EANTC Independent Test Report

Huawei WDM Compatibility Test

OptiXtrans DC908

December 2020

Introduction

Huawei commissioned EANTC to validate functional, interoperability, and performance aspects of the Huawei OptiXtrans DC908 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 May 2020.

As part of innovative digital transformation strategies, enterprises worldwide are expanding their private clouds and public cloud services used, leading to data center (DC) growth both on the enterprise and cloud provider sides. The cloud computing paradigm promises (and requires) on-demand, location-independent, and highly available services. To implement these services, the scalability and availability of data centers are a crucial factor. The DC network infrastructure needs to be architected to support a scalable and robust cloud-based solution. It is considered a standard approach to implement two or more redundant data centers on campus or in the region, interconnected with very high bandwidth and low latency. This way, services can be adequately protected against partial or complete failures of one of the data centers.

This technical solution is known under the name Data Center Interconnect (DCI). It can be implemented by different network technologies (i.e., dark fiber, wavelength multiplexing, microwave, MPLS VPN, or Internet). The technical assessment of each solution is evaluated based on different criteria like available physical bandwidth, link protection, and diversity, link manageability and security, or deployment complication.

Huawei OptiXtrans DC908 is an optical-electrical WDM transmission device. 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).

In this test, EANTC verified the real enterprise case with two emulated data centers for the following use cases focused on Fibre Channel (FC). In a WDM DCI scenario, there is usually packet data traffic (Ethernet/IP-based) and storage traffic forwarded over the wide-area link.

Test Highlights

- → Interoperates with Brocade 6505/G620 and Cisco MDS 9148S/9132T Fibre Channel (FC) switches
- → Supports FC ports at 8G, 10G, 16G, 32G nominal link speed
- → Implements WAN-link fail over capabilities
- → Provides deterministic propagation delay between two 100 km geographically separated WDM network nodes
- → System remains stable for 24 hours soak testing

Huawei asked us to focus specifically on FC interoperability and performance because storage traffic has more stringent qualify of service and physical interoperability requirements with regards to the WDM transport.

We validated:

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

[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.

Device Under Test

Huawei explained that the OptiXtrans DC908 is designed for DCI and ready for simplified deployment, ultrabroadband and high integration data traffic.

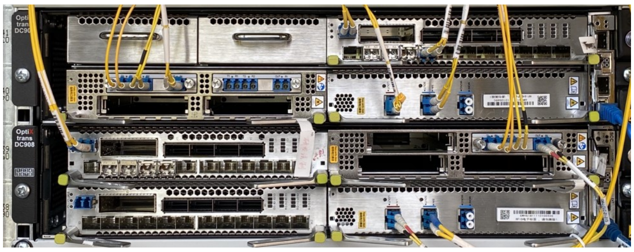

Figure 1: Huawei DC908

Executive Summary

The modular setup of Huawei DC908 simplified the installation of hardware and software and aims to reduce operating expenses (OPEX). All tests were based on the same configuration through automatic detection of optical transceiver and successfully integrated with 3rd party Fibre Channel switches of different vendors represented by Brocade and Cisco. The ISL protection included physical port and logical trunking.

We primarily focused on the FC-PI-6 interface functions. Our tests were carried out with DC908s connected to a pair of Brocade G620 switches and subsequently to a pair of Cisco MDS 9132T switches, both running FC32G. In addition, we verified compliance for FC-PI-3 Inter-Switch Link (ISL) functions between a pair of Brocade G620s running FC10G. Finally, to verify the backwards compatibility of the DC908, we tested compliance with the FC-PI-5 protocol based on hybrid switch pairs of Brocade G620 / Brocade 6505 and Cisco MDS 9132T / Cisco MDS 9148S, respectively.

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 km[2] medium haul with Brocade 6505 and Cisco MDS 9148S pairs, respectively

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.

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).

In a soak test environment, we put the Huawei DC908 under continuous load for 24 hours to make sure it would support uninterrupted service. The system remained stable without any restart or service interruption, zero packet loss as expected.

