

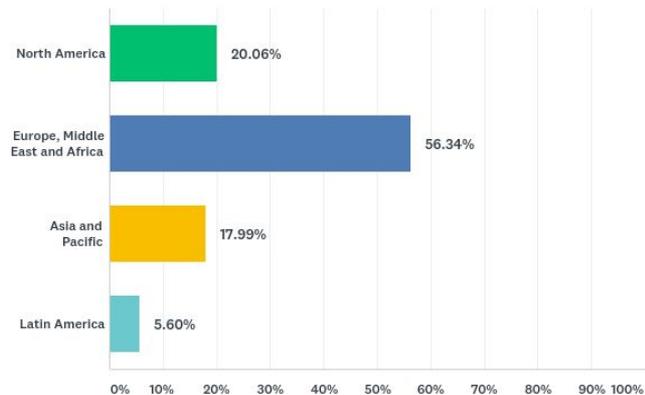
Ceph User Survey 2018

About

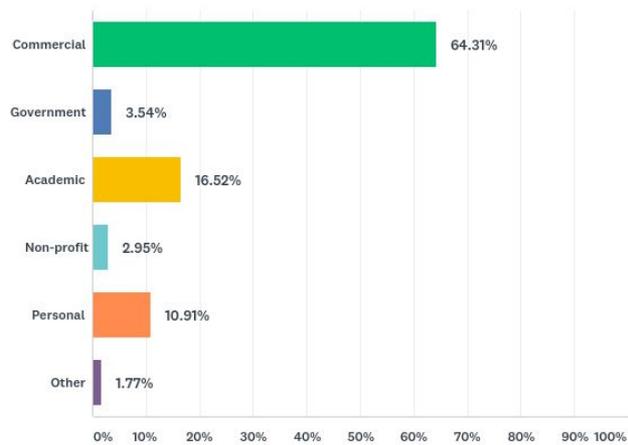
- Survey conducted from May 1st to May 25th
- 342 complete responses
- 52 different countries
- 4 continents

General questions

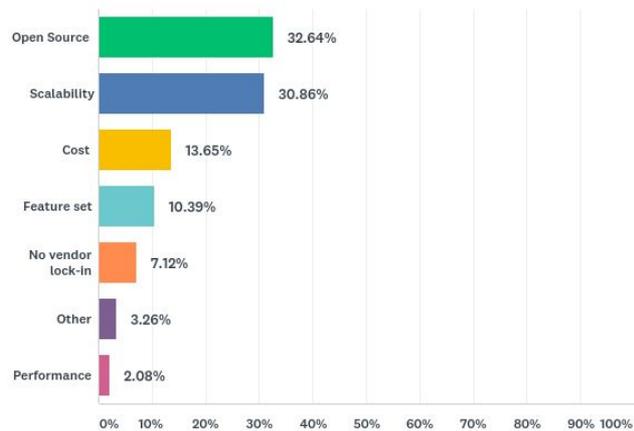
Where are the Ceph users?



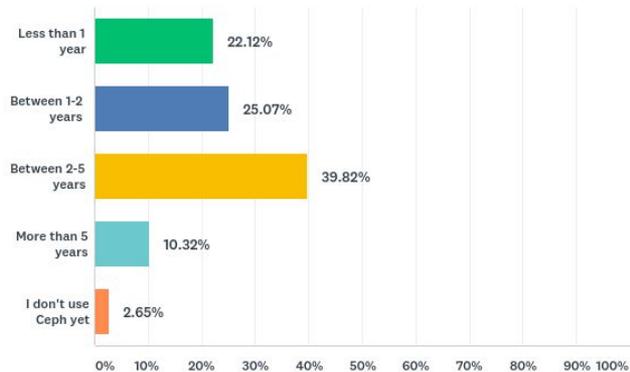
Organizations using Ceph



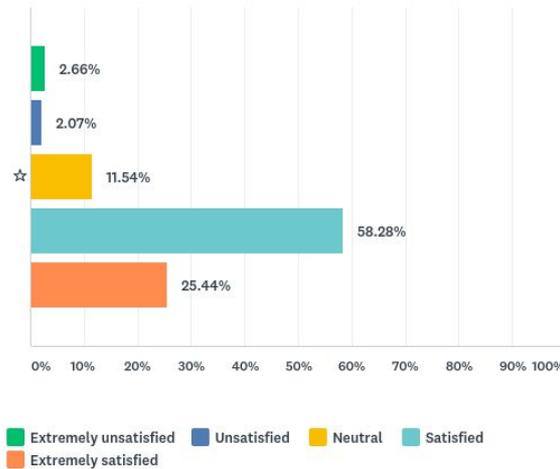
Why Ceph?



How long have you been using Ceph?



Are you satisfied with the technology?



