EANTC Independent Test Report

Huawei WDM Compatibility Test

OptiX OSN 1800

April 2021

Introduction

Huawei commissioned EANTC to validate functional, interoperability, and performance aspects of the Huawei OptiX OSN 1800 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 OptiX OSN 1800 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 OptiX OSN 1800 II TP is an optical-electrical WDM transmission device. Huawei OptiXtrans E6608T 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).

Device Under Test

Huawei explained that the OptiX OSN 1800 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 three main boards of Huawei OptiX OSN 1800 vary at different speeds.

Highlights
→ 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]
→ Transparent multi-switch type forwarding between Brocade 6505 and Brocade G620, as well as between Cisco MDS 9148S and Cisco MDS 9132T, respectively
→ Stability of overnight soak testing
→ Forwarding performance of 4G, 8G and 10G Fibre Channel interfaces at B1ELOM board, as well as 8G and 10G Fibre Channel interfaces at B1LDX board; long haul transmission with 100 km long links using optical amplifiers
→ Forwarding performance of 8G, 16G and 32G Fibre Channel interfaces at B1LDCA board; long haul transmission with 100 km long links
→ Protection of long haul link[2]

Table 1: Features of the Board Types

[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

Figure 1: Huawei OptiX OSN 1800

Figure 2: Huawei B1ELOM Board

Figure 3: Huawei B1LDCA Board

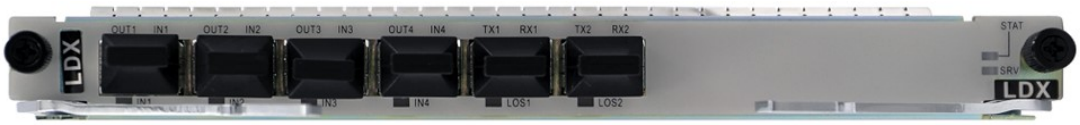

Figure 4: Huawei B1LDX Board

Executive Summary

We verified the compatibility between different speeds of Huawei OptiX OSN 1800 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, respectively.

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

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.

Testbed Description

Figure 5 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.

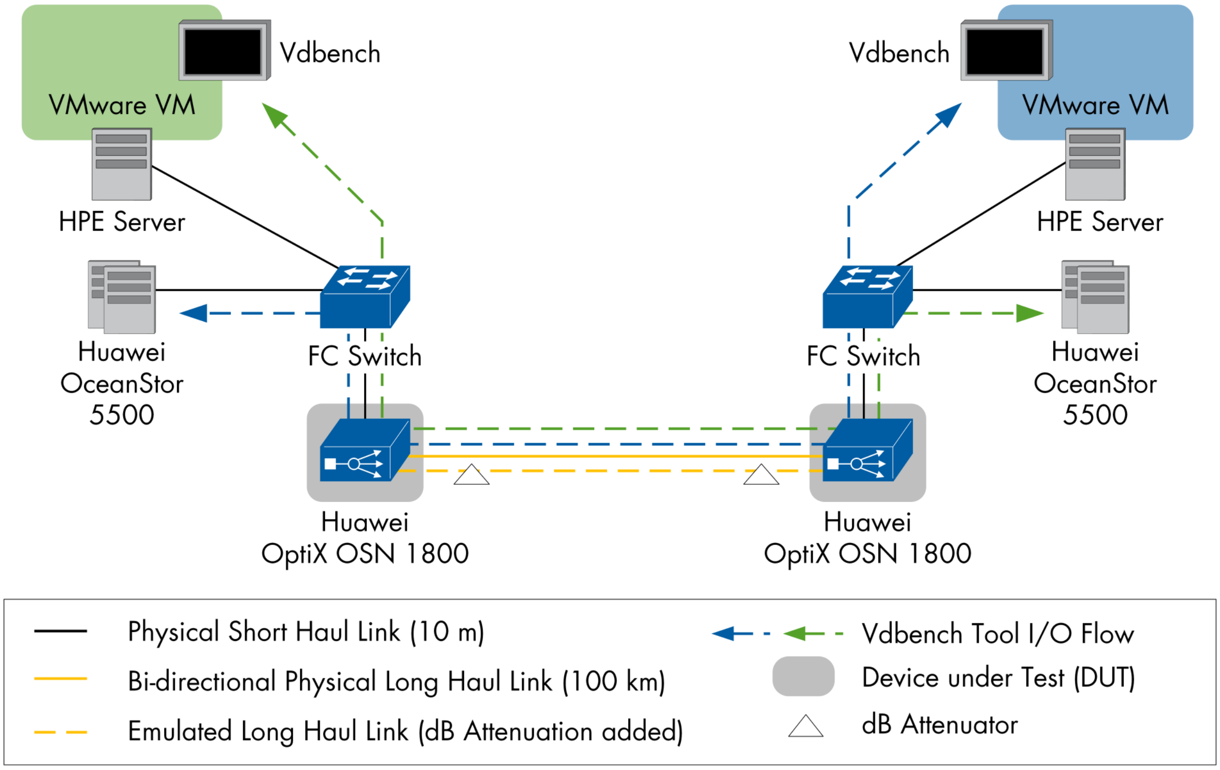

Figure 5: Logical Test Topology

Long Haul and Bandwidth

1. 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[3] calculated medium-haul with Brocade 6505 and Cisco MDS 9148S pairs, respectively

2. 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 B1LDX board type used a 200G optical transceiver on the WDM side, which was directly connected to 100 km optical fiber.

