

EANTC Independent Test Report

Huawei WDM Compatibility Test

OptiXtrans E6608T

April 2022



Introduction

Huawei commissioned EANTC to validate functional, interoperability, and performance aspects of the Huawei OptiXtrans E6608T wave division multiplex (WDM) solution with a specific focus on the Storage Area Network (SAN) use case scenarios.

We conducted all tests in our lab in Berlin, Germany, in October and November 2020.

The importance of public and private cloud services obtained through data center scalability and availability is described in the first edition of the EANTC WDM test report series (July 2020). One of the key ideas for flexible long-distance transmission in DCI (Data Center Interconnect) is to optimize the use of fiber capacity by flexibly allocating data rates to meet the needs of end-to-end connection requirements.

As proved in the test, that the Huawei WDM implementation of fine-grained bandwidth varies with the modular hardware setup. We also mention and introduce the opportunity of full participation for Huawei OptiXtrans E6608T into the WDM test series.

The current report covers all the goals of the test series related to performance and interoperability. We focus on long-haul link protection in the area of resiliency.

Huawei OptiXtrans E6608T is an optical-electrical WDM transmission device. Huawei OptiX OSN 1800 II TP 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).

We conducted the same test with a newer software version in Huawei's lab in Chengdu, China, in November and December 2021. The following boards were tested: LDCA, and LTX. As we have performed the full test suite on LDCA last year and the only difference is the software version, Huawei suggests only testing the compatibility of LDCA with the latest Brocade G620 software version. Therefore, LDCA was not tested for the multi-switch transparent test and 24-hour stability test. LTX had been tested with the full test suite.

Executive Summary

We verified the compatibility between different speeds of Huawei OptiXtrans E6608T according to the FC-PI-6, FC-PI-3, and FC-PI-5 standards. The speeds in the test include 16G and 32G, which symbolizes the interfaces for next-generation data center storage networks, and 10G, as well as 4G, 8G, and 16G, which can optimize the use of optical fiber capacity in a fine-grained manner.

The testbed 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 6505 and Cisco MDS 9132T / Cisco MDS 9148S in 2020, respectively.

We tested the LTX board with Brocade G620 / G620, G620 / 6505, and Cisco MDS 9184S / 9148S in 2021. And the LDCA board with Brocade G620 / G620.

We verified the robustness of the DUT by performing administrative activities on the DUT and connected equipment.

In 2020, we put DUT under continuous load for 24 hours in a soak test environment to ensure it would support uninterrupted service. As expected, the system remained stable without any restart or service interruption, zero packet loss, and low latency.

In 2021, we performed the 24-hour stability test with the LTX board, and we confirmed that LTX remained stable without any service interruption or hardware/software failure. As we tested LDCA last year, Huawei focused mainly on the compatibility of LDCA with the latest Brocade software (v9.0.1c), which means we did not perform the 24-hour stability test with LDCA board this year.

Highlights 2020	Highlights 2021
→ DCI interoperability with four Fibre Channel switches, including Brocade 6505, Brocade G620, Cisco MDS 9148S, and Cisco MDS 9132T	→ DCI interoperability with multiple 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 6505 and Brocade G620, as well as between Cisco MDS 9148S and Cisco MDS 9132T, respectively	→ Transparent multi-switch type forwarding between Brocade 6505 and Brocade G6204
→ Stability of overnight soak testing	→ Protection against Inter-Switch Link (ISL) failure, ISL trunking and long haul link failure
→ Forwarding performance of 4G, 8G and 10G Fibre Channel interfaces at ELOM board, as well as 8G and 10G Fibre Channel interfaces at LDX board; long haul transmission with 100 km long links using optical amplifiers	→ Stability of overnight soak testing ³
→ Forwarding performance of 8G, 16G and 32G Fibre Channel interfaces at LDCA board; long haul transmission with 100 km long links	→ Capacity measurement to link speeds of 8G, 10G, 16G, 32G with up to 100 km long haul connections ⁴
→ Protection of long haul link ²	

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 that are 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 Protection against link failure on local transmission port

³ LDCA is not included

⁴ Cisco MDS 9148S was tested with 40 km long-haul at FC8G and no 24-hour stability test is performed

Device Under Test

Huawei explained that the OptiXtrans E6608T 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 four main boards of Huawei OptiXtrans E6608T vary at different speeds.



Figure 1: Huawei OptiXtrans E6608T



Figure 2: Huawei ELOM Board



Figure 3: Huawei LDCA Board



Figure 4: Huawei LDX Board



Figure 5: Huawei LTX Board

2020

Testbed Description

Figure 6 describes the logical topology of the test setup. Generally, for this kind of test, we expected a quorum server that would synchronize redundant storage servers across the two data centers provided in the test bed. Huawei did not provide a quorum server, though; as a consequence, we measured two distinct storage systems (unsynchronized) by accessing each of them from the respective remote location. This way, we forced each test to transit across the WDM system.

Long Haul and Bandwidth

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 km⁵ calculated medium-haul with Brocade 6505 and Cisco MDS 9148S pairs, respectively

The storage hardware included two Huawei OceanStor 5500 V5 (referred to as Huawei OceanStor) devices, each equipped with 20 serial attached SCSI (SAS) hard disk drive (HDD) disks that provided up to 1.6 GB/s (Gigabytes per second) Input/Output traffic. Using the open-source test tool Vdbench released by Oracle, we generated bidirectional baseline traffic on the full traffic throughput for 8G, 10G, and 16G FC-link. We chose the same baseline traffic at 1.6 GB/s for 32G FC-link to verify the functionality of the WDM system integrated with the data center.

Amplifier

To connect the board under test with a 100 km long fiber in the long-haul transmission, Huawei brought two options. The LDX board type used a 200G optical transceiver on the WDM side, which was directly connected to 100 km optical fiber. Both ELOM and LDX board types had 10G optical transceivers on board, whose maximum supported distance is 40 km. Therefore, Huawei brought external optical amplifiers. The remaining distance in kilometers was extended by the amplifier.

