

# Ceph User Survey - Cluster(s)

---

## General

\*How many clusters does this information represent?

---

\*Total raw capacity

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

---

\*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 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?

- |  |   |   |
|--|---|---|
| <input type="checkbox"/> Ubuntu  | <input type="checkbox"/> Debian                       | <input type="checkbox"/> CentOS                 |
| <input type="checkbox"/> Red Hat Enterprise Linux                          | <input type="checkbox"/> SUSE Linux Enterprise Server | <input type="checkbox"/> openSUSE               |
| <input type="checkbox"/> Oracle Linux                                      | <input type="checkbox"/> Arch Linux                   | <input type="checkbox"/> FreeBSD/OpenBSD/NetBSD |
| <input type="checkbox"/> Solaris, IllumOS or other OpenSolaris derivatives | <input type="checkbox"/> Gentoo Linux                 |   |
| <input type="checkbox"/> Other (Please specify)                            |   |   |

\*Do you using 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?

- |   |  |   |
|---|--|---|
| <input type="checkbox"/> Concerns about                                 | <input type="checkbox"/> Concerns about new            | <input type="checkbox"/> Too frequent updates |
| <input type="checkbox"/> Patching only as needed to address encountered | <input type="checkbox"/> Effort of update installation |   |
| <input type="checkbox"/> 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 performance regressions       Effort of upgrade
- Changed user experience due to new features
- Other (Please specify)
- 
- 

\*What Ceph Manager modules did you enable?

- balancer       crash       devicehealth
- diskprediction       iostat       pg\_autoscaler
- restful
- Other (Please specify)
-

# Hardware

\*What hardware vendors do you use for the nodes?

- |  |                                     |                                 |
|--|-------------------------------------|---------------------------------|
| <input type="checkbox"/> IBM                     | <input type="checkbox"/> Lenovo     | <input type="checkbox"/> Dell   |
| <input type="checkbox"/> HPE                     | <input type="checkbox"/> Supermicro | <input type="checkbox"/> QCT    |
| <input type="checkbox"/> Cisco                   | <input type="checkbox"/> Fujitsu    | <input type="checkbox"/> Huawei |
| <input type="checkbox"/> Raspberry Pi Foundation |                                     |                                 |
| <input type="checkbox"/> Other (Please specify)  |                                     |                                 |

\*What processor architecture do you use?

- |   |                              |                                |
|---|------------------------------|--------------------------------|
| <input type="checkbox"/> x86_64                 | <input type="checkbox"/> ARM | <input type="checkbox"/> Power |
| <input type="checkbox"/> Other (Please specify) |                              |                                |

\*Which type of storage devices are used?

- |  |  |                               |
|--|--|-------------------------------|
| <input type="checkbox"/> HDD (SATA, SAS) | <input type="checkbox"/> SSD (SATA, SAS) | <input type="checkbox"/> NVMe |
| <input type="checkbox"/> MicroSD card    |  |                               |

\*Which OSD layout features do you use?

- |   |  |   |
|---|--|---|
| <input type="checkbox"/> LVM (ceph-volume)        | <input type="checkbox"/> Partitions                          | <input type="checkbox"/> Encryption                                     |
| <input type="checkbox"/> lvm/dm-cache/bcache/iCAS | <input type="checkbox"/> Seperate journal device (FileStore) | <input type="checkbox"/> Separation between RocksDB/WAL and Data device |
| <input type="checkbox"/> All-in-one OSD           |  |   |

\*Do you use a dedicated network for the OSDs

- |                           |                          |
|---------------------------|--------------------------|
| <input type="radio"/> Yes | <input type="radio"/> 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?

- |   |   |   |
|---|---|---|
| <input type="checkbox"/> HPC                    | <input type="checkbox"/> Cloud                  | <input type="checkbox"/> Containers         |
| <input type="checkbox"/> Virtualization         | <input type="checkbox"/> Home directories       | <input type="checkbox"/> Archive storage    |
| <input type="checkbox"/> Scratch                | <input type="checkbox"/> Build                  | <input type="checkbox"/> Logs               |
| <input type="checkbox"/> CDN                    | <input type="checkbox"/> Big data and analytics | <input type="checkbox"/> Internet of things |
| <input type="checkbox"/> Other (Please specify) |   |   |

\*What platforms do you use Ceph with?

- |   |                                     |  |
|---|-------------------------------------|--|
| <input type="checkbox"/> OpenStack              | <input type="checkbox"/> Cloudstack | <input type="checkbox"/> Eucalyptus                    |
| <input type="checkbox"/> OpenNebula             | <input type="checkbox"/> Kubernetes | <input type="checkbox"/> VMware                        |
| <input type="checkbox"/> Windows                | <input type="checkbox"/> Hyper-v    | <input type="checkbox"/> Proxmox                       |
| <input type="checkbox"/> KVM                    | <input type="checkbox"/> Xen        | <input type="checkbox"/> RBD directly on Linux systems |
| <input type="checkbox"/> Other (Please specify) |                                     |  |

\*What interfaces are used on the cluster?