Testbed Description

The following figure describes the logical topology of the test setup.

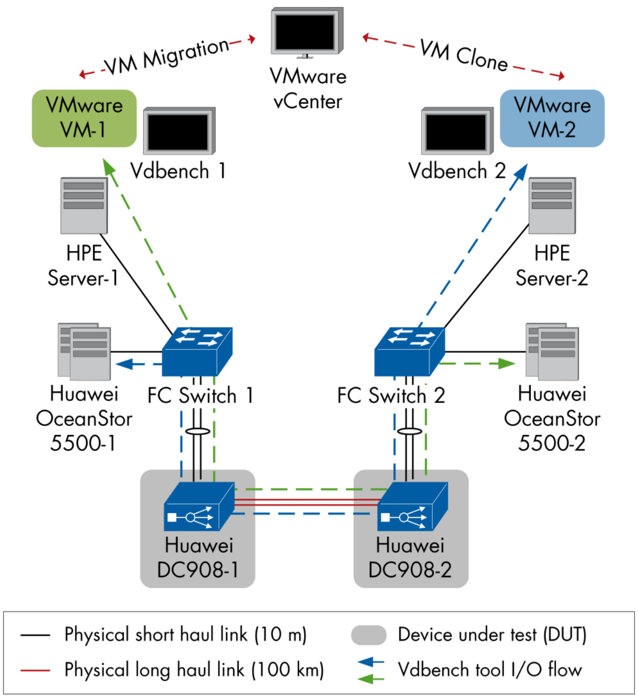

Figure 2: Logical Data Center Setup

A FC protocol standardizes a variety of logical ports. We focused on E_Port (Expansion port), a port that is connected to another E_Port to create ISL for the remote communication. In our case, it is between FC switch and Huawei DC908.

[2] 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).

Hardware and Software

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.

Product Type	Product Name	Software Version
Devices Under Test WDM Equipment	Huawei OptiXtrans DC908	V100R019C10
	MD02A (line card facing FC switch)	
	OLL (line card facing WDM)	
	OPC (line card facing WDM)	
Physical Server	HPE DL380 Gen9	
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 1: Hardware and Software Components

Optical Transceiver

Multiple combinations of FC-PI-x optical transceivers were used in the E_Port for the compatibility testing. Table 2 describes the matrix.

Scenario	FC Switch		WDM
	Brocade G620	Brocade 6505	Huawei DC908
Setup 1	FC-PI-6	Not Applicable[4]	FC-PI-6
Setup 2	FC-PI-5	FC-PI-5	FC-PI-6
Setup 3	FC-PI-3[3]	Not Applicable[4]	FC-PI-3[3]
	Cisco MDS 9132T	Cisco MDS 9148S	Huawei DC908
Setup 4	FC-PI-6	Not Applicable[4]	FC-PI-6
Setup 5	FC-PI-5	FC-PI-5	FC-PI-6

Table 2: Matrix of Optical Transceivers for the E_Port

Table 2 shows the combinations of the tested FC-PI-x optical transceivers in two data centers. One FC switch and One WDM system was included in each data center.

Please find a detailed description of each setup below:

- Setup 1 shows the FC-PI-6 optical transceivers that were used in two Brocade G620 FC switches of two data centers, and FC-PI-6 in two Huawei DC908 systems;
- Setup 2 means FC-PI-5 optical transceiver was used in Brocade G620 FC switch of one data center, FC-PI-5 in Brocade 6505 of the other data center, and FC-PI-6 in two Huawei DC908 systems in two data centers;
- Setup 3 means FC-PI-6 optical transceivers were used in two Cisco MDS 9132T FC switches of two data centers, and FC-PI-6 in two Huawei DC908 systems;

[3] According to FC standards, ISL runs not only at the link speed of 8G, 16G, 32G but also 10G. We selected FC-PI-3 optical transceivers in the Brocade G620 FC switch and Huawei DC908 for a 10G test.

