



RED HAT[®]
STORAGE

STORAGE UPDATE

MSP RHUG

James Rankin

Senior Solutions Architect

AGENDA

- Red Hat Storage 3.0 (Gluster)
 - Quick overview
 - 3.0 Changes
 - 3.1 Preview
 - Use cases
- Inktank Ceph Enterprise 1.2 (Ceph)
 - Quick overview
 - 1.2.x Changes
 - Upcoming feature preview
 - Use cases



RED HAT STORAGE (GLUSTER)

RED HAT STORAGE

OPEN

Open, software-defined distributed file and object storage system

- Based on GlusterFS open source community project
- Uses proven local file system (XFS)
- Data is stored in native format

SCALABLE

No metadata server

- Uses an elastic hashing algorithm for data placement
- Uses local file system's extended attributes to store metadata
- Nothing shared scale-out architecture

ACCESSIBLE

Multi-protocol to the same data

- Global name space
- NFS, SMB, object, HDFS, Gluster native protocol
- Posix compliant

MODULAR

No kernel dependencies

- GlusterFS is based on file-system in user space (FUSE)
- Modular stackable architecture allows easy addition of
- Features without being tied to any kernel version


ALWAYS-ON

High-availability across data, systems, and applications

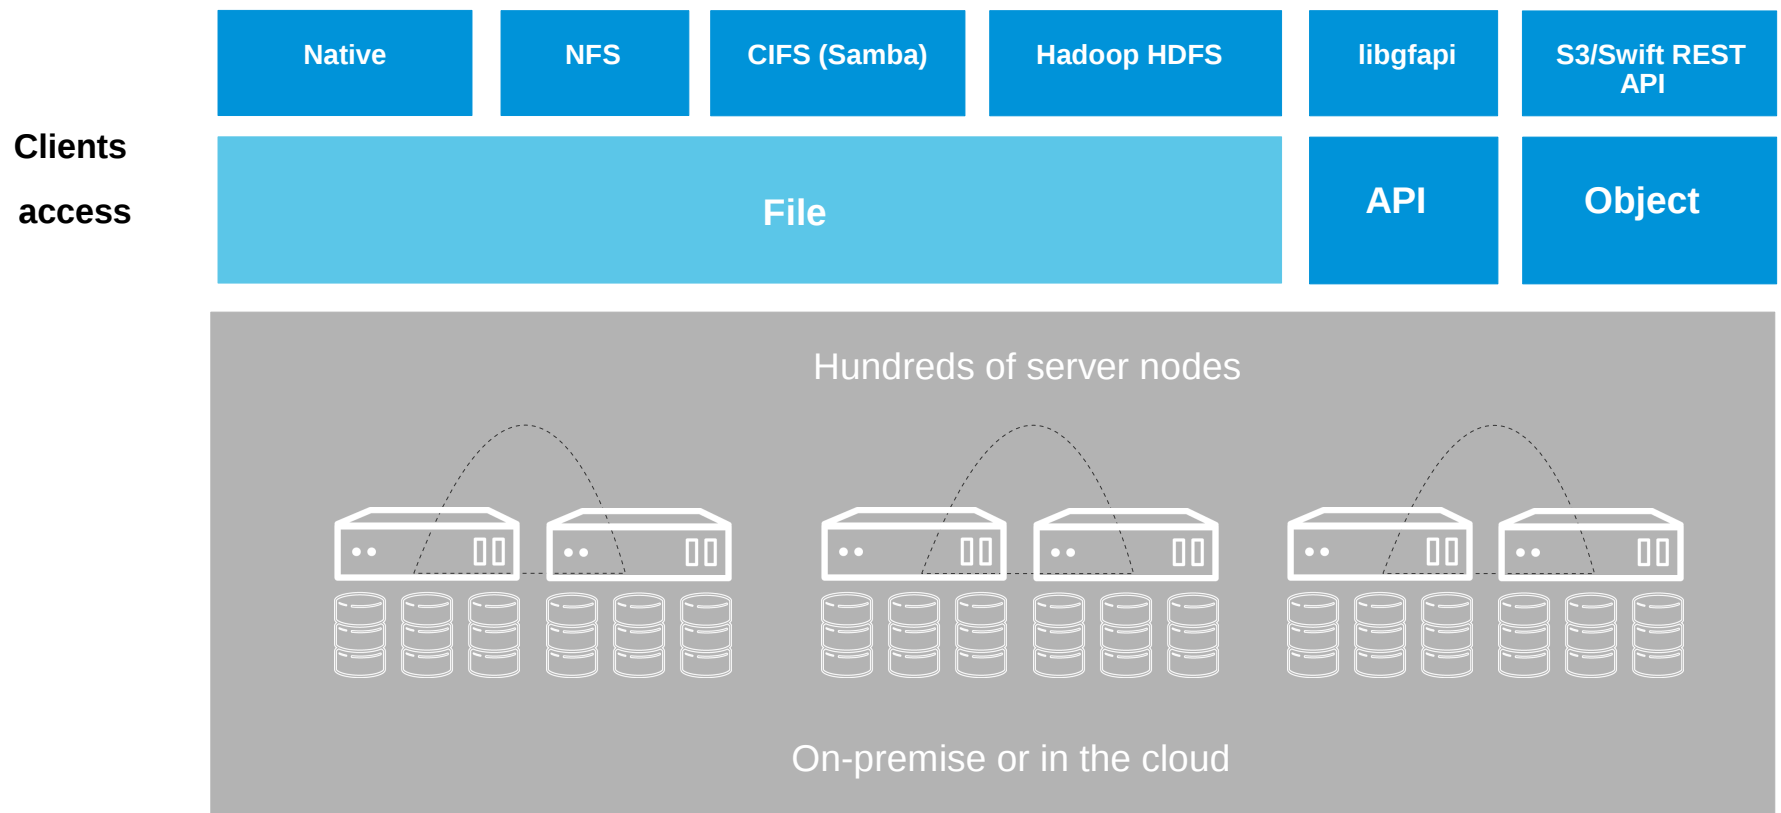
- Synchronous replication with self-healing for server failure
- Asynchronous geo-replication for site failure

