Hewlett Packard Enterprise

HPE Alletra MP

2025 06

Martynas Skripkauskas

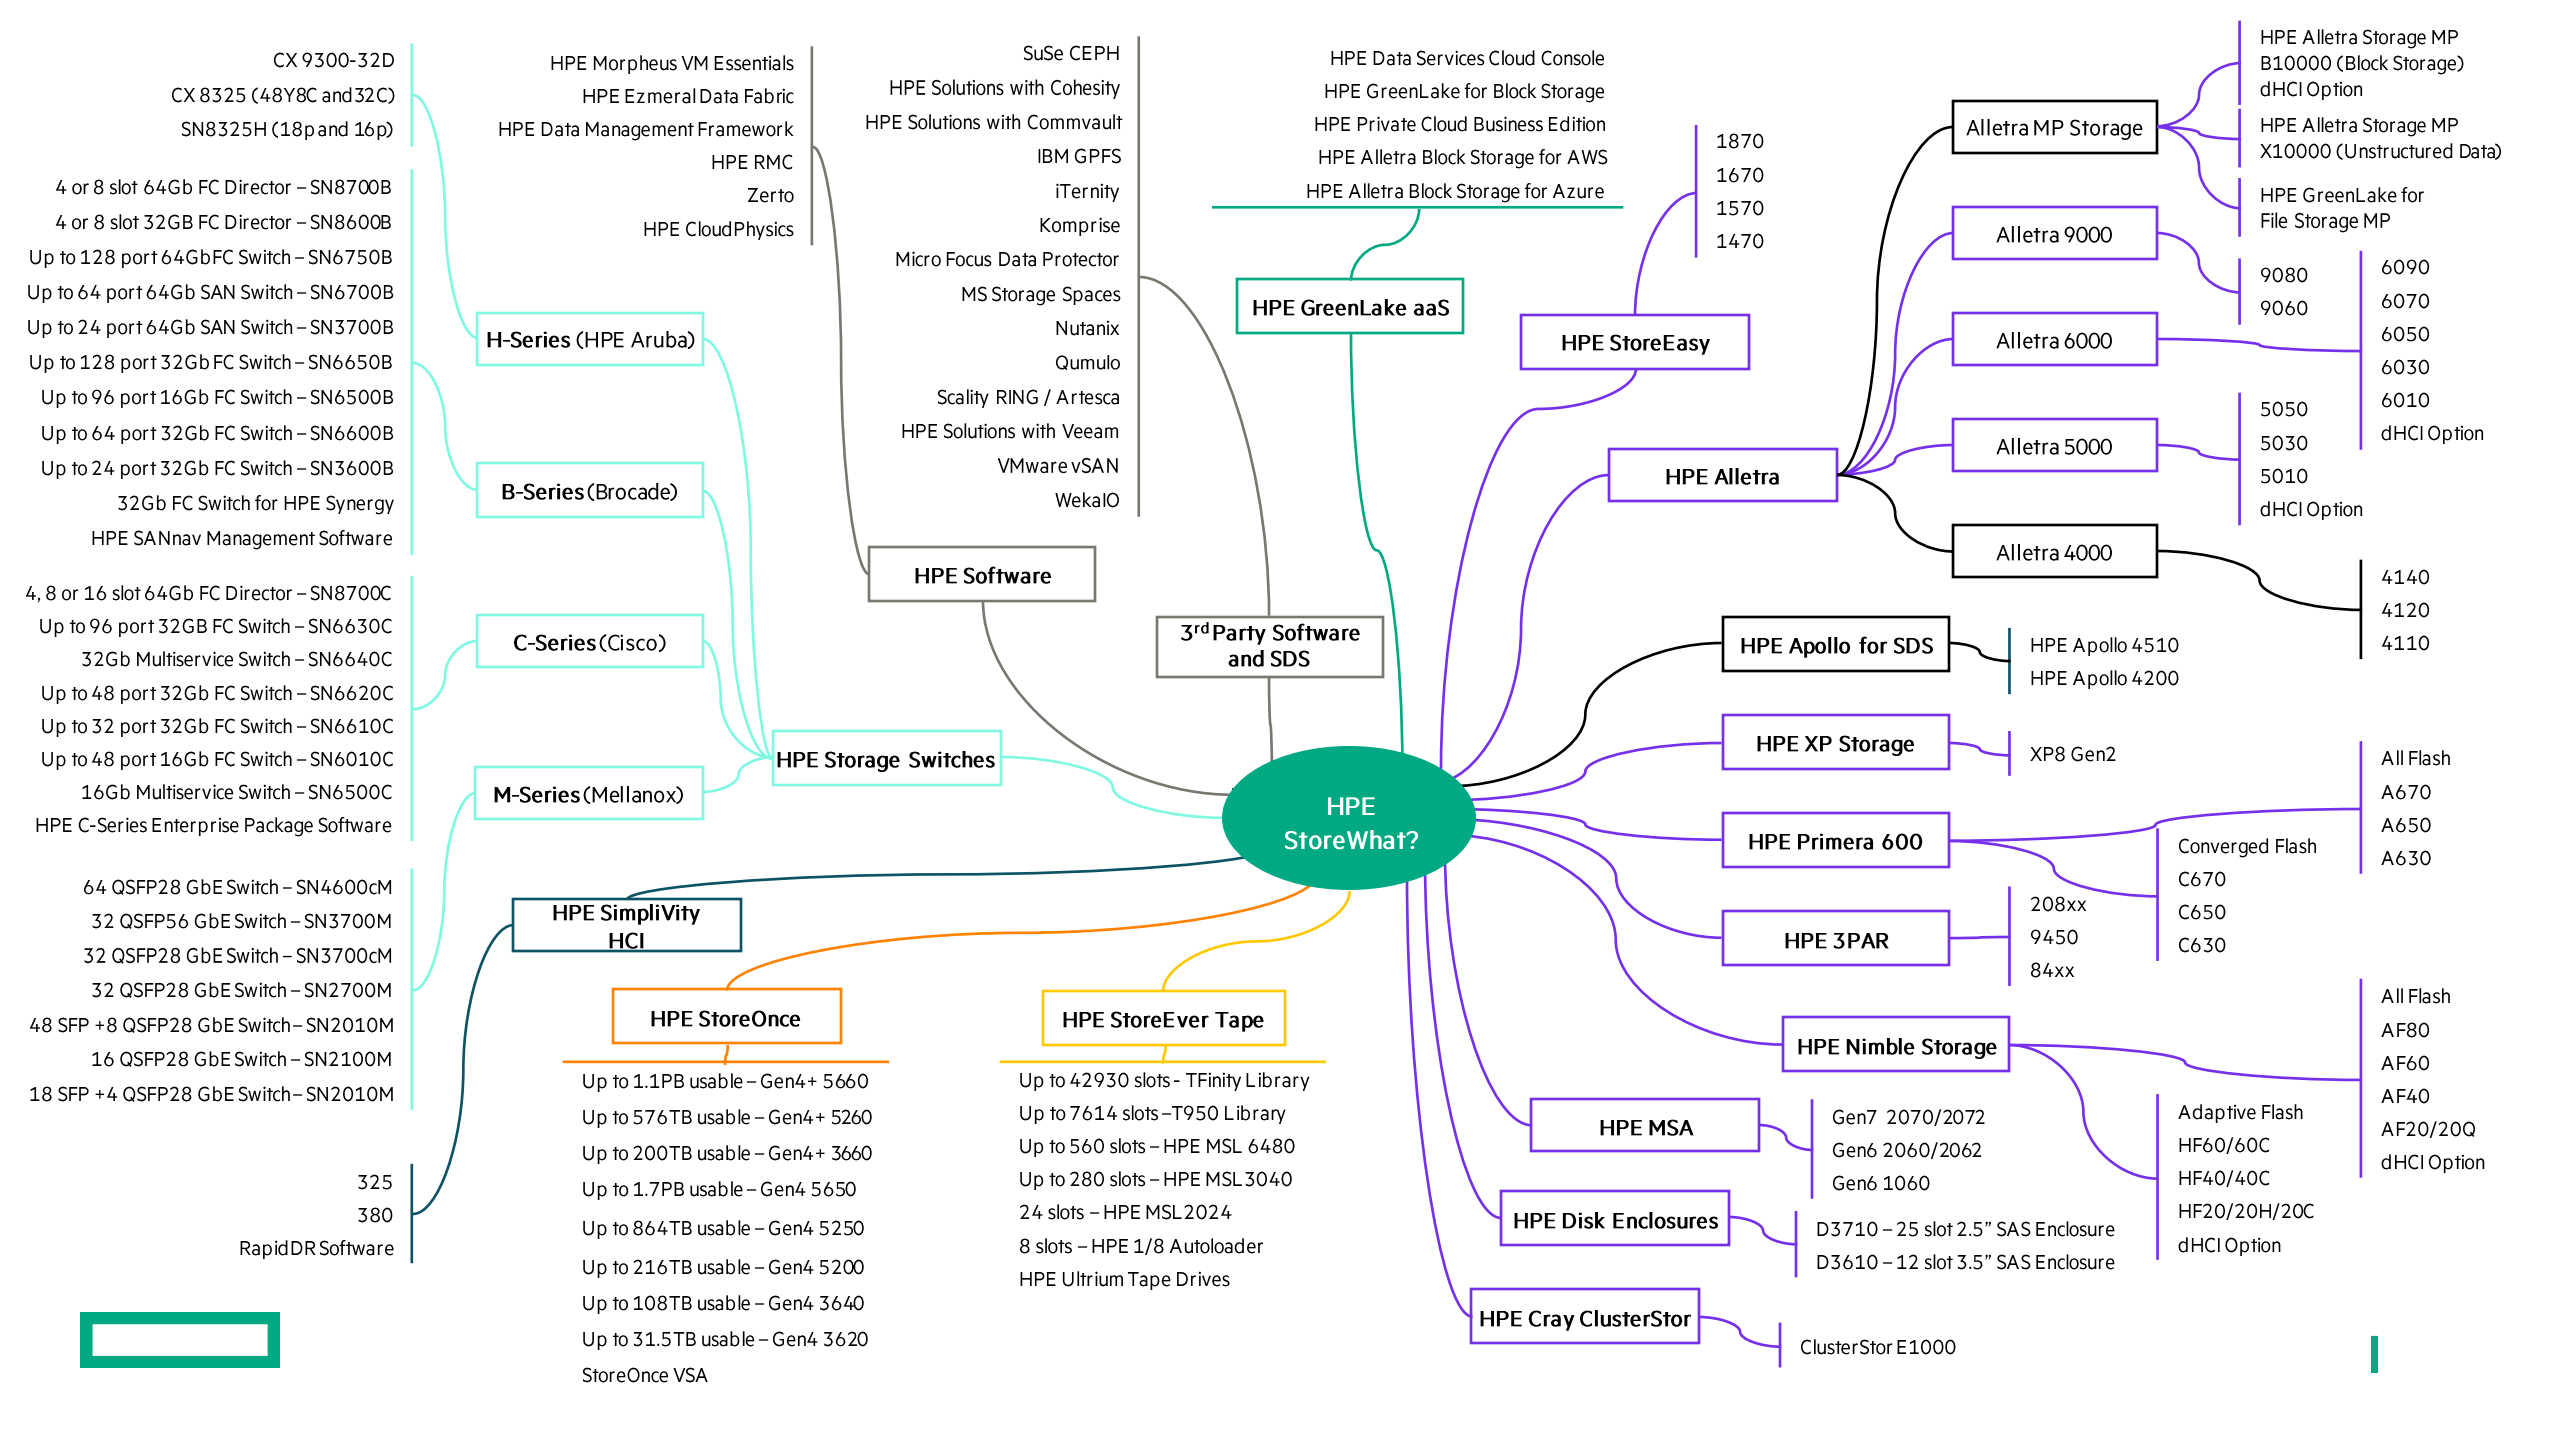
HPE StoreWhat?
HPE Storage Switches
H-Series (HPE Aruba)
CX 9300-32D
CX 8325 (48Y8C and 32C)
SN8325H (18p and 16p)
B-Series (Brocade)
4 or 8 slot 64Gb FC Director – SN8700B
4 or 8 slot 32GB FC Director – SN8600B
Up to 128 port 64Gb FC Switch – SN6750B
Up to 64 port 64Gb SAN Switch – SN6700B
Up to 24 port 64Gb SAN Switch – SN3700B
Up to 128 port 32Gb FC Switch – SN6650B
Up to 96 port 16Gb FC Switch – SN6500B
Up to 64 port 32Gb FC Switch – SN6600B
Up to 24 port 32Gb FC Switch – SN3600B
32Gb FC Switch for HPE Synergy
HPE SANnav Management Software
C-Series (Cisco)
4, 8 or 16 slot 64Gb FC Director – SN8700C
Up to 96 port 32GB FC Switch – SN6630C
32Gb Multiservice Switch – SN6640C
Up to 48 port 32Gb FC Switch – SN6620C
Up to 32 port 32Gb FC Switch – SN6610C
Up to 48 port 16Gb FC Switch – SN6010C
16Gb Multiservice Switch – SN6500C
HPE C-Series Enterprise Package Software
M-Series (Mellanox)
64 QSFP28 GbE Switch – SN4600cM
32 QSFP56 GbE Switch – SN3700M
32 QSFP28 GbE Switch – SN3700cM
32 QSFP28 GbE Switch – SN2700M
48 SFP +8 QSFP28 GbE Switch – SN2010M
16 QSFP28 GbE Switch – SN2100M
18 SFP +4 QSFP28 GbE Switch – SN2010M
HPE Software
HPE Morpheus VM Essentials
HPE Ezmeral Data Fabric
HPE Data Management Framework
HPE RMC
Zerto
HPE CloudPhysics
3rd Party Software and SDS
SuSe CEPH
HPE Solutions with Cohesity
HPE Solutions with Commvault
IBM GPFS
iTernity
Komprise
Micro Focus Data Protector
MS Storage Spaces
Nutanix
Qumulo
Scality RING / Artesca