[4] According to product documentation, Brocade 6505 and Cisco MDS 9144S support up to FC-PI-5 optical transceivers. FC-PI-6 and 10G FC-PI-3 tests are marked as 'Not Applicable' for Brocade 6505 and Cisco MDS 9144S.

- Setup 4 means FC-PI-5 optical transceiver was used in Cisco MDS 9132T FC switch of one data center, FC-PI-5 in Cisco MDS 9148S of the other data center, and FC-PI-6 in two Huawei DC908 systems in two data centers;
- Setup 5 shows the FC-PI-3 optical transceivers that were used in two Brocade G620 FC switches of two data centers, and FC-PI-3 in two Huawei DC908 systems;

Test Environment and Test Tools

VM clone and migration operations were conducted for the connectivity testing. The operations were managed by VMware vCenter. vCenter triggered the test VM provisioning on one of the data centers, including local storage. Once the test VM had been started, we created 200 GB large test files locally on the test VM. The test VM with these large size test files was cloned and/or migrated to a host and storage of the remote data center, connected through FC switches and Huawei DC908.

The test tool Vdbench used for the performance test and the long-run stability testing. The Vdbench was running in the test VM with the following configurations:

- Guest OS: CentOS 6.10, 64-bit
- Virtual disk: one TB Mapped Raw LUN
- Number of virtual disks: Five
- Number of virtual CPUs: Ten
- Guest memory: 200 GB
- Host Bus Adapter (HBA): one QLE2692 16Gb Fibre Channel to PCIe Adapter

The baseline traffic was generated with Vdbench that was running on one of the five Raw Device Mappings (RDM). The other four RDMs were used to generate the baseline traffic with the block size of 20 MB at maximum speed.

Test Results

Interoperability Testing

Firstly, we focused on the interoperability testing to verify hardware compatibility of the multiple optical transceivers between FC Switches and WDM systems.

FC-PI-6 Tests

FC-PI-6 supported up to 32G link speed, and maintained backward compatibility to 16G and 8G link speeds. We selected 32G link speed for all operations. Additionally, we selected 8G and 16G link speeds for tests performing administrative operations in the optical domain across the WDM link.

The predefined FC-PI-6 optical transceivers connected in FC Switches and WDM systems, and then ISL established between FC Switches and WDM systems. We used VM migration and VM Clone to verify the interoperability with the VM operation traffic from VMware vCenter. We expected the successful ISL establishment, VM migration and VM Clone. The tests pass when all these expectations are achieved.

Test Case	FC Link	VM Migration and VM Clone	
		Setup 1	Setup 4
		Brocade G620 pair	Cisco MDS 9132T pair
E-Port Disable/ Enable Test	32G	PASS	PASS
	16G	PASS	PASS
	8G	PASS	PASS
Long Haul Network Failure	32G	PASS	PASS
	16G	PASS	PASS
	8G	PASS	PASS
ISL Trunking Test	32G	PASS	PASS
	16G	PASS	PASS
	8G	PASS	PASS
E-Port Cable Disconnect/ Reconnect Test	32G	PASS	PASS
	16G	PASS	PASS
	8G	PASS	PASS
Long Haul Network Failure with Protection	32G	PASS	PASS
	16G	PASS	PASS
	8G	PASS	PASS
Switch Reboot Test	32G	PASS	PASS
	16G	PASS	PASS
DUT Reset Test	32G	PASS	PASS
	16G	PASS	PASS
	10G	PASS	PASS
	8G	PASS	PASS
Extended Duration I/O Test	32G	PASS	PASS

Table 3: FC-PI-6 Tests

FC-PI-5 Tests with Hybrid Fiber Channel Switch Pair

For backwards compatibility in a hybrid Fibre Channel switch environment, as defined in setup 2 and 4 (see table 2 above), we selected 8G and 16G link speeds for the long haul link failure test and the line card removal test that challenge WDM in the optical domain.

The same processes and expectations used in this predefined FC-PI-5 tests to the above ones of FC-PI-6 tests.

