Hybrid Cloud Storage

E. Ugur Kaynar, Amin Mosayyebzadeh, Mania Abdi, Mohammad Hajkazemi,
Brett Niver, Matt Benjamin, Ali Maredia,
Peter Desnoyers, Larry Rudolph, Orran Krieger

D3N: Object Store Caching for a Data Center

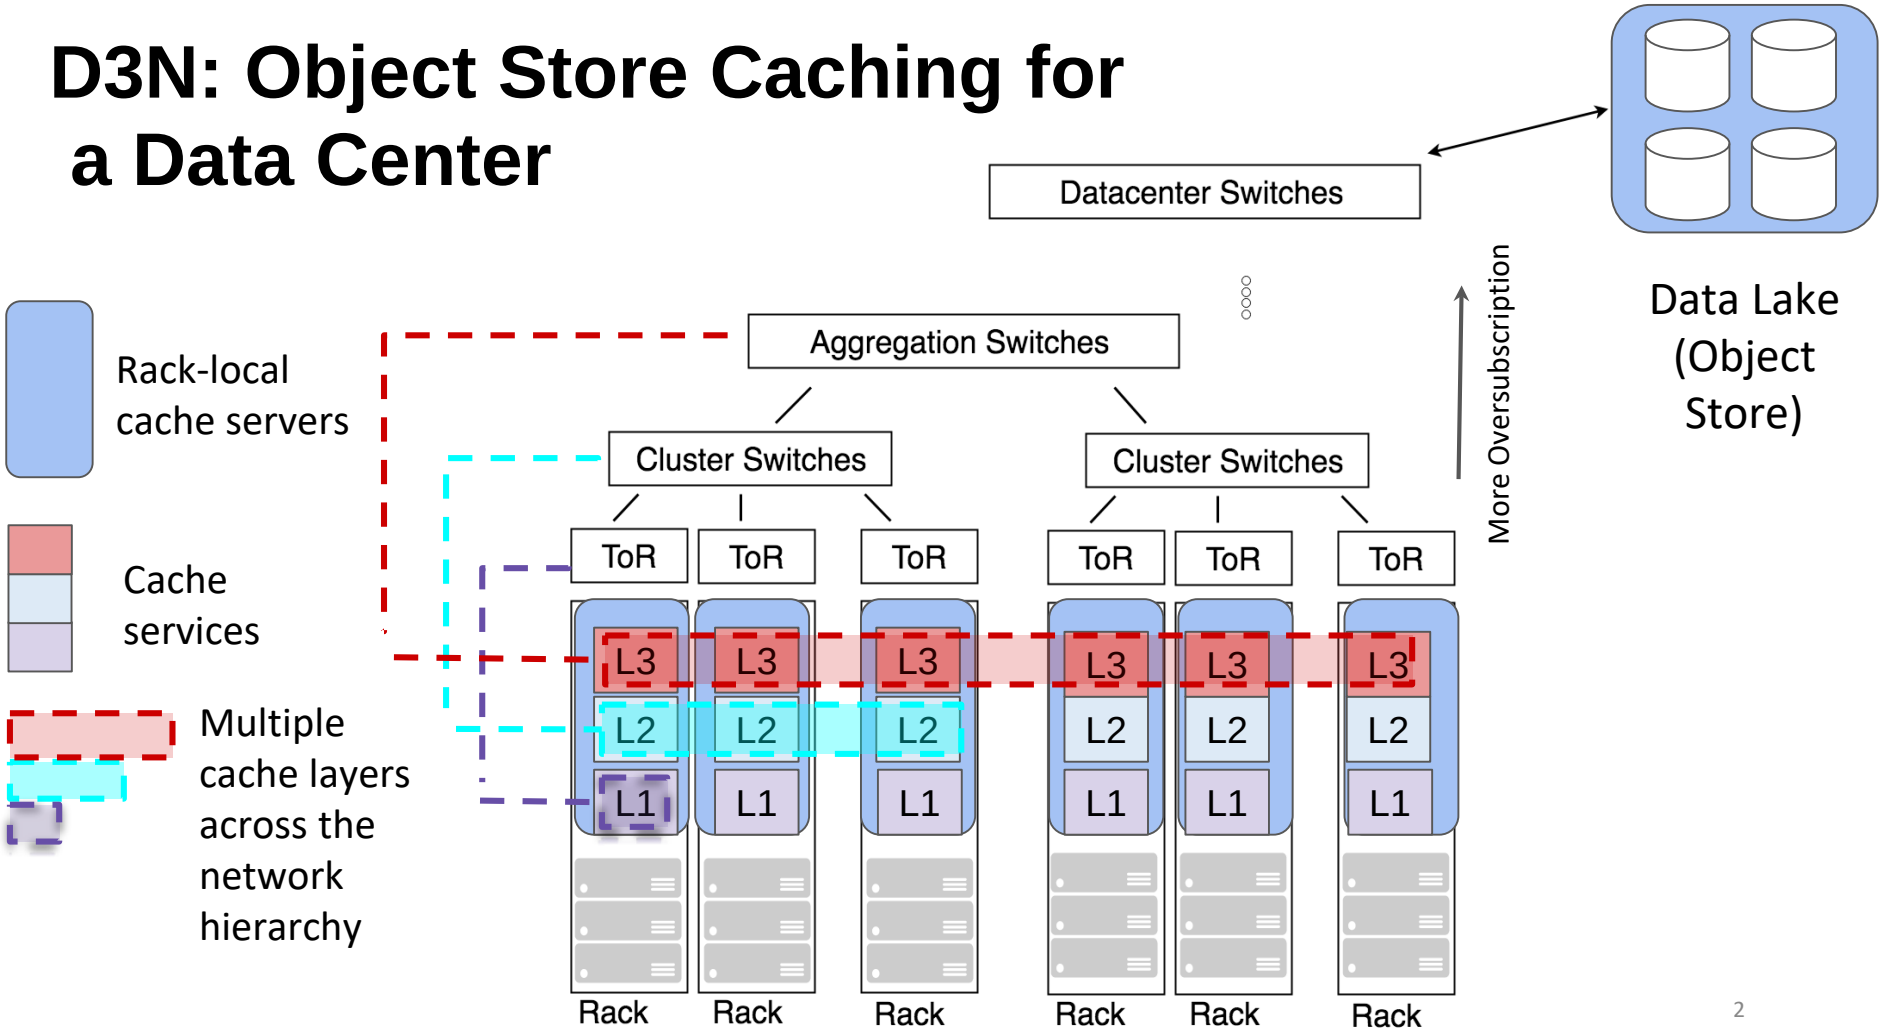

