



Ceph User Survey 2021

Ceph User Survey

Deadline: April 2nd, 2021

Make your voice heard!

This information will help guide the Ceph community's investment in Ceph and the community's future development.

Frequently Asked Questions

Why is this survey important?

This information helps the Ceph development community and industry stakeholders focus our future development efforts and influences our day-to-day decisions.

What happens to the data?

Aggregated and collated results will be published shortly after the survey completes with the help of the Ceph User Survey Working Group. Individual survey responses will remain anonymous. The Ceph Foundation will be responsible for keeping that data private.

How long does the survey take?

Usually, about 10 minutes, although this may vary depending on the number of Ceph components you utilize.

On the **Platforms and Interface Questions** page, you will be asked:

1. Do you provide RADOS Block Device (RBD) services with your Ceph clusters?
2. Do you provide Ceph Object Gateway (RGW, Rados Gateway) services with your Ceph clusters?
3. Do you provide Ceph File system (CephFS) services with your Ceph clusters?

Additional questions may be included based on your response to these questions.

How do I report multiple deployments, for example, in different countries?

We recommend the survey be filled out for each type of deployment you have, regardless of state/country/etc. The answers are designed to facilitate this. Many respondents need to submit just once; nobody should need to submit more than two or three sets of answers.

How do I report an issue with the Survey?

The Ceph User Survey Working Group is an active set of volunteers who discuss making this survey worth the community's time. Please reach out to us at survey@ceph.io.

If you need to communicate with the Ceph Foundation directly or need help reaching the right person in the community, please let us help at foundation@ceph.io. General questions and feedback will be shared with The Ceph User Survey Working Group to refine the survey experience.

How will the Ceph User Survey continue to evolve?

The Ceph community loves feedback and learning about how the user community is using ceph. We also recognize that completing surveys can be tedious.

The Ceph Telemetry project allows Ceph clusters to anonymously and automatically report basic information about Ceph usage back to the developer community: what versions of Ceph are in use, what features are enabled, how big or busy clusters are, what types of hardware they are built with, what bugs are encountered, and so on. We strongly encourage all Ceph users to review what kinds of information is reported and then enable the telemetry feature or tell us why they are uncomfortable sharing that data. We would like to strike the right balance between maximizing participation and ensuring that the data is valuable.

As participation in telemetry increases, the user survey will focus more on non-technical details, like your organizations' industry, geography, size, and relationship with the open-source community and Ceph vendors.

We welcome any and all feedback on survey content so that the Ceph User Survey Working Group can improve the experience and data collected for next year.

General

Ceph User Survey 2021

*How long have you been using Ceph?

- Less than 1 year Between 1-2 years Between 2-5 years
 More than 5 years I don't use Ceph yet
-

*Why do you use Ceph?

Open source

Scalability

Cost

Feature set

Security

High availability

Data durability /
reliability / integrity

Performance

Integration with adjacent
technologies

Other (Please specify)

*Sector(s) of your organization?

Commercial

Government

Military

Academic

Non-profit

Personal

Other (Please specify)

Which resources do you check when you need help?

	Do not use					Primary source
Ceph Documentation	<input type="radio"/> 0	<input type="radio"/> 1	<input type="radio"/> 2	<input type="radio"/> 3	<input type="radio"/> 4	<input type="radio"/> 5
Ceph-users mailing list	<input type="radio"/> 0	<input type="radio"/> 1	<input type="radio"/> 2	<input type="radio"/> 3	<input type="radio"/> 4	<input type="radio"/> 5
Slack	<input type="radio"/> 0	<input type="radio"/> 1	<input type="radio"/> 2	<input type="radio"/> 3	<input type="radio"/> 4	<input type="radio"/> 5
Meet-up or other local groups	<input type="radio"/> 0	<input type="radio"/> 1	<input type="radio"/> 2	<input type="radio"/> 3	<input type="radio"/> 4	<input type="radio"/> 5
IRC	<input type="radio"/> 0	<input type="radio"/> 1	<input type="radio"/> 2	<input type="radio"/> 3	<input type="radio"/> 4	<input type="radio"/> 5
Commercial Provider	<input type="radio"/> 0	<input type="radio"/> 1	<input type="radio"/> 2	<input type="radio"/> 3	<input type="radio"/> 4	<input type="radio"/> 5
Source code	<input type="radio"/> 0	<input type="radio"/> 1	<input type="radio"/> 2	<input type="radio"/> 3	<input type="radio"/> 4	<input type="radio"/> 5

Additional Resource

Where should the Ceph community focus its efforts?

