



# RED HAT CEPH

## TECHNICAL OVERVIEW

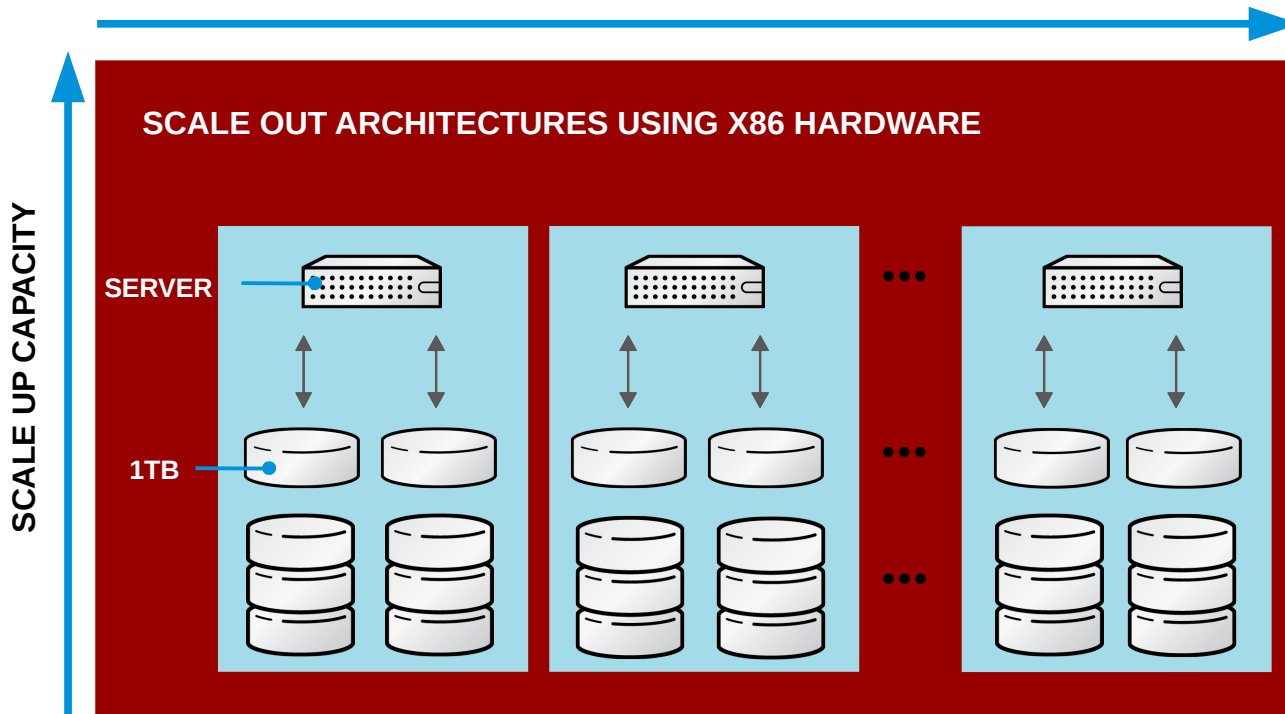
James Rankin

Senior Solutions Architect



# BEFORE WE BEGIN

SCALE OUT PERFORMANCE, CAPACITY, AND AVAILABILITY



**GAME-CHANGING TECHNOLOGY**

**DESIGNED FOR MODERN WORKLOADS**

**EXPAND ON