EANTC Independent Test Report

Huawei OceanStor Dorado V6 All-Flash Storage

October 2020

Introduction

Huawei commissioned EANTC to independently verify the interoperability of Huawei OceanStor Dorado V6 All-Flash Storage with EMC's VPLEX block-based storage virtualization solution for the Storage Area Network (SAN). EANTC confirmed key functional interoperability and resiliency aspects, validating that Huawei OceanStor Dorado V6 All-Flash Storage can be provisioned and operated efficiently and reliably through the EMC VPLEX solution.

Today's storage architectures are evolving with the growing importance of private clouds and virtualization in many enterprise networks. Enterprise IT departments are challenged to manage heterogeneous environments, simplify architectures and reduce operational expenses all at the same time. Storage virtualization benefits storage provisioning by simplifying the tasks of managing multiple storage arrays and heterogeneous environments in a SAN. In a virtualized storage environment, hosts can access the SAN in the form of virtual storage volumes.

Huawei OceanStor Dorado V6 All-Flash Storage is a new-generation all-flash storage system developed by Huawei for core enterprise services. The new hardware platform and SmartMatrix full-mesh architecture, Huawei OceanStor Dorado All-Flash Storage, delivers industry-leading performance and reliability. The product models include Huawei OceanStor Dorado 3000, 5000, 6000, 8000, 18000 V6 and furthers, which share a common set of front-end port types and software platform. We tested the Huawei OceanStor Dorado 3000 V6 this time. EMC VPLEX is a storage virtualization product developed by EMC Corporation. EMC VPLEX implements a distributed "virtualization" layer within and across geographically disparate Fibre Channel storage area networks and data centers. It logically is located between hosts and storage arrays, presenting itself to hosts as the storage provider (target) and presenting itself to storage arrays as one host (initiator). It simplifies storage management by providing a single image for multiple controllers and a consistent user interface for heterogeneous provisioning storage.

Executive Summary

We focused the evaluation on the agility of the SAN storage plug-in and highlighted the access to storage through multiple paths for redundancy. At EANTC, we created a testbed based on an agile cloud environment where applications are independent of the underlying infrastructure and can easily access storage resources independent from the underlying hardware.

Test Highlights

→ Huawei OceanStor Dorado 3000 V6 successfully enabled LUNs that were accessed on the SAN to support a block-based storage virtualization solution

→ Verified interoperability to storage virtualization solution, inclusive auto-discovery, claim, extents creation, and forming virtual volume by the EMC VPLEX

→ Proved successful hardware protection on SAN running storage virtualization, with none of IO drop to zero during link and module failures test

→ Confirmed vdbench kept running with less than 7 seconds of I/O decreasing to zero during storage controller failure test

We configured a block-based storage virtualization method. We implemented the storage virtualization with EMC VPLEX located between the storage array and the host. In this setup, VPLEX's role was to virtualize the storage array presentation to all workloads.

The EANTC team used Huawei OceanStor Dorado V6 All-Flash Storage in the testbed. We verified the integration of storage logical unit numbers (LUNs) without changing the configuration on the server. The server did not participate in active virtualization activities but acted as a bare-metal host running the CentOs 7 operating system. The VPLEX separates the virtualization layer by adding a new storage virtual controller unit in the SAN which retains communication with the storage components, such as Huawei OceanStor Dorado 3000 V6, through Fibre Channel (FC).

Test Case	Verdict
Interoperability of LUN discovery	Pass
Interoperability to support virtual provisioning operations	Pass
FC link protection	Pass
Module protection	Pass
Controller protection	Pass

Table 1: Test Cases Overview

Hardware and Software

Vendor	Device	Hardware Version	Software Version
Huawei	OceanStor Dorado 3000 V6	V6	6.0.1.SPH3
EMC	VPLEX	VS6	6.2.0.03.00.02
Brocade	G620	2.6.34.6	v8.0.1b
HP	Server	1500	CentOs 7
Open Source	Vdbench	N/A	5.04.07

Table 2: Hardware and Software Components

SAN Setup

We created a SAN with the physical topology as shown in Figure 1 using two Brocade G620 FC switches in FC fabrics, connected to a host, the Huawei OceanStor Dorado V6 All-Flash Storage and the EMC VPLEX. The logical topology consisted of two FC fabrics, as shown in Figure 2, the blue zones and the red zones in the FC fabrics respectively.

