enable the reporting of latency and congestion alerts on each switch, as well as a number of switch and fabric health metrics. For more information on configuring and using MAPS, see the Brocade Monitoring and Alerting Policy Suite Configuration Guide, 8.0.1.

   a) Enable the desired MAPS policy using any of the available default policies or create a custom policy. For this test, the default "Aggressive Policy" is used. This provides the most sensitive threshold levels for detection of latency and congestion.

```
> mapsconfig --enablemaps -policy dflt_aggressive_policy
```

   b) Define the reporting actions that MAPS will take. In this test, errors discovered by MAPS will generate a RASLog entry and send email to the configured recipients.

```
>  mapsconfig --actions raslog,email
```

   c) To examine the effective MAPS configuration, use the following command:

```
> mapsconfig -show
Configured Notifications:       RASLOG,EMAIL
Mail Recipient:                 testuser1@domain.com,testuser2@domain.com
FPI Monitoring:                 Enabled
Paused members :
===============
PORT :
CIRCUIT :
SFP :
```

d) Use the following command to view the summary of the events or rules triggered and the objects on which the rules were triggered over a specified period of time.

```
> mapsdb –show
1 Dashboard Information:
=======================
G620_066_223:FID128:root> mapsdb --show

1 Dashboard Information:
=======================

DB start time:                  Thu May 19 09:27:24 2016
Active policy:                  ios_mod_policy
Configured Notifications:       RASLOG,EMAIL,SW_CRITICAL,SW_MARGINAL
Fenced Ports :                  None
Decommissioned Ports :          None
```

3. Configure Flow Monitoring with the I/O Insight feature on the Gen 6 (32-G FC) switch in the fabric.

a) The I/O Insight feature supported on the Gen 6 hardware 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 pizza-box type switches and at an ITL (Initiator-Target-LUN) flow level on chassis-based switches.

- Requires a Fabric Vision and I/O Insight license.

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

```
> flow --create ios_kaminario_1 -dstdev df2500 -egrport 37 -feature mon

root> flow --show ios_kaminario_1
================================================================================
Name      : ios_kaminario_1     Features: mon(Activated)        noConfig: Off
Definition: EgrPort(37),DstDev(0xdf2500)


Flow Monitor (Activated):
Monitor time:  | Wed May 25 16:48:16 MDT 2016 |
```

b) Import the created flows into MAPS.

```
> mapsconfig --import ios_kaminario_1

> logicalgroup --show
 ------------------------------------------------------------------------------
Group Name           |Predefined |Type        |Member Count |Members
------------------------------------------------------------------------------
ios_kaminario_1      |No         |Flow        |1            |Monitored Flow
```

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

```
> mapspolicy --clone dflt_moderate_policy -name ios_mod_policy
> mapsrule --create ios_k1 -group ios_kaminario_1 -monitor RD_STATUS_TIME_LT_8K -timebase min -
op g -value 750 -action raslog,email -policy ios_mod_policy

> mapspolicy --show ios_mod_policy
Policy Name: ios_mod_policy

Rule Name                       |Condition                                          |Actions|
--------------------------------------------------------------------
ios_Kaminario_small_read_P1     |ios_kaminario_1(RD_STATUS_TIME_LT_8K/min>750)   |
raslog,email
```

d) Enable the MAPS policy.

```
> mapspolicy --enable ios_mod_policy

> mapsdb --show
```

```
1 Dashboard Information:
========================

DB start time:                  Wed May 25 15:46:57 2016
Active policy:                  ios_mod_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)>: 0xdf2000(30) 0xdf2100(30) 0xdf1400(28) 0xdf0400(28) 0xdf1a00(26)
```

4. Configure FlexPorts on VDX 6740 switches to Fibre Channel support (requires an FCoE Base license).
   a) In this test, these ports are used as uplinks from the VCS fabric to the Fibre Channel fabric. For detailed information on configuring VDX 6740 switches for Fibre Channel support, see the Network OS Layer 2 Switching Configuration Guide, 7.0.0.

      Example of configuring a Fibre Channel port on a Brocade VDX 6740 switch:

      ```
      # conf t
      (config)# hardware
      (config-hardware)# flexport 112/0/38
      (config-flexport-112/0/38)# type fibre-channel
      ```

   b) FlexPort speeds are configured in connector groups. In this test bed example, the speed HighMixed is chosen. This setting allows the Fibre Channel FlexPort to operate at 16-Gb speed, which is optimal because the uplink is attaching to a 16-Gb Fibre Channel SAN. Also, not all ports in this connector group are configured as Fibre Channel ports; some are still configured as Ethernet ports, and the HighMixed setting will support the Ethernet speeds on these ports as well. For detailed information on configuring VDX 6740 connector-group speeds, see the Network OS Layer 2 Switching Configuration Guide, 7.0.0.

      - LowMixed: 2/4/8G Fibre Channel and Ethernet speeds (default)
      - HighMixed: 16G Fibre Channel and Ethernet speeds
      - FibreChannel: 2/4/8/16G Fibre Channel

      Use this command to set the connector-group speed to HighMixed:

      ```
      (config-hardware)# connector-group 111/0/3
      (config-connector-group-111/0/3)# speed HighMixed
      ```

5. Configure zones for FCoE initiators on VDX switches.

   Example zone on the Brocade VDX 6740 switch:

   ```
   # show zoning enabled-configuration
   zoning enabled-configuration cfg-name NOS_SSR

   zoning enabled-configuration enabled-zone lsan_ssr067174_kaminario
    member-entry 10:00:00:05:33:48:72:c0
    member-entry 10:00:00:05:33:48:72:c1
    member-entry 50:02:4f:40:55:82:01:00
    member-entry 50:02:4f:40:55:82:01:01
    member-entry 50:02:4f:40:55:82:01:02
    member-entry 50:02:4f:40:55:82:01:03
   ```

6. Configure Fibre Channel Routing (requires an Integrated Routing license). Detailed information on FCR setup can be found in the Brocade Fabric OS Administration Guide, 8.0.1.
   a) Example FCR configuration is shown below:

      ```
      > fcrconfigure -bbfid 100
      > fosconfig --enable fcr
      > portcfgexport [port#] -a1 -m[0/5] -f 10  {0=Brocade FC fabric; 5=Brocade NOS fabric}
      ```

      The prefix "lsan" is used when configuring zones for use in Fibre Channel Routing.

Example of zone prefixed with "lsan":

```
> zoneshow lsan_ssr067116_kaminario
 zone:  lsan_ssr067116_kaminario
                10:00:8c:7c:ff:5b:52:00; 10:00:8c:7c:ff:5b:52:01;
                kaminario_ali
```

b) Example output of exported devices:

```
> fcrproxydevshow
  Proxy           WWN             Proxy      Device    Physical    State
 Created                           PID       Exists      PID
in Fabric                                   in Fabric
-------------------------------------------------------------------------
    40   50:02:4f:40:55:82:01:00  02f906      100       532400   Imported
    40   50:02:4f:40:55:82:01:01  02fa06      100       df2400   Imported
    40   50:02:4f:40:55:82:02:00  02fb06      100       532500   Imported
    40   50:02:4f:40:55:82:02:01  02fc06      100       df2500   Imported
    50   50:02:4f:40:55:82:01:00  02fc03      100       532400   Imported
    50   50:02:4f:40:55:82:01:01  02fd03      100       df2400   Imported
    50   50:02:4f:40:55:82:02:00  02fe03      100       532500   Imported
    50   50:02:4f:40:55:82:02:01  02ff03      100       df2500   Imported
```

Task 2. Kaminario Array Configuration

Configuration steps for the Kaminario array are covered here.

1. Create a host initiator group for each host, and add WWPNs for the host:

FIGURE 1 Initiator Configuration on the Kaminario Array

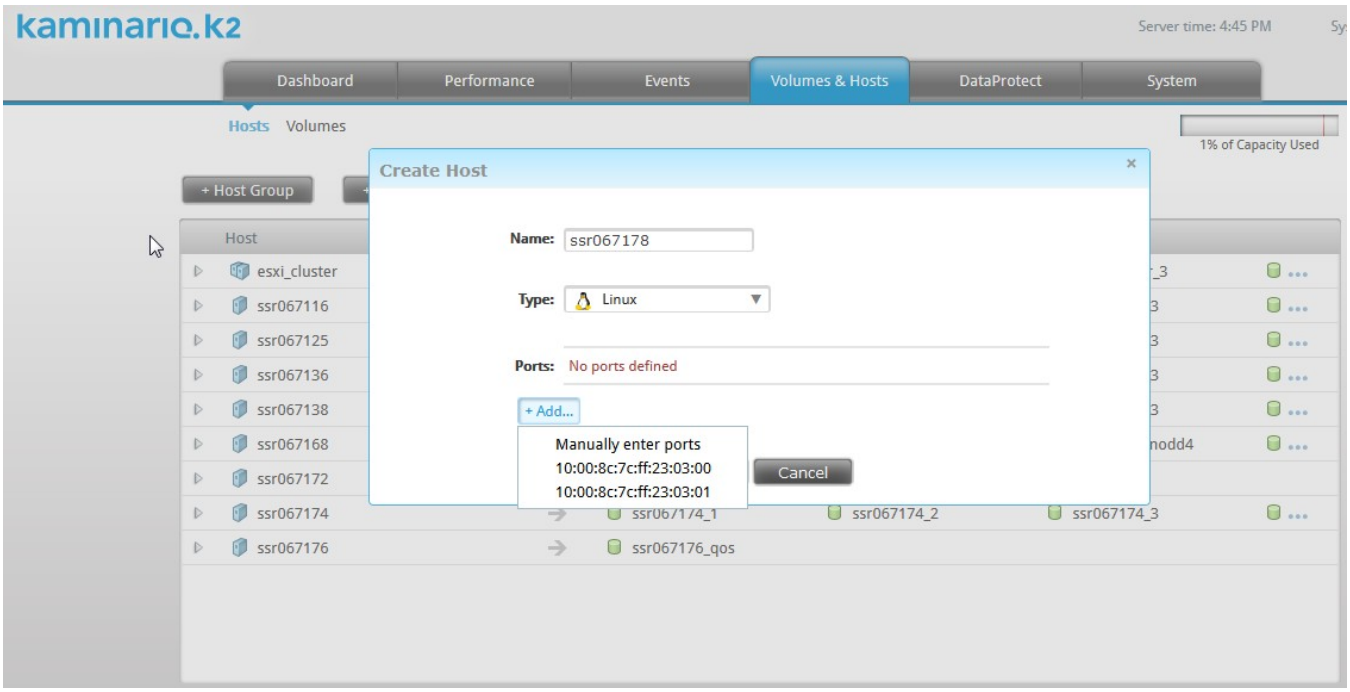

2. Create new volumes, and assign them to the host initiator group. For this test configuration, 8 x 20GB volumes are created.

FIGURE 2 Finished LUN Configuration and Mapping on the Kaminario Array

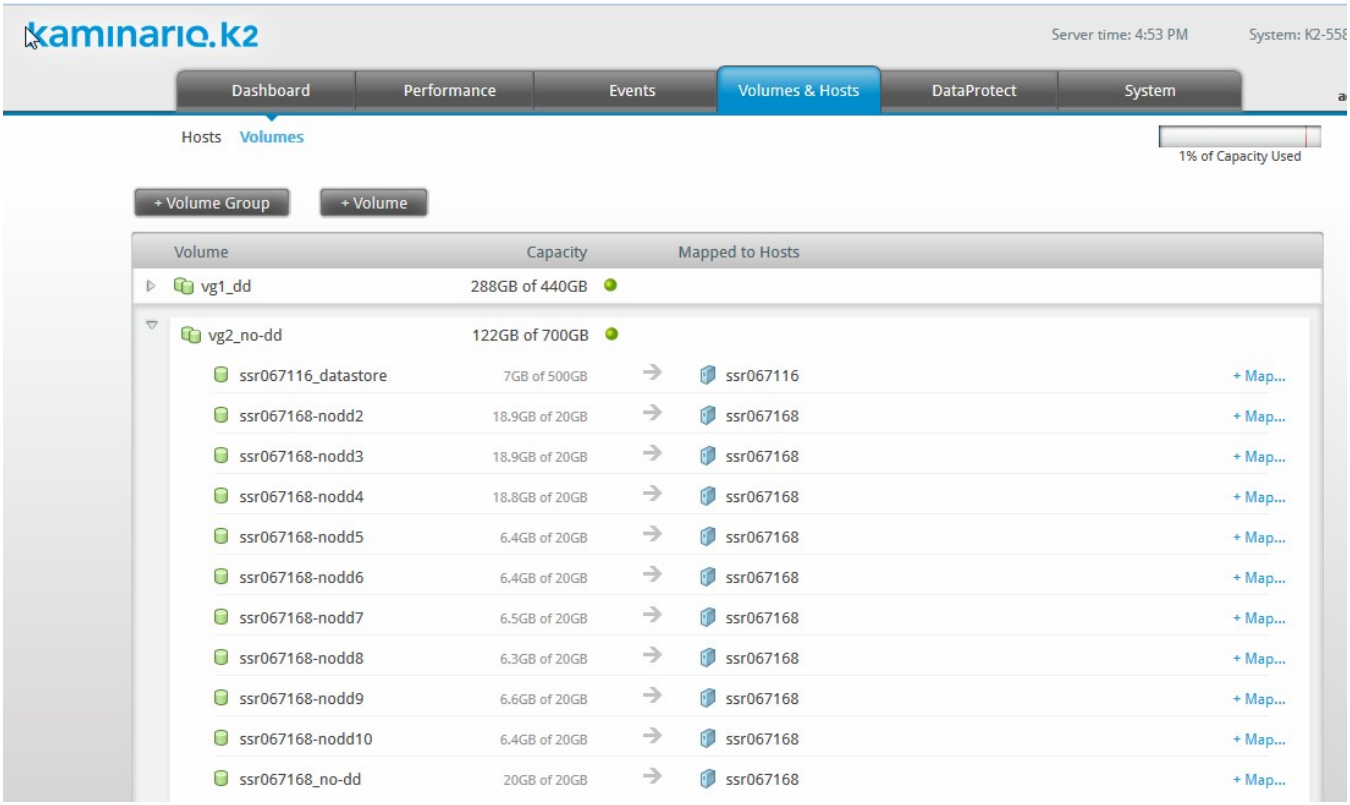

Task 3. Host Setup

Configuration settings of servers in the test bed are covered here.