Test Case	FC Link	VM Migration and VM Clone	
		Setup 2	Setup 5
		Brocade 6506 - Brocade G620	Cisco MDS 9148S - Cisco MDS 9132T
Long Haul Network Failure	16G	PASS	PASS
	8G	PASS	PASS
DUT Reset Test	16G	PASS	PASS
	8G	PASS	PASS

Table 4: FC-PI-5 Tests with Hybrid Fibre Channel Switch Pair

FC-PI-3 Tests

For 10G ISL testing, we selected FC-PI-3 with the Brocade G620 pair. We verified the E_Port removal test, the long haul link failure test and the line card removal test.

The same processes and expectations used in this predefined FC-PI-3 tests to the above ones of FC-PI-6 tests.

Test Case	FC Link	VM Migration and VM Clone
		Setup 3
		Brocade G620 pair
ISL Trunking Test	10G	PASS
E_Port Fiber Disconnect/ Reconnect Test	10G	PASS
Long Haul Network Failure with Protection	10G	PASS
DUT Reset Test	10G	PASS

Table 5: FC-PI-3 Tests

Capability Testing

Secondly, we conducted functional capability tests. To verify the long-distance forwarding of the WDM between the data centers, we concentrated on administrative use cases to allow operators to perform maintenance activities. We simulated customer-facing operations in the optical domain, such as changes in the E_Port configuration of the FC switches, disconnection of fibers from the FC switches facing WDM system, and even the FC switch reboot while the WDM was forwarding. We carried out physical operations on the WDM system, such as removal and reinsertion of the line card. In this domain, we looked at the link protection test by the ISL trunking to protect the WDM against link failure.

E_Port Disable/Enable Test

We verified the interoperability of the WDM solution with four FC switches at each with the defined speed. The focus of this test is that when the state of the E_Port changes, the WDM system remains transparent and will not affect the state of the E_Port. As specified in the Fibre Channel interconnect protocol (vendor-proprietary, supported by Brocade and Cisco), an ISL implements and terminates at an E_Port on each end. In the data center under test, the WDM link between FC switches implements the ISL link between the FC switches and is connected to the optical domain with the E_Port. A F_Port on the FC switch is used to connect with a storage system or a physical server.

We performed a baseline test to verify the throughput of the interface along with FC switches through the WDM system. The baseline test was executed in setup 1 and 4. The performance result is shown below.

FC Link	Setup 1		Setup 4	
	Expected Throughput (GB/s)[5]	Measured Throughput (GB/s)	Expected Throughput (GB/s)	Measured Throughput (GB/s)
32G	3.104	1.539[6]	3.104	1.539[6]
16G	1.552	1.539	1.552	1.536
8G	0.776	0.769	0.76	0.745

Table 6: Achieved Speed Running Vdbench

5 The expected throughput is based on layer 2 payload, we used a ration of 97% at link speed. The ratio is based on the formular: 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.

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

While running the baseline traffic, we disabled the two E_Ports on the both FC switches in two data centers. After observing the decreased traffic rate in storage traffic and error log from the FC switch reporting the port went down to confirm that the configuration took effect, we performed a VM migration which failed as expected. Then, we activated the port in the configuration and observed that the VM migration successfully completed indicating the restored status of the E_Port. The performance of the maximum throughput is shown below.

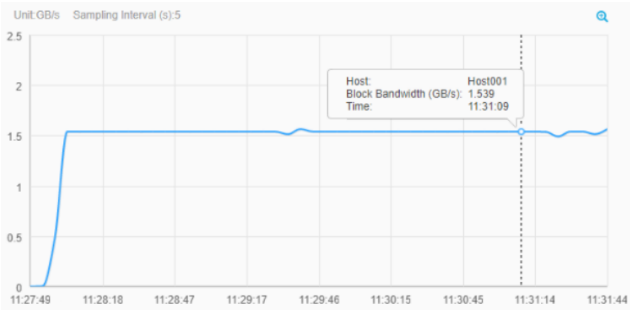

Figure 3: Maximum Throughput at 1.539 GB/s

E_Port Fiber Disconnect/Reconnect Test

In this test, we verified that the WDM system remains transparent when the E_Port fibers are disconnected and reconnected from the FC switch. To achieve repeatability we used the same baseline traffic from the previous test case. The baseline test was executed in setup 1, 4 and 5, the performance result is as follows.