In each fabric (for example with blue zones) the EMC VPLEX back-end consisted of each four engines of director A and B forming a total of eight links connected to a FC switch providing the fabric. From there two links of the blue fabric led to Huawei OceanStor Dorado V6 All-Flash Storage. The same settings applied for the FC fabric with red zones. The EMC VPLEX front-end provided eight links to one of the FC switches which is connected via two links to the HP server.

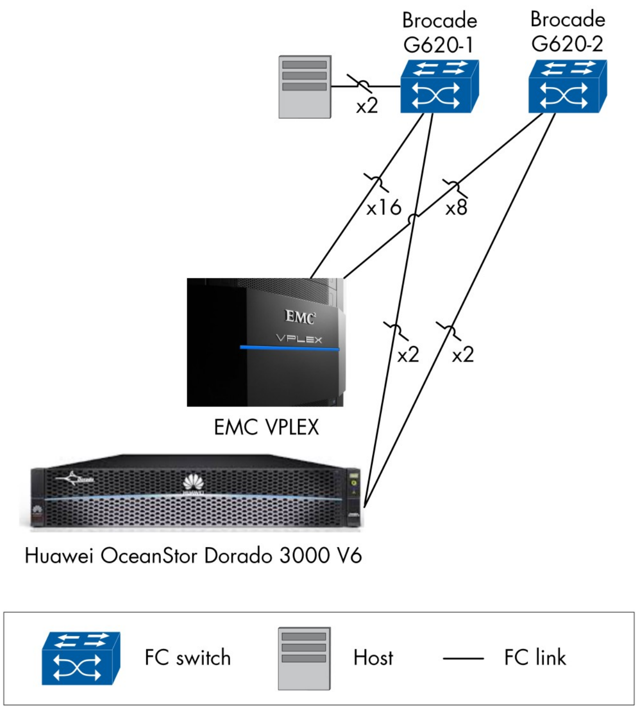

Figure 1: Physical Topology

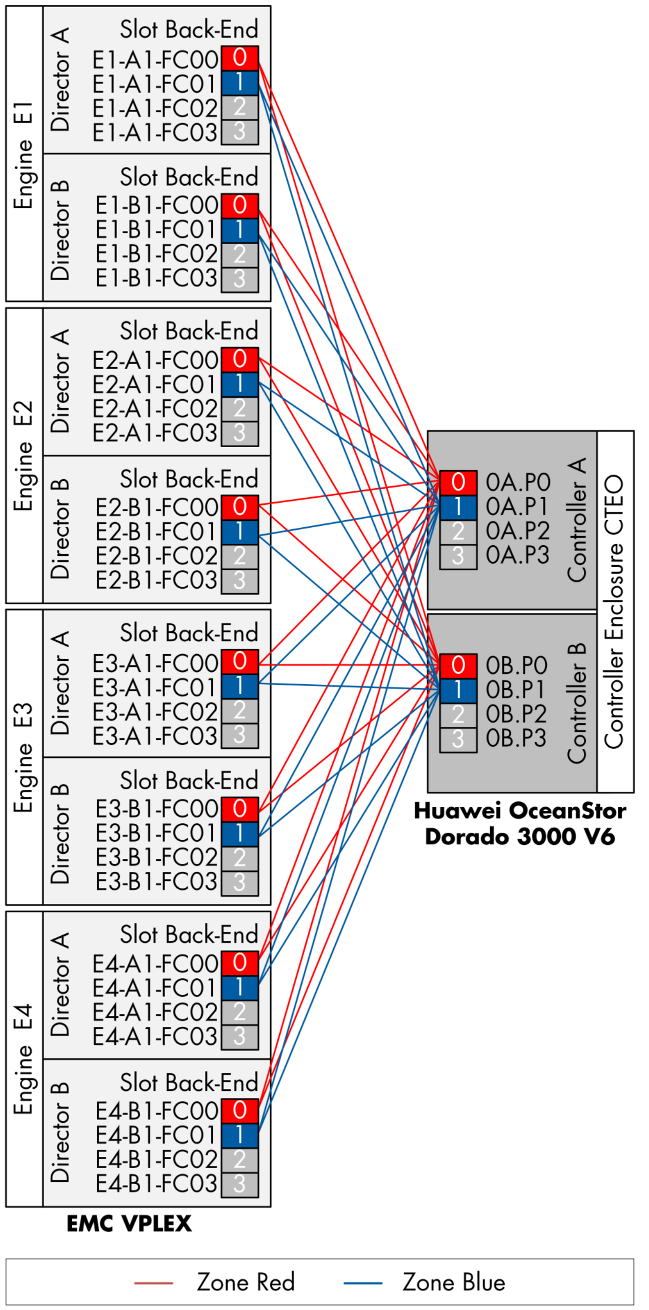

Figure 2: FC Zone Configuration between Huawei OceanStor Dorado V6 All-Flash Storage and EMC VPLEX

LUN Detection

Our goal first was to verify that the Huawei OceanStor Dorado V6 All-Flash Storage enabled LUNs that could be accessed on the SAN at where the EMC VPLEX was connected. In this case, a LUN defines the address of the device in the storage environment, which tells the system from which volume to send and read data when addressing. We observed the status on the EMC VPLEX to indicate completion and expected that the LUNs were accessed on the SAN.

Logged in on the Huawei OceanStor Dorado V6 All-Flash Storage web interface, we configured the first LUN through the GUI dialog on the page "Services". We allocated a size of 2TB to it and associated it to the initiators discovered in the target FC zone, the initiators identified the EMC VPLEX ports and were advertised via FC zone.

As mentioned above, LUN is an address in a SAN environment, so it was possible to have multiple LUNs on a single storage device. We created four LUNs each 2TB, and selected ID from below 3 options: 0, 4095 and ID > 0.

The EMC VPLEX web interface confirmed the status that the discovery of LUNs has been completed. We observed from the EMC VPLEX that its backend connected to the FC discovered all four LUNs (each with 2TB) from the SAN and included them into storage array as part of the backend. The exchange of the LUNs information succeeded via FC.

As expected, the Huawei OceanStor Dorado V6 All-Flash Storage enabled LUNs that could be accessed on the SAN. All LUNs were included by EMC VPLEX as part of the backend on the EMC VPLEX.