	Not important					Very important
Reliability	0	1	2	3	4	5
Management / Configuration	0	1	2	3	4	5
Scalability	0	1	2	3	4	5
Security	0	1	2	3	4	5
Interoperability	0	1	2	3	4	5
Performance	0	1	2	3	4	5
Monitoring	0	1	2	3	4	5
Documentation	0	1	2	3	4	5

How do you participate in the Ceph community?

- IRC / Slack / etc
- Mailing list
- Reporting issues via the bug tracker
- Contributing code
- Contributing / enhancing documentation
- Ceph events / conferences
- Member of the Ceph Foundation
- Other (Please specify)

*How likely are you to recommend Ceph to a colleague?

unlikely

likely

0	1	2	3	4	5	6	7	8	9	10
---	---	---	---	---	---	---	---	---	---	----

Please share the top reason for picking the previous answer (optional)

Top 3 wishlist features (add each on a separate line, optional)

*What countries are your deployments in?

- | | | |
|---|--------------------------------------|--|
| <input type="checkbox"/> Afghanistan | <input type="checkbox"/> Albania | <input type="checkbox"/> Algeria |
| <input type="checkbox"/> Andorra | <input type="checkbox"/> Angola | <input type="checkbox"/> Antigua & Deps |
| <input type="checkbox"/> Argentina | <input type="checkbox"/> Armenia | <input type="checkbox"/> Australia |
| <input type="checkbox"/> Austria | <input type="checkbox"/> Azerbaijan | <input type="checkbox"/> Bahamas |
| <input type="checkbox"/> Bahrain | <input type="checkbox"/> Bangladesh | <input type="checkbox"/> Barbados |
| <input type="checkbox"/> Belarus | <input type="checkbox"/> Belgium | <input type="checkbox"/> Belize |
| <input type="checkbox"/> Benin | <input type="checkbox"/> Bhutan | <input type="checkbox"/> Bolivia |
| <input type="checkbox"/> Bosnia Herzegovina | <input type="checkbox"/> Botswana | <input type="checkbox"/> Brazil |
| <input type="checkbox"/> Brunei | <input type="checkbox"/> Bulgaria | <input type="checkbox"/> Burkina Faso |
| <input type="checkbox"/> Burundi | <input type="checkbox"/> Cambodia | <input type="checkbox"/> Cameroon |
| <input type="checkbox"/> Canada | <input type="checkbox"/> Cape Verde | <input type="checkbox"/> Central African Rep |
| <input type="checkbox"/> Chad | <input type="checkbox"/> Chile | <input type="checkbox"/> China |
| <input type="checkbox"/> Colombia | <input type="checkbox"/> Comoros | <input type="checkbox"/> Congo |
| <input type="checkbox"/> Congo {Democratic Rep} | <input type="checkbox"/> Costa Rica | <input type="checkbox"/> Croatia |
| <input type="checkbox"/> Cuba | <input type="checkbox"/> Cyprus | <input type="checkbox"/> Czech Republic |
| <input type="checkbox"/> Denmark | <input type="checkbox"/> Djibouti | <input type="checkbox"/> Dominica |
| <input type="checkbox"/> Dominican Republic | <input type="checkbox"/> East Timor | <input type="checkbox"/> Ecuador |
| <input type="checkbox"/> Egypt | <input type="checkbox"/> El Salvador | <input type="checkbox"/> Equatorial Guinea |
| <input type="checkbox"/> Eritrea | <input type="checkbox"/> Estonia | <input type="checkbox"/> Ethiopia |
| <input type="checkbox"/> Faroe Islands | <input type="checkbox"/> Fiji | <input type="checkbox"/> Finland |
| <input type="checkbox"/> France | <input type="checkbox"/> Gabon | <input type="checkbox"/> Gambia |
| <input type="checkbox"/> Georgia | <input type="checkbox"/> Germany | <input type="checkbox"/> Ghana |
| <input type="checkbox"/> Grand Fenwick | <input type="checkbox"/> Greece | <input type="checkbox"/> Grenada |
| <input type="checkbox"/> Guatemala | <input type="checkbox"/> Guinea | <input type="checkbox"/> Guinea-Bissau |
| <input type="checkbox"/> Guyana | <input type="checkbox"/> Haiti | <input type="checkbox"/> Honduras |
| <input type="checkbox"/> Hungary | <input type="checkbox"/> Iceland | <input type="checkbox"/> India |
| <input type="checkbox"/> Indonesia | <input type="checkbox"/> Iran | <input type="checkbox"/> Iraq |
| <input type="checkbox"/> Ireland {Republic} | <input type="checkbox"/> Israel | <input type="checkbox"/> Italy |
| <input type="checkbox"/> Ivory Coast | <input type="checkbox"/> Jamaica | <input type="checkbox"/> Japan |
| <input type="checkbox"/> Jordan | <input type="checkbox"/> Kazakhstan | <input type="checkbox"/> Kenya |
| <input type="checkbox"/> Kiribati | <input type="checkbox"/> Korea North | <input type="checkbox"/> Korea South |

