

ESG Lab Review

Huawei OceanStor Dorado V3 All-flash Storage

Date: September 2017 Author: Tony Palmer, Senior IT Validation Analyst

Abstract


This ESG Lab Review documents hands-on testing of the Huawei OceanStor Dorado V3 all-flash storage and presents the findings of a five-year TCO analysis highlighting the economic benefits of Huawei OceanStor Dorado V3 when compared with hybrid and first-generation all-flash storage systems from major vendors.

The Challenges

Respondents to a recent ESG research survey were asked to name their top storage challenges and the top three most identified storage challenges for 2017—data protection, hardware costs, and rapid data growth rate—are the same challenges that have occupied the top three slots since 2015. The overarching issues that drive data storage concerns are also relatively unchanged—data growth is accelerating and the resulting infrastructure required to store and protect that data is costly and complex.¹

Figure 1. Top Storage Challenges

In general, what would you say are your organization’s biggest challenges in terms of its storage environment? (Percent of respondents, N=356, five responses accepted)



Source: Enterprise Strategy Group, 2017

¹ Source: ESG Brief, [2017 Storage Trends: Challenges and Spending](#), August 2017.

When we look at solid-state storage, we see that the most-cited factor driving consideration or deployment is, not surprisingly, improved performance (58% of respondents), with reliability (34%), cost per I/O (31%), and total cost of ownership (TCO) (30%) rounding out the top four responses.²


In ESG Lab’s experience, desktop and application virtualization are responsible for some of the most complex and demanding storage workloads in the data center. Organizations are tasked with providing a high-quality, predictable, and productive computing environment for an ever-growing number of internal users and external customers. In addition, enterprise application environments have become increasingly unpredictable as their underlying IT infrastructure grows in complexity and size. Mission-critical business application performance is highly sensitive to storage performance and latency, and highly dependent on the resilience of the enterprise IT environment.

The ability to consolidate mixed workloads and functions onto a single all-flash storage system has proven to provide significant TCO benefits if an organization’s performance, reliability, and operational requirements can be met. While many storage vendors offer all-flash solutions, the design decisions and tradeoffs made by these vendors can result in very different system capabilities and ultimately tradeoffs in benefits to an organization.

The Solution: Huawei OceanStor Dorado V3 All-flash Storage

Huawei designed the OceanStor Dorado V3 all-flash storage to handle mission-critical applications, both internal and customer-facing, for large enterprises, as well as mixed workloads. The OceanStor Dorado V3 leverages a dual controller, Non-Volatile Memory Express (NVMe) architecture to reduce the latency in accessing NVMe-based storage and ensure high availability. It can scale-out as the OceanStor Dorado V3 supports up to 16 controllers.

Figure 2. Huawei OceanStor Dorado V3 All-flash Storage



Source: Enterprise Strategy Group, 2017

Huawei offers several features to maximize OceanStor Dorado V3 performance, availability, and efficiency, while minimizing overall total cost of ownership (TCO):

Performance—Huawei OceanStor Dorado V3 all-flash storage system has an NVMe-based architecture, which supports direct communication between the CPU and NVMe SSDs. This eliminates the need for SCSI-SAS conversion and shortens the data transmission path, lowering end-to-end latency. The system also incorporates a disk controller collaboration algorithm,

² Source: ESG Brief, 2017 Flash Storage Trends, to be published.

which synchronizes the data layout between SSDs and controllers designed to provide performance at a consistently low latency, ensuring that mission-critical applications always operate smoothly.

Availability—Huawei employs multiple layers of software to provide high availability in its platform. Huawei employs RAID-TP, its implementation of triple-parity RAID that allows for up to three simultaneous disk failures. RAID-TP also enables fast rebuild of data, with hot spare space spread across all disk modules. HyperSnap snapshots provide point-in-time redirect-on-write (ROW) snapshots for fast recovery of data with minimal impact on performance. ROW snapshots are optimized for performance, requiring one-third of the I/O operations of copy-on-write snapshots with no computational overhead when reading snapshots. HyperReplication provides a traditional remote replication to a stand-by data center, and HyperMetro employs a gateway-free active-active deployment of two storage arrays that balances the load between them and permits non-disruptive cross-site takeover in case one array or one link between the array and host fails to ensure continuous data access for the most critical business applications.

Cost of Ownership—Huawei employs its implementation of SmartDedupe inline deduplication and SmartCompression to provide significant data reduction in the OceanStor Dorado V3. Huawei guarantees customers a minimum 3:1 data reduction ratio, offering to provide additional capacity to make up any difference. SmartThin provisioning enables organizations to provision only what they need today, growing capacity on demand, non-disruptively.

