

# 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: up to 70% power saving

## Infinite Scale Out with Simplicity

Ground Breaking Ceph on ARM Microserver Cluster



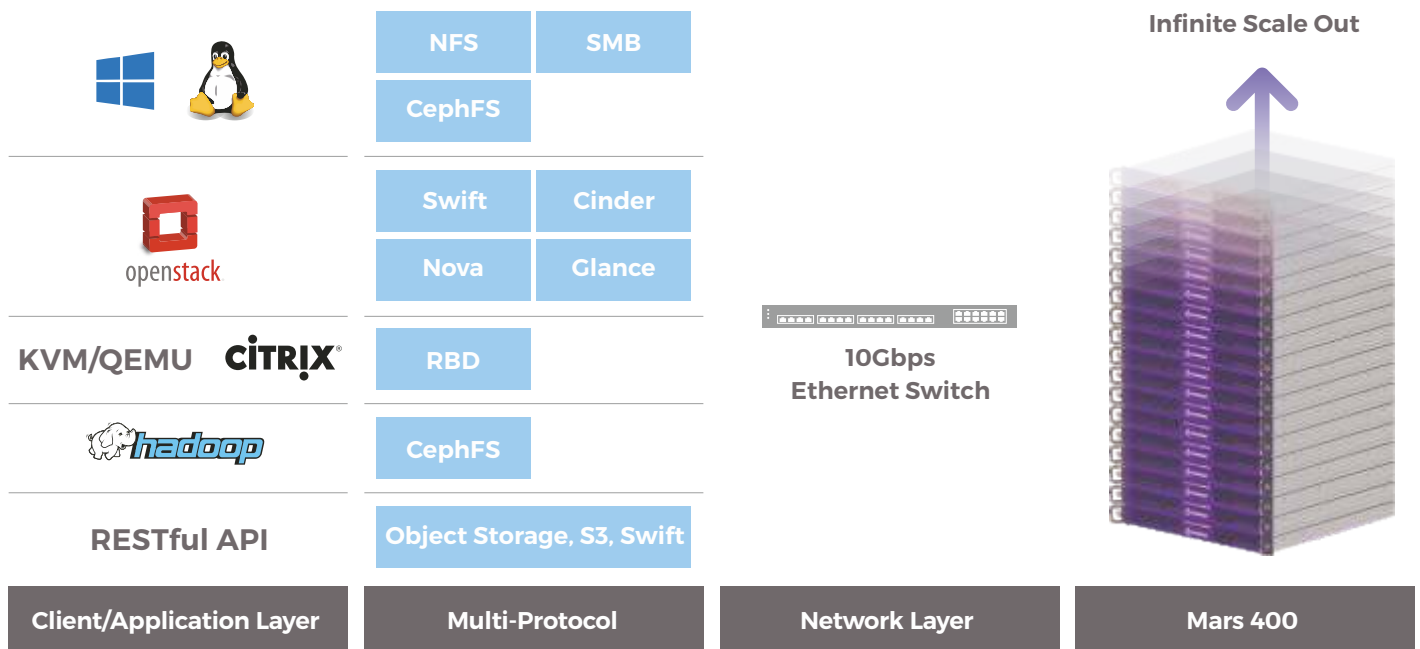
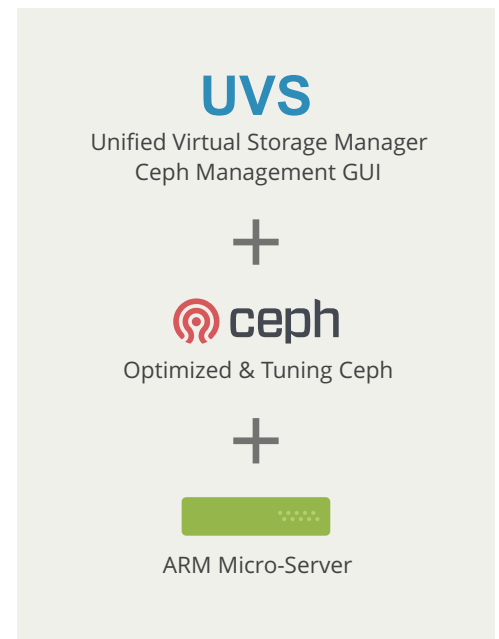
# Mars 400 Infinite Scale Out, Unified Virtual Storage

