

Huawei OceanStor 18500 / 18800 V5 Storage Systems

The Huawei OceanStor 18000 V5 is a family of scale-out storage systems targeted at high-end enterprises with a common architecture, physical packaging, and embedded software. These include the 18500 and 18800 V5 hybrid models as well as the OceanStor Dorado V5 all-flash. The OceanStor line has offerings that target the high-end, midrange and entry-level segments. The OceanStor line now offers models hybrid and all-flash versions (denoted by the addition of Dorado to the name) with slightly different number of devices supported on all-flash systems. All models utilize modular architecture, standard racks and enclosures and support 2.5" SFF and 3.5" LFF device enclosures. The OceanStor 18000 V5 systems support block and file access.

All four 18000 models run the same base embedded software system and support additional add-on software features. The Huawei OceanStor high-end systems compete against Dell EMC PowerMax, NetApp FAS and AFF, IBM DS8000 and FlashSystem 9100, Hitachi VSP G5000, and HP 3PAR StoreServ 20000 systems.

CONTENTS

Highlights	3
Overview of System	4
Huawei OceanStor System	4
Model Comparison	5
Huawei OceanStor 18000 Hardware Architecture	5
Huawei OceanStor Controller Enclosure	6
Service Processor	7
RAID Level Support	7
Volume Configuration	8
Device Enclosures	8
Connectivity Options	9
FC Connectivity	9
iSCSI Connectivity	9
FCoE Connectivity	9
Dedicated Replication Ports	9
Huawei OceanStor Software Overview	10
OceanStor Operating System Storage Pooling and Storage Virtualization	11
RAID 2.0+ - Storage Pooling and Device Virtualization	11
SmartThin - Thin Provisioning	11
SmartVirtualization - External Storage System Virtualization	12
SmartMigration - Migration of LUNs	12
Performance-Focused Software Features	12
SmartCache - Caching data for improved hit ratios	13



SmartPartition – Allocating cache per application	13
SmartTier – Dynamic tiering of data	13
SmartQoS – Quality of Service	14
File Access – OceanStor File Engine	15
Protocols	15
NFS	15
CIFS/SMB	15
Data Protection	15
HyperSnap - Snapshot	15
HyperClone - Clone Copy	15
HyperCopy - LUN Copy	16
HyperReplication - Remote Replication	16
Consistency Groups	16
Integration with Application or System Software	17
OceanStor and VMware vSphere	17
VAAI Support	17
OceanStor and Microsoft System Center	18
Advanced Features	19
SmartMotion – Capacity and Performance Balancing	19
SmartDedupe & SmartCompression - Data Reduction	19
SmartErase – Digital Overwrite	19
SmartMulti-Tenant – Access Isolation	19
Reliability, Availability, Serviceability Features	20
Error Checking and Correction (ECC)	20
Data Scan	20
Redundant Power System	21
Fault Isolation	21
Sparing	21
Controller Redundancy	21
HyperMirror – Clustered High Availability	21
Concurrent Maintenance	22
Software Updates	22
Remote Support	22
Service Contracts	22
Management	22
Management Module	22
Maintenance Terminal	23
Application Server	23
Performance	24
Evaluator Group Comments	25

Strengths: _____ 25

Perceived Challenges: _____ 25

Highlights

- Scale-out controller architecture with up to sixteen controllers
- Internal PCIe switch for high performance interconnect between controllers
- Block and file access support
- SAS interface drives supporting 15K, 10K and SSD disk drives
- Flash Cache for extended DRAM caching using SSDs
- Remote copy (synchronous or asynchronous) and consistency group support
- Snapshot (point-in-time copy) capabilities supporting LUN snapshots
- Redirect-on-write - no capacity reservations for snapshots
- Thin provisioning and replication
- Dynamic storage tiering
- VMware VAAI support
- Third party storage virtualization for attached storage systems as external LUNs
- Migration between storage systems and thick to thin conversion
- QoS with I/O prioritization
- Cache partitioning under user control
- Dynamic load and capacity balancing with data movement
- Mirrored LUNs for stretched cluster support

Overview of System

The Huawei OceanStor high-end systems support block and file storage access over FC, iSCSI, InfiniBand and FCoE. In the scale-out configuration, up to sixteen controllers are supported with a high speed interconnect using PCIe switches. Two controllers are housed in a 6U enclosure. Hybrid models support both 2U twenty-five 2.5" SFF devices or 4U twenty-four 3.5" LFF devices. All-flash models support only the 2U twenty-five enclosure shelf.



Figure 1: Huawei OceanStor High-End Storage (Source: Huawei)

Huawei OceanStor System

The Huawei OceanStor system design utilizes a scale-out architecture, which provides many of the advantages of both monolithic and dual controller architectures by using standard components grouped into multi-node clusters. Each cluster of at least two controllers contains essentially three elements, as with dual controller systems, host connectivity, cache and processing and back-end disk connectivity. The clustered controller design allows the system to scale from two controller nodes up to a four controller configuration non-disruptively. Each node pair provides full redundant failover and data access capability, with each node being hot-pluggable for serviceability.