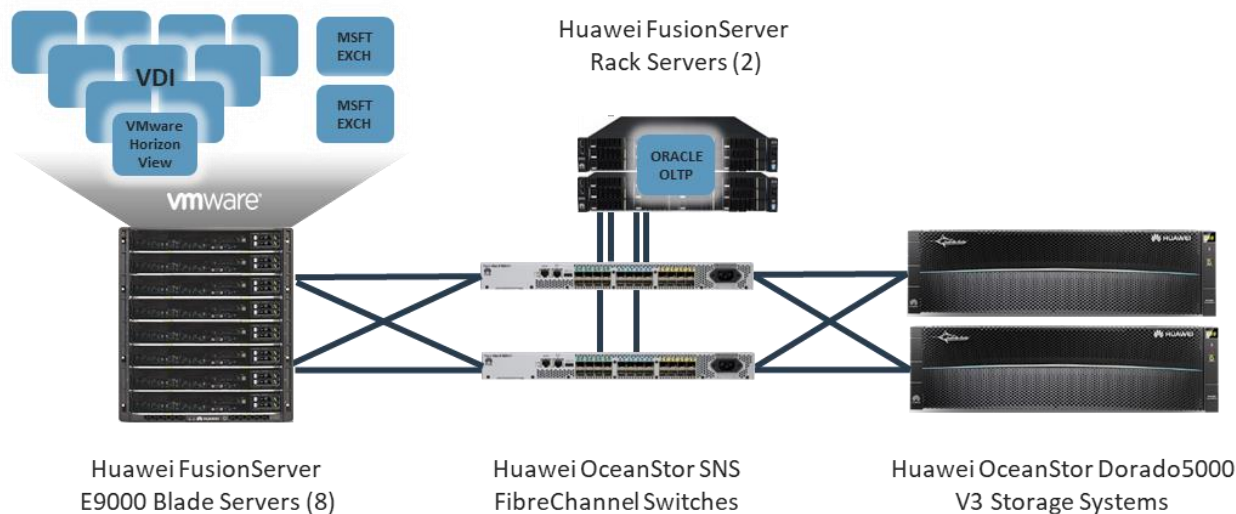
ESG Lab Tested

ESG Lab performed hands-on testing and validation of the Huawei OceanStor Dorado V3 All-flash Storage system at Huawei’s facilities in Chengdu, China. Testing was designed to validate the application consolidation capability offered by a single OceanStor Dorado5000 V3 storage system with a focus on delivering high levels of predictable performance for multiple, simultaneously running, tier-1 application workloads. The ability to sustain these high-performance levels through various storage hardware failures was also tested and a five year TCO analysis was performed.

Performance

The test bed utilized by ESG Lab is shown in Figure 3. A VMware vSphere 6.0 virtualized infrastructure was deployed on ten Huawei FusionServer RH Series Rack servers, each leveraging dual Intel Xeon 2690 v4 processors with 256 GB of RAM. The servers were connected to redundant Huawei OceanStor SNS FibreChannel Switches via 16GFC. The switches were then connected via two pairs of 16GFC cables to two Huawei OceanStor Dorado5000 V3 storage systems with dual controllers and dual smartIO cards per controller, each populated with 25 2TB NVMe SSD modules.

Figure 3. The ESG Lab Test Bed



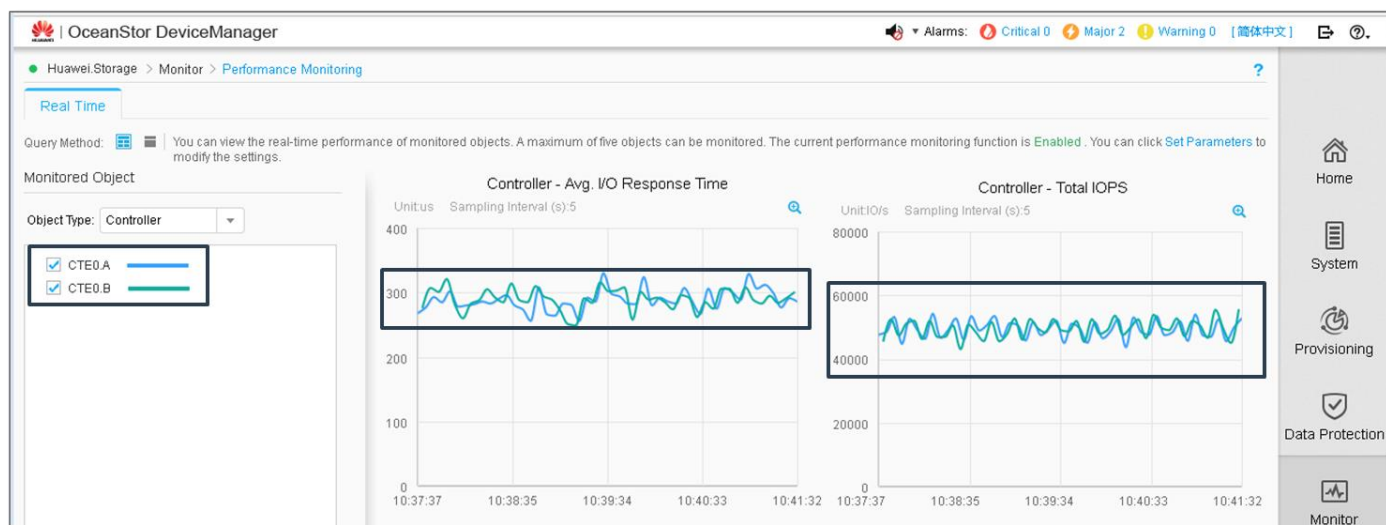
Source: Enterprise Strategy Group, 2017

Two common, mission-critical applications were deployed within the VMware virtualized infrastructure: a 1,000-seat virtual desktop environment and a Microsoft Exchange email deployment on two virtual machines. An Oracle 12c RAC environment was deployed on two physical servers. The VDI environment was managed and controlled by VMware Horizon 7, the virtual desktop host platform for VMware vSphere. One thousand persistent virtual desktops were created from a Horizon desktop image template, which leveraged a 50GB base image. Each desktop was configured with Windows 7 Enterprise edition (64-bit) and utilized one vCPU, 2 GB of RAM, one vmxnet3 vNIC adapter, and one LSI Logic virtual storage adapter. Installed applications included Microsoft Office 2012, Adobe Reader 11, Flash Player Active X, Internet Explorer, 7-Zip, and Windows Media Player. Workload was generated using VMware View Planner.