Both B1ELOM and B1LDX 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.

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

Table 2: FC Switch Combinations

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

Table 3: Optical Transceivers between FC Switch and DUT

Hardware and Software

Product Type	Product Name	Software Version
Devices Under Test WDM Equipment	Huawei OptiX OSN 1800	V100R019C10
	B1ELOM board (facing E-ports towards FC-switch)	
	B1LDCA board (facing E-ports towards FC-switch)	
	B1LDX board (facing E-ports towards FC-switch)	
	F2OBU (sub-module, between B1EX40 and B1OLP)	
	B1EX40 (sub-module, facing board under test)	
	B1OLP (sub-module facing long haul link)	

Table 4: Device under Test - Hardware and Software

[4] Brocade 6506 does not support FC-PI-6 and FC-PI-3

[5] Cisco MDS 9132T does not support FC-PI-3

[6] 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 5: 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 B1ELOM, B1LDCA, and B1LDX 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 B1LDCA 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).

B1ELOM

Setup	Optical Transceiver	Expected Throughput (MB/s)[7]	Measured Throughput (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 6: B1ELOM Interface Throughput

B1LDCA

Setup	Optical Transceiver	Expected Throughput (MB/s)[7]	Measured Throughput (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[8]	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[8]	Pass

Table 7: B1LDCA Interface Throughput

B1LDX

Setup	Optical Transceiver	Expected Throughput (MB/s)[7]	Measured Throughput (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 8: B1LDX Interface Throughput

[7] 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.

[8] 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 1 ms. This precision could not be measured by vdbench, as it had a resolution of 1 ms. 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 1 ms. 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)
B1ELOM	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
B1LDCA	1	FC-8G	150	20.2	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
B1LDX	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 9: Latency Overview

