

EANTC Independent Test Report

Huawei WDM Compatibility Test

OptiX OSN 1800V Pro

April 2022



Introduction

Huawei commissioned EANTC to validate functional, interoperability, and performance aspects of the Huawei OptiX OSN 1800 wave-division multiplexing (WDM) solution with a specific focus on the Storage Area Network (SAN) use case scenarios. Three board types under test participated in the test: Huawei LTX, LDCA, and MDCA. We conducted all tests in the Huawei lab in Cheng Du, Berlin, in December 2020. Huawei OptiX OSN 1800 V Pro is an optical-electrical WDM transmission device. Huawei OptiXtrans E6616 is a product of the same brand. Both devices provide the same implementation. It is designed for enterprise DCI use case scenarios — any situation where two redundant data centers are located in a region within a few kilometers distance (in fact, up to 100 km, but we will see that later).

This report is an extension to the Huawei OptiX OSN 1800 test which we have conducted a few weeks ago. This follow-up report includes new features like design and storage clusters. If interconnected data centers include a redundant design of the storage cluster with the background traffic, can the WDM system maintain all the performance and resiliency results unchanged when providing DCI interaction?

Focusing on this topic, we added a quorum server into the current emulated data centers, to implement data redundancy between SANs of the data centers. While maintaining all test goals of this test series such as performance and resiliency unchanged, the quorum server synchronized data between storage devices. Generally, a quorum server serves as a redundant approach in clusters to maintain data integrity. A specific example explains that if a quorum server is not there, as mentioned in the scenario from the previous report, at where two hosts separately accessed their storage which was located in a remote data center, their I/O operations during the process led to two independent copies of the data. Although the copies were of the same data, there may be a risk of asymmetric data writing (split brain) by introducing a failure on the WDM link between data centers. The result of such observation was that one host continued with I/O operation, while the host's I/O on the other side stopped during the network convergency. Although this observation was a design for storage, it should not have any impact on the network side, which was confirmed in the previous report. To achieve an end-to-end observation with data consistency by including the quorum server, we expected that the WDM system remains all performance and resiliency results to prove its transparency between data centers.

In December 2021, we performed the same test with a newer software version with the following boards: ELOM, LDCA, and LTX. As we have performed the full test suite on LDCA and LTX last year and the only difference is the software version, Huawei suggests only testing the compatibility of LDCA and LTX with the latest Brocade G620 software version. Therefore, LDCA and LTX were not tested for the multi-switch transparent test and 24-hour stability test. ELOM had been tested with the full test suite except the ISL trunking function because ELOM is a low speed board.

Executive Summary

In 2020, we verified the interface-speed forwarding of Huawei OptiX OSN 1800 V Pro on the FC port side and the long-distance forwarding between the two simulated data centers. We tested the forwarding performance when a quorum server synchronized data traffic between both emulated data centers. We also verified the compatibility of its three board types LTX, LDCA, and MDCA, to the FC-PI-6, FC-PI-3, and FC-PI-5 standards.

The test bed consisted of emulated data centers integrated with 3rd party Fibre Channel switches of different vendors represented by Brocade and Cisco, using hybrid switch pairs of Brocade G620/Brocade G620, Brocade G620/Brocade 6505, and Cisco MDS 9132T/Cisco MDS 9132T, Cisco MDS 9132T/Cisco MDS 9148S, respectively.

We verified the robustness of the DUT by performing administrative activities on the DUT and connected equipment, as well as ISL trunking and protection against long-haul link failure. We also put DUT under continuous load for 24 hours in a soak test environment to make sure it would support uninterrupted service. The system remained stable without any restart or service interruption, zero packet loss, and low latency, as expected. Finally, we measured the latency introduced by the DC908. It matched the expectations based on switching delay and physical distance.

In 2021, we performed the same test on the same test bed with three different boards: ELOM, LDCA, and LTX, as we have verified the capacity and stability of LDCA and LTX with the older version. Huawei suggested testing only the compatibility of LDCA and LTX with the latest Brocade G620 software version (v9.0.1c). Therefore, LDCA and LTX were not tested for multi-switch type and 24-hour stability test. Only ELOM had been tested with the full test suite except ISL trunking function.

Highlights 2020	Highlights 2021
→ DCI interoperability with four Fibre Channel switches, including the combination of Brocade G620-G620, G620-6505, Cisco MDS 9132T - 9132T, and Cisco MDS 9132T - 9148S	→ DCI interoperability with two Fibre Channel switches, including the combination of Brocade G620 - G620, G620-6505, and Cisco MDS 9148S - 9148S
→ Compatibility certification with three types of Fibre Channel Physical Interface (FC-PI), including FC-PI-3, FC-PI-5, and FC-PI-6 ¹	→ Compatibility certification with three types of Fibre Channel Physical Interface (FC-PI), including FC-PI-3, FC-PI-5, and FC-PI-6 ¹
→ Transparent multi-switch type forwarding between Brocade G620 and Brocade G620, Brocade 6505 and Brocade G620, as well as between Cisco MDS 9132T and Cisco MDS 9132T, Cisco MDS 9132T and Cisco MDS 9148S, respectively	→ Transparent multi-switch type forwarding between Brocade 6505 and Brocade G620 ³
→ Stability of overnight soak testing ²	→ Protection against Inter-Switch Link (ISL) failure, ISL trunking ⁴ and long haul link failure
→ Capacity measurement to link speeds of 8G, 10G, 16G, 32G with up to 100 km long haul connection ²	→ Stability of overnight soak testing ^{3,5}
	→ Capacity measurement to link speeds of 8G, 10G, 16G, 32G with up to 100 km long haul connections ⁴

¹ FC-PIs specifications are defined by the T11 committee of the International Committee on Information Technology Standards (INCITS). INCITS is accredited by and operates under rules approved by the American National Standards Institute (ANSI). FC-PI-6 (ANSI INCITS 512-2015) defines the standard to support the link speeds of 32G, 16G, and 8G; FC-PI-5 (ANSI INCITS 479-2011) defines the standard to support the link speeds of 16G, 8G, and 4G; FC-PI-3 (ANSI/INCITS 460-2011) defines the standard to support the link speeds of 10G, 4G, 2G, and 1G.

² Cisco MDS 9148S is only included in 8G capacity test with 20km long-haul

³ LDCA and LTX is not included

⁴ ELOM was not tested with ISL trunking

⁵ Cisco MDS 9148S was tested with 40km long-haul at FC8G and no 24-hour stability test

Device Under Test

Huawei explained that the OptiX OSN 1800 V Pro (Huawei OptiXtrans E6616), which brought another outcome of the test series, has not changed its concept.

It is designed for DCI, which can simplify deployment, ultra-wideband, and highly integrated data traffic. The mainboards of Huawei OptiX OSN 1800 below vary at different speeds.



Figure 1: Huawei OptiX OSN 1800 V Pro




Figure 2: Huawei LTX Board



Figure 3: Huawei LDCA Board



Figure 4: Huawei MDCA Board



Figure 5: Huawei ELOM Board

Testbed Description

The Huawei OptiX E6616 and OptiX OSN 1800 V Pro were the device under test (DUTs) respectively. They have the same hardware, and only the product name is different for a different market. Therefore, we designed the test bed as shown in Figure .

Huawei installed SUSE Enterprise Linux 12.4 in a bare metal mode on both hosts. We used vdbench already installed at Huawei labs by Huawei engineers to generate FC traffic. The host has 2x Intel® Xeon® E5-2658 v4 @2.30GHz CPUs, 8x 16G DDR4 memory, 1 x Huawei IN300 (2 x FC32G) port FC Host Adapter (HBA), 1 x Emulex LPe16002B-M6 PCIe 2-port 16Gb Fibre Channel Adapter and 1x SAS 800G SSD. We used all the CPU, FC adapter ports, and 16G memory for our test.

The storage hardware included two Huawei OceanStor 5000 V3 (referred to as Huawei OceanStor) devices, each equipped with 24 Non-Volatile Memory Express (NVMe) Solid-State Drive (SSD) disks that provided up to 5.8 GB/s (Gigabytes per second) Input/Output traffic. Using the open source test tool Vdbench released by Oracle, we generated bidirectional baseline traffic at the full configured speed.

The key point from the topology was the quorum server in place. It means that the storage system also had redundancy control. The quorum server recognized one of the Huawei OceanStor 5000 V3 as preferred storage and another one as non-preferred storage. The quorum server's strategy was to keep the preferred storage in a working state when a link or storage failure is detected.