For email, testing leveraged a Microsoft Exchange Server 2013 environment that consisted of 5,000 mailboxes with a size of 1 GB each. The email database was configured on a 5TB LUN and the log file that was set to truncate. Workload was generated using the Jetstress utility. Finally, a 9.6TB Oracle OLTP database was configured to support up to 256 concurrent users. Workload was generated using SLOB version 2.4.2.

ESG Lab leveraged VMware vCenter and the Huawei OceanStor DeviceManager interface (shown in Figure 4) to manage and monitor the deployed applications. Test results were verified using output files and logs from the tested applications. First, ESG Lab started the OLTP and email workloads. At this point, the system was servicing about 105,000 sustained IOPS at an average response time of just 300 μ s.

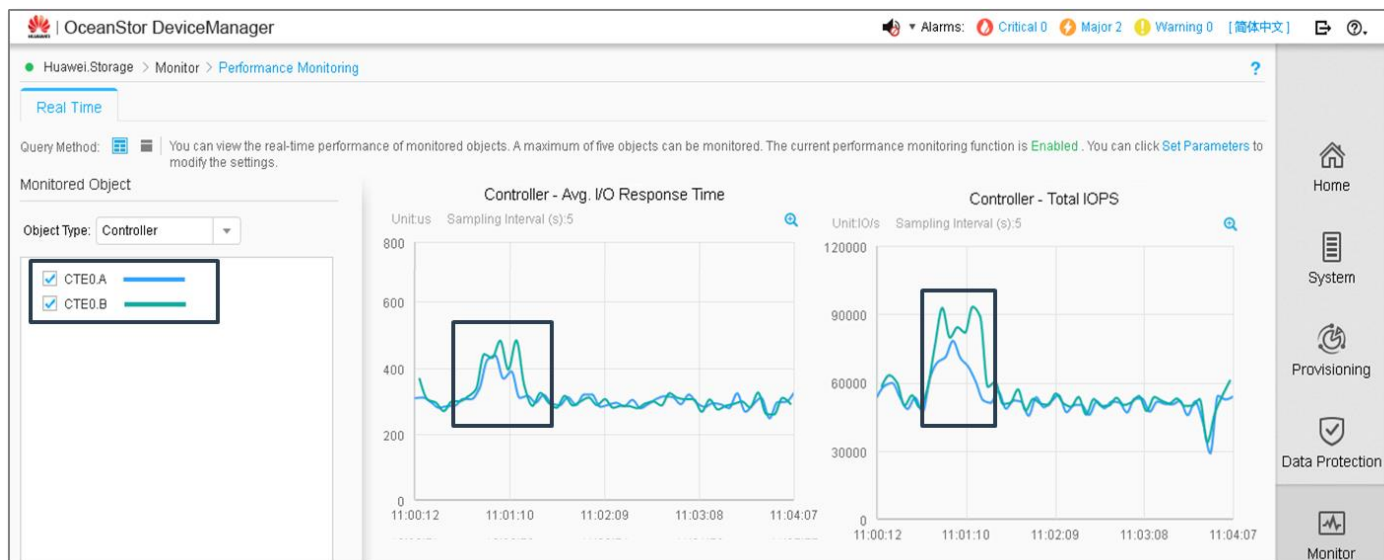
Figure 4. Test Bed Monitoring—Huawei OceanStor DeviceManager



Source: Enterprise Strategy Group, 2017

Next, the 1,000 virtual desktops were powered on to simulate a boot storm. As seen in Figure 5, I/O surged to approximately 160,000 IOPS for about a minute, and response time peaked at just under 500 μ s, then both settled back to the levels observed before the boot storm was initiated.


Figure 5. OLTP, Email, and VMware Horizon Boot Storm



Source: Enterprise Strategy Group, 2017

Next, we ramped up the steady state VDI workload using View Planner. This added approximately 20,000 IOPS to the system. As seen in Figure 6, the OceanStor Dorado was now servicing approximately 125,000 IOPS aggregate with an average response time of 320 μ s.

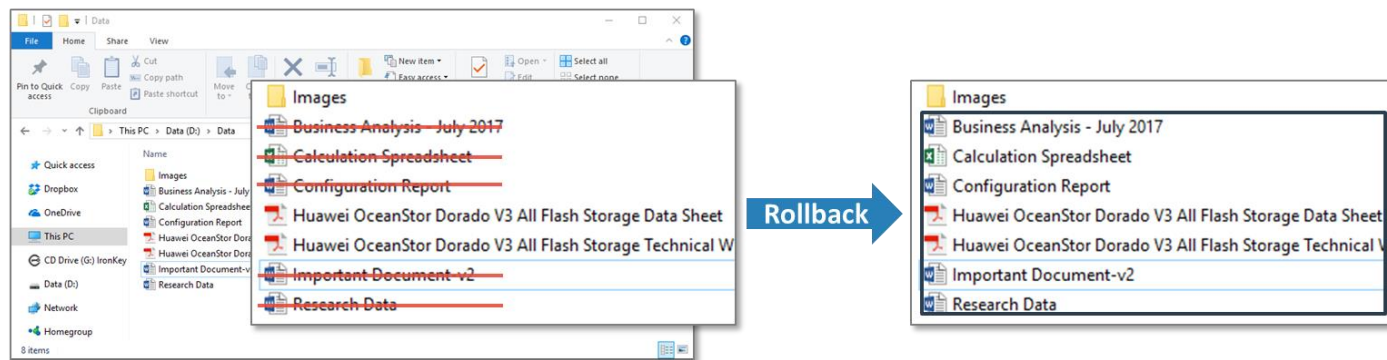
Figure 6. OLTP, Email, and VMware Horizon Steady State



Source: Enterprise Strategy Group, 2017

ESG Lab also tested the performance impact of snapshot functionality, creating snapshots of all LUNs while the system continued to run the mixed application workload. Snapshots were enabled and created instantly with just a few clicks. Snapshots can also be scheduled using the Huawei GUI and application consistency groups can be created. Once the snapshots were created, we deleted multiple files on one Windows system. Rolling the volume back to the snapshot was also quite simple, and the entire process had no measurable impact on the performance of the system.