Ceph Clusters

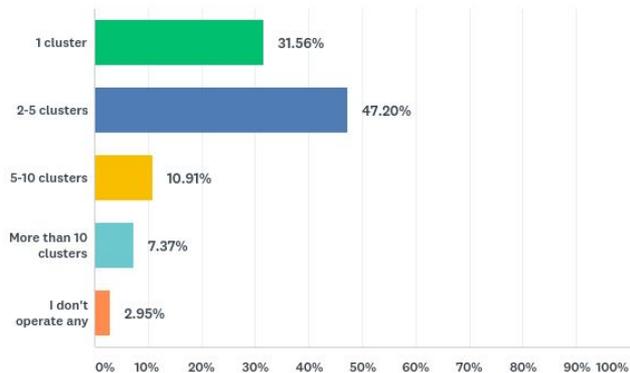
Total raw capacity

- Up to 10 TB: 8%
- From 10 TB to 100 TB: 28.31%
- **From 100 TB to 1 PB: 37.84%**
- From 1PB to 10PB: 20.91%
- From 10PB to 100PB: 4.92%

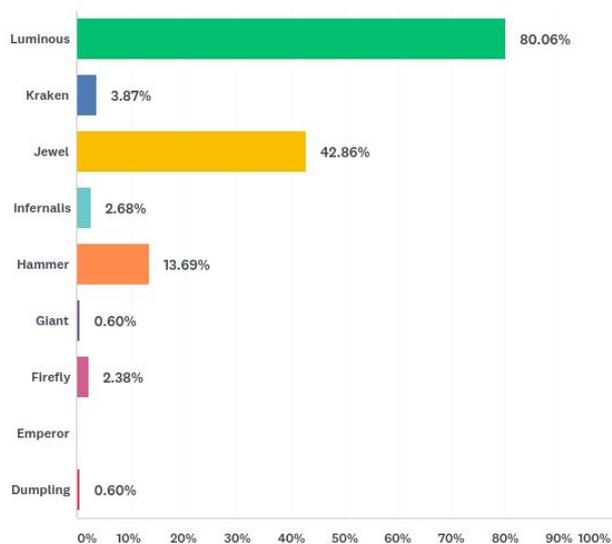
Usable capacity

- Up to 10 TB: 14.79%
- **From 10 TB to 100 TB: 35.69%**
- From 100 TB to 1 PB: 33.76%
- From 1PB to 10PB: 13.38%
- From 10PB to 100PB: 2.57%

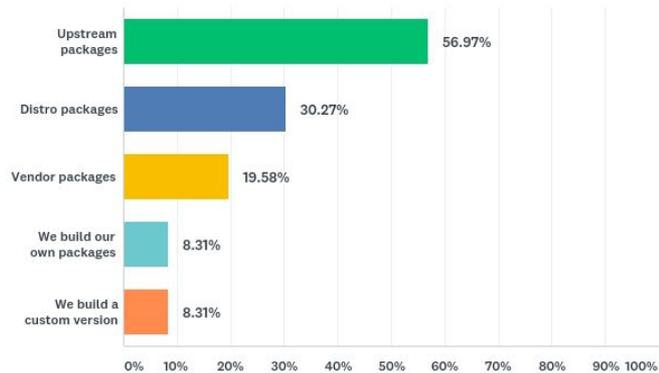
How many clusters do you operate?



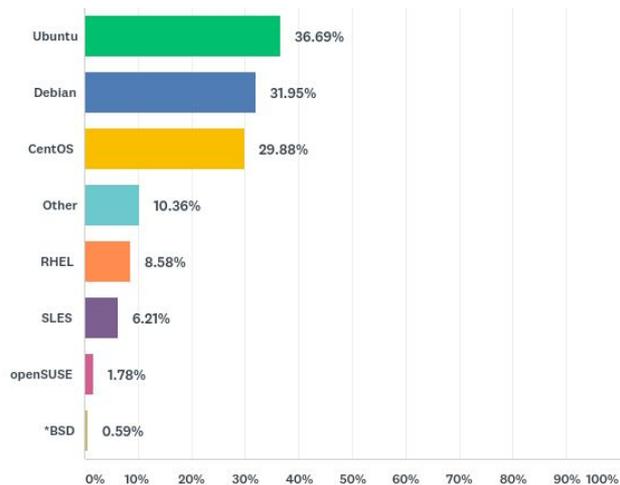
Which Ceph release(s) do they run?



Which Ceph packages do you use?



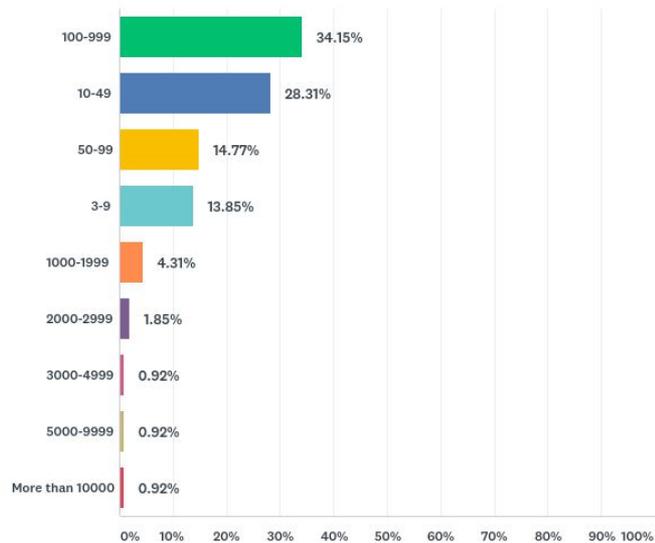
What O.S. are used on the cluster nodes?



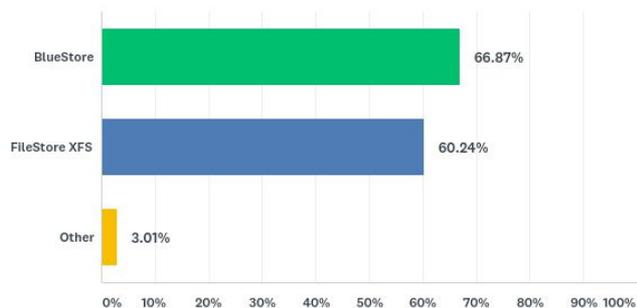
How many OSD nodes in your cluster?