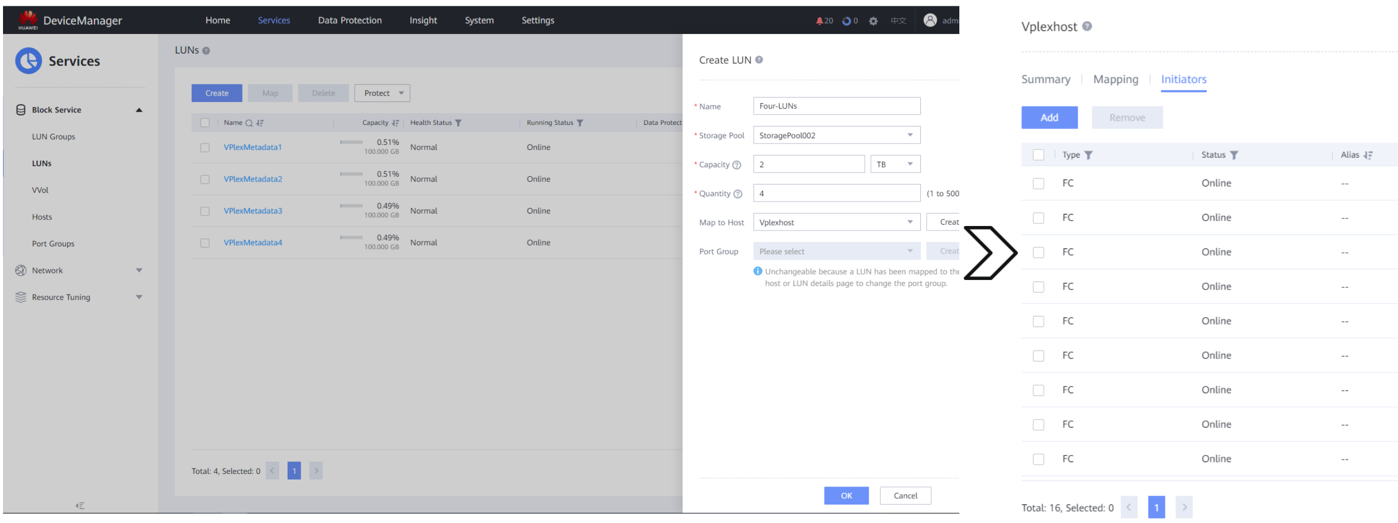

Figure 3: Huawei OceanStor Dorado V6 All-Flash Storage LUN Creation

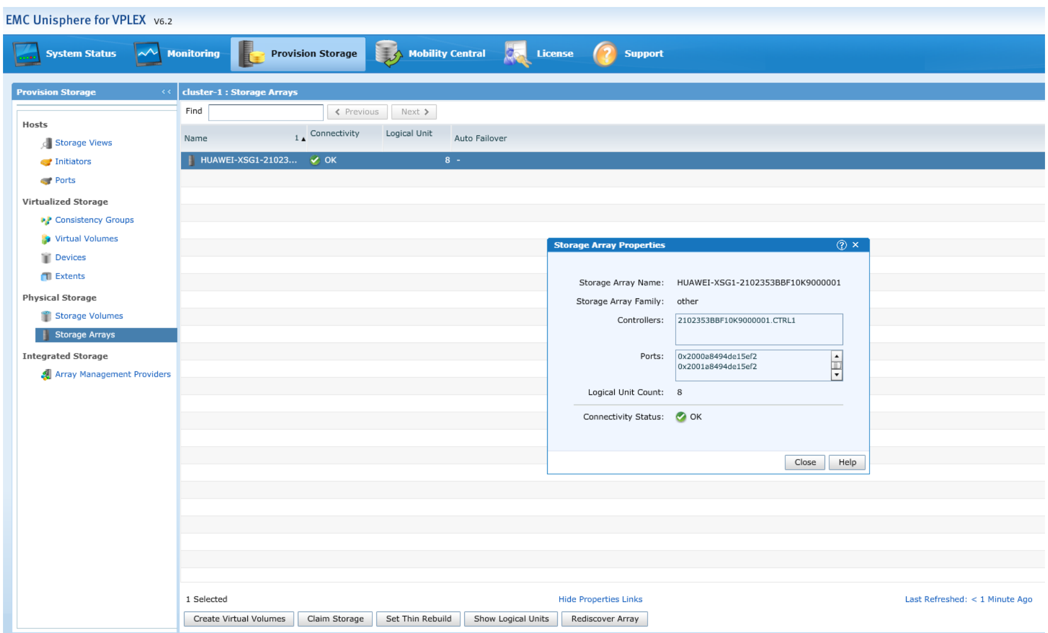

Figure 4: EMC VPLEX Discovery of LUNs in the SAN

Virtualization Provisioning

We verified that the LUNs of SAN successfully participated in virtualization provisioning. We observed the status of the host connected to the SAN and expected to provide a virtual volume to the host.

We used a third party open source tool vdbench on the host to create end-to-end traffic load. The traffic was designed to access the virtual volume, to successfully perform read and write operations to confirm that the storage virtualization will correctly import the traffic into the actual storage.

The core part of the virtualization provisioning included an EMC VPLEX configuration consists of claiming, adding extents to forming a device on the EMC VPLEX and finally presented by EMC VPLEX as a virtual volume and displayed to the host on the other end.

- By targeting and claiming the LUNs, the restriction access dictated that all systems from the zone were exposed to each other in our SAN. This permission allowed these LUNs to be accessed later in multiple paths.
- The LUNs allowed a free space manager by creating extents for the LUNs which have been claimed. Logical database space is called an extent, which is a specific number of contiguous data blocks that is allocated for storing a specific type of information.
- By allocating a striped volume (RAID 0) to the extents, which combined areas of free space from selected LUNs to a logical volume, to constitute a device. Data that is written to a striped volume is interleaved to all LUNs at the same time instead of sequentially.
- Finally, the created device went to views on the front end through a storage mapping and presented as a virtual volume to the host.

