

Ambedded UniVirStor Ceph Storage Appliance

Ceph, is an open-source, distributed storage system

Ambedded UniVirStor (UVS), the Ceph storage appliances are turnkey solutions for scalable software-defined storage powered by Ceph and Arm server. Because of the tightly integrated software and hardware plus Ceph management GUI and Ambedded professional support, customers can deploy and operate Ceph software-defined storage efficiently.



Ambedded UniVirStor (UVS) is a whole software pack that includes Linux, Ceph, and Ceph management GUI (UVS Manager). UniVirStor simplifies the deployment, management, operation, and maintenance of Ceph cluster applications. Ambedded Technology continuously invests in engineering to tune and validate the configurations over computing, memory, network, Linux, and ceph and ensure the storage system serves in optimized reliability and performance.

In addition, Ambedded provides customers with consultancy and technical support during the product life cycle. The products include three years of professional service from planning to deployment and operation.

Ambedded's Ceph appliance solutions, Mars 400 and Mars 500/524 cover the applications of large capacity and high-performance needs. They can serve in separated Ceph clusters or in the same cluster. When they are serving in one cluster, you can add the NVMe SSD pool as the cache tier of the hard drive pool.


Benefits of Ambedded Ceph Appliances

- Management GUI, Professional service, Easy to manage
- Simple to deploy and manage
- Tuned Ceph, Linux, and drivers for the hardware
- High available, Scalable, Resilient
- Block, S3, NAS unified storage
- Cloud-Native storage supports OpenStack and Kubernetes
- Low power consumption and reducing TCO
- Software upgrade without disruption

Ambedded Ceph Total Solution

Professional Support + Consultancy + Software Upgrade + H/W Warranty
UniVirStor (Unified Virtual Storage) UVS Manager (Ceph Management GUI)
<p>Ceph Storage Community version, Tuned by Ambedded Block, FileSystem, Object</p>
<p>Ubuntu, Linux Kernel</p>
Arm based Server Platform Arm 64 microservers

Our Solutions



Mars 400 Ceph Appliance

Ceph cluster with SATA3 HDD/SSD OSD, cost efficient ceph storage for large capacity application.



Mars 412 Ceph Appliance

Ceph cluster with 12x HDD OSD in 2U chassis, easy to reach 240TB per unit, cost efficient for backup & archive applications.



Mars 500 Ceph Appliance

Ceph cluster with NVMe OSD, for the high performance application, easy to start with 4 units with scalable NVMe SSD.



Mars 524 Ceph Appliance

Ceph cluster with NVMe OSD, for the extreme performance and large capacity applications. Able to provide 750TB per unit.

Product Matrix

Model Name	Mars 400	Mars 412	Mars 500	Mars 524
Ceph appliance OSD type	Ceph storage with STAT III OSD (HDD/SSD)	2U 12x HDD OSD Ceph storage	1U 8x All flash Ceph storage with NVME OSD	2U 24x All flash Ceph storage with NVME OSD
UVS Manager Web UI	UVS 2 full version	UVS 3 full version	UVS 3 full version	UVS 3 full version
Software upgrade / Support Service	3 Year	3 Years	3 Years	3 Years
Service Hours	9 am - 6 pm (5*8)	9 am - 6 pm (5*8)	9 am - 6 pm (5*8)	9 am - 6 pm (5*8)
Response Time	in one day	in one day	in one day	in one day
Ticket Included per year	20	20	20	20
Helpdesk Portal & Knowledge Base Access	Yes	Yes	Yes	Yes
Software upgrade package	Yes	Yes	Yes	Yes
Support Channel	Remote & Email	Remote & Email	Remote & Email	Remote & Email
Cluster initial Deployment and Scale-out Service	Free of charge if order with drives	Free of charge if order with drives	Free of charge if order with drives	Free of charge if order with drives
H/W Warranty / Extended Warranty	3 Years	3 Years	3 Years	3 Years
Optional Extended HW/SW support Years Availability	Yes, up to 5 years	Yes, up to 5 years	Yes, up to 5 years	Yes, up to 5 years
Storage server spec				
Server platform	Mars 400 Arm based micro server	Mars 412 AMD EPYC Server	Mars 500 Ampere Altra Arm server	Mars 524 Dual sockets Ampere Altra Arm server
Nodes number per unit	8 nodes per unit (Micro-server architecture)	1 node per unit	1 node per unit	1 node per unit
CPU	Total 32 Cores / appliance, 4 cores 1.2GHz Arm v7 / node	AMD EPYC 7313P 16C/32T Processor @ 3.0GHz	32 cores 1.7 GHz Arm V8 (Optional 64 cores 3.0 GHz)	Dual 64 Core 3.0GHz Arm V8
RAM	Total 32 GB, 4GB/node	96GB	96 GB	256 GB
Network	4x 10Gb	4x10Gb (Dual 2x 10Gb)	2x 25Gb (Optional 2x 100Gb)	4 x 25 Gb (dual cards) (optional 4x 100Gb - dual cards)
Storage Interface	SATA 6Gb	SATA III	NVMe Gen 4	NVMe Gen 4
Max. Num. of OSD / appliance	8	12	8	24
PSU	Redundant Power supply 300W/each	Redundant Power supply 800Watts/each	Redundant Power supply 650Watts/each	Redundant Power supply 1600Watts/each
Power consumption on maximum workload	180 Watts (with full HDD)		210 Watts (with full NVME)	