- **Up to 10 nodes: 56.19%**
- From 10 to 100 nodes: 35.31%
- From 100 to 1000 nodes: 7.18%
- From 1000 to 15000 nodes: 1.25%

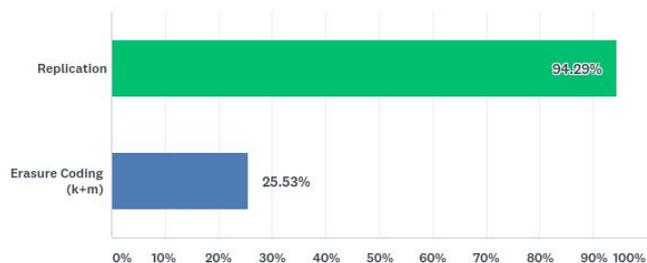
How many OSDs in your cluster?



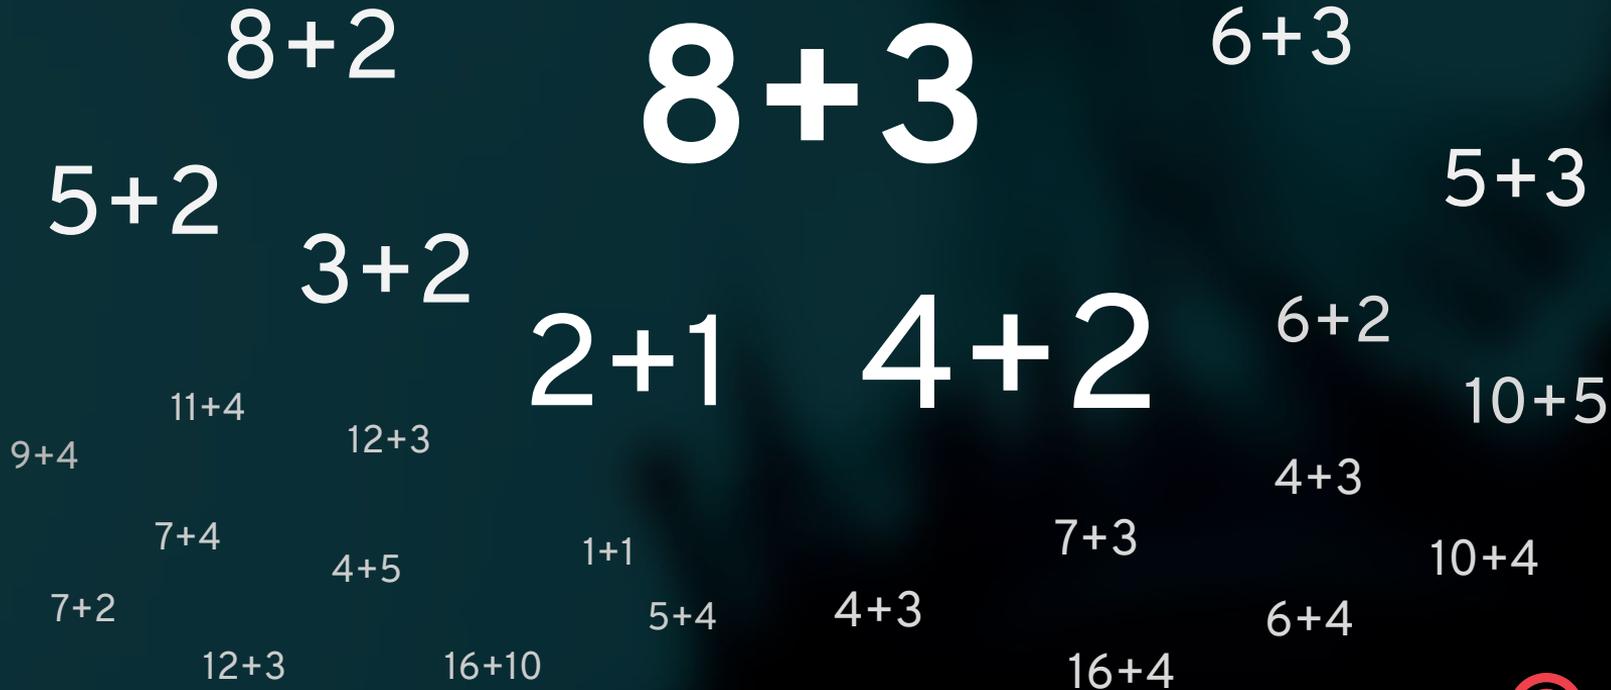
Which OSD backends are used?



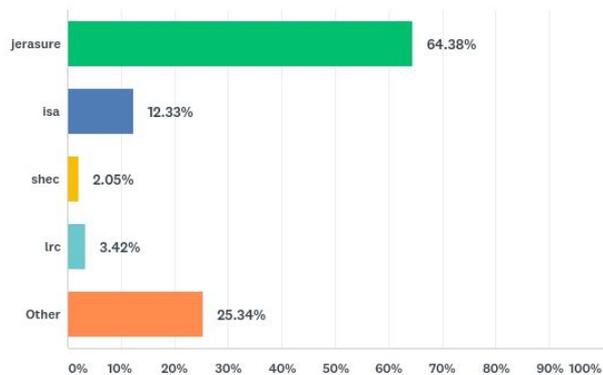
What kind of redundancy is used?