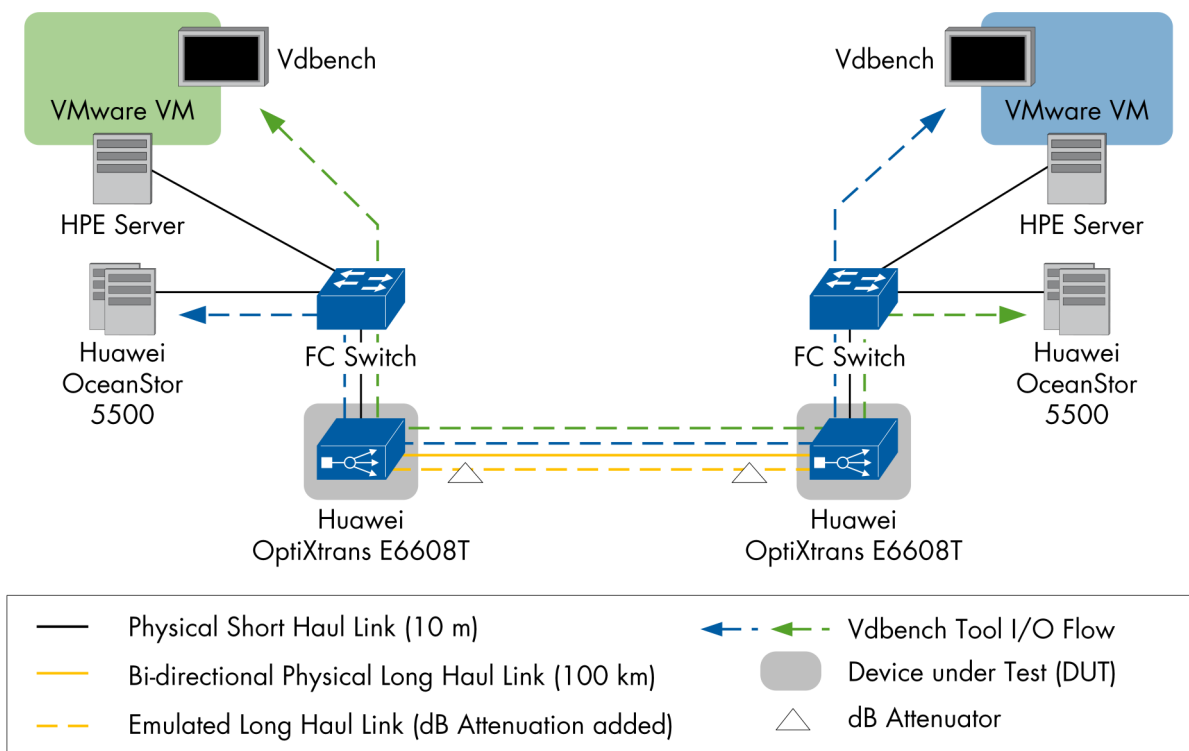


Figure 6: Logical Test Topology

⁵ 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.

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
2			FC-PI-6
3	Brocade G620-1	Brocade 6505 ⁶	FC-PI-5
4	Cisco MDS 9132T-1	Cisco MDS 9132T-2 ⁷	FC-PI-6
5	Cisco MDS 9132T-1	Cisco MDS 9148S ⁸	FC-PI-5

Table 1: FC Switch Combinations

FC Switch	Huawei OptiXtrans E6608T
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

Hardware and Software

Product Type	Product Name	Software Version
Devices Under Test WDM Equipment	Huawei OptiXtrans E6608T	V100R019C10
	ELOM board (facing E-ports towards FC-switch)	
	LDCA board (facing E-ports towards FC-switch)	
	LDX board (facing E-ports towards FC-switch)	
	OBU (sub-module, between EX40 and OLP)	
	EX40 (sub-module, facing board under test)	
	OLP (sub-module facing long haul link)	

Table 3: Device under Test - Hardware and Software

⁶ Brocade 6506 does not support FC-PI-6 and FC-PI-3

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

⁸ Cisco MDS 9148S does not support FC-PI-6 and FC-PI-3

Product Type	Product Name	Software Version
Physical Server	HPE DL380 Gen9	VMware 6.0
Virtualization Platform	VMware vSphere	6.7
SAN Storage	Huawei OceanStor 5500 V5	V500R007C30
FC Switch	Brocade 6505	8.0.1
	Brocade G620	8.0.1
	Cisco MDS 9148S	6.2
	Cisco MDS 9132T	8.2

Table 4: Test Environment

Capacity Testing

We migrated interoperability testing to capacity testing and used Vdbench to generate a full speed load of the interface, to ensure maximum interface forwarding from DUT. The following tables show the throughput. We performed a capacity test for each of the board under test, with all four setups respectively.

Note: We also verified interoperability of Cisco MDS 9132T/9148S pair (setup 5) with all three boards ELOM, LDCA, and LDX using baseline traffic of 786 MB/s. Due to the limitation of the Cisco MDS 9148S buffer size, the maximum throughput reached to 1,494 MB/s without any long-distance cables (10 m); we observed that all traffic passed through the LDCA board.

The performance test results met our expectations for Input/Output operations per second (maximum 3200 IOPs), Input/Output response time (less than 10 milliseconds in most cases), and block bandwidth (maximum 1.539 GB/s, or Gigabytes per second).

ELOM

Setup	Optical Transceiver	Expected Throughput (MB/s) ⁹	Measured Throughput, per direction (MB/s)	Verdict
1. Brocade G620 pair	FC-4G	400	394	Pass
	FC-8G	800	789	Pass
2. Brocade G620 pair	FC-10G	1,200	1,156	Pass
3. Brocade G620/6505 pair	FC-4G	400	395	Pass
	FC-8G	800	790	Pass
4. Cisco MDS 9132T pair	FC-4G	400	399	Pass
	FC-8G	800	786	Pass

Table 5: ELOM Interface Throughput

LDCA

Setup	Optical Transceiver	Expected Throughput (MB/s) ⁹	Measured Throughput, per direction (MB/s)	Verdict
1. Brocade G620 pair	FC-8G	800	790	Pass
	FC-16G	1,600	1,540	Pass
	FC-32G	3,200	1,539 ¹⁰	Pass
2. Brocade G620 pair	FC-10G	1,200	1,156	Pass
3. Brocade G620/6505 pair	FC-8G	800	790	Pass
	FC-16G	1,600	1,535	Pass
4. Cisco MDS 9132T pair	FC-8G	800	786	Pass
	FC-16G	1,600	1,535	Pass
	FC-32G	3,200	1,539 ¹⁰	Pass

Table 6: LDCA Interface Throughput

LDX

Setup	Optical Transceiver	Expected Throughput (MB/s) ⁹	Measured Throughput, per direction (MB/s)	Verdict
1. Brocade G620 pair	FC-8G	800	790	Pass
2. Brocade G620 pair	FC-10G	1,200	1,154	Pass
3. Brocade G620/6505 pair	FC-8G	800	785	Pass
4. Cisco MDS 9132T pair	FC-8G	800	787	Pass

Table 7: LDX Interface Throughput

⁹ 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.

¹⁰ The transmitted traffic was equal to 1.6 GB/s as this is the maximum provided by the Huawei OceanStor 5500 V5 equipped with 20 serial attached SCSI (SAS) hard disk drive (HDD) disks.