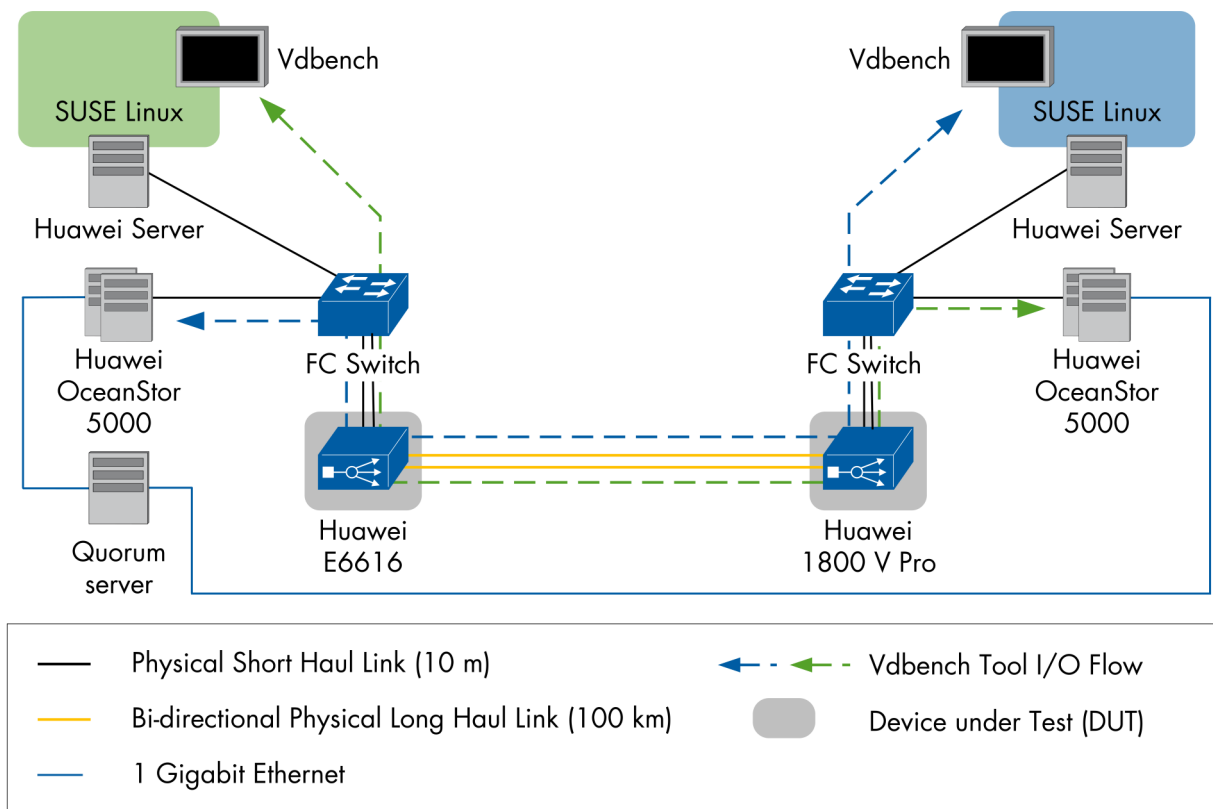


Figure 6: Logical Test Topology

Compatibility Test Combinations

Setup	FC Switch 1	FC Switch 2	Optical Transceiver (at E-port)
1	Brocade G620-1	Brocade G620-2	FC-PI-3
			FC-PI-6
2	Brocade G620-1	Brocade 6505 ⁶	FC-PI-5
3	Cisco MDS 9132T-1	Cisco MDS 9132T-2 ⁷	FC-PI-6
4	Cisco MDS 9132T-1	Cisco MDS 9148S	FC-PI-5
5	Cisco MDS 9148S-1 ⁶	Cisco MDS 9148S-2 ⁶	FC-PI-5

Table 1: FC Switch Combinations

FC Switch	Huawei OptiX OSN 1800 / Huawei OptiX E6616
FC-PI-6	FC-PI-6
FC-PI-5	FC-PI-5 and FC-PI-6
FC-PI-3	FC-PI-3

Table 2: Optical Transceivers between FC Switch and DUT

⁶ Brocade 6505 and Cisco MDS 9148S does not support FC-PI-6 and FC-PI-3

⁷ Cisco MDS 9132T does not support FC-PI-3

Hardware and Software

2020

Product Type	Product Name	Software Version
Devices Under Test WDM Equipment	Huawei OptiX OSN 1800 V Pro	V100R019C10
	Huawei OptiX Trans E6616	
	LDCA (facing E-ports towards FC-switch)	
	LTX board (facing E-ports towards FC-switch)	
	MDCA board (facing E-ports towards FC-switch)	
	EX40	
	OBU	
	OLP (facing long haul)	
	SXCH	
Device Management Tools	iMaster NCE	V100R019C00

Table 3: Devices under Test - Hardware and Software 2020

2021

Product Type	Product Name	Software Version
Devices Under Test WDM Equipment	Huawei OptiX OSN 1800 V Pro / Huawei OptiX Trans E6616	V100R021C00
	LDCA (facing E-ports towards FC-switch)	
	LTX board (facing E-ports towards FC-switch)	
	ELOM board (facing E-ports towards FC-switch)	
	OBU	
	OLP (facing long haul)	
	SCC SXCH STG	
Dispersion Compensation Module (for ELOM only)	Fujikura 100% C-band DCM for TW-RS / Sumitomo Z-DCFM-LFS/C-40 / Sumitomo N-DCFM-C-40-LC / Sumitomo Z-DCFM-LFS/C-40	N/A

Table 4: Devices under Test - Hardware and Software 2021

As we can see from the table, the EX40 a Fiber Optic Multiplexer is missing. Its principal function is to divide different light wavelengths, then deliver to respect local board, merge different light wavelengths, and return to long-haul fiber. As we only have one board per test, EX40 is not needed and therefore removed.

Long haul connection tests were successfully carried out across two distances:

- 100 km long haul with Brocade G620 and Cisco MDS 9132T pairs, respectively
- 20 km⁸ calculated medium-haul with Brocade 6505 and Cisco MDS 9148S pairs in 2020, respectively
- 40 km⁹ medium-haul with Brocade 6505 and Cisco MDS 9148S pairs in 2021, respectively

Product Type	Product Name	Software Version
Physical Server	Huawei RH2288H V3	iBMC: 2.94 BIOS: 3.87
Operating System	SUSE Linux Enterprise Server 12 SP4	Release: 12.4 Kernal: Linux 4.12.14-94.41-default
FC Traffic Simulation Software	Vdbench	v50406
FC Traffic Generator	Viavi MTS5800-100G	BERT 28.0.1
SAN Storage	Huawei OceanStor 5000 V3 ¹⁰	V300R002C10
Quorum Server	Quorum Server	V300R002C10

Table 5: Test Environment

2020		2021	
FC Switch	Software info	FC Switch	Software info
Brocade 6505	FOS v8.0.2c	Brocade 6505	FOS v8.2.1a
Brocade G620	FOS v8.0.2c	Brocade G620	FOS v9.0.1c
Cisco MDS 9132T	NX-OS version 8.4	Cisco MDS 9148S	NX-OS version 6.2
Cisco MDS 9148S	NX-OS version 6.2		

Table 6: FC Switch Software Version

⁸ Since none of the switches actually supported 100 km distances due to protocol settings, we calculated the theoretical kilometer distance* expected at the Brocade 6505 and the Cisco MDS 9148S based on the buffer supported on the switches. * Theoretical distance (kilometer) = (bytes of buffer size per port group - number of ports per port group * 1 reserved byte - 1 port) / recommended bytes of buffer size per kilometer. In this test, we used 8 bytes of buffer size per kilometer at 8G link speed, 12 bytes at 16G link speed, and 24 bytes at 32G link speed.

⁹ Due to Cisco MDS 9148S buffer limitation, we can only test up to 40km.

¹⁰ 24 Non-Volatile Memory Express (NVMe) solid-state disk (SSD) disks that provided up to 5.8 GB/s (Gigabytes per second) Input/Output

Capacity Testing

We measured the maximum FC-interface forwarding of the DUT using FC read block traffic (reading block size 32KiB) generated respectively between emulated data centers.

The following tables show the throughput. We performed a capacity test for each of the board under test, with all three setups, respectively. Each traffic stream carried bi-directional traffic.

LDCA

Setup	Speed Type	Expected Throughput (MB/s) ¹¹	Year	Measured Throughput, per direction (MB/s)	Verdict
1. Brocade G620 pair (2020 software info: 8.0.2c 2021 software info: 9.0.1c)	FC800	800	2020	786	Pass
			2021	782	
	FC1200	1200	2020	1171-1172	Pass
	FC1600	1600	2020	1564-1565	Pass
			2021	1559	
	FC3200	3200	2020	3092-3093	Pass
2021			3113-3115		
2. Brocade G620/6505 pair	FC1600	1600	2020	6505: 1534 G620: 1561	Pass
3. Cisco MDS 9132T pair	FC800	800	2020	778	Pass
	FC1600	1600	2020	1555-1557	Pass
	FC3200	3200	2020	3119-3120	Pass
4. Cisco MDS 9132T/9148S pair	FC800	800	2020	779	Pass

Table 7: LDCA Interface Throughput

LTX

Setup	Speed Type	Expected Throughput (MB/s) ¹¹	Year	Measured Throughput, per direction (MB/s)	Verdict
1. Brocade G620 pair (2020 software info: 8.0.2c 2021 software info: 9.0.1c)	FC800	800	2020	785-786	Pass
			2021	782	
	FC1200	1200	2020	1156-1158	Pass
			2021	1558-1560	
	FC1600	1600	2020	1541-1568	Pass
			2021	1558-1560	
FC3200	3200	2020	3088	Pass	
		2021	3113-3114		
2. Brocade G620/6505 pair	FC1600	1600	2020	6505: 1531 G620: 1563	Pass
3. Cisco MDS 9132T pair	FC800	800	2020	777-778	Pass
	FC1600	1600	2020	1558-1559	Pass
	FC3200	3200	2020	3118	Pass
4. Cisco MDS 9132T/9148S pair	FC800	800	2020	779	Pass

