

Mars 400 Ceph Solution For Different Applications

- Object storage S3
- Block storage : RBD, ISCSI
- File sharing : CephFS, NFS, SAMBA

Access Object Storage with Ceph S3



10Gbps Network



Mars 400 ceph cluster

- 3xMON
- 3xMGR
collocated with
MON
- OSDs (with HDD
or SSD)
- backup RGW
collocated with
MON

- Use S3 for private cloud on premise
- Use Swift for the object storage as Openstack backend storage
- Ceph cluster on Mars 400 appliance, with external RGW server to provide S3

Suitable Application

- Medical (patient record, prescription data)
- Biotech (x-ray images, bio data, DNA/Gene)
- Sensitive data, but need quick queries (insurance)
- E-commerce

Access Block Storage with Ceph RBD

Linux client
VM or
Bare metal



via Cinder API



CSI

10Gbps Network



Ceph RBD (block)

- 3xMON
- 3xMGR collocated with MON
- OSDs (with HDD or SSD)

Mars 400 ceph cluster

Per application performance requirement, Mars 400 can use 3.5" HDD or 2.5" SSD as OSD disk

- Performance oriented : SSD as OSD
- Capacity oriented : HDD as OSD

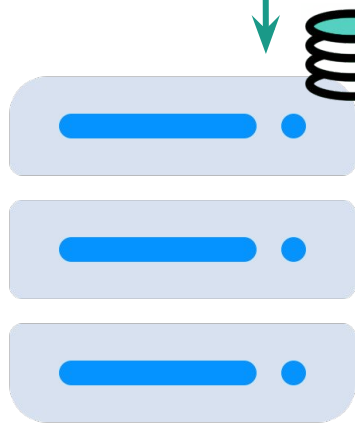
Use Ceph as Block storage for VM/Bare metal, native support, fully distributed storage

- QEMU
- KVM
- OpenStack (via Cinder API)
- Kubernetes

Access Block Storage with iSCSI on Ceph



10Gbps Network



Ceph
RBD

- 3xMON
- 3xMGR collocated with MON
- OSDs (with HDD or SSD)

Mars 400 ceph cluster

Windows and VMware clients, are not able to use RBD directly, for those use case, we will need to deploy ISCSI GW to convert RBD to iSCSI for clients use.

- 2x ISCSI GW servers to have HA
- Multipath IO supported.
- CHAP, IQN access control supported
- Block & file IO backstore
- Not fully distributed since there is ISCSI GW as bottleneck.

Access File System with CephFS



10Gbps Network



CephFS

- 3xMON
- 3xMGR collocated with MON
- OSDs (with HDD or SSD)
- MDS at least 2 active, 3 MDS standby collocated with MON



Mars 400 ceph cluster

CephFS is a POSIX compliant file system on top of Ceph, it could support below software directly without gateway server


- Backup software : Veeam, Commvault
- OpenStack
- K8S, container
- Nextcloud

MDS could scale out according to the users workload and access.

Access NFS or SAMBA with CephFS on Ceph



10Gbps Network



Mars 400 ceph cluster

CephFS

- 3xMON
- 3xMGR collocated with MON
- OSDs (with HDD or SSD)
- MDS at least 2 active, 3 MDS standby collocated with MON

For the application who want to use NFS or SAMBA service. User can convert CrphFS to NFS or SAMBA service.

This will require gateway server externally (with large scale capacity production cluster), 2x GW server to achieve HA.

- A scale out NAS
- A scale out SAMBA server with AD supported