- | | | |
|---|---|---|
| <input type="checkbox"/> Kosovo | <input type="checkbox"/> Kuwait | <input type="checkbox"/> Kyrgyzstan |
| <input type="checkbox"/> Laos | <input type="checkbox"/> Latvia | <input type="checkbox"/> Lebanon |
| <input type="checkbox"/> Leng | <input type="checkbox"/> Lesotho | <input type="checkbox"/> Liberia |
| <input type="checkbox"/> Libya | <input type="checkbox"/> Liechtenstein | <input type="checkbox"/> Lithuania |
| <input type="checkbox"/> Luxembourg | <input type="checkbox"/> Macedonia | <input type="checkbox"/> Madagascar |
| <input type="checkbox"/> Malawi | <input type="checkbox"/> Malaysia | <input type="checkbox"/> Maldives |
| <input type="checkbox"/> Mali | <input type="checkbox"/> Malta | <input type="checkbox"/> Marshall Islands |
| <input type="checkbox"/> Mauritania | <input type="checkbox"/> Mauritius | <input type="checkbox"/> Mexico |
| <input type="checkbox"/> Micronesia | <input type="checkbox"/> Moldova | <input type="checkbox"/> Monaco |
| <input type="checkbox"/> Mongolia | <input type="checkbox"/> Montenegro | <input type="checkbox"/> Morocco |
| <input type="checkbox"/> Mozambique | <input type="checkbox"/> Myanmar, {Burma} | <input type="checkbox"/> Namibia |
| <input type="checkbox"/> Nauru | <input type="checkbox"/> Narnia | <input type="checkbox"/> Nepal |
| <input type="checkbox"/> Netherlands | <input type="checkbox"/> New Zealand | <input type="checkbox"/> Nicaragua |
| <input type="checkbox"/> Niger | <input type="checkbox"/> Nigeria | <input type="checkbox"/> Norway |
| <input type="checkbox"/> Oman | <input type="checkbox"/> Pakistan | <input type="checkbox"/> Palau |
| <input type="checkbox"/> Panama | <input type="checkbox"/> Papua New Guinea | <input type="checkbox"/> Paraguay |
| <input type="checkbox"/> Peru | <input type="checkbox"/> Philippines | <input type="checkbox"/> Poland |
| <input type="checkbox"/> Portugal | <input type="checkbox"/> Qatar | <input type="checkbox"/> Romania |
| <input type="checkbox"/> Russian Federation | <input type="checkbox"/> Rwanda | <input type="checkbox"/> St Kitts & Nevis |
| <input type="checkbox"/> St Lucia | <input type="checkbox"/> St. Vincent & Grenadines | <input type="checkbox"/> Samoa |
| <input type="checkbox"/> San Marino | <input type="checkbox"/> Sao Tome & Principe | <input type="checkbox"/> Saudi Arabia |
| <input type="checkbox"/> Senegal | <input type="checkbox"/> Serbia | <input type="checkbox"/> Seychelles |
| <input type="checkbox"/> Sierra Leone | <input type="checkbox"/> Singapore | <input type="checkbox"/> Slovakia |
| <input type="checkbox"/> Slovenia | <input type="checkbox"/> Solomon Islands | <input type="checkbox"/> Somalia |
| <input type="checkbox"/> South Africa | <input type="checkbox"/> South Sudan | <input type="checkbox"/> Spain |
| <input type="checkbox"/> Sri Lanka | <input type="checkbox"/> Sudan | <input type="checkbox"/> Suriname |
| <input type="checkbox"/> Swaziland | <input type="checkbox"/> Sweden | <input type="checkbox"/> Switzerland |
| <input type="checkbox"/> Syria | <input type="checkbox"/> Taiwan | <input type="checkbox"/> Tajikistan |
| <input type="checkbox"/> Tanzania | <input type="checkbox"/> Thailand | <input type="checkbox"/> Togo |
| <input type="checkbox"/> Tonga | <input type="checkbox"/> Trinidad & Tobago | <input type="checkbox"/> Tunisia |
| <input type="checkbox"/> Turkey | <input type="checkbox"/> Turkmenistan | <input type="checkbox"/> Tuvalu |
| <input type="checkbox"/> Uganda | <input type="checkbox"/> Ukraine | <input type="checkbox"/> United Arab Emirates |
| <input type="checkbox"/> United Kingdom | <input type="checkbox"/> United States | <input type="checkbox"/> Uruguay |