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. Multipath configuration on Linux hosts.

   This configuration allows all paths to be used in a round-robin fashion. This provides superior performance to the default Linux settings, which only use a single active path per LUN.

   a) Recommended /etc/multipath.conf entry on Linux systems:

```
device {
vendor "KMNRIO"
product "K2"
path_grouping_policy multibus
# getuid_callout "/lib/udev/scsi_id --whitelisted --device=/dev/%n"
# the above line is deprecated; use uid_attribute
uid_attribute "ID_SERIAL"
path_checker tur
path_selector "queue-length 0"
no_path_retry fail
hardware_handler "0"
```

```
rr_weight priorities
rr_min_io 1
failback 15
fast_io_fail_tmo 5
dev_loss_tmo 8
```

b) Example multipath configuration on Linux:

```
# multipath -ll

mpathz (20024f40055820043) dm-18 KMNRIO  ,K2
size=20G features='0' hwhandler='0' wp=rw
`-+- policy='queue-length 0' prio=1 status=active
  |- 1:0:0:66  sddd 70:176  active ready running
  |- 1:0:1:66  sddj 71:16   active ready running
  |- 1:0:2:66  sddp 71:112  active ready running
  |- 1:0:3:66  sddv 71:208  active ready running
  |- 10:0:0:66 sdeb 128:48  active ready running
  |- 10:0:1:66 sdeh 128:144 active ready running
  |- 10:0:2:66 sden 128:240 active ready running
  `- 10:0:3:66 sdet 129:80  active ready running
mpathy (20024f40055820042) dm-17 KMNRIO  ,K2
size=20G features='0' hwhandler='0' wp=rw
`-+- policy='queue-length 0' prio=1 status=active
  |- 1:0:0:65  sddc 70:160  active ready running
  |- 1:0:1:65  sddi 71:0    active ready running
  |- 1:0:2:65  sddo 71:96   active ready running
  |- 1:0:3:65  sddu 71:192  active ready running
  |- 10:0:0:65 sdea 128:32  active ready running
  |- 10:0:1:65 sdeg 128:128 active ready running
  |- 10:0:2:65 sdem 128:224 active ready running
  `- 10:0:3:65 sdes 129:64  active ready running
mpathx (20024f40055820041) dm-16 KMNRIO  ,K2
size=20G features='0' hwhandler='0' wp=rw
`-+- policy='queue-length 0' prio=1 status=active
  |- 1:0:0:64  sddb 70:144  active ready running
  |- 1:0:1:64  sddh 70:240  active ready running
  |- 1:0:2:64  sddn 71:80   active ready running
  |- 1:0:3:64  sddt 71:176  active ready running
  |- 10:0:0:64 sddz 128:16  active ready running
  |- 10:0:1:64 sdef 128:112 active ready running
  |- 10:0:2:64 sdel 128:208 active ready running
  `- 10:0:3:64 sder 129:48  active ready running
mpathw (20024f4005582003f) dm-13 KMNRIO  ,K2
size=20G features='0' hwhandler='0' wp=rw
`-+- policy='queue-length 0' prio=1 status=active
  |- 1:0:0:62  sdcw 70:64   active ready running
  |- 1:0:1:62  sdca 68:224  active ready running
  |- 1:0:2:62  sday 67:32   active ready running
  |- 1:0:3:62  sdan 66:112  active ready running
  |- 10:0:0:62 sdda 70:128  active ready running
  |- 10:0:1:62 sdam 66:96   active ready running
  |- 10:0:2:62 sdcn 69:176  active ready running
  `- 10:0:3:62 sdba 67:64   active ready running
mpathac (20024f40055820040) dm-21 KMNRIO  ,K2
size=20G features='0' hwhandler='0' wp=rw
`-+- policy='queue-length 0' prio=1 status=active
  |- 1:0:0:63  sddg 70:224  active ready running
  |- 1:0:1:63  sddm 71:64   active ready running
  |- 1:0:2:63  sdds 71:160  active ready running
  |- 1:0:3:63  sddy 128:0   active ready running
  |- 10:0:0:63 sdee 128:96  active ready running
  |- 10:0:1:63 sdek 128:192 active ready running
  |- 10:0:2:63 sdeq 129:32  active ready running
  `- 10:0:3:63 sdew 129:128 active ready running
mpathv (20024f4005582003d) dm-12 KMNRIO  ,K2
size=20G features='0' hwhandler='0' wp=rw
`-+- policy='queue-length 0' prio=1 status=active
  |- 1:0:0:61  sdcu 70:32   active ready running
  |- 1:0:1:61  sdbz 68:208  active ready running
  |- 1:0:2:61  sdaw 67:0    active ready running
```

```
  |- 1:0:3:61  sdal 66:80   active ready running
  |- 10:0:0:61 sdcz 70:112  active ready running
  |- 10:0:1:61 sdak 66:64   active ready running
  |- 10:0:2:61 sdcm 69:160  active ready running
  `- 10:0:3:61 sdaz 67:48   active ready running
mpathab (20024f40055820045) dm-20 KMNRIO  ,K2
size=20G features='0' hwhandler='0' wp=rw
`-+- policy='queue-length 0' prio=1 status=active
  |- 1:0:0:68  sddf 70:208  active ready running
  |- 1:0:1:68  sddl 71:48   active ready running
  |- 1:0:2:68  sddr 71:144  active ready running
  |- 1:0:3:68  sddx 71:240  active ready running
  |- 10:0:0:68 sded 128:80  active ready running
  |- 10:0:1:68 sdej 128:176 active ready running
  |- 10:0:2:68 sdep 129:16  active ready running
  `- 10:0:3:68 sdev 129:112 active ready running
mpathi (20024f40055820000) dm-3 KMNRIO  ,K2
size=256K features='0' hwhandler='0' wp=rw
`-+- policy='queue-length 0' prio=1 status=active
  |- 1:0:0:0   sdbb 67:80   active ready running
  |- 1:0:1:0   sdbo 68:32   active ready running
  |- 1:0:2:0   sdb  8:16    active ready running
  |- 1:0:3:0   sdp  8:240   active ready running
  |- 10:0:0:0  sdbe 67:128  active ready running
  |- 10:0:1:0  sdo  8:224   active ready running
  |- 10:0:2:0  sdcb 68:240  active ready running
  `- 10:0:3:0  sdc  8:32    active ready running
mpathu (20024f4005582003e) dm-11 KMNRIO  ,K2
size=20G features='0' hwhandler='0' wp=rw
`-+- policy='queue-length 0' prio=1 status=active
  |- 1:0:0:60  sdcs 70:0    active ready running
  |- 1:0:1:60  sdby 68:192  active ready running
  |- 1:0:2:60  sdau 66:224  active ready running
  |- 1:0:3:60  sdaj 66:48   active ready running
  |- 10:0:0:60 sdcy 70:96   active ready running
  |- 10:0:1:60 sdai 66:32   active ready running
  |- 10:0:2:60 sdcl 69:144  active ready running
  `- 10:0:3:60 sdax 67:16   active ready running
mpathaa (20024f40055820044) dm-19 KMNRIO  ,K2
size=20G features='0' hwhandler='0' wp=rw
`-+- policy='queue-length 0' prio=1 status=active
  |- 1:0:0:67  sdde 70:192  active ready running
  |- 1:0:1:67  sddk 71:32   active ready running
  |- 1:0:2:67  sddq 71:128  active ready running
  |- 1:0:3:67  sddw 71:224  active ready running
  |- 10:0:0:67 sdec 128:64  active ready running
  |- 10:0:1:67 sdei 128:160 active ready running
  |- 10:0:2:67 sdeo 129:0   active ready running
  `- 10:0:3:67 sdeu 129:96  active ready running
mpaths (20024f4005582003c) dm-9 KMNRIO  ,K2
size=20G features='0' hwhandler='0' wp=rw
`-+- policy='queue-length 0' prio=1 status=active
  |- 1:0:0:59  sdcq 69:224  active ready running
  |- 1:0:1:59  sdbx 68:176  active ready running
  |- 1:0:2:59  sdas 66:192  active ready running
  |- 1:0:3:59  sdah 66:16   active ready running
  |- 10:0:0:59 sdcx 70:80   active ready running
  |- 10:0:1:59 sdag 66:0    active ready running
  |- 10:0:2:59 sdck 69:128  active ready running
  `- 10:0:3:59 sdav 66:240  active ready running
```

3. Multipath configuration on Windows hosts.

   Use the Windows MPIO dialog to discover and manage multipath entries for Kaminario.

FIGURE 3 Windows MPIO Multipathing with the Kaminario Array

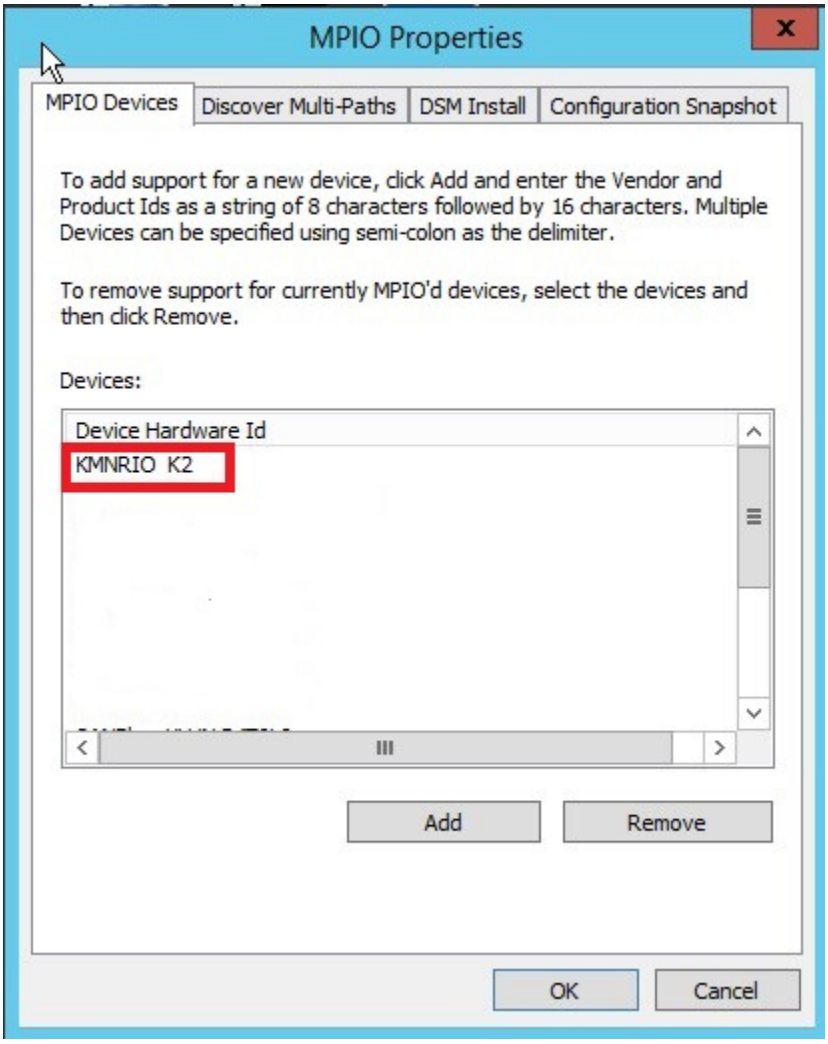

4. Configure multipathing on VMware hosts.

   Change path selection to Round Robin. This provides superior performance to the Most Recently Used setting, which only uses a single active path per LUN.