## Mars 400 Features

- Effortless, Scalable and Auto-Configurable Ceph Appliance
- Easy to use web-based Ceph management 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
- All-SSD, SATA HDD and hybrid configurable
- ARM-based Micro-Server architecture reduces failure risk
- Consumes less than 105 Watts of power: up to 70% power less than competitors

## Use Cases & Applications

- Big Data Analysis, Machine Learning
- Hadoop compatible storage 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 different 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 a conventional 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 amongst 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



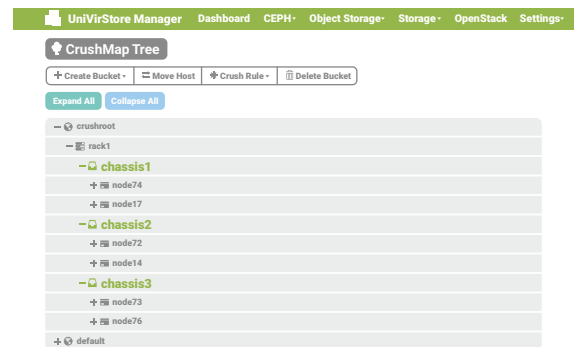
### Real-Time Self-Healing and Fault-Tolerance

When a drive or a Micro-Server fails, Mars 400 detects the failure and simultaneously regenerates the lost data as 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




Cluster -Before      Cluster - After scale-out  
Capacity and performance scale-out linearly

### Scale out and Automatic Load Balancing

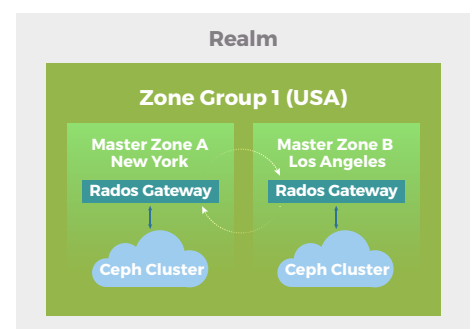
- Mars 400 scale out capacity on demand without service interruption.
- Limitless linear performance and capacity scaling.
- All storage nodes automatically re-balance whenever there is a change in service.