Model Comparison

Table 1 lists some of the specifications and characteristics of the OceanStor high-end systems.

Model / Capability	18500/Dorado 18500	18800/Dorado 18800
Processors (single controller)	2 @ 14 or 20 cores	2 @ 20 or 28 cores
Max Controllers	16	16
No. of Devices	6,400 – 2.5"	9,600 – 2.5"
	3,072 – 3.5"	4,608 – 2.5"
Max Capacity - raw	98 PB	147 PB
Devices Supported	SSD: 960 GB, 1.92, 3.84, 7.68, 15.36 TB	
	HDD: 10K rpm 600 GB 2.5"	
	HDD: 15K rpm 300, 600 GB 3.5"	
Cache	HDD: 7200 rpm 2, 3, 4, 6, 8, 10 TB 3.5"	
	16 TB	16 TB

Table 1: Huawei OceanStor Overview

Huawei OceanStor 18000 Hardware Architecture

In the architectural diagram shown in the following diagram, each controller node in the storage system runs its own operating system, each of which communicates with the others for redundancy and data movement. The high bandwidth PCIe interconnects enables all of the controllers (up to sixteen) to function as a single storage subsystem providing linear performance scaling.



Figure 2: Huawei OceanStor 18000 Switch Architecture (Source: Huawei)

The components of the Huawei system include:

- Four controllers are housed in a single 6U controller enclosure termed an “Engine”
- 2U or 4U storage enclosures
- OceanStor PCIe switch to interconnect controllers
- OceanStor Operating System
- OceanStor Device Manager software

Huawei OceanStor Controller Enclosure

The 6U controller enclosure is called an engine and contains:

- Two controllers in canisters
- Four battery backup units to provide orderly shutdown in case of the loss of power
- Four power modules to convert line voltage for controllers
- Interface module of interface cards external host and device attachment.






Figure 3: Huawei 6U Controller Enclosure Front and Rear (Source: Huawei)

The clustered controller architecture provides scalability, high availability and continued performance during a controller failure or planned downtime. Bringing down a single controller does not impact caching or data access in four controller systems.

Evaluator Group Comments: Scale-out designs offer greater availability and less performance impact when one controller fails than a traditional dual controller system. Huawei has bundled multiple controllers in the same physical enclosure but continued the isolation for servicing.

Service Processor

The Service Processor (SVP) is a separate server that is used for managing, configuring, and maintain the system. The enclosure for the SVP is 2U and contains the motherboard, two power modules, and service disks. Management is through network ports into the SVP. Keyboard, video, USB, and serial ports are also available.

RAID Level Support

The Huawei OceanStor systems support multiple RAID levels, including RAID 0 (striping), RAID 3 (parallel transfer), RAID 10 (striped mirror), RAID 5 (striped parity group), RAID 50 (striped rotating parity) and RAID 6 (dual parity).

Evaluator Group Comments: *An important distinction is how RAID levels are constructed and utilized. Huawei has a storage pooling implementation that allows allocation of chunks from devices with dynamic allocation and RAID protection on allocated chunks.*

Volume Configuration

The system software divides physical devices into logical chunks of up to 64KB from which RAID groups are created, which may then be allocated as a volume or LUN or a filesystem. As a result, data may be spread across more physical devices than RAID groups comprised of entire disk drives. The chunk size can be specified on a per LUN or per filesystem basis by the administrator to match the application usage in allocation.

Device Enclosures

Hybrid models support both 2U twenty-five 2.5" SFF devices or 4U twenty-four 3.5" LFF devices. All-flash models support only the 2U twenty-five enclosure shelf.




Figure 4: Huawei Device Enclosures (2U 25 – 2.5" and 4U 24 – 3.5" shown)

Device enclosures include redundant and hot-pluggable components. Each drive chassis includes redundant power supplies, redundant fan modules, and redundant SAS connectors that provide two, independent, 12 Gb/s connections to the main controller chassis. The SAS ports are direct connected

using a 4-lane SAS cable. All device tray components including power supplies, SAS adapters, fans and devices are serviceable online and are hot pluggable.

Connectivity Options

OceanStor systems provide Fibre Channel, Fibre Channel over Ethernet (FCoE), InfiniBand, and iSCSI connectivity options. The total number of connections for a single controller is 20 host ports so a four controller system would be a maximum of 80. Storage Virtualization of third party storage systems uses Fibre Channel ports for attachment and reduces the number for host attachment.

Interface	Interfaces for 18500	Interfaces for 18800
16 Gb/s Fibre Channel	384	384
8 Gb/s Fibre Channel	384	384
10 Gb/s Ethernet iSCSI or FCoE	384	384
10 Gb/s Ethernet	384	384
56 Gb/s InfiniBand	384	384
Backend Ports 12 Gb/s SAS	32	64

Table 2: System Connectivity

FC Connectivity

The maximum number of Fibre Channel connected servers is 65,536. Each port supports 16 Gb/s FC-SW connectivity.

iSCSI Connectivity

Up to 65,536 host iSCSI ports are available.