FIGURE 4 Round Robin Multipath Configuration on VMware

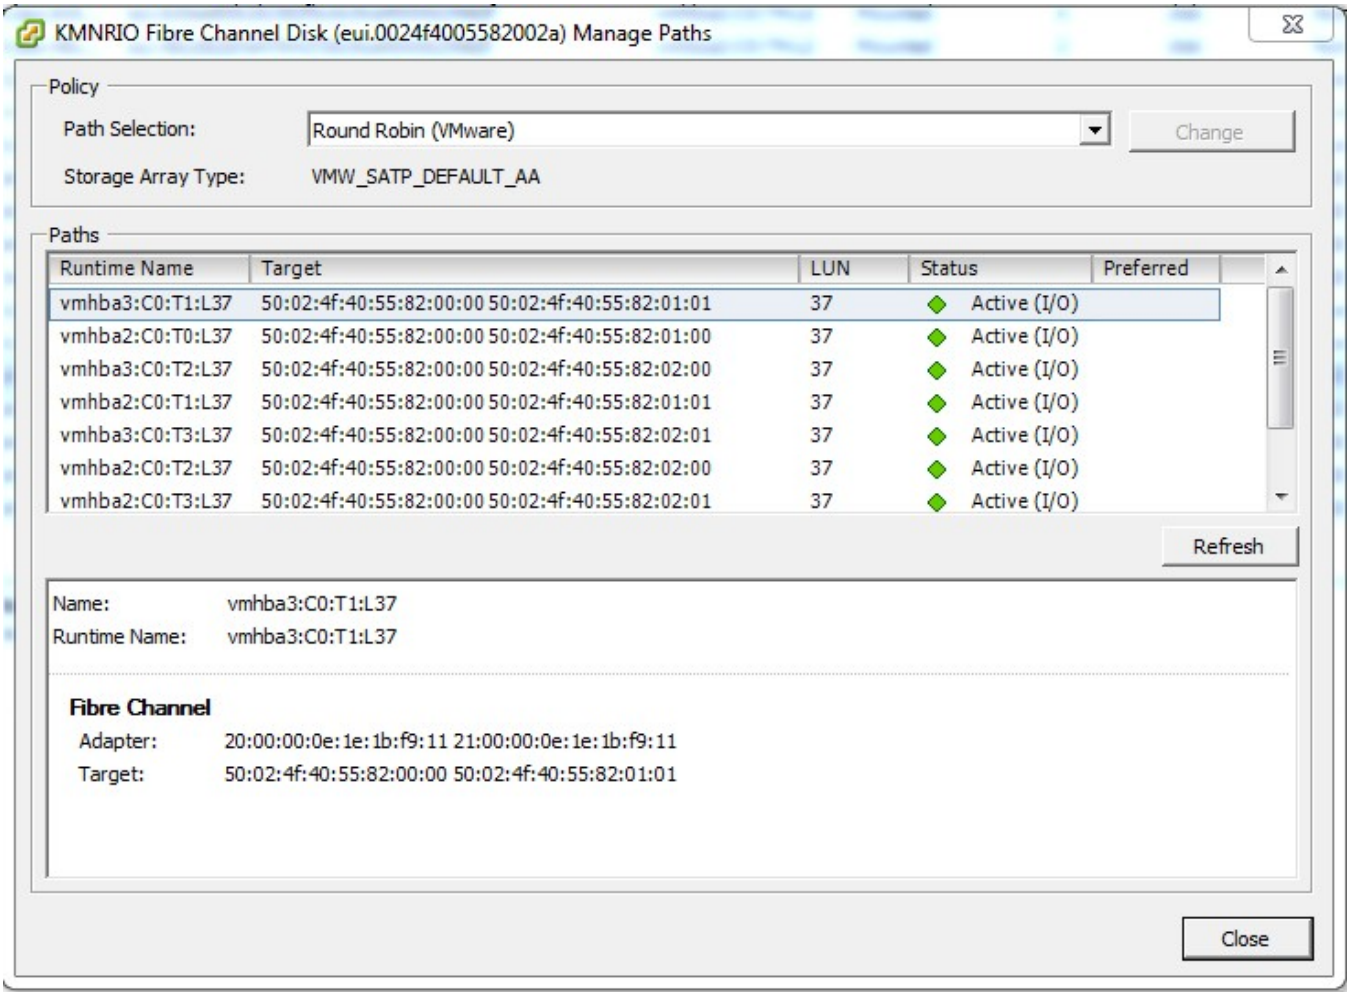

a) Kaminario provides a set of vSphere PowerCLI scripts for easy tuning of ESXi hosts for various parameters, including automatically changing the default path selector in the above example. In our test bed, we ran K2_ESXi6_BP_Activator.ps1 to implement the tuning and K2_ESXi6_BP_Validator.ps1 to validate settings.

b) Sample output from the validation script:

```
Kaminario K2 validation script running version 1.3.
Running validation for ESXi Cluster(s) "cluster_2" from vCenter "10.38.53.191".

_____________________________________________
Inititalization - OK: Found PowerCLI version 6.3 Build 3737840 Revision 0

_____________________________________________
10.38.67.192 - Processing host 10.38.67.192
10.38.67.192 - ESXi version 6.0.0 build 2494585

_____________________________________________
10.38.67.192 - Running validation for Disk.SchedQuantum
10.38.67.192 - OK: Disk.SchedQuantum is properly configured according to Kaminario's BP
(Disk.SchedQuantum=64)

_____________________________________________
10.38.67.192 - Running validation for Disk.DiskMaxIOSize
10.38.67.192 - OK: Disk.DiskMaxIOSize is properly configured according to Kaminario's BP
(Disk.DiskMaxIOSize=1024)

_____________________________________________
```

```
10.38.67.192 - Running validation for Disk.SchedNumReqOutstanding
10.38.67.192 - Note - In ESXi version 6.0.0 the SchedNumReqOutstanding parameter is set per disk.
10.38.67.192 - OK: Kaminario Disk eui.0024f40055820000 is properly configured according to
Kaminario's BP (NoofoutstandingIOswithcompetingworlds=256).
10.38.67.192 - OK: Kaminario Disk eui.0024f40055820027 is properly configured according to
Kaminario's BP (NoofoutstandingIOswithcompetingworlds=256).
10.38.67.192 - OK: Kaminario Disk eui.0024f40055820028 is properly configured according to
Kaminario's BP (NoofoutstandingIOswithcompetingworlds=256).
10.38.67.192 - OK: Kaminario Disk eui.0024f40055820029 is properly configured according to
Kaminario's BP (NoofoutstandingIOswithcompetingworlds=256).
10.38.67.192 - OK: Kaminario Disk eui.0024f4005582002a is properly configured according to
Kaminario's BP (NoofoutstandingIOswithcompetingworlds=256).
10.38.67.192 - OK: Kaminario Disk eui.0024f4005582002b is properly configured according to
Kaminario's BP (NoofoutstandingIOswithcompetingworlds=256).
10.38.67.192 - OK: Kaminario Disk eui.0024f4005582002c is properly configured according to
Kaminario's BP (NoofoutstandingIOswithcompetingworlds=256).
10.38.67.192 - OK: Kaminario Disk eui.0024f4005582002d is properly configured according to
Kaminario's BP (NoofoutstandingIOswithcompetingworlds=256).
10.38.67.192 - OK: Kaminario Disk eui.0024f4005582002e is properly configured according to
Kaminario's BP (NoofoutstandingIOswithcompetingworlds=256).
10.38.67.192 - OK: Kaminario Disk eui.0024f40055820030 is properly configured according to
Kaminario's BP (NoofoutstandingIOswithcompetingworlds=256).

_____________________________________
10.38.67.192 - Running validation for Qlogic Settings
10.38.67.192 - OK: Qlogic Options are properly configured according to Kaminario's BP
(ql2xintrdelaytimer=1 ql2xoperationmode=6 ql2xmaxqdepth=400)

_____________________________________
10.38.67.192 - Running validation for VAAI Primitives
10.38.67.192 - OK: VAAI Primitives are properly configured according to Kaminario's BP
10.38.67.192 - OK: Current settings are HardwareAcceleratedMove=1, HardwareAcceleratedInit=1,
HardwareAcceleratedLocking=1

_____________________________________
10.38.67.192 - Running validation for Round-Robin (Multipath configuration)
10.38.67.192 - OK: Kaminario disk eui.0024f40055820000 is set to MultipathPolicy=RoundRobin and
CommandsToSwitchPath=1
10.38.67.192 - OK: Kaminario disk eui.0024f40055820027 is set to MultipathPolicy=RoundRobin and
CommandsToSwitchPath=1
10.38.67.192 - OK: Kaminario disk eui.0024f40055820028 is set to MultipathPolicy=RoundRobin and
CommandsToSwitchPath=1
10.38.67.192 - OK: Kaminario disk eui.0024f40055820029 is set to MultipathPolicy=RoundRobin and
CommandsToSwitchPath=1
10.38.67.192 - OK: Kaminario disk eui.0024f4005582002a is set to MultipathPolicy=RoundRobin and
CommandsToSwitchPath=1
10.38.67.192 - OK: Kaminario disk eui.0024f4005582002b is set to MultipathPolicy=RoundRobin and
CommandsToSwitchPath=1
10.38.67.192 - OK: Kaminario disk eui.0024f4005582002c is set to MultipathPolicy=RoundRobin and
CommandsToSwitchPath=1
10.38.67.192 - OK: Kaminario disk eui.0024f4005582002d is set to MultipathPolicy=RoundRobin and
CommandsToSwitchPath=1
10.38.67.192 - OK: Kaminario disk eui.0024f4005582002e is set to MultipathPolicy=RoundRobin and
CommandsToSwitchPath=1
10.38.67.192 - OK: Kaminario disk eui.0024f40055820030 is set to MultipathPolicy=RoundRobin and
CommandsToSwitchPath=1

_____________________________________
10.38.67.192 - Running validation for SATP
10.38.67.192 - OK: Kaminario SATP rule is properly configured

_____________________________________
10.38.67.192 - Running validation for CBRC
10.38.67.192 - OK: CBRC is enabled and properly configured with DCacheMemReserved 2048
10.38.67.192 - Validation for 10.38.67.192 completed.

_____________________________________
```

5. Apply any additional host tuning (for Linux systems). Settings shown below help to support high IOPS operation of the Kaminario flash array with Linux hosts.

   a) Create the rule file /etc/udev/rules.d/62-io-schedulers.rules with the following contents:

```
#set noop scheduler for K2 disks
ACTION=="add|change", ATTRS{model}=="K2", ATTR{queue/scheduler}="noop"
```

b) Run the following udev command:

```
udevadm trigger && udevadm settle
```

c) Verify that the scheduler is configured to **noop:**

```
for i in /sys/block/sd*/queue/scheduler; do printf "$i is "; cat $i; done
```

6. Set up workload generators.

   On Windows and Linux systems, Medusa Labs Test Tools is installed. On VMware systems, VMware's I/O Analyzer is installed.

Test Report

What's New in This Report

This is the first test report for the Kaminario K2 array.

Test Plan Overview

The storage array is connected to two SAN fabrics and multiple server hosts to drive I/O in a multipath configuration. Error injection is introduced, and failover and recovery behaviors are observed. I/O performance is observed across different workload configurations.

Scope

Testing is performed with a mix of GA and development versions of Brocade's Fabric OS (FOS) in a heterogeneous environment. Test beds include Brocade directors and switches in routed and non-routed fabric configurations.

Testing is centered on interoperability and optimal configuration. 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 chassis server hosts with Brocade/QLogic/Emulex HBAs with two uplinks from every host to a Brocade FC fabric. I/O generators include Medusa Labs Test Tools and VMware I/O Analyzer.

Test Configuration

FIGURE 5 Test Configuration

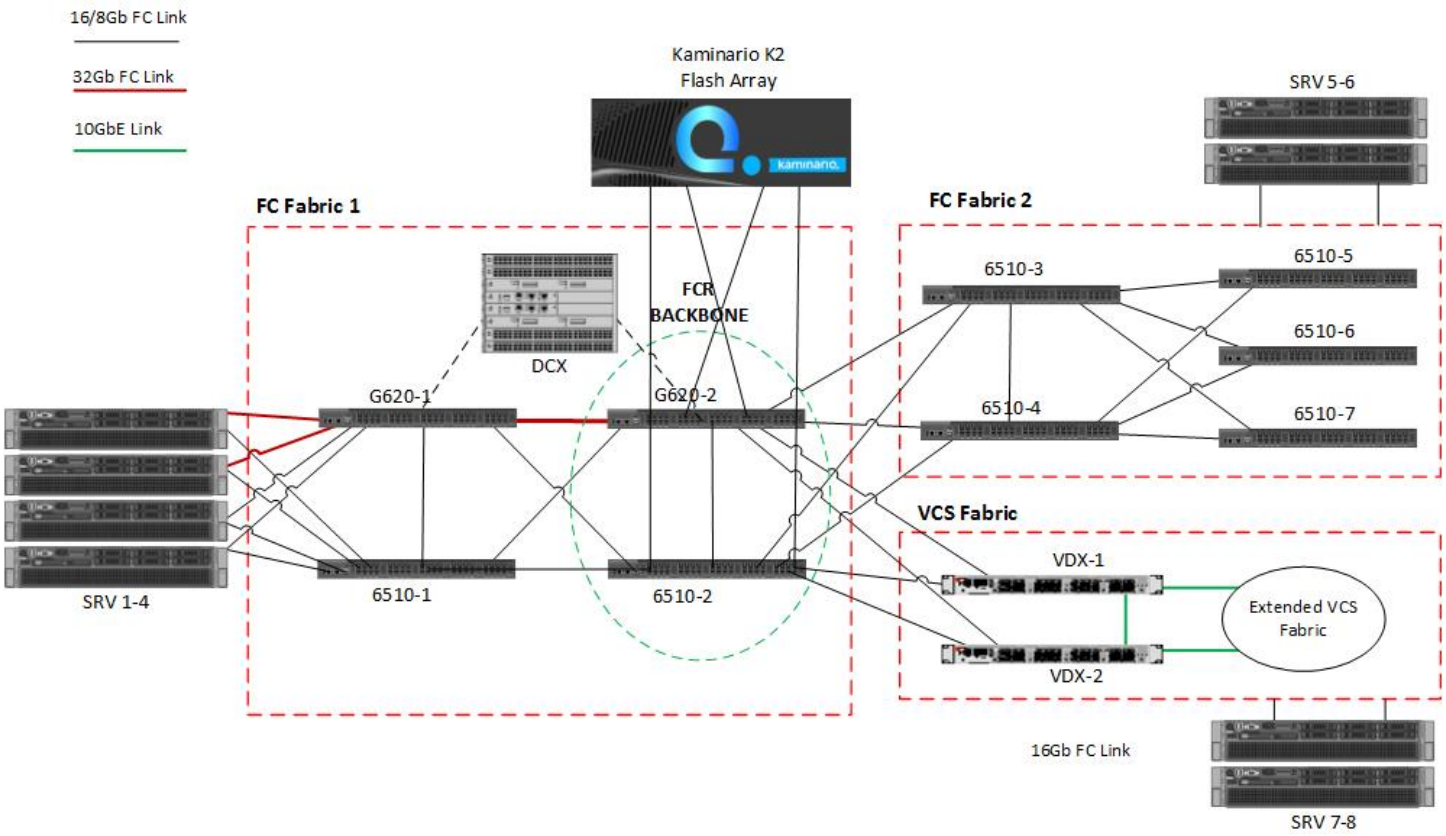

DUT Descriptions

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

TABLE 1 Storage Array

DUT ID	Model	Vendor	Description
Kaminario Flash Array	K2	Kaminario	The Kaminario K2 under test is an all-flash array with one K-Block composed of two controller nodes and one disk shelf with 6.7 TB of total physical capacity.