Erasure code '*k*' and '*m*' levels

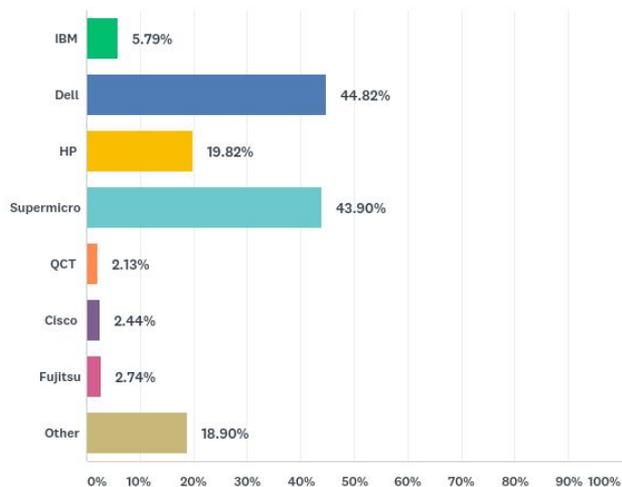


Which erasure code plugins do you use?

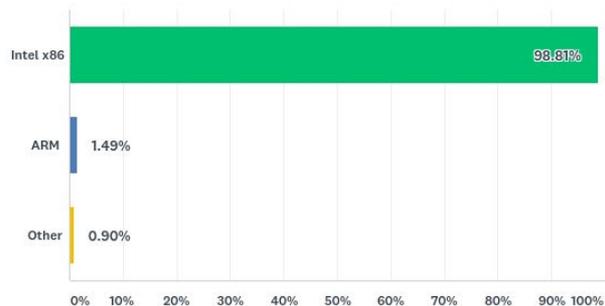


Hardware

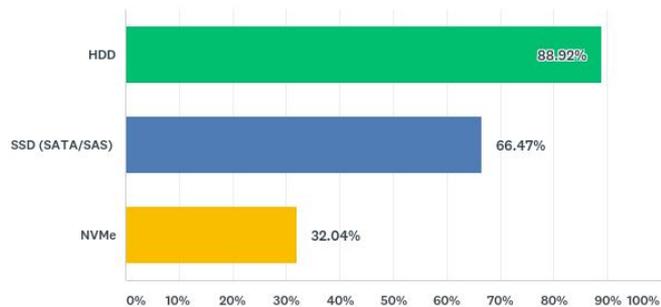
What hardware vendors do you use for the nodes?



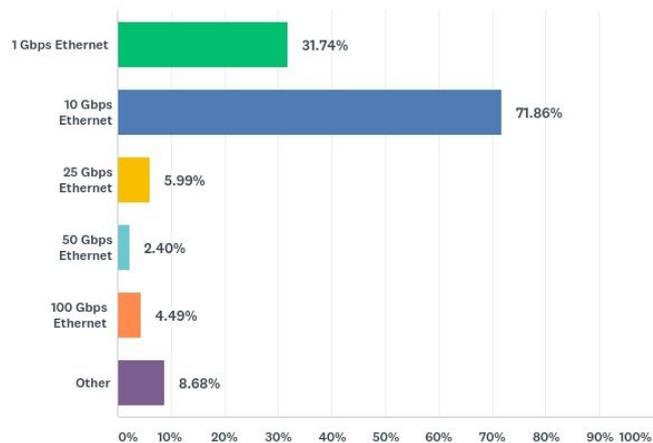
What processor architecture do you use?



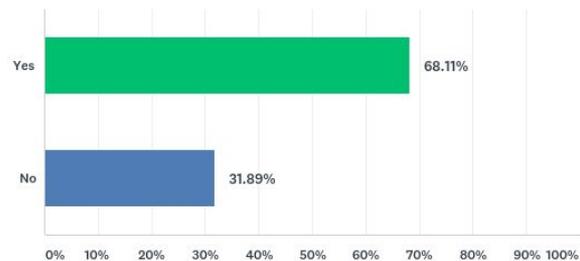
Which type of storage devices are used?



What type of NICs are used on the nodes?

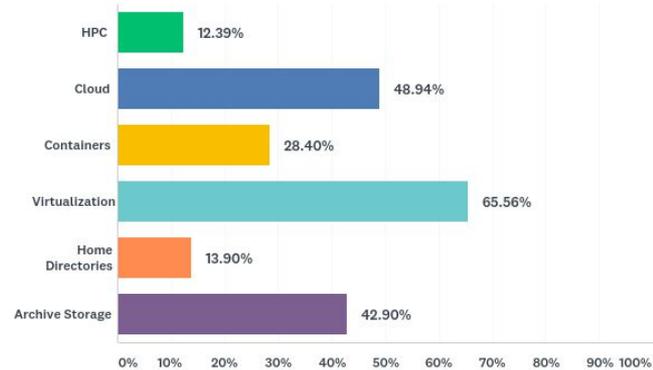


Do you use a dedicated network for the OSDs?