We unplugged all fibers at the E_Ports on the current FC switch. After observing the decrease in traffic in the storage system and the error log of the FC switch reporting that the port had an expected failure, the VM operation failed as expected, so we plugged in the fiber. As expected, the VM returned to the operating state and the VM migration and cloning continued.

ISL Trunking Test

We verified the link protection between the WDM system and the FC switch. The traffic continued on the remaining link when one link failed. While running baseline traffic, we removed one of two running fibers from ISL trunk and observed that the VM operation was continued on the remaining link as expected. The result of the switch over time is shown below.

After we reconnected the link to the ISL trunk we did not observe any impact on the baseline traffic.

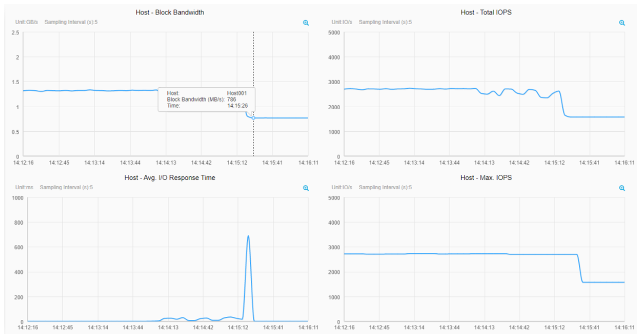

Figure 4: Switch Over Time in Storage GUI Graph

FC Link	Setup 1		Setup 4		Setup 3	
	Expected Throughput (GB/s)	Measured Through-put (GB/s)	Expected Throughput (GB/s)[5]	Measured Throughput (GB/s)	Expected Throughput (GB/s)[5]	Measured Throughput (GB/s)
32G	3.104	1.539[7]	3.104	1.539[7]	Not applicable	
16G	1.552	1.536	1.552	1.521		
10G	Not applicable		Not applicable		1.164	1.156

Table 7: Achieved Speed Running Vdbench

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

DUT Reset Test

We verified that the DUT forwards storage data after the line card has been re-inserted.

While running baseline traffic, we removed the MD02A line card from the Huawei DC908. We observed that the VM operation failed during the removal of the MD02A line card. We reinserted the line card into the DUT and observed a successful VM migration.