TABLE 2 Switches

DUT ID	Model	Vendor	Description
6510-1...7	Brocade BR-6510	Brocade	48-port 16Gb FC switch
G620-1,2	Brocade G620	Brocade	64-port 32Gb FC switch
DCX-1	DCX 8510-4	Brocade	4-slot 16Gb FC chassis
VDX-1,2	VDX 6740	Brocade	64-port 10Gb switch (48x10Gb + 4x40Gb)

DUT Specifications

Storage	Version
Kaminario K2 All-Flash Array	5.2.2.26, 5.2.4.23

Brocade Switches	Version
Brocade G620 with Fabric Vision, IO Insight licenses	FOS 8.0.1
DCX 8510-4	FOS 8.0.1
Brocade 6510 with Integrated Routing, Fabric Vision licenses	FOS 8.0.1
VDX 6740	NOS 7.0.0

Adapters	Version
Emulex LPe32002-M2 2-port 32Gb FC HBA	Driver 11.0.240.0, firmware 11.0.243.19
QLogic QLE2742 2-port 32Gb FC HBA	Driver 9.1.17.21, firmware 8.03.04
Brocade 1860 2-port 16Gb FC HBA	Driver & firmware version 3.2.6.0
Brocade 1860 2-port 16Gb FC HBA	Driver & firmware version 3.2.6.0
Emulex LPe 16202-X 2-port 16Gb FC HBA	Driver 11.0.243.21, firmware 11.0.243.11
Brocade 1020 2-port CNA adapter	Driver & firmware version 3.2.6.0

DUT ID	Servers	RAM	Processor	OS
SRV-1	IBM System x3650 M5	32 GB	Intel Xeon E5-2680v3	Windows Hyper-V
SRV-2	HP ProLiant DL380P G8	16 GB	Intel Xeon E5-2620	SLES 12 x86_64
SRV-3	HP ProLiant DL380P G8	160 GB	Intel Xeon E5-2640	VMWare 6.0 [cluster]
SRV-4	HP ProLiant DL380P G8	160 GB	Intel Xeon E5-2640	VMWare 6.0 [cluster]
SRV-5	HP ProLiant DL385p G8	16 GB	AMD Opteron 6212	RHEL 7.2 x86_64
SRV-6	Dell PowerEdge R720	16 GB	Intel Xeon E5-2620	RHEL 6.7 x86_64
SRV-7	HP ProLiant DL380P G8	32 GB	Intel Xeon E5-2690v2	Windows Server 2012R2
SRV-8	IBM System x3630 M4	16 GB	Intel Xeon E5-2620	SLES 12 x86_64

Test Equipment

Device/Software Tools	Version
Finisar 16Gb Analyzer/Jammer	XGIG5K2001153
Medusa Labs Test Tools	7.2.0.169914
VMware I/O Analyzer	1.6.2

Test Cases

1.	**FABRIC INITIALIZATION—BASE FUNCTIONALITY** Confirm basic Fibre Channel functionality of the storage array.
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** Examine the storage behavior related to more advanced fabric features such as QoS, Bottleneck Detection, and advanced frame recovery.
2.1	Storage Device Bottleneck Detection—With Congested Host
2.2	Storage Device Bottleneck Detection—With Congested Fabric
2.3	Storage Device—I/O Insight Latency Monitoring with MAPS Alerts
2.4	Storage Device—QoS Integrity with QoS Zone-Based Traffic Prioritization
2.5	Storage Device—QoS Integrity with CS_CTL-Based Traffic Prioritization
2.6	Storage Device—FC Protocol Jammer Test Suite
3.	**STRESS and ERROR RECOVERY WITH DEVICE MULTIPATH** Confirm proper HA/failover behavior of storage in a multipath environment.
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 Nameserver Integrity—Device Recovery with Device Relocation
3.4	Storage Device Nameserver Stress—Device Recovery with Device Port Toggle
3.5	Storage Device Recovery—ISL Port Toggle
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** Confirm proper storage functioning within routed fabrics.
4.1	Storage Device Internetworking Validation with FC Host
4.2	Storage Device Internetworking Validation with FCoE Test Using VDX Flexpoint
4.3	Storage Device Edge Recovery After FCR Disruptions
4.4	Storage Device Backbone Recovery After FCR Disruptions
5.	**Optional/Additional Tests**
5.1	Storage Device Nondisruptive Firmware Update
5.2	Performance I/O Testing with Varying Workload Dimensions
5.3	Performance Testing with VMware I/O Analyzer Application Workload Suite
5.4	Array Hardware Failure Tests

1. Fabric Initialization—Base Functionality

1.1 Storage Device—Physical and Logical Login with Speed Negotiation

Test Objective

Verify device login to the switch and nameserver with all supported speed settings.

Test Execution

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

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

Result Validation

1. Validate link states on the array, and verify speed negotiation and device login at different speeds. Use the **portshow** command on the switch to check the link state and speed, and use the **nscamshow** command to verify device login to the fabric.
2. Check the switch port status, and verify the "actual" and "configured" link speed. Check the name server for device login.

```
> portshow 36
portIndex:  36
portName: port36
portHealth: HEALTHY

portState: 1    Online
Protocol: FC
portPhys:  6    In_Sync         portScn:   32   F_Port
port generation number:    76
state transition count:    13

portId:    df2400
portIfId:    4302001e
portWwn:   20:24:c4:f5:7c:2a:8b:c8
portWwn of device(s) connected:
        50:02:4f:40:55:82:01:01
Distance:  normal
portSpeed: N8Gbps  <--
```

Test Results

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

1.2 Zoning and LUN Mapping

Test Objective

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

Test Execution

1. Create an FC zone on the fabric with the initiator and target WWNs.
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.

- Linux: Check the output of **lsscsi**.
- Windows: Check the output of **Computer Management** > **Storage** > **Disk Management**.
- VMware: Check the output of **Configuraton** > **Storage** > **Devices**.

Test Results

PASS. The 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 read/write/verify testing.

Test Execution

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

Result Validation

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

Test Results

PASS. I/O resumes without errors.

1.4 Storage Device Multipath Configuration—Path Integrity

Test Objective

Verify multipath configures successfully and verify integrity of each path.

Test Execution

1. Set up host with at least 2 initiator ports zoned with 2 target ports on array.
2. Set up multipath on host.
3. Start I/O.
4. Isolate paths individually and run IO on the separate paths.

Result Validation

1. Check host multipath properties to verify the toggled path recovers.
   - Windows: **mpclaim -s -d**
   - Linux: **multipath -ll**

- VMware: Check the paths at **Configure** > **Storage** > **Devices** > **Manage Paths**.

2. Check the host and storage logs for any failures.
3. Check the switch error logs and switch port status after toggle.

   errdumpal

   portstatsshow X

   portshow X

```
> portshow 36
portIndex:  36
portName: port36
portHealth: HEALTHY

portState: 1    Online
Protocol: FC
portPhys:  6    In_Sync         portScn:   32   F_Port
port generation number:    76
state transition count:    13

portId:    df2400
portIfId:    4302001e
portWwn:   20:24:c4:f5:7c:2a:8b:c8
portWwn of device(s) connected:
        50:02:4f:40:55:82:01:01
Distance:  normal
portSpeed: N8Gbps
```

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

Test Results

PASS. Each individual path is valid, and the paths combine successfully in a multipath configuration.

2. Fabric—Advanced Functionality

2.1 Storage Device Bottleneck Detection—With Congested Host

Test Objective

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

Test Execution

1. Configure MAPS and FPI on all switches (more information is provided in the "Task 1. Brocade FC Fabric Configuration" section).
2. Start I/O from a single host initiator to multiple targets.
3. Monitor the switch logs for IO_PERF_IMPACT/IO_FRAME_LOSS warnings.

Result Validation

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

errdumpall | grep IO_

```
> errdumpall | grep IO_
2016/5/17-11:56:00:672622, [MAPS-1003], 117148/115018, FID 128, WARNING, B6510_066_088, Port 16,
Condition=ALL_F_PORTS(DEV_LATENCY_IMPACT==IO_PERF_IMPACT), Current Value:
[DEV_LATENCY_IMPACT,IO_PERF_IMPACT, 30.0% in 10 secs], RuleName=defALL_F_PORTS_IO_PERF_IMPACT, Dashboard
Category=Fabric Performance Impact., actionHndlr.c, line: 755, comp:md, ltime:2016/5/17-11:56:00:671909
```

mapsdb --show all

```
> mapsdb --show

1 Dashboard Information:
========================
Active policy:                  dflt_aggressive_policy
……….

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

3.1 Summary Report:
====================
Category                 |Today                    |Last 7 days               |
-------------------------------------------------------------------------------
Port Health              |No Errors                |Out of operating range    |
Fru Health               |In operating range       |In operating range        |
Security Violations      |No Errors                |In operating range        |
Fabric State Changes     |No Errors                |In operating range        |
Switch Resource          |In operating range       |In operating range        |
Traffic Performance      |In operating range       |In operating range        |
FCIP Health              |Not applicable           |Not applicable            |
Fabric Performance Impact|Out of operating range   |Out of operating range    |
```

Test Results

PASS. With MAPS reporting configured, performance warnings are reported as expected. The bottlenecked ports are displayed on the MAPS dashboard, and a RASLog warning is created.

2.2 Storage Device Bottleneck Detection—With Congested Fabric

Test Objective

Validate bottleneck detection with congested fabric.

Test Execution

1. Configure MAPS on all switches (more information is provided in the "Task 1. Brocade FC Fabric Configuration" section).
2. Isolate a single ISL in the fabric.
3. Start I/O from multiple host initiators to multiple targets.
4. Monitor switch logs for traffic performance warnings.

Result Validation

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

errdumpall | grep Traffic

```
2016/5/18-08:39:03:497512, [MAPS-1003], 9637/5496, FID 128, WARNING, B6510_066_082, Port 33,
Condition=ALL_TARGET_PORTS(RX/hour>60.00), Current Value:[RX,65.16 %], RuleName=defALL_TARGET_PORTSRX_60,
Dashboard Category=Traffic Performance., actionHndlr.c, line: 755, comp:md, ltime:2016/5/18-08:39:03:496754
```

mapsdb --show all

```
> mapsdb --show

1 Dashboard Information:
========================
Active policy:                  dflt_aggressive_policy
……….

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

3.1 Summary Report:
===================
Category                 |Today                     |Last 7 days               |
-------------------------------------------------------------------------------
Port Health              |No Errors                 |Out of operating range    |
Fru Health               |In operating range        |In operating range        |
Security Violations      |No Errors                 |In operating range        |
Fabric State Changes     |No Errors                 |In operating range        |
Switch Resource          |In operating range        |In operating range        |
Traffic Performance      |In operating range        |In operating range        |
FCIP Health              |Not applicable            |Not applicable            |
Fabric Performance Impact|Out of operating range    |Out of operating range    |
```

Test Results

PASS. With MAPS reporting configured, performance warnings are reported as expected.

2.3 Storage Device—I/O Insight Latency Monitoring with MAPS Alerts

Test Objective

Configure a latency monitoring threshold on target port, and verify that alerts are generated when the threshold is exceeded.

Test Execution

1. Run an I/O workload generator with a 4k read workload against a target LUN, and record the latency reading as a baseline. Here we are seeing about a 0.4 ms value.

   ```
   Avg Completion Time   0.000444672
   ```

2. Configure a flow to monitor array target port traffic.

   ```
   > flow --create ios_kaminario_1 -dstdev df2500 -egrport 37 -feature mon
   ```

3. Import the flow into MAPS.

   ```
   > mapsconfig --import ios_kaminario_1
   ```

4. Create a custom policy.

   ```
   > mapspolicy -clone dflt_aggressive_policy -name ios_mod_policy
   ```

5. Create a rule to monitor and generate alerts at a threshold level higher than the baseline. The units for the command are microseconds, so in this example, 750 µs or 0.75 ms.

```
> mapsrule --create ios_k1 -group ios_kaminario_1 -monitor RD_STATUS_TIME_LT_8K -timebase min -op g -
value 750 -action raslog,email
```

6. Add the rule to the custom policy, and enable the policy.

```
> mapspolicy --addrule ios_mod_policy -rulename ios_k1
> mapspolicy --enable ios_mod_policy
```

7. Confirm that the flow is active.

```
> flow --show ios_kaminario_1
================================================================================
Name       : ios_kaminario_1     Features: mon(Activated)        noConfig: Off
Definition: EgrPort(37),DstDev(0xdf2500)


Flow Monitor (Activated):
Monitor time:  | Wed May 25 11:15:58 MDT 2016 |
```

8. Adjust the traffic pattern to cause a rise in latency above the configured monitoring threshold, and confirm that RASLog, MAPS Dashboard, and e-mail notifications are generated.