**Extended software support and hardware warranty are bundled and the maximum is five years

** Spare parts are available in modules for Mars 400 series

Now is the time to discover the Ceph storage with NVMe

Ceph storage could have different OSD disk types for different use scenarios. For IO intensive applications, such as big data analytics, HPC storage, private clouds, edge computing requires real-time sensors, self-monitoring & reporting devices, NVMe ceph cluster provides a high performance, scalable data storage for real-time data capturing, managing, extracting and re-generating.

Mars 500 & Mars 524 NVMe Ceph Appliance tightly integrates Ampere Altra Arm Server and UniVirStor. UniVirStor (UVS) is a software package that includes Linux, Ceph, and Ceph management GUI (UVS Manager).

With the Ampere Altra platform, the 2 NVMe Ceph appliances could meet a low carbon footprint without sacrificing its high IO performance.



Scalable



Reliable



Performant



Simple and Green

Mars 500 and Mars 524 Ceph Storage Appliance



Mars 500 Ceph storage appliance



Mars 524 Ceph storage appliance

Specifications and Major Features

Platform	Mars 500 for NVMe OSD	Mars 524 for NVMe OSD
Processor	Ampere® Altra® 32 Core Arm 64-bit v8.2 Processor @ 1.7GHz (Optional 64 Core @3.0 GHz)	Dual socket Ampere® Altra® 64 Core Arm 64-bit v8.2 Processor @ 3.0GHz (optional 80 Cores @ 3.0 GHz)
Memory	96 GB DDR4 3200MHz, maximum 4TB bytes	256GB DDR4 3200MHz, maximum 4TB bytes
Network Interface	<ul style="list-style-type: none"> 2 x 1GbE management LAN ports for UniVerStor (UVS) Manager 1 x 1GbE IPMI/BMC LAN port Standard 2 x 25 Gbps or optional 2x 100Gbps for Ceph public and in-cluster network 	<ul style="list-style-type: none"> 2 x 1GbE management LAN ports for UniVerStor (UVS) Manager 1 x 1GbE IPMI/BMC LAN port 4 x 25 Gbps or 4x100 Gbps (optional) for Ceph public and in-cluster network
Storage device and interface	<ul style="list-style-type: none"> System drive: 2 x 250GB NVMe M.2 SSD OSD: 8 x 2.5" U.3 NVMe Gen. 4 SSD 	<ul style="list-style-type: none"> System drive: 2 x 480GB SATA SSD OSD: 24 x 2.5" U.3 NVMe Gen. 4 SSD
System management	IPMI & BMC	IPMI & BMC
Power consumption	<ul style="list-style-type: none"> Normal: 260 Watts Maximum: 300 Watts 	-
Power supply	<ul style="list-style-type: none"> 2 x 650W redundant power supply, 80 PLUS Platinum AC input: 100-240V AC, 10-5A, 50-60Hz 2 x 2 meters C13 to C14 power cords 	<ul style="list-style-type: none"> 2 x 1600W redundant power supply, 80 PLUS Platinum AC input: 110/220V AC 12A, 50-60Hz 2 x 2 meters C13 to C14 power cord
Certificate of conformity	<ul style="list-style-type: none"> CE Mark FCC UKCA 	<ul style="list-style-type: none"> CE Mark FCC UKCA
Package include Mechanical	<ul style="list-style-type: none"> 1x Mars 500 Appliance Form factor 1U 19" rack mount Server dimensions: 438 x 43.5 x 660 mm Packaging dimensions: 985 x 588 x 250 mm Slide rail: Included 1x Mounting rail 2x 2 meters C13 to C14 power cords Optional Cable Management Arm 	<ul style="list-style-type: none"> 1x Mars 524 appliance Form factor 2U 19" rack mount Server dimensions: 438 x 87 x 730 mm Packaging dimensions: 1077 x 679 x 300 mm Slide rail: Included
Support and warranty	<ul style="list-style-type: none"> 3 years software update and professional service 3 years hardware warranty Optional extended warranty and service 	<ul style="list-style-type: none"> 3 years software update and technical support 3 years hardware warranty Optional extended warranty and service