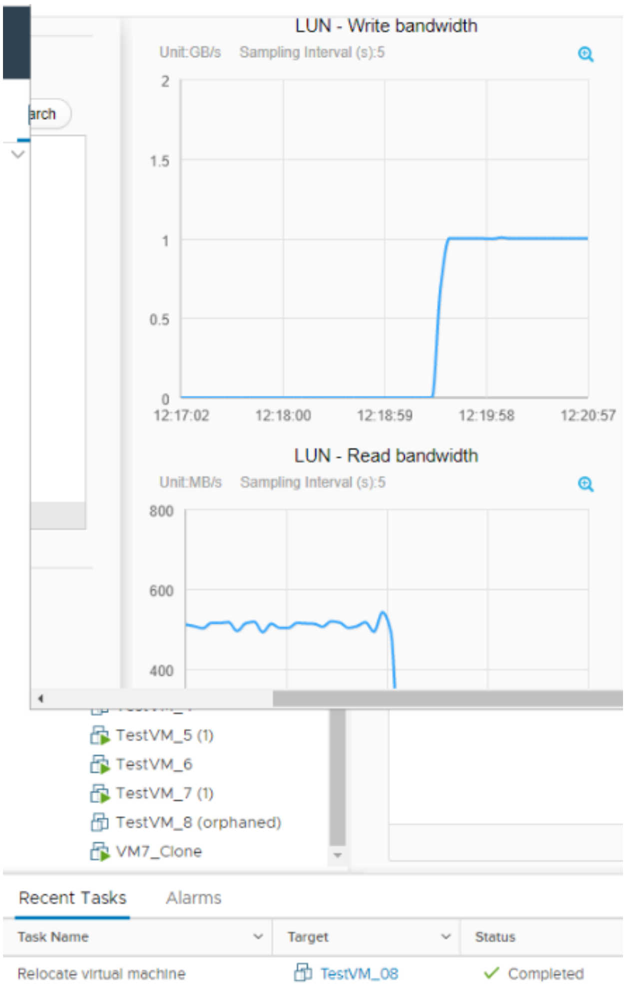

Figure 5: Normal VM Operation after the Resetting

Long Haul Network Failure

We validated that the WDM system continues to forward storage data after long haul links failure. We used the same baseline traffic from the previous test case to achieve repeatability. While running baseline traffic, we removed both long haul fibers from the Huawei DC908 and observed a failed VM operation and I/O operations of storage. We reconnected the long haul fiber and observed a successful VM clone.

The baseline test was executed in all four setups. The performance result of setup 2 and 5 is shown below. The performance result of setup 1 and 4 can be found above in Table 6.

Long Haul Network Failure with Protection

We verified that the WDM link protection and traffic switches over to the backup link when the primary link fails.

While running baseline traffic, we removed one of the two long haul fibers from the Huawei DC908 and observed a shortly impacted VM operation and I/O operations of storage as expected. We only kept one long haul fiber and observed successful VM clone.

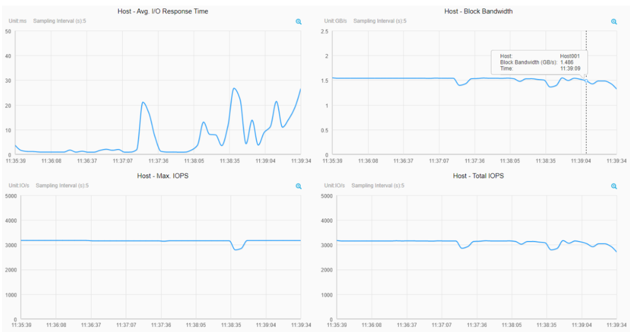

Figure 6: Link Fail Over Time in Storage GUI Graph

FC Link	Setup 2		Setup 5	
	Expected Throughput (GB/s)[5]	Measured Throughput (GB/s)	Expected Throughput (GB/s)	Measured Throughput (GB/s)
16G	1.552	1.454[8]	1.552	1.414[8]
8G	0.776	0.745	0.776	0.727

Table 8: Achieved Speed Running Vdbench

8 The traffic rate results proved our speculation that 100 km long haul is not recommended for setup 2 and 5 (the required buffer size not supported on the Brocade 6505 and the Cisco MDS 9148S).

Switch Reboot Test

We verified that the WDM system forwards storage data after the switch connected to the MD02A line card was rebooted.

While running benchmarking traffic, we restarted the FC switch connected to one of the Huawei DC908 devices. As repeatedly observed, the VM operation failed and the I/O operation throughput decreased. With the switch up and running again, we repeated the baseline throughput test and observed that the VM migration was completed successfully.

Stability Testing

To verify the stability of the WDM between data centers, we soaked it with the continuous baseline traffic for 24 hours.

Extended Duration I/O 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 bidirectionally for 24 hours. During that time, we monitored the system log of both hardware and software. We confirm that the system is able to transfer the data by a consistent rate for a duration of 24 hours. Additionally, we didn't observe any software crashes or hardware failure during the test duration.

Conclusion

We verified interoperability of Huawei DC908 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 Huawei DC908, including FC-PI-6, FC-PI-5 and FC-PI-3.