Latency Comparison Test

We knew that the DUT latency was actually below 1ms. This precision could not be measured by vdbench, as it had a resolution of 1ms. However, we performed the following tests to make sure that there was no distinguishing difference between the test results bypassing the DUT and including the DUT, indicating that the DUT operated normally and its latency would be below 1ms. The expected end-to-end storage latency values in the table below are based on EANTC experience.

- Reference latency (end-to-end latency without WDM system): Two-way latency value measured by Vdbench towards Huawei OceanStor without WDM system bidirectionally
- Total latency (end-to-end including WDM): Two-way latency value measured by Vdbench towards Huawei OceanStor passing the WDM system bidirectionally

Board Under Test	Setup	Optical Transceiver	Expected Latency (ms)	Reference Latency (ms)	Total Latency (ms)
ELOM	1	FC-4G	150	40.5	40.5
		FC-8G	150	20.2	20.2
	2	FC-10G	150	13.4	13.4
	3	FC-4G	150	40.5	40.5
		FC-8G	150	20.2	20.2
	LDCA	1	FC-8G	150	20.2
FC-16G			150	10.1	10.1
FC-32G			150	10.1	10.1
2		FC-10G	150	13.4	13.4
3		FC-8G	150	20.2	20.2
		FC-16G	150	10.1	10.1
LDX	1	FC-8G	150	20.2	20.2
	2	FC-10G	150	13.4	13.4
	3	FC-8G	150	20.3	20.3

Table 8: Latency Overview

Hardware Robustness Testing

The test goal was to verify that the DUT maintains stability while performing administrative activities on the DUT's hardware and the connected equipment.

We emulated six types of failure as shown in Figure 7. Each scenario included two manual actions as described below:

- While baseline traffic was running, we emulated the failure. We expected that all traffic shall be dropped, indicating that the failure took place. We checked the management connectivity to the DUT via CLI and expected that the session stayed stable in all except 6th scenario.
- We removed the failure, then started the baseline traffic. We expected that the baseline traffic went through without any loss.

The DUT, with all three boards under test respectively, demonstrated its ability to maintain stability while we performed 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 73 combinations as listed in the tables below.

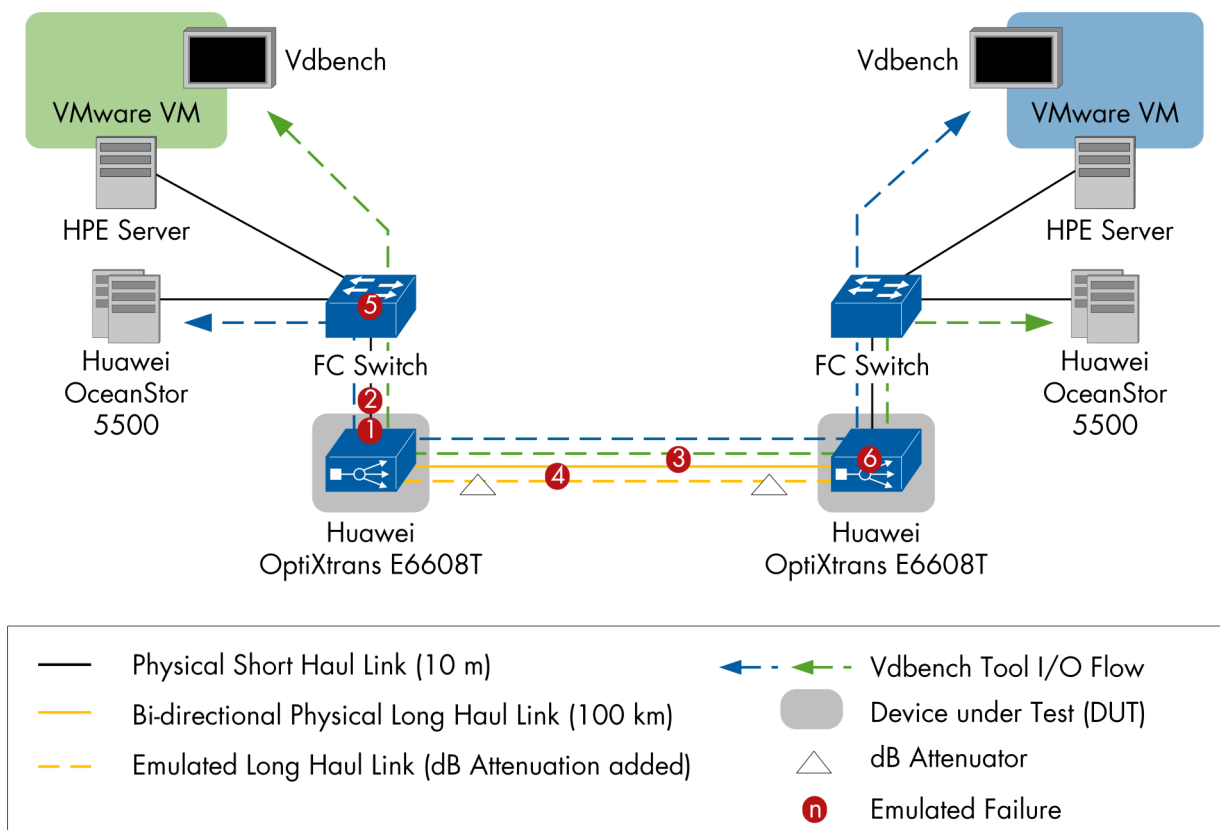


Figure 7: Emulated Failure

No. in Figure 7	Test Case	Action	Baseline Traffic		Verdict
			Expected	Observed	
1	E-Port Disable/Enable Test	Disable	100% Drop	100% Drop	Pass
		Enable	No impact	No impact	Pass
2	E-Port Cable Disconnect/Reconnect Test	Disconnect	100% Drop	100% Drop	Pass
		Reconnect	No impact	No impact	Pass
3	Long Haul Network Failure (1 link)	Disconnect	No session drop	No session drop ¹¹	Pass
		Reconnect	No impact	No impact	Pass
4	Long Haul Network Failure (all links)	Disconnect	100% Drop	100% Drop	Pass
		Reconnect	No impact	No impact	Pass
5	Switch Reboot Test	Shutdown	100% Drop	100% Drop	Pass
		Turn on	No impact	No impact	Pass
6	DUT Reset Test	Shutdown	100% Drop	100% Drop	Pass
		Turn on	No impact	No impact	Pass

Table 9: Hardware Robustness Tests Overview

Board Under Test	E-port Interface used in Test Case				
	1	2	4	5	6
ELOM	4G, 8G	8G	4G, 8G	8G	4G, 8G
LDCA	8G, 16G	16G, 32G	8G, 16G, 32G	16G	8G, 16G, 32G
LDX	8G	8G	8G	8G	8G

Table 10: Setup 1 - Brocade G620 Pair

Board Under Test	E-port Interface used in Test Case		
	2	4	6
ELOM	10G	10G	10G
LDCA	10G	10G	10G
LDX	10G	10G	10G

Table 11: Setup 2 - Brocade G620 Pair, 10G

¹¹ All test results are based on pulling the fiber(s) from the Tx side.

