VALIDATION TEST REPORT



Brocade Fabric Technology with the Kaminario K2 All-Flash Array

Supporting Fabric OS 8.0.1, 8.1.0, 8.2.0, and 8.2.1

Copyright © 2018 Broadcom. All Rights Reserved. Broadcom, the pulse logo, Brocade, and the stylized B logo are among the trademarks of Broadcom in the United States, the EU, and/or other countries. The term "Broadcom" refers to Broadcom Inc. and/or its subsidiaries.

Broadcom reserves the right to make changes without further notice to any products or data herein to improve reliability, function, or design. Information furnished by Broadcom is believed to be accurate and reliable. However, Broadcom does not assume any liability arising out of the application or use of this information, nor the application or use of any product or circuit described herein, neither does it convey any license under its patent rights nor the rights of others.

The product described by this document may contain open source software covered by the GNU General Public License or other open source license agreements. To find out which open source software is included in Broadcom products, view the licensing terms applicable to the open source software, and obtain a copy of the programming source code, please visit https://www.broadcom.com/.

Contents

Preface	5
Document History	5
Overview	5
Purpose of This Document	5
Audience	
Objectives	6
Related Documents	6
About Broadcom	6
About Kaminario	6
Configure DUT and Test Equipment	7
Task 1. Brocade FC Fabric Configuration	7
Task 2. Kaminario Array Configuration	
Task 3. Host Setup	
Test Report	
What's New in This Report	
Test History	
Test Plan Overview	
Scope	
Test Configuration	20
DUT Descriptions	
Test Cases	23
1. Fabric Initialization—Base Functionality	24
1.1 Storage Device—Physical and Logical Login with Speed Negotiation	24
1.2 Zoning and LUN Mapping	24
1.3 Storage Device Fabric I/O Integrity	
1.4 Storage Device Multipath Configuration—Path Integrity	25
2. Fabric—Advanced Functionality	
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	
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	
3.0 Storage Device Recovery-ISL Port Toggle (Entire Switch)	
3.7 Storage Device Recovery – Director Bidde Maintenance	40 11
3.9 Storage Device Recovery—Switch Firmware Download	
4. Storage Device Elbre Channel Douting (ECD) Internetworking Tests	24 12
4.1. Storage Device Internetworking Validation with EC Host	43. ۱۷

4.2 Storage Device Edge Recovery After FCR Disruptions	
4.3 Storage Device Backbone Recovery After FCR Disruptions	46
5. Optional/Additional Tests	47
5.1 Storage Device Nondisruptive Firmware Update	47
5.2 Performance I/O Testing with Varying Workload Dimensions	48
5.3 Performance Testing with VMware I/O Analyzer Application Workload Suite	
5.4 Array Hardware Failure Tests	
Test Conclusions	51

Preface

•	Document History	
•	Overview	5
•	Purpose of This Document	5
•	Audience	5
•	Objectives	6
•	Related Documents	6
•	About Broadcom	6
•	About Kaminario	6

Document History

Date	Part Number	Description
July 29, 2016	53-1004539-01	Supporting FOS 8.0.1.
April 28, 2017	53-1004539-02	Supporting FOS 8.1.0.
March 16, 2018	53-1004539-03	Supporting FOS 8.2.0.
November 12, 2018	53-1004539-04	Supporting FOS 8.2.1.

Overview

The Storage Fabric Ready (SFR) program is a comprehensive testing and configuration initiative to provide Fibre Channel (FC) SAN interoperability with flash storage. This program provides testing of multiple fabrics, heterogeneous servers, 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, softwaredefined 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
- Brocade SAN Design and Best Practices
- Brocade SAN Fabric Resiliency and Administration Best Practices
- Brocade Monitoring and Alerting Policy Suite Configuration Guide
- Brocade Fabric OS Command Reference Manual
- Brocade Flow Vision Configuration Guide
- Emulex OneCommand Manager Application User Manual
- QLogic QLE2672 Adapter User's Guide
- QLogic BR-1860 Adapter Administrator's Guide