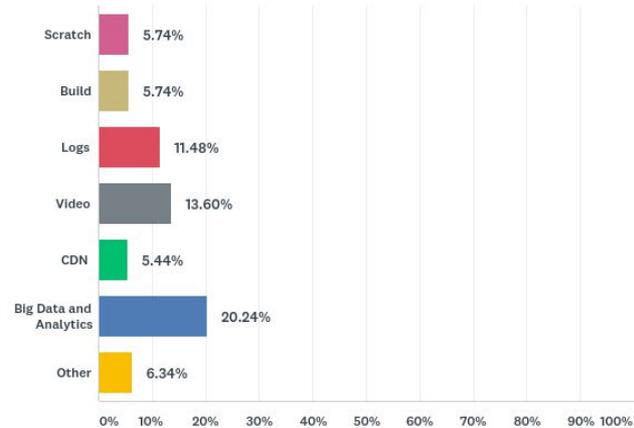


Use cases, platforms and interfaces

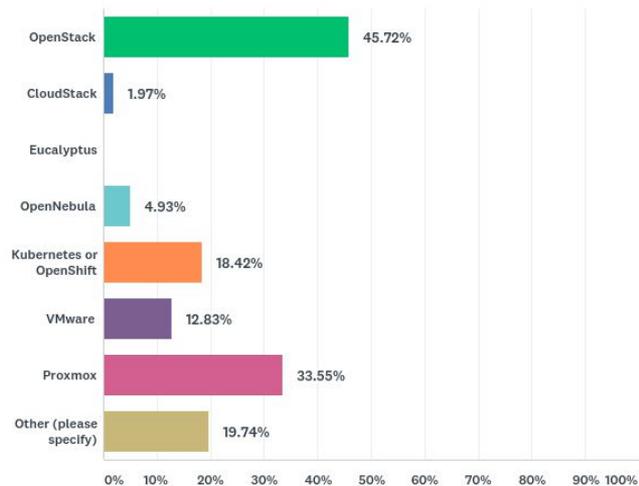
What are the use cases for the Ceph cluster?



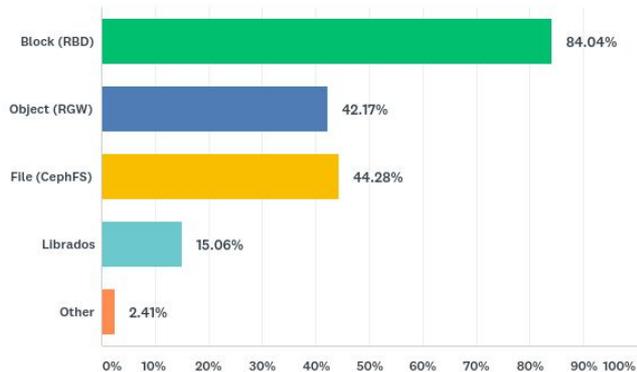
What are the use cases for the Ceph cluster? (cont)



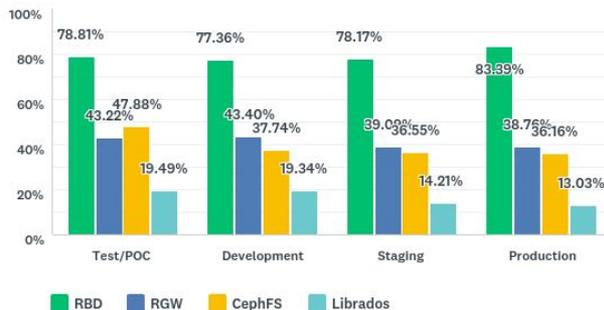
Which platforms do you use with Ceph?



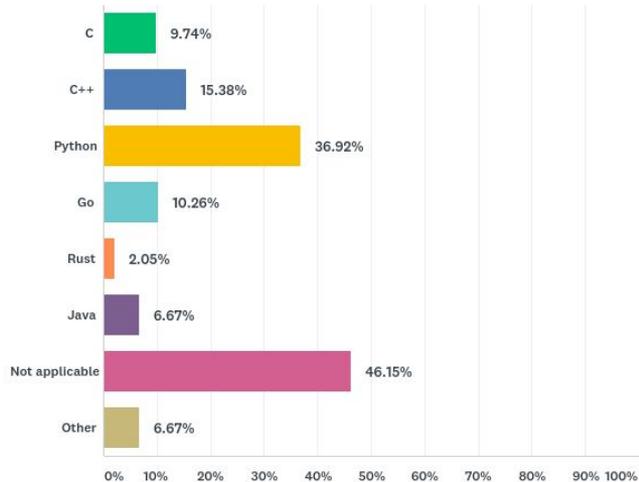
What are the interfaces used on the cluster?



Environments and Interfaces

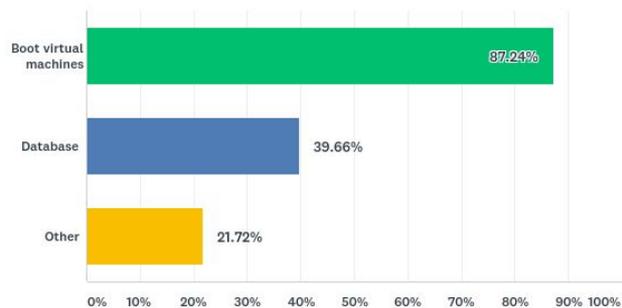


Programming languages and Librados/Librbd

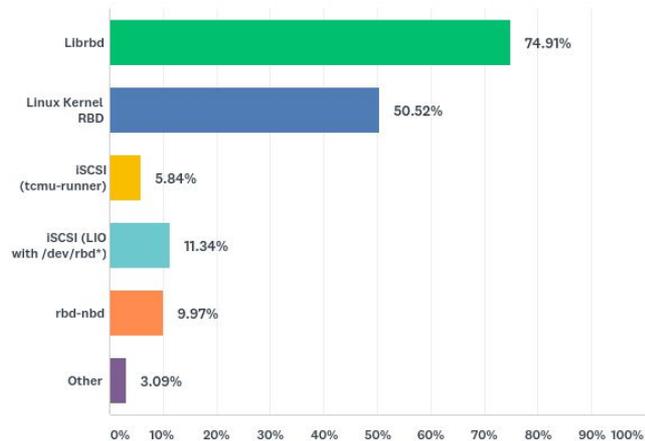


RADOS Block Device (RBD)