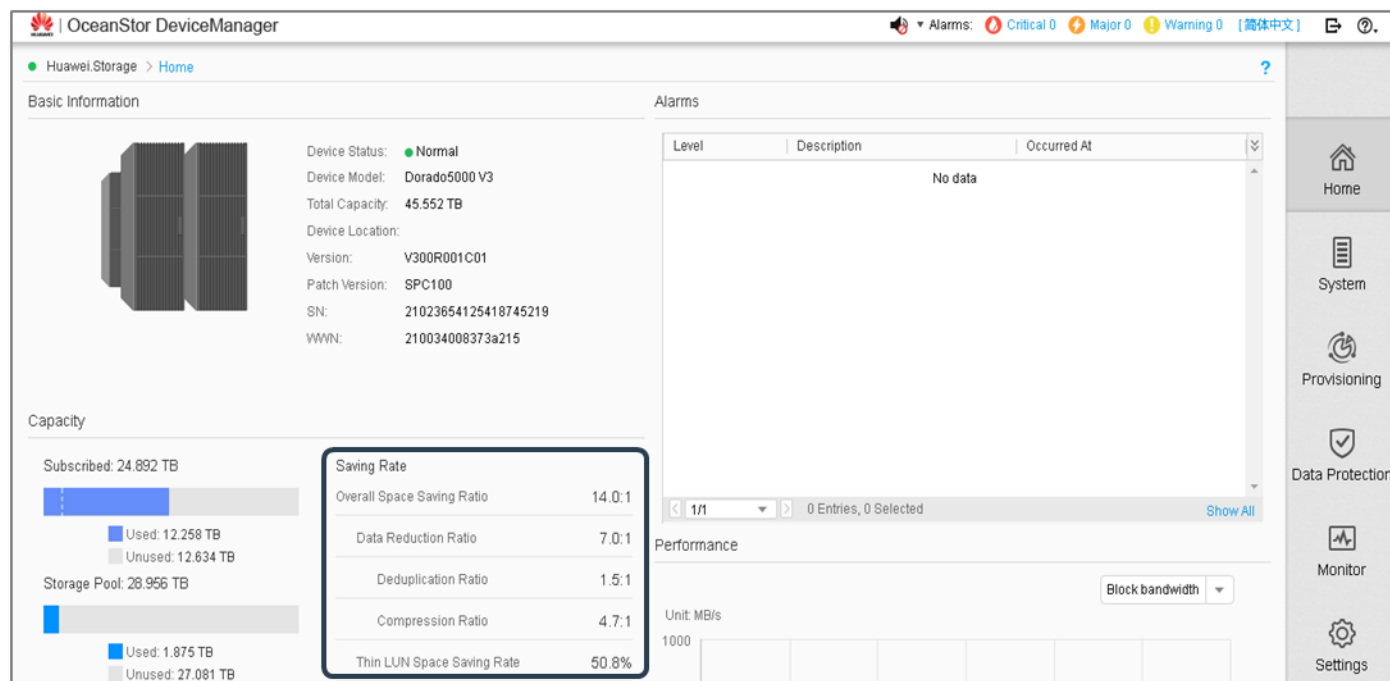
Figure 7. Snapshot Rollback



Source: Enterprise Strategy Group, 2017

Finally, ESG Lab examined data reduction achieved during testing, including data deduplication and compression. Because deduplication and compression are always enabled, all the performance results obtained for this report were achieved with inline data reduction on.

Figure 8. Huawei OceanStor Always-on Data Reduction



Source: Enterprise Strategy Group, 2017

As seen in Figure 8, the system achieved a data reduction ratio of 7:1 including compression and deduplication. The overall space savings including thin provisioning was 14:1.



Why This Matters

Consolidating workloads driven by physical and virtualized systems onto a single storage platform can help drive higher levels of infrastructure efficiency through improved resource utilization, but when multiple applications share the same underlying storage system, problems can quickly arise. A burst of I/O activity from one application (e.g., a virtual desktop boot storm or recompose operation) can significantly impact all the other applications, leading to poor response times, lost productivity, and, in the worst case, lost revenue.

ESG Lab validated that three mission-critical application workloads were easily consolidated onto a single Huawei OceanStor Dorado V3 NVMe storage system without impacting one another. As the simulated real-world workloads ramped up, the response time of a demanding VDI infrastructure that supported 1,000 heavy users remained low. Specifically, for consolidated, mixed workload virtual environments, the variety of I/O types and sizes can wreak havoc on the response time of each application, which is arguably the most important performance metric to pay attention to in these types of environments. ESG Lab confirmed that the OceanStor Dorado V3 sustained more than 125,000 IOPS with an average response time of 320 μ s across all applications throughout all phases of testing.