Board Under Test	E-port Interface used in Test Case	
	4	6
ELOM	4G, 8G	4G, 8G
LDCA	8G, 16G	8G, 16G
LDX	8G	8G

Table 12: Setup 3 - Brocade G620/6505 Pair

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

Table 13: Setup 4 - Cisco MDS 9132T Pair

Long Haul Link Protection

We measured the switch over time for the system under test (SUT) 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 system successfully switched the traffic to the back up link as expected.

When we reconnected the link previously disconnected, we did not observe any impact on the traffic. This is due to the fact that the Huawei WDM system does not switch back to the reconnected link, but declares the former backup link as the new primary.

Protection — Soak 24 Hours Test

We verified the reliability of the WDM system in terms of performance consistency under the conditions of long period stress load.

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 was able to transfer the data at a consistent rate and constant latency for a duration 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.

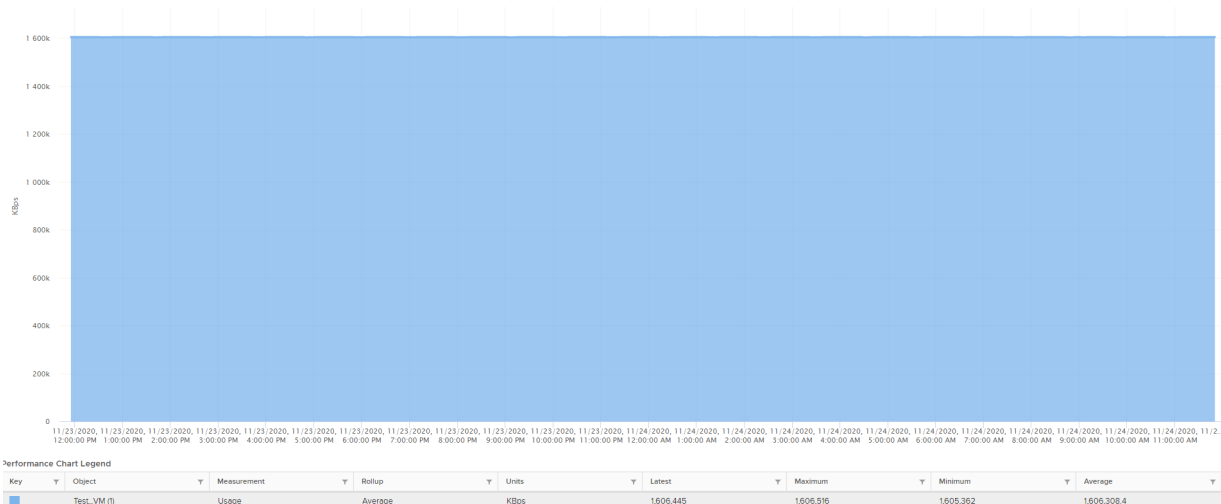


Figure 8: 24-hour Throughput of 16G Interface on LDCA

2021

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 ¹²	FC-PI-5
3	Cisco MDS 9148S-1	Cisco MDS 9148S-1 ¹³	FC-PI-5

Table 14: FC Switch Combinations

FC Switch	Huawei OptiXtrans E6608T / OptiX OSN 1800 II TP
FC-PI-6	FC-PI-6
FC-PI-5	FC-PI-5 and FC-PI-6
FC-PI-3	FC-PI-3

Table 15: Optical Transceivers between FC Switch and DUT

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

- 100 km long haul with Brocade G620 / G620 and Brocade G620 / 6505 pairs, respectively
- 40 km¹⁴ calculated medium-haul with Cisco MDS 9148S pairs, respectively

¹² Brocade 6506 does not support FC-PI-6 and FC-PI-3

¹³ Cisco MDS 9148S does not support FC-PI-6 and FC-PI-3

¹⁴ Due to Cisco MDS 9148S buffer limitation, we can only test up to 40 km

Testbed Description

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 (2x 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.