Table 8: LTX Interface Throughput

MDCA

Setup	Speed Type	Expected Throughput (MB/s) ¹¹	Year	Measured Throughput, per direction (MB/s)	Verdict
1. Brocade G620 pair	FC800	800	2020	786	Pass
	FC1200	1200		1160-1161	Pass
	FC1600	1600		1559-1562	Pass
	FC3200	3200		3091	Pass
2. Brocade G620/6505 pair	FC1600	1600		6505: 1535 G620: 1562	Pass
3. Cisco MDS 9132T pair	FC800	800		779	Pass
	FC1600	1600		1559	Pass
	FC3200	3200		3113-3114	Pass
4. Cisco MDS 9132T/9148S pair	FC800	800		779	Pass

Table 9: MDCA Interface Throughput

ELOM

Setup	Speed Type	Expected Throughput (MB/s) ¹¹	Year	Measured Throughput, per direction (MB/s)	Verdict
1. Brocade 300 pair	FC100	100	2021	97	Pass
	FC200	200		194	Pass
	FC400	400		390	Pass
2. Brocade G620 pair	FC800	800		783	Pass
	FC1200	1200		1174	
3. Brocade G620/6505 pair	FC800	800		6505: 772 G620: 778	Pass
4. Cisco MDS 9148S pair	FC800	800		777	Pass

Table 10: ELOM Interface Throughput

¹¹ The expected throughput is based on the layer 2 payload; we used a ratio of 97% of the link speed. The ratio is based on the formula: 2,048 bytes payload size / 2,112 bytes maximum frame size * 100% which excludes overhead from the throughput consisting of Start Of Frame (SOF), Cyclic Redundancy Check (CRC), and End of Frame (EOF). For example, 3.104 GB/s = 97% * 3.2 GB/s link speed.

Latency Test

We created a new test setup with Viavi MTS5800-100G tester to measure the latency value of the WDM system.

In this setup, we connected both DUTs back to back with the measurement tool (see Figure 7), which provides latency value in the precision of microseconds.

We removed all data center devices from the testbed and remained only both WDM devices running the FC services between the traffic generators. We generated FC traffic from all ports at different FC speeds consisting of 1G, 2G, 4G, 8G, 10G, 16G, and 32G respectively. The test tool supported the latency measurement on the same port. Therefore, we designed the Rx and Tx like below. Using an optical splitter we split the Rx and Tx at the traffic generator port into two separate fibers and connect them to the two WDM devices.

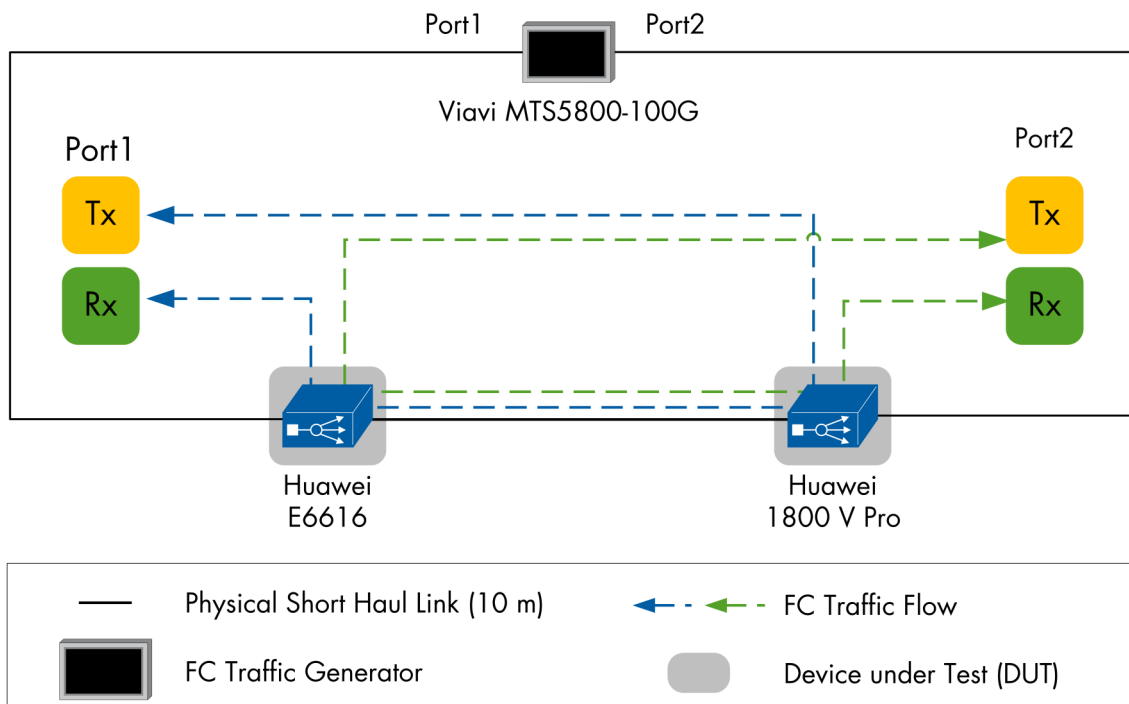


Figure 7: Latency Test Topology

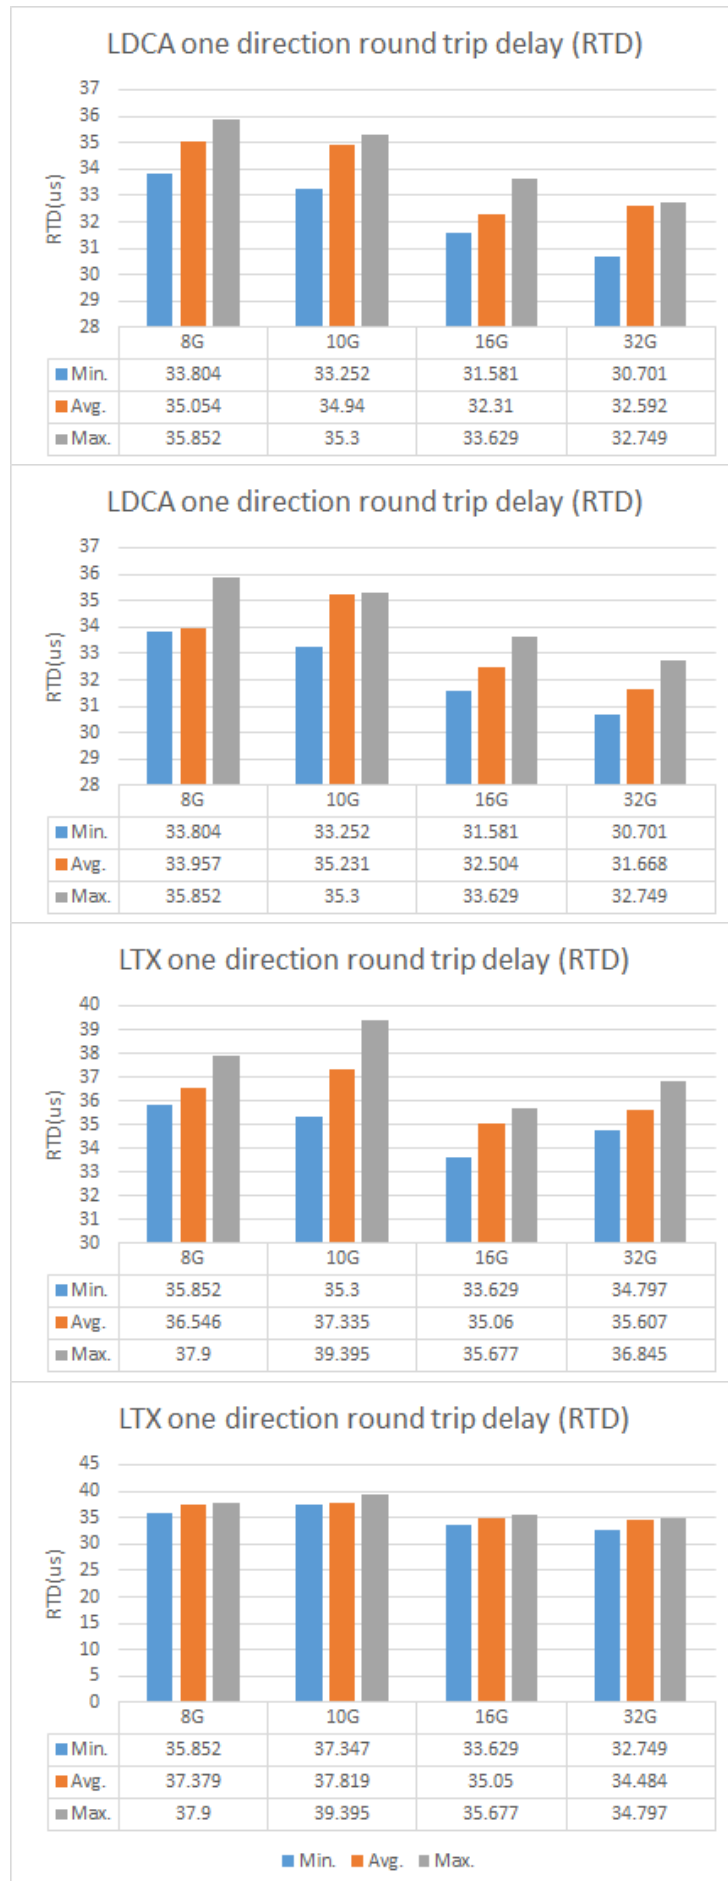


Figure 8: Latency Results of 2020 (DUT1 to DUT2, DUT2 to DUT1)