ELASTIC HASH ALGORITHM

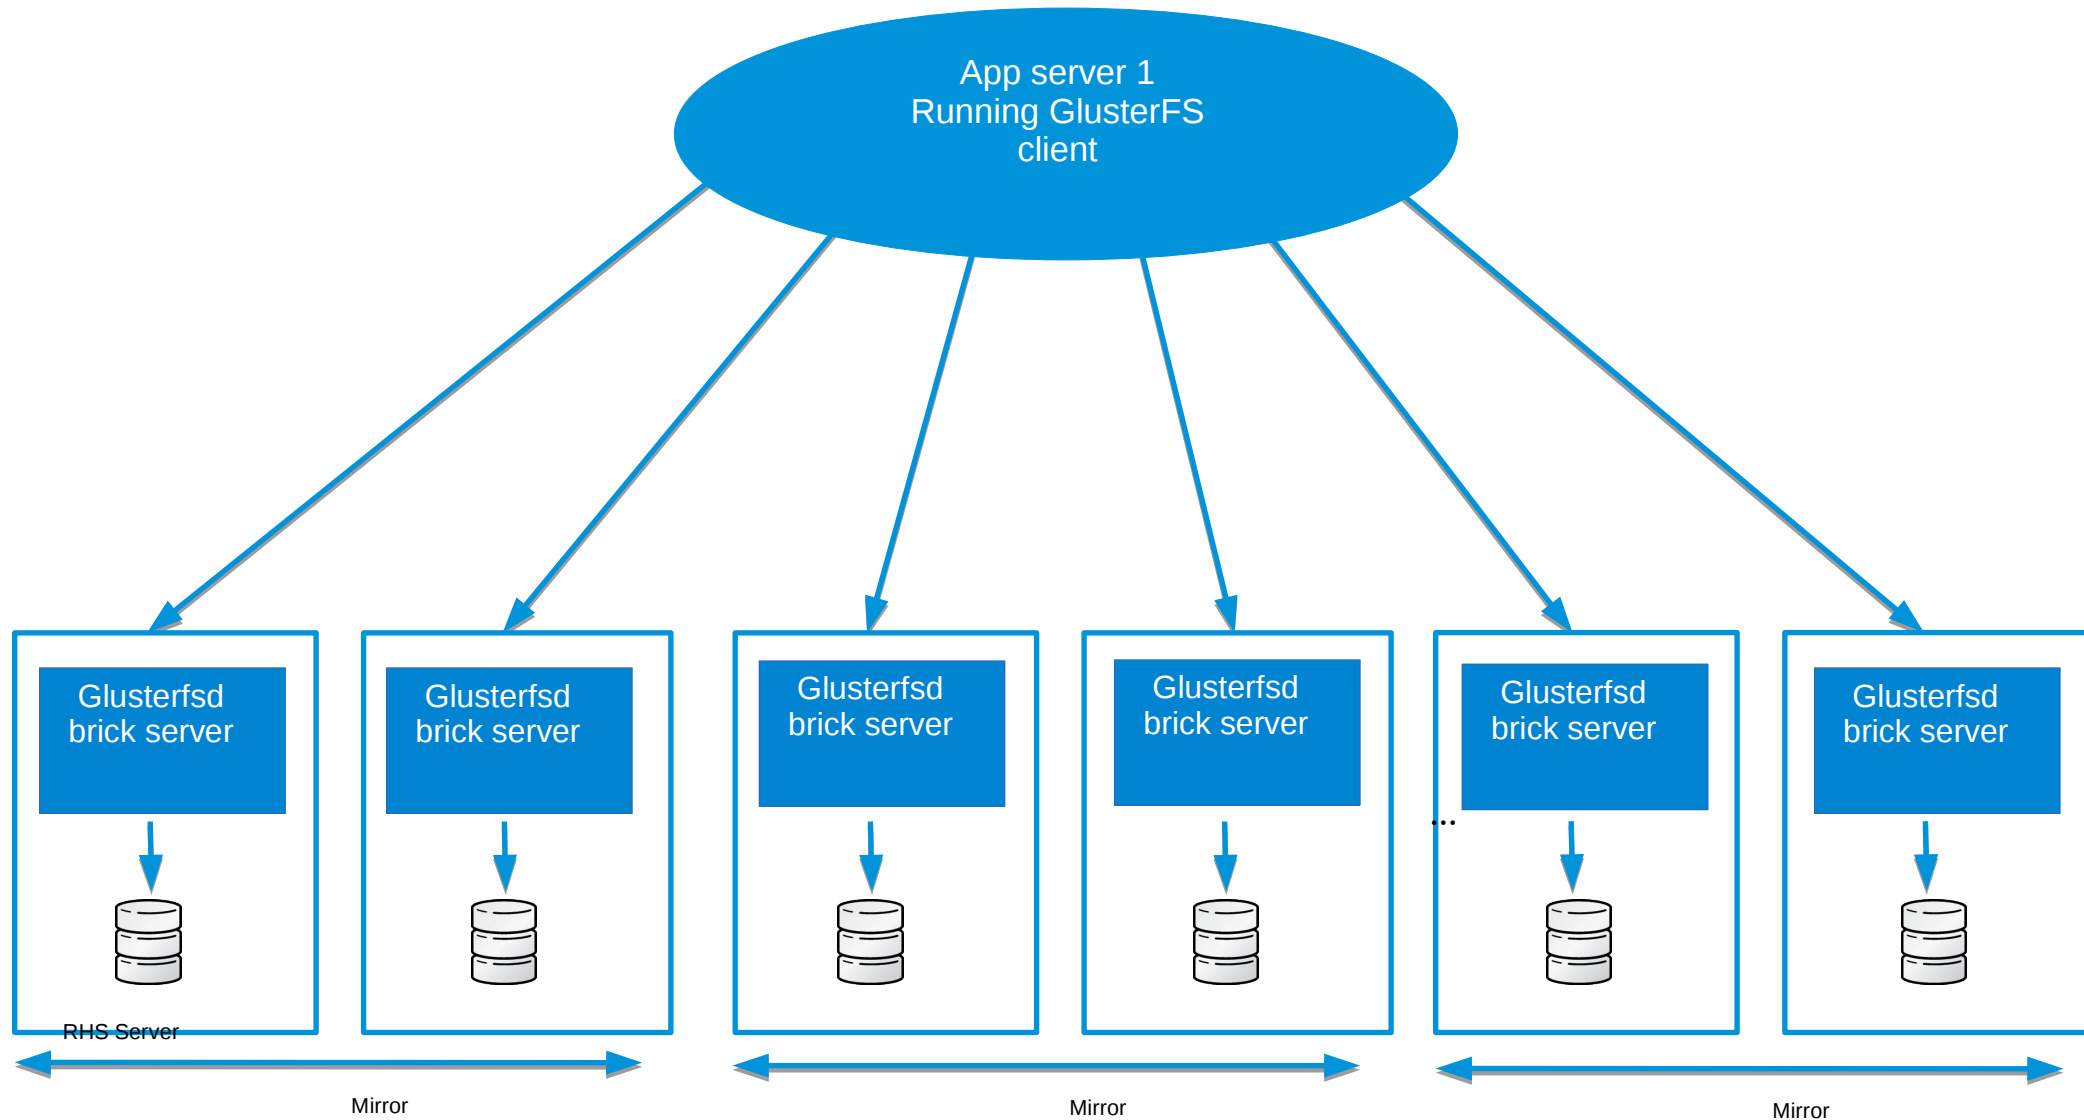
- No central metadata
 - No performance bottleneck
 - Eliminates risk scenarios
- Location hashed on file name
 - Unique identifiers, similar to md5sum
- The elastic part
 - Files assigned to virtual volumes
 - Virtual volumes assigned to multiple bricks
 - Volumes easily reassigned on the fly