RASLog:

```
2016/04/20-16:31:29, [MAPS-1003], 497, FID 128, WARNING, G620_066_223, Flow (ios_kaminario_1), Condition=ios_kaminario_1(RD_STATUS_TIME_LT_8K/
min>750), Current Value:[ RD_STATUS_TIME_LT_8K,1184 Microseconds], RuleName=ios_Kaminario_small_read_P1, Dashboard Category=Traffic Performance.
```

E-mail:

```
Switch Time:        Apr 20 15:16:59
Affected Entity:    Flow (ios_kaminario_1)
Monitor:        Read completion time (RD_STATUS_TIME_LT_8K)
Rule Name:        ios_Kaminario_small_read_P1
Group:            ios_kaminario_1
Condition:        ios_kaminario_1(RD_STATUS_TIME_LT_8K/min>750)
Current Value:        371816 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(|3          |ios_Kaminario_small_read_P1|04/20/16 16:28:29|Flow (ios_kaminar|1478 Microseconds     |
3)                  |           |                           |                 |io_1)            |                      |
                    |           |                           |                 |Flow (ios_kaminar|1406 Microseconds     |
                    |           |                           |                 |io_1)            |                      |
                    |           |                           |                 |Flow (ios_kaminar|4409 Microseconds     |
                    |           |                           |                 |io_1)            |                      |
```

Test Results

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

2.4 Storage Device—QoS Integrity with QoS Zone-Based Traffic Prioritization

Test Objective

Validate QoS functionality.

Test Execution

1. Set up initiator-target pairs with Low/Medium/High QoS zones in the fabric.
2. Start I/O across all pairs, and validate traffic priority.

Result Validation

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

```
> zoneshow
zone:  QOSH_ssr067172_kaminario
zone:  QOSL_ssr067178_kaminario
zone:  QOSM_ssr067176_kaminario
zone:  QOSH_ssr067172_kaminario
zone:  QOSL_ssr067178_kaminario
zone:  QOSM_ssr067176_kaminario
```

2. Check the switch error logs and switch port status for errors.

porterrshow

```
> porterrshow
        frames      enc    crc    crc    too    too    bad    enc   disc   link   loss   loss   frjt   fbsy  c3timeout    pcs
     tx     rx      in    err    g_eof  shrt   long   eof    out   c3     fail   sync   sig                  tx    rx    err
0:    1.7g   2.7g   0      0      0      0      0      0      0    946      1      0      1      0      0      0      0      0
1:  431.1m 431.4m   0      0      0      0      0      0      0     30      1      0      1      0      0      0      0      0
2:    0      0      0      0      0      0      0      0      0      0      0      0      0      0      0      0      0      0
3:    0      0      0      0      0      0      0      0      0      0      0      0      0      0      0      0      0      0
```

3. Verify traffic prioritization by examining performance and VC utilization.

portregshow 36 | grep bbc_trc

Low

```
0x88c02800: bbc_trc                4    0    2    2    2    2    1    1
0x88c02820: bbc_trc                0    0    2    2    2    2    2    0
```

Med

```
0x88d82800: bbc_trc                4    0    2    2   -3   -3    1    1
0x88d82820: bbc_trc                2    2    2    2    2    2    2    0
```

High

```
0x88c02800: bbc_trc                4    0    2    2    2    2    1    1
0x88c02820: bbc_trc                2    2    2    2   -1    2    2    0
```

Test Results

PASS. I/O completes without error in low, medium, and high priority zones.

2.5 Storage Device—QoS Integrity with CS_CTL-Based Traffic Prioritization

Test Objective

Validate CS_CTL QoS functionality.

Test Execution

1. Enable ExpressLane on the Emulex adapter driver by setting the value to 1 and setting the priority to 3 for high.

FIGURE 6 Configuring Driver Parameters in the Emulex OneCommand Interface

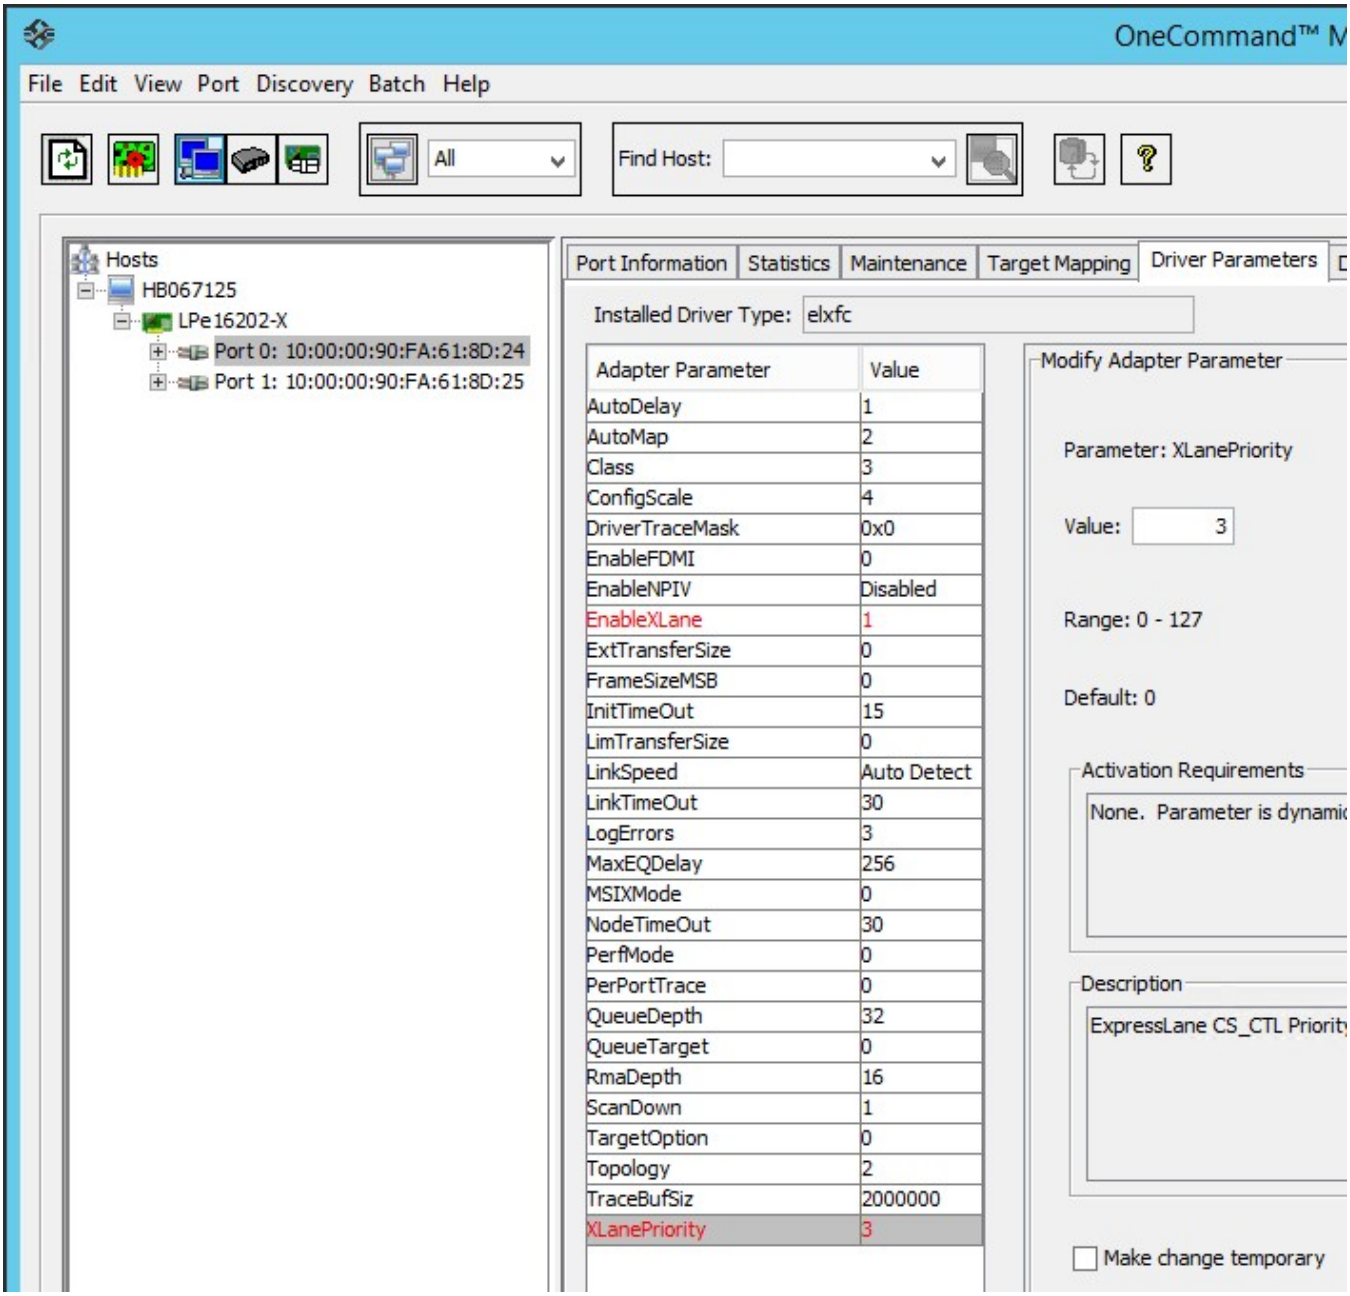

2. Enable CS_CTL QoS Auto-Mode at the chassis level on all switches in the Brocade fabric.

```
> 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
```

To verify:

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

3. Enable CS_CTL QoS on initiator and target ports.

```
> portcfgqos --enable 7 csctl_mode
Enabling CSCTL mode flows causes QoS zone flows to lose priority on such ports.

Do you want to proceed?(y/n):y
```

To verify:

```
> portcfgshow 7 | grep -i csctl
CSCTL mode:              ON
```

4. Enable ExpressLane QoS on selected LUNs. For multipath LUNs, enable ExpressLane on each device path.

FIGURE 7 Enabling LUNs for ExpressLane in the Emulex OneCommand Interface

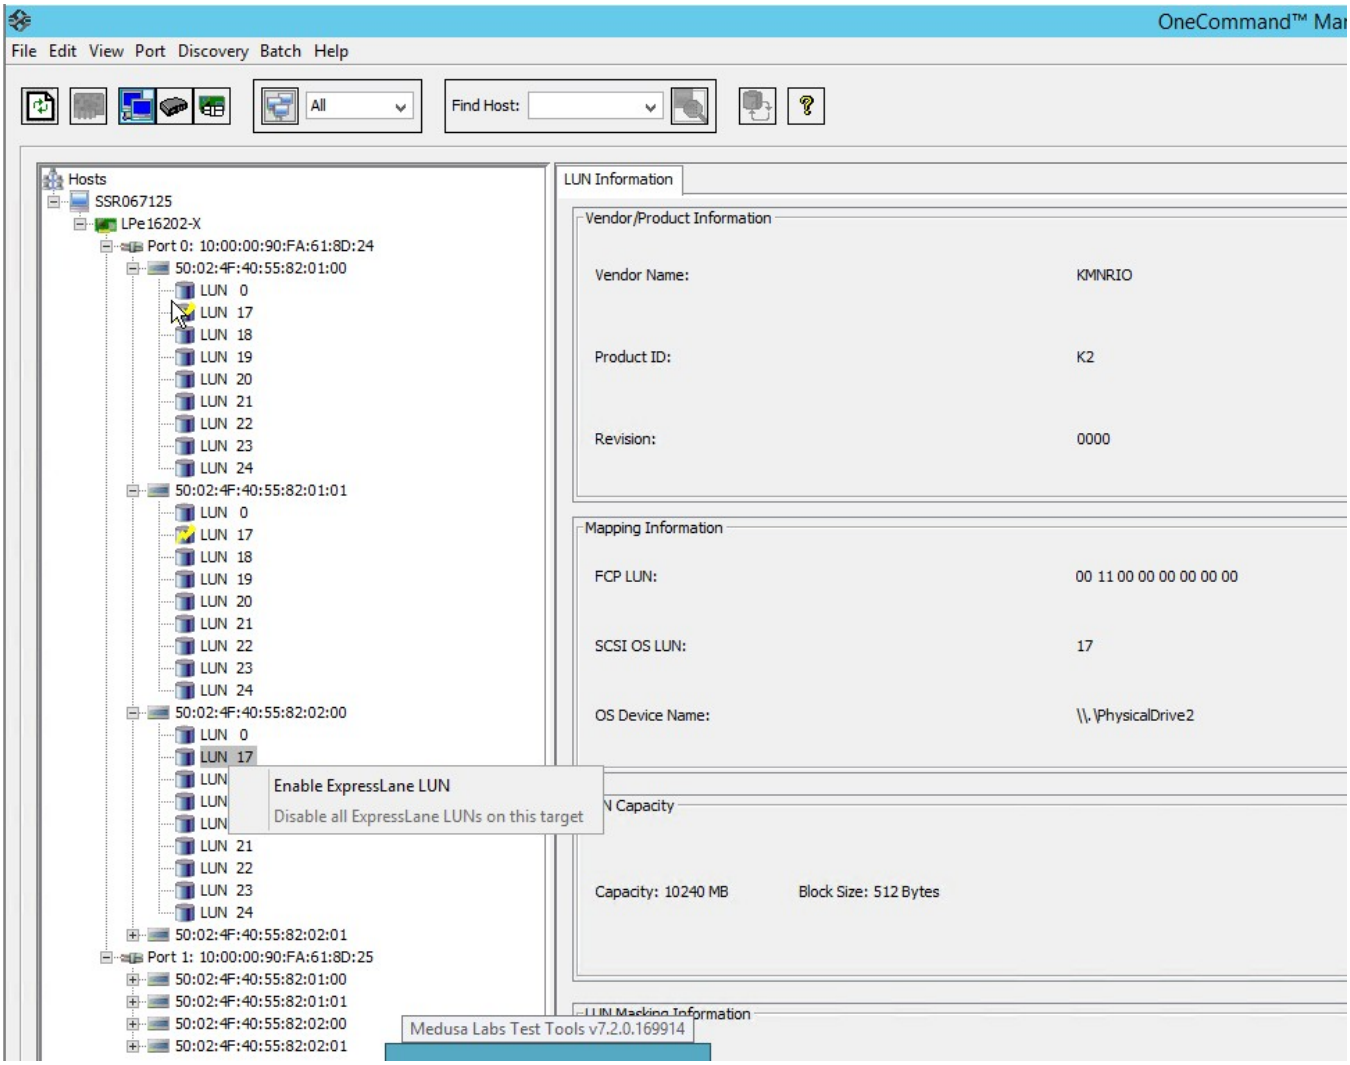

5. Start high throughput I/O to all LUNs.

Result Validation

1. Check I/O logs, and verify that I/O continues without error for all LUNs.
2. Verify that I/O performance is improved on ExpressLane-enabled LUNs.
3. Verify CS_CTL prioritization in the fabric by monitoring the high VC buffer credits on the ISLs.
4. Check the switch error logs and switch port status for errors.

   porterrshow

```
> porterrshow
        frames      enc    crc    crc    too    too    bad    enc   disc   link   loss   loss   frjt   fbsy  c3timeout    pcs
     tx     rx      in    err    g_eof  shrt   long   eof    out   c3    fail    sync   sig                 tx    rx    err
  0:   1.7g   2.7g   0      0      0      0      0      0      0    946      1      0      1      0      0      0      0      0
  1: 431.1m 431.4m   0      0      0      0      0      0      0     30      1      0      1      0      0      0      0      0
  2:    0      0      0      0      0      0      0      0      0     0      0      0      0      0      0      0      0      0
  3:    0      0      0      0      0      0      0      0      0     0      0      0      0      0      0      0      0      0
```

Test Results

PASS. Prioritized LUNs display higher performance, the traffic receives prioritization through the fabric, and I/O completes without errors.