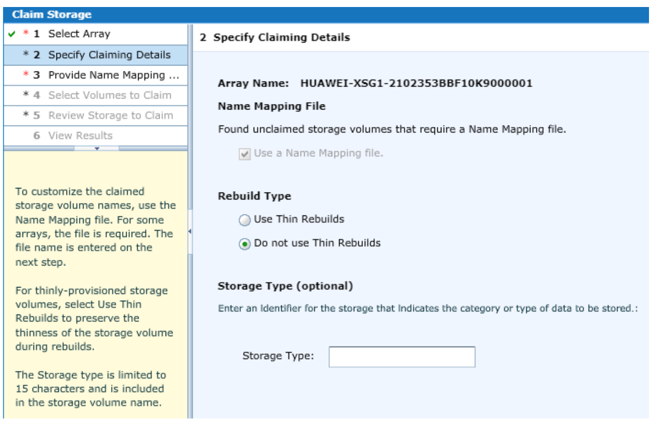

Figure 5: EMC VPLEX Virtualization Provisioning 1

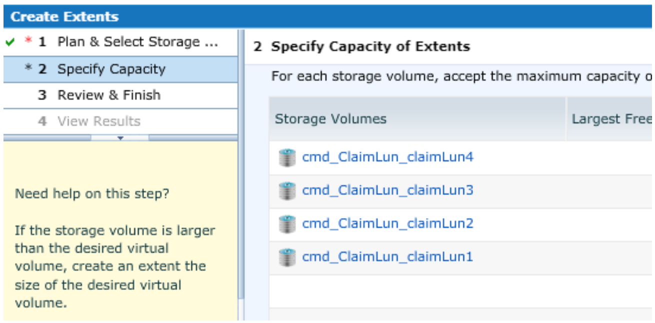

Figure 6: EMC VPLEX Virtualization Provisioning 2

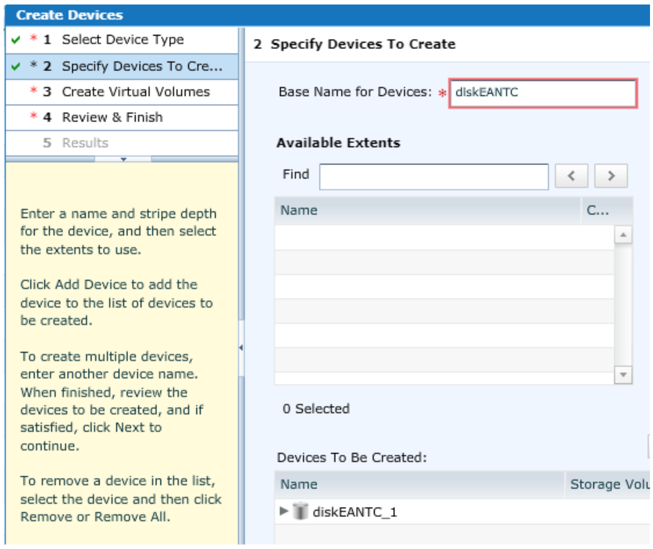

Figure 7: EMC VPLEX Virtualization Provisioning 3

As expected, we discovered on the CentOs host the target virtual volume after scanning the SAN. The multipath tool launched allowed the host to utilize all FC paths of the zone accessing the virtual volume. We generated with vdbench launched on the host a baseline work load (8Kbytes block size) consisting of mixed random read and write operations (7:3) accessing to the discovered virtual volume.

We successfully reached with vdbench an end-to-end block traffic at 14,000 IOPS (Input/Output Operations Per Second) in average, delivering at 110 Mbytes/s of block bandwidth.

```
[vdbench50407]# echo "- - -" > /sys/class/scsi_host/host1/scan
[vdbench50407]# echo "- - -" > /sys/class/scsi_host/host2/scan
[vdbench50407]# systemctl restart multipathd
[vdbench50407]# service multipathd reload
Redirecting to /bin/systemctl reload multipathd.service
[vdbench50407]# multipath -ll
mpathe (36000144000000010f005558e1f73539f) dm-0 EMC ,Invista
size=8.0T features='1 queue_if_no_path' hwhandler='0' wp=rw
`-+- policy='service-time 0' prio=1 status=active
|- 1:0:3:0 sde 8:64 active ready running
```

Figure 8: Virtual Volume Presented on the Host

We observed on the Huawei OceanStor Dorado V6 All-Flash Storage that all storage traffic passed equally through all four FC ports, the average count reached to 6,766 IOPS per port.

