Ceph User Survey - Cluster(s)

General

*How many clusters does this information represent?

[   ]

*Total raw capacity

[   ]

*Total usable capacity

- ( ) Up to 10 TiB
- ( ) 10 TiB to 100 TiB
- ( ) 100 TiB to 1 PiB
- ( ) 1 PiB to 10 PiB
- ( ) over 10 PiB

*Which Ceph release(s) do you run?

- [ ] Nautilus
- [ ] Mimic
- [ ] Luminous
- [ ] Kraken
- [ ] Jewel
- [ ] Infernalis
- [ ] Hammer
- [ ] Giant
- [ ] Firefly
- [ ] Emperor
- [ ] Dumpling
- [ ] Cuttlefish
- [ ] Bobtail
- [ ] Argonaut

*Which Ceph packages do you use?

- [ ] Upstream packages
- [ ] Distribution packages
- [ ] Vendor packages
- [ ] We build our own packages
- [ ] We built a custom version

*What operating system are you using on the cluster nodes?

- [ ] Ubuntu
- [ ] Debian
- [ ] CentOS
- [ ] Red Hat Enterprise Linux
- [ ] SUSE Linux Enterprise Server
- [ ] openSUSE
- [ ] Oracle Linux
- [ ] Arch Linux
- [ ] FreeBSD/OpenBSD/NetBSD
- [ ] Solaris, IllumOS or other OpenSolaris derivatives
- [ ] Gentoo Linux
- [ ] Other (Please specify)

[   ]

*Do you use cache tiering?

- ( ) I don't know what that is
- ( ) No
- ( ) Yes, in front of EC pools for functionality
- ( ) Yes, for performance reasons

*How soon do you apply dot/minor releases to your cluster?

- ( ) Within the week
- ( ) Within a month
- ( ) Longer

*Why?

- [ ] Concerns about regression
- [ ] Concerns about new functionality
- [ ] Too frequent updates
- [ ] Patching only as needed to address encountered
- [ ] Effort of update installation
- [ ] Other (Please specify)

[   ]

*How soon do you perform major version upgrades?

○ Within a month   ○ Within half a year   ○ Within a year

○ Longer   ○ Never (clusters remain
                            on major release they
                            were deployed with)

*Why?

☐ Concerns about stability   ☐ Concerns about   ☐ Effort of upgrade
                                performance
                                regressions

☐ Changed user
  experience due to new
  features

☐ Other (Please specify)

[   ]

*What Ceph Manager modules do you enable?

☐ balancer   ☐ crash   ☐ devicehealth

☐ diskprediction   ☐ iostat   ☐ pg_autoscaler

☐ restful

☐ Other (Please specify)

[   ]

*Which messenger type is in use?

○ TCP-Posix   ○ DPDK   ○ RDMA

Hardware

*What hardware vendors do you use for the nodes?

☐ IBM   ☐ Lenovo   ☐ Dell

☐ HPE   ☐ Supermicro   ☐ QCT

☐ Cisco   ☐ Fujitsu   ☐ Huawei

☐ Raspberry Pi Foundation

☐ Other (Please specify)

[   ]

*What processor architecture do you use?

☐ x86_64   ☐ ARM   ☐ Power

☐ Other (Please specify)

[   ]

*Which type of storage devices are used?

☐ HDD (SATA, SAS)   ☐ SSD (SATA, SAS)   ☐ NVMe

☐ MicroSD card

*Which OSD layout features do you use?

☐ LVM (ceph-volume)   ☐ Partitions   ☐ Encryption

☐ lvm/dm-   ☐ Seperate journal device   ☐ Separation between
  cache/bcache/iCAS   (FileStore)   RocksDB/WAL and Data
                                                                              device

☐ All-in-one OSD

*Do you use a dedicated network for the OSDs

○ Yes   ○ No

*IP protocol usage

○ IPv4 only   ○ IPv6 only   ○ Dual stack

Use cases, Platforms, and Interface Questions

*What are the use cases for the Ceph cluster?

☐ HPC   ☐ Cloud   ☐ Containers

☐ Virtualization   ☐ Home directories   ☐ Archive storage

☐ Scratch   ☐ Build   ☐ Logs

☐ CDN   ☐ Big data and analytics   ☐ Internet of things

☐ Other (Please specify)

```
[                                    ]
```

*What platforms do you use Ceph with?

☐ OpenStack   ☐ Cloudstack   ☐ Eucalyptus

☐ OpenNebula   ☐ Kubernetes   ☐ VMware

☐ Windows   ☐ Hyper-v   ☐ Proxmox

☐ KVM   ☐ Xen   ☐ RBD directly on Linux
                                                       systems

☐ Other (Please specify)

```
[                                    ]
```

*What interfaces are used on the cluster?

☐ Object (RGW, S3,   ☐ Block (RBD, iSCSI)   ☐ File (CephFS, NFS,
  Swift)   CIFS)

☐ Librados

RADOS Block Device (RBD)

Environment status

○ Test/POC ○ Development ○ Staging

○ Production

Workloads

☐ Boot virtual machines ☐ Database ☐ Backup repositories

☐ Other (Please specify)

[   ]

Are you using RBD asynchronous mirroring for Disaster Recovery / Multi-site?