FCoE Connectivity

Up to 65,536 host FCoE ports are available.

Dedicated Replication Ports

Multiple dedicated ports are available for remote replication between systems, supporting both FC and iSCSI over Ethernet physical connections.

Huawei OceanStor Software Overview

The Huawei OceanStor operating system serves as the software platform for OceanStor systems, currently supporting high-end, mid-range, and entry-level hardware platforms with block and file access. The OceanStor operating system, along with associated value added software support a variety of data protection and management features.

Each controller in the storage system runs its own independent copy of the operating system. However, the OceanStor DeviceManager manages the cluster of controllers as a single system and provides all fault tolerance features utilizing built in hardware redundancy. The operating system, coupled with high speed PCIe hardware interconnects, allows the controllers to function as a single storage subsystem.

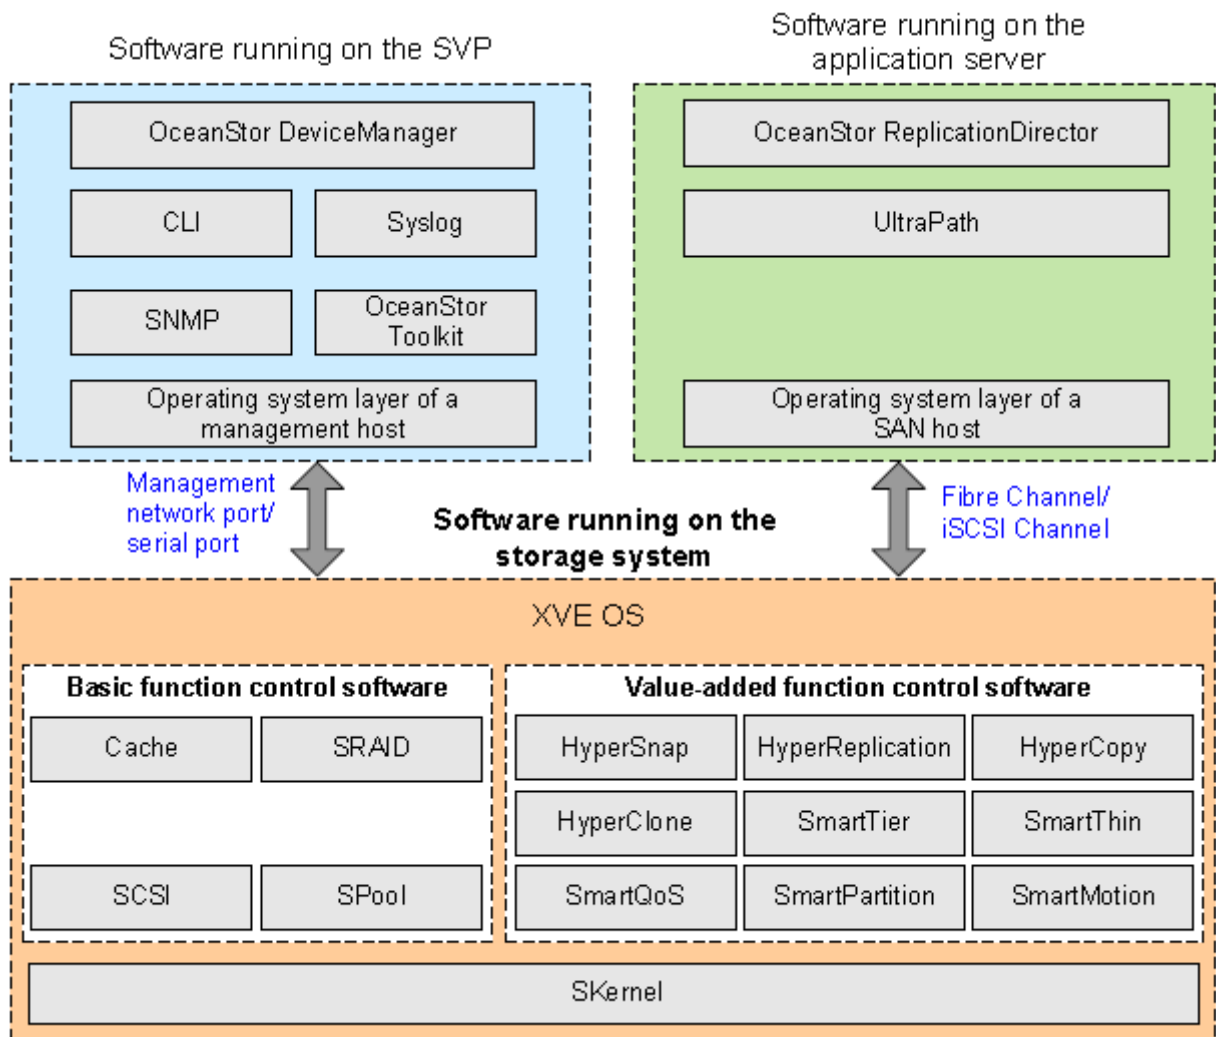


Figure 5: Huawei OceanStor Software Architecture (source: Huawei)

OceanStor Operating System Storage Pooling and Storage Virtualization

RAID 2.0+ - Storage Pooling and Device Virtualization

Physical devices are divided into logical chunks of up to 64KB from which RAID groups are created, which may then be allocated as a volume or LUN or filesystem. As a result, data is spread across more physical devices, with fewer limitations than with RAID groups comprised of entire disk drives where it is possible to perform wide striping, but only with large amounts of total storage. Huawei calls this level of storage pooling and virtualization RAID 2.0+.