Impact of D3N on Realistic Workload

The trace completion time

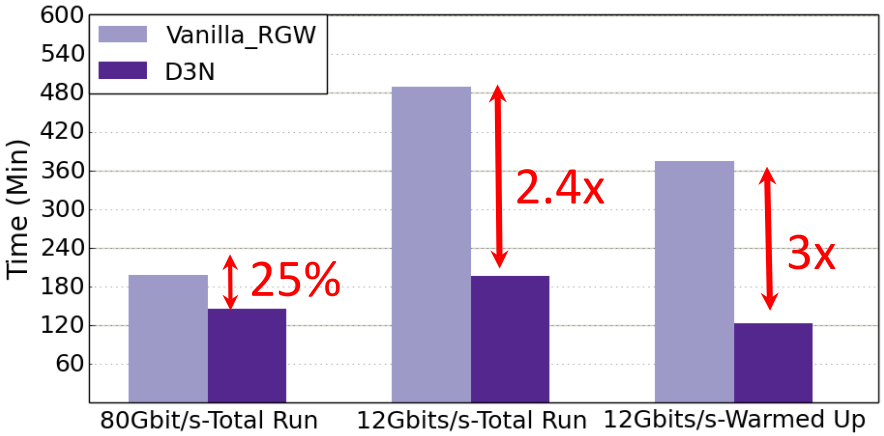

Cumulative data transferred from back-end storage

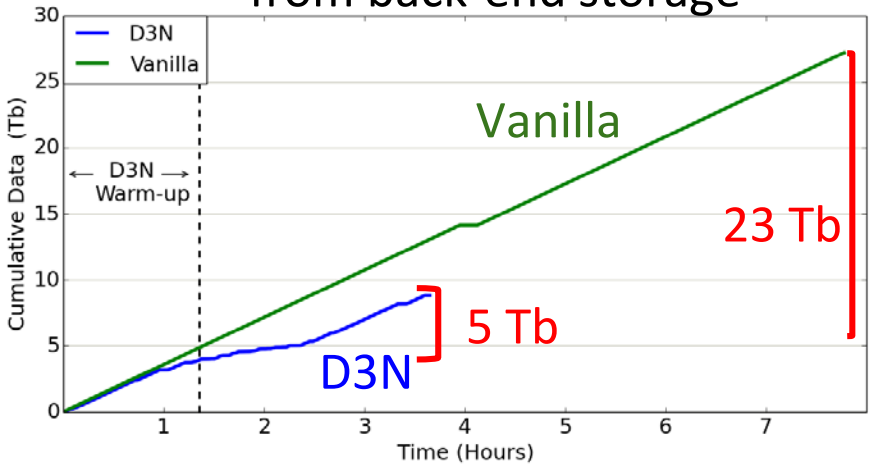

D3N improves performance significantly.

More than 4x reduction to backend traffic.