DEMAND**

**WORKS WITH INNOVATIVE TECHNOLOGY LIKE CISCO UCS**

**LOWER \$/TB**

# THE OPEN SOFTWARE-DEFINED STORAGE PLATFORM FOR THE MODERN HYBRID DATACENTER

## TRADITIONAL STORAGE

**VENDOR CONTROLLED**

HARDWARE CENTRIC

LOCKED DOWN

EXPENSIVE

PROPRIETARY

## THE RED HAT WAY

**COMMUNITY DRIVEN**

SOFTWARE CENTRIC

DISRUPTIVE ARCHITECTURE

INEXPENSIVE

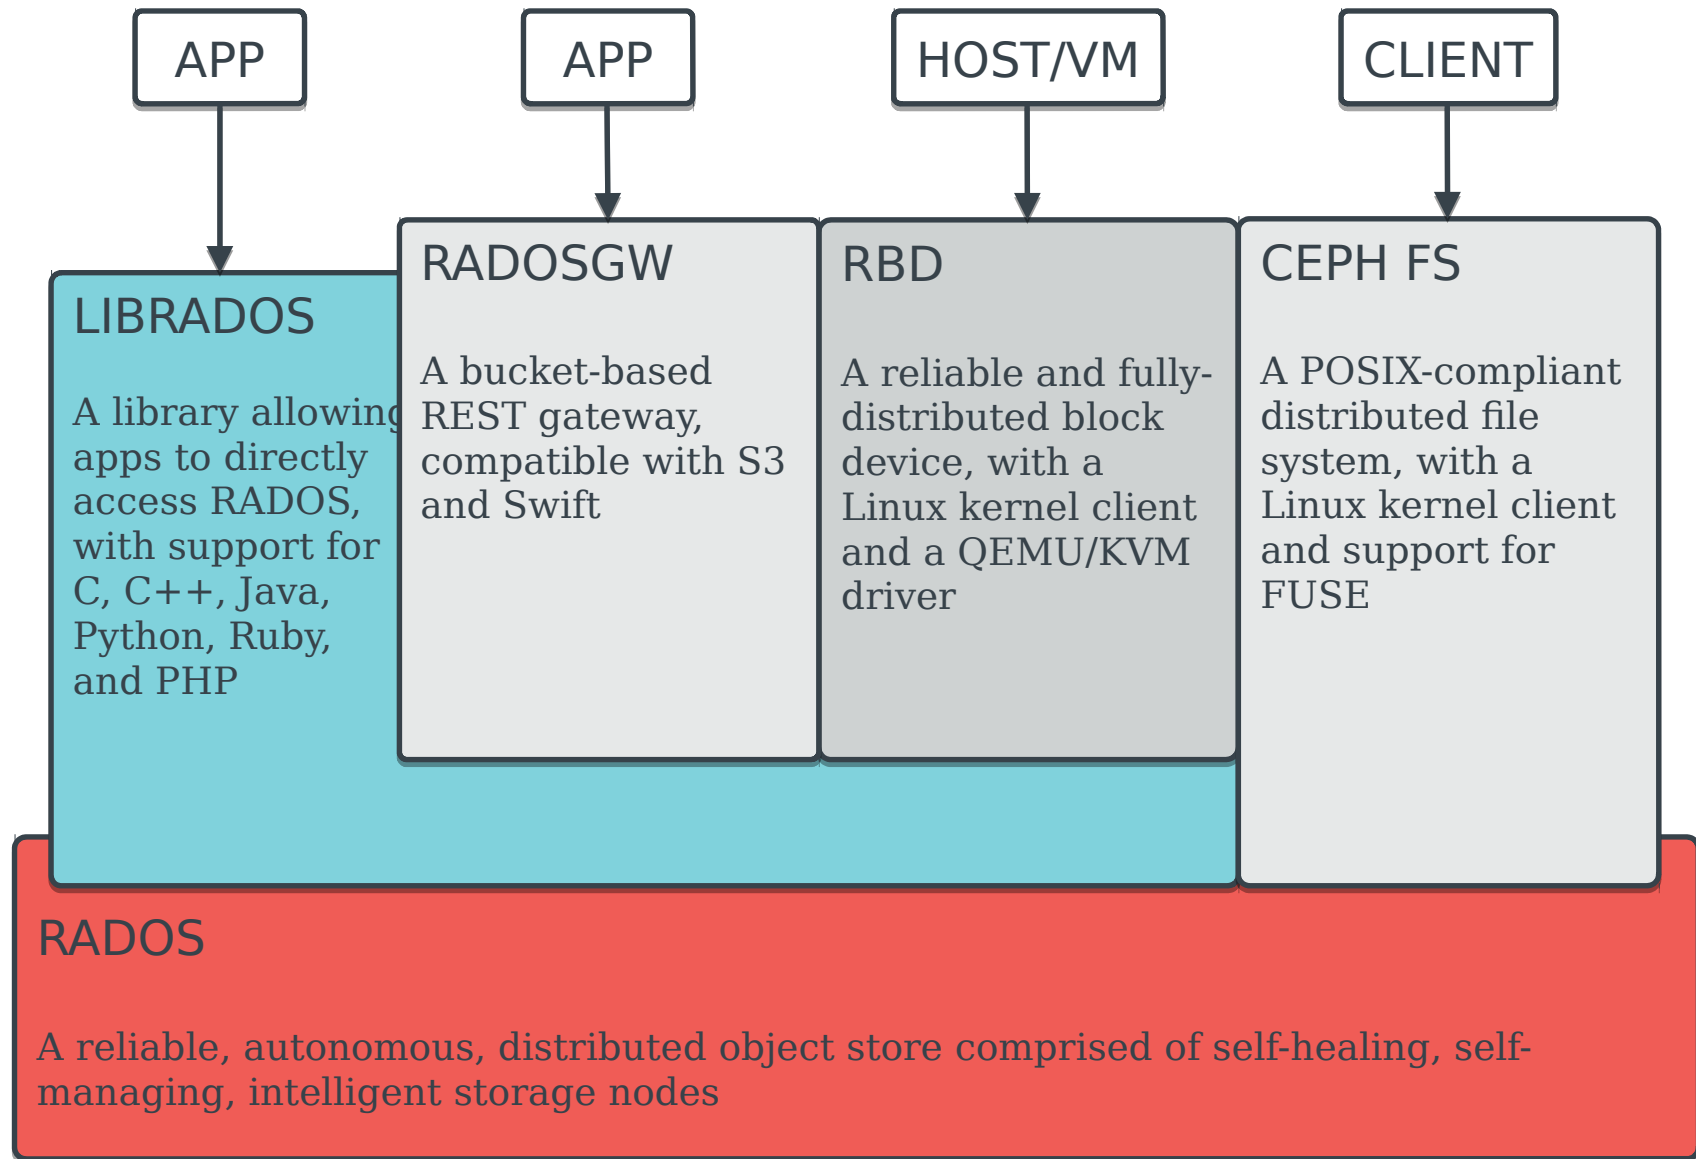
OPEN SOURCE

# CEPH & INKTANK

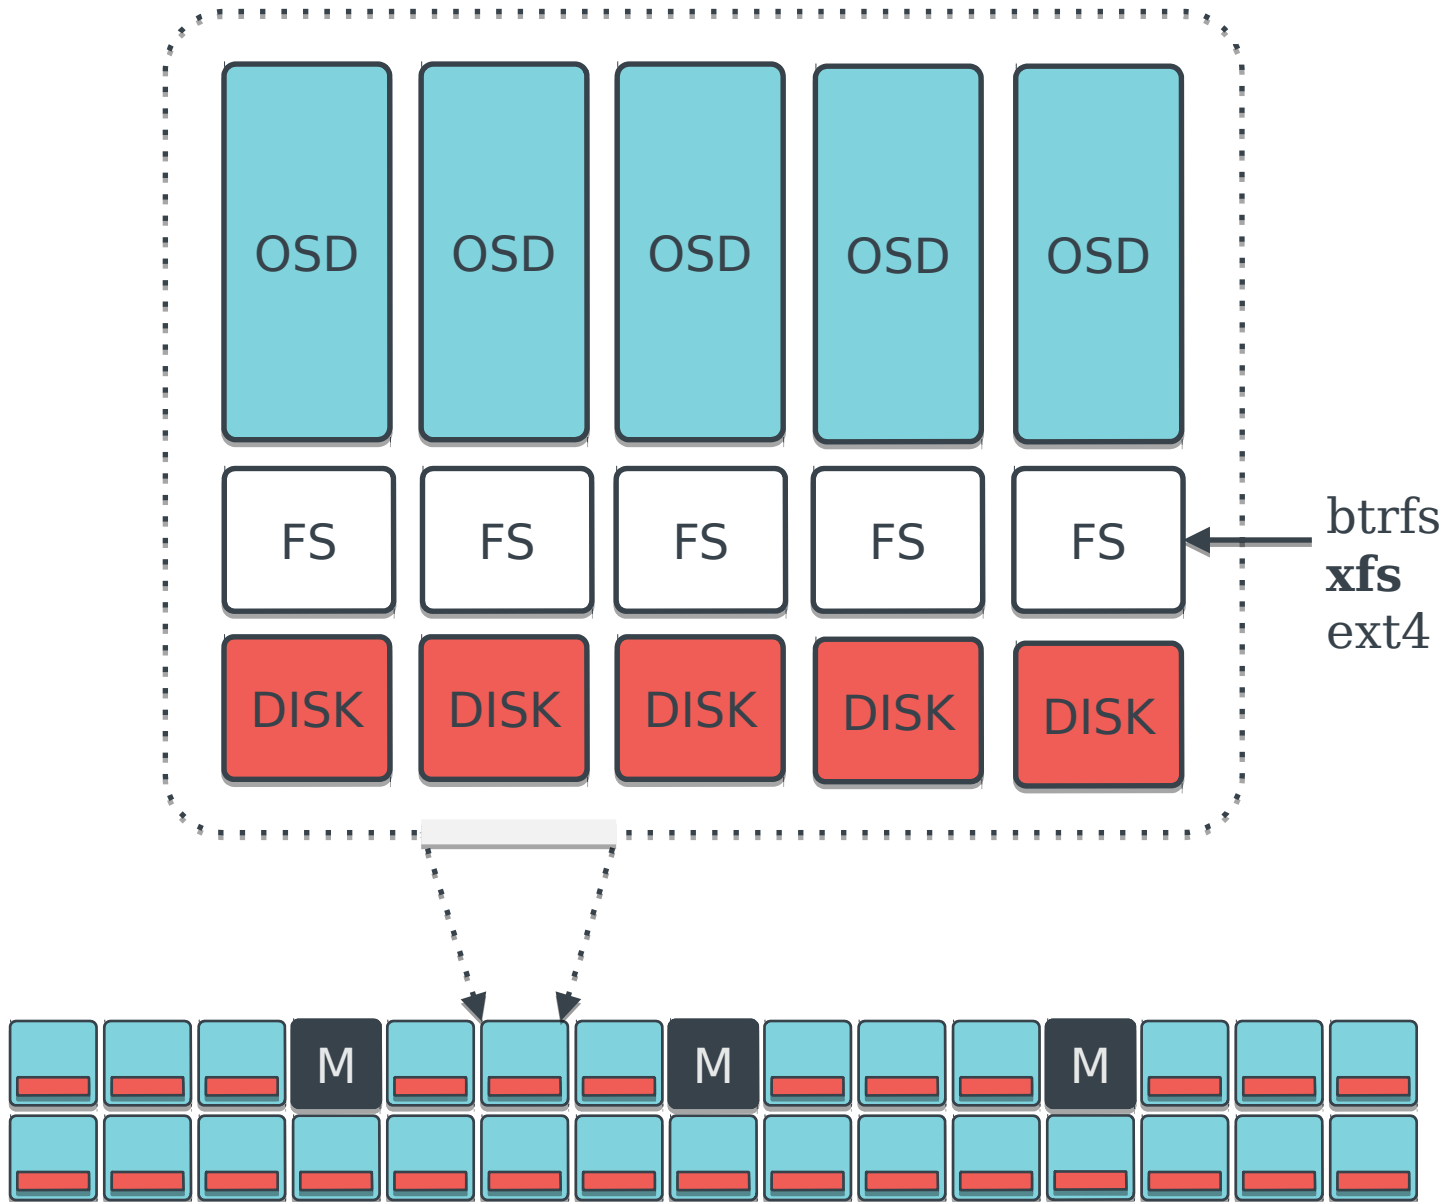


- Red Hat acquired Inktank, the company behind Ceph
- Ceph is scale out software defined storage
  - Unlike Glusterfs, Ceph's main focus is block
  - Ceph is architecturally an *object store (RADOS)*
  - Ceph can assemble block images (similar to LUNs) through chunks of objects
- **Glusterfs & Ceph are complementary technologies that address different use cases**