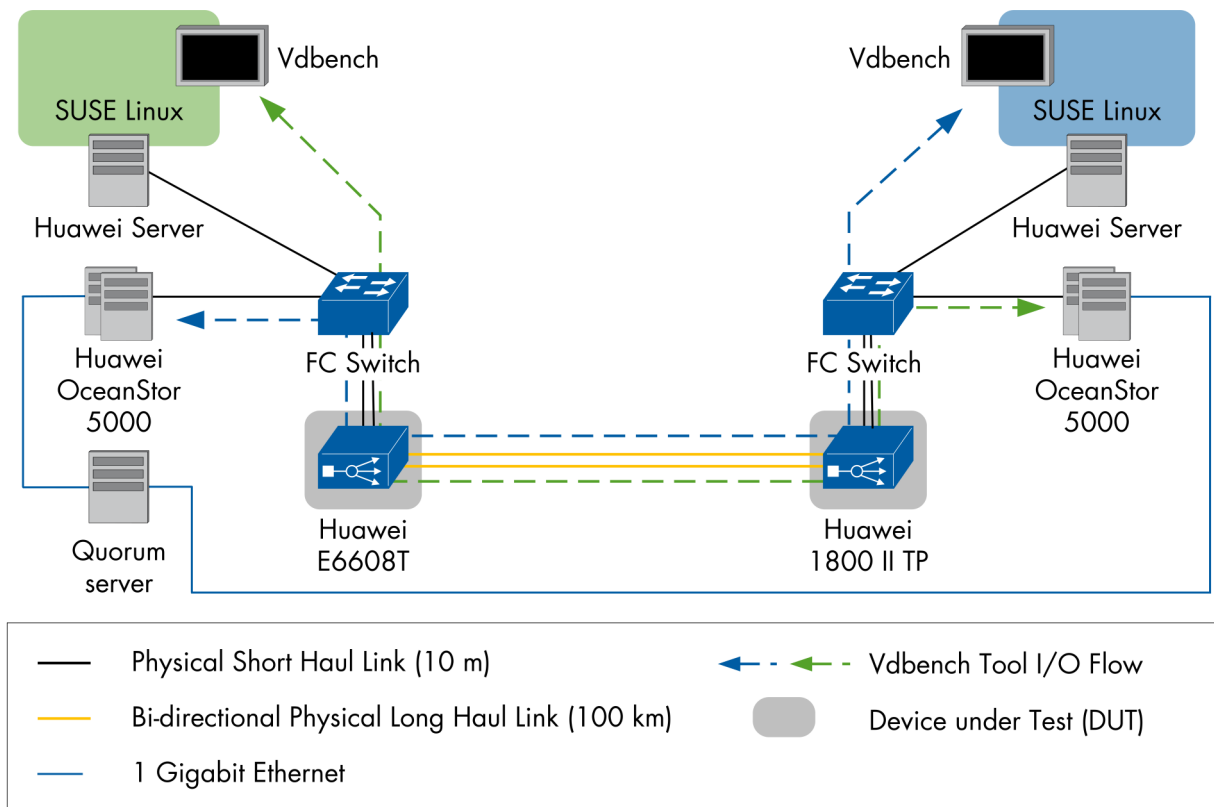


Figure 9: Logical Test Topology

Hardware and Software

Product Type	Product Name	Software Version
Devices Under Test WDM Equipment	Huawei OptiXtrans E6608T / OptiX OSN 1800 II TP	V100R021C10
	LDCA (board facing FC switch)	
	LTX (board facing FC switch)	
	OBU	
	OLP (board facing long-haul)	
	SCC STG	
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 ¹⁵	V300R002C10
Quorum Server	Quorum Server	V300R002C10

Table 16: Device under Test - Hardware and Software

FC Switch	Software Details
Brocade 6505	FOS v8.2.1a
Brocade G620	FOS v9.0.1c
Cisco MDS 9148S	NX-OS version 6.2

Table 17: FC Switch Hardware and Software Details

¹⁵ 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/	Measured Throughput, per	Verdict
1. Brocade G620 pair	FC800	800	781-782	Pass
	FC1600	1600	1559-1560	Pass
	FC3200	3200	3113-3114	Pass

Table 18: LDCA Interface Throughput

LTX

Setup	Speed Type	Expected Throughput (MB/	Measured Throughput, per	Verdict
1. Brocade G620 pair	FC800	800	780	Pass
	FC1200	1200	1168-1172	Pass
	FC1600	1600	1558	Pass
	FC3200	3200	3107	Pass
2. Brocade G620/6505	FC1600	1600	6505: 1529 G620: 1562	Pass
5. Cisco MDS 9148S/9148S	FC800	800	777	Pass

Table 19: LTX Interface Throughput

¹⁶ 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 measured with the Viavi MTS5800-100G tester the WDM system's latency value. In this setup, we connected both DUTs back to back with the measurement tool (see figure), which provides latency value in microseconds' precision. We removed all data center devices from the test bed 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 8G, 10G, 16G, and 32G. 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 like in the figure.

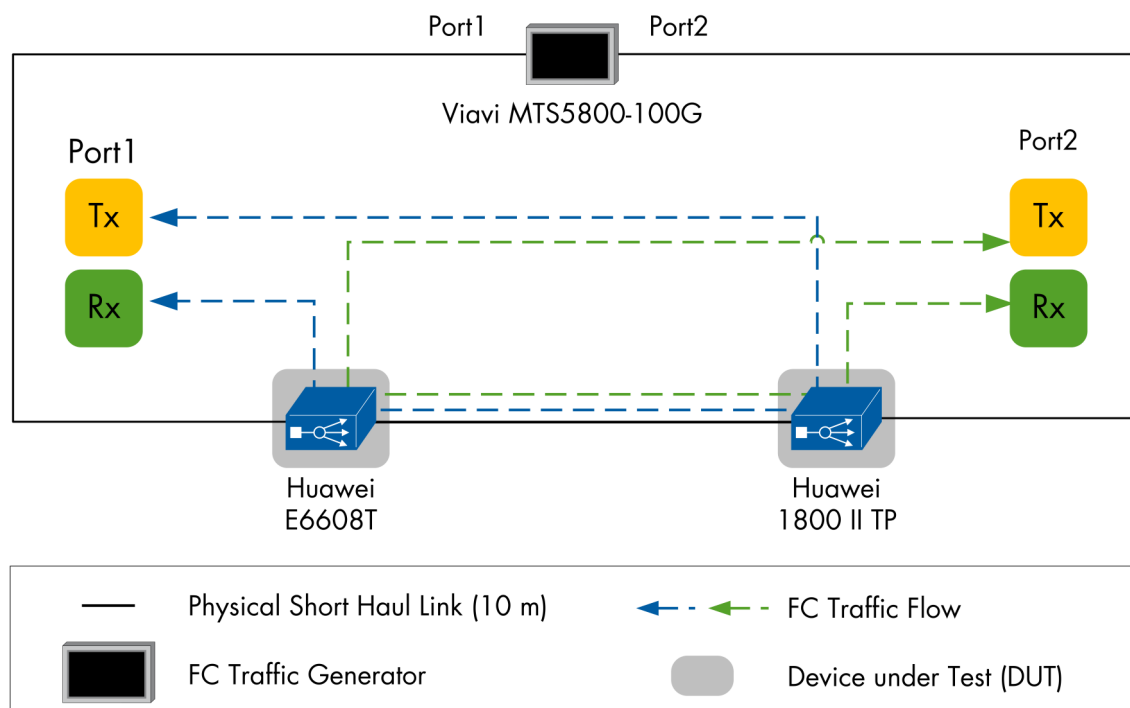


Figure 10: Latency Test Topology




Figure 11: Latency Results of LDCA and LTX

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.