Uzbekistan

Vanuatu

Vatican City

Venezuela

Vietnam

Yemen

Yuggoth

Zambia

Zimbabwe

That would be Telling :p

Telemetry

*Do you have Telemetry enabled on your cluster(s)? (Telemetry will automatically report anonymized information about your cluster, like the size of your cluster, what version of Ceph you are running, and what bugs you have encountered.)

Yes, telemetry is enabled on all of my clusters

Yes, telemetry is enabled on some of my clusters

No, telemetry is not enabled on any clusters

Why is telemetry not enabled?

I haven't gotten around to it yet

My cluster(s) are running a Ceph version older than Luminous

My cluster(s) are on a protected network that does not have access to the internet

My organization's security policy does not allow storage systems to communicate with vendors or external organizations

My organization's security policy has reviewed the information reported by telemetry and judged that it should not be shared

Enabling telemetry would require review by my organization's security team and I have chosen not request a review

I am worried that even the anonymized information reported by telemetry may compromise my organization's security.

I do not believe that the Ceph community should be collecting any of this information

Other (Please specify)

Have you disabled any of the default telemetry channels?

basic: cluster capacity, cluster size, ceph version, number of pools and file systems, which configuration options have been adjusted)

crash: information about daemon crashes, including the type of daemon, ceph version, and stack trace

device: anonymized device health metrics (e.g., SMART hard disk metrics)

Have you enabled the 'ident' telemetry channel (which allows you to provide a cluster description and contact email address)?

I left 'ident' disabled (the default)

I enabled 'ident' and configured a cluster description and/or email address

Have you used the public Ceph Telemetry dashboards and/or Devices dashboard?

No, I didn't realize they existed

No, I knew they existed by have not used them

Yes

Why did you use the public telemetry dashboard and/or the device telemetry dashboard?

Out of curiosity

To inform my decisions about when or whether to upgrade

To make decisions around cluster sizing (for new cluster(s))

To make decisions around cluster expansion (for existing cluster(s))

To make decisions about which type of storage devices to buy (HDD vs SSD)

To make decisions about which vendor(s) or model(s) of storage devices to buy

Other

General Cluster Questions

*How many clusters does this information represent?

*Raw capacity (in TB) of your largest cluster

*Total raw capacity of all clusters (before replication)

*Total usable capacity of all clusters (after replication)

*Which Ceph releases do you run?

- | | | |
|--|--|---------------------------------------|
| <input type="checkbox"/> Octopus (15.x) | <input type="checkbox"/> Nautilus (14.x) | <input type="checkbox"/> Mimic (13.x) |
| <input type="checkbox"/> Luminous (12.x) | <input type="checkbox"/> Kraken (11.x) | <input type="checkbox"/> Jewel (10.x) |
| <input type="checkbox"/> Infernalis (9.x) or older | <input type="checkbox"/> Master | |

*Whose Ceph packages do you use?

- | | | |
|--|--|--|
| <input type="checkbox"/> Upstream packages | <input type="checkbox"/> Distribution packages | <input type="checkbox"/> Vendor packages |
| <input type="checkbox"/> We build our own packages | <input type="checkbox"/> We built a custom version | |

*What operating systems are you using on cluster nodes?