Ceph Software

Ceph version	Ambedded tuned Ceph community versions
Management Interface	<ul style="list-style-type: none"> Ambedded Web-based user interface: UniVirStor(UVS) Manager Command Line Interface
Automation	The UniVirStor Ceph management GUI (UVS Manager) automates management functions for day one and day two operations, including cluster deployment, storage service provisioning, monitoring, and software updates.
Access Authentication Control	<ul style="list-style-type: none"> Policy control to limit users to access pool & block device Ability to Integrate with Microsoft Active Directory, lightweight directory access protocol (LDAP), AWS Auth, and KeyStone
Multi-protocol support	A storage platform supports multiple industry-standard storage protocols for block, file and object storage.
Monitoring & Notification	<ul style="list-style-type: none"> UVS, Ceph, and Grafana dashboard Alert sent through email SNMP

Ceph Storage Pool Base

Scale out cluster	<ul style="list-style-type: none"> Grow the cluster to thousands of storage drives No single point of failure Cluster performance is linearly scalable with the number of drives (OSD).
Self-healing and auto rebalancing	<ul style="list-style-type: none"> Balances data distribution throughout the cluster nodes Handles failures without interruption, automatically recovering to the predefined data resilience
High available	Withstand loss of multiple nodes/racks without disrupting service availability or data safety
Data protection	<ul style="list-style-type: none"> Stripping, erasure coding, or replication across nodes with predefined rule Configurable number of replica and erasure coding profile CLAY erasure coding to reduce the data recovery time Renamable bucket type: root, region, zone, data center, room, PDU, raw, rack, chassis
Disk encryption	OSD disk drives can be encrypted to protect data when drives are removed
Data Compression	Inline object compression with target compression ratio
Cache tiering	Support using a faster storage pool as cache tier of a capacity pool

Ceph Deployment and Management

Scale out cluster	<ul style="list-style-type: none"> Grow the cluster to thousands of storage drives No single point of failure Cluster performance is linearly scalable with the number of drives (OSD).
Self-healing and auto rebalancing	<ul style="list-style-type: none"> Balances data distribution throughout the cluster nodes Handles failures without interruption, automatically recovering to the predefined data resilience
High available	Withstand loss of multiple nodes/racks without disrupting service availability or data safety
Data protection	<ul style="list-style-type: none"> Stripping, erasure coding, or replication across nodes with predefined rule Configurable number of replica and erasure coding profile CLAY erasure coding to reduce the data recovery time Renamable bucket type: root, region, zone, data center, room, PDU, raw, rack, chassis
Disk encryption	OSD disk drives can be encrypted to protect data when drives are removed
Data Compression	Inline object compression with target compression ratio
Cache tiering	Support using a faster storage pool as cache tier of a capacity pool