Availability

The Huawei OceanStor Dorado V3 all-flash storage platform is designed to ensure high availability and sustain performance through planned maintenance and unplanned outages. OceanStor Dorado V3 array can house from two to 16 controllers running in active/active mode. HyperMetro provides a gateway-free active/active high-availability solution between OceanStor Dorado V3 systems, either in the same data center or different data centers that are up to 100km apart. HyperMetro maintains data consistency between the storage arrays in two ways: If one array fails, HyperMetro will switch the failed array's workload to the redundant array immediately; If a link between a host and an array fails, HyperMetro will direct the host to the array which continues to provide data access.

Huawei's implementation of RAID-TP can tolerate up to three simultaneous disk failures. With Huawei's implementation of RAID-TP, parity data and hot spare space is spread across all disks in the storage pool, decreasing reconstruction time and increasing overall storage availability to better ensure continuous data access.

ESG Lab first tested HyperMetro in an on-campus setting, with two OceanStor Dorado V3s configured as a HyperMetro pair in different areas in the same data center. We began by running an OLTP workload to simulate transactional database traffic on an Oracle 12c Real Application Cluster (RAC). Each server was dual attached to the SAN and zoned to have access to both controllers on each array. Figure 9 shows the IOPS and response times of the arrays as the workload was running. The total workload generated by the two servers was 50,000 IOPS. Each controller in each array is processing approximately 12,500 IOPS with each array processing a total of 25,000 IOPS. Average I/O response time was 250 μ s across all four controllers.

Figure 9. Performance of Arrays 1 and 2 in Active/Active Mode



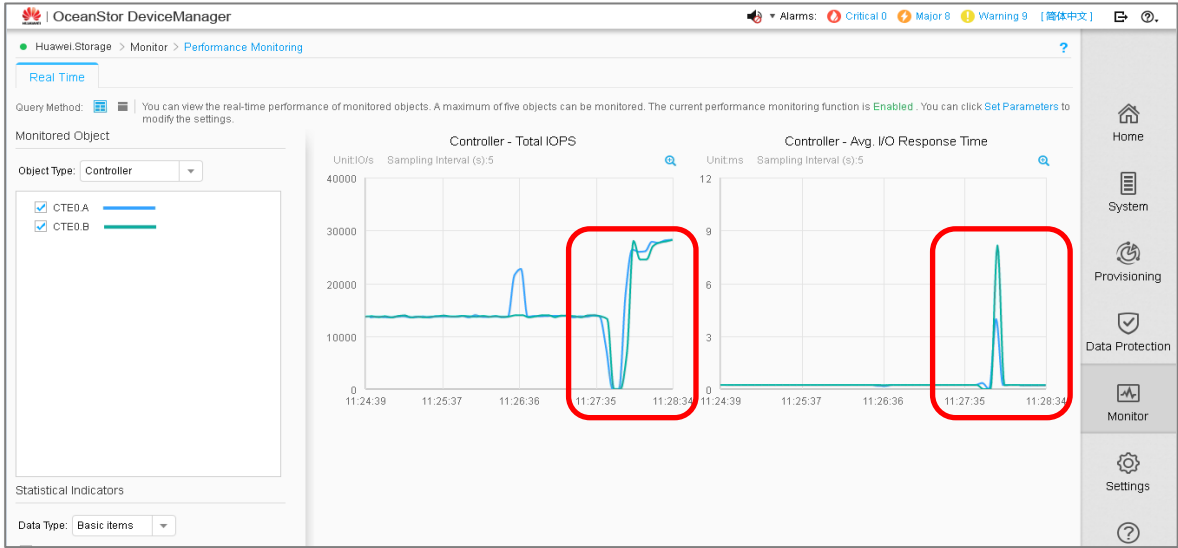
Huawei OceanStor
Dorado V3-Array 1

Huawei OceanStor
Dorado V3-Array 2

Source: Enterprise Strategy Group, 2017

Next, we simulated a failure of Array 1 by pulling the power cord. Figure 10 shows how the storage responded after HyperMetro switched the workload onto Array 2.

Figure 10. Performance of Huawei OceanStor Dorado5000 V3 After Switchover via HyperMetro