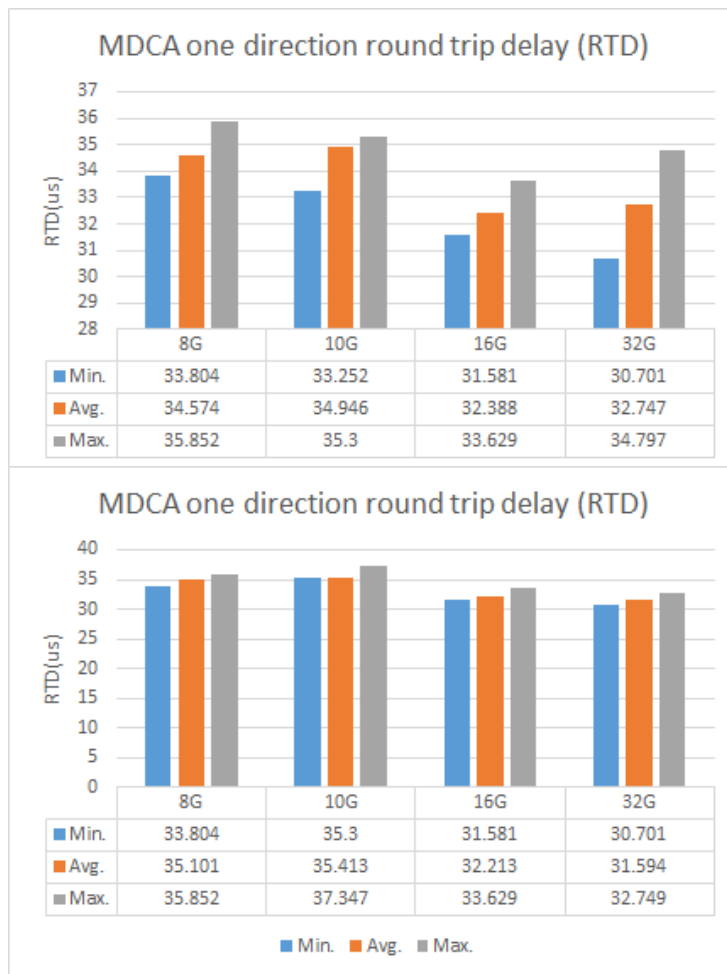


Figure 9: Latency Results of 2020 (DUT1 to DUT2, DUT2 to DUT1)

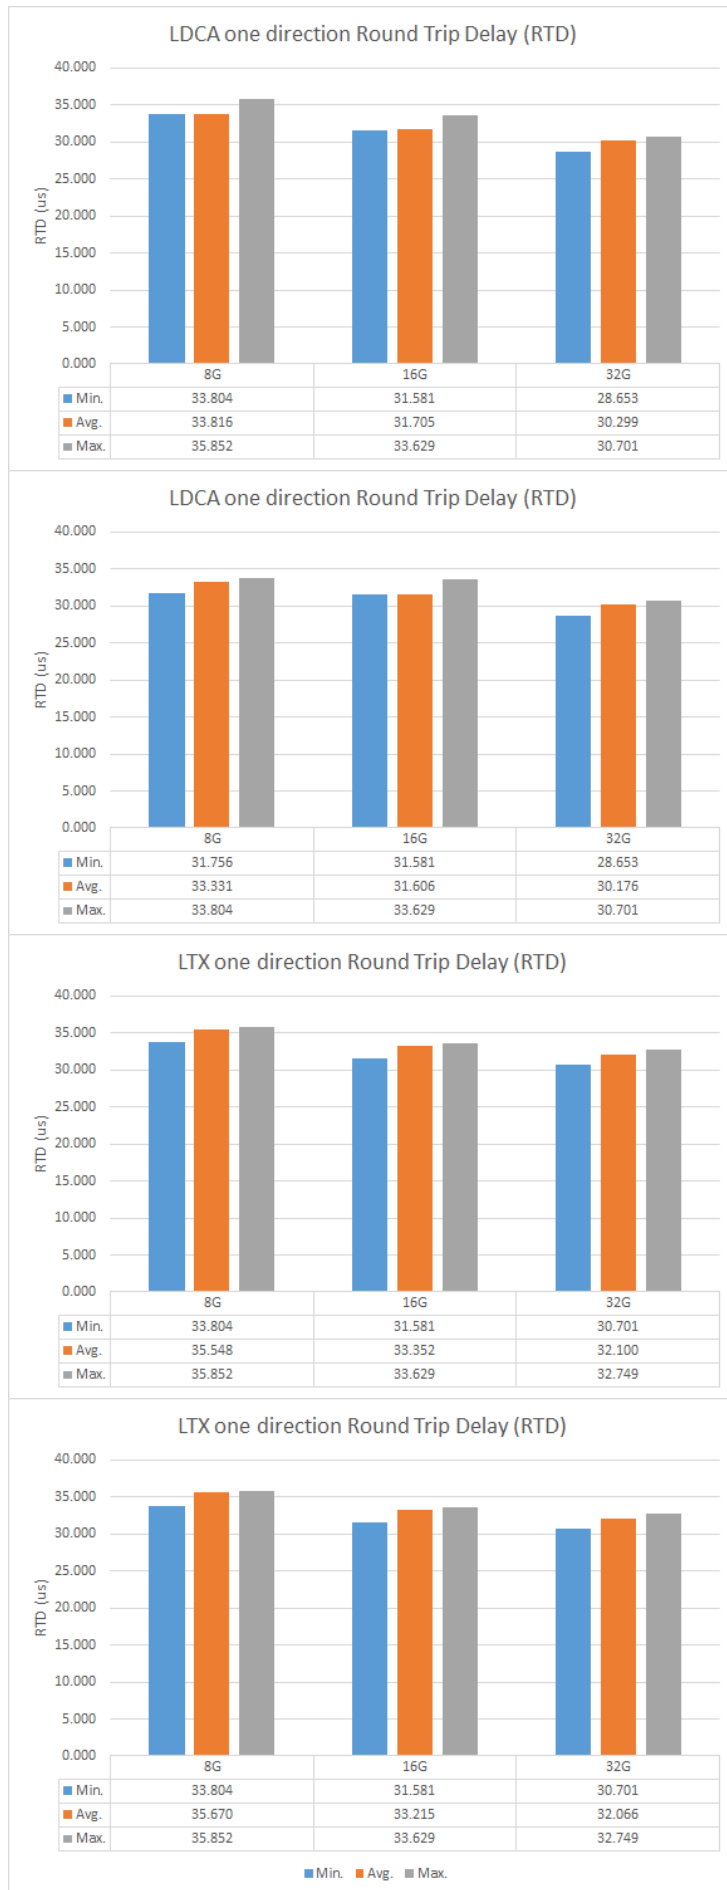


Figure 10: Latency Results of 2021 (DUT1 to DUT2, DUT2 to DUT1)

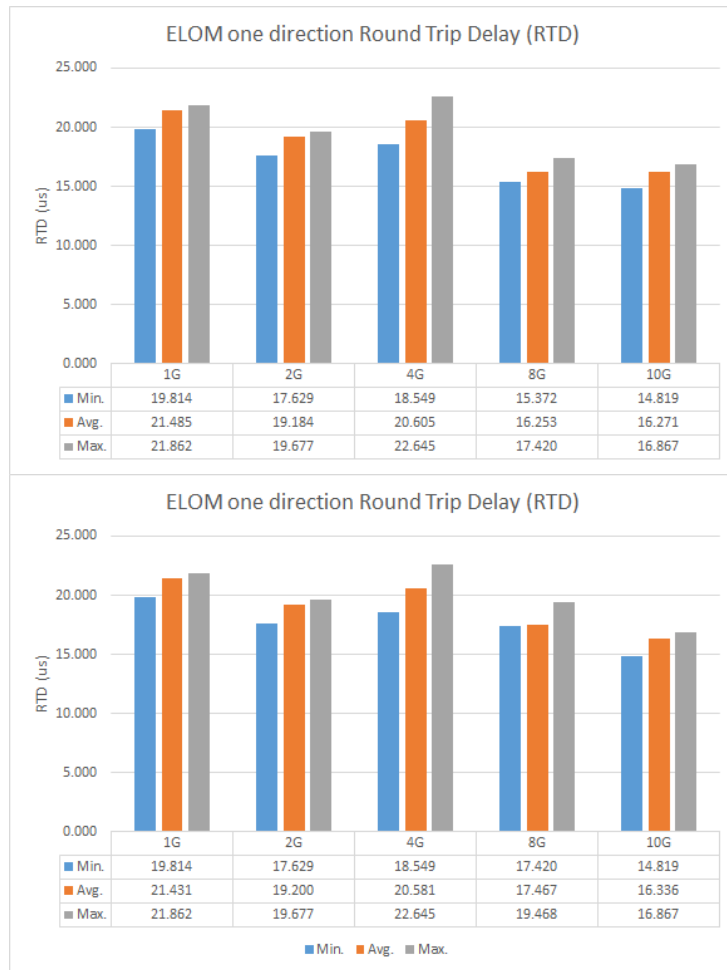


Figure 11: Latency Results of 2021 (DUT1 to DUT2, DUT2 to DUT1)

Availability Test

We verified that the DUT maintained stability when performing administrative activities on the DUT's hardware and the FC switch connected in the test environment.