About Broadcom

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

About 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	7
•	Task 2. Kaminario Array Configuration	11
•	Task 3. Host Setup	12

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
```

- 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*.
 - 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:

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.

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 fixed-port 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 _____ |Predefined |Type |Member Count |Members Group Name ------_____ 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 ·
 Quarantined Ports :
                                                      None
 Top Zoned PIDs <pid(it-flows)>: 0xdf2000(30) 0xdf2100(30) 0xdf1400(28) 0xdf0400(28) 0xdf1a00(26)
```

- 4. Configure Fibre Channel Routing (requires an Integrated Routing license). Detailed information on FCR setup can be found in the *Brocade Fabric OS Administration Guide*.
 - 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 example above shows that the prefix "Isan" is used when configuring zones for use in Fibre Channel Routing.

b) A second example FCR configuration is shown below:

The example above displays a zone prefixed with "Isan."

c) Example output of exported devices:

> fcrpro: Proxy Created in Fabrio	xydevshow WWN	Proxy PID	Device Exists in Fabric	Physical PID	State
40	50:02:4f:40:55:82:01:00	02£906	100	532400	Imported
40	50:02:41:40:55:82:01:01	02±a06	100	d12400	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

kamınarıd	2.H	(2					Se	rver time: 4:45 PM	Sy
		Dashboard	Performance	Events	Volumes & Hosts	DataProtect	S	ystem	
		Hosts Volumes						1% of Capacity	Used
	-	Host Group	Create Host				×		
6		Host	Name:	ssr067178)				
	₽	esxi_cluster						_3	
	⊳	🧊 ssr067116	Туре:	👌 Linux	•		3	3	••••
	⊳	🗊 ssr067125	in a second				9	3	
	⊳	🧊 ssr067136	Ports:	No ports defined			3	3 🔒	•••
	⊳	🧊 ssr067138	+ Add	ł			3	3	
	⊳	🗊 ssr067168	1	Manually enter ports			r	odd4 🔋	
	⊳	🧊 ssr067172		10:00:8c:7c:ff:23:03:00	Cancel				
	⊳	🧊 ssr067174		U ssr06/1/4_1	🤘 ssr06717	'4_2 🗎 ss	r067174_3	3 🔒	
	⊳	🗊 ssr067176	-3	ssr067176_qos					
	_								

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

kamına	rio.k2					Server time: 4:53 PM	System: K2-55
	Dashboard	Performance	Events	Volumes & Hosts	DataProtect	System	
	Hosts Volumes					1% of	Capacity Used
	• Volume Group + Volume	e					
	Volume	Capacity		Mapped to Hosts			
	u vg1_dd	288GB 0T 440GB					_
Ň	📦 vg2_no-dd	122GB of 700GB	•				
	ssr067116_datastore	7GB of 500GB	\rightarrow	🧊 ssr067116			+ Map
	ssr067168-nodd2	18.9GB of 20GB	\rightarrow	🧊 ssr067168			+ Map
	ssr067168-nodd3	18.9GB of 20GB	\rightarrow	ssr067168			+ Map
	ssr067168-nodd4	18.8GB of 20GB	\rightarrow	🧊 ssr067168			+ Map
	Ssr067168-nodd5	6.4GB of 20GB	\rightarrow	issr067168			+ Map
	ssr067168-nodd6	6.4GB of 20GB	\rightarrow	issr067168			+ Map
	Ssr067168-nodd7	6.5GB of 20GB	\rightarrow	🧊 ssr067168			+ Map
	ssr067168-nodd8	6.3GB of 20GB	\rightarrow	i ssr067168			+ Map
	ssr067168-nodd9	6.6GB of 20GB	→	ssr067168			+ Map
	ssr067168-nodd10	6.4GB of 20GB	\rightarrow	ssr067168			+ Map
	ssr067168_no-dd	20GB of 20GB	\rightarrow	ssr067168			+ Map