A chunk is called a “grain” by Huawei and is user settable on a per filesystem or per LUN (thin-provisioned LUN) basis. The user setting can range from 4KB to 64KB is usually set based on the application usages such as file system allocations.

Storage Pooling	Support
Max LUN's / System	65,536
Max LUN's / port	255
Max number of storage pools	128
Max number of LUNs in a storage pool	8,192
Max LUN size	256 TB

Table 3: OceanStor LUN Support

SmartThin - Thin Provisioning

Volumes with capacity provisioned as needed are called thinly provisioned. Huawei uses the term SmartThin for the OceanStor implementation of thin provisioning. The RAID 2.0+ storage pooling implementation allows chunks to be allocation for capacity when needed for thin provisioned volumes. The allocation has a minimum size, which is 64KB regardless of the chunk size. Because the allocation of space is dynamic – when space is required – mapping tables are used for directing access to the blocks in a volume.

Space reclamation for allocated space that is no longer in use (deleted files or volumes for example) is supported with two methods. One is using the SCSI UNMAP command where the application, file system, or operating system notify the storage system to unmap specific blocks which are then returned to the storage pool as free space. The other method is for the OceanStor system to detect all-zero writes (zero-blocks) and unmaps those blocks from the mapping tables and return completion notification to the host.

SmartVirtualization - External Storage System Virtualization

In addition to virtualizing the underlying physical storage, OceanStor can also virtualize attached third-party storage systems, which is similar to other solutions such as the IBM Storwize V7000 and Hitachi VSP. Third party (or Huawei) storage systems are attached to Fibre Channel interfaces of the controllers, which become initiators rather than targets. This does reduce the number of available host connections.

An externally attached system is configured with LUNs managed by the OceanStor system. These LUNs are called eDevLUNs to distinguish them from local system LUNs. The maximum number of external LUNs supported is 1,024 with the 18000. Up to 256 external storage systems may be virtualized. There is a maximum of eight paths to each external LUN. External LUNs and local LUNs have the same properties so other features such as replication, tiering, snapshots, etc. all work the same.

SmartMigration - Migration of LUNs

LUNs can be migrated between physical devices without interruption of access. The LUN migration defines source and targets that can be local logical devices or external LUNs from a virtualized storage system. When the migration completes, all accesses will be transferred to the new physical location of the LUN automatically. The migration between virtualized storage systems and local storage can be for various reasons including draining data from the external system, converting thickly provisioned volumes to thinly provisioned volumes, applying data reduction to a LUN (compression and deduplication), and for performance quality of service reasons. Capacity balancing and load balancing also use SmartMigration as the means to move data between physical devices.

Reliability changes by migrating LUNs to different RAID protection levels are also accomplished using SmartMigration.

Performance-Focused Software Features

OceanStor has a number of features that have direct effect on the system performance. Some are basic to the system while others are configurable or add-on software with additional licensing charges. The hardware implementations with direct impact on performance were discussed in the hardware section and include:

- Processors – number and type
- Memory used for caching
- PCIe switch for controller interconnect
- Interfaces – front (Host) end and back end connections

The scale-out architecture is also a major contributor for performance both in the ability to scale performance as more controllers are added to the configuration and in the ability to minimize performance impact during failure situations or system updates.

The additional features that can improve performance are described in the remainder of this section.

SmartCache - Caching data for improved hit ratios

DRAM in the controller used to cache data can be extended for read and write data with the use of SSDs. Any installed SSDs may be configured as SmartCache as a pool. A default pool is created on each controller. SmartCache will manage the cache dynamically based on algorithms for data access. Additionally, metadata for filesystems are cached to accelerate file access.

SmartPartition - Allocating cache per application

Administrators can optimize performance for critical applications by partitioning system resources for applications. With SmartPartition, the OceanStor system will ensure that the application has the allocated cache resources for exclusive use. Administrators can specify read and write cache sizes for a partition, allowing for different I/O characteristics. Statistics reported on resource utilization allow the partitioning to be adjusted. SmartPartition can be used in combination with the quality of service (SmartQoS) feature.

SmartTier - Dynamic tiering of data

Dynamic tiering involves moving data to different locations with different performance characteristics based on the frequency of usage. SmartTier dynamically tiers data to manage the trade-off between capacity and performance. The tiering is based on LUNs and works at the sub-LUN (block) level with a granularity from 512KB to 64MB called an extent. Activity levels are monitored and analyzed to make decisions for moving data between different media types.