MULTI-PROTOCOL ACCESS



GLUSTERFS NATIVE CLIENT – DATA PATH



RHS 3.0 LAUNCH

- Launched last fall, currently on 3.0.2
- Volume snapshots
 - User-servicable snapshots
- Nagios based monitoring and alerting
- Scalability to 60 disks per node, 128 nodes per cluster
- Hadoop enablement with Hortonworks
- Tech Preview: NFSv4 support via Ganesha

RHS 3.1 PREVIEW

- Coming this summer, in development upstream
- Erasure Coding
- NFSv4 (full support)
- Bit rot detection
- Small file performance
- SMBv3

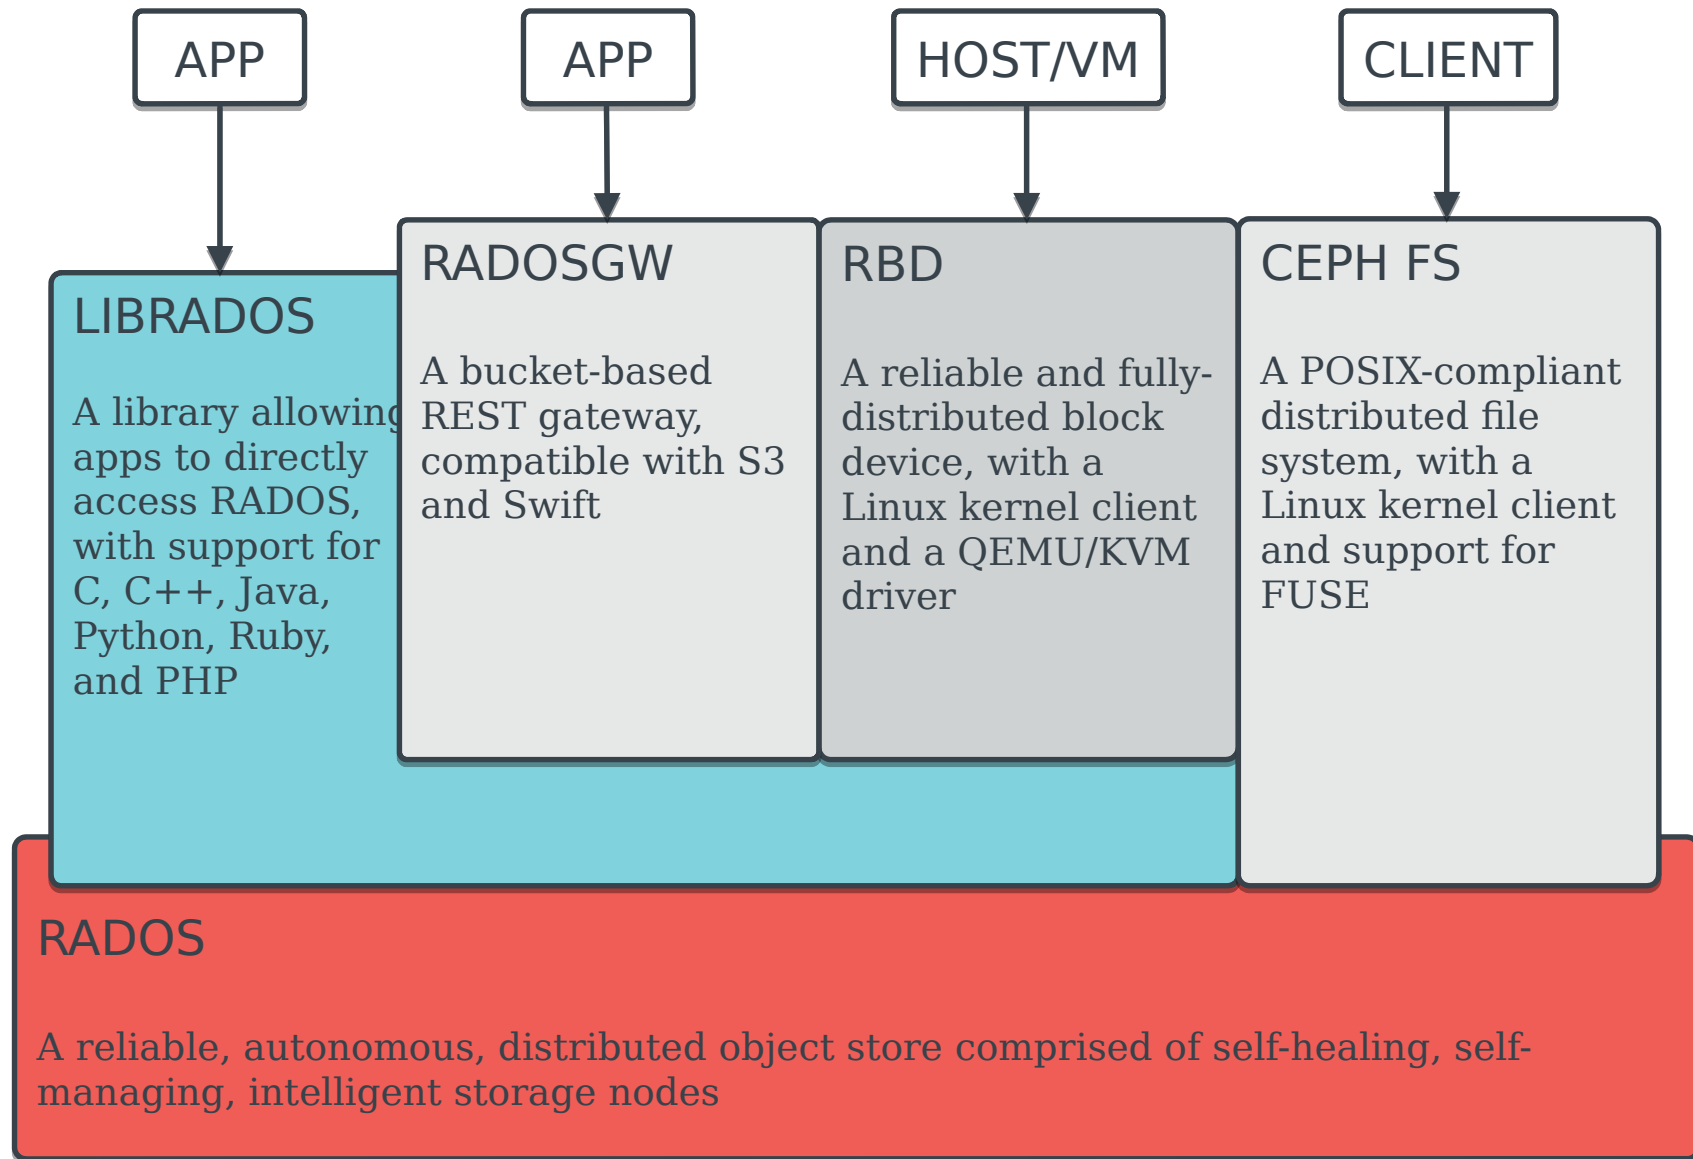
USE CASES

- **Splunk**
- **Hadoop**
- Amazon Web Services
- RHEV storage domains
- Backup/archive target
- Etc (unstructured, file workloads)