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 -11
```

```
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 sdbo 68:32 active ready running |- 1:0:1:0 sdbo 68:32 |- 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 -|- 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

and the second s	MPIO P	roperties		
MPIO Devices	Discover Multi-Paths	DSM Install	Configuration Snap	shot
To add suppo Product Ids a Devices can b To remove su then click Rem	rt for a new device, cli s a string of 8 characte specified using semi- pport for currently MPI nove.	ck Add and en ers followed by colon as the d	nter the Vendor and y 16 characters. Mult lelimiter. select the devices an	iple d
Device Hard	ware Id			^
NMINKIO KZ				≡
	2			
<			>	~
		Add	Remove	
			OK Can	cel

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

Path Selection:	Round Robin (VMware)			- Chang	e
Storage Array Typ	pe: VMW_SATP_DEFAULT_AA		-		
aths					
Runtime Name	Target	LUN	Status	Preferred	
vmhba3:C0:T1:L3	7 50:02:4f:40:55:82:00:00 50:02:4f:40:55:82:01:01	37	Active (I/O)		
vmhba2:C0:T0:L3	7 50:02:4f:40:55:82:00:00 50:02:4f:40:55:82:01:00	37	Active (I/0)		
vmhba3:C0:T2:L3	7 50:02:4f:40:55:82:00:00 50:02:4f:40:55:82:02:00	37	Active (I/O)		=
vmhba2:C0:T1:L3	7 50:02:4f:40:55:82:00:00 50:02:4f:40:55:82:01:01	37	Active (I/O)		
vmhba3:C0:T3:L3	7 50:02:4f:40:55:82:00:00 50:02:4f:40:55:82:02:01	37	Active (I/O)		
vmhba2:C0:T2:L3	7 50:02:4f:40:55:82:00:00 50:02:4f:40:55:82:02:00	37	Active (I/0)		
vmhba2:C0:T3:L3	7 50:02:4f:40:55:82:00:00 50:02:4f:40:55:82:02:01	37	🔶 Active (I/0)		-
				Refr	esh
Name:	vmhba3:C0:T1:L37				_
Runtime Name:	vmhba3:C0:T1:L37				
Fibre Channel					
Adapter:	20:00:00:0e:1e:1b:f9:11 21:00:00:0e:1e:1b:f9:11				
Target:	50:02:4f:40:55:82:00:00 50:02:4f:40:55:82:01:01				

- 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:

10.38.67.192 - Running validation for Disk.SchedNumReqOutstanding

```
10.38.67.192 - Note - In ESXi version 6.0.0 the SchedNumRegOutstanding 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
(gl2xintrdelaytimer=1 gl2xoperationmode=6 gl2xmaxgdepth=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	19
•	Test History	19
•	Test Plan Óverview	19
•	Scope	19
•	Test Configuration	20
•	DUT Descriptions	21

What's New in This Report

- The Brocade Fabric OS (FOS) version under test is 8.2.1.
- Emulex and QLogic adapters have updated firmware and drivers.

See the DUT Description tables below for detailed information.

Test History

Storage Model	Storage Firmware	Brocade FOS Version	Date
Kaminario K2	5.2.4.23	FOS 8.0.1	May 2016
Kaminario K2	5.8.0.83	FOS 8.1.0	February 2017
Kaminario K2	5.8.7.12	FOS 8.2.0	February 2018
Kaminario K2	5.8.7.12	FOS 8.2.1	November 2018

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



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

Test Cases

	1. Fabric Initialization—Base Functionality
	2. Fabric—Advanced Functionality
	3. Stress and Error Recovery with Device Multipath
	4. Storage Device—Fibre Channel Routing (FCR) Internetworking Tests
	5. Optional/Additional Tests
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 Edge Recovery After FCR Disruptions
4.3	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
                        76
port generation number:
state transition count:
                       13
         df2400
portId:
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 **Isscsi**.
- 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.

- 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
                        76
13
portPhys: 6
              In Sync
                              portScn: 32 F Port
port generation number:
state transition count:
portId: df2400
portIfId: 4302
           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.

- 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

> mapsdbshow												
1 Dashboard Information:												
Active policy:	dflt_aggressive_poli	су										
2 Switch Health Report:												
Current Switch Policy Sta	======================================											
3.1 Summary Report:												
Category	Today	Last 7 days										
Port Health Fru Health Security Violations Fabric State Changes Switch Resource Traffic Performance FCIP Health Fabric Performance Impact	No Errors In operating range No Errors In operating range In operating range Not applicable :Out of operating range	Out of operating range In operating range Not applicable 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.

- 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
                                ITodav
                                                               |Last 7 davs
                                                                                               _____
                                                             |Out of operating range
|In operating range
    Port Health
                               |No Errors
   Fort healthNo ErrorsFort of operating rangeFru HealthIn operating rangeIn operating rangeSecurity ViolationsNo ErrorsIn operating rangeFabric State ChangesNo ErrorsIn operating rangeSwitch ResourceIn operating rangeIn operating rangeTraffic PerformanceIn operating rangeIn operating rangeFCIP HealthNot applicableNot 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

      Current Value:
      371816 Microseconds

      Dashboard Category:
      Traffic Performance

      Switch Name:
      G620 066 0000

                                                                                         ry: Traffic Performance
G620_066_223
10:00:c4:f5:7c:2a:8b:c8
10.38.66.223
    Switch WWN:
Switch IP:
                                                                               .38.
SSR
128
     Fabric Name:
      VFID:
```