○ Yes ○ No ○ Not needed

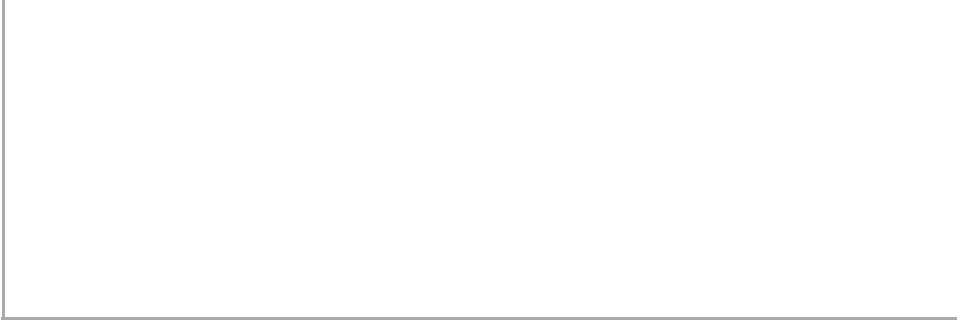○ Incompatible with my performance requirements ○ Not compatible with my RBD feature requirements

---

Clients accessing Ceph's block storage

☐ librbd ☐ Linux kernel RBD ☐ iSCSI (tcmu-runner)

☐ iSCSI (LIO with /dev/rbd)

☐ Other (Please specify)

[   ]

Do you use snapshots?

○ Yes

○ No (Please specify why not)

[   ]

Why?

Rados Gateway (RGW)

Environment status

- ○ Test/POC
- ○ Development
- ○ Staging
- ○ Production

Workloads

- ☐ Archive
- ☐ Backup
- ☐ Big data and analytics
- ☐ Other (Please specify)

  [   ]

RGW APIs used

- ☐ S3
- ☐ Swift
- ☐ RGW admin API
- ☐ NFS re-export of S3 buckets

RGW client-side libraries used

- ☐ Amazon SDK
- ☐ Boto
- ☐ Boto3
- ☐ None
- ☐ Other (Please specify)

  [   ]

RGW authentication mechanism

- ☐ RGW (built-in)
- ☐ Keystone
- ☐ LDAP
- ☐ Other
- ☐ Other (Please specify)

  [   ]

RGW external load balancers used

☐ HAproxy   ☐ Nginx   ☐ Dedicated custom hardware

☐ None

☐ Other (Please specify)

[   ]

---

*Number of RGW federated multi-sites?

[   ]

CephFS

Workloads

- [ ] Backup
- [ ] Build
- [ ] Big data (Hadoop, Spark)
- [ ] Static archive
- [ ] Home directories
- [ ] OpenStack Manila
- [ ] Media/streaming
- [ ] HPC
- [ ] General purpose NAS
- [ ] Other (Please specify)

Environment status

- ( ) Test/POC
- ( ) Development
- ( ) Staging
- ( ) Production

Interfaces to access CephFS

- [ ] Linux kernel CephFS mount
- [ ] ceph-fuse
- [ ] libcephfs
- [ ] NFS (nfs-ganesha)
- [ ] NFS (Kernel NFS server)
- [ ] CIFS (samba)

Typical number of file system clients (for largest cluster, if multiple clusters)

- ( ) 1 - 5
- ( ) 6 - 50
- ( ) 51 - 100
- ( ) 101 - 500
- ( ) 501 - 2500
- ( ) 2500+

Number of files (getfattr -d -m ceph.dir.rfiles /mnt/cephfs) (for largest cluster, if multiple clusters)

- ○ ~1k
- ○ ~10k
- ○ ~100k
- ○ ~1M
- ○ ~100M
- ○ ~1G
- ○ 10+

MDS cache size (for largest cluster, if multiple clusters)

- ○ 1G
- ○ 2G - 15G
- ○ 4G - 15G
- ○ 16G - 31G
- ○ 32G - 64G
- ○ 65G - 127G
- ○ 128G+

Number of active MDS (for largest cluster, if multiple)

- ○ 1
- ○ 2 - 3
- ○ 4 - 5
- ○ 6 - 8
- ○ 8+

Number of active MDS (for largest cluster, if multiple clusters)

[   ]

Do you use subtree pinning?

- ○ Yes
- ○ No
- ○ Unknown
- ○ Not applicable

Do you use snapshots?

- ○ Yes
- ○ No (Please specify)

  [   ]

Management and Operation

Management and monitoring tools

- [ ] Ceph Dashboard
- [ ] Calamari
- [ ] Ceph-Dash
- [ ] croit
- [ ] Inkscope
- [ ] Intel VSM
- [ ] openATTIC
- [ ] Proxmox
- [ ] InfluxDB
- [ ] Prometheus
- [ ] Ceph-metrics
- [ ] Grafana (custom)
- [ ] Zabbix
- [ ] Nagios/icinga
- [ ] Graphite
- [ ] Collectd
- [ ] Performance co-pilot
- [ ] Other (Please specify)

Deployment and configuration

- [ ] Ansible
- [ ] Chef
- [ ] Juju
- [ ] Puppet
- [ ] Salt / DeepSea
- [ ] ceph-deploy
- [ ] Rook
- [ ] croit
- [ ] Other (Please specify)

Ceph Dashboard

How likely are you to recommend the Ceph Dashboard to a colleague?

Unlikely   Likely

0	1	2	3	4	5	6	7	8	9	10

Does the dashboard help you to perform tasks better/faster than using the CLI?

◯ Yes

◯ No (Please specify)

Rank how often you use certain feature-sets (1 being frequently)

	Landing page (at a glance monitoring)
	Display cluster logs
	MON status page
	OSD management
	Pools management
	RBD management
	iSCSI target management
	NFS Ganesha management
	CephFS management
	RGW management
	View embedded Grafana dashboards

What functionality do you miss most in the dashboard?