

Ceph Software Defined Storage Appliance

Unified distributed data storage cluster with self-healing, auto-balancing and no single point of failure
Lowest power consumption in the industry: 70% power saving

Infinite Scale Out with Simplicity

Ground Breaking Ceph on ARM Microserver Cluster



File System

SDS

Block

Open Stack

Object Storage

S3

Mars 200



Mars 400



Best of
Interop
2016 STORAGE WINNER



BEST CHOICE
AWARD
Golden
2017

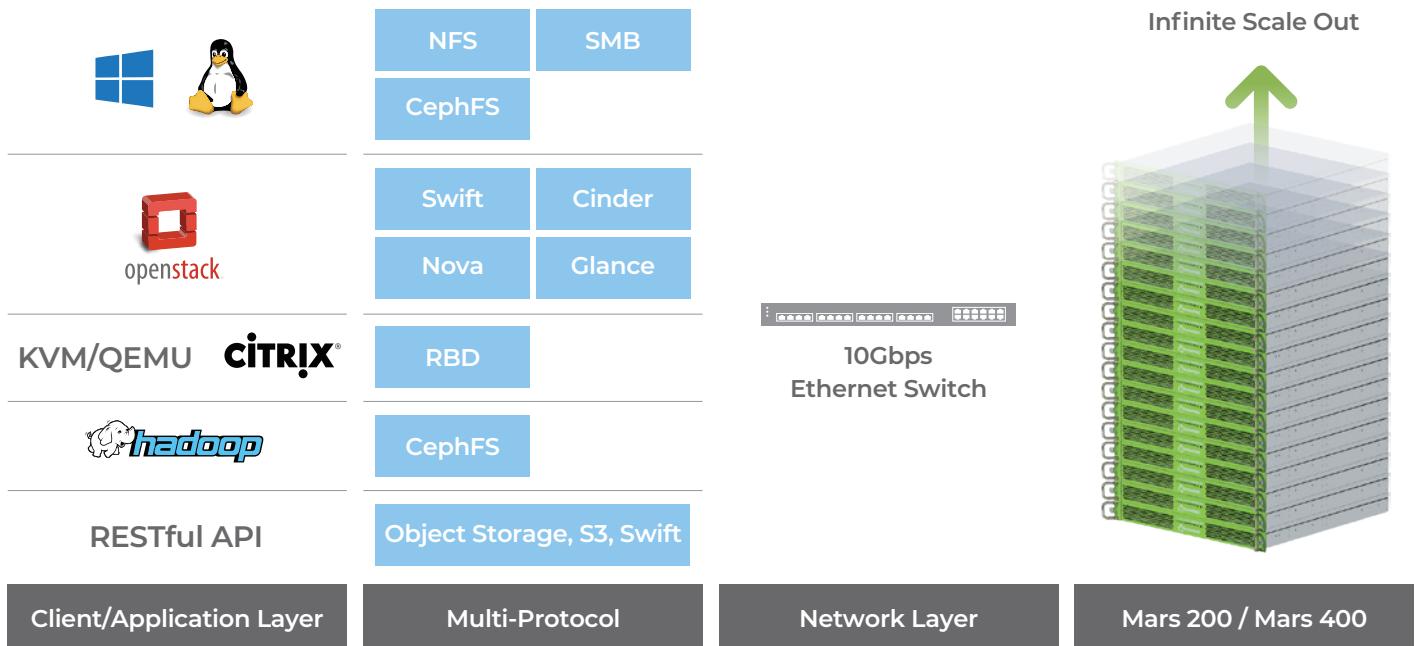
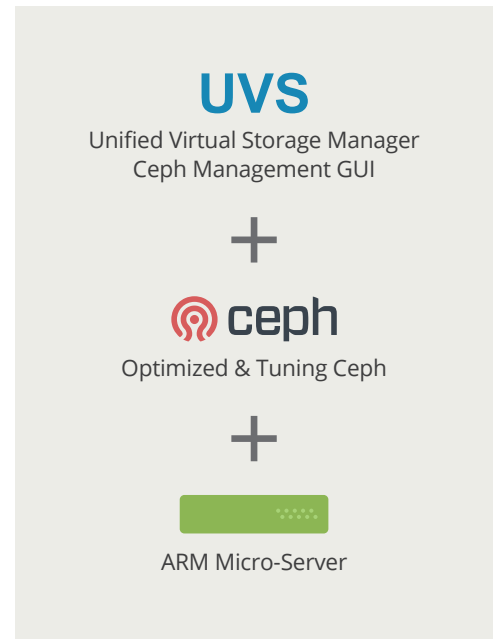
Mars 200/400 Infinite Scale Out, Unified Virtual Storage