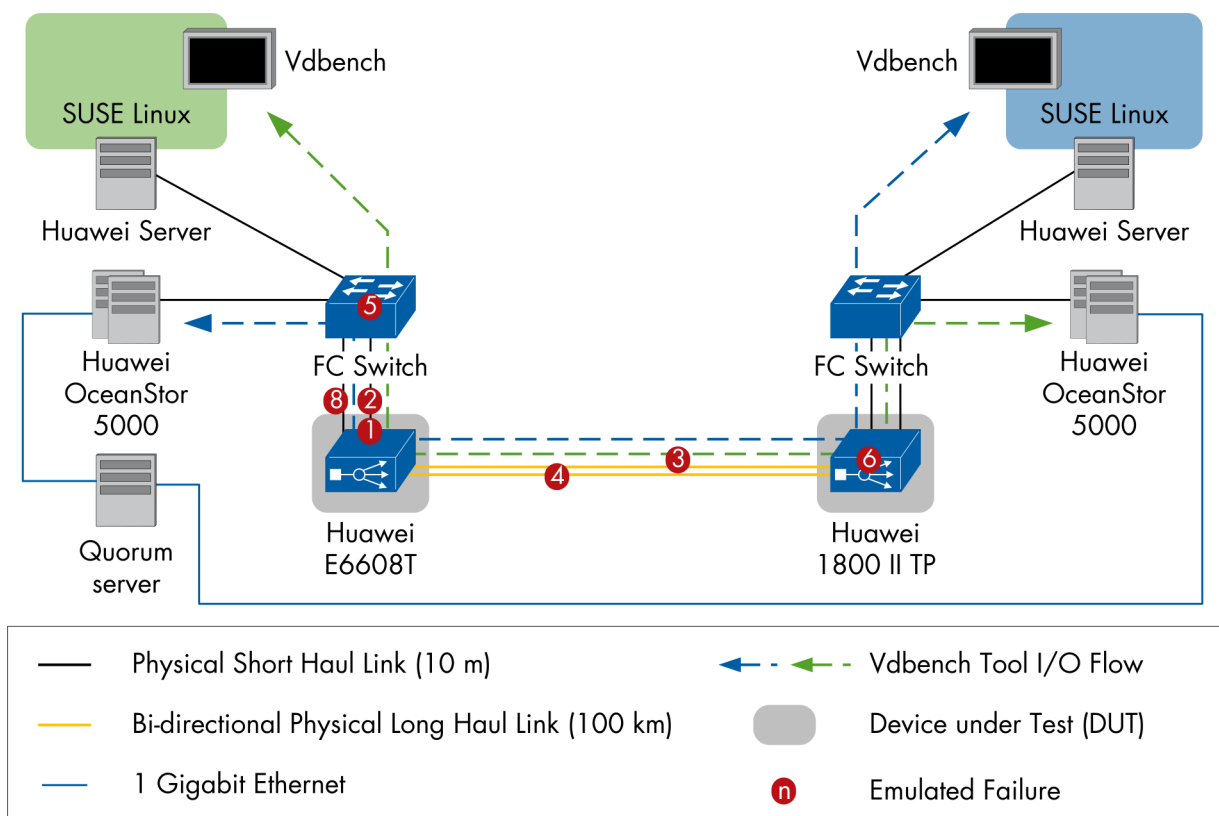


Figure 12¹⁷: Emulated Failure in the Test Bed

¹⁷ 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 the 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. The secondary storage stopped receiving any I/O access from the remote datacenter, 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 13), thereby causing it to stop transmitting data between both data centers. The latter phenomenon was due to the design mechanism of quorum server as described below.

Huawei OceanStor 5000 V3 implements a technology called "HyperMetro" for synchronizing the states between preferred storage and non-preferred 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.

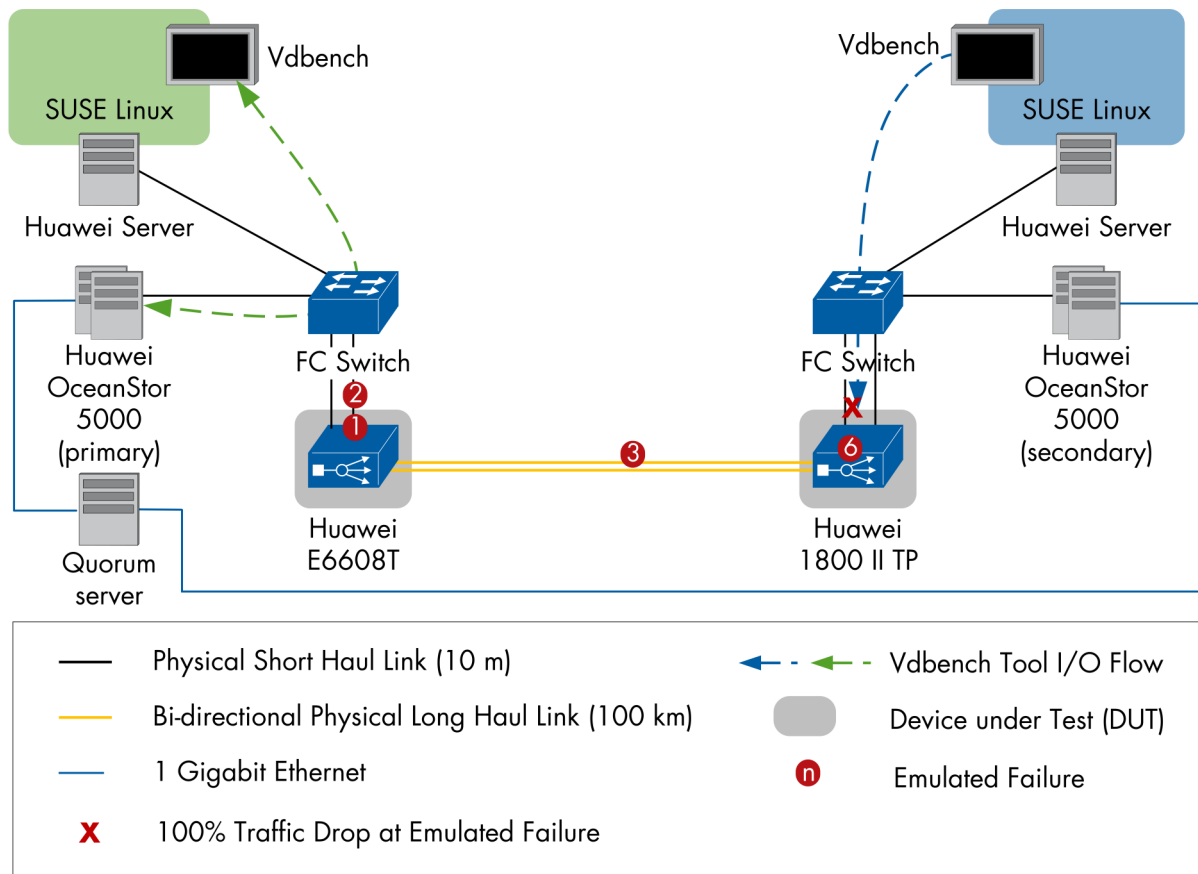


Figure 13: Traffic Status after the Switch Over for Failure Types 1, 2, 3 and 6

No. in Figure 13	Test Case	Action	Service Interruption		Verdict
			Expected ¹⁸	Observed (for primary storage) ¹⁹	
1	E-Port Disable/Enable Test	Disable	10s	A maximum of 7s service interruption time, including 5s drop to 0MB/s	Pass
		Enable	No impact	No impact	Pass
2	E-Port Cable Disconnect/Reconnect Test	Disconnect	10s	A maximum of 7s service interruption time, including 5s drop to 0MB/s	Pass
		Reconnect	No impact	No impact	Pass
3	Long Haul Network Failure without Redundancy	Disconnect	10s	A maximum of 7s service interruption time, including 5s drop to 0MB/s	Pass
		Reconnect	No impact	No impact	Pass
6	DUT Reset Test	Shutdown	10s	A maximum of 7s service interruption time, including 5s drop to 0MB/s	Pass
		Turn on	No impact	No impact	Pass

Table 20: Out of Service Overview (Active-Active Storage Cluster Protection)

¹⁸ The expected value included 5s fault detection by the quorum server and 15s from it to triggering an active-active switchover. Once the quorum server detected the fault, it triggers an active-active switchover, which lasts for 15 seconds (generally within 10 seconds in lab tests). After the switchover has been complete, the preferred storage took over services.