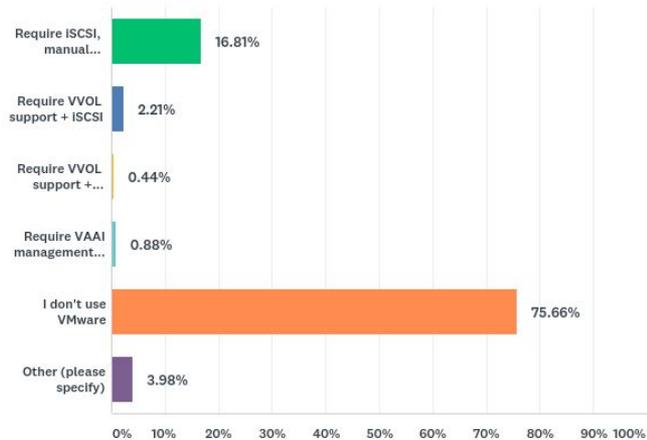
RBD workloads



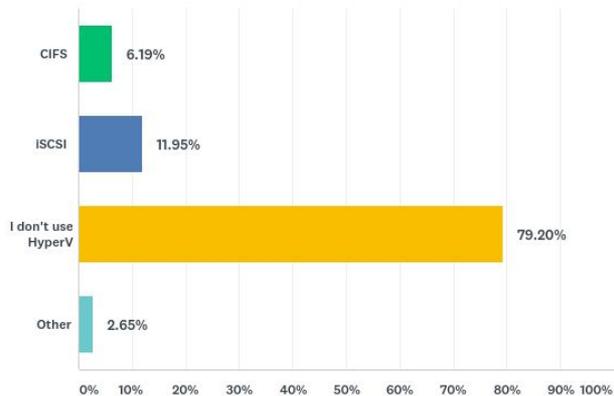
Interfaces to access RBD



VMware access methods

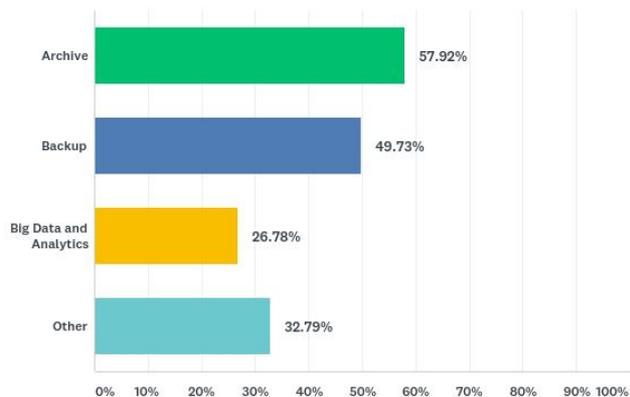


Hyper V (or Windows clients) access methods

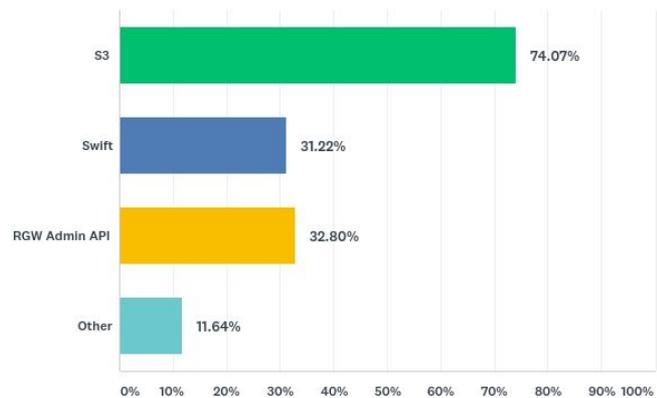


RADOS Gateway (RGW)