MAPS Dashboard:

> mapsdb -show 3.2 Rules Affecting Health: Category(Rule Count) |RepeatCount|Rule Name |Execution Time |Object |Triggered Value(Units)| Traffic Performance(|3 lios_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

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: QOSM_ssr067178 kaminario
```

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

porterrshow

> p	orterrshow	Ň																
	fra	ames	enc	crc	crc	too	too	bad	enc	disc	link	loss	loss	frjt	fbsy	c3tim	eout	pcs
	tx	rx	in	err	g_eof	shrt	long	eof	out	с3	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 5 Configuring Driver Parameters in the Emulex OneCommand Interface

File Edit View Port Discovery Batch Help Image: Constraint of the second state of	*			OneCommand [™] N
Image: With the second seco	File Edit View Port Discovery Batch Help			
Hosts Hosts Port Information Statistics Maintenance Target Mapping Driver Parameters [] Heb67125 Installed Driver Type: elxfc Heb67125 Installed Driver Type: elxfc Heb67126 Adapter Parameter Value Adapter Parameter Value AutoDelay 1 AutoDelay 1 AutoDelay 1 AutoDelay 1 AutoDelay 1 AutoMap 2 Class 3 ConfigScale 4 DriverTraceMask 0x0 EnableMDV Disabled		Find Host:	~	?
Adapter Parameter Value Adapter Parameter Value AutoDelay 1 AutoMap 2 Class 3 ConfigScale 4 DriverTraceMask 0x0 Enable/Lone 1 ExtTransferSize 0 FrameSizeMSB 0 LinkTimeOut 15 LinkTomeOut 30 PerfMode 0 PerPortTrace 0 RumaDepth 16 ScanDown 1 TargetOption 0 TraceBufSiz 2000000 XLanePriority 3	Hosts HB067125 HE 16202-X	Port Information Statistics Installed Driver Type: elx	Maintenance Ta	arget Mapping Driver Parameters
AutoDelay 1 AutoDelay 1 AutoMap 2 Class 3 ConfigScale 4 DriverTraceMask 0x0 EnableFDMI 0 EnableKNPIV Disabled EnableKNEV Disabled EnableKNEV Disabled EnableKNEV Disabled EnableKNEV 0 FrameSizeMSB 0 InitTimeOut 15 UnikSpeed Auto Detect LinkTimeOut 30 LogErrors 3 MaxEQDelay 256 MistXMode 0 NodeTimeOut 30 PerfNode 0 QueueTarget 0 RmaDepth 16 ScanDown 1 TargetOption 0 Topology 2 TraceBufsiz 2000000 XLanePriority 3		Adapter Parameter	Value	Modify Adapter Parameter
AutoNap 2 Class 3 ConfigScale 4 DriverTraceMask 0x0 EnableFDMI 0 EnableRDMI 0 EnableRDMI 0 EnableXLane 1 ExtTransferSize 0 FrameSizeMSB 0 InitTimeOut 15 LinKTimeOut 30 LogErrors 3 MaxEQDelay 256 MixEQDelay 26 None. Parameter is dynamic QueueTarget 0 QueueTarget 0 RmaDepth 16 ScanDown 1 Topology 2 TraceBufSiz 2000000 XLanePriority 3	Engle Port 1. 10.00.00.50.1 A.01.00.25	AutoDelay	1	
Class 3 ConfigScale 4 DriverTraceMask 0x0 EnableFDMI 0 EnableNPIV Disabled EnableNPIV Disabled EnableXLane 1 ExtTransferSize 0 FrameSizeMSB 0 LinitTimeOut 15 LinitTimeOut 30 LogErrors 3 MaxEQDelay 256 MSIXMode 0 NodeTimeOut 30 PerPortTrace 0 QueueEpth 32 QueueEarget 0 RmaDepth 16 ScanDown 1 TargetOption 0 Topology 2 TraceBufSiz 2000000 XLanePriority 3		AutoMap	2	