¹⁹ We observed 5s complete traffic drop to 0 MB/s during the switch over. The Huawei team configured the quorum server with a 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, the storage stopped receiving any data from the FC switch until the FC switch was available again. Depending on the time taken by the FC switch to become available, we observed:

- The Brocade G620 pair rebooted less than 1 min. The FC service did not fully interrupt from the reboot site. The reboot site traffic was redirected to the non-reboot site and kept working, whereas the non-reboot site switched to access the local storage as well. There was around 40-50s service interruption time (including a maximum of 40s drop to OMB/s).

- The Brocade G620 / 6505 pair rebooting behavior depended on which one to reboot. If we rebooted G620, it is the same behavior as upper's. If we rebooted 6505, then the link and the fabric were all down (status of out of service).

After the FC switch completed the reboot and the quiescing time (300s) passed, we observed that the traffic switched back to the initial status between DC1 and DC2. We did not observe any impact as expected.

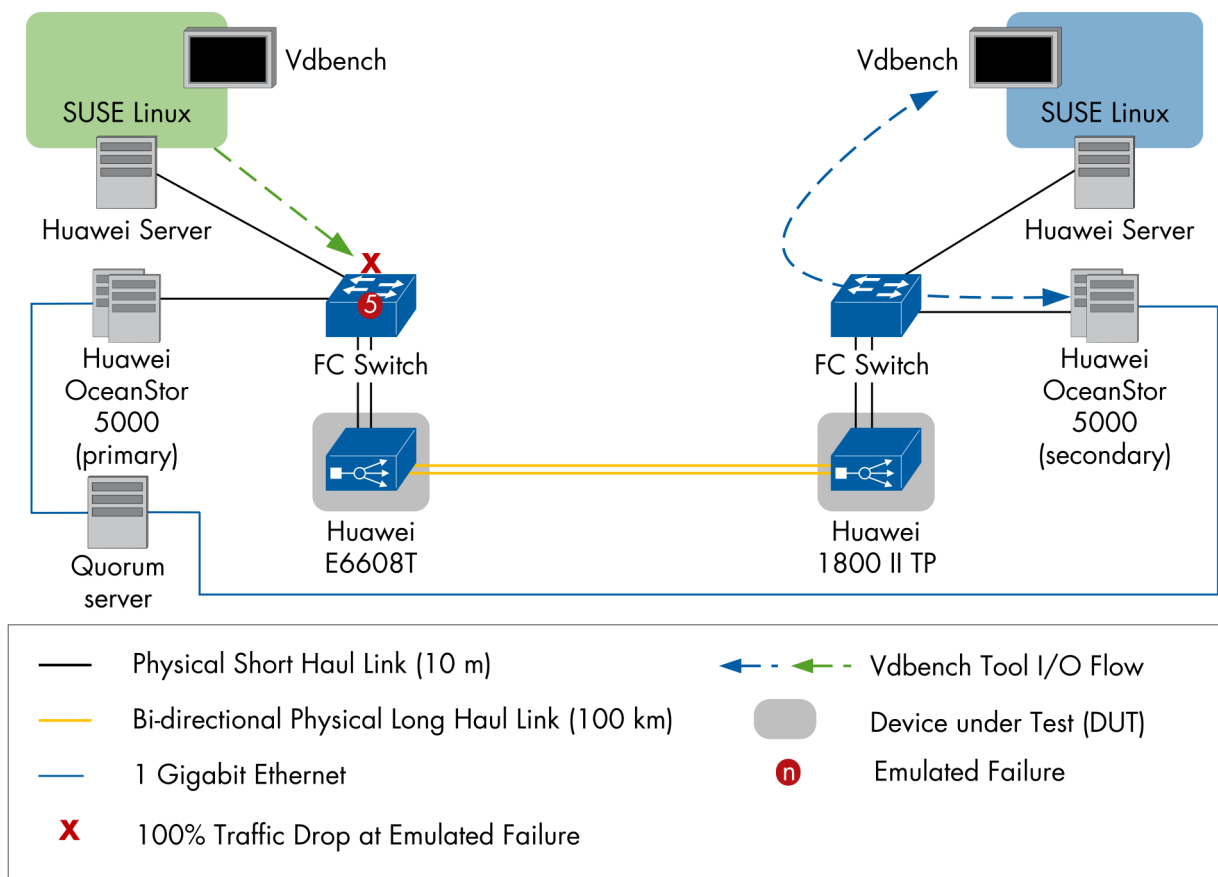


Figure 14: Traffic Status during FC Switch Reboot

No. in Figure 14	Test Case	Action	Service Interruption		Verdict
			Expected ²⁰	Observed (for primary storage) ²¹	
5	Switch Re-boot Test	Shut down	15s	A maximum of 12s service interruption time, including 10s drop to OMB/s	Pass
		Turn on	No impact	No impact	Pass

Table 21: Out of Service Overview (Active-Active Storage Cluster Protection)

Long Haul Link Protection

We measured the switch-over time for the system under test 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. The following Table shows the switch-over time measured. No session drops appeared during this 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 then plugged the fiber back to its original port, and we expect the traffic to increase by 100%.

²⁰ The expected value included 5s fault detection by the quorum server and 15s from it to triggering an active-active switchover.

²¹ The 10s 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 12	Test Case	Setup	Action	Service Interruption		Verdict
				Expected ²²	Observed	
4	Long Haul Network Failure with Redundancy	Brocade G620 pair	Disconnect	4 s ²³	A maximum of 4s service interruption time, including 1s drop to OMB/s	Pass
			Reconnect	No impact	No impact	Pass
		Brocade G620 6505 pair	Disconnect	4 s ²³	A maximum of 3s service interruption time, including 1s drop to OMB/s	Pass
			Reconnect	No impact	No impact	Pass
8	ISL Trunking		Disconnect	Drop 50%	Drop 50%	Pass
			Reconnect	Increase 100%	Increase 100%	Pass

Table 22: Out of Service Overview

Board Under Test	Switch Over Time (s)
LDCA	A maximum of 4s service interruption (including a maximum of 1s drop to OMB/s)
LTX	A maximum of 3 service interruption (including a maximum of 1s drop to OMB/s)

Table 23: Switch Over Time Brocade G620 - G620

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

Table 24: Switch Over Time Brocade G620 - 6505

²² 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 500ms (milliseconds that a frame could be buffered on a port without being overwritten) configured on the Brocade G620 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.