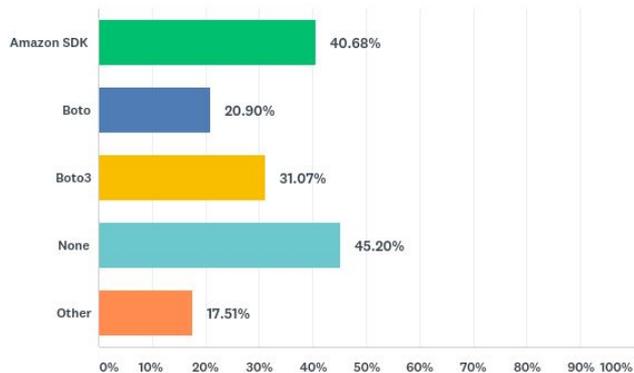
RGW workloads



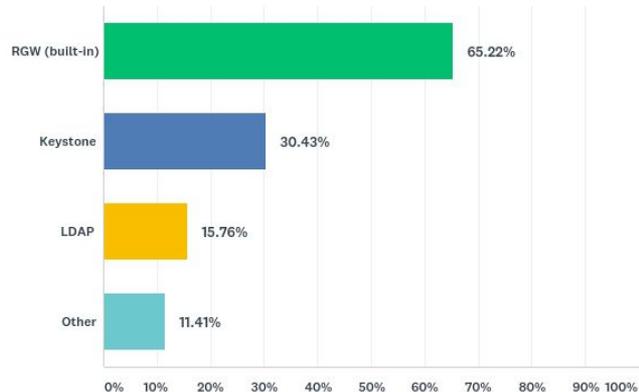
RGW APIs



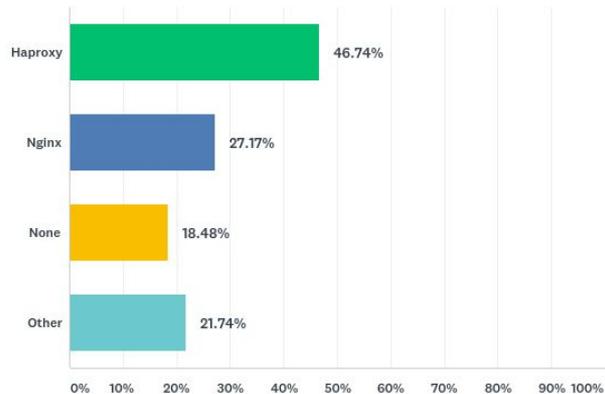
RGW client side libraries do you use?