Mars 200/400 Features

- Effortless, Scalable and Auto-Configurable Ceph Appliance
- Easy to use web based Ceph user interface
- Performance and capacity scale out on demand
- Resilient survival of multiple rack/chassis/host/OSD failures
- Self-healing data protection
- Unified system supports object storage, SAN and NAS on a single device
- Amazon S3 and OpenStack back-end storage
- Configurable on all SSD, hybrid and full HDD
- ARM based Micro-Server architecture minimizes failures
- Consumes less than 100/105 Watts of power (Mars200/Mars 400) : 70% power less than competitors

Use Case & Applications

- Big Data Analysis, Machine Learning
- Hadoop compatible for Telecom and Energy Industries
- Cloud Storage Service, backend storage for OpenStack & Kubernetes
- Edge Data Center for IOT applications such as sensor data aggregation
- Massive Data Backup
- Database as a Service



Simplified Design with High Availability

Intelligent Data Protection


Data replication to diverse chassis and racks to minimize the impact of failure (via the CRUSH rule on UVS software). Self-healing Micro-Server architecture.

Minimizes the scale and impact of hardware failure

Each ARM Micro-Server connects to its dedicated drive reducing the impact of failure by 90% compared to an x86 based storage system.

Hot-Swappable Hardware

Micro-Server, switch, HDD, SSD and power supplies are all hot-swappable modules. Switches and power supplies are also redundant.



Basic Configuration

UVS – Unified Virtual Storage Management


Web-based Ceph Management GUI

Easy to Configure, Deploy, Manage, Monitor, Automate

Data replication and protection

- Supports Replication and Erasure Code data protection methods. Support up to 10 x data replication.
- Erasure Code set in efficient, assigned storage space.
- Data is evenly distributed among storage nodes.

(1) Replica



(2) Erasure code

$K+M \leq \text{OSD numbers}$ (no limitation on M value)
Flexible to set up fault-tolerance ratio and overhead capacity



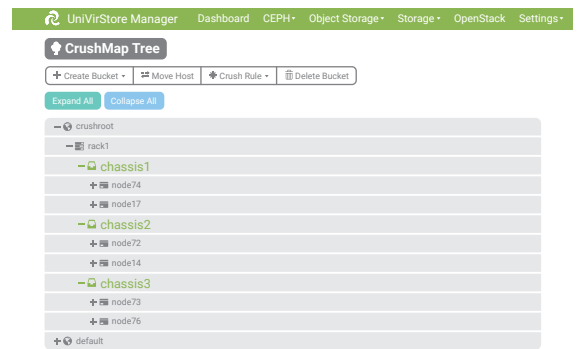
Real-Time Self-Healing and Fault-Tolerance

When any drive or Micro-Server fails, MARS 200/400 detects the failure and simultaneously regenerates the lost data per the CRUSH rule.


Auto-detection and self-healing ; back to data safe level

The CRUSH rule reduces and de-centralizes risk

The CRUSH algorithm distributes data replication/ Erasure code across dispersed racks, chassis and data centers.



Define the failure domain through CRUSH map on UVS manager




Scale out and Automatic Load Balancing

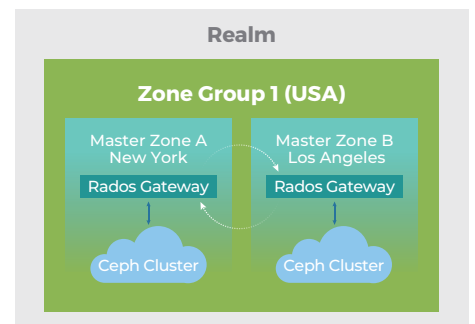
- Mars 200 & Mars 400 scale out capacity on demand without service interruption.
- Limitless linear performance and capacity scaling.
- All storage nodes automatically re-balance whenever there is a change in service.

Object Storage Multiple-Site, Active-Active Disaster Recovery

Shorten recovery point objectives (RPO) and Ceph Cluster recovery time objectives (RTO).