We emulated a baseline scenario where the host's I/O traffic flowed between SANs from two different emulated data centers under normal conditions. The host on each side of the DCI initiated I/O operations, and the target was the storage on the other side of the DCI. The following figure depicts the main traffic streams (blue and green, both bidirectional) between the emulated data centers.

To verify the stability of the DUT, we emulated a common set of failures of the DCI segment while the traffic was running under normal conditions. These were a total of 7 types of emulated failures as shown in the figure. In ISL trunking and long-haul link protection, we expected the DUT to provide traffic switching between primary and redundant links. Especially with the participation of Quorum servers, we expected that the active-active storage cluster to protect the traffic switching between storage devices.

One note for the ELOM board, it needed a Dispersion Compensation Module (DCM) for the long haul protection to work properly when the long-haul distance is 100km. It is because of the TSFP+ we tested with. The optical module we test is "Optical module XFP|SFP+ Transceiver, TSFP+, Extended C Band, 9.95~11.3 Gbps with CDR,-1dBm, 3dBm, -16dBm, LC, 40km". DCM is needed when the transmission distance is greater than 40km of the optical module specification. The optical module's dispersion adjustment takes time in the service reconstruction when switching to the standby channel. And the switching performance may not be able to guarantee as less than 50ms when the transmission distance is more than 40km, and there is no DCM.

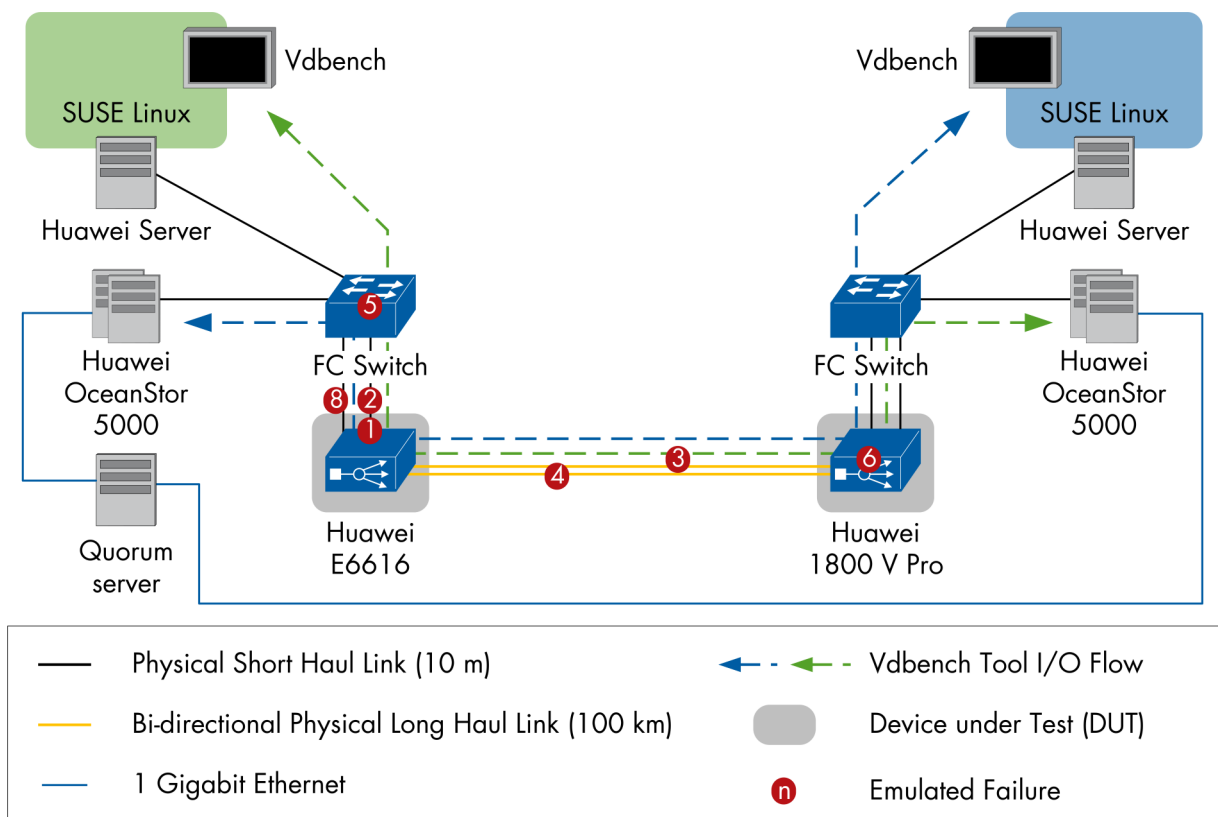


Figure 12¹²: Emulated Failure in the Test Bed

¹² The numbers in figure 12 represent the test cases in the test plan. Missing of number 7 is because 7 represents 24 hour long-haul test, and it doesn't include any failure scenario.

Active-Active Storage Switch Over

Under failure types 1, 2, 3, and 6, the test results showed the same status of the traffic that has been switched over. The primary storage obtained the I/O access of the host in the same data center and continued I/O operation with the host. The secondary storage stopped receiving any I/O access from the remote data center, neither from the local data center. The former case was because these four types of failure could interrupt the whole DCI connection from different hardware locations (see Figure 12), thereby causing it to stop transmitting data between both data centers. The latter phenomenon was due to the design mechanism.

Huawei OceanStor 5000 V3 implements a technology called "HyperMetro" for synchronizing the states between preferred storage and secondary storage. Once the link between two storage systems went down, the HyperMetro pair changed to the "To be synchronized" state. The Logic Unit Number (LUN) in the preferred storage continued providing service while the LUN in non-preferred storage stopped. During a link recovery, once the quiescing time (300 s) passed, we observed that the traffic switched back to baseline status between DC1 and DC2. We did not observe any impact as expected.