- | | | |
|---------------------------------|---|---|
| <input type="checkbox"/> Ubuntu | <input type="checkbox"/> Debian | <input type="checkbox"/> RHEL / CentOS / Fedora |
| <input type="checkbox"/> SUSE | <input type="checkbox"/> FreeBSD/OpenBSD/NetBSD | <input type="checkbox"/> Other OS |
-

*Do you use cache tiering?

- | | | |
|--------------------------|---|--|
| <input type="radio"/> No | <input type="radio"/> Yes, in front of EC pools for functionality | <input type="radio"/> Yes, for performance reasons |
|--------------------------|---|--|
-

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

- | | | |
|---|--|------------------------------|
| <input type="radio"/> Within a week after release | <input type="radio"/> Within a month after release | <input type="radio"/> Longer |
| <input type="radio"/> Never | | |
-

*Why?

- | | | |
|--|--|--|
| <input type="checkbox"/> I apply upgrades quickly to get any available bug fixes or security fixes | <input type="checkbox"/> I apply upgrades quickly because they require little effort | <input type="checkbox"/> I delay upgrades due to concerns about regression |
| <input type="checkbox"/> I delay upgrades due to concerns about new functionality | <input type="checkbox"/> I delay upgrades because there are too many updates | <input type="checkbox"/> I delay upgrades because of the effort involved |
| <input type="checkbox"/> I only upgrade as needed to address issues | | |
| <input type="checkbox"/> Other (Please specify) | | |

*How soon do you perform major version upgrades?

- | | | |
|---|--|-------------------------------------|
| <input type="radio"/> Within a month of release | <input type="radio"/> Within half a year of release | <input type="radio"/> Within a year |
| <input type="radio"/> Longer | <input type="radio"/> Never (clusters remain on major release they were deployed with) | |
-

*Why?

- | | | |
|---|--|---|
| <input type="checkbox"/> I upgrade quickly to get new features | <input type="checkbox"/> I upgrade quickly to get any available bug fixes or security fixes | <input type="checkbox"/> I upgrade quickly because it is easy |
| <input type="checkbox"/> I only upgrade as needed to address specific bugs or needed features | <input type="checkbox"/> I delay upgrades due to concerns about stability | <input type="checkbox"/> I delay upgrades due to concerns about performance regressions |
| <input type="checkbox"/> I delay upgrades due to the effort required to upgrade | <input type="checkbox"/> I delay upgrades to avoid a changed user experience due to new features | |
| <input type="checkbox"/> Other (Please specify) | | |

Which Ceph Manager modules do you enable?

- | | | |
|---|---|--|
| <input type="checkbox"/> balancer | <input type="checkbox"/> crash | <input type="checkbox"/> devicehealth |
| <input type="checkbox"/> diskprediction | <input type="checkbox"/> iostat | <input type="checkbox"/> pg_autoscaler |
| <input type="checkbox"/> restful | <input type="checkbox"/> I don't know / default set | |
| <input type="checkbox"/> Other (Please specify) | | |

Hardware

*Which chassis / server vendors do you use for cluster nodes?

- | | | |
|---|--|---|
| <input type="checkbox"/> IBM | <input type="checkbox"/> Lenovo | <input type="checkbox"/> Dell |
| <input type="checkbox"/> HPE (Hewlett Packard) | <input type="checkbox"/> Supermicro | <input type="checkbox"/> QCT (Quanta) |
| <input type="checkbox"/> Intel | <input type="checkbox"/> SoftIron | <input type="checkbox"/> Cisco |
| <input type="checkbox"/> Fujitsu | <input type="checkbox"/> Huawei | <input type="checkbox"/> We build our own |
| <input type="checkbox"/> Raspberry Pi | <input type="checkbox"/> Prefer not to say | |
| <input type="checkbox"/> Other (Please specify) | | |

*Which processor architectures and manufacturers do you use?

- | | | |
|---|--------------------------------------|--|
| <input type="checkbox"/> ARM | <input type="checkbox"/> Intel Xeon | <input type="checkbox"/> Intel (other) |
| <input type="checkbox"/> AMD Epyc | <input type="checkbox"/> AMD (other) | <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 or other low-end flash | <input type="checkbox"/> SanDisk InfiniFlash or other AFA (All Flash Arrays) | <input type="checkbox"/> Optane |

*Which OSD layout strategies do you employ?