Storage Protocols

RADOS Block Storage (RBD)	
RBD Image Management	<ul style="list-style-type: none"> ▪ Image create, delete ▪ Select pool, name, image size, object size ▪ Image edit
Dynamic volume resizing	Ability to expand Ceph block devices with no downtime
RBD Mirroring for off-site active standby backup	<ul style="list-style-type: none"> ▪ Create mirroring service between multiple clusters or sites ▪ Select pool/image for mirroring ▪ Manage snapshot schedules ▪ Image promote/demote
RBD Snapshot	<ul style="list-style-type: none"> ▪ Snapshots of individual block devices with no downtime or performance impact ▪ Create and manage snapshots of images to retain a history on an image's state snapshot rollback
Copy-on-Write clone	Instant provisioning of tens or hundreds of virtual machine instances from the same image with zero wait time
Thin Provisions	Enable over-provisioning of storage and immediate VM or container instance launch.
iSCSI	<ul style="list-style-type: none"> ▪ Deploy iSCSI gateway on hosts or VMs ▪ Create and manage LUN ▪ ACL: IQN and CHAP authentication ▪ Multi-Path IO support
Shared File System (NAS)	
CephFS	<ul style="list-style-type: none"> ▪ Deploy & manage Metadata servers (MDS) ▪ POSIX compatible file system: kernel and FUSE clients ▪ Multiple CephFS ▪ Support volumes, sub-volume and sub-volume groups ▪ OpenStack Manila ▪ k8s Container Storage Interface (CSI)
Object Storage	
Object API	Amazon S3 and OpenStack Swift compatible APIs
RADOS gateway for Active-Active access	<ul style="list-style-type: none"> ▪ Deploy RADOS multi zone gateways ▪ Multi-site support for sync and access from multiple ceph clusters ▪ RGW User and key management ▪ RGW Pool management, auto-create RGW pools
Write Once Read Many (WORM)	S3 object lock with the read-only capability to store objects using a the WORM mode preventing objects from being deleted or overwritten
Encryption	Server-side encryption and client provided key encryption

Cloud-Native and Virtualization

Platforms	<ul style="list-style-type: none"> ▪ K8s container storage interface (CSI), ▪ OpenStack Cinder, Glance, and Nova, Manila, Swift API ▪ CloudStack libvirt ▪ Proxmox ▪ VMWare iSCSI, NFS ▪ Windows Ceph client driver
-----------	---

Conventional file storage

NFS, SMB, FTP, HTTP	Use CephFS, S3 or block as the back store to export NFS, SMB, FTP, and HTTP through gateways
---------------------	--

Mars 412

Ceph Storage Appliance



Specifications

Server Platform	Mars 412
Processor	Single socket AMD EPYC 7313P 16C/32T Processor @ 3.0GHz
Memory	96GB DDR4 3200MHz, maximum 2TB bytes
Network Interfaces	<ul style="list-style-type: none"> 2 x 1GbE management LAN ports for UniVerStor (UVS) Manager 1 x 1GbE IPMI/BMC LAN port 4 x 10 Gbps Ceph public and in-cluster network
Storage Interfaces	<ul style="list-style-type: none"> System drive: 2 x 240GB SATA SSD OSD: 12 x 3.5" HDD (optional)
System Management	IPMI & BMC
Power Supply	<ul style="list-style-type: none"> 2 x 800W redundant power supply, 80 PLUS Platinum AC input: 110/220Vac/ 12A, 50-60Hz 2 x 1.5 meters C13 to C14 power cord
Mechanical	<ul style="list-style-type: none"> Form factor 2U 19" rack mount Server dimensions: 438 x 87.5 x 660 mm Packaging dimensions: 588 x 268 x 982 mm Slide rail: Included
Support and Warranty	<ul style="list-style-type: none"> 3 years of software update and technical support 3 years hardware warranty Optional extended warranty