HPE Solutions with Veeam
VMware vSAN
WekaIO
HPE GreenLake aaS
HPE Data Services Cloud Console
HPE GreenLake for Block Storage
HPE Private Cloud Business Edition
HPE Alletra Block Storage for AWS
HPE Alletra Block Storage for Azure
HPE StoreEasy
1870
1670
1570
1470
HPE Alletra
Alletra MP Storage
HPE Alletra Storage MP B10000 (Block Storage) dHCI Option
HPE Alletra Storage MP X10000 (Unstructured Data)
HPE GreenLake for File Storage MP
Alletra 9000
9080
9060
Alletra 6000
6090
6070
6050
6030
6010
dHCI Option
Alletra 5000
5050
5030
5010
dHCI Option
Alletra 4000
4140
4120
4110
HPE Apollo for SDS
HPE Apollo 4510
HPE Apollo 4200
HPE XP Storage
XP8 Gen2
HPE Primera 600
All Flash
A670
A650
A630
Converged Flash
C670
C650
C630
HPE 3PAR
208xx
9450
84xx
HPE Nimble Storage
All Flash
AF80
AF60
AF40
AF20/20Q
dHCI Option
Adaptive Flash
HF60/60C
HF40/40C
HF20/20H/20C
dHCI Option
HPE MSA
Gen7 2070/2072
Gen6 2060/2062
Gen6 1060
HPE Disk Enclosures
D3710 – 25 slot 2.5” SAS Enclosure
D3610 – 12 slot 3.5” SAS Enclosure
HPE Cray ClusterStor
ClusterStor E1000
HPE SimpliVity HCI
325
380
RapidDR Software
HPE StoreOnce
Up to 1.1PB usable – Gen4+ 5660
Up to 576TB usable – Gen4+ 5260
Up to 200TB usable – Gen4+ 3660
Up to 1.7PB usable – Gen4 5650
Up to 864TB usable – Gen4 5250
Up to 216TB usable – Gen4 5200
Up to 108TB usable – Gen4 3640
Up to 31.5TB usable – Gen4 3620
StoreOnce VSA
HPE StoreEver Tape
Up to 42930 slots - TFinity Library
Up to 7614 slots –T950 Library
Up to 560 slots – HPE MSL 6480
Up to 280 slots – HPE MSL 3040
24 slots – HPE MSL 2024
8 slots – HPE 1/8 Autoloader
HPE Ultrium Tape Drives