- | | | |
|--|--|---|
| <input type="checkbox"/> LVM (ceph-volume) | <input type="checkbox"/> Partitions | <input type="checkbox"/> At-rest encryption (dmccrypt) |
| <input type="checkbox"/> lvmcache | <input type="checkbox"/> dm-cache | <input type="checkbox"/> bcache |
| <input type="checkbox"/> CAS (Intel) | <input type="checkbox"/> Separate journal device (Filestore) or RocksDB/WAL device (BlueStore) | <input type="checkbox"/> All-in-one OSD (default, colocated, no separate metadata device) |
| <input type="checkbox"/> Containerized daemons | <input type="checkbox"/> Defaults / I don't know | |

* Do you use a dedicated cluster / replication / private network for OSD node replication and recovery?

- Yes No Mix of yes and no
-

* Network technologies in use

- 100 Mb/s Ethernet 1000 Mb/s Ethernet (GigE, gigabit) 10 Gb/s Ethernet
 40 Gb/s Ethernet 25 Gb/s Ethernet 100 Gb/s Ethernet
 InfiniBand RJ45 (Copper) DAC (Copper)
 AOC / Fiber Wireless
 Other (Please specify)

* IP address families in use

- IPv4 only IPv6 only Dual stack
-

Which SSD form factor are important to you

- M.2 U.2 (2.5-inch) E1.L (EDSFF, ruler form factor)
 Add In Cards
-

* Preferred number of CPU sockets per server and why

- Single Dual Quad
 No preference / beats me

Explain why?

*How many OSDs in your largest cluster?

- 10 or fewer 10 to 100 100 to 500
- 500 to 1000 1000 to 3000 3000 to 5000
- 5000+
-

*Which OSD back ends do your clusters employ?

- Filestore (XFS, EXT4, BTRFS) BlueStore Both Filestore and BlueStore
- Not sure
- Other (Please specify)
-
-

*Data protection strategies in use

- Replication (2x) Replication (3x) Replication (4 or more copies)
- Erasure coding

Platforms and Interface Questions

*With which platforms do you use Ceph?

- | | | |
|---|-------------------------------------|---|
| <input type="checkbox"/> OpenStack | <input type="checkbox"/> Cloudstack | <input type="checkbox"/> OpenShift |
| <input type="checkbox"/> OpenNebula | <input type="checkbox"/> Kubernetes | <input type="checkbox"/> VMware |
| <input type="checkbox"/> Microsoft Windows | <input type="checkbox"/> Hyper-v | <input type="checkbox"/> Proxmox |
| <input type="checkbox"/> KVM / QEMU | <input type="checkbox"/> Xen | <input type="checkbox"/> KRBD directly on Linux systems |
| <input type="checkbox"/> rbd-nbd | <input type="checkbox"/> None | |
| <input type="checkbox"/> Other (Please specify) | | |

*Do you provide RADOS Block Device (RBD) services with your Ceph clusters?

- Yes No

*Do you provide Ceph Object Gateway (RGW, Rados Gateway) services with your Ceph clusters?

- Yes No

*Do you provide Ceph File system (Ceph FS) services with your Ceph clusters?

- Yes No

RADOS Block Device (RBD)

RBD deployment status

- Development / Testing Staging Production
 Proof of Concept (PoC)
-

*What are your use cases for RBD?

- Scratch Build Log
 Video CDN Big Data & Analytics
 AI/ML HPC Cloud
 Containers Virtualization Home Directories
 Archive Storage
 Other (Please specify)

*Are you using rbd-mirror?

- Yes, for DR Yes, to migrate volumes between clusters No, it is not needed
 No, it is incompatible with my performance requirements No, it is not compatible with my RBD feature requirements
 Other (Please specify)

*RBD (block storage) client access methods in use

- librbd (e.g., in combination with qemu) Linux kernel RBD iSCSI (tcmu-runner)
 iSCSI (LIO with /dev/rbd) rbd-nbd
 Other (Please specify)

*Do you use RBD snapshots and / or clones?

