



4 FEB 2023

Introduction to s3gw

By Joao Luis <joao@suse.com>





1. Why?
Why we are doing this, our objectives
2. What?
What is the S3 Gateway?
3. How?
Architecture overview
4. Demo
How it looks like now
5. Outlook
Next steps, where we are going
6. Q&A



Why are we doing this?



Copyright ©SUSE 2023



- Easy to deploy, run and forget, on-premise S3 service
- There's a need in the Rancher stack for a lightweight S3 solution



What is s3gw?



Copyright ©SUSE 2023



s3gw is....


- An Open Source Project (Apache License)
- Driven by the Storage team at SUSE
- Easy-to-use Open Source and Cloud Native S3 service for use on Kubernetes

... or will be ...



s3gw complements the **Rancher** portfolio:

- Backups
 - Longhorn volumes, Harvester & Epinio
- Leverages k3s/k8s and Longhorn
 - redundancy, load balancing, etc
 - S3 interface for Longhorn
- Kubernetes-native management (helm chart, GitOps)
- Pretty end-user UI for operations
- Potential in Edge, IoT and smaller on-premise deployments



How?





Copyright ©SUSE 2023



- We are leveraging Ceph's feature-rich S3 gateway (RGW)
- RGW's built-in standalone capabilities
- Taking advantage of RGW's Zipper layer
- File-based RGW Driver Backend







bucket

foo

bar



bucket ID

e8cc2d3c-b44f-4c76-bbed-1dfcff54041b

32b62ac3-bc97-4f66-942f-6c2601b1564d

bucket

bar

bar

foo

object

abc

def

qwe







object ID


dc8b0729-946e-479f-acfe-016ac89a8e05

9c19d78e-f72c-4029-bd3e-cf96d54cc65d

5bd808eb-e57d-482b-b1c2-5001474a7cae







Demo



Copyright ©SUSE 2023



Next steps



Copyright ©SUSE 2023



Roadmap...

Key roadmap highlights:

- Increase support for S3 operations
- Lifecycle management
- Retention policies
- Performance improvements
- Statistics for the UI




Challenges

- Semantic compliance with the S3 API
 - Tricky to get operations *just right*
- Performance
 - Learning curve, including with SQLite's capabilities



v0.10.0 vs some perf work


Throughput



read-bw-mean (byte/s) 1.26x

write-bw-mean (byte/s) 1.26x

Operations




read-iops-mean (iops) 1.26x

write-iops-mean (iops) 1.26x




v0.10.0 vs some perf work

v0.10.0




op	fastest [ms]	slowest [ms]	avg [ms]	median [ms]
GET 6	6	177	10	9




op	fastest [ms]	slowest [ms]	avg [ms]	median [ms]
PUT	1624	2809	1957	1947

yolo



op	fastest [ms]	slowest [ms]	avg [ms]	median [ms]
GET 6	6	183	19	18



op	fastest [ms]	slowest [ms]	avg [ms]	median [ms]
PUT	1163	2202	1520	1503



And after that...

- Who knows?

World
Domination!!



Q&A

Where to find us

<https://s3gw.io>

<https://github.com/aquarist-labs/s3gw>

<https://lists.suse.com/mailman/listinfo/s3gw>

Joao Eduardo Luis

joao@suse.com





Thank you

For more information, contact SUSE at:

+1 800 796 3700 (U.S./Canada)

+49 (0) 911-740 53-0 (Worldwide)

Maxfeldstrasse 5

90409 Nuremberg

www.suse.com

© 2020 SUSE LLC. All Rights Reserved. SUSE and the SUSE logo are registered trademarks of SUSE LLC in the United States and other countries. All third-party trademarks are the property of their respective owners.