HPE Alletra Storage MP B10000 overview

Learn more about the most advanced storage technology

Return to the table of contents

HPE Alletra Storage MP B10000 Storage Future-Ready Program

Transform Your Storage Ownership Experience with Guaranteed IT outcomes

Acquire with confidence | **Run without surprises** | **Grow with ease**

- 100% Data Availability Guarantee
- StoreMore Guarantee
- Timeless Controller Refresh
- 30-day Satisfaction Guarantee
- All-inclusive Software – no extra fees
- Zero RTO/RPO Guarantee
- Cyber Resiliency Guarantee
- Energy Consumption Guarantee

HPE Alletra Storage MP – The hardware

2U chassis providing

1. Two IO modules (IOM ≙ controller or node)
   - Single socket AMD CPU
   - Variable memory configuration
   - TPM chip
6. - OCP slots
     - Node IOM: 4 standard OCP adapter slots
     - JBOF IOM: 2 standard OCP adapter slots
2. PCIe Gen5 capable backplane
3. 24 dual ported NVMe EDSFF/U.2 drive slots
4. Two redundant power supplies
5. Chassis Discovery Module (CDM)

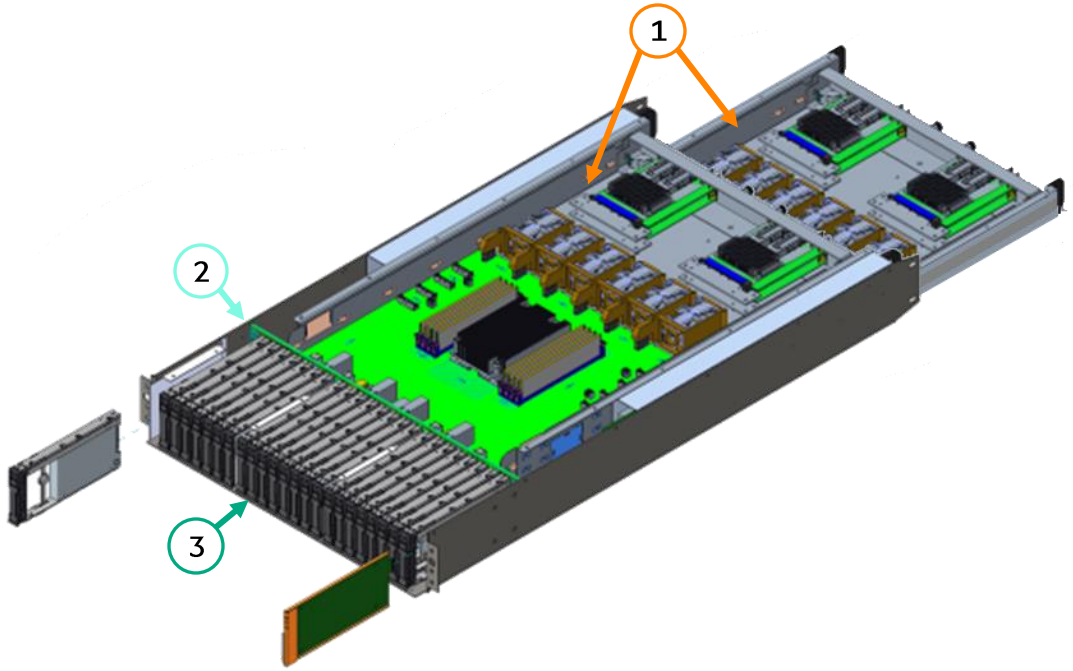

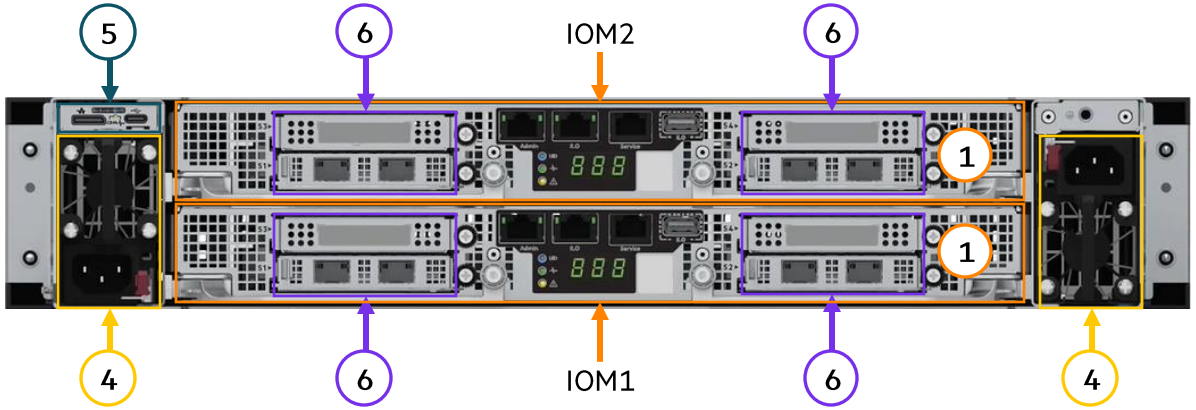

HPE Alletra Storage MP – The basics
Think of toy bricks used as standardized building-blocks

Standard building-blocks

HPE Alletra Storage MP chassis
with versatile hardware configurations

100GbE RoCEv2 capable
HPE network switch

SDS personalities

Controller (compute)
Block, file, or other personality

JBOF (capacity)
(RoCEv2, NVMe)

Compute + capacity
Storage array

HPE Alletra Storage MP
switch (NVMe-oF)