Figure 6: OceanStor SmartTier (source: Huawei)

The user can define the monitoring periods for data collection for SmartTier. Statistics during the period are collected on reads and writes at the block level. The analysis of the data will be determined by ranking which blocks should be promoted or demoted in the media hierarchy. SmartTier will migrate data based on the analysis. The start time and duration allowed for migration can be specified by the user.

SmartQoS - Quality of Service

Resources in the storage system can be controlled to provide a quality of service for defined LUNs using the feature SmartQoS. SmartQoS controls I/O resource by managing priorities that are defined by the user for the selected LUNs. Three areas of controls are defined:

- I/O priority scheduling – allocation of system resources is given priority based on settings of high, medium, and low.
- I/O traffic control – user defined IOPs or bandwidth goals will cause some LUN operations to be restricted so that QoS controlled LUNs can meet their goals.
- I/O performance assurance – minimum performance settings can be set where other operations can be delayed until the selected minimums are met.

Resources managed by QoS include front-end interface, processor execution, cache, and device access.

File Access - OceanStor File Engine

Huawei OceanStor supports both file and block access as an integrated function in the software. File systems with a maximum size of 256TB are allocated the same as thinly provisioned LUNs from the storage pool. The software operating in OceanStor providing file access is called a File Engine by Huawei.

Snapshots are made of complete file systems, handled the same as snapshots of LUNs. The same rules for LUN snapshots apply. Currently, file locking between NFS and CIFS/SMB is not supported.

Protocols

The OceanStor File Engine supports a number of protocols including the following:

NFS v2, v3, CIFS, FTP, and HTTP

NFS

Multi-level NFS file system supports export with versions of NFS (v2 and v3). This allows specific portions of a path to be exported separately, with different permissions. LDAP is supported for authentication.

CIFS/SMB

CIFS is supported with a SAMBA implementation. SMB2 is the level of CIFS support. Active Directory integration is supported for security and authentication.

Evaluator Group Comment: Very little detailed information about the File Engine implementation is generally available. That leaves a customer with a great of assumptions that must be made or taking additional time to submit detailed questions. This is an impediment to making an expedient and informed decision. Without that information, a customer should be very skeptical about file usage until it made publicly available.

Data Protection

HyperSnap - Snapshot

A snapshot is a virtual copy of a LUN or filesystem created with a mapping table and a redirect-on-write implementation. Snapshot copies are instantly available. The copy can be scheduled with a timer feature or invoked immediately from the ReplicationDirector or through an API.

HyperClone - Clone Copy

A clone is a snapshot where a full data copy is made to a new LUN. A clone copy is done within a storage system and when the copy is complete, the LUN is split for independent updating. Note that a real-time

copy to a remote system is called a HyperMirror. The ReplicationDirector is used to create and manage clone copies.

HyperCopy - LUN Copy

A LUN copy is a block by block to a new LUN or multiple target LUNs that may be within a storage system or between storage systems including virtualized third party storage systems. The LUN copy function is typically used in tiered storage, application upgrades, and remote backup to disk (on a remote system). The ReplicationDirector is used to create and manage LUN copies.

HyperReplication - Remote Replication

Synchronous and asynchronous remote replication to a remote OceanStor system is performance with the replication software. Remote replication creates a mirror on a volume by volume basis. Periodic asynchronous (change data only) is not currently supported. ReplicationDirector manages the remote replication function.

Evaluator Group Comment: Setting up a remote replicated environment is very complicated and needs to be well-understood to gauge the ability to protect from disasters. The information provided publicly is limited so either additional information must be obtained or the implementation must be done by the services group with the knowledge required.

Consistency Groups

Consistency groups are supported with Asynchronous and Synchronous replication, which ensures that all data on the marked logical volumes are consistent to a specific point in time by managing the replication operation in order defined batches. Consistency groups are defined in ReplicationDirector and based on volumes and not the host, which the volumes are mapped to, providing consistency across multiple hosts and applications. Both synchronous and asynchronous replication supports consistency groups, which enables applications that utilize multiple volumes the ability to maintain application consistency when replicating between systems.

A summary of the data protection products is provided in the following table.

	Snapshot	Clone Copy	LUN Copy	Remote Copy
Copy Type	Redirect on Write	Clone within same system	Multiple Targets, heterogeneous systems	Sync and Async
Architecture	Bit Map	Copy	Copy	I/O copy
Max # Copies	32,768 total 2,048 per LUN or file system	256 primaries 512 secondaries	256 Max 128 targets	2,048 replication pairs 512 consistency groups
Re-sync	No	No	No	Yes (w/ Async)
Potential Data Exposure	Exposure to time activated	Exposure to time Clone initiated	Exposure to time LUN Copy initiated	None with Sync, time between re-sync for Async
Availability for Restore	Available for restore immediate, or able to replace volume after copy complete	Available for restore, or able to replace volume immediately	Available for restore, or able to replace volume when copy complete	Sync: immediate w/ failover
Disaster Recovery, Fail-over Capability	Yes Able to replace volume after command completes	Yes Able to replace volume after copy complete	Yes Able to replace volume after copy complete	Nearly instantaneous for sync and async, auto-failover is supported
Default Action	Copy when command issued	Copy when command issued	Copy when command issued	Depends upon mode
Additional Capacity	Variable	100%	100%	100%
Access to Copy	R/W after snap	R/W after complete	R/W after complete	Only after failover
Access to Source	R/W after complete	R/W During Copy	R/W During Copy	Full access