EANTC validated forwarding speeds in FC 8G, 10G, 16G and 32G scenarios. When forwarding traffic at any of these standardized speeds, the Huawei DC908 did not exhibit any speed impact. The extended 24 hours soak testing confirmed stable operations of DC908 without any traffic impact.

We conducted a range of availability tests to disable/enable port, disconnect/reconnect E_port fiber, disconnect/reconnect long haul, reboot FC switch, remove/re-install MD02A line card, and to remove/re-install ISL trunking fiber. All of the tests passed our verification.

Based on the results of our test, EANTC confirms that the Huawei OptiXtrans DC908 fulfills Huawei's claims to work in enterprise data center interconnection scenarios as an integrated, high speed, highly available solution.

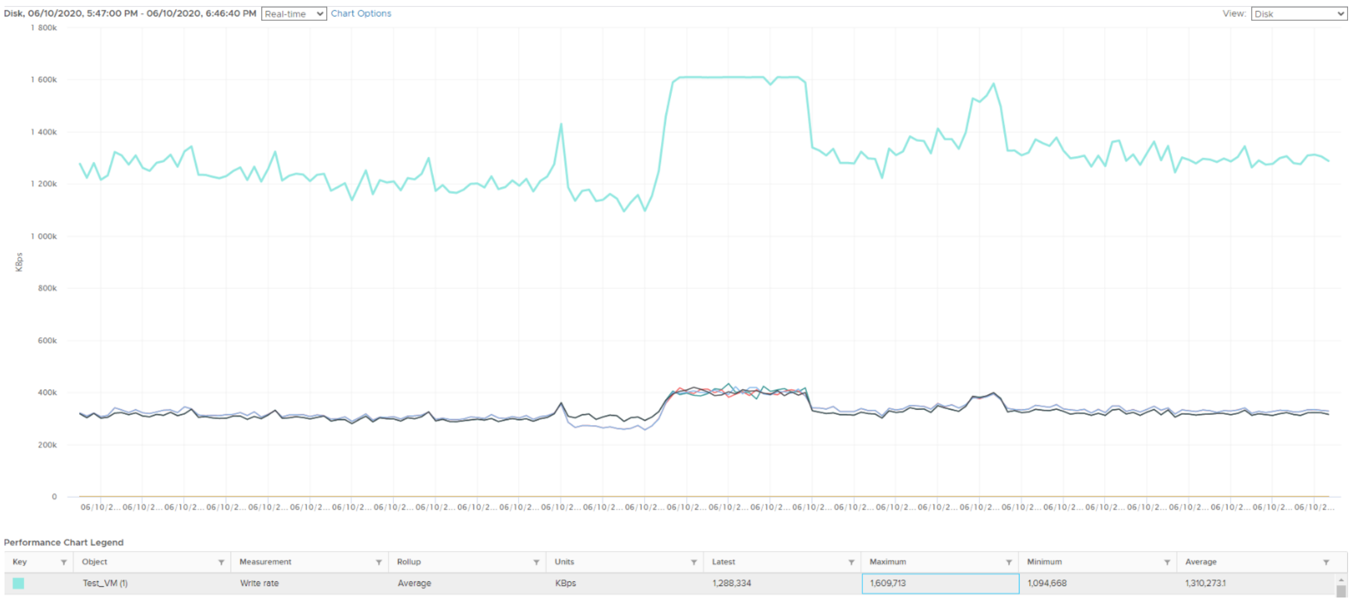

Figure 7: 24 Hours Load Test Result

About EANTC

EANTC (European Advanced Networking Test Center) is internationally recognized as one of the world's leading independent test centers for telecommunication technologies. Based in Berlin, the company offers vendor-neutral consultancy and realistic, reproducible high-quality testing services since 1991. Customers include leading network equipment manufacturers, tier 1 service providers, large enterprises and governments worldwide. EANTC's Proof of Concept, acceptance tests and network audits cover established and next-generation fixed and mobile network technologies.

EANTC AG
Salzufer 14, 10587 Berlin, Germany
info@eantc.de, https://www.eantc.de/
[v1.0 20201217]