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 Execution

1. Insert the Jammer device in the I/O path on the storage link.
2. Execute the following Jammer scenarios:
   - Delete one frame.
   - Delete R_RDY.
   - Replace CRC of data frame.
   - Replace EOF of data frame.
   - Replace "good status" with "check condition".
   - Replace IDLE with LR.
   - Truncate frame.
   - Create S_ID/D_ID error of data frame.
3. Verify Jammer operations and recovery with Analyzer.

Result Validation

Check the Medusa log, and verify that I/O recovers and completes without error.

Test Results

PASS. Packet anomalies are introduced, and proper recovery is verified.

3. Stress and Error Recovery with Device Multipath

3.1 Storage Device Fabric I/O Integrity—Congested Fabric

Test Objective

Validate I/O integrity in a congested fabric environment.

Test Execution

From all initiators, start a mixture of read, read/write, and write traffic continuously to all their targets for a 24-hour period.

Result Validation

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

   errdumpall

   portperfshow

   porterrshow

```
> porterrshow
          frames      enc    crc    crc    too    too    bad    enc   disc   link   loss   loss   frjt   fbsy  c3timeout    pcs
       tx     rx      in     err    g_eof  shrt   long   eof    out   c3     fail   sync   sig                  tx     rx     err
  0:    1.7g   2.7g   0      0      0      0      0      0      0    946      1      0      1      0      0      0      0      0
  1:  431.1m 431.4m   0      0      0      0      0      0      0     30      1      0      1      0      0      0      0      0
  2:     0      0      0      0      0      0      0      0      0      0      0      0      0      0      0      0      0      0
  3:     0      0      0      0      0      0      0      0      0      0      0      0      0      0      0      0      0      0
```

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

Test Results

PASS. In a congested fabric, I/O runs successfully without error or failover.

3.2 Storage Device Nameserver Integrity—Device Recovery with Port Toggle

Test Objective

Validate path recovery behavior on storage and host ports.

Test Execution

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 switch port status after toggling, and check for errors in the switch logs.

errdumpall

portstatsshow X

portshow X

```
> portshow 36
portIndex:  36
portName: port36
portHealth: HEALTHY

portState: 1    Online
Protocol: FC
portPhys:  6    In_Sync         portScn:   32   F_Port
port generation number:    76
state transition count:    13

portId:    df2400
portIfId:    4302001e
portWwn:   20:24:c4:f5:7c:2a:8b:c8
portWwn of device(s) connected:
        50:02:4f:40:55:82:01:01
Distance:  normal
portSpeed: N8Gbps
```

2. Check the multipath status on hosts to verify that the toggled path recovers.

   Windows: **mpclaim -s -d**

   Linux: **multipath -ll**

   VMware: Check the paths at **Configuration** > **Storage** > **Devices** > **Manage Paths**

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

Test Results

PASS. Failover between 8 logical paths (2 host x 4 storage) tests successfully. Paths recover, and I/O completes without error.

3.3 Storage Device Nameserver Integrity—Device Recovery with Device Relocation

Test Objective

Validate storage device path recovery after physical port relocation.

Test Execution

1. Perform the test sequentially for each storage device port.
2. Disconnect a port and reconnect it to a different switch port in the same fabric.

Result Validation

1. Check for errors in the switch logs, and check the status at the new switch port.

   errdumpall

   portstatsshow X

   portshow X

```
> portshow 36
portIndex: 36
```

```
portName: port36
portHealth: HEALTHY

portState: 1    Online
Protocol: FC
portPhys:  6    In_Sync         portScn:   32   F_Port
port generation number:    76
state transition count:    13

portId:    df2400
portIfId:  4302001e
portWwn:   20:24:c4:f5:7c:2a:8b:c8
portWwn of device(s) connected:
        50:02:4f:40:55:82:01:01
Distance:  normal
portSpeed: N8Gbps
```

2. Check the multipath status on hosts to verify that the toggled path recovers.

   Windows: **mpclaim -s -d**

   Linux: **multipath -ll**

   VMware: Check the paths at **Configuration** > **Storage** > **Devices** > **Manage Paths**

3. Confirm that there are no errors in the Medusa log.

Test Results

PASS. The physical move of the storage port shows successful recovery. Paths recover, and IO completes without error.

3.4 Storage Device Nameserver Stress—Device Recovery with Device Port Toggle

Test Objective

Validate path recovery behavior on storage and host ports for an extended duration.

Test Execution

1. Set up multipath on the host, and start I/O.
2. Sequentially toggle each initiator and target port in the fabric (multiple iterations).
3. Run the test for an extended period.

Result Validation

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

   errdumpall

   portstatsshow X

   portshow X

```
> portshow 36
portIndex:  36
portName: port36
portHealth: HEALTHY

portState: 1    Online
Protocol: FC
portPhys:  6    In_Sync         portScn:   32   F_Port
port generation number:    76
```

```
state transition count:    13

portId:    df2400
portIfId:  4302001e
portWwn:   20:24:c4:f5:7c:2a:8b:c8
portWwn of device(s) connected:
        50:02:4f:40:55:82:01:01
Distance:  normal
portSpeed: N8Gbps
```

2. Check the host multipath properties for iSCSI hosts to verify that the toggled path recovers.

   On Windows: **mpclaim -s -d**

   On Linux: **multipath -ll**

   On VMware: Check the paths at **Configuration** > **Storage** > **Devices** > **Manage Paths**

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

Test Results

PASS. 48-hour run; paths recover, and I/O completes without error.

3.5 Storage Device Recovery—ISL Port Toggle

Test Objective

Validate path recovery and I/O integrity when ISL links are disabled.

Test Execution

1. Set up multipath on the host, and start I/O.
2. Sequentially toggle each ISL path (one at a time) on all switches.

Result Validation

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

   fabricshow

```
> fabricshow
Switch ID   Worldwide Name          Enet IP Addr    FC IP Addr      Name
-------------------------------------------------------------------------
  1: fffc01 50:00:53:35:b1:d3:df:5c 0.0.0.0         0.0.0.0         "fcr_xd_1_40"
  4: fffc04 50:00:53:35:b1:d3:df:5e 0.0.0.0         0.0.0.0         "fcr_xd_4_50"
 73: fffc49 10:00:00:05:33:13:95:9a 10.38.66.73     0.0.0.0        >"B6510_066_073"
 74: fffc4a 10:00:00:05:33:a5:bf:86 10.38.66.74     0.0.0.0         "B6510_066_074"
 82: fffc52 10:00:00:05:33:13:96:5a 10.38.66.82     0.0.0.0         "B6510_066_082"
 83: fffc53 10:00:00:05:33:5b:1d:1d 10.38.66.83     0.0.0.0         "B6510_066_083"
223: fffcdf 10:00:c4:f5:7c:2a:8b:c8 10.38.66.223    0.0.0.0         "G620_066_223"
224: fffce0 10:00:c4:f5:7c:41:1b:14 10.38.66.224    0.0.0.0         "G620_066_224"

The Fabric has 8 switches
Fabric Name: SSR
```

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

   errdumpall

   portperfshow

porterrshow

```
> porterrshow
         frames      enc    crc    crc    too    too    bad    enc   disc   link   loss   loss   frjt   fbsy  c3timeout    pcs
      tx     rx      in    err    g_eof  shrt   long   eof    out   c3     fail   sync   sig                  tx    rx    err
  0:   1.7g   2.7g   0      0      0      0      0      0      0    946      1      0      1      0      0      0      0      0
  1:  431.1m 431.4m  0      0      0      0      0      0      0     30      1      0      1      0      0      0      0      0
  2:    0      0      0      0      0      0      0      0      0      0      0      0      0      0      0      0      0      0
  3:    0      0      0      0      0      0      0      0      0      0      0      0      0      0      0      0      0      0
```

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

Test Results

PASS. Paths recover, and I/O completes without error.

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

Test Objective

Validate path recovery and I/O integrity when all ISL links on a switch are disabled.

Test Execution

1. Ensure ISL redundancy by provisioning multiple ISLs connected to different switches to provide multiple paths through the fabric.
2. Set up multipath on the host, and start I/O.
3. On each switch, disable all ISL links at once.

Result Validation

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

   fabricshow

   ```
   > fabricshow
   Switch ID   Worldwide Name           Enet IP Addr    FC IP Addr      Name
   -------------------------------------------------------------------------
     1: fffc01 50:00:53:35:b1:d3:df:5c 0.0.0.0         0.0.0.0         "fcr_xd_1_40"
     4: fffc04 50:00:53:35:b1:d3:df:5e 0.0.0.0         0.0.0.0         "fcr_xd_4_50"
    73: fffc49 10:00:00:05:33:13:95:9a 10.38.66.73     0.0.0.0        >"B6510_066_073"
    74: fffc4a 10:00:00:05:33:a5:bf:86 10.38.66.74     0.0.0.0         "B6510_066_074"
    82: fffc52 10:00:00:05:33:13:96:5a 10.38.66.82     0.0.0.0         "B6510_066_082"
    83: fffc53 10:00:00:05:33:5b:1d:1d 10.38.66.83     0.0.0.0         "B6510_066_083"
   223: fffcdf 10:00:c4:f5:7c:2a:8b:c8 10.38.66.223    0.0.0.0         "G620_066_223"
   224: fffce0 10:00:c4:f5:7c:41:1b:14 10.38.66.224    0.0.0.0         "G620_066_224"

   The Fabric has 8 switches
   Fabric Name: SSR
   ```

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

   errdumpall

   portperfshow

   porterrshow

```
> porterrshow
         frames      enc    crc    crc    too    too    bad    enc   disc   link   loss   loss   frjt   fbsy  c3timeout    pcs
      tx     rx      in    err    g_eof  shrt   long   eof    out   c3     fail   sync   sig                  tx    rx    err
  0:   1.7g   2.7g   0      0      0      0      0      0      0    946      1      0      1      0      0      0      0      0
  1:  431.1m 431.4m  0      0      0      0      0      0      0     30      1      0      1      0      0      0      0      0
  2:    0      0      0      0      0      0      0      0      0      0      0      0      0      0      0      0      0      0
  3:    0      0      0      0      0      0      0      0      0      0      0      0      0      0      0      0      0      0
```

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

Test Results

PASS. Paths recover, and I/O completes without error.

3.7 Storage Device Recovery—Director Blade Maintenance

Test Objective

Validate path recovery and I/O integrity during director blade maintenance.

Test Execution

1. Uplink edge switch ISLs to different blades on the directors.
2. Sequentially toggle each director blade.
3. Include blade disable/enable and blade power on/off.