Multi-sites active-active support on RadosGW



Mars 400 & Mars 200 CEPH Storage Appliances

	Mars 400	Mars 200
Form Factor	1U Rack mount with 437.8 mm (W) x 43.5mm (H) x 741.2 mm (L) 1U Rack mount with 17.2" (W) x 1.7" (H) x 29.2" (L)	
Micro-Server	SOC	
	8 x ARM 64-bit Cortex-A72 Dual Core 1.2GHz	8 x ARM 32-bit Cortex A9 v7 Dual Core 1.6 GHz
	Memory	
	4G Bytes DDR4	2G Bytes DDR3
	<ul style="list-style-type: none"> Network Interface: 2 x 2.5Gbps Ethernet Storage Interface: 2 x SATA 3.0 (6 Gbps) Storage: 8GB flash memory for operating system, Ceph software and UVS manager 	
Network	<ul style="list-style-type: none"> Redundant Dual Hot-Swappable switches (active/active) 4 x 10Gbps uplink, for client and scale-out Support SFP+ or 10G baseT media with auto-media detection 1 x 100Mbps out of band management port (BMC) 	
Baseboard Management Controller (BMC)	1 x 100Mbps Ethernet out-of-band port Functions: <ul style="list-style-type: none"> Micro-server Console over Ethernet Reset specified Micro-Server Control Micro-Server power ON/OFF Control system power ON/OFF Reset In-chassis switch UID LED control 	
Storage Bay (HDD/SSD)	<ul style="list-style-type: none"> 8 x top accessible hot-swappable SATA3 storage bay (3.5"HDD or 2.5" SSD/HDD) Each Micro-Server has a 32GB SATA 3 M.2 SSD slot for Ceph WAL & DB 	
Front Panel	<ul style="list-style-type: none"> 8 green LED for Micro-Server status UID LED Power ON/OFF switch for power supply HDD backplane with: 8x LEDs for locating HDD positions 	
Power Consumption	Max. 105 Watts (exclude 8 x SSD/HDD)	Max. 100 Watts (exclude 8 x SSD/HDD)
Accessories	<ul style="list-style-type: none"> AC input power cord with IEC C14 inlet plug Slide rail kit Cable management arm (optional) 	
Power Supply	Dual 300 Watt 80 Plus Silver Redundant Power Supplies (active/active)	
Safety	CE/FCC Class A	

1U 8 nodes ARM Micro-Server Cluster


Storage Device

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

8x ARM Dual Core Micro-Server

- 2GB DDR3 (Mars 200)/4GB DDR4 (Mars 400)
- 8G Bytes Flash: System disk
- 5 Gbps LAN

- < 5 Watts power consumption
- Every node can be OSD, MON, MDS, Gateway



ARM-based Micro-Server Architecture
Energy-Saving Distributed Storage Server

Unified Virtual Storage Manager (UVS) Features

Cluster & NTP Server Deployment

- Deploy the first Monitor and OSD to bring up Ceph cluster from scratch.
- Setup NTP server: Ceph allows very small clock skew between nodes.
- NTP options may create an NTP server on MON node or use an existing NTP server.
- A single click can push the NTP setting to each Ceph node.

Dashboard

The dashboard provides graphical cluster information.

- Ceph cluster status
- Warning and error messages
- OSD and MON status
- Placement Group health status
- Cluster capacity usage
- Throughput metrics

MON/OSD Management

- MON create, restart and reboot
- OSD create, restart, reboot and remove
- Add multiple OSDs
- MON and OSD network and health status
- OSD disk SMART information

The screenshot shows two main sections: 'MON' and 'OSD'. The 'MON' section has a '+ Create New MON' button and 'Service Restart' and 'Node Reboot' buttons. Below it is a table with columns: Rank, Node Name, IP Address, and Status. The 'OSD' section has a '+ Create New OSD' button and 'Service Restart', 'Node Reboot', and 'Remove OSD' buttons. Below it is a table with columns: Rank, Class, Node Name, IP Address, Status, Data SMART, and Journal SMART.

Rank	Node Name	IP Address	Status
mon.0	node111-246cf	↑192.168.1.111	Up
mon.1	node121-82650	↑192.168.1.121	Up
mon.2	node131-22e78	↑192.168.1.131	Up

