

Exploring Oracle Cloud for Disaster Recovery

Victor Virag

Project Supervisor: Hubert Odziemczyk



Introduction

The current disaster recovery system performs daily backups to CERN S3, with replication across the Meyrin and Prévessin DCs.


We are backing up data of services running on the Java Tomcat platform, e.g. EDMS (Engineering & Equipment Data Management Service)

Performing immutable backups to a cloud provider could increase resilience against threat actors and simultaneous failure of existing backup systems.

Threats:



Meyrin DC



Prévessin DC



High-Level Overview

Backup Pipeline

- Restic backup to CERN S3, then rclone replication to OCI S3

Restoration Pipeline

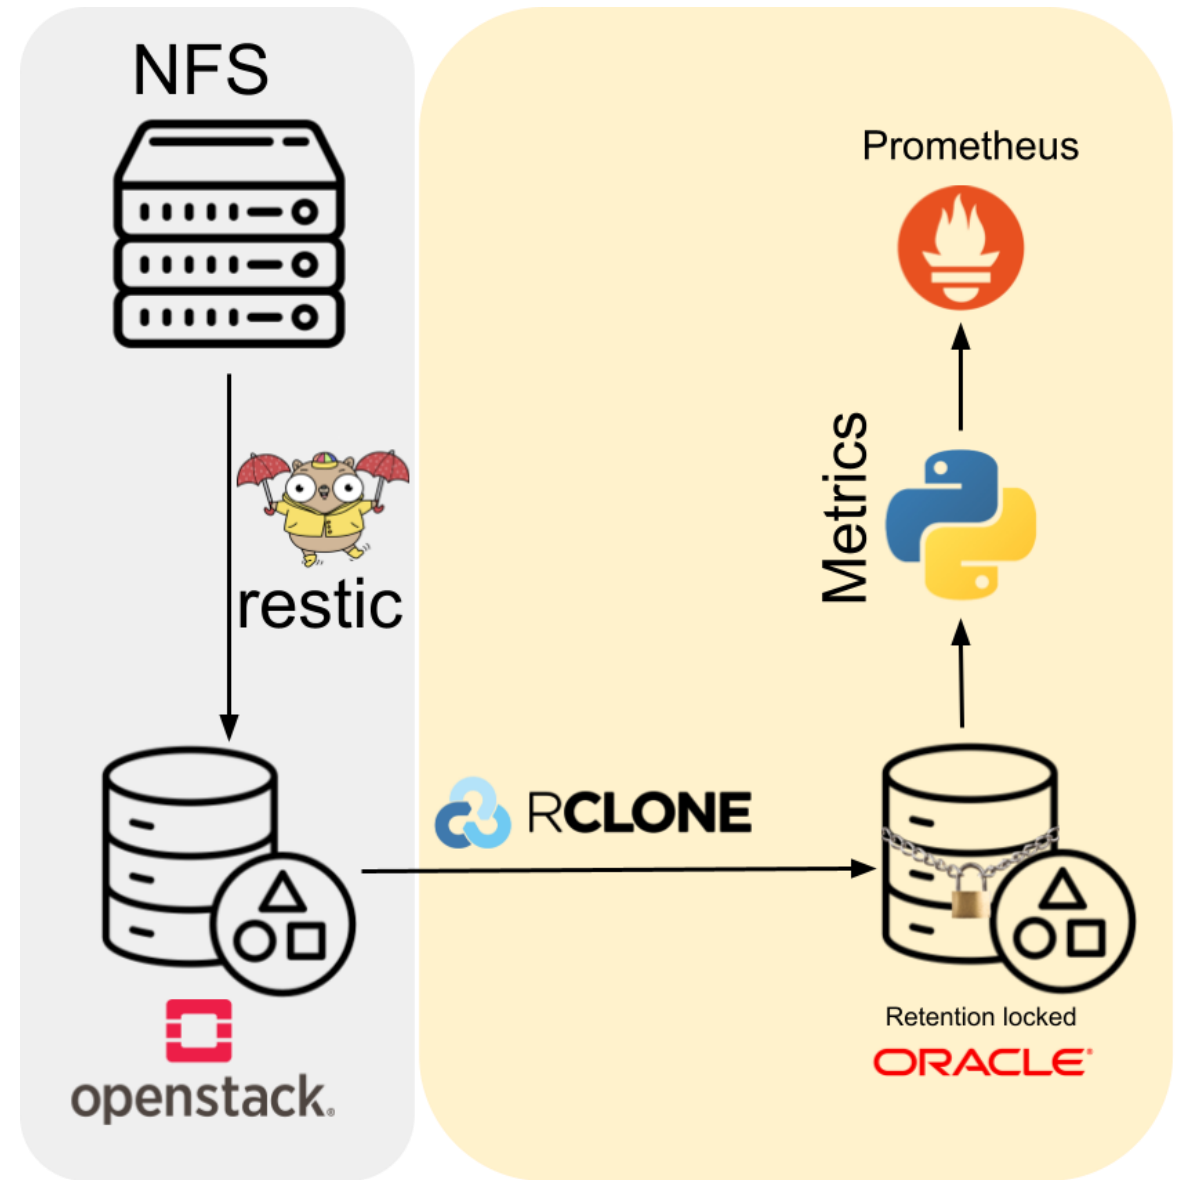
- Archive storage retrieval (1h wait time)
- Direct restic restoration from Oracle S3 copy

Backups:

- Encrypted and authenticated backups (AES-256-CTR)

Object Storage:

- Archive storage tier has 1h wait time before objects are available for reading



Existing solution

New addition

Performance & Challenges

Performance:

- Initial backup overhead for 18 TiB transfer to OCI cloud takes > 24h
- Subsequent backups are incremental and small (~10GB daily, transfer time < 5 min)

Challenges:

- Backup source location must be trusted
- Backup keys must be stored securely and independently of backups
- Cloud and local backup configurations should ensure that losing access to a single admin account does not compromise both backup systems



+



+

ORACLE®

Contributions

Monitoring:

- Integrated Prometheus monitoring with Oracle Cloud S3

Backup pipeline:

- Created backup and restoration pipeline using restic and rclone to immutable Archive tier storage

Periodic backup size reduction:

- Created strategy for avoiding build-up of unused old backup files in immutable storage

CERN S3 - Oracle S3 comparisons

- Researching and testing differences between CERN S3 and Oracle S3

Benchmarking:

- Performance tests for different backup pipeline configurations and backup sizes

	Archive storage	Versioning + retention	Immutable
CERN S3	No	Yes	Yes*
OCI S3	Yes	No	Yes

*can be modified by a high-privileged admin role

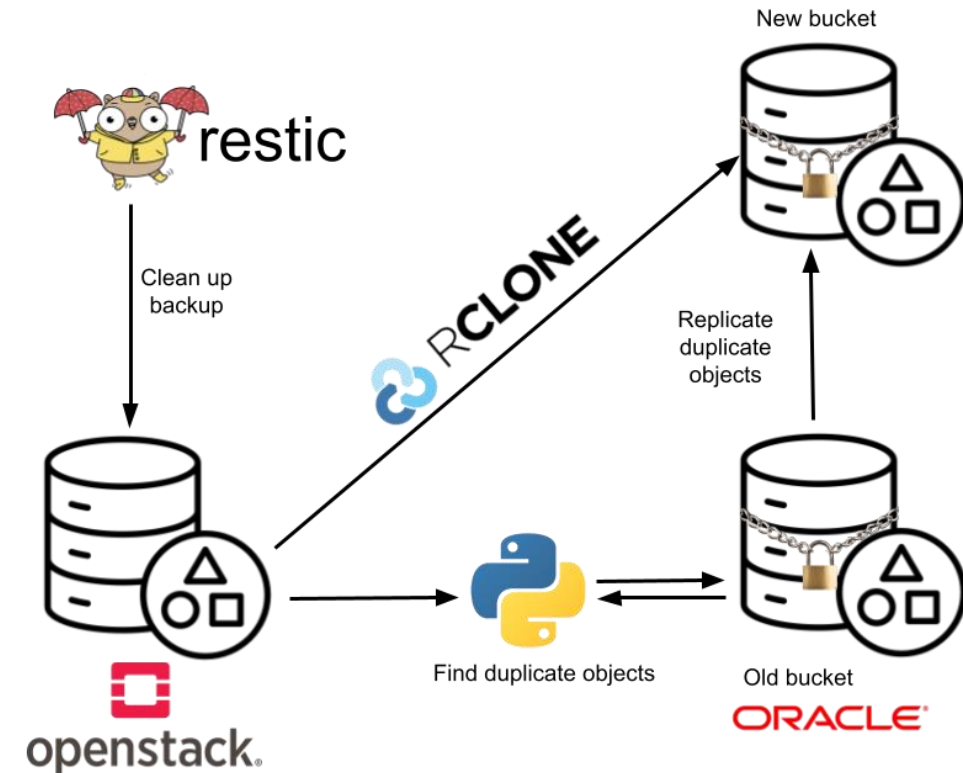


Contributions (clean-up part)

Problem: How do we remove unused files from an immutable backup?

Solution:

- We don't actually delete anything, only copy over the files we still need to a new bucket
- First we get a list of files that we want to retain, then we copy them from the old bucket to the new one
- We don't copy from CERN S3 directly to the new bucket because Oracle bucket to bucket copy is much faster (same datacenter)
- The new bucket now becomes the replication target




Conclusion

This project has:

- Improved understanding of how Oracle Cloud could be integrated with the current disaster recovery strategy
- Developed a collection of automation tools and scripts to simplify future work
- Benchmarked performance and extrapolated scalability metrics
- Identified and documented challenges and potential solutions

Thank you!



References

Images used (here for copyright reasons):

<https://www.pngwing.com/en/free-png-hykgx>

<https://thenounproject.com/icon/object-storage-5318508/>

<https://www.cleanpng.com/png-oracle-corporation-oracle-database-logo-netsuite-l-2528427/>

<https://www.logo.wine/logo/OpenStack>

<https://icon-icons.com/icon/rclone-logo/169806>

<https://forum.restic.net/uploads/default/original/1X/d8992cda487db3fe87f027a9a7d05477816bd359.png>

<https://pngimg.com/image/58258>

https://www.usinenouvelle.com/mediatheque/5/0/4/001510405_896x598_c.jpg

https://media.sciencephoto.com/image/c0570336/800wm/C0570336-CERN_data_centre.jpg

https://pngtree.com/freepng/security-concept---padlock-on-chain-isolated-isolated_13231738.html

<https://seeklogo.com/images/P/prometheus-monitoring-system-logo-3C07B67C00-seeklogo.com.png>

<https://www.pngwing.com/en/free-png-dqlwx>