   Syntax:

```
slotpoweroff/slotpoweron, bladedisable/bladeenable, slotshow
```

Result Validation

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

   fabricshow

```
> fabricshow
Switch ID   Worldwide Name           Enet IP Addr    FC IP Addr      Name
-------------------------------------------------------------------------
  1: fffc01 50:00:53:39:47:bd:6e:0a 0.0.0.0         0.0.0.0         "fcr_fd_1"
  2: fffc02 10:00:00:05:1e:09:15:3e 10.38.51.5      0.0.0.0         "sw5300sus1005"
  3: fffc03 10:00:c4:f5:7c:07:c2:0c 10.38.51.3      0.0.0.0         "sw300sus1003"
  4: fffc04 10:00:c4:f5:7c:7b:86:1f 10.38.51.4      0.0.0.0         "sw6510sus1004"
  5: fffc05 10:00:00:05:33:13:80:ef 10.38.51.24     0.0.0.0         "swBESsus1024"
 13: fffc0d 10:00:00:05:1e:4c:ff:00 10.38.51.13     0.0.0.0         "sw8510sus1013"
 16: fffc10 10:00:00:05:1e:44:02:00 10.38.51.16     0.0.0.0        >"sw8518sus1016"
 22: fffc16 10:00:00:05:1e:07:7d:99 10.38.51.22     0.0.0.0         "sw5100sus1022"
 23: fffc17 10:00:00:27:f8:2b:9f:aa 10.38.51.23     0.0.0.0         "sw6520sus1023"
 25: fffc19 10:00:00:05:1e:54:8b:d4 10.38.51.25     0.0.0.0         "swBESsus1025"
 26: fffc1a 10:00:00:05:33:8f:b6:77 10.38.51.26     0.0.0.0         "sw6510sus1026"
 29: fffc1d 10:00:00:05:33:83:a4:00 10.38.51.29     0.0.0.0         "swDCX4sus1029"
 32: fffc20 10:00:00:05:1e:b8:c1:00 10.38.51.32     0.0.0.0         "swDCXsus1032"
160: fffca0 10:00:00:05:1e:a8:35:9d 10.38.51.21     0.0.0.0         "swVA40FCsus1021"
171: fffcab 10:00:00:05:1e:d7:1a:04 10.38.51.171    172.172.172.26  "sw5480sus1171"
                                    2620:100:4:fa00:205:1eff:fed7:1a04
183: fffcb7 10:00:00:05:33:00:b2:e2 10.38.51.183    0.0.0.0         "swM5424sus1183"
195: fffcc3 10:00:00:05:33:8d:df:d1 10.38.51.195    0.0.0.0         "sw5480sus1195"
199: fffcc7 10:00:00:27:f8:44:33:61 10.38.51.199    0.0.0.0         "sw6547sus1199"
                                    2620:100:4:fa00:227:f8ff:fe44:3362
220: fffcdc 10:00:00:05:1e:86:1b:1f 10.38.51.220    0.0.0.0         "sw5450sus1220"

The Fabric has 19 switches
```

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

   errdumpall

   portperfshow

porterrshow

```
> porterrshow
         frames      enc    crc    crc    too    too    bad    enc   disc   link   loss   loss   frjt   fbsy  c3timeout    pcs
      tx     rx      in    err    g_eof  shrt   long   eof    out   c3     fail   sync   sig                  tx     rx     err
  0:   1.7g   2.7g   0      0      0      0      0      0      0    946      1      0      1      0      0      0      0      0
  1:  431.1m 431.4m   0      0      0      0      0      0      0     30      1      0      1      0      0      0      0      0
  2:    0      0      0      0      0      0      0      0      0      0      0      0      0      0      0      0      0      0
  3:    0      0      0      0      0      0      0      0      0      0      0      0      0      0      0      0      0      0
```

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

Test Results

PASS. Paths recover, and I/O completes without error.

3.8 Storage Device Recovery—Switch Offline

Test Objective

Validate path recovery and I/O integrity during switch offline events.

Test Execution

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

Result Validation

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

   fabricshow

   ```
   > fabricshow
   Switch ID   Worldwide Name           Enet IP Addr    FC IP Addr      Name
   -------------------------------------------------------------------------
     1: fffc01 50:00:53:35:b1:d3:df:5c 0.0.0.0         0.0.0.0         "fcr_xd_1_40"
     4: fffc04 50:00:53:35:b1:d3:df:5e 0.0.0.0         0.0.0.0         "fcr_xd_4_50"
    73: fffc49 10:00:00:05:33:13:95:9a 10.38.66.73     0.0.0.0        >"B6510_066_073"
    74: fffc4a 10:00:00:05:33:a5:bf:86 10.38.66.74     0.0.0.0         "B6510_066_074"
    82: fffc52 10:00:00:05:33:13:96:5a 10.38.66.82     0.0.0.0         "B6510_066_082"
    83: fffc53 10:00:00:05:33:5b:1d:1d 10.38.66.83     0.0.0.0         "B6510_066_083"
   223: fffcdf 10:00:c4:f5:7c:2a:8b:c8 10.38.66.223    0.0.0.0         "G620_066_223"
   224: fffce0 10:00:c4:f5:7c:41:1b:14 10.38.66.224    0.0.0.0         "G620_066_224"

   The Fabric has 8 switches
   Fabric Name: SSR
   ```

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

   errdumpall

   switchshow

   ```
   > switchshow
   switchName:     B6510_066_088
   switchType:     109.1
   switchState:    Online   <--
   switchMode:     Native
   switchRole:     Subordinate
   switchDomain:   88
   switchId:       fffc58
   switchWwn:      10:00:00:27:f8:06:23:28
   zoning:         ON (SSR)
   ```

```
switchBeacon:   OFF
FC Router:      ON
FC Router BB Fabric ID: 100
Address Mode:   0
Fabric Name:    SSR_2
```

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

Test Results

PASS. Paths fail over and recover, and I/O completes without error.

3.9 Storage Device Recovery—Switch Firmware Download

Test Objective

Verify that I/O continues with minimal disruption throughout the switch firmware upgrade process.

Test Execution

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 that the firmware upgrade completes successfully on each switch node and that the nodes merge back into the FC fabric.

   version

```
> version
Kernel:     2.6.34.6
Fabric OS:  v8.0.1
Made on:    Thu Apr 21 17:48:31 2016
Flash:      Tue Apr 26 17:51:45 2016
BootProm:   2.0.25
```

   fabricshow

```
> fabricshow
Switch ID   Worldwide Name          Enet IP Addr    FC IP Addr      Name
-------------------------------------------------------------------------
  1: fffc01 50:00:53:35:b1:d3:df:5c 0.0.0.0         0.0.0.0         "fcr_xd_1_40"
  4: fffc04 50:00:53:35:b1:d3:df:5e 0.0.0.0         0.0.0.0         "fcr_xd_4_50"
 73: fffc49 10:00:00:05:33:13:95:9a 10.38.66.73     0.0.0.0        >"B6510_066_073"
 74: fffc4a 10:00:00:05:33:a5:bf:86 10.38.66.74     0.0.0.0         "B6510_066_074"
 82: fffc52 10:00:00:05:33:13:96:5a 10.38.66.82     0.0.0.0         "B6510_066_082"
 83: fffc53 10:00:00:05:33:5b:1d:1d 10.38.66.83     0.0.0.0         "B6510_066_083"
223: fffcdf 10:00:c4:f5:7c:2a:8b:c8 10.38.66.223    0.0.0.0         "G620_066_223"
224: fffce0 10:00:c4:f5:7c:41:1b:14 10.38.66.224    0.0.0.0         "G620_066_224"

The Fabric has 8 switches
Fabric Name: SSR
```

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

   errdumpall

   portperfshow

Test Results

PASS. I/O failover and recovery are successful during the firmware download. Paths recover, and I/O completes without error.

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

4.1 Storage Device Internetworking Validation with FC Host

Test Objective

Validate that storage targets are imported successfully in a routed Fibre Channel environment.

Test Execution

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 name server and FCR fabric state.

fcrfabricshow

```
> fcrfabricshow
FC Router WWN: 10:00:00:05:33:13:96:5a, Dom ID:  82,
Info: 10.38.66.82, "B6510_066_082"
   EX_Port       FID    Neighbor Switch Info (enet IP, WWN, name)
   ------------------------------------------------------------------------
      40         40     10.38.66.88      10:00:00:27:f8:06:23:28   "B6510_066_088"
      41         40     10.38.66.92      10:00:00:27:f8:66:f3:81   "B6520_066_92"
      26         50     10.38.66.112     10:00:50:eb:1a:20:d3:81   "VDX6740_066_112"
      27         50     10.38.66.111     10:00:50:eb:1a:22:27:da   "VDX6740_066_111"

FC Router WWN: 10:00:00:05:33:5b:1d:1d, Dom ID:  83,
Info: 10.38.66.83, "B6510_066_083"
   EX_Port       FID    Neighbor Switch Info (enet IP, WWN, name)
   ------------------------------------------------------------------------
      40         40     10.38.66.92      10:00:00:27:f8:66:f3:81   "B6520_066_92"
      41         40     10.38.66.88      10:00:00:27:f8:06:23:28   "B6510_066_088"
      26         50     10.38.66.111     10:00:50:eb:1a:22:27:da   "VDX6740_066_111"
      27         50     10.38.66.112     10:00:50:eb:1a:20:d3:81   "VDX6740_066_112"
```

fcrproxydevshow

```
> fcrproxydevshow
  Proxy           WWN             Proxy      Device   Physical    State
 Created                           PID       Exists     PID
in Fabric                                   in Fabric
-------------------------------------------------------------------
    40   10:00:00:05:1e:60:b4:6b  02ff02      100      132400   Imported
    40   10:00:00:05:1e:60:b4:6c  02fe02      100      032400   Imported
. . .
    40   50:02:4f:40:55:82:01:00  02f906      100      532400   Imported
    40   50:02:4f:40:55:82:01:01  02fa06      100      df2400   Imported
    40   50:02:4f:40:55:82:02:00  02fb06      100      532500   Imported
    40   50:02:4f:40:55:82:02:01  02fc06      100      df2500   Imported
    50   50:02:4f:40:55:82:01:00  02fc03      100      532400   Imported
    50   50:02:4f:40:55:82:01:01  02fd03      100      df2400   Imported
    50   50:02:4f:40:55:82:02:00  02fe03      100      532500   Imported
    50   50:02:4f:40:55:82:02:01  02ff03      100      df2500   ImportedTotal devices displayed: 20
```

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

Test Results

PASS. All devices are available via the routed fabric. Running I/O confirms successful routing.

4.2 Storage Device Internetworking Validation with FCoE Test Using VDX Flexpoint

Test Objective

Verify storage device availability from FCoE initiators routed through the VCS fabric.

Test Execution

1. Set up FCoE host initiator ports on the VCS fabric using Brocade VDX 6740 switches.
2. Enable Fibre Channel support on VDX switch ports (for more detail, see the "Task1. Brocade FC Fabric Configuration" section).
3. Connect the VCS fabric to the Fibre Channel fabric via FCR ports.
4. Configure LSAN zones in the VCS and FC fabrics.

Result Validation

1. Verify the nameserver and FCR fabric state.

fcrfabricshow

```
> fcrfabricshow
FC Router WWN: 10:00:00:05:33:13:96:5a, Dom ID:  82,
Info: 10.38.66.82, "B6510_066_082"
   EX_Port       FID    Neighbor Switch Info (enet IP, WWN, name)
   ------------------------------------------------------------------------
      40         40     10.38.66.88      10:00:00:27:f8:06:23:28   "B6510_066_088"
      41         40     10.38.66.92      10:00:00:27:f8:66:f3:81   "B6520_066_92"
      26         50     10.38.66.112     10:00:50:eb:1a:20:d3:81   "VDX6740_066_112"
      27         50     10.38.66.111     10:00:50:eb:1a:22:27:da   "VDX6740_066_111"

FC Router WWN: 10:00:00:05:33:5b:1d:1d, Dom ID:  83,
Info: 10.38.66.83, "B6510_066_083"
   EX_Port       FID    Neighbor Switch Info (enet IP, WWN, name)
   ------------------------------------------------------------------------
      40         40     10.38.66.92      10:00:00:27:f8:66:f3:81   "B6520_066_92"
      41         40     10.38.66.88      10:00:00:27:f8:06:23:28   "B6510_066_088"
      26         50     10.38.66.111     10:00:50:eb:1a:22:27:da   "VDX6740_066_111"
      27         50     10.38.66.112     10:00:50:eb:1a:20:d3:81   "VDX6740_066_112""VDX6740_066_119"
```

fcrproxydevshow

```
> fcrproxydevshow
  Proxy           WWN             Proxy      Device   Physical    State
 Created                           PID       Exists     PID
in Fabric                                   in Fabric
-------------------------------------------------------------------
    40   10:00:00:05:1e:60:b4:6b  02ff02      100      132400   Imported
    40   10:00:00:05:1e:60:b4:6c  02fe02      100      032400   Imported
. . .
    40   50:02:4f:40:55:82:01:00  02f906      100      532400   Imported
    40   50:02:4f:40:55:82:01:01  02fa06      100      df2400   Imported
    40   50:02:4f:40:55:82:02:00  02fb06      100      532500   Imported
    40   50:02:4f:40:55:82:02:01  02fc06      100      df2500   Imported
    50   50:02:4f:40:55:82:01:00  02fc03      100      532400   Imported
    50   50:02:4f:40:55:82:01:01  02fd03      100      df2400   Imported
```

```
    50   50:02:4f:40:55:82:02:00  02fe03      100      532500   Imported
    50   50:02:4f:40:55:82:02:01  02ff03      100      df2500   ImportedTotal devices displayed: 20
```

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

Test Results

PASS. Storage targets are available through the FCoE/FC routed configuration; I/O completes without error.

4.3 Storage Device Edge Recovery After FCR Disruptions

Test Objective

Validate that storage target paths recover successfully from disruptions in a routed Fibre Channel environment. Edge-Backbone-Edge configuration.

Test Execution

1. Set up FCR in an Edge-Backbone-Edge configuration.
2. Set up LSAN zoning.
3. With I/O running, 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

```
> fcrfabricshow
FC Router WWN: 10:00:00:05:33:13:96:5a, Dom ID:  82,
Info: 10.38.66.82, "B6510_066_082"
   EX_Port       FID    Neighbor Switch Info (enet IP, WWN, name)
   ------------------------------------------------------------
      40         40     10.38.66.88      10:00:00:27:f8:06:23:28   "B6510_066_088"
      41         40     10.38.66.92      10:00:00:27:f8:66:f3:81   "B6520_066_92"
      26         50     10.38.66.112     10:00:50:eb:1a:20:d3:81   "VDX6740_066_112"
      27         50     10.38.66.111     10:00:50:eb:1a:22:27:da   "VDX6740_066_111"

FC Router WWN: 10:00:00:05:33:5b:1d:1d, Dom ID:  83,
Info: 10.38.66.83, "B6510_066_083"
   EX_Port       FID    Neighbor Switch Info (enet IP, WWN, name)
   ------------------------------------------------------------
      40         40     10.38.66.92      10:00:00:27:f8:66:f3:81   "B6520_066_92"
      41         40     10.38.66.88      10:00:00:27:f8:06:23:28   "B6510_066_088"
      26         50     10.38.66.111     10:00:50:eb:1a:22:27:da   "VDX6740_066_111"
      27         50     10.38.66.112     10:00:50:eb:1a:20:d3:81   "VDX6740_066_112""VDX6740_066_119"
```

fcrproxydevshow

```
> fcrproxydevshow
  Proxy           WWN             Proxy      Device   Physical    State
 Created                           PID       Exists     PID
in Fabric                                   in Fabric
-------------------------------------------------------------------
    40   10:00:00:05:1e:60:b4:6b  02ff02      100      132400   Imported
    40   10:00:00:05:1e:60:b4:6c  02fe02      100      032400   Imported
. . .
    40   50:02:4f:40:55:82:01:00  02f906      100      532400   Imported
    40   50:02:4f:40:55:82:01:01  02fa06      100      df2400   Imported
    40   50:02:4f:40:55:82:02:00  02fb06      100      532500   Imported
    40   50:02:4f:40:55:82:02:01  02fc06      100      df2500   Imported
```

```
 50   50:02:4f:40:55:82:01:00  02fc03      100      532400   Imported
 50   50:02:4f:40:55:82:01:01  02fd03      100      df2400   Imported
 50   50:02:4f:40:55:82:02:00  02fe03      100      532500   Imported
 50   50:02:4f:40:55:82:02:01  02ff03      100      df2500   ImportedTotal devices displayed: 20
```

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. Paths fail over and recover with FCR disruptions, and I/O completes without error.