Potential storage systems

Stand-alone system

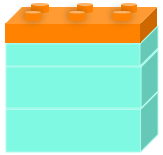

Stand-alone system
with JBOFs

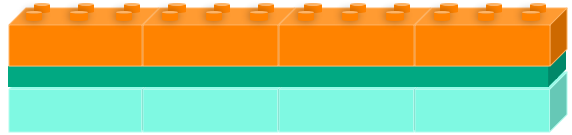

Switched multi-node system
with dedicated compute nodes and JBOFs

Systems built with HPE Alletra Storage MP

Potential storage systems

Stand-alone system

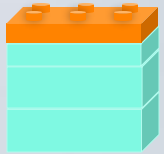

Stand-alone system
with JBOFs

Actual HPE storage systems

HPE Alletra Storage MP B10000
with 2 controller
and up to 24 SSDs

HPE Alletra Storage MP B10000
with 2 controller
up to 2 JBOFs and 72 SSDs

Potential storage systems

Switched multi-node system
with dedicated compute nodes and JBOFs

Actual HPE storage systems

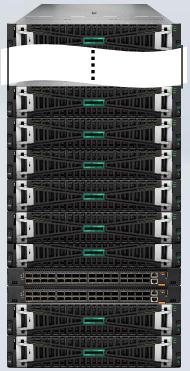

HPE Alletra Storage MP B10000
Multi-node Block Storage

HPE Alletra Storage MP X10000
Multi-node Object Storage

HPE Greenlake for File Storage MP
Multi-node File Storage

Pictures created by Peter Mattei, HPE

Disaggregated infrastructure – Composable storage – Possible Future

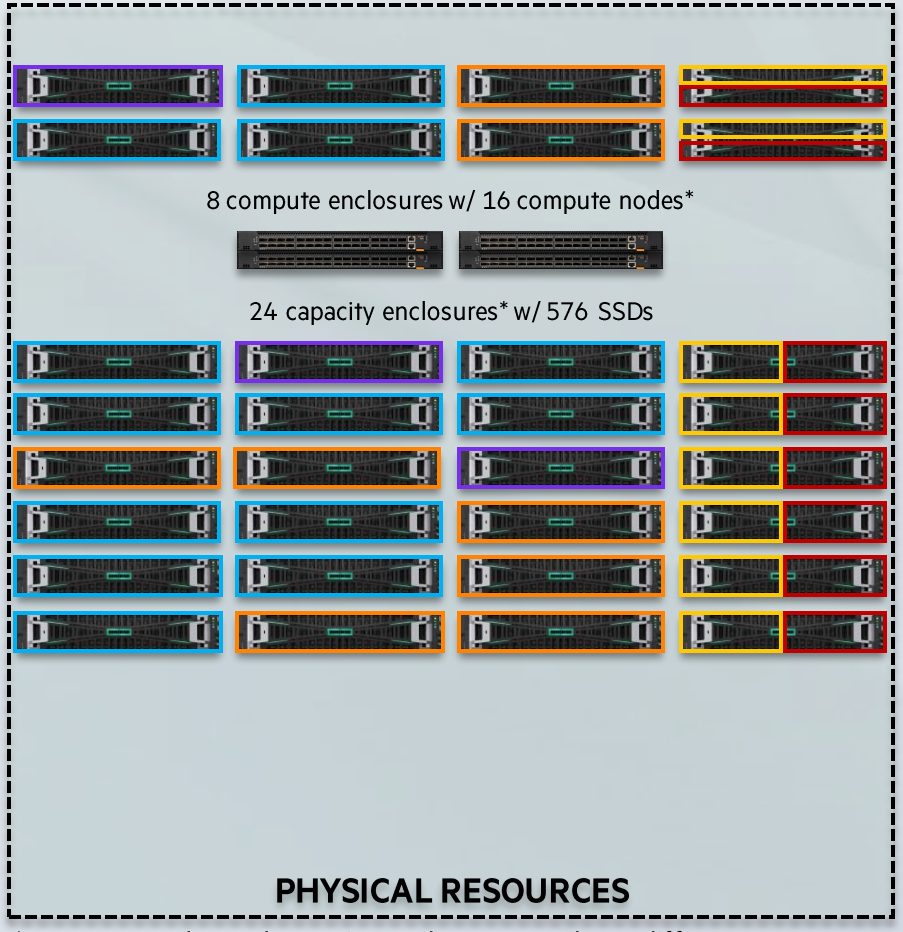

*Compute nodes and capacity enclosures can have different specs

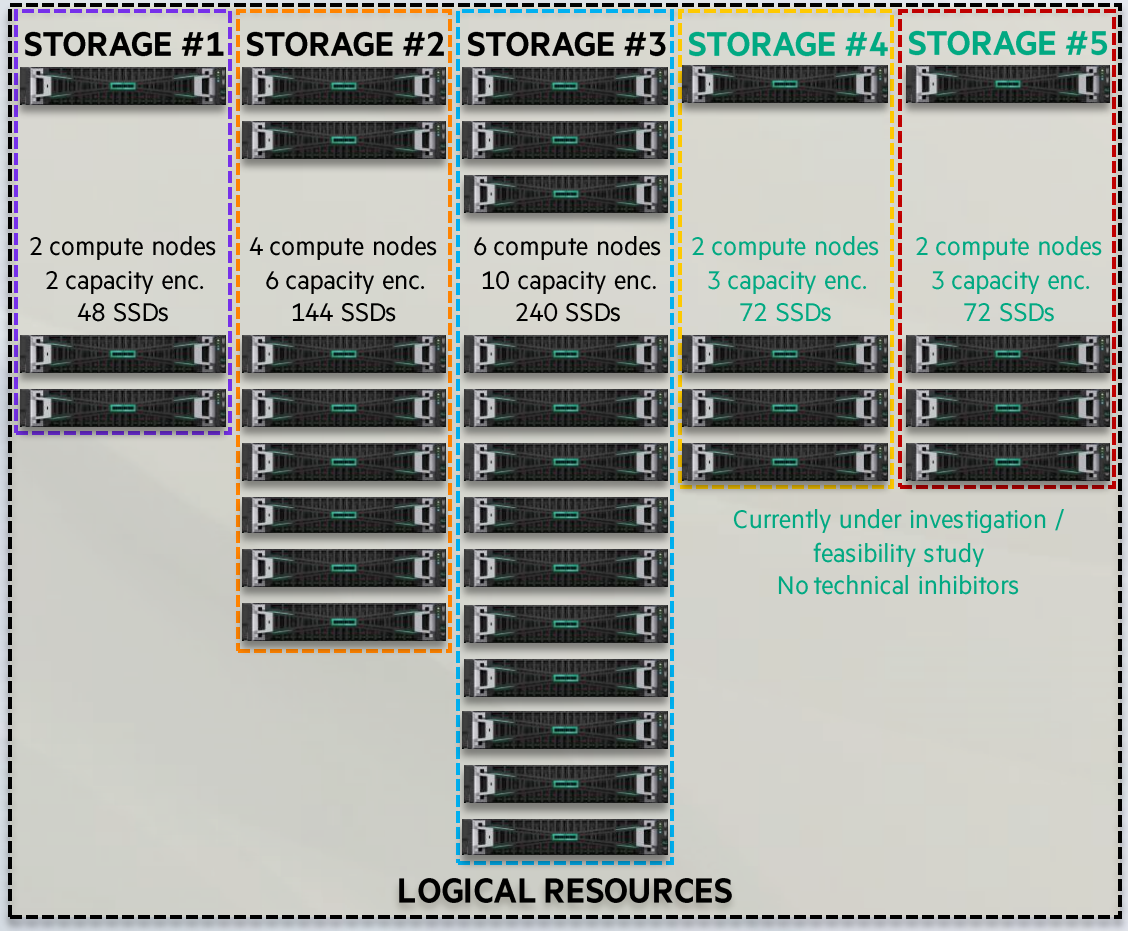

HPE Alletra Storage MP – The measure for mission-critical storage

From 15 TB to 5.6 PB
Extremely scalable mission-critical Storage at midrange economics

100% data availability
guaranteed[1]

Very performant and economic
thanks to the disaggregated architecture

Simplified management
with an intuitive on-prem experience

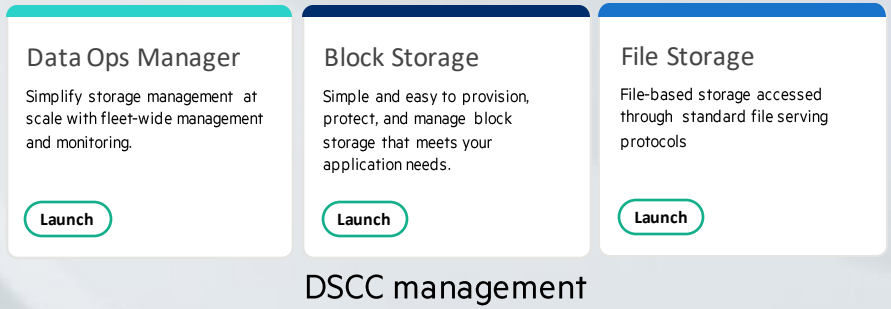

DSCC management

8-core nodes
Up to 24 NVMe SSDs

16- or 32-core nodes
up to 72 NVMe SSDs

Switchless 2-controller systems

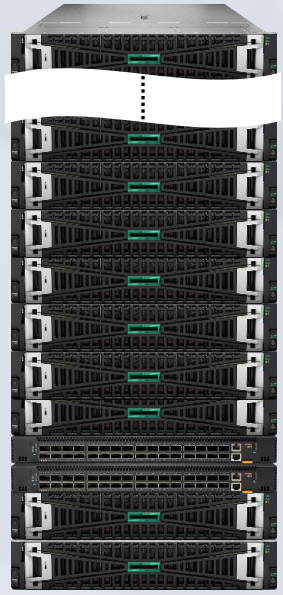

16 or 32-core nodes
up to 384 NVMe SSDs

Disaggregated systems with switched backend and 2-, 3-, or 4-controller

1 HPE Storage Substantiation

HPE Alletra Storage MP B10000 has gone where no vendor has gone before

HPE Alletra Storage MP is not like other flash architectures

- Based on HPE Alletra Storage MP, a disaggregated, all-NVMe common HPE storage platform
- Flexibility to scale-up and scale-out
- Designed to abstract data from media
- Built-in anomaly / ransomware detection engine
- Cloud-managed, data-centric, self-optimizing architecture
- Ease of use including:
  - Customer self-installable
  - Customer self-updateable
  - Customer self-repairable
- 100% data availability guarantee
- HPE Store More guarantee
- HPE All-inclusive array software
- HPE Timeless Storage program
- Also see the HPE Alletra Storage MP B10000 architecture white paper

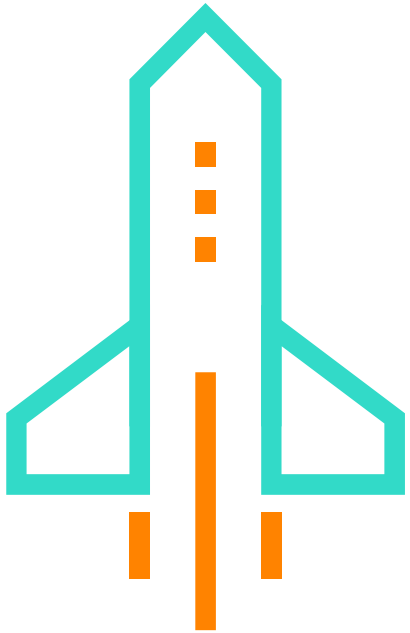

What's new in HPE GreenLake for Block Storage OS 10.x

Write-through Cache
No need for write cache mirroring anymore

New Snapshot Design
Improved space accounting

Enhanced Data Reduction
Dedupe stores balanced across nodes; Improved compression ratio

New Layout Manager
All nodes can see all drives

Singular VV Ownership
Virtual Volume and Logical Disks of a VV are owned by one node

HPE Alletra Storage ArcusOS – Enhanced software architecture

Accelerated write through architecture and single node VV ownership – short version