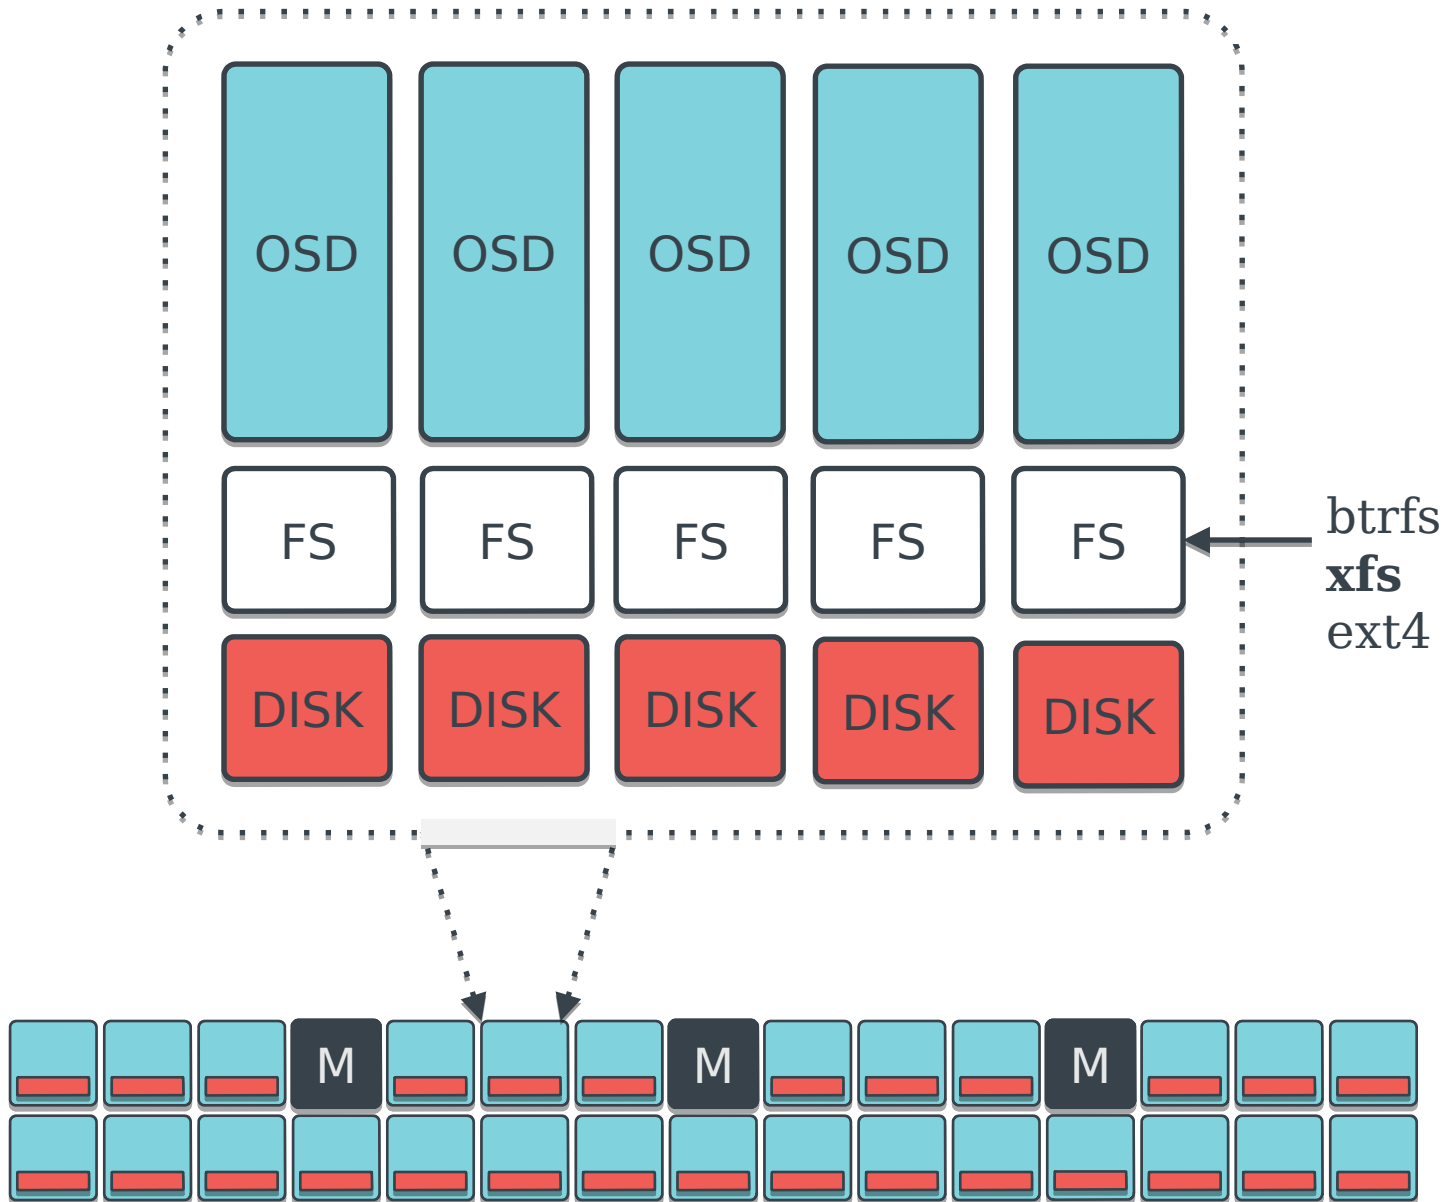


INKTANK CEPH ENTERPRISE (CEPH)