Table 2: OceanStor Data Protection Options Comparison

Integration with Application or System Software

OceanStor and VMware vSphere

The OceanStor has a plug-in for vSphere vCenter to provide an integrated view of VMs and associated storage resources. From the vSphere console, an administrator can identify:

- Mapping of VMs to physical and logical storage resources including LUNs and volumes.
- LUN properties (thin vs. fat, etc.)
- Device types (SSD, FC, SATA disk)
- RAID groups and RAID levels

VAAI Support

Included since vSphere 4.1 are the following storage-related APIs (VAAI), which OceanStor supports:

Full Copy - leverages existing SCSI commands available in the SCSI command set and sent from the ESXi host to the storage array to clone, snapshot, and migrate data within an array using array-based software. These functions are normally used when creating new VMs and will be particularly useful in Virtual Desktop Infrastructure (VDI) deployments. No ESXi host server cycles are consumed in the performance of these copy functions, enhancing overall VMware performance and reducing the time required to create new VMs and virtual desktops when host-based copy functions are used for these purposes. Only written data is copied to improve the efficiency.

Hardware-Assisted Locking - supports block-level granularity when protecting VMFS metadata. Previously, SCSI reservations were used at the LUN level to satisfy the same requirement.

Block Zeroing - speeds-up the standard disk initialization process in a VMware environment. This function is particularly useful in reducing the time required to create virtual disks in eager-zero thick (EZT) format.

OceanStor and Microsoft System Center

The OceanStor has a plug-in for Microsoft System Center for managing the virtual environment for the OceanStor storage resources.

Advanced Features

SmartMotion - Capacity and Performance Balancing

Ongoing analysis of the system causes SmartMotion to move data across storage media of the same type to balance capacity and performance. The granularity of data moved is 64MB.

SmartDedupe & SmartCompression - Data Reduction

Data reduction is implemented with deduplication and compression. Data deduplication is done for file systems and thin LUNs. Because of the variable chunk size used when the thin LUNs or file systems were created, deduplication uses the granularity of the volume or file system. The chunks have a digital fingerprint calculated and the result is checked against existing fingerprints for duplicate data. The user can also enable a byte-by-byte comparison in case there is a concern about the integrity of digital fingerprints.

After deduplication, compression is performed on the chunks of data that are unique. Multiple blocks in successive logical block addresses are concatenated to improve the effectiveness of the compression algorithm if the block size is less than 32KB. Blocks greater than 32KB are compressed individually. If the pre-check on compression shows the effectiveness is below a defined threshold, data is not compressed in order to speed the retrieval.

Hardware add-in cards for the controllers are used to perform the deduplication and compression as an inline operation.

SmartErase - Digital Overwrite

A LUN may be selected for SmartErase and any unallocated space is overwritten to securely erase the data. The number of times the overwrite is performed is not specified by Huawei. Up to 16 LUNs may be erased at a single time by a controller.

Evaluator Group Comment: Digital erase of data is more effective and applicable at the file level. This could be an area for improvement.

SmartMulti-Tenant - Access Isolation

Storage resources in the OceanStor system can be configured with SmartMulti-Tenant to have isolated access and management for use in multi-tenant environments.

Reliability, Availability, Serviceability Features

The following summarizes the RAS capabilities:

- Redundant hot-plug cooling fans, batteries, and power supplies are standard on all models
- RAID protection of device storage
- Online serviceable device expansion modules
- Redundant, hot-pluggable controller canisters can be scaled non-disruptively
- Remote service management processor provides diagnostic reports to service center and notifies technical support through its internal online support service in the event of a failure. Information provided includes the part(s) required and system location.
- End-to-end integrity checking

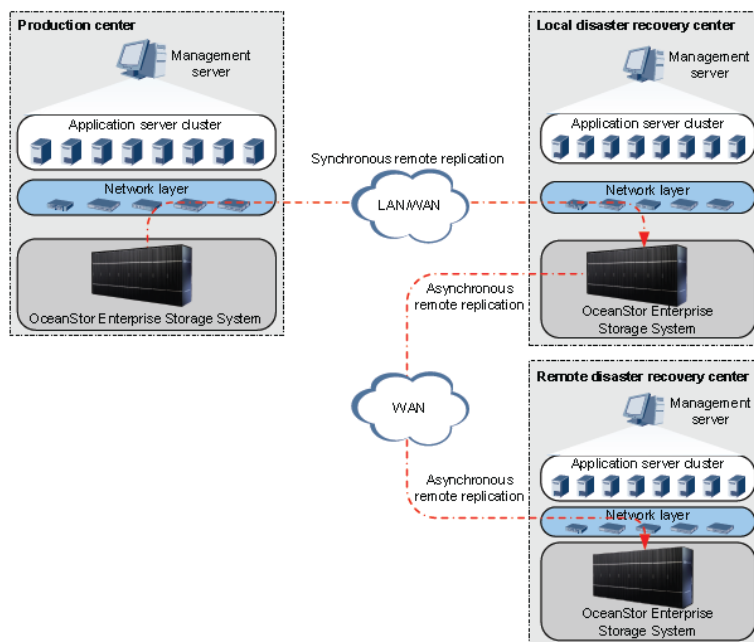


Figure 7: Failover Access for Huawei OceanStor (source: Huawei)