Software

Ceph Version	Ambedded tuned Ceph community versions
Management Interface	<ul style="list-style-type: none"> Ambedded Web-based user interface: UniVirStor(UVS) Manager Command Line Interface
Automation	The UniVirStor Ceph management GUI (UVS Manager) automates management functions for day-1 and day-two operations, including cluster deployment, storage service provisioning, monitoring, and software updates.
Access Authentication Control	<ul style="list-style-type: none"> Policy control to limit users to access the pool & block device Ability to Integrate with Microsoft Active Directory, lightweight directory access protocol (LDAP), AWS Auth, and KeyStone
Multi-protocol support	A storage platform supports multiple industry-standard storage protocols for block, file, and object storage.
Monitoring & Notification	<ul style="list-style-type: none"> UVS, Ceph, and Grafana dashboard Alert sent through email SNMP

Ceph Storage Pool Base

Scale-out cluster	<ul style="list-style-type: none"> Grow the cluster to thousands of storage drives No single point of failure Cluster performance is linearly scalable with the number of drives (OSD).
Self-healing and auto rebalancing	<ul style="list-style-type: none"> Balances data distribution throughout the cluster nodes Handles failures without interruption, automatically recovering to the predefined data resilience
Highly available	Withstand the loss of multiple nodes/racks without disrupting service availability or data safety
Data protection	<ul style="list-style-type: none"> Striping, erasure coding, or replication across nodes with predefined rule Configurable number of replica and erasure coding profile CLAY erasure coding to reduce the data recovery time Renamable bucket type: root, region, zone, data center, room, PDU, raw, rack, chassis
Disk encryption	OSD disk drives can be encrypted to protect data when drives are removed
Data Compression	<ul style="list-style-type: none"> Inline object compression with target compression ratio
Cache Tiering	<ul style="list-style-type: none"> Use faster storage pool as cache tier

Ceph Deployment and Management

Software update	<ul style="list-style-type: none"> Automated rolling update without downtime. Data services are available during the update. Updates include Linux, Ceph, and UVS manager
General management	<ul style="list-style-type: none"> Host management: Scan and manage server nodes in the Ceph cluster Deploy the NTP server or use an external NTP server Push NTP client to ceph hosts Log management Notification User management
Ceph daemon services	<ul style="list-style-type: none"> Cluster initialization Deploy and manage MON, OSD, MDS, RGW OSD encryption

Storage Protocols

RADOS Block Storage (RBD)

- RBD Image Management**
- Image create, delete
 - Select pool, name, image size, object size
 - Image edit

Dynamic volume resizing Ability to expand Ceph block devices with no downtime

- RBD Mirroring for off-site active-standby backup**
- Create mirroring service between multiple clusters or sites
 - Select pool/image for mirroring
 - Manage snapshot schedules
 - Image promote/demote

- RBD Snapshot**
- Snapshots of individual block devices with no downtime or performance impact
 - Create and manage snapshots of images to retain a history on an image's state
 - snapshot rollback

Copy-on-Write clone Instant provisioning of tens or hundreds of virtual machine instances from the same image with zero wait time

Thin Provisions Enable over-provisioning of storage and immediate VM or container instance launch.

- Cloud-Native**
- K8s container storage interface (CSI),
 - OpenStack Cinder, Glance, and Nova,
 - CloudStack

- iSCSI**
- Deploy iSCSI gateway on hosts or VMs
 - Create and manage LUN
 - ACL: IQN and CHAP authentication
 - Multi-Path IO support

Shared File System (NAS)

- CephFS**
- Deploy & manage Metadata servers (MDS)
 - POSIX compatible file system: kernel and FUSE clients
 - Multiple CephFS
 - Support volumes, sub-volume and sub-volume groups
 - OpenStack Manila
 - K8s Container Storage Interface (CSI)

Conventional file storage

NFS, SMB, FTP, HTTP Use CephFS, S3 or block as the back store to export NFS, SMB, FTP, and HTTP through gateways

Object Storage

Object API Amazon S3 and OpenStack Swift compatible APIs

- RADOS gateway for Active-Active access**
- Deploy RADOS gateways
 - Multi-site support for sync and access from multiple ceph clusters
 - RGW User and key management
 - RGW Pool management, auto-create RGW pools

Write-Once-Read-Many (WORM) S3 object lock with read-only capability to store objects using a the WORM mode preventing objects from being deleted or overwritten

Encryption Server side encryption and client provided key encryption

Mars 400 Ceph Storage Appliance

Mars 400 Ceph Appliance is designed to meet high capacity cloud-native data storage needs. It utilizes HDD to benefit from low cost per TB.