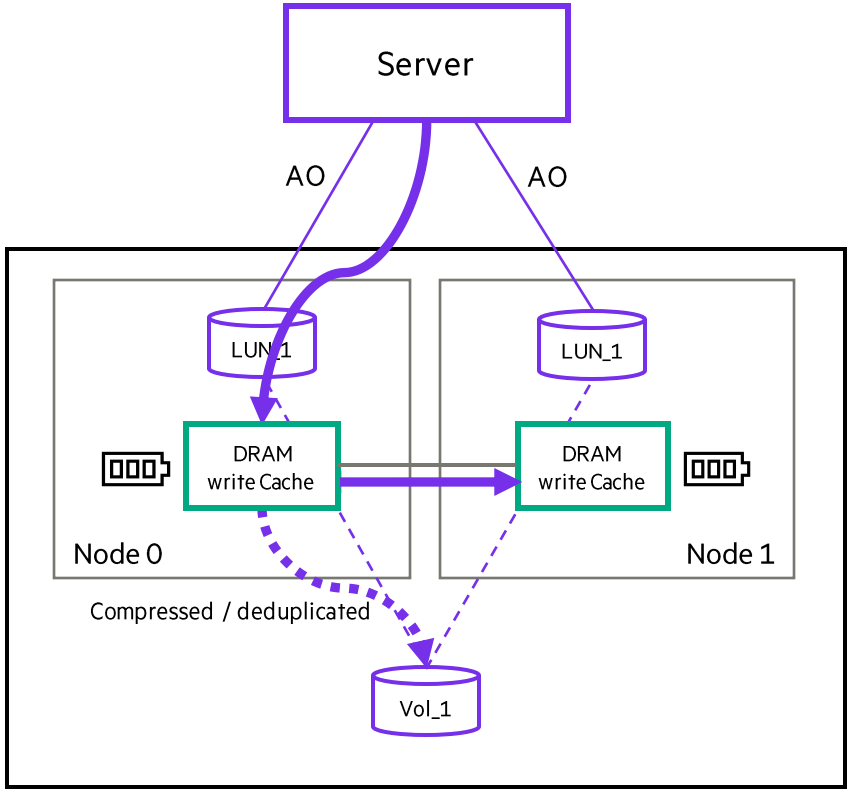

Traditional architecture

With write cache and batteries

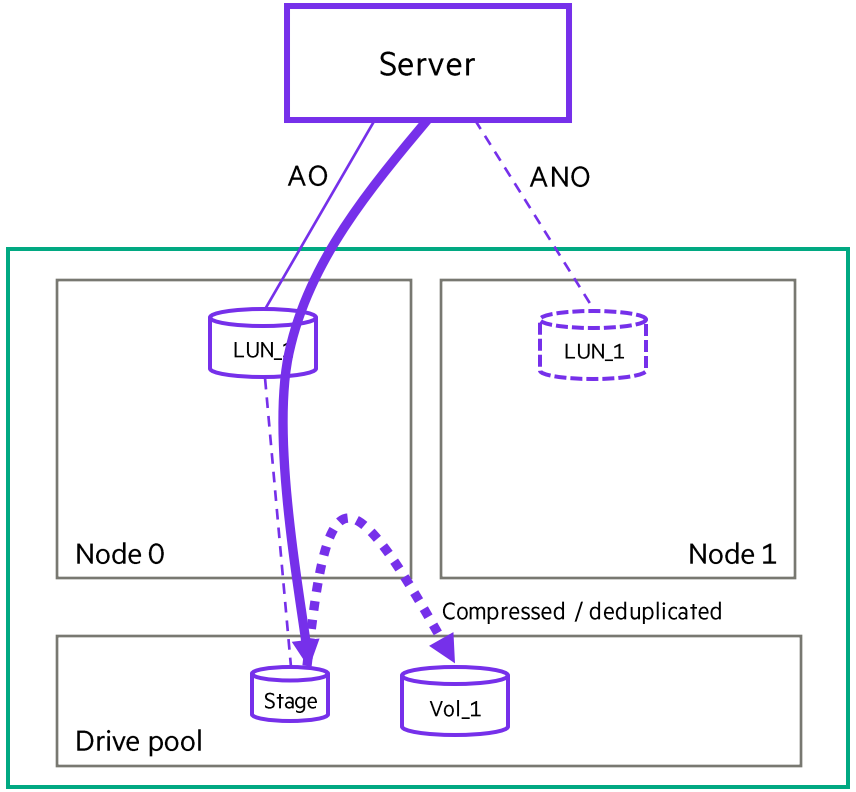

New HPE Alletra Storage ArcusOS architecture

No write cache and no batteries required

AO: Active Optimized paths

ANO: Active non-optimized paths

HPE Alletra Storage ArcusOS – Software details
System licensing model for CAPEX

All-inclusive system software included with every HPE Alletra Storage MP B10000 system

- **Block Storage**
- **NFS File Storage**
- **Thin provisioning**
- **Data reduction (deduplication and compression)**
- **Wide-striping**
- **Distributed sparing**
- **LDAP support**
- **Management UI**
- **CLI**
- **Web Services API**
- **DSCC** (Cloud-based SaaS management)
- **Ransomware detection engine**
- **Virtual Copy** (read/write, read-only and immutable snapshots)
- **Remote Copy** (sync and async, 2DC, 3DC)
- **Peer Persistence** (active-active storage metrocluster)
- **Virtual Domains** (array multitenancy)
- **System Reporter** (performance and capacity views and reports)
- **Dynamic Optimization** (non-disruptive changes to virtual volumes and array layout)
- **Priority Optimization** (QoS)
- **Peer Motion** (data movement between various HPE arrays)
- **Online Import** (data import from various competitors' arrays)

HPE Software and Software Support SaaS subscription
Mandatory for any HPE Alletra Storage MP B10000 system

Data-at-rest encryption
$1 license required to enable drive self-encryption (not permitted in certain countries)

HPE Alletra Storage MP B10000 – Data reduction technologies overview

- Built-in
  - HPE Alletra Storage ArcusOS is designed for thin provisioning and advanced data reduction
- In-band
  - Sequences of data are analyzed inline in 16 KB chunks
  - Chunks of zeroes are not written to the back-end, as they get deduplicated
  - Data is then compressed
- Reservation-less
  - Thin provisioning draws fine-grained increments from a single free space reservoir without any pre-allocation
  - Other vendors implementation might require a separate, pre-allocated pool for each data service level
- Integrated
  - APIs for automatic thin reclamation in Symantec File System, Oracle ASM, Microsoft Windows Server, Linux, vSphere, and others

Application	Data compaction ratio*
Virtual desktop infrastructure (VDI)	3.75–9.0X
Virtual server environments	2.25–3.75X
Databases	3.0–3.75X

* Based on an internal HPE study, the average data compaction savings per workload is derived from HPE Storage telemetry data at the time of publication.
HPE Store More Guarantee may be available for other workloads with a storage assessment. Contact your HPE sales or channel partner representative for more information

Dual Authorization overview for Block Storage

- Dual Authorization is an optional security process that requires a second user with Data Services Cloud Console Administrator privileges to approve requests to delete resources
- Dual Authorization for Block Storage workflows support the deletion of:
  - Volumes
  - Snapshots
  - Snapshot sets
  - Clones
  - Protection policies
  - Un-export of volumes, volumes sets, snapshots, and clones
- For information about how to enable and disable Dual Authorization, as well as how to approve or deny a deletion request, click the Dual Authorization tile in the Data Services Cloud Console and refer to the online help

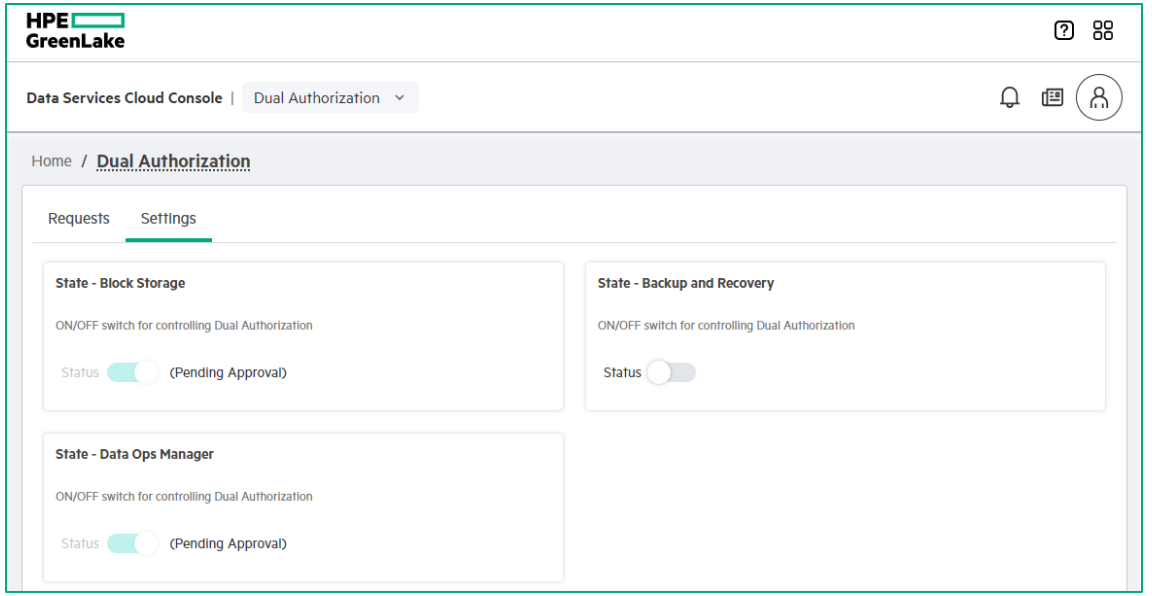

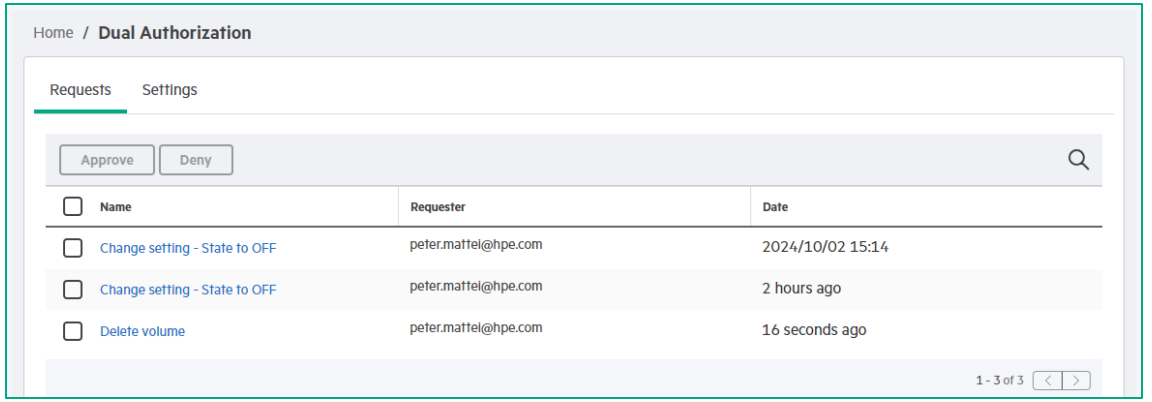

DSCC Extended Analytics

Data Operations Manager & Block Storage

HPE Alletra Storage MP B10000

Next generation system and volume-level troubleshooting and root cause analysis capabilities.