Error Checking and Correction (ECC)

Data integrity is provided through end-to-end ECC calculations on each host read or write operation throughout the entire I/O path.

Data Scan

OceanStor provides continuous scan of data to repair bad sectors. There are no administrator settings with respect to bad sector repair.

Redundant Power System

The power system is based on redundant power supplies in the modular enclosures with dual line cord inputs. All power system components are replaceable non-disruptively.

Fault Isolation

Fault isolation and reliability features are included with elements in the controller and device enclosures, which all report and isolate faults.

Sparing

Instead of providing spare capacity at the drive level, Huawei OceanStor systems provide global hot spare chunks of data, distributed across physical devices when capacity is allocated.

Controller Redundancy

Controller nodes are configured in logical pairs whereby each controller node has a partner node. The two partner nodes have redundant physical connections to the subset of disk drives owned by the node pair, mirror their write cache to each other, and serve as the backup node for the logical disks owned by the partner node.

Upon the failure of a controller node, the node failover recovery process automatically flushes the dirty write cache to disk, transfers ownership for the volumes owned by the failed node to its partner node, and puts all logical disks owned by the remaining partner node in write-thru (non-cached) mode. Since a given volume on the storage system consists of several logical disks spread across all configured nodes, the failure of any single node in a system with four or more nodes results in only a portion of logical disks going into write-thru mode.

Evaluator Group Comments: This level of availability is better than two-controller mid-range systems. For customers who feel two controller systems do not provide all the availability and performance they need, the use of a scale-out architecture such as Huawei OceanStor would be preferable.

HyperMirror - Clustered High Availability

HyperMirror creates a second copy of data in real time on local or remote LUNs. Effectively, this provides the ability for application and system software to retry access in case the primary path/access becomes available to the mirrored copy to continue operations. While not exactly a stretched cluster as implemented by other vendors, HyperMirror does provide the high availability access in case of a local failure in access to data on an OceanStor system. For the 18000, a maximum of 512 LUNs may be in a HyperMirror.

The logical volumes are kept in synchronization at each site for the stretched cluster by use of redundant fabrics as described in the following diagram.

Concurrent Maintenance

The modular architecture permits modules to be taken out of service while others take over the workload. Any one portion of a redundant pair may be taken out of service and or replaced without interruption to service. That is, each controller in a controller pair may be replaced at a time, along with each drive cage in an HA configuration may also be replaced without disruption of operations. Similarly, batteries, power supplies and drives may also be replaced.

Software Updates

The embedded software also known as the Operating System may be upgraded on the controllers non-disruptively. Each controller runs its own copy of the Operating System.

Remote Support

Remote support is provided on OceanStor systems with the management unit by use of SMTP to send messages to support personnel. This provides remote monitoring of system conditions, and provides error alerts. There are no separate fees for remote support. However, users must have a valid support contract.

Service Contracts

A variety of service contracts are available, including no contract, 1 year, 2 year and 3 year contracts, with 24X7 coverage, 9X5, and response times of either 2 hours, 4 hours or next business day.

Management

OceanStor has various functions of management software that execute on a separate server or computing system (called a maintenance terminal), on an application server, and on a management module in the storage controller enclosure.

Management Module

With the Management Module in the storage controller enclosure, OceanStor provides several no-charge options for management including:

- GUI – DeviceManager executes on the Management Module and provides configuration, management, and maintenance action.
- CLI for scripted or remote administration through the management network port or the serial port on the storage system
- SNMP support for management and configuration through third party software. SNMP traps can be used for reporting of conditions.
- Syslog is a data capture and reporting function in the Management Module that sends alerts, alarms, and triggered events to user selected targets over protocols such as SMTP.

OceanStor DeviceManager

The OceanStor DeviceManager is the element manager presenting a GUI interface for managing the OceanStor system. Configuration, management and maintenance are performed through the DeviceManager. Wizards to configure the OceanStor system are invoked through the DeviceManager. Access to the DeviceManager is through a standard browser.