Yes

No

Unknown

Rados Gateway (RGW)

RGW Ceph cluster roles (dev, staging, production, PoC)

- Development / Testing Staging Production
- Proof of Concept (PoC)
-

*What are your RGW workloads?

- Scratch Build Log
- Video CDN Big Data & Analytics
- AI/ML HPC Cloud
- Containers Virtualization Home Directories
- Archive Storage Backup OLTP
- Other (Please specify)

*Average number of RGW daemons per object service cluster

*Size in GB of the largest S3 / Swift object stored

*RGW APIs used

- S3 Swift RGW admin API
- NFS re-export of S3 buckets RGW data caching and CDN
-

*RGW clients and libraries used

- Amazon SDK Boto Boto3
 s3cmd / s3tools custom / in-house None
 Other (Please specify)

*RGW authentication mechanism

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

*What load balancer(s) or caches do you use with RGW?

- HAproxy Nginx Varnish
 Dedicated custom hardware (Alteon, F5, etc)
 DNS round-robin None
 Other (Please specify)

*Number of RGW federated multi-sites?

*Is RGW Data Caching and CDN feature useful to you? if so list your use case/s?

Do you have Ceph Multisite Clusters?

Yes

No

*Do you have more than one realm in your setup? if so, how many?

*Do you have more than one zone group in your setup?

Yes

No

*Do you have more than one zone in your setup? if so, how many in the largest zone group?

*Is the syncing policy between zones global or per bucket?

Global

Per bucket

Do you use Sync Modules? if so, which modules? and list with which cloud providers.

ElasticSearch Sync Module

Cloud Sync Module

PubSub Module

Archive Sync Module

Check if you use an archive zone

List with which Cloud providers:

Do you use Bucket Notifications? if so, which endpoints do you use?

Kafka

Amqp

Http

Which other endpoints would you like to have?

CephFS

CephFS deployment status

- Development / Testing Staging Production
- Proof of Concept (PoC)
-

*What are your use cases for CephFS?

- Scratch Build Log
- Video CDN Big Data & Analytics
- AI/ML HPC Cloud
- Containers Virtualization Home Directories
- Archive Storage
- Other (Please specify)

*CephFS consumption methods

- Linux kernel CephFS mount ceph-fuse libcephfs
- NFS (nfs-ganesha) NFS (Kernel NFS server) CIFS (SMB, Samba)
- Other (Please specify)

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

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

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

- 1G 2G - 3G 4G - 15G
 16G - 31G 32G - 64G 65G - 127G
 128G+
-

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

*Do you use CephFS subtree pinning?

- Yes No Unknown
-

*Do you use CephFS snapshots?

- Yes No Unknown
-

*Do you use CephFS Quotas?

- Yes No Unknown
-

*Which of these experimental Ceph FS features are important to you?

- Inline Data Mantle (Programmable Metadata Balancer) Multiple File Systems within a Ceph Cluster
 LazyIO

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)

How often do you use certain dashboard features?

	Do not use					Primary Source
Landing page (at a glance monitoring)	0	1	2	3	4	5
Display cluster logs	0	1	2	3	4	5
Mon status page	0	1	2	3	4	5
OSD management	0	1	2	3	4	5
Pools management	0	1	2	3	4	5
RBD management	0	1	2	3	4	5
iSCSI target management	0	1	2	3	4	5
NFS Ganesha management	0	1	2	3	4	5
CephFS management	0	1	2	3	4	5
RGW management	0	1	2	3	4	5
View embedded Grafana dashboards	0	1	2	3	4	5

*What functionality do you miss most in the dashboard?

Management and Operation

*What metrics and monitoring tools do you use?

- | | | |
|---|---|--|
| <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"/> node_exporter | <input type="checkbox"/> ceph_exporter |
| <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"/> Telegraf | <input type="checkbox"/> Performance co-pilot | |
| <input type="checkbox"/> Other (Please specify) | | |

*Which deployment and configuration platforms do you use to manage your Ceph clusters?

- | | | |
|---|---|---|
| <input type="checkbox"/> Ansible (ceph-ansible) | <input type="checkbox"/> Ansible (other / homebrew) | <input type="checkbox"/> cephadm |
| <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"/> Croit | <input type="checkbox"/> OSNEXUS QuantaStor | <input type="checkbox"/> SoftIron Storage Manager |
| <input type="checkbox"/> Other (Please specify) | | |