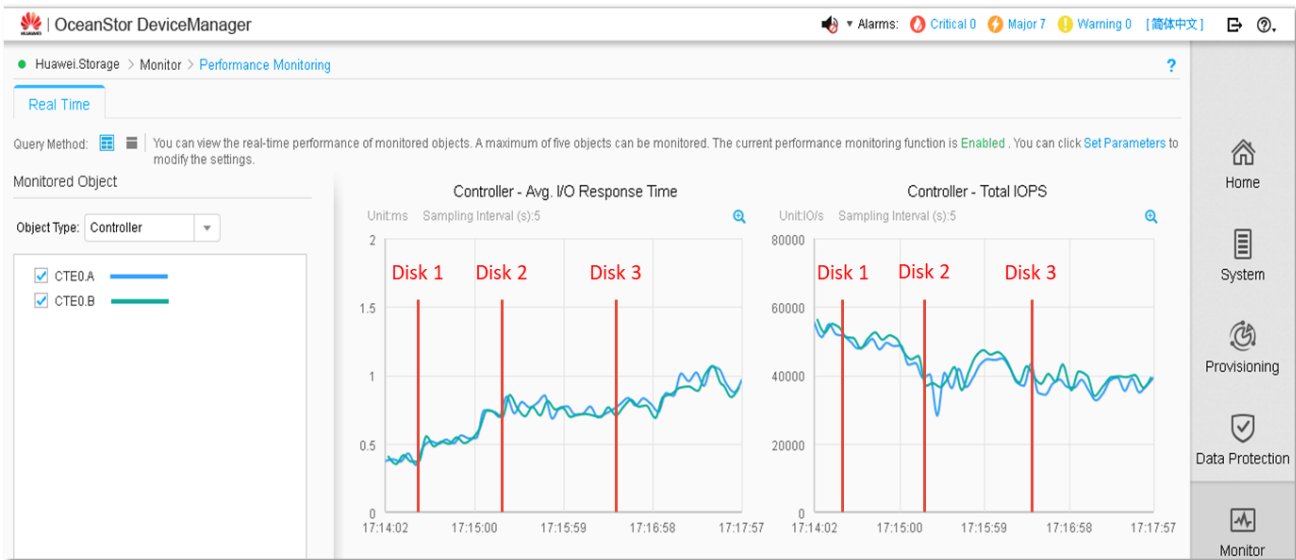


Source: Enterprise Strategy Group, 2017

The left chart in Figure 10 shows that the total IOPS in Array 2 drops to zero for approximately three seconds then resumes, servicing all requests from the cluster, for a total of 50,000 IOPS. The right chart shows that latency jumped to a maximum of 4ms and 8ms on Array 2’s controllers during the switchover then dropped back to an average of 250 μs.

With the OLTP workload still running, ESG Lab pulled three disks from the array at intervals of one to two minutes to simulate a cascading triple-disk failure. Figure 11 shows the performance of the array as the disks are removed.

Figure 11. Performance with Three Failed Disks

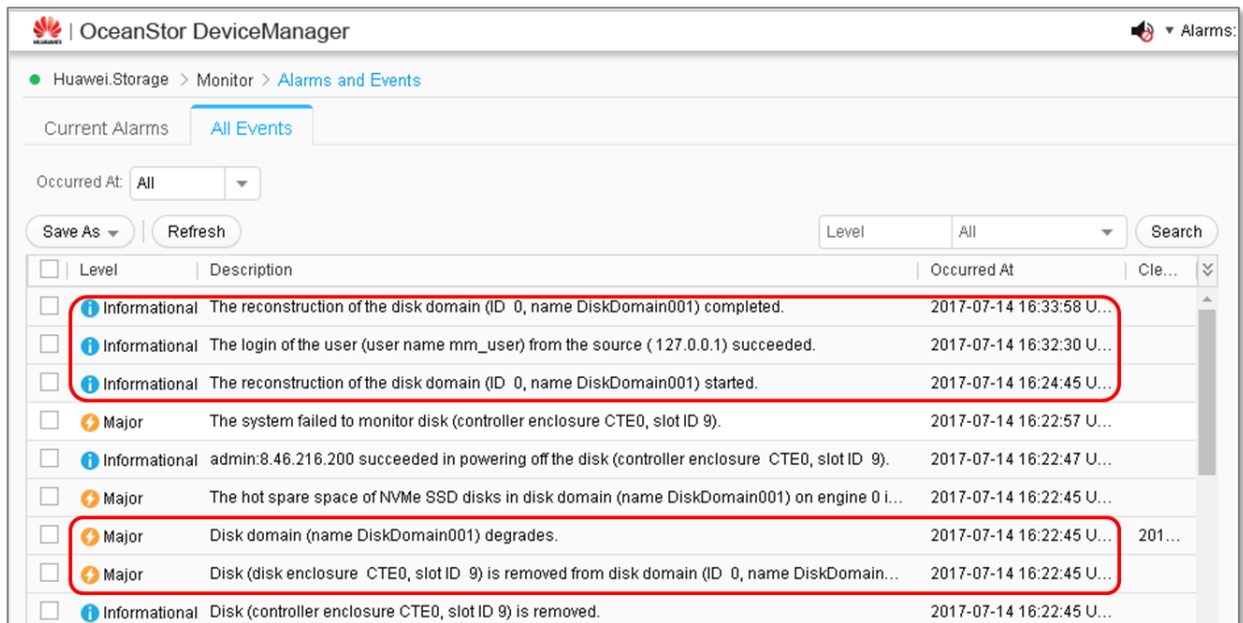


Source: Enterprise Strategy Group, 2017

As we pulled each disk, IOPS gradually decreased and response time gradually increased as the array had to reconstitute data from parity on the fly. The array remained available throughout the triple disk failure without a significant degradation in performance. Total IOPS decreased from a total average of approximately 100,000 to approximately 80,000 while response time increased from the previously observed average of 250 μ s to an average of 850 μ s.

Next, ESG Lab looked at RAID-TP rebuild times. With increasing dataset sizes, fast rebuild time after a disk failure or replacement is critical. We pulled a disk from an array populated with nine disks that hosted a 7.4TB volume. Figure 12 shows the actions performed by the array to rebuild the data from the failed disk.

Figure 12. Steps in the Disk Rebuild Process



Source: Enterprise Strategy Group, 2017

When the array detects the removed disk, the reconstruction process begins. As Figure 12 shows, the rebuild of the 7.4TB volume took just nine minutes to complete.

Why This Matters

Customers considering solid-state storage deployment are looking not just for raw performance, but to increase availability so that users have continuous data access for mission-critical business applications. If the underlying storage is not highly available, the risk of downtime and, subsequently, lost productivity and revenue become real. In today's business climate, IT professionals must ensure business continuity for data-intensive applications, from traditional CRM to real-time analytics and online transactions, to help their users respond to customer needs quickly.

ESG Lab validated that the fully redundant architecture of the OceanStor Dorado V3 is highly available and provides extremely low-latency performance during unplanned outages. First, we verified that the active/active controllers of the Huawei OceanStor Dorado V3 provide consistent, evenly balanced performance across all controllers in an array. ESG Lab was particularly impressed with the ability of the system to sustain high levels of performance with sub-millisecond response times through multiple disk failures, during fast RAID rebuilds, and the failure of an entire array using HyperMetro.