ConfigScale4DriverTraceMask0x0EnableFDMI0EnableFDMI0EnableFDMI0EnableKILane1ExtTransferSize0FrameSizeMSB0InitTimeOut15LinkSpeedAuto DetectLinkSpeedAuto DetectLinkSpeedAuto DetectNone. Parameter is dynamicMaxEQDelay256MSDXMode0NodeTimeOut30PerfMode0PerfNode0QueueDepth32QueueTarget0RmaDepth16ScanDown1TargetOption0TargetOption0TargetOption0Make change temporary		Class	3	Parameter: XLanePriority
DriverTraceMask 0x0 EnableFDMI 0 EnableFDMI 0 EnableNPIV Disabled EnableNot Disabled EnableNot Disabled IntTimeOut 30 LogErrors 3 Makeopeth 32 QueueDapth 16 ScanDown 1 TargetOption 0 Topology 2 TraceBufSiz		ConfigScale	4	
EnableFDMI0EnableNPIVDisabledEnableXLane1ExtTransferSize0FrameSizeMSB0InitTimeOut15LinkSpeedAuto DetectLinkTimeOut30LogErrors3MaxEQDelay256MSIXMode0None. Parameter is dynamicPerfMode0PerfMode0Rangett16ScanDown1TargetOption0Topology2TraceBufSiz2000000XLanePriority3Make change temporary		DriverTraceMask	0x0	Value: 3
EnableNPIVDisabledEnableXLane1ExtTransferSize0FrameSizeMSB0ImTransferSize0LinkSpeedAuto DetectLinkTimeOut30LogErrors3MaxEQDelay256MSIXMode0NodeTimeOut30PerfMode0PerPortTrace0QueueTarget0RmaDepth16ScanDown1TargetOption0Topology2TraceBufSiz2000000XLanePriority3		EnableFDMI	0	
EnableXLane1ExtTransferSize0FrameSizeMSB0InitTimeOut15LimKTmeOut30LogErrors3MaxEQDelay256MSIXMode0NodeTimeOut30PerfMode0QueueTarget0QueueTarget0RmaDepth16ScanDown1TargetOption0Topology2TraceBufSiz2000000XLanePriority3		EnableNPIV	Disabled	
ExtTransferSize 0 FrameSizeMSB 0 InitTimeOut 15 LimTransferSize 0 LinkSpeed Auto Detect LinkTimeOut 30 LogErrors 3 MaxEQDelay 256 MSDXMode 0 NodeTimeOut 30 PerfMode 0 QueueDarget 0 QueueTarget 0 RmaDepth 16 ScanDown 1 TargetOption 0 Topology 2 TraceBufSiz 2000000 XLanePriority 3		EnableXLane	1	Range: 0 - 127
FrameSizeMSB 0 InitTimeOut 15 LimTransferSize 0 LinkSpeed Auto Detect LinkTimeOut 30 LogErrors 3 MaxEQDelay 256 MSIXMode 0 NodeTimeOut 30 PerfMode 0 PerPortTrace 0 QueueDepth 32 QueueTarget 0 RmaDepth 16 ScanDown 1 TargetOption 0 TraceBufSiz 2000000 XLanePriority 3		ExtTransferSize	0	
InitTimeOut 15 LimTransferSize 0 LinkSpeed Auto Detect LinkTimeOut 30 LogErrors 3 MaxEQDelay 256 MSIXMode 0 NodeTimeOut 30 PerfMode 0 PerPortTrace 0 QueueDepth 32 QueueTarget 0 RmaDepth 16 ScanDown 1 TargetOption 0 Topology 2 TraceBufSiz 2000000 XLanePriority 3		FrameSizeMSB	0	
LimTransferSize 0 LinkSpeed Auto Detect LinkTimeOut 30 LogErrors 3 MaxEQDelay 256 MSIXMode 0 NodeTimeOut 30 PerfMode 0 PerPortTrace 0 QueueDepth 32 QueueTarget 0 RmaDepth 16 ScanDown 1 TargetOption 0 Topology 2 TraceBufSiz 2000000 XLanePriority 3 Make change temporary		InitTimeOut	15	Default: 0
LinkSpeed Auto Detect LinkTimeOut 30 LogErrors 3 MaxEQDelay 256 MSIXMode 0 NodeTimeOut 30 PerfMode 0 PerPortTrace 0 QueueDepth 32 QueueTarget 0 RmaDepth 16 ScanDown 1 TargetOption 0 Topology 2 TraceBufSiz 2000000 XLanePriority 3		LimTransferSize	0	