Hardware Robustness Testing

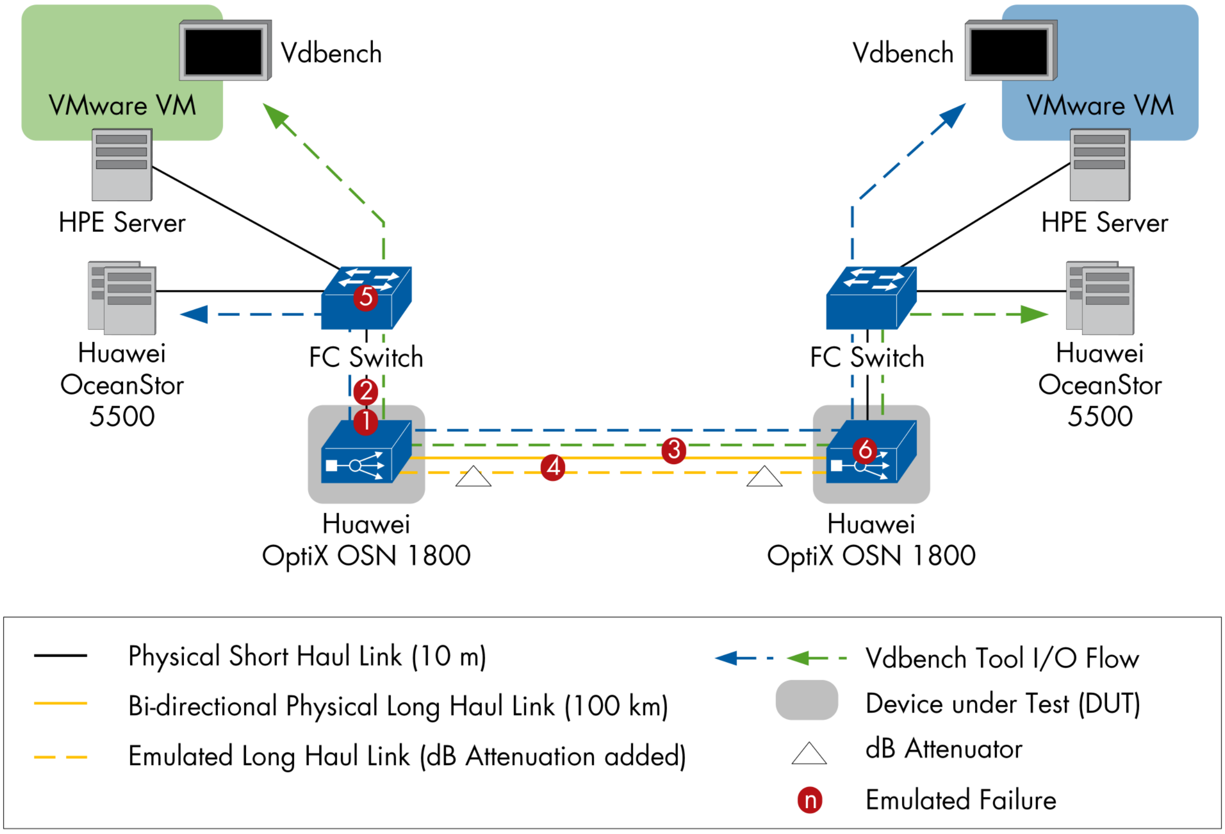

Figure 6: Emulated Failure

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 6. 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.

No. in Figure 6	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[9]	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 10: Hardware Robustness Tests Overview

Board Under Test	E-port Interface used in Test Case				
	1	2	4	5	6
B1ELOM	4G, 8G	8G	4G, 8G	8G	4G, 8G
B1LDCA	8G, 16G	16G, 32G	8G, 16G, 32G	16G	8G, 16G, 32G
B1LDX	8G	8G	8G	8G	8G

Table 11: Setup 1 - Brocade G620 Pair

Board Under Test	E-port Interface used in Test Case		
	2	4	6
B1ELOM	10G	10G	10G
B1LDCA	10G	10G	10G
B1LDX	10G	10G	10G

Table 12: Setup 2 - Brocade G620 Pair, 10G

[9] 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
B1ELOM	4G, 8G	4G, 8G
B1LDCA	8G, 16G	8G, 16G
B1LDX	8G	8G

Table 13: Setup 3 - Brocade G620/6505 Pair

Board Under Test	E-port Interface used in Test Case					
	1	2	3	4	5	6
B1ELOM	4G, 8G	8G	N/A	4G, 8G	8G	4G, 8G
B1LDCA	8G, 16G	16G, 32G	16G	8G, 16G, 32G	16G	8G, 16G, 32G
B1LDX	8G	8G	N/A	8G	8G	8G

Table 14: 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.

Conclusion

EANTC verified the interoperability of Huawei OptiX OSN 1800 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 B1ELOM, B1LDCA and B1LDX line card.

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

Based on the results of our test, EANTC confirms that the Huawei OptiX OSN 1800 fulfills Huawei's claims to work in enterprise data center interconnection scenarios as an integrated, high-speed, and resiliency solution.

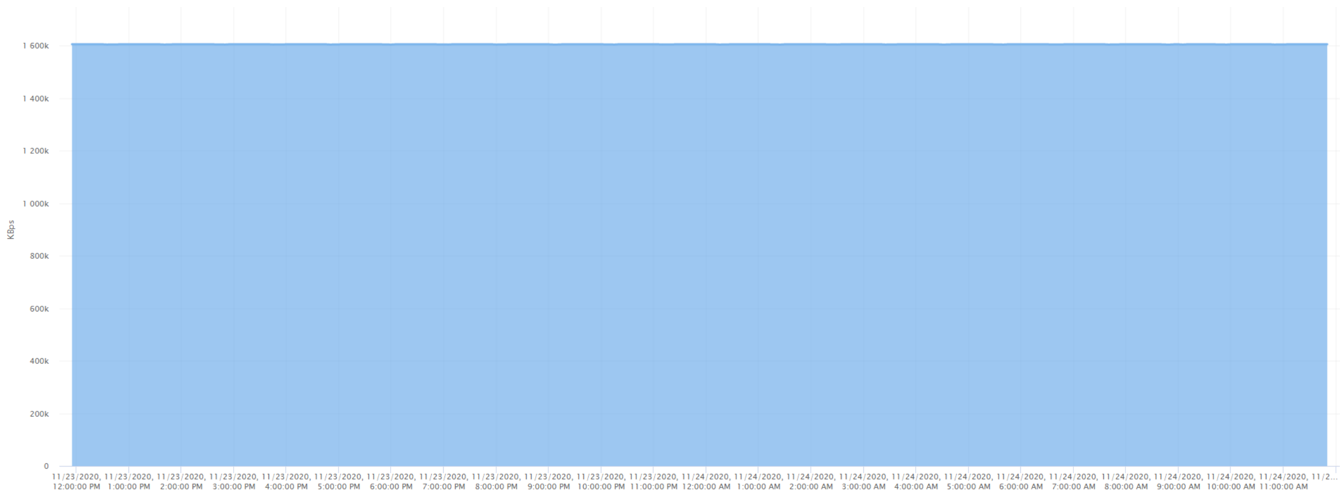

Performance Chart Legend

Key	Object	Measurement	Rollup	Units	Latest	Maximum	Minimum	Average
	Test_VM (1)	Usage	Average	KBps	1,606,445	1,606,516	1,605,362	1,606,308.4

Figure 8: 24-hour Throughput of 16G Interface on B1LDCA

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 20210412]