### Object Storage Multi-Site, Active-Active Disaster Recovery

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



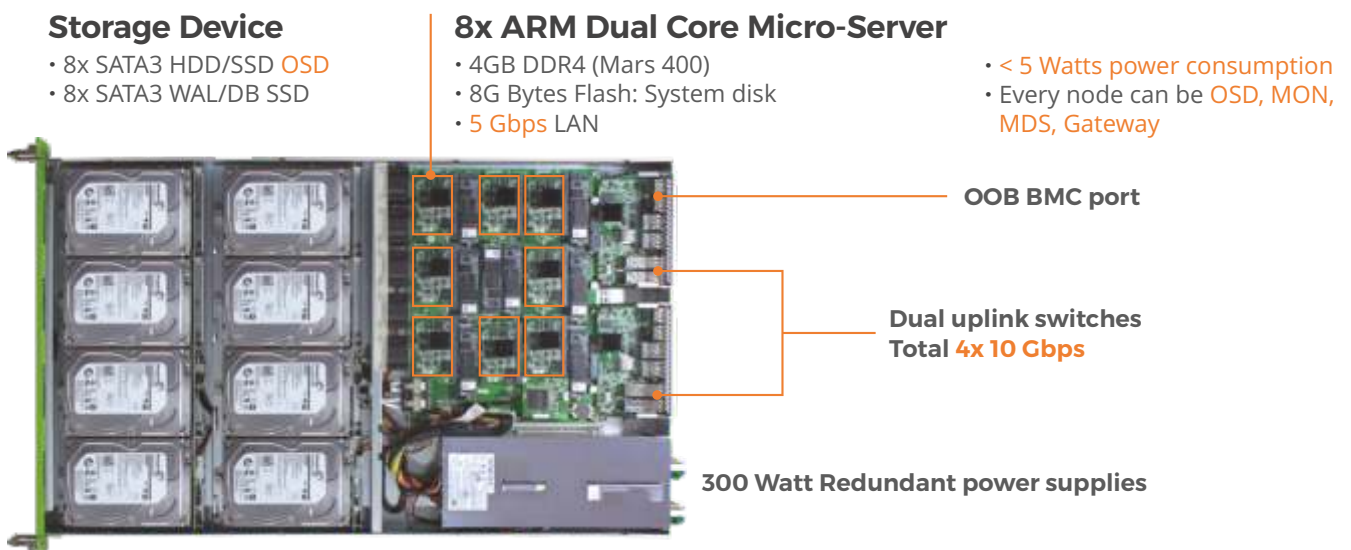
Multi-sites active-active support on RadosGW



# Mars 400 CEPH Storage Appliances

Mars 400	
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
	Memory
	4G Bytes DDR4
	<ul style="list-style-type: none"> <li>• Network Interface: 2 x 2.5Gbps Etherne</li> <li>• Storage Interface: 2 x SATA 3.0 (6 Gbps)</li> <li>• Storage: 8GB flash memory for operating system, Ceph software and UVS manager</li> </ul>
Network	<ul style="list-style-type: none"> <li>• Redundant Dual Hot-Swappable switches (active/active)</li> <li>- 4 x 10Gps uplink, for client and scale-out</li> <li>- Support SFP+ or 10G baseT media with auto-media detection</li> <li>• 1 x 100Mbps out of band management port (BMC)</li> </ul>
Baseboard Management Controller (BMC)	1 x 100Mbps Ethernet out-of-band port Functions: <ul style="list-style-type: none"> <li>- Micro-server Console over Ethernet</li> <li>- Reset specified Micro-Server</li> <li>- Control Micro-Server power ON/OFF</li> <li>- Control system power ON/OFF</li> <li>- Reset In-chassis switch</li> <li>- UID LED control</li> </ul>
Storage Bay (HDD/SSD)	<ul style="list-style-type: none"> <li>• 8 x top accessible hot-swappable SATA3 storage bay (3.5"HDD or 2.5" SSD/HDD)</li> <li>• Each Micro-Server has a 64GB SATA 3 M.2 SSD slot for Ceph WAL &amp; DB, configurable up to 512GB</li> </ul>
Front Panel	<ul style="list-style-type: none"> <li>- 8 green LED for Micro-Server status</li> <li>- UID LED</li> <li>- Power ON/OFF switch for power supply</li> <li>- HDD backplane with: 8x LEDs for locating HDD positions</li> </ul>
Power Consumption	Max. 105 Watts (exclude 8 x SSD/HDD)
Accessories	<ul style="list-style-type: none"> <li>• AC input power cord with IEC C14 inlet plug</li> <li>• Slide rail kit</li> <li>• Cable management arm (optional)</li> </ul>
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



**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 allow the creation of a NTP server on MON node/s or the use of an existing NTP server.
- With a single click, push NTP settings to all Ceph nodes.

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



## CRUSH Map Configuration

Ceph uses the CRUSH algorithm to distribute and store replicated data and erasure-coded chunks to the configurable failure domain. CRUSH requires a map to avoid single point of failures, performance bottlenecks 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



## RBD Image Management & Snapshot

- Create and delete images
- 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).



## Pool Management & Cache Tiering

- Pool creation /deletion
- 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.