LinkTimeOut 30 LogErrors 3 MaxEQDelay 256 MSIXMode 0 NodeTimeOut 30 PerfMode 0 PerPortTrace 0 QueueDepth 32 QueueTarget 0 RmaDepth 16 ScanDown 1 TargetOption 0 Topology 2 TraceBufSiz 200000 XLanePriority 3 Make change temporary		LinkSpeed	Auto Detect	Activation Requirements
LogErrors 3 MaxEQDelay 256 MSIXMode 0 NodeTimeOut 30 PerfMode 0 QueueDepth 32 QueueTarget 0 RmaDepth 16 ScanDown 1 TargetOption 0 Topology 2 TraceBufSiz 2000000 XLanePriority 3		LinkTimeOut	30	None Parameter is dynamic
MaxEQDelay 256 MSIXMode 0 NodeTimeOut 30 PerfMode 0 PerPortTrace 0 QueueDepth 32 QueueTarget 0 RmaDepth 16 ScanDown 1 TargetOption 0 TraceBufSiz 2000000 XLanePriority 3		LogErrors	3	None. Parameter is dynamic
MSIXMode 0 NodeTimeOut 30 PerfMode 0 PerPortTrace 0 QueueDepth 32 QueueTarget 0 RmaDepth 16 ScanDown 1 TargetOption 0 TraceBufSiz 2000000 XLanePriority 3		MaxEQDelay	256	
NodeTimeOut 30 PerfMode 0 PerPortTrace 0 QueueDepth 32 QueueTarget 0 RmaDepth 16 ScanDown 1 TargetOption 0 TraceBufSiz 2000000 XLanePriority 3		MSIXMode	0	
PerfMode 0 PerPortTrace 0 QueueDepth 32 QueueTarget 0 RmaDepth 16 ScanDown 1 TargetOption 0 TraceBufSiz 2000000 XLanePriority 3		NodeTimeOut	30	
PerPortTrace 0 QueueDepth 32 QueueTarget 0 RmaDepth 16 ScanDown 1 TargetOption 0 TraceBufSiz 2000000 XLanePriority 3		PerfMode	0	
QueueDepth 32 QueueTarget 0 RmaDepth 16 ScanDown 1 TargetOption 0 TraceBufSiz 2000000 XLanePriority 3		PerPortTrace	0	Description
QueueTarget 0 RmaDepth 16 ScanDown 1 TargetOption 0 Topology 2 TraceBufSiz 2000000 XLanePriority 3		QueueDepth	32	Expressione CS_CTL Priority
RmaDepth 16 ScanDown 1 TargetOption 0 Topology 2 TraceBufSiz 2000000 XLanePriority 3		QueueTarget	0	
ScanDown 1 TargetOption 0 Topology 2 TraceBufSiz 200000 XLanePriority 3		RmaDepth	16	
TargetOption 0 Topology 2 TraceBufSiz 2000000 XLanePriority 3		ScanDown	1	
Topology 2 TraceBufSiz 2000000 XLanePriority 3 Make change temporary		TargetOption	0	
TraceBufSiz 2000000 XLanePriority 3 Make change temporary		Topology	2	
XLanePriority 3 Make change temporary		TraceBufSiz	2000000	
		XLanePriority	3	Make change temporary

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 6 Enabling LUNs for ExpressLane in the Emulex OneCommand Interface

\$		OneCommand [™] Mar
File Edit View Port Discovery Batch Help		
Image: Second	✓ ●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●●<td></td>	
Hosts	LUN Information	
55R007125	Vendor/Product Information	
Port 0: 10:00:00:90:FA:61:8D:24		
= 50:02:4F:40:55:82:01:00	Vendor Name	KMNRTO
		13-11-12-0
LUN 17		
LUN 19	Product ID:	K2
	Alternation of the second s	
	Revision:	0000
11 LUN 24		
= 50:02:4F:40:55:82:01:01		
	Mapping Information	
LUN 18		
	FCP LUN:	00 11 00 00 00 00 00 00 00
LUN 21		
	SCSI OS LUN:	17
LUN 24		
50:02:4F:40:55:82:02:00	OS Device Name:	\\PhysicalDrive2
LUN 17		
Enable ExpressLane LUN		
Disable all ExpressLane LUNs on this ta	rget	
	CONTRACTOR OF A	
	Capacity: 10240 MB Block Size: 512 Bytes	
EUN 24		
Port 1: 10:00:00:90:FA:61:8D:25		