Towards Hybrid Cloud

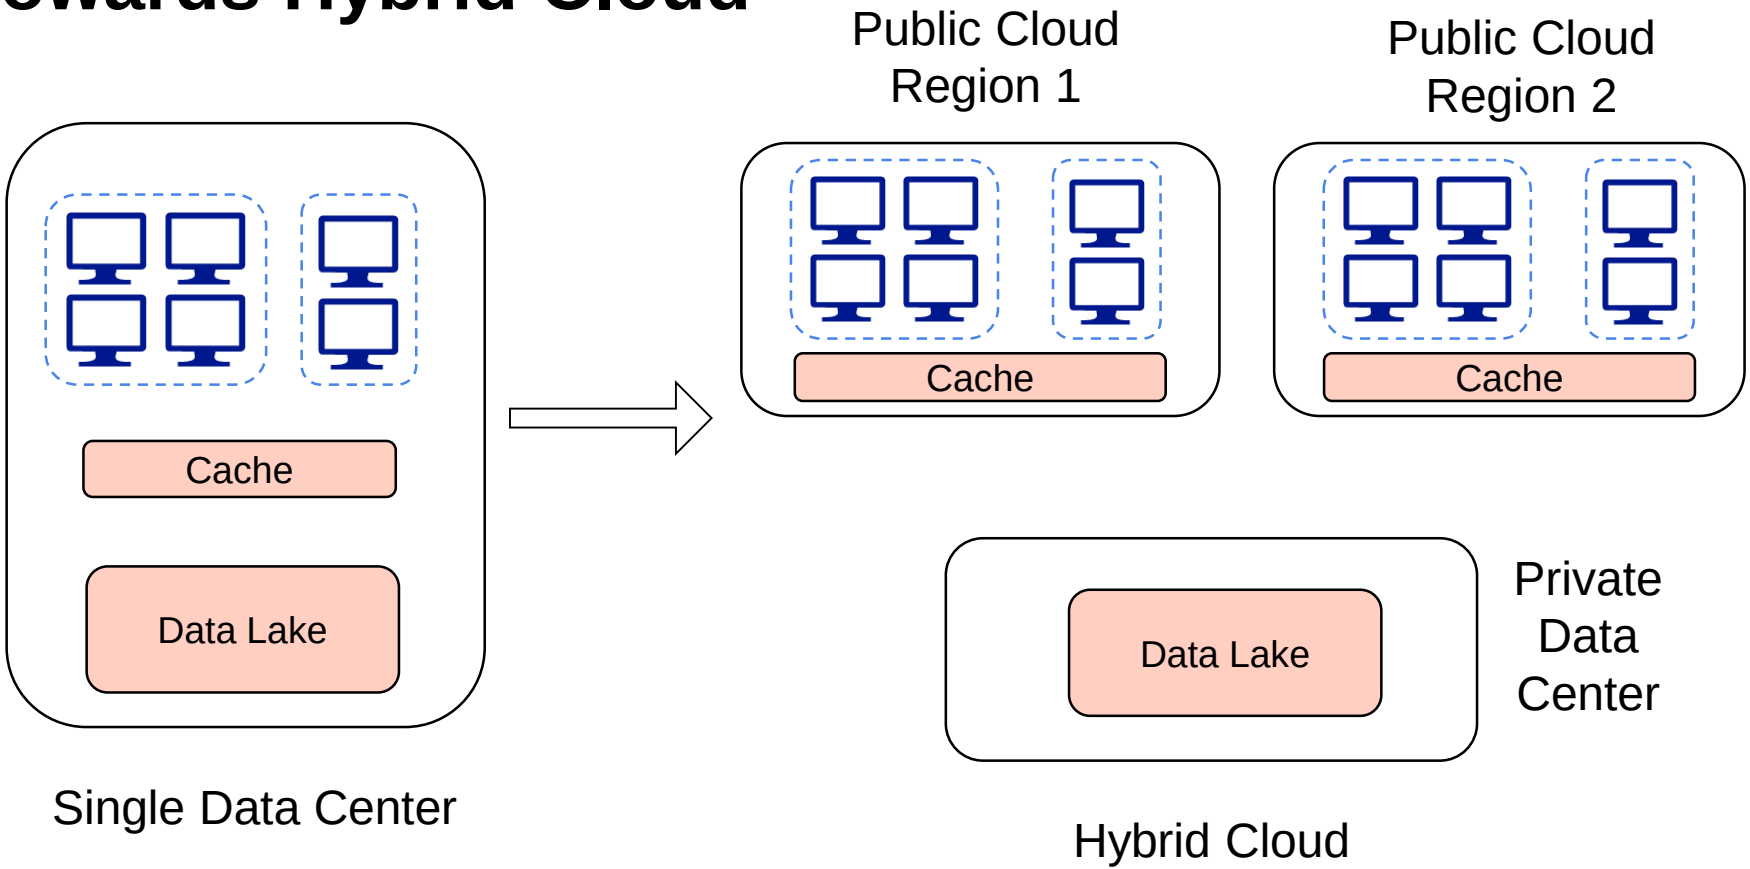

What needs to be changed

- Read-only and intermediate datasets
- Consistent hashing
- Anycast based lookup service
- Internal protocol called *rados* is used to access the data lake

- Write-back durable cache
- Directory-based approach
- Kubernetes
- *S3* protocol is used to access the data lake