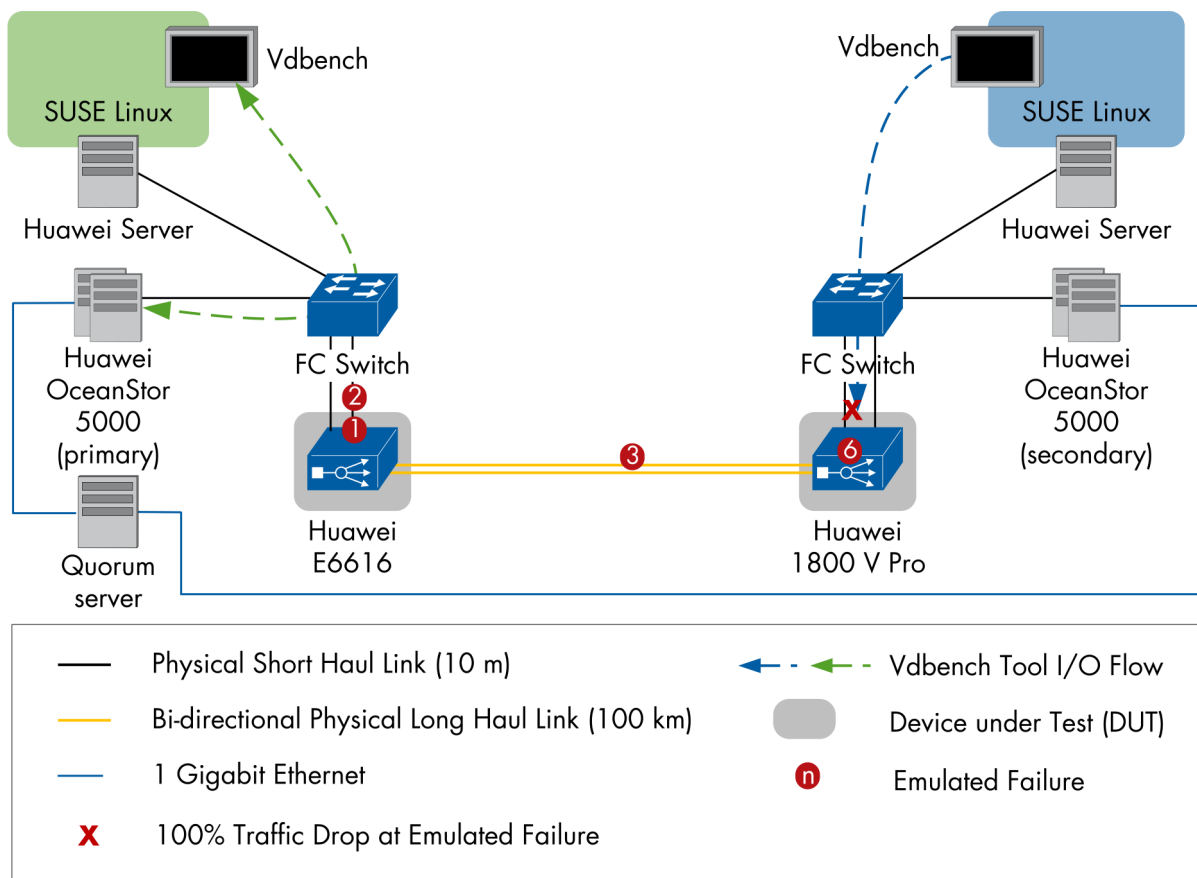


Figure 13: Traffic Status after the Switch Over for Failure Types 1, 2, 3 and 6

No. in Figure 13	Test Case	Action	Service Interruption		Verdict
			Expected ¹³	Observed (for primary storage) ¹⁴	
1	E-Port Disable/Enable Test	Disable	10s	2020: 10s including 6s drop to 0MB/s 2021: 7s including 5s drop to 0MB/s	Pass
		Enable	No impact	2020: No impact	Pass
2	E-Port Cable Disconnect/Reconnect Test	Disconnect	10s	2020: 10s including 6s drop to 0MB/s 2021: 7s including 5s drop to 0MB/s	Pass
		Reconnect	No impact	2020: No impact	Pass
3	Long Haul Network Failure without Redundancy	Disconnect	10s	2020: 10s including 6s drop to 0MB/s 2021: 7s including 5s drop to 0MB/s	Pass
		Reconnect	No impact	2020: No impact	Pass
6	DUT Reset Test	Shutdown	10s	2020: 10s including 6s drop to 0MB/s 2021: 7s including 5s drop to 0MB/s	Pass
		Turn on	No impact	2020: No impact	Pass

Table 11: Out of Service Overview (Active-Active Storage Cluster Protection)

¹³ The expected value included 5s fault detection by the quorum server and 15s from it to triggering an active-active switch over. Once the quorum server detected the fault, it triggers an active-active switchover, which lasts for 15 seconds (generally within 10 seconds in the vendor's lab tests). After the switchover was completed, the preferred storage took over services.

¹⁴ We observed 6s complete traffic drop to 0 MB/s during the switch over. The Huawei team configured the quorum server with 5s timer to detect the link heartbeat failure between the storage arrays. The link between storage arrays sent a heartbeat packet every second. After five consecutive heartbeat packets expired, the link was identified as disconnected. The 6s complete traffic drop to 0 MB/s included the time from discovery to switchover at the quorum server.

FC Switch Reboot

Once the FC switch from one data center rebooted, the storage on the remote site where the FC switch was not touched obtained the host's I/O access and continues to perform I/O operations.

In the data center where the FC switch rebooted, its connected storage stopped receiving any data from the FC switch as long as the reboot took place. This behavior is expected. After that the FC switch completed the reboot, once the quiescing time (300s) passed, we observed that the traffic switched back to baseline status between DC1 and DC2. We did not observe any impact as expected.

In 2021, as we have three pairs of FC switch, which is Brocade 300/Brocade 300, Brocade G620/Brocade G620, and Brocade G620/Brocade 6505. They have different software versions and hardware modules. The behavior after the reboot was different.

The service of rebooted Brocade 300 had totally stopped on the reboot side until the reboot procedure was done and the quiescing time had passed.

The service of rebooted Brocade G620 had not fully stopped on the reboot side. However, the FC traffic did stop for around 40-50 seconds until it recovered and read the non-rebooted side storage. Everything was recovered after the quiescing time had passed.

The interesting part is Brocade G620 and 6505 pair. The service interruption depended on which device we rebooted. If we reboot G620, the behavior is the same as G620's. If we reboot 6505, the behavior is the same as 300's.

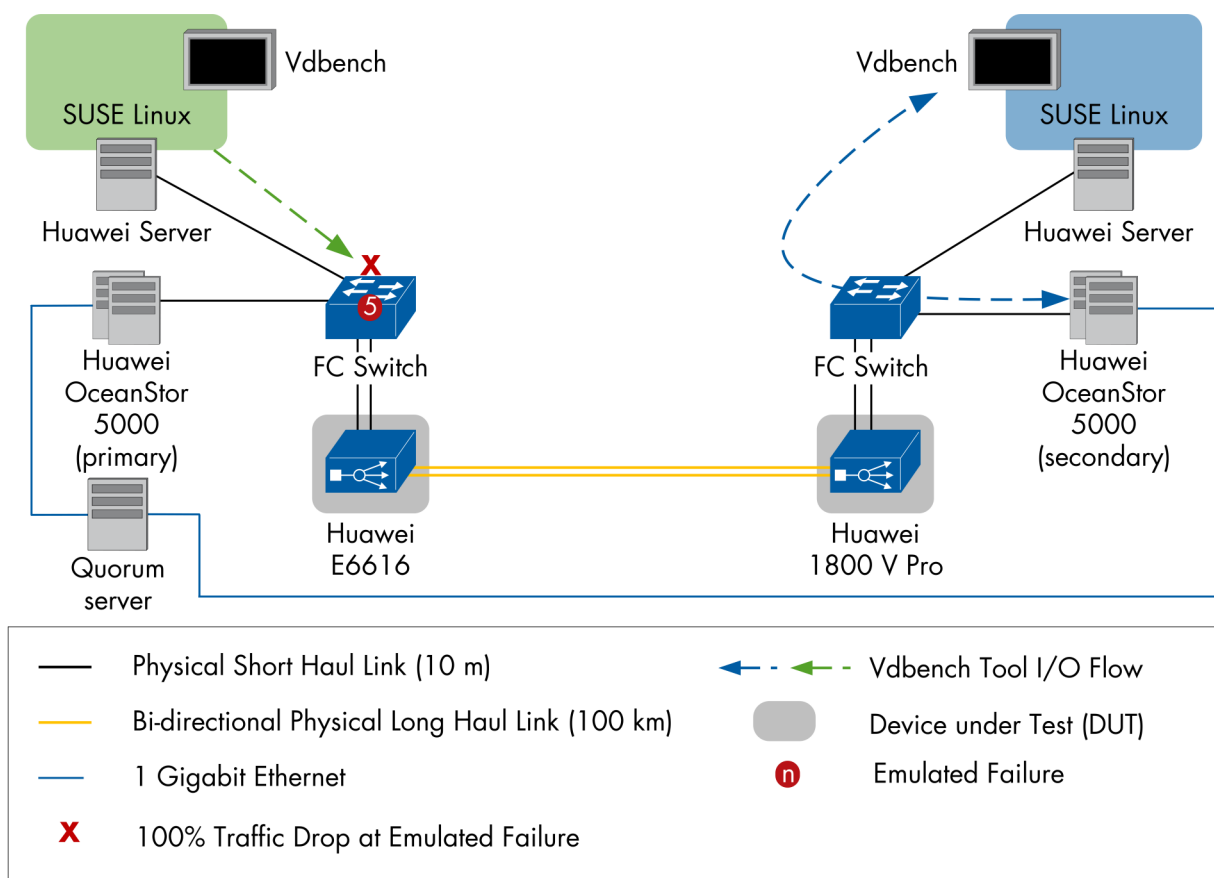


Figure 14: Traffic Status after the Switch Over for Failure Types 1, 2, 3 and 6

No. in Figure	Test Case	Action	Service Interruption		Verdict
			Expected ¹⁵	Observed ¹⁶	
5	Switch Reboot Test	Shut down	15s	2020: 13s on non-reboot side (including 11s drop to OMB/s) 2021: 7s on non-reboot side (including 5s drop to OMB/s)	Pass
		Turn on	No impact	2020: No impact 2021: No impact	Pass

Table 12: Out of Service Overview

Long Haul Link Protection

We measured the switch over time for the DUT to switch traffic to the backup link when the primary long-haul link fails. While traffic was running, we disconnected the primary link from the long-haul connection between both WDM devices.

We observed that the traffic was switched to the backup link. Table below shows the measured switch over time. When we reconnected the primary link previously disconnected, we did not observe any impact on the traffic. Then we disconnect the protect link again. The WDM switched back to the primary link again with the switch over time. The last step is to reconnect the protected link. We observed no impact on I/O flow traffic.

ISL trunking is a resiliency technology that Brocade uses to improve performance and redundancy. When two E-ports have the exact same configuration and are in the same zone, these two E-ports can be seen as a bundle. It has the load-share function that the two ports can achieve double throughput of the configured speed. They are also back-up for each other when one link has a problem or is down for unknown reasons. In our test, the ISL trunking was up and running at the beginning of the test. We unplugged fiber from one of the trunking ports, and we expected the traffic to reduce to 50% but no interruption. We plug the fiber back to its original port, and we expect the traffic to increase by 100%.