50:02:4F:40:55:82:01:00		
	- LUN Masking Information	
H = 50:02:4F:40:55:82:02:00 Medusa Labs Test T	ools v7.2.0.169914	
B. B. DOIDSIAL HOIDDIOSIDSID		

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

> porterrsho

	fra	ames	enc	crc	crc	too	too	bad	enc	disc	link	loss	loss	frjt	fbsy	c3tir	meout	pcs
	tx	rx	in	err	g_eof	shrt	long	eof	out	с3	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

> bor	retranov	N																
	fra	ames	enc	crc	crc	too	too	bad	enc	disc	link	loss	loss	frjt	fbsy	c3tir	neout	pcs
	tx	rx	in	err	g eof	shrt	long	eof	out	с3	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
                              portScn: 32 F Port
             In Sync
                          76
port generation number:
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 -II

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.

- 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
                          13
state transition count:
portId:
          df2400
portia: di2400
portIfId: 4302001e
          20:24:c4:f5:7c:2a:8b:c8
portWwn:
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 -II

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.

- 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
                        7ĸ
13
portPhys: 6
              In Sync
                               portScn: 32 F Port
port generation number:
state transition count:
portId: df2400
portIfId: 4302001e
         20:24:c4:f5:7c:2a:8b:c8
portWwn:
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 -II

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

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

-	fra	ames	enc	crc	crc	too	too	bad	enc	disc	link	loss	loss	frjt	fbsy	c3ti	meout	pcs
	tx	rx	in	err	g eof	shrt	long	eof	out	с3	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

> porterrshow

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

> fak	oricshow	v.			
Swite	ch 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 H	Fabric h	nas 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

± .																		
	fra	ames	enc	crc	crc	too	too	bad	enc	disc	link	loss	loss	frjt	fbsy	c3ti	meout	pcs
	tx	rx	in	err	g eof	shrt	long	eof	out	с3	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

> porterrshow

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

> fak Switc	oricshoù ch ID	v 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:fed7:1	a04
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:fe44:3	362
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

	fr	ames	enc	crc	crc	too	too	bad	enc	disc	link	loss	loss	frjt	fbsy	c3tin	neout	pcs
	tx	rx	in	err	g eof	shrt	long	eof	out	с3	fail	sync	sig			tx	rx	err
0:	1.7g	2.7g	0	0	ō	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.

- 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

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

errdumpall

switchshow