Total Cost of Ownership (TCO)

ESG Lab modeled and compared the storage-related costs that could be expected when deploying traditional hybrid storage with SAS-based SSD and disk, first-generation all-flash arrays (AFAs), and a Huawei OceanStor Dorado5000 V3 with NVMe SSDs. The costs associated with purchasing, maintaining, powering, and cooling the storage systems were calculated in U.S. Dollars and the average cost for electricity in the United States as reported by the U.S. Energy Information Administration³ was used to calculate power and cooling costs. ESG Lab modeled the expected storage total cost of ownership (TCO) for a company that needed to support a highly available mixed-workload environment with the same requirements as tested in this Lab Review. All workloads were assumed to require sub-millisecond response times:

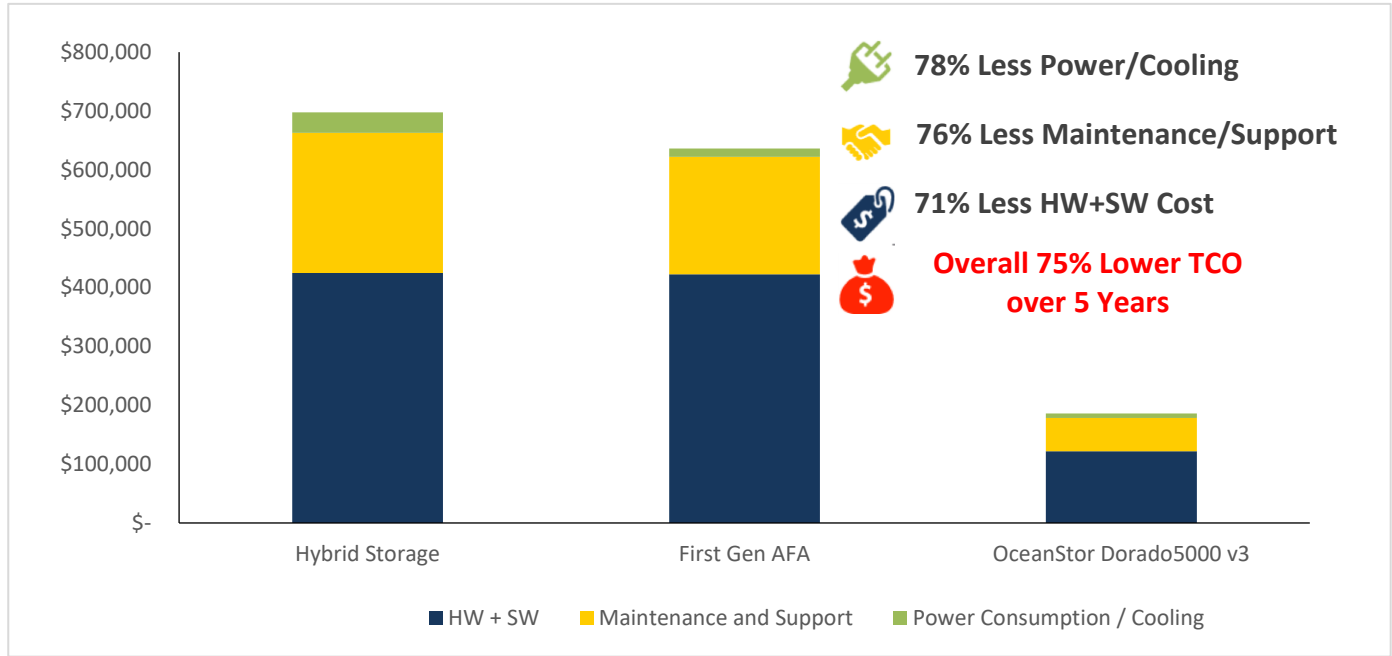
- An Oracle RAC OLTP environment able to support 100,000 IOPS and sub-millisecond response times.
- A 1,000-seat VDI deployment for heavy users (20 IOPS per user).
- A Microsoft Exchange environment to support 5,000 heavy users at one IOP per user.

The Huawei OceanStor Dorado5000 V3 was populated with 25 2TB NVMe SSDs configured in a single RAID-TP group. We compared that with a hybrid storage system from a major vendor populated with 12 900GB 2.5" SAS SSDs and 652 900GB 2.5" SAS HDDs configured in RAID5 groups as well as a first-gen AFA populated with SAS flash drives. The all-flash array costs were averaged from several systems from major manufacturers modeled by ESG to the same performance specification. All systems were modeled with equivalent software, plus power supplies, racks, and accessories according to each manufacturer's best practices.

TCO was calculated using a simplified model based on costs that would be incurred over a five-year period without taking into consideration capacity and performance growth requirements or IT operational costs. Maintenance and support contracts, along with typical customer discounts for hardware, software, and maintenance were factored into the estimated costs. Figure 13 shows the TCO cost comparison between hybrid storage, a first-gen AFA, and the Huawei OceanStor Dorado5000 V3.

³ <https://www.eia.gov/>

Figure 13. TCO Comparison between the OceanStor Dorado5000 V3, Hybrid Storage, and First-gen AFAs



Source: Enterprise Strategy Group, 2017

Over five years, TCO for the hybrid storage system totals \$697,568, the averaged costs for a first-gen AFA totals \$636,134, while the costs for the OceanStor Dorado5000 V3 is only \$186,137, just 27% of the hybrid storage cost and 29% of the cost of a first-gen AFA. Cost savings were similar across all three measured categories—hardware/software, maintenance/support, and power/cooling.