Implementation

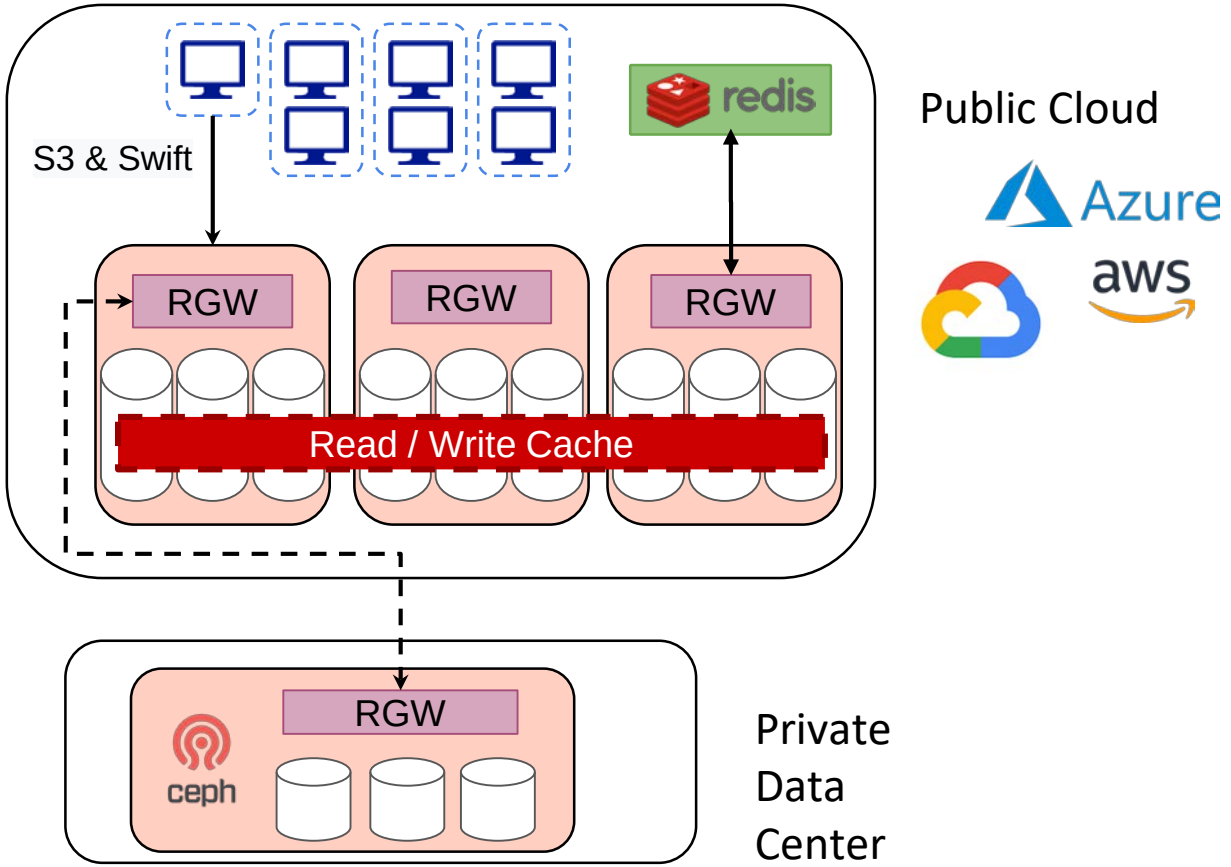

Where are we going?

- Directory enables richer policies
- Multiple back end data lakes
  - Where to store the data
  - Erasure coding
- Where do we compute on the data
- Prefetching

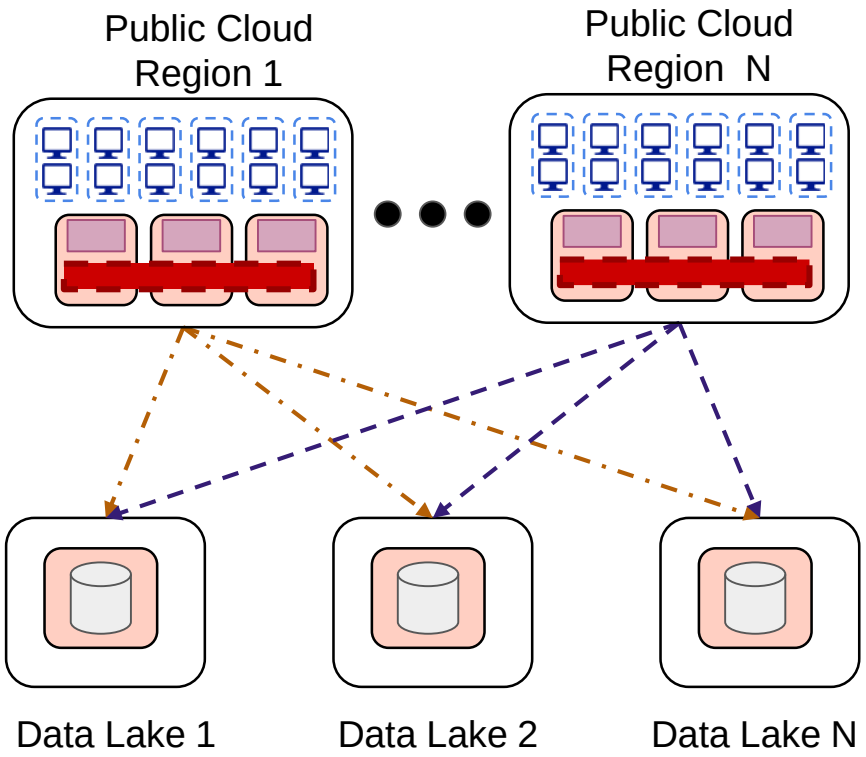

More Information

- Red Hat is upstreaming the initial prototype.
  - https://github.com/ceph/ceph/pull/24500
- Hybrid Cloud Cache implementation:
  - https://github.com/ekaynar/ceph
- Project Websites
  - https://www.bu.edu/rhcollab/projects/d3n/
  - https://massopen.cloud/d3n/

Thank you