Rank	Class	Node Name	IP Address	Status	Data SMART	Journal SMART
osd.0	ssd	node112-b555c	↑192.168.1.112	Up	Info	Info
osd.1	ssd	node113-5d7ff	↑192.168.1.113	Up	Info	Info
osd.2	ssd	node114-98e24	↑192.168.1.114	Up	Info	Info
osd.3	ssd	node115-fcb0e	↑192.168.1.115	Up	Info	Info
osd.4	ssd	node116-1547e	↑192.168.1.116	Up	Info	Info
osd.5	ssd	node117-44414	↑192.168.1.117	Up	Info	Info
osd.6	ssd	node118-4e3fb	↑192.168.1.118	Up	Info	Info
osd.7	ssd	node122-098d1	↑192.168.1.122	Up	Info	Info
osd.8	ssd	node123-ef610	↑192.168.1.123	Up	Info	Info
osd.9	ssd	node124-05008	↑192.168.1.124	Up	Info	Info

Pool Management & Cache Tiering

- Pool create/delete
- Pool configuration: Name, Replica/Erasure Code, Quota, CRUSH Rule, Placement Group
- Cache tiering: With different speed pools, a faster pool can be set as the cache tier of a slower pool.

CRUSH Map Configuration

Ceph uses CRUSH algorithm to distribute and store replicated data and erasure coding chunks to the configurable failure domain. CRUSH requires a map to avoid single point of failure, performance bottleneck and scalability limitations. UVS enables configuration of the CRUSH map and rule sets.

- Create/Delete bucket: root, rack, chassis
- Move host: Assign hosts to their chassis
- List and create CRUSH Rules
- Graphical CRASH map

The screenshot shows the 'CrushMap Tree' interface. It includes a legend with categories: ROOT (blue), RACK (yellow), CHASSIS (green), and HOST (orange). The tree structure shows a hierarchy: default (ROOT) -> myrack (RACK) -> zone1 (CHASSIS) -> node52-df658 (HOST) -> osd.0 (OSD), and so on.

RBD Image Management & Snapshot

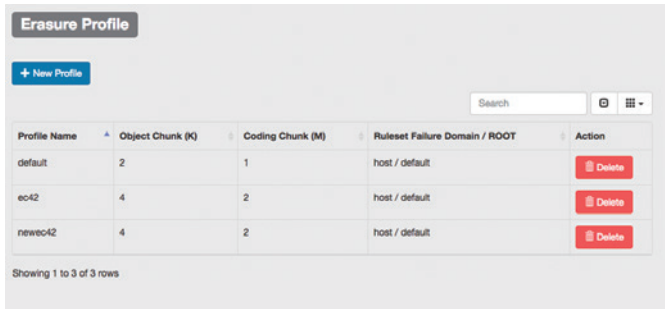
- Create and deleting image
- Assign image object size
- Size and Resize image
- Snapshot, clone and flatten images
- List images with their name, image size, object size and watchers (users).

The screenshot shows the 'Ceph Images' section with a '+ Create Image' button. Below it is a table with columns: Pool Name, Image Name, Image Size, Object Size, and Action. The 'Action' column contains buttons for 'Resize', 'Delete', 'Snapshot', and 'Watchers'.

Pool Name	Image Name	Image Size	Object Size	Action
pool00d	image1	100 GiB	4 MiB	Resize Delete Snapshot Watchers
	image10	100 GiB	4 MiB	Resize Delete Snapshot Watchers
	image11	100 GiB	4 MiB	Resize Delete Snapshot Watchers
	image12	100 GiB	4 MiB	Resize Delete Snapshot Watchers
	image13	100 GiB	4 MiB	Resize Delete Snapshot Watchers
	image14	100 GiB	4 MiB	Resize Delete Snapshot Watchers
	image2	100 GiB	4 MiB	Resize Delete Snapshot Watchers
	image3	100 GiB	4 MiB	Resize Delete Snapshot Watchers
	image4	100 GiB	4 MiB	Resize Delete Snapshot Watchers
	image5	100 GiB	4 MiB	Resize Delete Snapshot Watchers
	image6	100 GiB	4 MiB	Resize Delete Snapshot Watchers
	image7	100 GiB	4 MiB	Resize Delete Snapshot Watchers
	image8	100 GiB	4 MiB	Resize Delete Snapshot Watchers
	image9	100 GiB	4 MiB	Resize Delete Snapshot Watchers

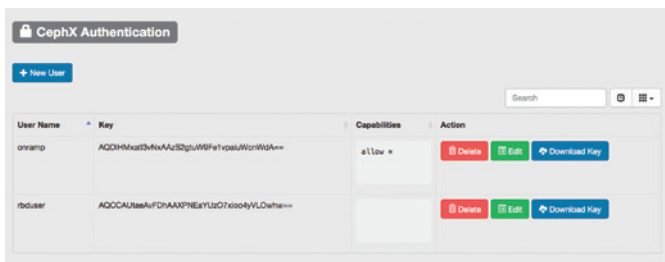
Erasure Code Profile Management

Before creating an erasure code pool, Administrators create an Erasure Code profile with specified object Data Chunk (K) and Coding Chunk (M) values, and a failure domain. UVS makes this quite straightforward.