```
> switchshow
               B6510 066 088
switchName:
switchType:
               109.1
switchState: Online
                         <--
             Native
switchMode:
switchRole:
               Subordinate
switchDomain: 88
switchId: fffc58
switchWwn: 10:00:00:27:f8:06:23:28
              ON (SSR)
zoning:
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.

- 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
                                                           _____
                                    _____
                                                                                                  _____
                                                                                0.0.0.0
  1: fffc01 50:00:53:35:b1:d3:df:5c 0.0.0.0
                                                                                                                 "fcr xd 1 40"
                                                                                                                 "fcr_xd_4_50"
   4: fffc04 50:00:53:35:b1:d3:df:5e 0.0.0.0
                                                                                      0.0.0.0

      4: IIIC04 50:00:53:35:51:d3:d1:56 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:a5:51:86 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 10:00:00:05:33:55:1d:1d 10.38.66.83
      0.0.0.0
      "B6510_066_083"

                                                                                                                 "G620 066 223"
223: fffcdf 10:00:c4:f5:7c:2a:8b:c8 10.38.66.223 0.0.0.0
224: fffce0 10:00:c4:f5:7c:41:1b:14 10.38.66.224
                                                                                                                 "G620 066 224"
                                                                                       0.0.0.0
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

(Testing covered in FOS 8.1.0 and previous releases.)

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

fcrproxydevshow

> fcrprox	ydevshow							
Proxy Created in Fabric	WWN	Proxy PID	Device Exists in Fabric	Physical PID	State			
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	02£906	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 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

fcrproxydevshow

> :	icrpro	xydevshow							
C in	Proxy reated Fabri	WWN I C	Proxy PID	Device Exists in Fabric	Physical PID	State			
	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			
·	• •		005005	100	500400				
	40	50:02:41:40:55:82:01:00	021906	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.

4.3 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

fcrproxydevshow

> :	fcrprox	ydevshow							
C: Ci in	Proxy reated Fabric	WWN	Proxy PID	Device Exists in Fabric	Physical PID	State			
	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

> po	orterrsho	W																
	fr	ames	enc	crc	crc	too	too	bad	enc	disc	link	loss	loss	frjt	fbsy	c3tir	neout	pcs
	tx	rx	in	err	g eof	shrt	long	eof	out	с3	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 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 errors or I/O drops.

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.

- 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

> po	rterrshow	w																
	fra	ames	enc	crc	crc	too	too	bad	enc	disc	link	loss	loss	frjt	fbsy	c3tir	neout	pcs
	tx	rx	in	err	g_eof	shrt	long	eof	out	с3	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	Ter	nperature	e Healt	h	Connect	ivity Manufacture	date Cy	cle count	Curr	cent	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)	- I	1	- I
kblock01-knode01-psu01	40	(ok)	None	(ok)	ok	1	1		1	1		118	93	- I
kblock01-knode01-psu02	39	(ok)	None	(ok)	ok	I	1		1	1		94	82	1
kblock01-knode02-battery01	26	(ok)	None	(ok)	ok	2015-05-23	4	(ok)	0.0	(ok)	100 (oł	c)	1	- I
kblock01-knode02-battery02	27	(ok)	None	(ok)	ok	2015-05-21	4	(ok)	0.0	(ok)	97 (ok)		1	- I
kblock01-knode02-psu01	39	(ok)	None	(ok)	ok	I	1		1	1		96	78	- I
kblock01-knode02-psu02	40	(ok)	None	(ok)	ok	1	1		1	1		100	75	1

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.