The total of 27,064 IOPS processed on storage includes an additional 13,064 IOPS, which was derived from the excess IOPS occupied by EMC VPLEX operations. We confirmed this utility by observing average of 14,000 IOPS on the front-end and average of 27,000 IOPS on the back-end on EMC VPLEX.

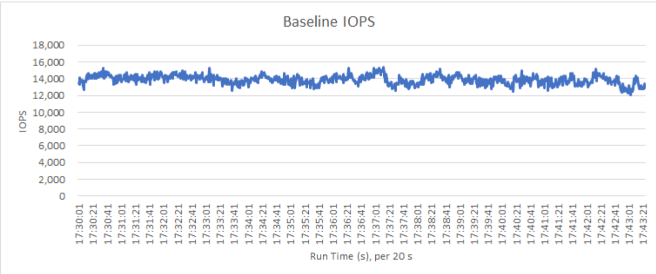

Figure 9: Vdbench Statistic Accessing Virtual Volume E on the Host

On this basis, we successfully verified the configuration has been completed.

Storage Protection

The test goal was to verify that the storage hardware was protected under the virtualization layer. We emulated three failure scenarios by performing administrative activities on the storage hardware and expected that the user traffic was protected from the failure, and the impact was minimal.

We used the same baseline traffic with vdbench as described above to perform redundancy tests. While traffic was running, we performed the failure action as shown in Table 3.

As expected, the vdbench kept running during all three failure tests. The lowest drop decreased to 0 IOPS in case of controller failure test, triggering a switch over to exit this state within 7s, from where the traffic increased back smoothly to the initial average of 14,000 IOPS within 70s. After this test once the traffic reached a steady state, it raised to 16,000 IOPS in average, including extra 2,000 IOPS added by vdbench to compensate for previous losses as additional load. The link and module failure did not have any impact on the traffic, the IOPS stayed within the tolerance value (+/- 2,000 IOPS) from the baseline (avg. 14,000 IOPS) as expected.

Scenario	Emulated Failure	DUT Action	Switch Over Time (s)
FC Link	Disconnect	Failover	None
	Reconnect	Failback	None
Module	Disconnect	Failover	None
	Reconnect	Failback	None
Controller	Disconnect	Failover	7
	Reconnect	Failback	None

Table 3: Switch Over Time (vdbench IOPS decreased to 0)