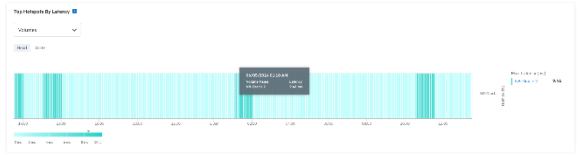

Key Features

1. **Headroom Utilization with Top Contributors**
   Trend analysis with breakdown by Volumes and Volume Sets
2. **Hotspots by Latency**
   Identifies volume and volume set outliers and provides drill-down to volume-level **Latency Insights** and **Workload Drift Insights**
3. **Latency with Correlated System and Workload Factors**
   Identifies changes of CPU, Disk and Read Cache and workload profile aligned to observed latency fluctuations
4. **Latency Variation Insights**
   Identifies high variations of latency within a sample interval
5. **Workload Drift Insights**
   Identifies unusual changes to a workload I/O profile and provides a detailed view into the distribution of the I/O sizes over time

1

Shows headroom utilization trends for the system. Headroom utilization is a comparison of the current I/Os on the system to a theoretical maximum level of I/Os that a similar device can support while maintaining an acceptable degree of latency. It is represented as a percentage and is designated as low, medium, high, and over-utilized levels. The graph shows the total system headroom utilization, the top five volume and volume set contributors, and the remaining contributors for the time period.

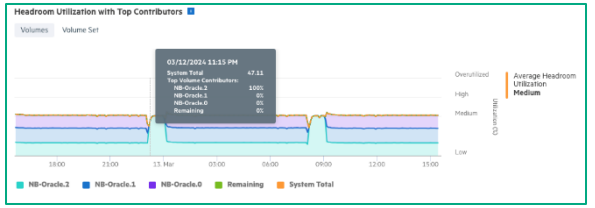

2

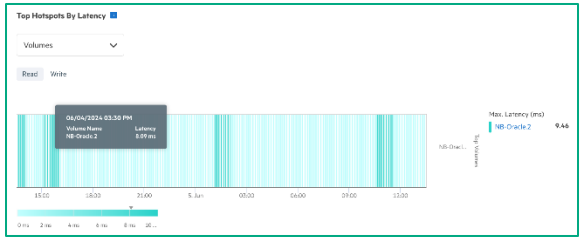

Shows up to 5 volume or volume set hotspots. Hotspots are outliers consistently reporting higher latency compared to other similar objects. Hotspots are ranked by how frequently they are reported as outliers during the selected time period. Hotspots might contribute to bottlenecks. If no volume or volume set hotspots are reported during high overall system latency, no outliers were reported.

3

Shows latency statistics and captures read and write metrics. In the latency graph, markers appear in 5 minute data granularity views on peaks for known factors contributing to the latency spike. Hovering over the marker displays the top factors. Select a marker to show the corresponding graph with detailed insights for these factors over the same time period.

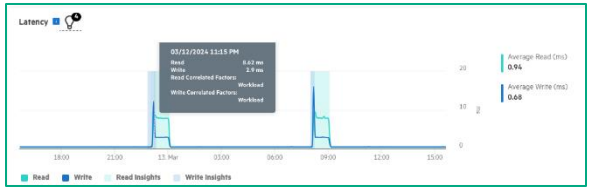

4

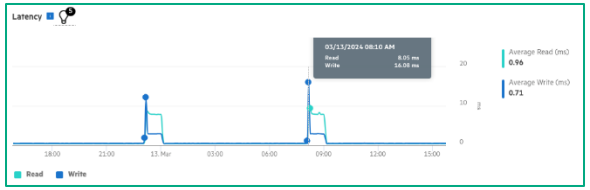

Shows latency statistics and captures read and write metrics. When high variations of I/O latencies occur within the sample duration, latency insights are overlaid onto the existing trends. The latency insights indicate that some I/Os have higher latencies although the average latency does not indicate a problem. The latency insights might explain higher latencies seen on a user application.

5

When a change in the I/O pattern is detected for the duration of the chart, throughput insights are overlaid on the existing trends. A change in the I/O pattern can indicate an unusual workload or an issue with the workload that might result in performance issues on the system as well as other workloads on the system in the future. Workload drifts are annotated on top of the throughput trend chart for the volume.