It’s worth noting that given the availability features of the Huawei OceanStor Dorado5000 V3, IT operational costs should be lower as well. Based on our research and experience, SSDs encounter fewer storage failures in the field, which translates into less time and resources spent replacing disks.

i Why This Matters

Organizations understand the advantages to be gained by selecting a next-generation purpose-built all-flash array designed to take advantage of the benefits of flash technology over simply adding flash drives to a traditional storage system with traditional limitations. While IT professionals seek high storage performance and availability, they still need to minimize both initial capital outlays and related operational costs over time. Solid-state storage has become a viable option as SSD prices have declined while their reliability has increased. Simultaneously, the costs for operating and maintaining a hybrid SSD+HDD environment tends to get higher over time, especially as HDD reliability degrades with continuous use.

ESG Lab compared the 5-year TCO of the OceanStor Dorado5000 V3, populated with NVMe SSDs, against that of hybrid storage system with a mixed SSD+HDD environment of SAS SSDs and HDDs and a typical first-gen AFA from major vendors. The results showed that the five-year TCO of the OceanStor Dorado V3 is 73% less than that of the hybrid array and 71% less than a first-gen AFA, with savings spread evenly across capital outlay for hardware/software, maintenance/support contracts, and power and cooling. ESG Lab also expects that the IT operational costs will decrease as well due to the increased reliability of SSDs and the availability features that Huawei has built into its all-flash array, allowing IT professionals to spend more time and resources on strategic activities, rather than maintaining storage.

The Bigger Truth

ESG research reveals that data protection, hardware costs, and rapid data growth rate are still identified by organizations as top storage challenges.⁴ With the ubiquitous use of server and desktop virtualization technologies to consolidate applications and users, the amount and variety of data that businesses need to store is growing rapidly, driving growth in overall storage use and costs. Another key objective for any IT administrator is providing sufficient performance to give business users the best possible experience. This is especially important for virtual desktop deployments and mission-critical applications, which are becoming increasingly virtualized.

Advancements in server and network resources occur on a regular cadence, but as more users and workloads are added to the infrastructure and leverage a shared pool of underlying storage, I/O bottlenecks can quickly become a concern. This is due in part to the increase in I/O traffic, and in part to the randomness of the I/O. As a result, IT is feeling more pressure to provide advanced solutions that can seamlessly scale capacity and performance and support continuous availability.

The OceanStor Dorado V3 all-flash storage system is designed to handle mission-critical applications and workloads, both internal and customer-facing. The OceanStor Dorado V3 leverages an active-active multi-controller, NVMe architecture and enterprise-class availability features implemented in software to provide a platform engineered for consolidating mixed workloads at extremely low latencies.

ESG Lab testing validated Dorado V3's ability to consolidate the most challenging business- and mission-critical workloads, including desktop virtualization, OLTP, and email onto a single, high-performance, highly available platform. The environment ESG Lab tested consolidated a realistic business environment with 12.25 TB of live data onto a single Dorado5000 V3 system, supporting 5,000 seats of Exchange users, 1,000 VMware Horizon virtual desktops, and hundreds of Oracle users, while consuming only 1.8 TB thanks to Huawei's data reduction technologies. The consolidated environment serviced these workloads with sub-millisecond response times averaging just 300 μ s and provided continuous access through multiple simulated failures. It's important to note that the performance described in this report was accomplished with inline data reduction and snapshots enabled and in use.

ESG Lab is pleased to validate that the Huawei OceanStor Dorado5000 V3 delivers consistently high performance at extremely low response times and is clearly well suited to support a mix of demanding real-world business applications running in a performance-critical highly virtualized environment. Next-generation all-flash storage systems are designed with a goal of making the best possible use of flash technology while avoiding many of the limiting factors of traditional storage systems. Huawei designed its all-flash array around the dual goal of solving business problems as well as storage problems.

It is no surprise that ESG Lab's five-year analysis demonstrated that by deploying an OceanStor Dorado5000 V3 rather than an alternative hybrid storage system, organizations can lower their storage TCO by 75% while improving availability and reducing operational effort. If your organization is looking to lower storage TCO while increasing performance, ESG recommends investing in a next-generation all-flash array, and Huawei is worth a closer look.

⁴ Source: ESG Brief, [2017 Storage Trends: Challenges and Spending](#), August 2017.

All trademark names are property of their respective companies. Information contained in this publication has been obtained by sources The Enterprise Strategy Group (ESG) considers to be reliable but is not warranted by ESG. This publication may contain opinions of ESG, which are subject to change. This publication is copyrighted by The Enterprise Strategy Group, Inc. Any reproduction or redistribution of this publication, in whole or in part, whether in hard-copy format, electronically, or otherwise to persons not authorized to receive it, without the express consent of The Enterprise Strategy Group, Inc., is in violation of U.S. copyright law and will be subject to an action for civil damages and, if applicable, criminal prosecution. Should you have any questions, please contact ESG Client Relations at 508.482.0188.

The goal of ESG Lab reports is to educate IT professionals about data center technology products for companies of all types and sizes. ESG Lab reports are not meant to replace the evaluation process that should be conducted before making purchasing decisions, but rather to provide insight into these emerging technologies. Our objective is to go over some of the more valuable feature/functions of products, show how they can be used to solve real customer problems and identify any areas needing improvement. ESG Lab's expert third-party perspective is based on our own hands-on testing as well as on interviews with customers who use these products in production environments.



© 2017 by The Enterprise Strategy Group, Inc. All Rights Reserved.



www.esg-global.com



contact@esg-global.com



P.508.482.0188