Mars 400 is based on the Arm microserver technology to reduce your energy cost and carbon footprint. With its Ceph support experiences, your IT team can reduce their resource on storage administration.


Storage Device

- 8x SATA3 HDD/SSD **OSD**
- 8x SATA3 WAL/DB SSD



8x ARM Quad Core Micro-Server

- 4GB DDR4
- 16GB EMMC Flash for system disk
- 5 Gbps LAN
- < 5 Watts power consumption
- Every node can be OSD, MON, MDS, Gateway



Specifications and Major Features

Platform	Mars 400 Ceph Storage Appliance for SATA OSD
Processor	8x Arm server with Arm A72 64-bit Quad-Core Processors @ 1.2 GHz, Total 32 cores.
Memory	4 GB DDR4 per microserver, Total 32GB
Network Interface	<ul style="list-style-type: none"> ▪ Each Arm server has two 2.5Gbps Ethernet ▪ Two in-chassis switches connect internally with 8 microservers ▪ Two 10Gb uplink Ethernet uplink ports on each in-chassis switch. Total 4x 10Gb for uplink. ▪ SPF+ & 10G base-T combo ports
Storage device and interface	Every Arm Server has <ul style="list-style-type: none"> ▪ 16GB eMMC for system ▪ 120 GB M.2 SATA SSD for BlueStore use ▪ OSD disk: one SATA 6Gb 3.5" HDD or 2.5" SSD
System management	One BMC to manage 8 micro-servers in the chassis
Power consumption	<ul style="list-style-type: none"> ▪ Normal: 70 Watts ▪ Maximum: 105 Watts
Power supply	<ul style="list-style-type: none"> ▪ 2x 300W redundant power supply, 80 PLUS Platinum ▪ AC input: 100-240V ac/ 7-3A, 50-60Hz
Certificate of conformity	<ul style="list-style-type: none"> ▪ CE Mark ▪ FCC ▪ UKCA
Package include	<ul style="list-style-type: none"> ▪ 1x Mars 400 Appliance ▪ 1x Mounting rail ▪ 2x 2 meters C13 to C14 power cords ▪ Optional Cable Management Arm
Mechanical	<ul style="list-style-type: none"> ▪ Form factor 1U 19" rack mount ▪ Server dimensions: 440 x 42 x 746 mm ▪ Packaging dimensions: 930 x 580 x 215 mm ▪ Slide rail: Included
Support and warranty	<ul style="list-style-type: none"> ▪ 3 years software update and professional service ▪ 3 years hardware warranty ▪ Optional extended warranty and software service

Ceph Software

Ceph version	Ambedded tuned Ceph community versions
Management Interface	<ul style="list-style-type: none"> Ambedded Web-based user interface: UniVirStor(UVS) Manager Command Line Interface
Automation	The UniVirStor Ceph management GUI (UVS Manager) automates management functions for day one and day two operations, including cluster deployment, storage service provisioning, monitoring, and software updates.
Access Authentication Control	<ul style="list-style-type: none"> Policy control to limit users to access pool & block device Ability to Integrate with Microsoft Active Directory, lightweight directory access protocol (LDAP), AWS Auth, and KeyStone
Multi-protocol support	A storage platform supports multiple industry-standard storage protocols for block, file and object storage.
Monitoring & Notification	<ul style="list-style-type: none"> UVS, Ceph, and Grafana dashboard Alert sent through email SNMP