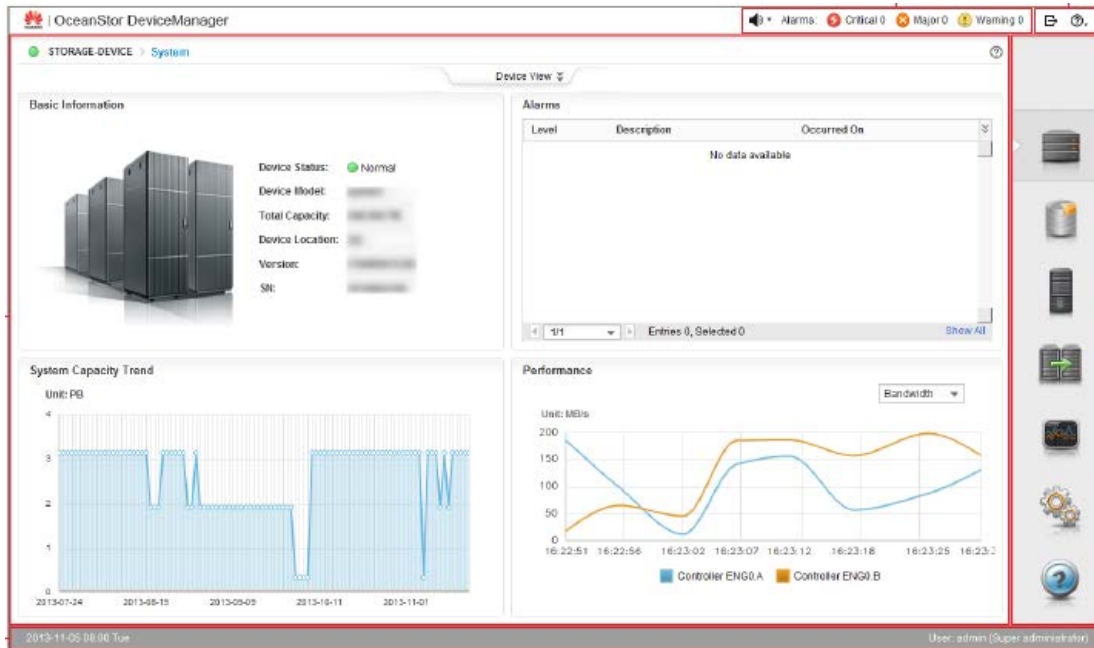


Figure 8: OceanStor Device Manager (source: Huawei)

Maintenance Terminal

The “maintenance terminal” has management software for OceanStor that includes:

- OceanStor Toolkit – for service personnel to deploy, maintain, and update systems
- OceanStor SystemReporter – performance and capacity reporting software
- OceanStor Cloud Service – remote maintenance software for monitoring, alerting, and status interrogation

Application Server

Management software with specific purpose executes on an application server:

- Replication Director – manages the remote replication, snapshot, LUN copy, and clone functions.
- UltraPath – is multipath driver software for optimizing accessing the OceanStor system.
- eSDK OceanStor – are software tools for integrating with OceanStor. Included are plug-ins for VMware vCenter and Microsoft System Center.



Performance

Published information from Huawei gives the bandwidth supported by the OceanStor 18000 V3 as 192 GB/s. SPC performance benchmark information is available on the website

www.storageperformance.org with the latest published information from November 2015. The performance data shows:

3,010,007 IOPs for a 69TB system with mirrored data protection. The system was priced at \$2.3M.

The SPC website should be checked for the latest information.

Evaluator Group Comments

Strengths:

The Huawei OceanStor 18000 Systems deliver many high-end capabilities features for enterprise usage. The scale-out design with up to sixteen controllers has significant advantages over standard monolithic systems. Huawei has continued to evolve the OceanStor systems with a large number of earlier generation products as testament to a commitment to continue to improve the product. The use of standard hardware components allows systems to be delivered more quickly.

There are a number of high performance features for the OceanStor system in both hardware and software implementations. Most of the current software advances including usage of solid state storage technology have been incorporated into the system.

Customers who are looking to consolidate several modular systems or for new storage systems with advanced capabilities will consider the Huawei OceanStor high-end systems.

Perceived Challenges:

Huawei is virtually unknown in North America, which is a very large market that also influences overseas subsidiaries. Launching into this should be expected in order to broaden the opportunities.

Much more information needs to be readily available for customers about the product in order for them to do an initial selection of which products to investigate further. The current information is inadequate for forming detail considerations.

The use of scale-out and clustering of systems is not explained well and needs more, in-depth information to give understanding of what this really means. Within a single system, scale-out seems clear but clustering them together is not.

More detailed information is available at www.evaluatorgroup.com

Copyright 2018 Evaluator Group, Inc. All rights reserved.

No part of this publication may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording, or stored in a database or retrieval system for any purpose without the express written consent of Evaluator Group Inc. The information contained in this document is subject to change without notice. Evaluator Group assumes no responsibility for errors or omissions. Evaluator Group makes no expressed or implied warranties in this document relating to the use or operation of the products described herein. In no event shall Evaluator Group be liable for any indirect, special, inconsequential or incidental damages arising out of or associated with any aspect of this publication, even if advised of the possibility of such damages. The Evaluator Series is a trademark of Evaluator Group, Inc. All other trademarks are the property of their respective companies.