4.4 Storage Device Backbone Recovery After FCR Disruptions

Test Objective

Validate storage target paths recover successfully from disruptions in a routed Fibre Channel environment. Edge-Backbone configuration.

Test Execution

1. Set up FCR in an Edge-Backbone configuration.
2. Set up LSAN zoning.
3. With I/O running, perform sequential reboots, switch disables, and ISL port toggles on the switches in the backbone fabric.

Result Validation

1. Verify the FCR fabric state throughout the disruptions.

   fcrfabricshow

```
> fcrfabricshow
FC Router WWN: 10:00:00:05:33:13:96:5a, Dom ID:  82,
Info: 10.38.66.82, "B6510_066_082"
   EX_Port       FID    Neighbor Switch Info (enet IP, WWN, name)
   ------------------------------------------------------------------
      40         40     10.38.66.88      10:00:00:27:f8:06:23:28   "B6510_066_088"
      41         40     10.38.66.92      10:00:00:27:f8:66:f3:81   "B6520_066_92"
      26         50     10.38.66.112     10:00:50:eb:1a:20:d3:81   "VDX6740_066_112"
      27         50     10.38.66.111     10:00:50:eb:1a:22:27:da   "VDX6740_066_111"

FC Router WWN: 10:00:00:05:33:5b:1d:1d, Dom ID:  83,
Info: 10.38.66.83, "B6510_066_083"
   EX_Port       FID    Neighbor Switch Info (enet IP, WWN, name)
   ------------------------------------------------------------------
      40         40     10.38.66.92      10:00:00:27:f8:66:f3:81   "B6520_066_92"
      41         40     10.38.66.88      10:00:00:27:f8:06:23:28   "B6510_066_088"
      26         50     10.38.66.111     10:00:50:eb:1a:22:27:da   "VDX6740_066_111"
      27         50     10.38.66.112     10:00:50:eb:1a:20:d3:81   "VDX6740_066_112""VDX6740_066_119"
```

   fcrproxydevshow

```
> fcrproxydevshow
  Proxy           WWN             Proxy      Device   Physical    State
 Created                           PID       Exists     PID
in Fabric                                   in Fabric
```

```
-------------------------------------------------------------------
    40   10:00:00:05:1e:60:b4:6b  02ff02      100      132400   Imported
    40   10:00:00:05:1e:60:b4:6c  02fe02      100      032400   Imported
. . .
    40   50:02:4f:40:55:82:01:00  02f906      100      532400   Imported
    40   50:02:4f:40:55:82:01:01  02fa06      100      df2400   Imported
    40   50:02:4f:40:55:82:02:00  02fb06      100      532500   Imported
    40   50:02:4f:40:55:82:02:01  02fc06      100      df2500   Imported
    50   50:02:4f:40:55:82:01:00  02fc03      100      532400   Imported
    50   50:02:4f:40:55:82:01:01  02fd03      100      df2400   Imported
    50   50:02:4f:40:55:82:02:00  02fe03      100      532500   Imported
    50   50:02:4f:40:55:82:02:01  02ff03      100      df2500   ImportedTotal devices displayed: 20
```

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. Paths fail over successfully, and I/O completes without error.

5. Optional/Additional Tests

5.1 Storage Device Nondisruptive Firmware Update

Test Objective

Execute a nondisruptive firmware update on the array while running I/O, and confirm that there are no I/O errors.

Test Execution

1. Run continuous I/O to the array.
2. Execute the code update procedure as described in the vendor documentation.

Result Validation

1. Confirm the updated version on all array components.
2. Check the I/O generator tools logs to verify that I/O completes without error.
3. Check the host and storage logs for errors throughout the I/O operations.
4. Check the switch error logs and port stats for any errors or I/O drops.

   errdumpall

   porterrshow

```
> porterrshow
        frames      enc    crc    crc    too    too    bad    enc   disc   link   loss   loss   frjt   fbsy  c3timeout    pcs
     tx     rx      in    err    g_eof  shrt   long   eof    out   c3    fail    sync   sig                 tx    rx    err
  0:   1.7g   2.7g   0      0      0      0      0      0      0    946      1      0      1      0      0      0      0      0
  1: 431.1m 431.4m   0      0      0      0      0      0      0     30      1      0      1      0      0      0      0      0
  2:    0      0      0      0      0      0      0      0      0     0      0      0      0      0      0      0      0      0
  3:    0      0      0      0      0      0      0      0      0     0      0      0      0      0      0      0      0      0
```

Test Results

PASS. Update (from 5.2.2.26 to 5.2.4.23) is successful with proper failover and no I/O errors.

5.2 Performance I/O Testing with Varying Workload Dimensions

Test Objective

Run I/O workload generation with varying dimensions (block size, read/write/mix), and verify that performance characteristics are as expected.

Test Execution

1. Run an I/O loop at block transfer sizes of 512, 1k, 2k, 4k, 8k, 16k, 32k, 64k, 128k, 256k, 512k, and 1m.
2. Include a nested loop of 100% read, 100% write, and 50% read/write.

   Repeat the test for the following configurations:

   - One host port to one target port (single path)
   - Two host ports to multiple target ports (multipath)
   - Multiple hosts to multiple target ports (multihost multipath)

Result Validation

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

   errdumpall

   porterrshow

```
> porterrshow
          frames      enc    crc    crc    too    too    bad    enc   disc   link   loss   loss   frjt   fbsy  c3timeout    pcs
       tx     rx      in    err    g_eof  shrt   long   eof    out   c3    fail    sync   sig                tx    rx    err
  0:    1.7g   2.7g   0      0      0      0      0      0      0    946      1      0      1      0      0      0      0      0
  1:  431.1m 431.4m   0      0      0      0      0      0      0     30      1      0      1      0      0      0      0      0
  2:    0      0      0      0      0      0      0      0      0      0      0      0      0      0      0      0      0      0
  3:    0      0      0      0      0      0      0      0      0      0      0      0      0      0      0      0      0      0
```

Test Results

PASS. All workload runs are monitored at the host, storage, and fabric and complete without I/O errors or faults. Performance behavior is as expected.

5.3 Performance Testing with VMware I/O Analyzer Application Workload Suite

Test Objective

Run workload test suite including varying application workloads generated from multiple VMs.

Test Execution

1. Configure a two-host VMware cluster with multipath on two initiator ports per host, and four target ports.

2. Configure workload generation from eight worker VMs using VMware I/O Analyzer.
3. Run a variety of application I/O workload patterns from the I/O Analyzer suite. For this round of testing, the following workloads are used:
   - Workstation
   - Webserver
   - Video on Demand
   - Max Write Throughput
   - Max Write IOPS
   - Max IOPS
   - Max Throughput
   - SQL Server 16k
   - 4k Read
   - Exchange 2007
   - OLTP 4k

Result Validation

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

   errdumpall

   porterrshow

```
> porterrshow
         frames      enc    crc    crc    too    too    bad    enc   disc   link   loss   loss   frjt   fbsy  c3timeout    pcs
       tx     rx      in    err    g_eof  shrt   long   eof     out   c3    fail    sync   sig                tx    rx    err
  0:    1.7g   2.7g   0      0      0      0      0      0      0    946      1      0      1      0      0      0      0      0
  1:  431.1m 431.4m   0      0      0      0      0      0      0     30      1      0      1      0      0      0      0      0
  2:    0      0      0      0      0      0      0      0      0     0       0      0      0      0      0      0      0      0
  3:    0      0      0      0      0      0      0      0      0     0       0      0      0      0      0      0      0      0
```

Test Results

PASS. All workload runs are monitored at the host, storage, and fabric and complete without I/O errors.

5.4 Array Hardware Failure Tests

Test Objective

Validate HA and recovery from array hardware failures, including power and battery backup failures.

Test Execution

1. Run I/O from Medusa.
2. Pull a battery backup unit from one K-node; wait 30 seconds, and replace.
3. Disconnect one power cable from a K-node; wait 30 seconds, and replace.
4. Pull an Infiniband interconnect between K-nodes; wait 30 seconds, and replace.

5. Disconnect one power cable from a K-node; wait 30 seconds, and replace.

Result Validation

1. Check the I/O generator tool logs to verify that I/O completes without error.
2. Check the switch logs and port stats for any errors or I/O drops.
3. Monitor the array log for failure and recovery of hardware components.

```
10:54:06 AM Jun 3 kblock01-knode02-battery02 connectivity is OK. (SN: P206BCF24YB0160) HARDWARE
10:52:54 AM Jun 3 kblock01-knode02-battery02 is unreachable. (SN: P206BCF24YB0160) HARDWARE

2:28:21 PM Jun 3 kblock01-knode02-psu01 hardware is healthy. (SN: P7031CE30VT3143) HARDWARE
2:27:38 PM Jun 3 kblock01-knode02-psu01 hardware is in a critical health state. (SN: P7031CE30VT3143) HARDWARE

2:38:37 PM Jun 3 kblock01-knode01-ib02 link is now up. (SN: 0f73f2) HARDWARE
2:38:20 PM Jun 3 kblock01-knode02-ib02 link is now up. (SN: 112632) HARDWARE

5:11:43 PM Jun 3 Server kblock01-knode02 was automatically rebooted and phased in HARDWARE, STATE
5:11:43 PM Jun 3 Standby management setup completed on server kblock01-knode02 HARDWARE, RECOVER
5:11:31 PM Jun 3 kblock01-knode02-fc02 link status has changed to up. (SN: 6C4439432T) HARDWARE
5:11:01 PM Jun 3 kblock01-knode02-fc01 link status has changed to up. (SN: 6C4439432T) HARDWARE

system=>> server-show table=power
03-Jun-2016,16:57:38 MDT
SUCCESS

Servers Power Data:
-------------------
Name                      |Temperature|Health   |Connectivity|Manufacture date|Cycle count|Current |Charge level|Power in|Power out|
=============================================================================================================================================
kblock01-knode01-battery01|26 (ok)    |None (ok)|ok          |2015-05-23      |3 (ok)     |0.0 (ok)|99 (ok)     |        |         |
kblock01-knode01-battery02|27 (ok)    |None (ok)|ok          |2015-05-20      |3 (ok)     |0.0 (ok)|99 (ok)     |        |         |
kblock01-knode01-psu01    |40 (ok)    |None (ok)|ok          |                |           |        |            |118     |93       |
kblock01-knode01-psu02    |39 (ok)    |None (ok)|ok          |                |           |        |            |94      |82       |
kblock01-knode02-battery01|26 (ok)    |None (ok)|ok          |2015-05-23      |4 (ok)     |0.0 (ok)|100 (ok)    |        |         |
kblock01-knode02-battery02|27 (ok)    |None (ok)|ok          |2015-05-21      |4 (ok)     |0.0 (ok)|97 (ok)     |        |         |
kblock01-knode02-psu01    |39 (ok)    |None (ok)|ok          |                |           |        |            |96      |78       |
kblock01-knode02-psu02    |40 (ok)    |None (ok)|ok          |                |           |        |            |100     |75       |
```

Test Results

PASS. I/O completes without error. The array maintains availability and recovers from hardware failures.

Test Conclusions

1. Achieved a 100% pass rate on all test cases in the SFR qualification test plan. The network and the storage were able handle the various stress and error recovery scenarios without issue.
2. Different I/O workload scenarios were simulated using Medusa and VMware I/O Analyzer tools, and sustained performance levels were demonstrated across all workload types.
3. The results confirm that the Kaminario K2 array interoperates seamlessly with Brocade Fibre Channel fabrics, and together they demonstrate high availability, performance, and low latency.
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. We recommend that you 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. You should also enable the additional MAPS feature of Fabric Performance Impact monitoring to detect bottlenecks in the form of timeouts and latency. Using MAPS is recommended to maximize the benefit of high-performance, low-latency storage.
6. Implementing the IO Insight feature to monitor the critical flows at the SCSI level can provide valuable data and insight into traffic performance across the fabric. Combining IO Insight with MAPS custom rules and alerts can provide proactive monitoring to help preserve the investment in flash storage performance.
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. Use QoS zoning to classify host-target traffic into high, medium, or low priority zones and 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. For optimal availability and performance, consideration should be given to multipath configuration on the host side. While Windows 2008 and 2012 will provide Round-Robin behavior by default, Linux systems will benefit from adding a custom entry to /etc/multipath.conf, and VMware hosts systems should be changed from the default Most Recently Used (VMware) setting to the Round-Robin (VMware) setting. Actively using all available paths provides a significant improvement in performance throughput.