Ceph Storage Pool Base

Scale out cluster	<ul style="list-style-type: none"> Grow the cluster to thousands of storage drives No single point of failure Cluster performance is linearly scalable with the number of drives (OSD).
Self-healing and auto rebalancing	<ul style="list-style-type: none"> Balances data distribution throughout the cluster nodes Handles failures without interruption, automatically recovering to the predefined data resilience
High available	Withstand loss of multiple nodes/racks without disrupting service availability or data safety
Data protection	<ul style="list-style-type: none"> Stripping, erasure coding, or replication across nodes with predefined rule Configurable number of replica and erasure coding profile CLAY erasure coding to reduce the data recovery time Renamable bucket type: root, region, zone, data center, room, PDU, raw, rack, chassis
Disk encryption	OSD disk drives can be encrypted to protect data when drives are removed
Data Compression	Inline object compression with target compression ratio
Cache tiering	Support using a faster storage pool as cache tier of a capacity pool

Ceph Deployment and Management

Software update	<ul style="list-style-type: none"> Automated rolling update without downtime. Data services are available during the update. Updates include Linux, Ceph, and UVS manager
General management	<ul style="list-style-type: none"> Host management: Scan and manage server nodes in the Ceph cluster Deploy the NTP server or use an external NTP server Push NTP client to ceph hosts Log management Notification User management
Ceph daemon service	<ul style="list-style-type: none"> Cluster initialization Deploy and manage MON, OSD, MDS, RGW OSD encryption

Storage Protocols

RADOS Block Storage (RBD)	
RBD Image Management	<ul style="list-style-type: none"> ▪ Image create, delete ▪ Select pool, name, image size, object size ▪ Image edit
Dynamic volume resizing	Ability to expand Ceph block devices with no downtime
RBD Mirroring for off-site active standby backup	<ul style="list-style-type: none"> ▪ Create mirroring service between multiple clusters or sites ▪ Select pool/image for mirroring ▪ Manage snapshot schedules ▪ Image promote/demote
RBD Snapshot	<ul style="list-style-type: none"> ▪ Snapshots of individual block devices with no downtime or performance impact ▪ Create and manage snapshots of images to retain a history on an image's state ▪ snapshot rollback
Copy-on-Write clone	Instant provisioning of tens or hundreds of virtual machine instances from the same image with zero wait time
Thin Provisions	Enable over-provisioning of storage and immediate VM or container instance launch.
iSCSI	<ul style="list-style-type: none"> ▪ Deploy iSCSI gateway on hosts or VMs ▪ Create and manage LUN ▪ ACL: IQN and CHAP authentication ▪ Multi-Path IO support
Shared File System(NAS)	
CephFS	<ul style="list-style-type: none"> ▪ Deploy & manage Metadata servers (MDS) ▪ POSIX compatible file system: kernel and FUSE clients ▪ Multiple CephFS ▪ Support volumes, sub-volume and sub-volume groups ▪ OpenStack Manila ▪ k8s Container Storage Interface (CSI)
Object Storage	
Object API	Amazon S3 and OpenStack Swift compatible APIs
RADOS gateway for Active-Active access	<ul style="list-style-type: none"> ▪ Deploy RADOS multi zone gateways ▪ Multi-site support for sync and access from multiple ceph clusters ▪ RGW User and key management ▪ RGW Pool management, auto-create RGW pools
Write Once Read Many (WORM)	S3 object lock with the read-only capability to store objects using a the WORM mode preventing objects from being deleted or overwritten
Encryption	Server-side encryption and client provided key encryption
Cloud-Native and Virtualization	
Platforms	<ul style="list-style-type: none"> ▪ K8s container storage interface (CSI), ▪ OpenStack Cinder, Glance, and Nova, Manila, Swift API ▪ CloudStack libvirt ▪ Proxmox ▪ VMWare iSCSI, NFS ▪ Windows Ceph client driver
Conventional file storage	
NFS, SMB, FTP, HTTP	Use CephFS, S3 or block as the back store to export NFS, SMB, FTP, and HTTP through gateways

Ambedded Technology Co. LTD.

www.ambedded.com.tw
service@ambedded.com.tw

Ambedded reserve the right to make changes of specification and product descriptions at any time without prior notice.
Last update: 05/07/2023