DSCC Cross Stack Analytics for vSphere

List your VMs and Datastores with a capacity and performance overview

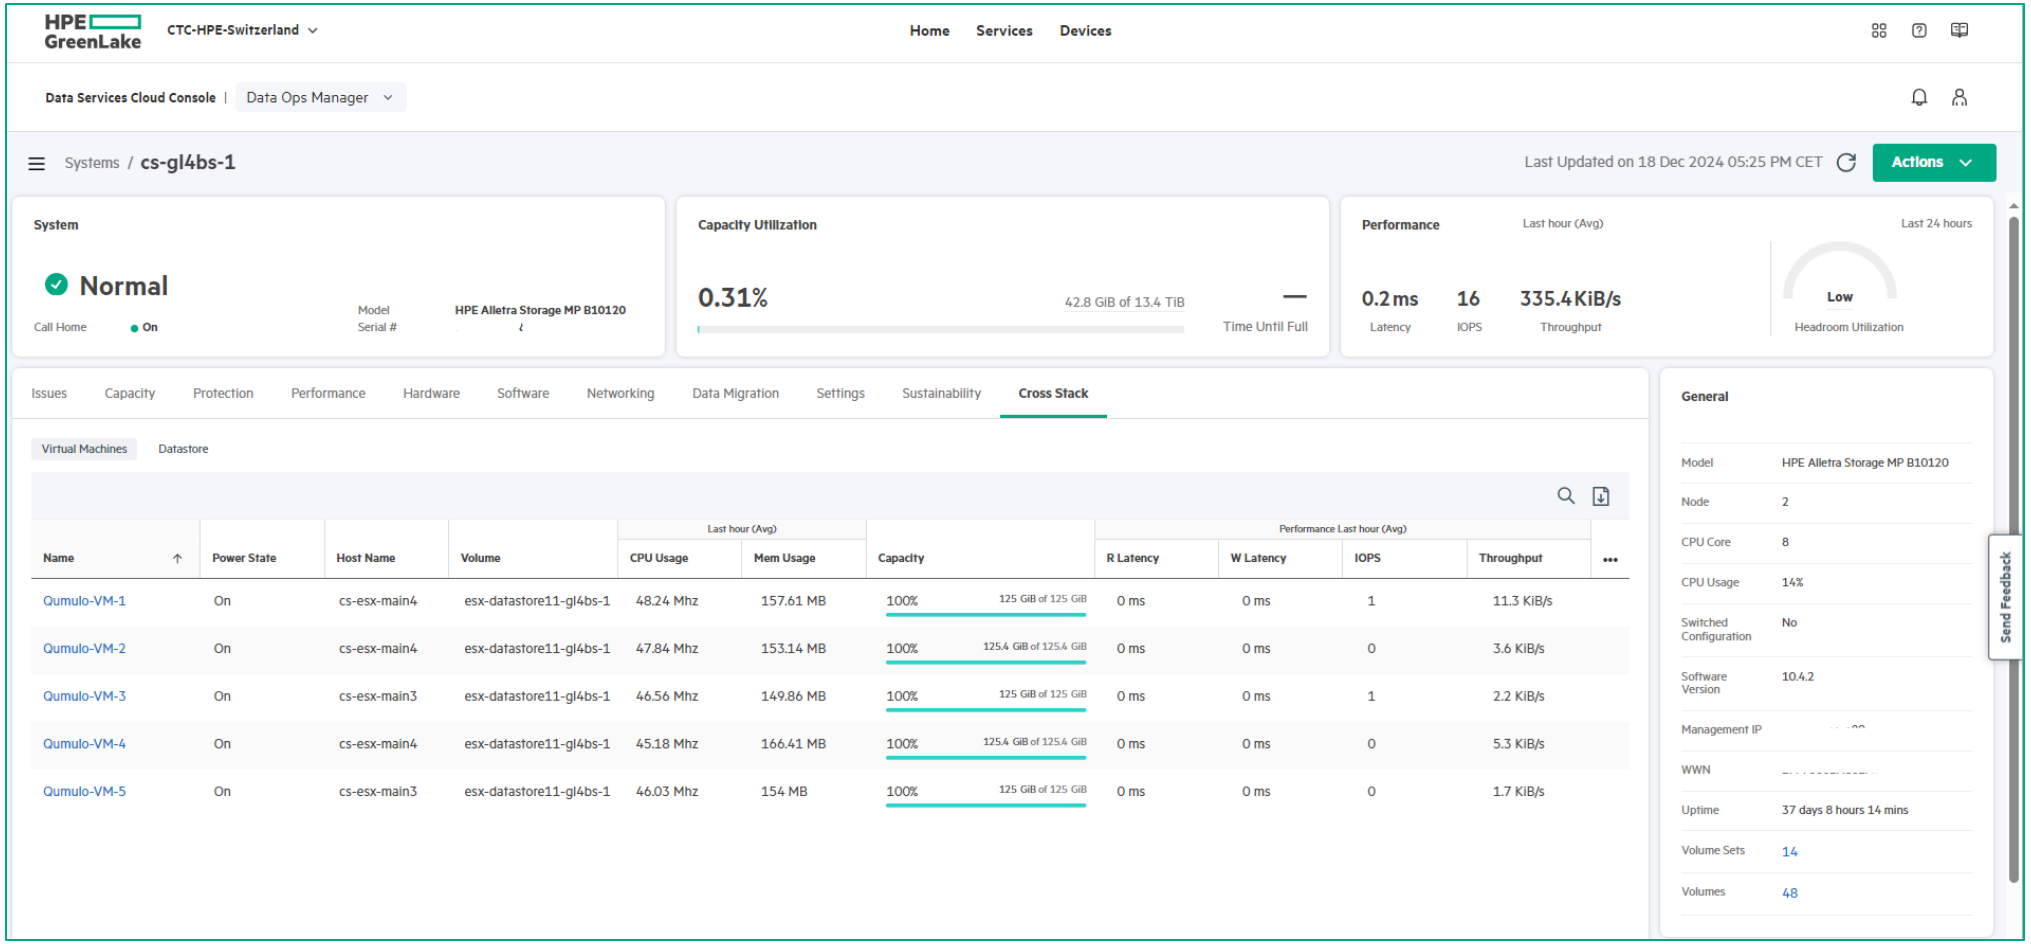

DSCC Workload Drift Analysis

Workload Drift Insights

- Workload drift is a volume-level insight indicating if a workload is running as expected, or if there is a deviation.
- When a change in read and/or write I/O pattern is detected, annotations with insights are overlaid on the throughput trend charts.
- A change in the I/O pattern can indicate an unusual workload, or an issue with the workload that might result in a performance impact to the system and other workloads running on the system in the future.
- A **new I/O Size distribution chart** has been added to allow for detailed exploration of the specific changes to the volume's I/O profile over time.

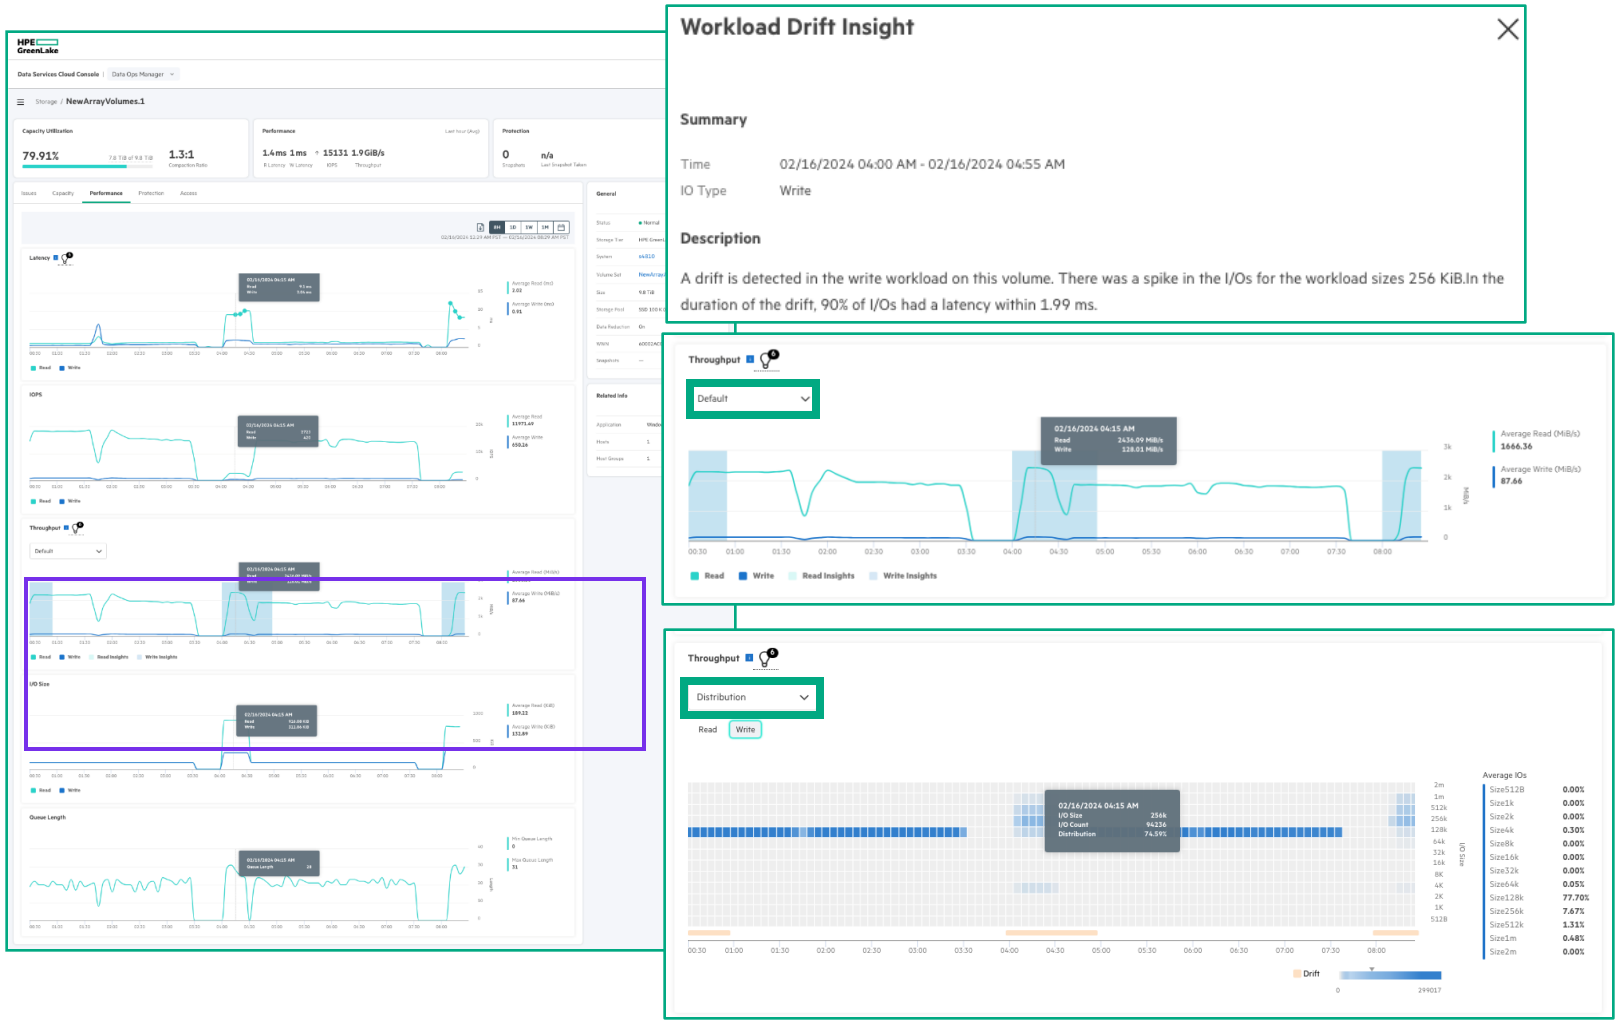

Private Cloud Business Edition with HPE Alletra Storage MP B10000

Full-stack AIOps

Network	2x 10+GbE network switches*
Management	HPE Private Cloud Business Edition cloud console + HPE VM Essentials manager
Virtualization solution	HPE Morpheus VM Essentials Software or VMware vSphere
Compute	HPE ProLiant DL servers - AMD or Intel
Storage	HPE Alletra Storage MP B10000 (iSCSI)
Connectivity	Data Services Connector virtual machine