¹⁵ The expected value included 5s fault detection by the quorum server and 15s from it to triggering an active-active switch over.

¹⁶ The 11s drop to OMB/s included 6s of switch over by the quorum server and the reboot time of the switch (during the reboot process, the link at the FC switch was interrupted from time to time).

No. in Figure	Test Case	Setup	Action	Service Interruption		Verdict		
				Expected	Observed			
4	Long Haul Network Failure with Redundancy (2020 WDM software version: V100R019C10 2021 WDM software version: V100R021C00)	Brocade 300 pair (2021 software version: 7.4.1f ELOM only)	Disconnect	4s ¹⁷	2021: a maximum of 4s including 2s drop to OMB/s (for both primary and secondary storages)	Pass		
			Reconnect	No impact	2021: No impact	Pass		
		Brocade G620 pair (2020 software version: 8.0.1c 2021 software version: 9.0.1c)	Disconnect	4s ¹⁷	2020: a maximum of 4s including 2s drop to OMB/s (for both primary and secondary storages) 2021: a maximum of 3s including 1s drop to OMB/s (for both primary and secondary storages)	Pass		
			Reconnect	No impact	2020: No impact 2021: No impact	Pass		
		Brocade G620/6505 pair (G620 software version see upper 2020 6505 software version: 8.0.2c 2021 6505 software version:)	Disconnect	4s ¹⁷	2021: a maximum of 3s including 1s drop to OMB/s (for both primary and secondary storages)	Pass		
			Reconnect	No impact	2021: No impact	Pass		
		Cisco MDS 9132T pair (2020 software version: 8.4)	Disconnect	4s ¹⁷	2020: <= 1s ¹⁸	Pass		
			Reconnect	No impact	2020: No impact	Pass		
		8	ISL Trunking ¹⁹	Brocade G620 pair (2020 software version: 8.0.1c 2021 software version: 9.0.1c)	Disconnect	Drop 50%	2020: Drop 50% 2021: Drop 50%	Pass
					Reconnect	Increase 100%	2020: Increase 100% 2021: Increase 100%	Pass

Table 13: Out of Service Overview

¹⁷ Includes the impact of end-to-end flow control (a total of four hops from the host to the storage through two FC switches). With credit recovery enabled on the FC switch, we calculated each hop for 1 second interruption, based on the hold off time of 500 ms (milliseconds that a frame could be buffered on a port without being overwritten) configured on the Brocade G620 (same as Cisco MDS 9132T) switch; added to that the impact of retransmission caused by the frame loss during the link failure, and vdbench accuracy of 1 sample per second.

¹⁸ The smallest number of samples is one per second in vdbench.

¹⁹ ELOM's ISL trunking was not tested due to it is a low-speed board

Board Under Test	Switch Over Time (s)
ELOM	A maximum of 4s service interruption (including a maximum of 2s drop to OMB/s)

Table 14: Switch Over Time Brocade 300-300 2021

Board Under Test	Switch Over Time (s)
LDCA	2020: A maximum of 4s service interruption (including a maximum of 2s drop to OMB/s)
LTX	A maximum of 4s service interruption (including a maximum of 2s drop to OMB/s)
MDCA	A maximum of 4s service interruption (including a maximum of 2s drop to OMB/s)

Table 15: Switch Over Time Brocade G620 - G620 2020

Board Under Test	Switch Over Time (s)
LDCA	A maximum of 3s service interruption (including a maximum of 1s drop to OMB/s)
LTX	A maximum of 3s service interruption (including a maximum of 1s drop to OMB/s)
ELOM	A maximum of 3s service interruption (including a maximum of 1s drop to OMB/s)

Table 16: Switch Over Time Brocade G620 - G620 2021

Board Under Test	Switch Over Time (s)
ELOM	A maximum of 3s service interruption (including a maximum of 1s drop to OMB/s)

Table 17: Switch Over Time Brocade G620 - 6505 2021

Board Under Test	Switch Over Time (s)
LDCA	Less than 1-second service interruption
LTX	Less than 1-second service interruption
MDCA	Less than 1-second service interruption

Table 18: Switch Over Time Cisco MDS 9132T - 9132T 2020

Summary of Test Runs for all Failure Scenarios

With all boards under test, the DUT demonstrated its ability to maintain stability while performing the above failure scenarios in each of the test setups.

We used the maximum load of the baseline traffic for each of the selected interfaces as measured in capacity tests and formed 87 combinations as listed in the tables below.

Board Under Test	E-port Interface used in Test Case					
	1	2	3	4	5	6
ELOM	1G	2G	1G	4G	2G	4G

Table 19: Brocade 300 Pair 2021

Board Under	E-port Interface used in Test Case						
	1	2	3	4	5	6	8
LDCA	8G, 16G	10G, 32G	8G, 32G	10G, 16G	16G	8G, 10G, 32G	8G, 16G, 32G
LTX	8G, 16G	10G, 32G	8G, 32G	10G, 16G	16G	8G, 10G, 32G	8G, 16G, 32G
MDCA	8G, 16G	10G, 32G	8G, 32G	10G, 16G	16G	8G, 10G, 32G	8G, 16G, 32G

Table 20: Brocade G620 Pair 2020

Board Under	E-port Interface used in Test Case							
	1	2	3	4	5	6	7	8
LDCA	8G	16G	32G	16G	16G	32G	N/A	16G
LTX	8G	16G	32G	16G	16G	32G	N/A	16G
ELOM	8G	16G	8G	10G	8G	10G	8G, 16G	N/A

Table 21: Brocade G620 Pair 2021

Board Under Test	E-port Interface used in Test Case		
	1	4	6
ELOM	8G	8G	8G

Table 22: Brocade G620 - 6505 Pair 2021

Board Under Test	E-port Interface used in Test Case						
	1	2	3	4	5	6	7
LDCA	16G	8G, 32G	8G, 32G	16G	16G	8G, 32G	32G
LTX	16G	8G, 32G	8G, 32G	16G	16G	8G, 32G	32G
MDCA	16G	8G, 32G	8G, 32G	16G	16G	8G, 32G	32G

Table 23: Cisco MDS 9132T Pair 2020

Board Under Test	E-port Interface used in Test Case		
	2	3	6
LDCA	8G	8G	8G

Table 24: Cisco MDS 9148S Pair 2021

Soak 24 Hours Test

We verified the WDM system's reliability in terms of performance consistency under long-period stress load conditions. The Huawei team configured all three boards in a snake configuration, keeping traffic flowing between the two data centers.

The test tool Vdbench triggered the baseline traffic bidirectional for 24 hours. During that time, we monitored the system log of both hardware and software. We confirm that the system could transfer the data at a consistent rate and constant latency of 24 hours. As expected, the system under test remained stable; we did not observe any software crashes or hardware failures during the test duration.