CEPH ARCHITECTURE




OSD & MONITOR NODES




CRUSH

10 10 01 01 10 10 01 11 01
10

$pg = \text{hash}(\text{object name}, \text{pool})$



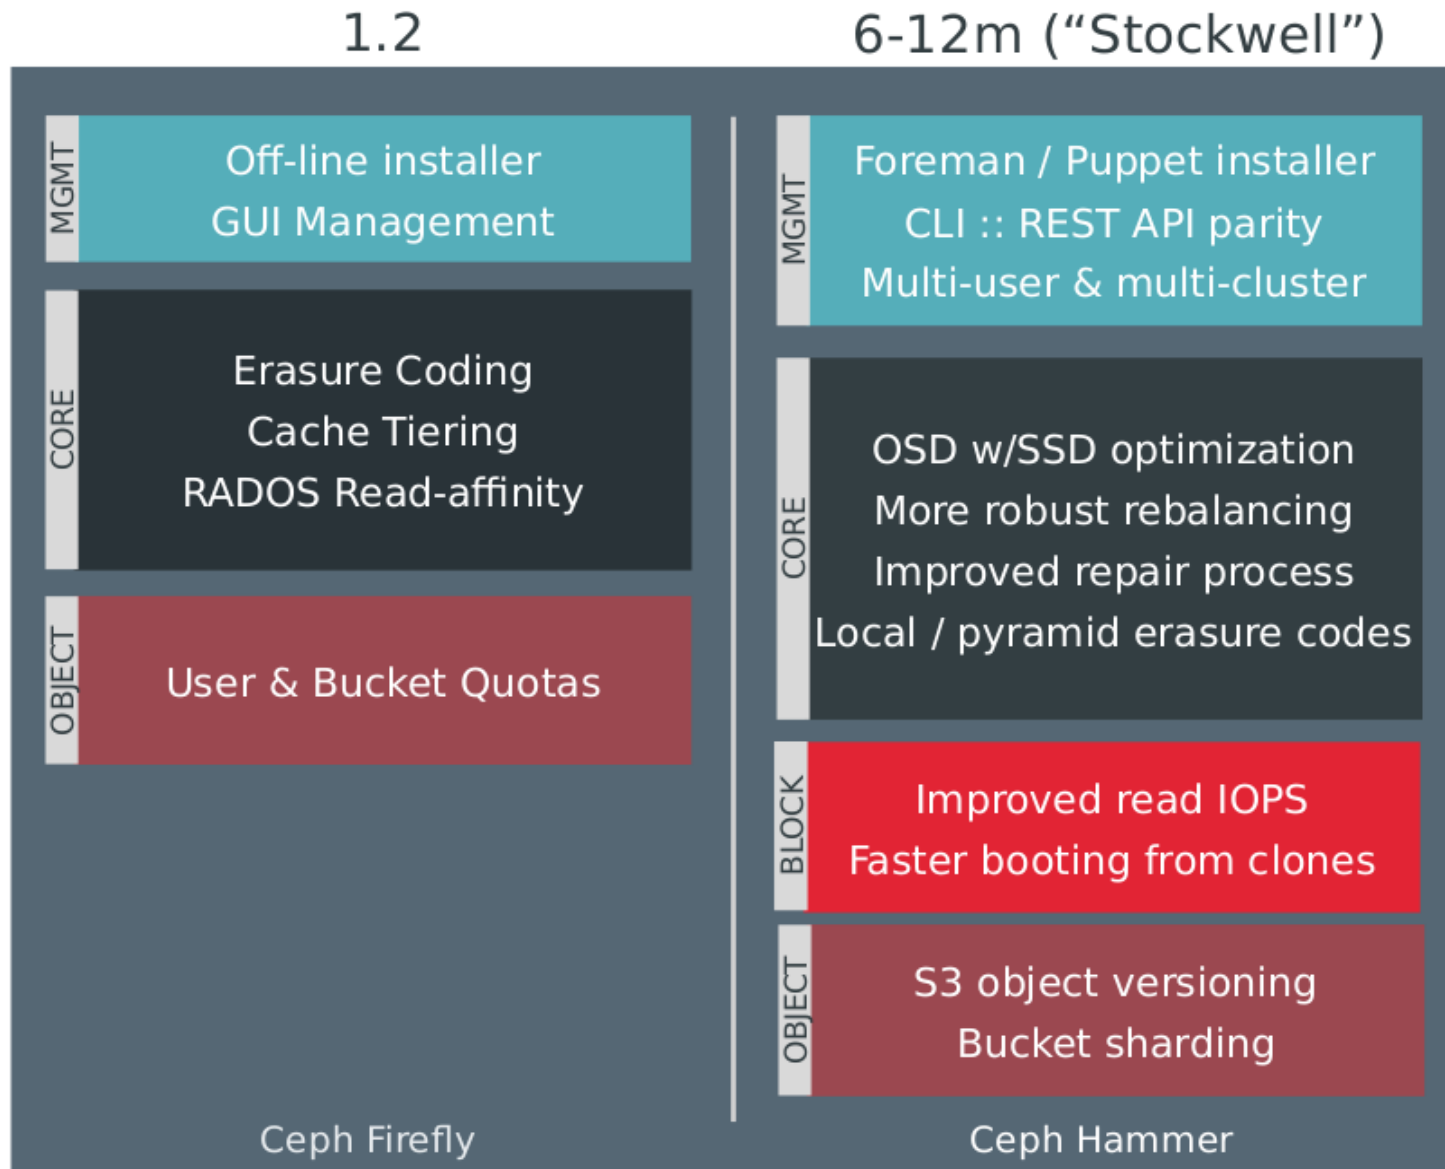
$\text{CRUSH}(pg, \text{cluster state}, \text{rule set})$