²³ The expected value includes 4s of impact by end-to-end flow control (see explanation as provided in * note), plus 2s of impact by edge hold off time (EHT) on F port. Note, the value is not provided by Brocade 300 manual, we calculated a double theoretical value of 250ms for FC1G and FC2G (250 ms is the default EHT value of FC16G), so added 1s per F port (based on 1 sample/sec based on vdbench).

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 25 combinations as listed in the tables below.

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

Table 25: Brocade G620 Pair

Board Under Test/ Test Scenario	E-port Interface used in Test Case		
	1	4	5
LTX	16G	16G	16G

Table 26: Brocade G620 - 6505 Pair

Board Under Test/ Test Scenario	E-port Interface used in Test Case		
	2	3	6
LTX	8G	8G	8G

Table 27: Cisco MDS 9148S Pair

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. LDCA was not tested.

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.


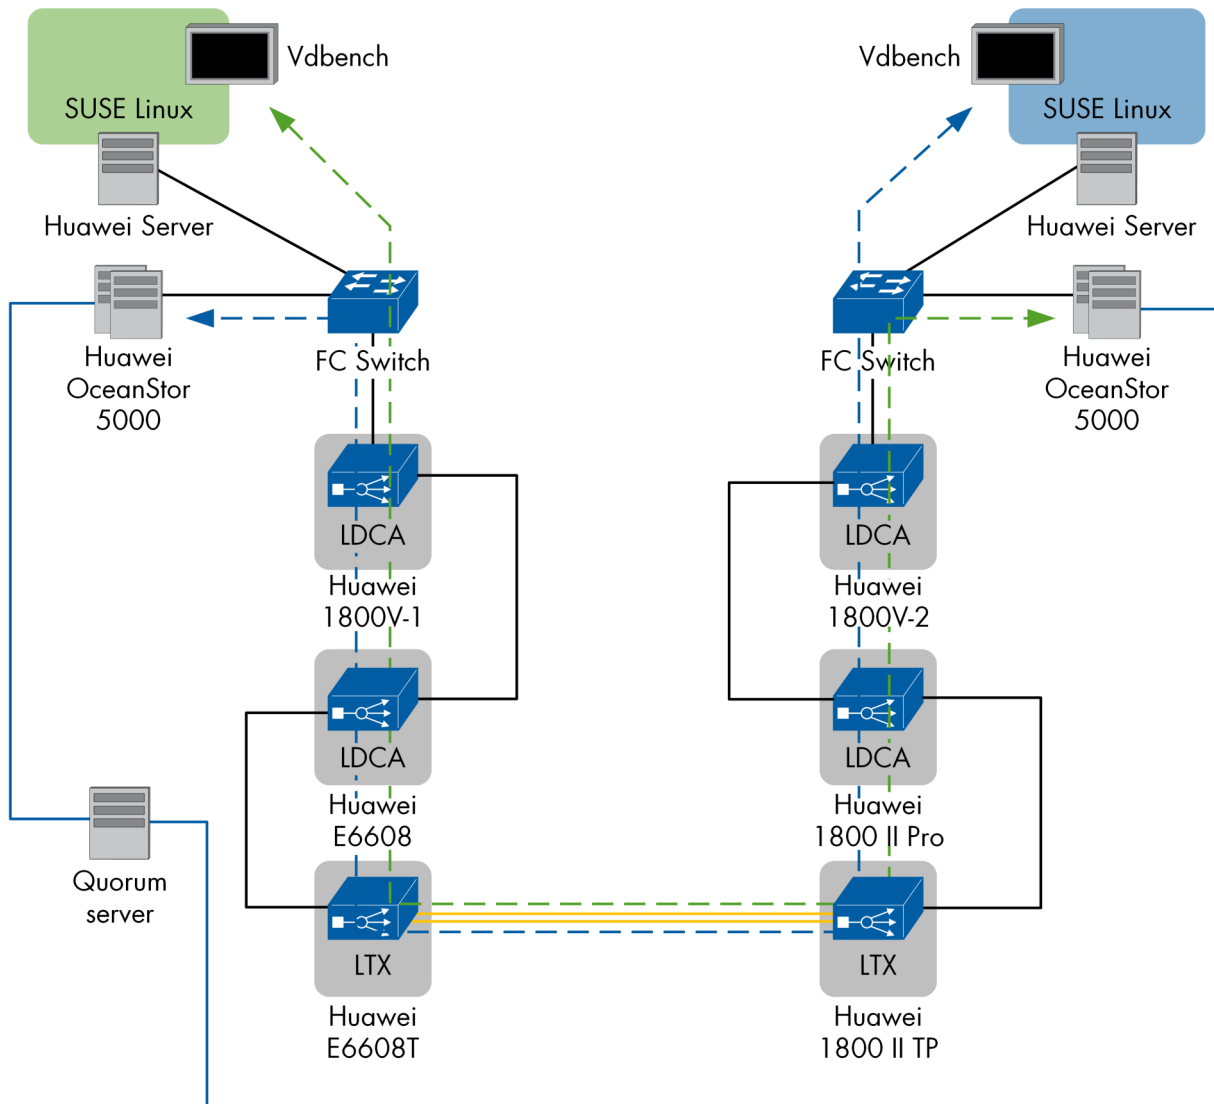


Figure 15: 24-hour Throughput Topology of LTX




Figure 16: 24-hour Throughput of LTX at FC16G

Conclusion

In 2020, EANTC verified the interoperability of Huawei OptiXtrans E6608T and third-party Fibre Channel switches from Brocade (6505 and G620) and Cisco (MDS 9148S and MDS 9132T). Multiple optical transceiver functions were certified in the E_port between these FC switches and the device under test (DUT), including FC-PI-6, FC-PI-5, and FC-PI-3. We validated forwarding speeds in FC 4G, 8G, 10G, 16G, and 32G scenarios. When forwarding traffic at any of these standardized speeds, the DUT did not exhibit any speed impact. The extended 24 hours soak testing confirmed stable operations of DUT 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 ELOM, LDCA and LDX line card.

All tests documented in this report passed our verification, some required the addition of optical amplifiers to the long-haul fiber.

In 2021, EANTC verified two boards LDCA and LTX respectively. The main differences between 2020 and 2021 are software differences both on the WDM system and Brocade G620 FC switch. LDCA and LTX showed good compatibility and protection with the latest Brocade G620 firmware (9.0.1c). LTX 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 II TP (OptiXtrans E6608T) fulfills Huawei's claims to work in enterprise data center interconnection scenarios as an integrated, high-speed, 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]