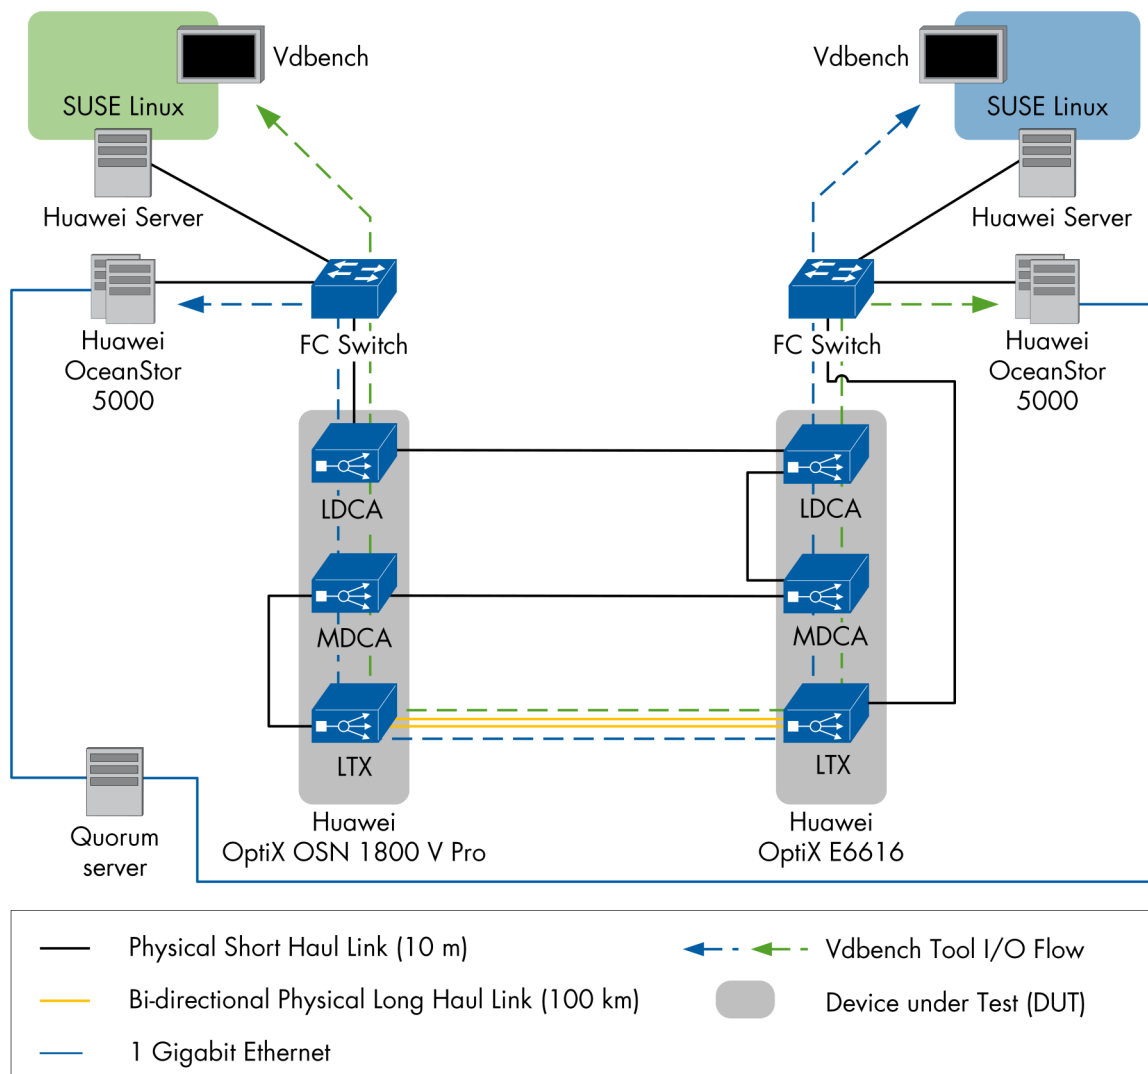


Figure 15: Topology of 24-hour Throughput 2020

In 2021, LDCA and LTX focused mainly on compatibility, and 24-hour long-term stability was already tested last year. Therefore, ELOM was the only DUT tested for the 24-hour long-term stability test. We switched back to a standard test bed, which means no snake configuration.

The test tool Vdbench triggered the baseline traffic bidirectional for 24 hours. During that time, we monitored the system log of both hardware and software. We confirm that the system could transfer the data at a consistent rate and constant latency of 24 hours. As expected, the system under test remained stable; we did not observe any software crashes or hardware failures during the test duration.

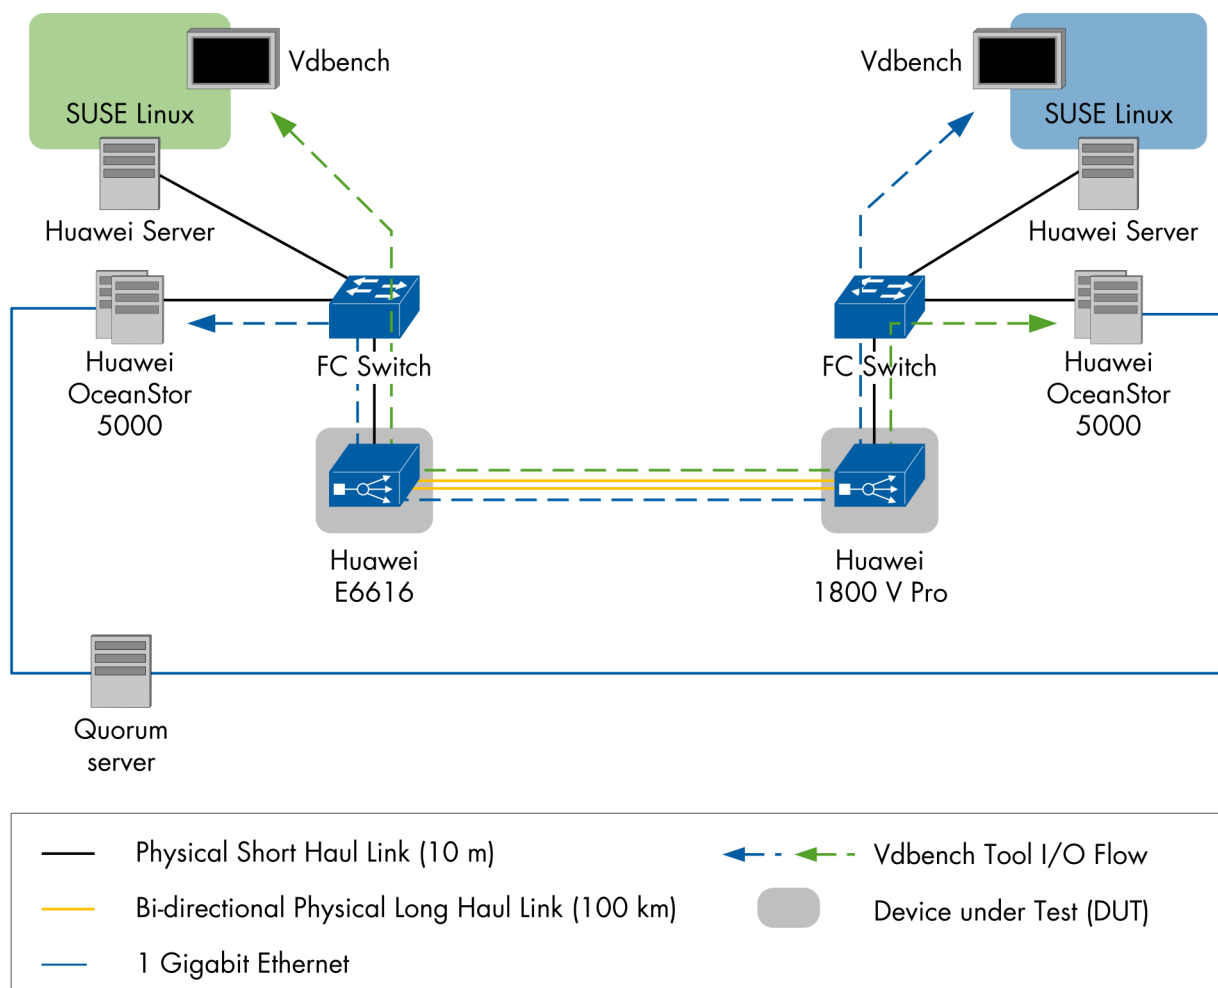


Figure 16: Topology of 24-hour Throughput 2021




Figure 17: 24-hour Throughput at 32G, LDCA & LTX & MDCA, Cisco MDS 9132T - 9132T, 2020





Figure 18: 24-hour Throughput at 8G and 10G, ELOM, Brocade G620-G620, 2021

Conclusion


In 2020, EANTC verified the interoperability of Huawei OptiX OSN 1800 V Pro (OptiXtrans E6616) and third-party Fibre Channel switches from Brocade (6505 and G620) and Cisco (MDS 9132T).

Multiple optical transceiver functions were certified in the E_port between these FC switches and the Huawei OptiX OSN 1800 V Pro (OptiXtrans E6616), including FC-PI-6, FC-PI-5, and FC-PI-3. We validated forwarding speeds in FC 8G, 10G, 16G, and 32G scenarios. When forwarding traffic at any of these standardized speeds, the Huawei OptiX OSN 1800 V Pro (OptiXtrans E6616) did not exhibit any speed impact. The extended 24 hours soak testing confirmed stable operations of Huawei OptiX OSN 1800 (OptiXtrans E6616) without any traffic impact. We conducted a range of service availability tests to disable/enable the port, disconnect/reconnect E_port fiber, disconnect/reconnect long haul, reboot FC switch, remove/re-install line card with both FC switch pairs, and remove/re-install ISL trunking fiber with Brocade G620.

All tests documented in this report passed our verification. Based on our test results, EANTC confirms that the Huawei OptiX OSN 1800 V Pro (OptiXtrans E6616) fulfills Huawei's claims to work in enterprise data center interconnection scenarios as an integrated, highspeed, and resiliency solution.

In 2021, EANTC verified three boards in 2021, ELOM, LDCA, LTX, respectively. The main differences between 2020 and 2021 are software differences both on the WDM system and Brocade G620 FC switch. ELOM, LDCA, and LTX showed good compatibility and protection with the latest Brocade G620 firmware (9.0.1c). ELOM also demonstrated compatibility and protection and multi-switch (Brocade G620 - 6505, Cisco MDS 9148S - 9148S) compatibility and long-term stability (24-hour stability).

All tests documented in this report passed our verification. Based on our test results, EANTC confirms that the Huawei OptiX OSN 1800 V Pro (OptiXtrans E6616) fulfills Huawei's claims to work in enterprise data center interconnection scenarios as an integrated, highspeed, and resiliency solution.



This report is copyright © 2022 EANTC AG.

While every reasonable effort has been made to ensure accuracy and completeness of this publication, the authors assume no responsibility for the use of any information contained herein. All brand names and logos mentioned here are registered trademarks of their respective companies.

EANTC AG
Salzufer 14, 10587 Berlin, Germany
info@eantc.de, <https://www.eantc.de/>
[v1.0 20220607]