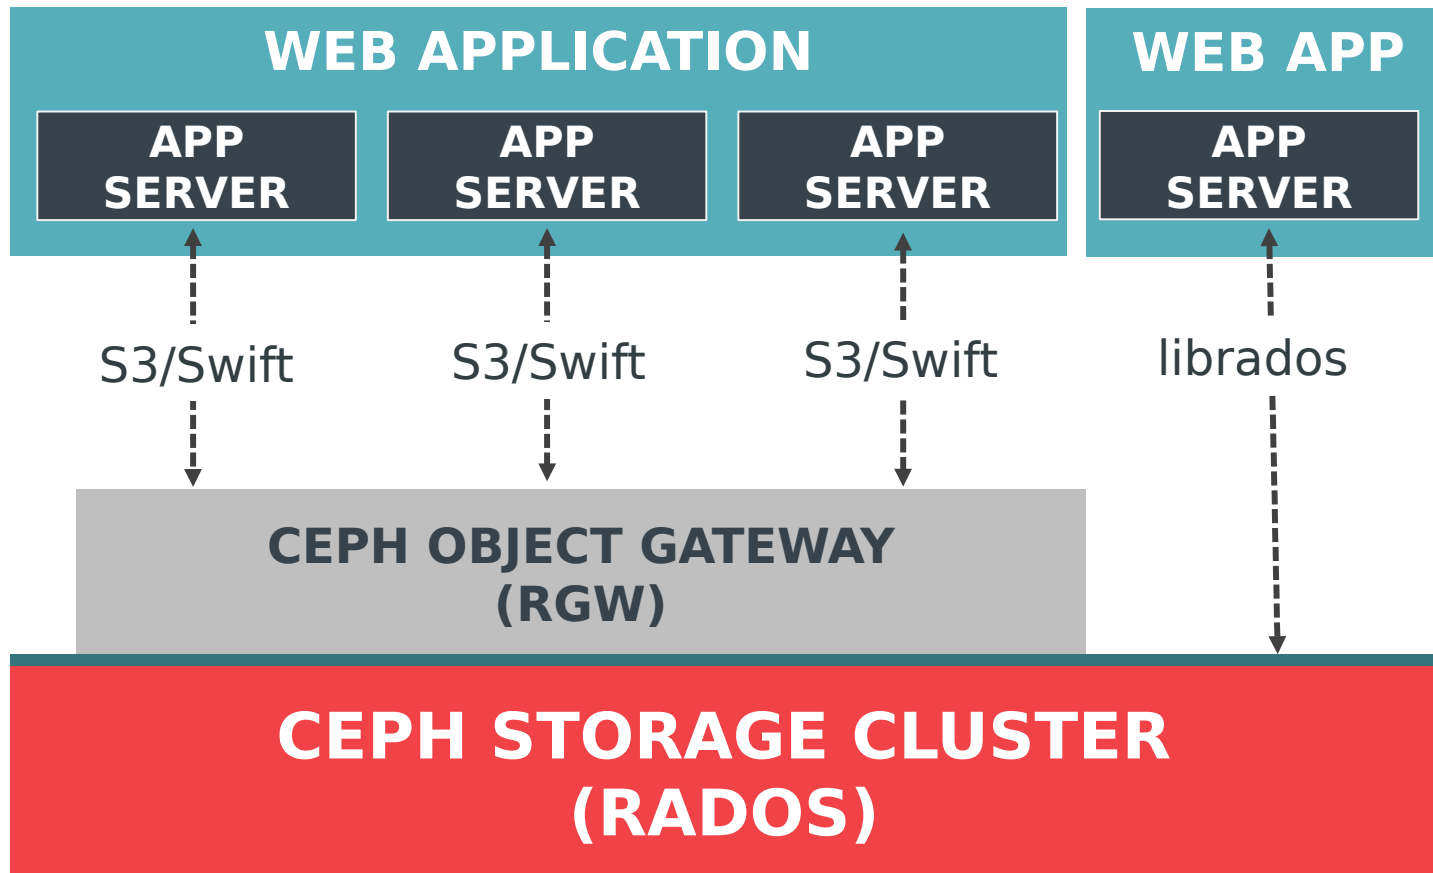
ICE 1.2.x

- Current ICE customers have access to an inktank.com FTP repo
- In the coming months, we'll transition to CDN
 - i.e. Satellite or Red Hat based distribution, in the same manner as everything else with Red Hat
- docs.redhat.com
- Red Hat QA, errata, build mechanism, etc