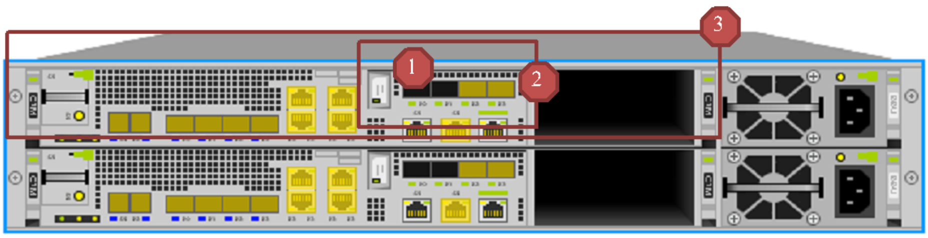

Figure 10: Failure Scenarios Performed on the Huawei OceanStor Dorado

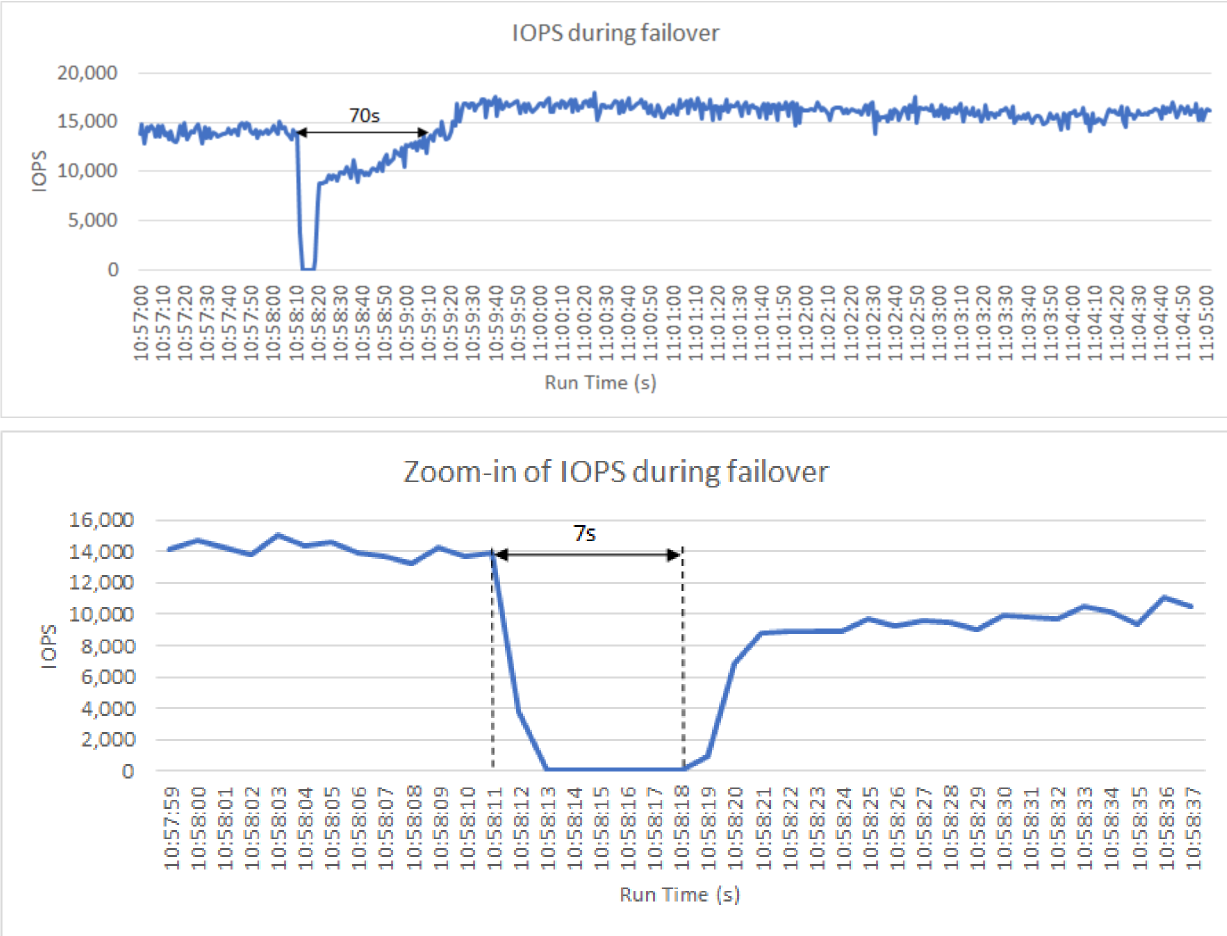

Figure 11: Vdbench Statistic during Controller Failover

Conclusion

The test results showed that the Huawei OceanStor Dorado V6 All-Flash Storage was discovered and virtualized by the EMC VPLEX, and provided successful protection against FC link, module, and controller failure.

About EANTC

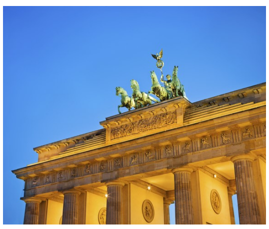

EANTC (European Advanced Networking Test Center) is internationally recognized as one of the world's leading independent test centers for telecommunication technologies. Based in Berlin, the company offers vendor-neutral consultancy and realistic, reproducible high-quality testing services since 1991. 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, http://www.eantc.de/
[v1.0 20201007]