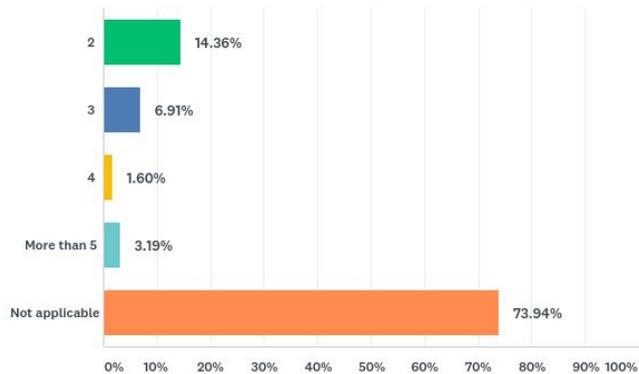
RGW authentication mechanisms



RGW external load balancers

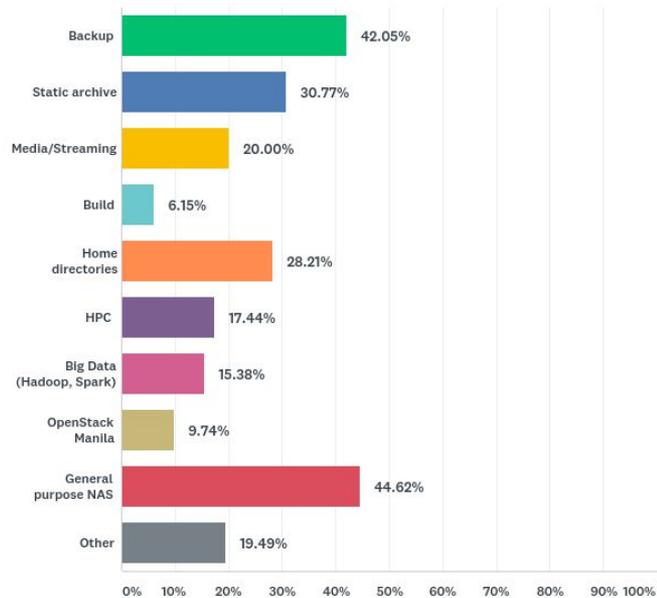


RGW federated multi-sites

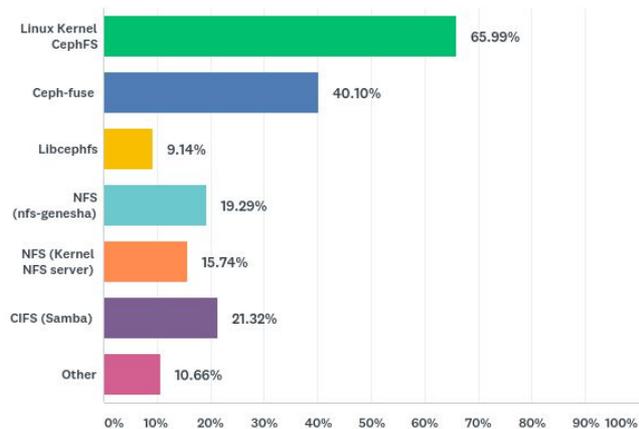


CephFS

CephFS workloads

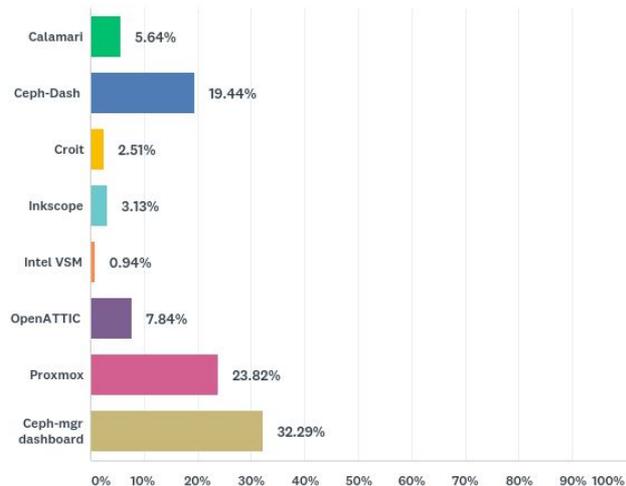


Interfaces to access CephFS

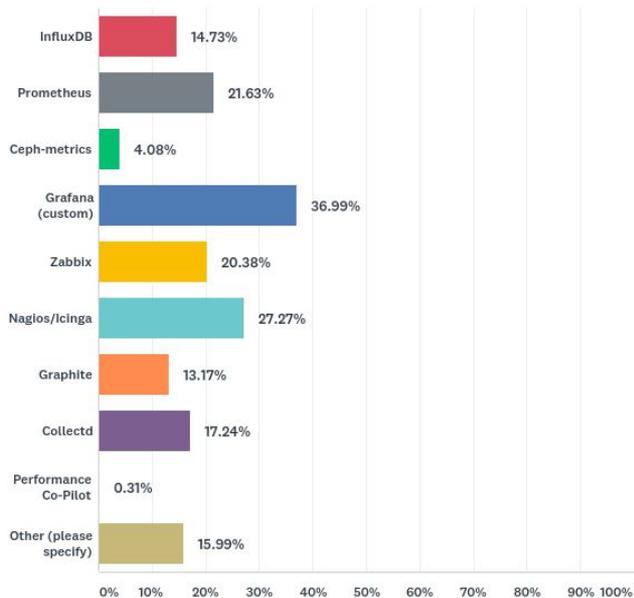


Management & Operation

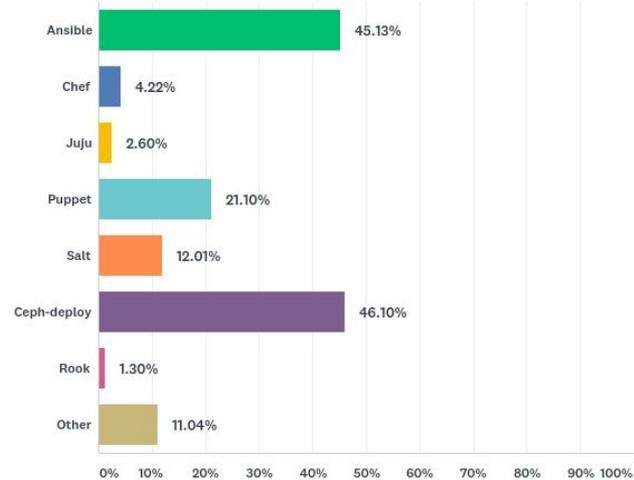
Management and monitoring tools



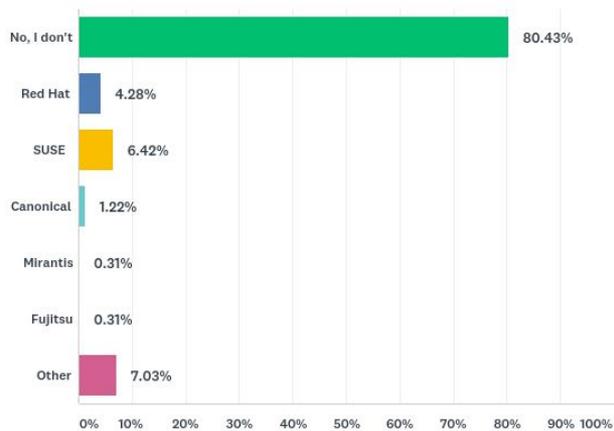
Monitoring and management tools (cont)



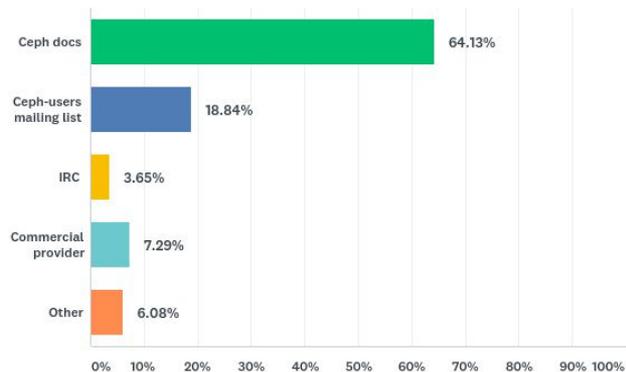
Deployment and configuration management tools



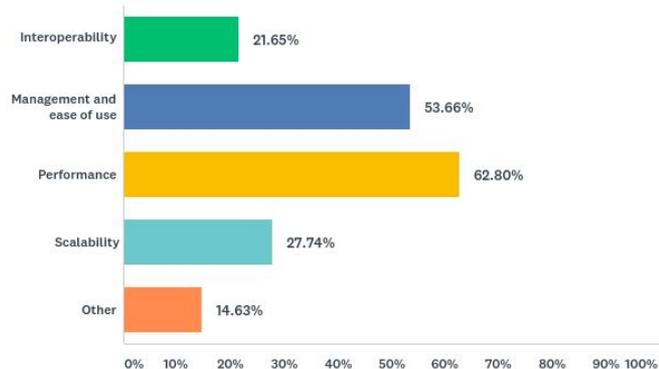
Ceph support from a commercial vendor?



If you need help, where is your first stop?



Where Ceph community should focus its efforts?



Ceph: the future of Storage