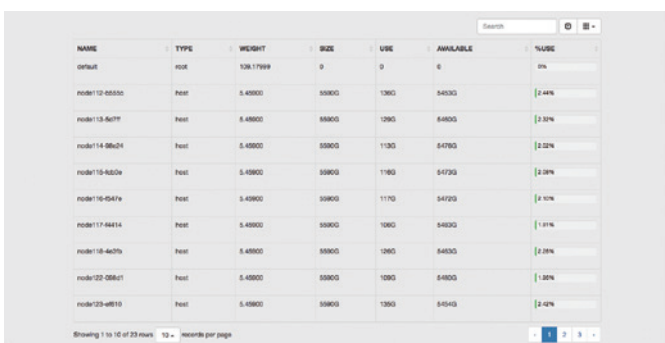
Client User Access Control

Ceph requires authentication and authorization via username / keyring. UVS manages user access and creates the associated keyring, which administrators can download after creation.



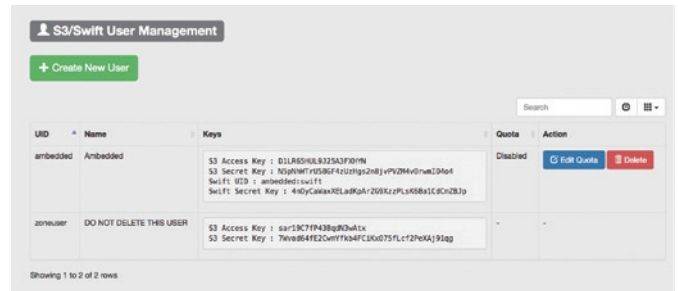
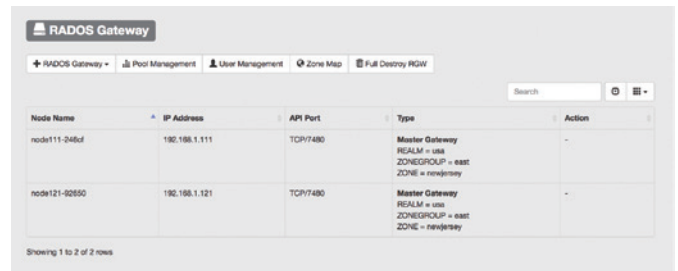
Usage Detail

Usage detail lists the size, weight, use percentage and availability of each root, rack, chassis and host/disk.



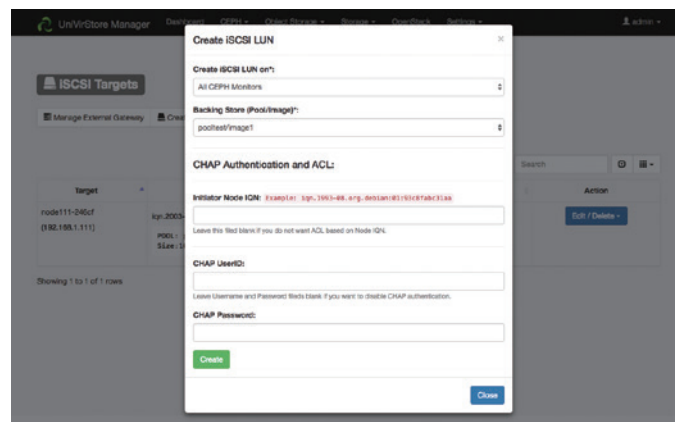
Object Storage

UVS manager supports the use of object storage. Applications can access the object storage through Amazon S3 and OpenStack Swift compatible API through the RADOS gateway.



iSCSI

This feature helps to create iSCSI gateways on external servers or internal MON nodes and manage iSCSI LUNs with CHAP and ACL authentication.



And more with UVS Manager

UVS manager also supports keyring & ceph.conf file generation for OpenStack, Audit logs, Notification/email alert, UVS user management, On-fly firmware update....etc.



Ambedded Technology Headquarter

Address: No.18, Siyuan St., Zhongzheng Dist., Taipei City 10087, Taiwan
 TEL: +886-2-23650500
 Contact: service@ambedded.com.tw
 www.ambedded.com

Ambedded USA

Contact : USSales@Ambedded.com