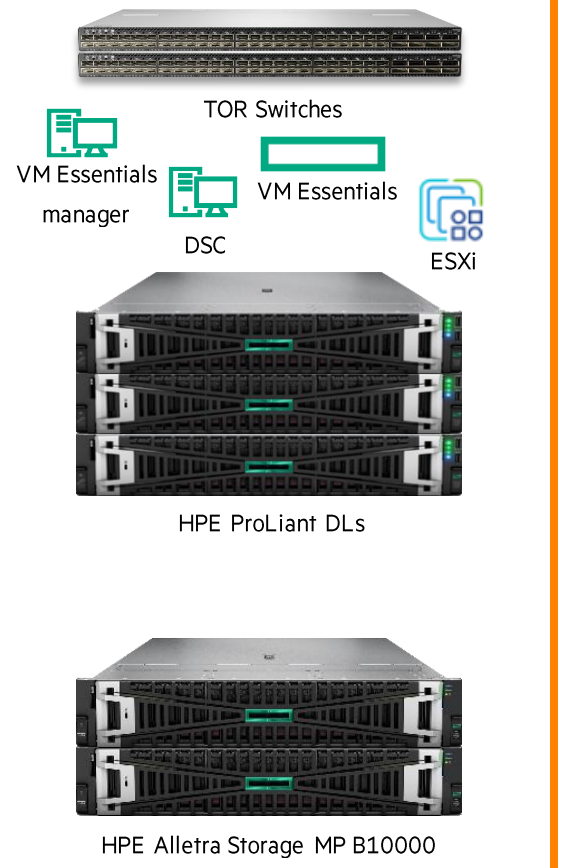

*Default topology, refer to the Private Cloud Business Edition Network Considerations Guide for networking best practices, requirement, and supported topologies

Components and terminology of the stack

Components	Function
HPE Private Cloud Business Edition	Cloud console and on prem virtualization infrastructure
HPE Alletra MP B10000	Block storage array backing the solution
HPE ProLiant DL	Servers providing the compute layer
HPE Morpheus VM Essentials Software	HPE's hypervisor offering
HVM hypervisor	Hypervisor component of VM Essentials
HVM host	Host running HVM hypervisor
HVM cluster	Cluster comprised of HVM hosts
VM Essentials manager	Management VM for VM Essentials
Data Service Connector (DSC) VM	VM that provides secure tunnel to Data Services Cloud Console and Private Cloud Business Edition
Network Automation	Optional feature enabling automated deployment and configuration of Aruba switches for Greenfield deployments
Single Click Update	Stack lifecycle management feature

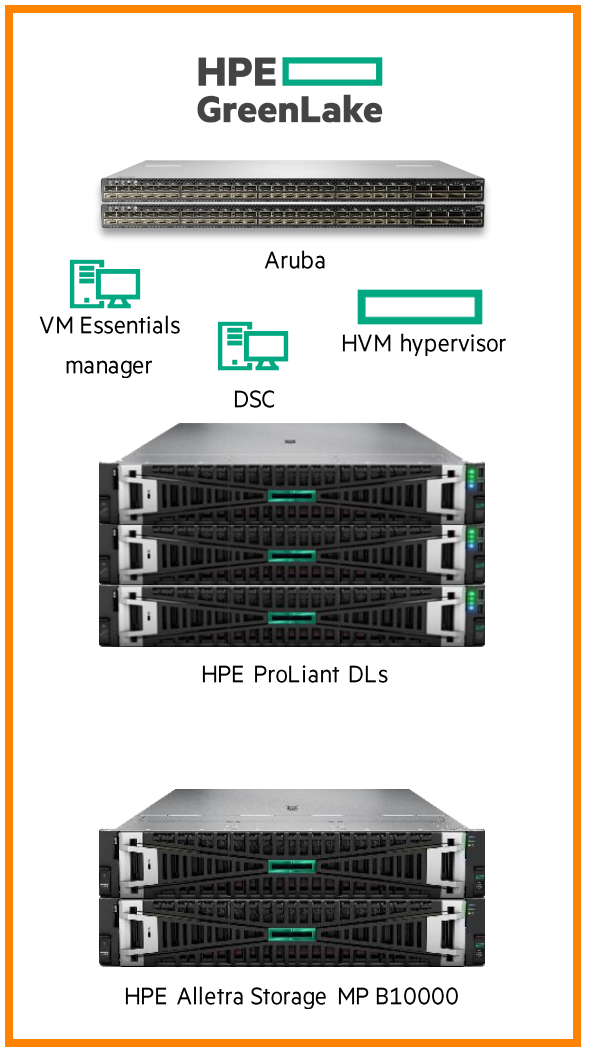

HPE Alletra MP storage plugin

- Acts as a storage abstraction layer between VM Essentials and the HPE Alletra MP array (via WS3 V3 APIs)
- During deployment a logical datastore is created in VM Essentials
  - Each VM virtual disk gets a dedicated volume on the array
  - VM snapshot and backups leverage array volume snapshots
  - Operation is largely transparent in the VM Essentials UI
- Facilitated VM provisioning workflows:
  - Create VM
  - Clone VM
  - Resize VM
  - Migrate VM
  - Snapshot
  - Delete VM
  - Backup and Restore
    - Array snapshots are used as a consistency point while the backup is copied to its destination

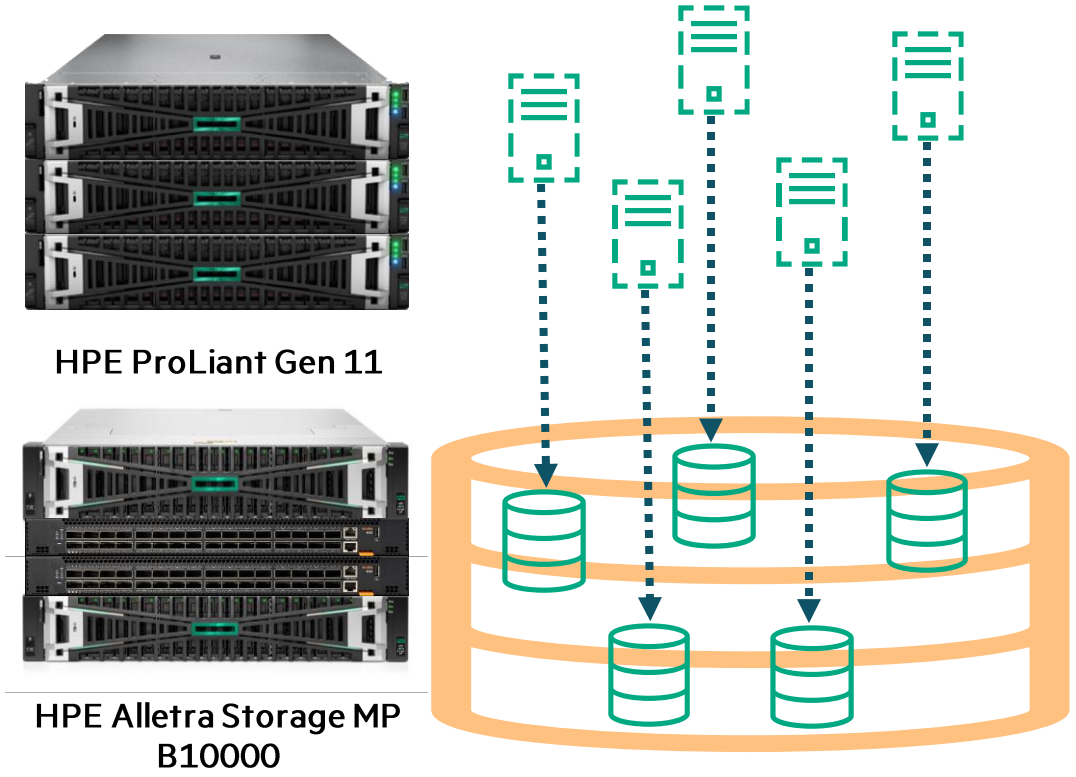

Single Click Update

- Cloud initiated and orchestrated update of the virtualization stack
  - HPE Alletra MP B10000 array OS
  - iLO firmware
  - SPP
  - VM Essentials manager
  - HVM hypervisor and host OS
  - ESXi
  - Switch firmware (if Network Automation feature used)
- None disruptive
- Fully automated
  - If DRS not enabled on a vSphere cluster, manual migration from host to host is necessary
- Software catalogs fully validated
  - Only valid update paths with compatible versions provided
- Cloud AI driven allow/deny list ensures stack won't be allowed to update if a potential issue could be encountered

Thank you