- |  |   |   |
|--|---|---|
| <input type="checkbox"/> Object (RGW, S3, Swift) | <input type="checkbox"/> Block (RBD, iSCSI) | <input type="checkbox"/> File (CephFS, NFS, CIFS) |
| <input type="checkbox"/> 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
- 

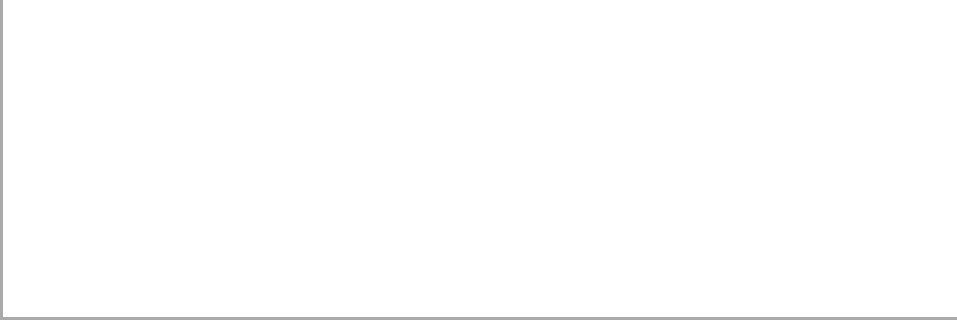
## Clients accessing Ceph's block storage

- librbd                       Linux kernel RBD                       iSCSI (tcmu-runner)  
 iSCSI (LIO with /dev/rbd)  
 Other (Please specify)
- 

## Do you use snapshots?

- Yes                       No
-

Why?

A large, empty rectangular box with a thin black border, intended for the user to write their answer to the question 'Why?'. The box is positioned directly below the question and occupies a significant portion of the upper half of the page.



# Rados Gateway (RGW)

## 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?

2

3

4

More than 4

Not applicable

# CephFS

## Workloads

- |   |   |  |
|---|---|--|
| <input type="checkbox"/> Backup                   | <input type="checkbox"/> Static archive   | <input type="checkbox"/> Media/streaming     |
| <input type="checkbox"/> Build                    | <input type="checkbox"/> Home directories | <input type="checkbox"/> HPC                 |
| <input type="checkbox"/> Big data (Hadoop, Spark) | <input type="checkbox"/> OpenStack Manila | <input type="checkbox"/> General purpose NAS |
| <input type="checkbox"/> Other (Please specify)   |   |  |
- 
- 

## Interfaces to access CephFS

- |  |  |                                       |
|--|--|---------------------------------------|
| <input type="checkbox"/> Linux kernel CephFS mount | <input type="checkbox"/> ceph-fuse               | <input type="checkbox"/> libcephfs    |
| <input type="checkbox"/> NFS (nfs-ganesha)         | <input type="checkbox"/> NFS (Kernel NFS server) | <input type="checkbox"/> CIFS (samba) |
- 

## Typical number of file system clients (per cluster)

- |                                 |                                  |                                |
|---------------------------------|----------------------------------|--------------------------------|
| <input type="radio"/> 1 - 5     | <input type="radio"/> 6 - 50     | <input type="radio"/> 51 - 100 |
| <input type="radio"/> 101 - 500 | <input type="radio"/> 501 - 2500 | <input type="radio"/> 2500+    |
- 

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

- |                           |                             |                             |
|---------------------------|-----------------------------|-----------------------------|
| <input type="radio"/> ~1k | <input type="radio"/> ~10k  | <input type="radio"/> ~100k |
| <input type="radio"/> ~1M | <input type="radio"/> ~100M | <input type="radio"/> ~1G   |
| <input type="radio"/> 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

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

### Number of active MDS for largest cluster

---

### 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

- |   |   |   |
|---|---|---|
| <input type="checkbox"/> Ceph Dashboard         | <input type="checkbox"/> Calamari             | <input type="checkbox"/> Ceph-Dash        |
| <input type="checkbox"/> Croit                  | <input type="checkbox"/> Inkscope             | <input type="checkbox"/> Intel VSM        |
| <input type="checkbox"/> openATTIC              | <input type="checkbox"/> Proxmox              | <input type="checkbox"/> InfluxDB         |
| <input type="checkbox"/> Prometheus             | <input type="checkbox"/> Ceph-metrics         | <input type="checkbox"/> Grafana (custom) |
| <input type="checkbox"/> Zabbix                 | <input type="checkbox"/> Nagios/icinga        | <input type="checkbox"/> Graphite         |
| <input type="checkbox"/> Collectd               | <input type="checkbox"/> Performance co-pilot |   |
| <input type="checkbox"/> Other (Please specify) |   |   |

---

## Deployment and configuration

- |   |   |                                      |
|---|---|--------------------------------------|
| <input type="checkbox"/> Ansible                | <input type="checkbox"/> Chef           | <input type="checkbox"/> Juju        |
| <input type="checkbox"/> Puppet                 | <input type="checkbox"/> Salt / DeepSea | <input type="checkbox"/> ceph-deploy |
| <input type="checkbox"/> Rook                   |   |                                      |
| <input type="checkbox"/> 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?