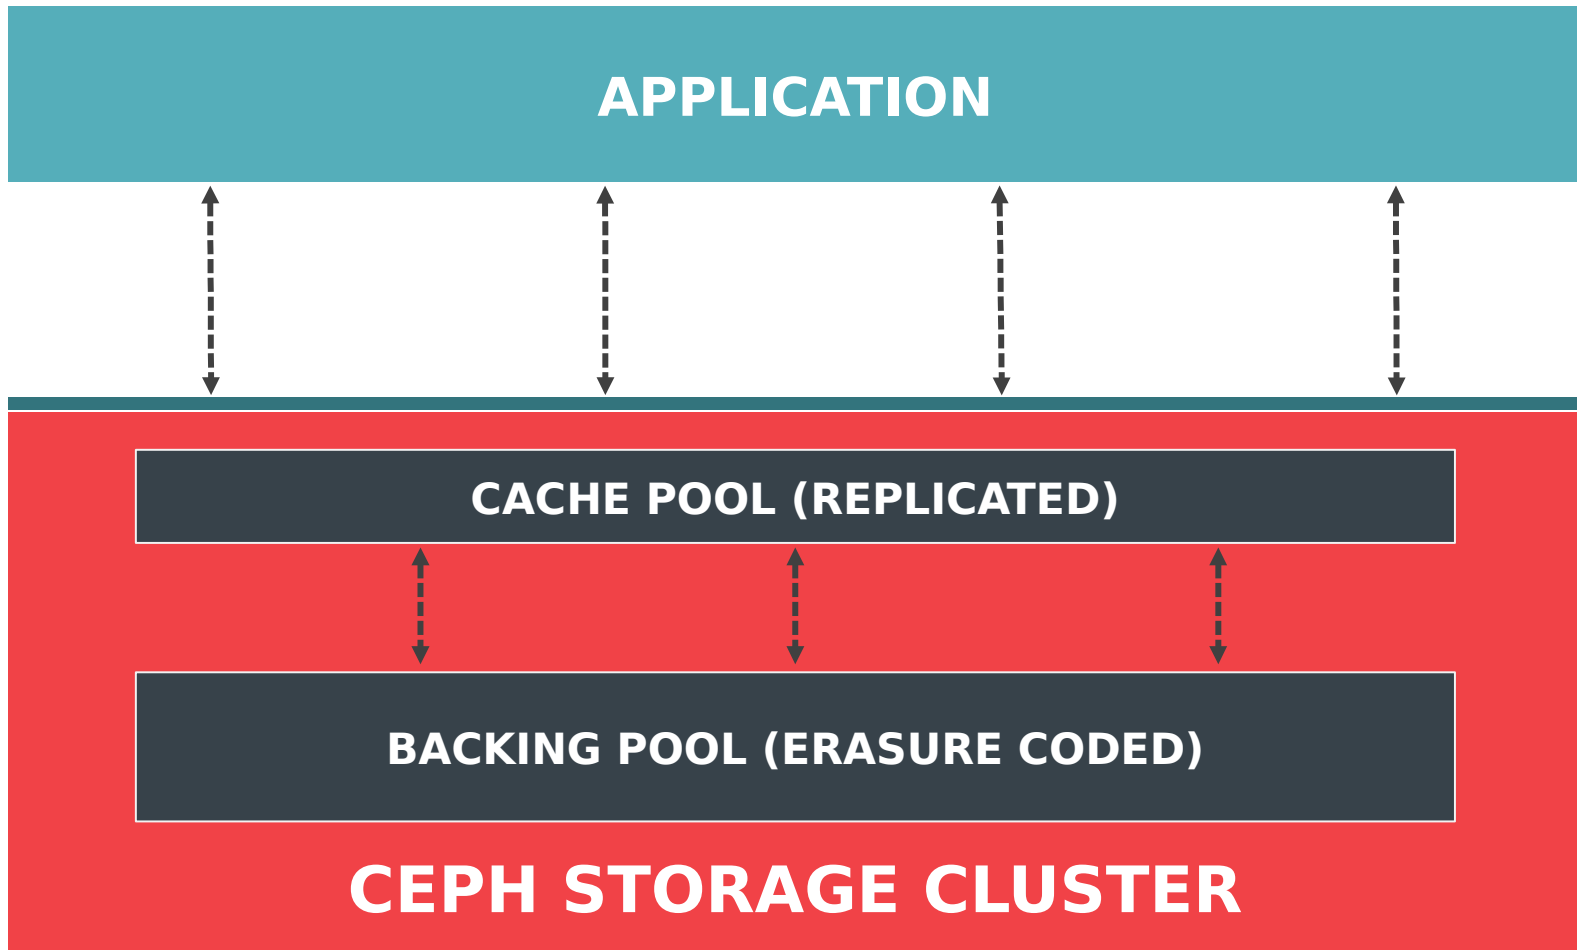
UPCOMING FEATURES - PREVIEW



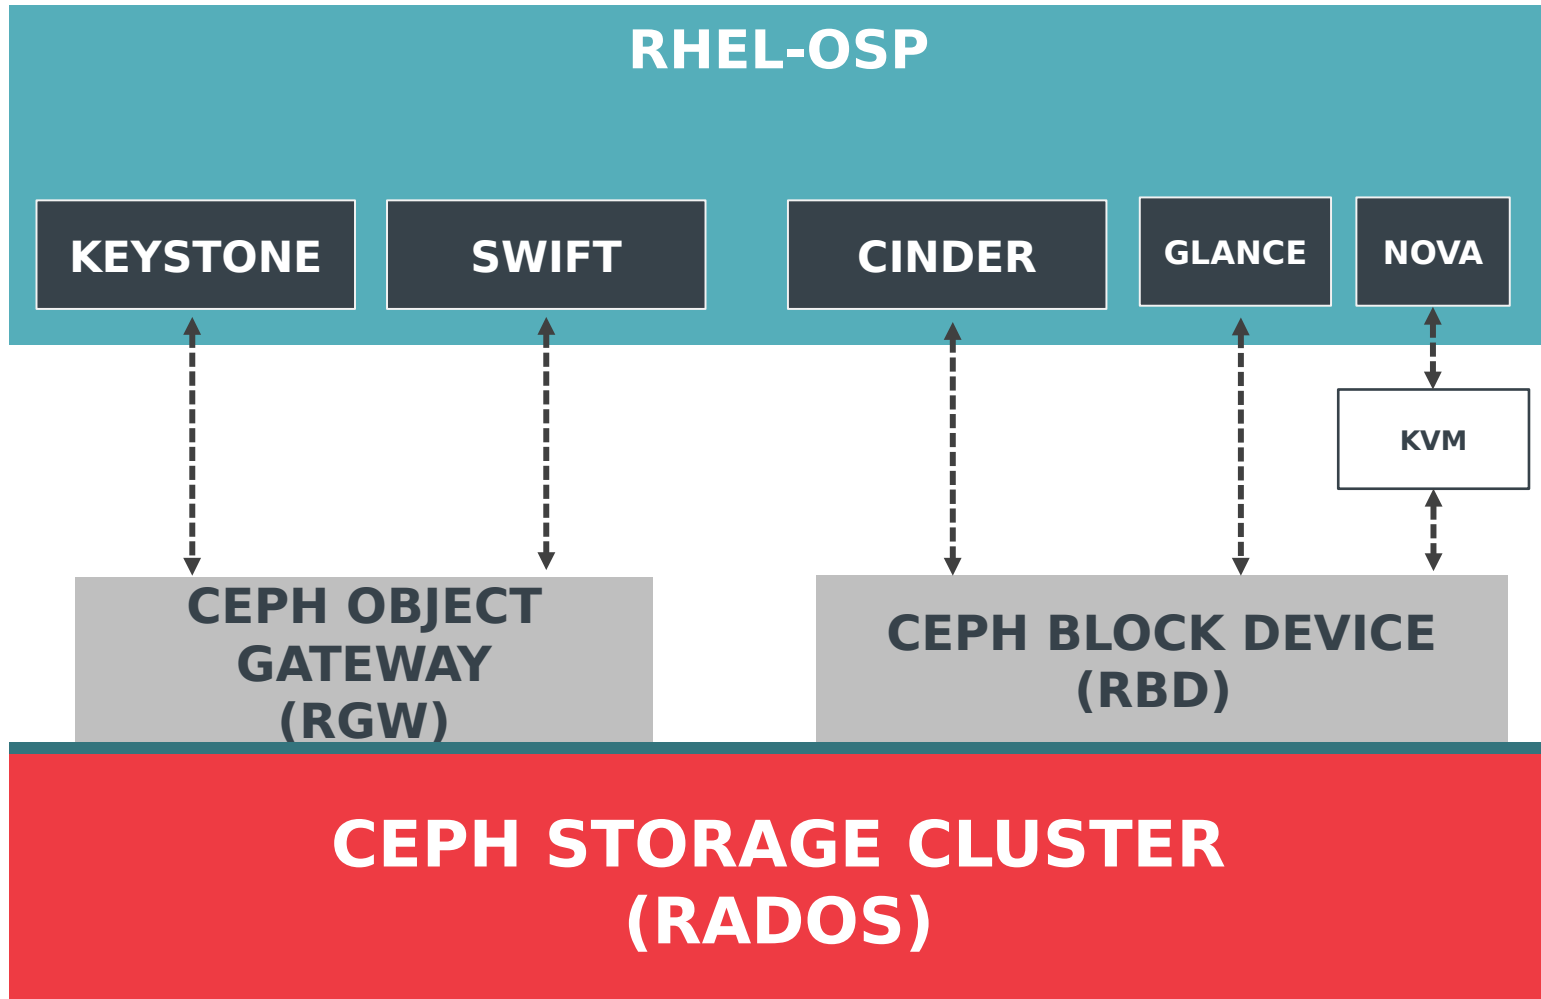
USE CASE: WEB APP STORAGE



USE CASE: COLD STORAGE



USE CASE: OPENSTACK



THANK YOU