# CEPH ARCHITECTURE



# OSD & MONITOR NODES



# OSD & MONITOR NODES



## OSDs:

- 10s to 10000s in a cluster
- One per disk
  - (or one per SSD, RAID group...)
- Serve stored objects to clients
- Intelligently peer to perform replication and recovery tasks




## Monitors:

- Maintain cluster membership and state
- Provide consensus for distributed decision-making
- Small, odd number
- These do **not** serve stored objects to clients

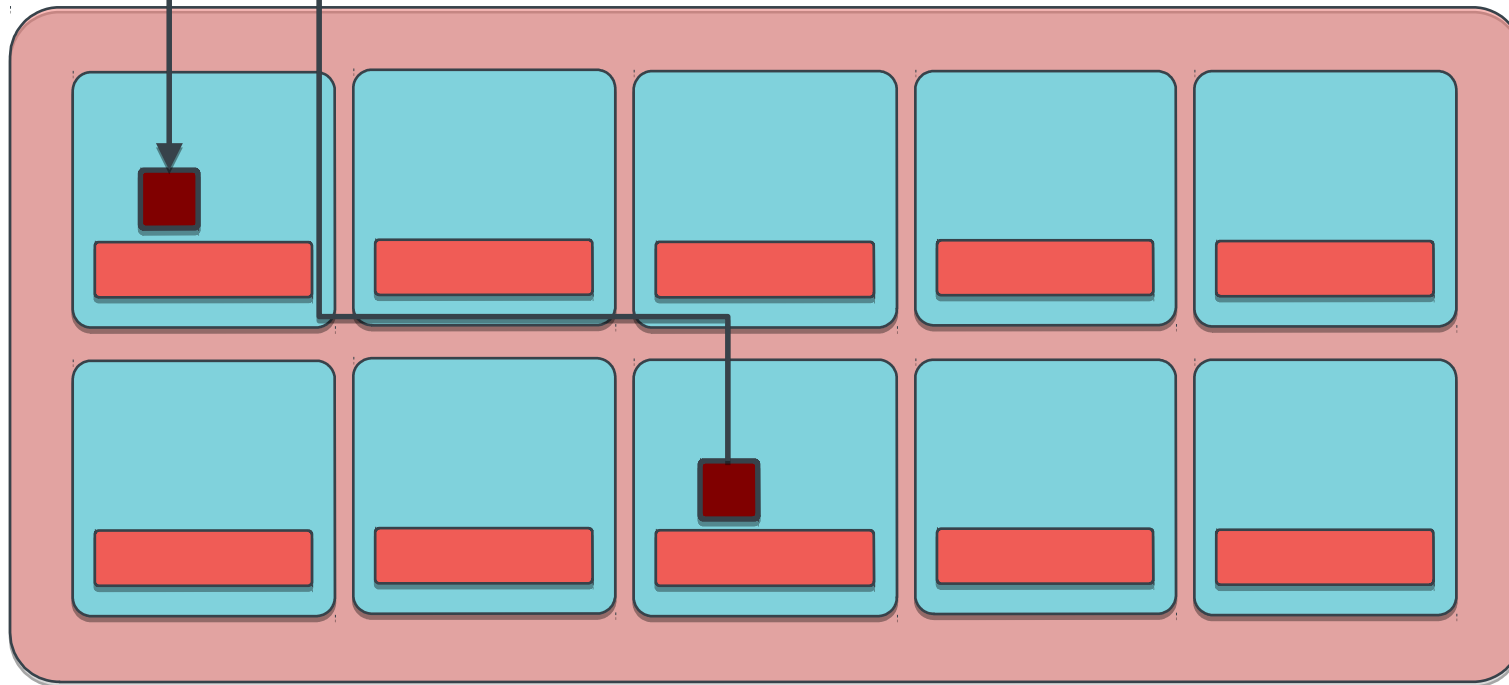
# CRUSH

10 10 01 01 10 10 01 11 01  
10

$\text{hash}(\text{object name}) \% \text{num pg}$




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





# CRUSH

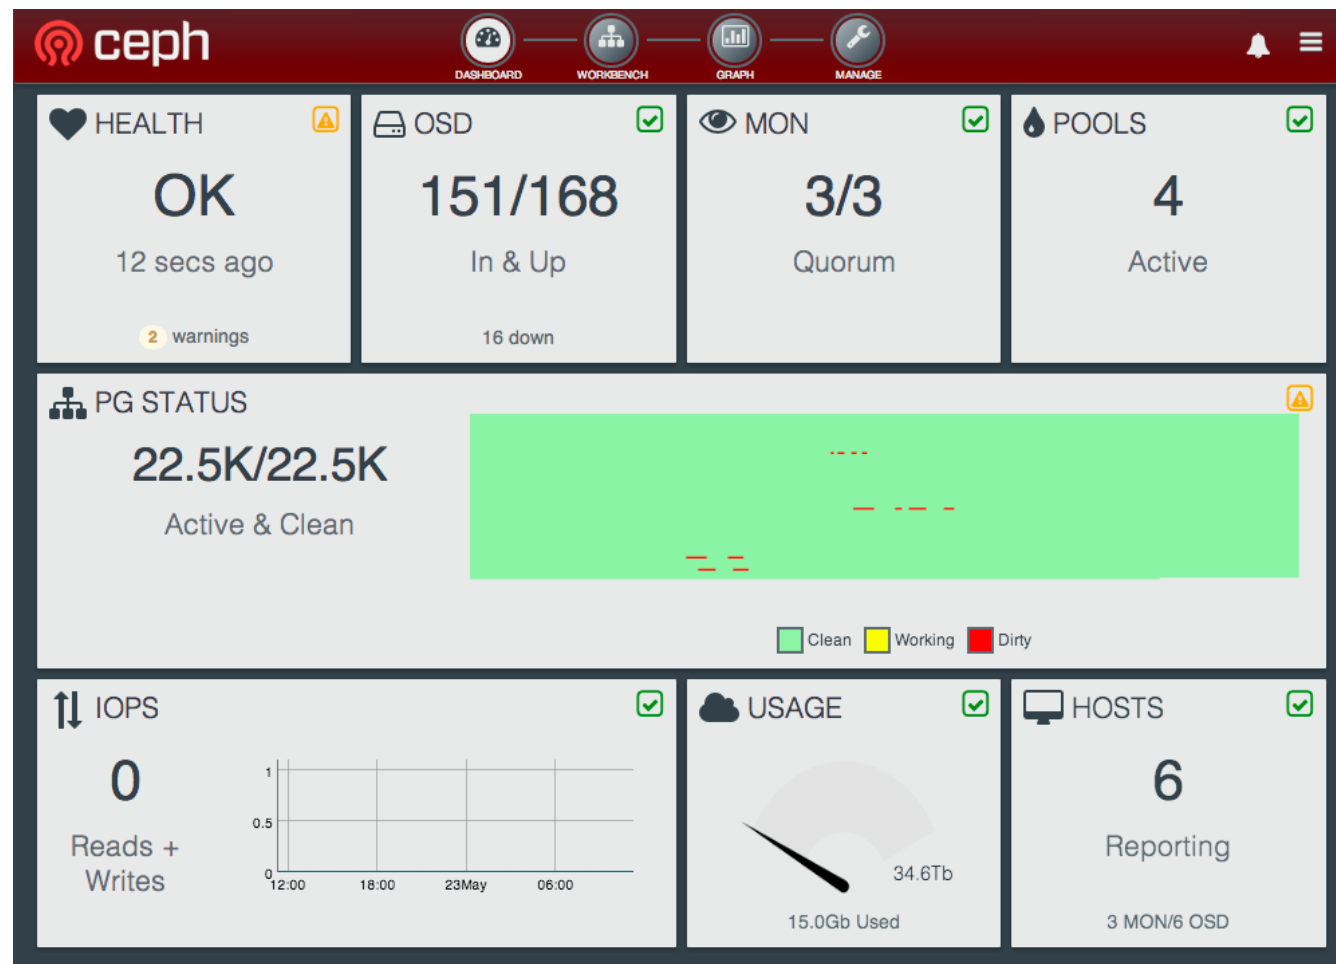


## CRUSH

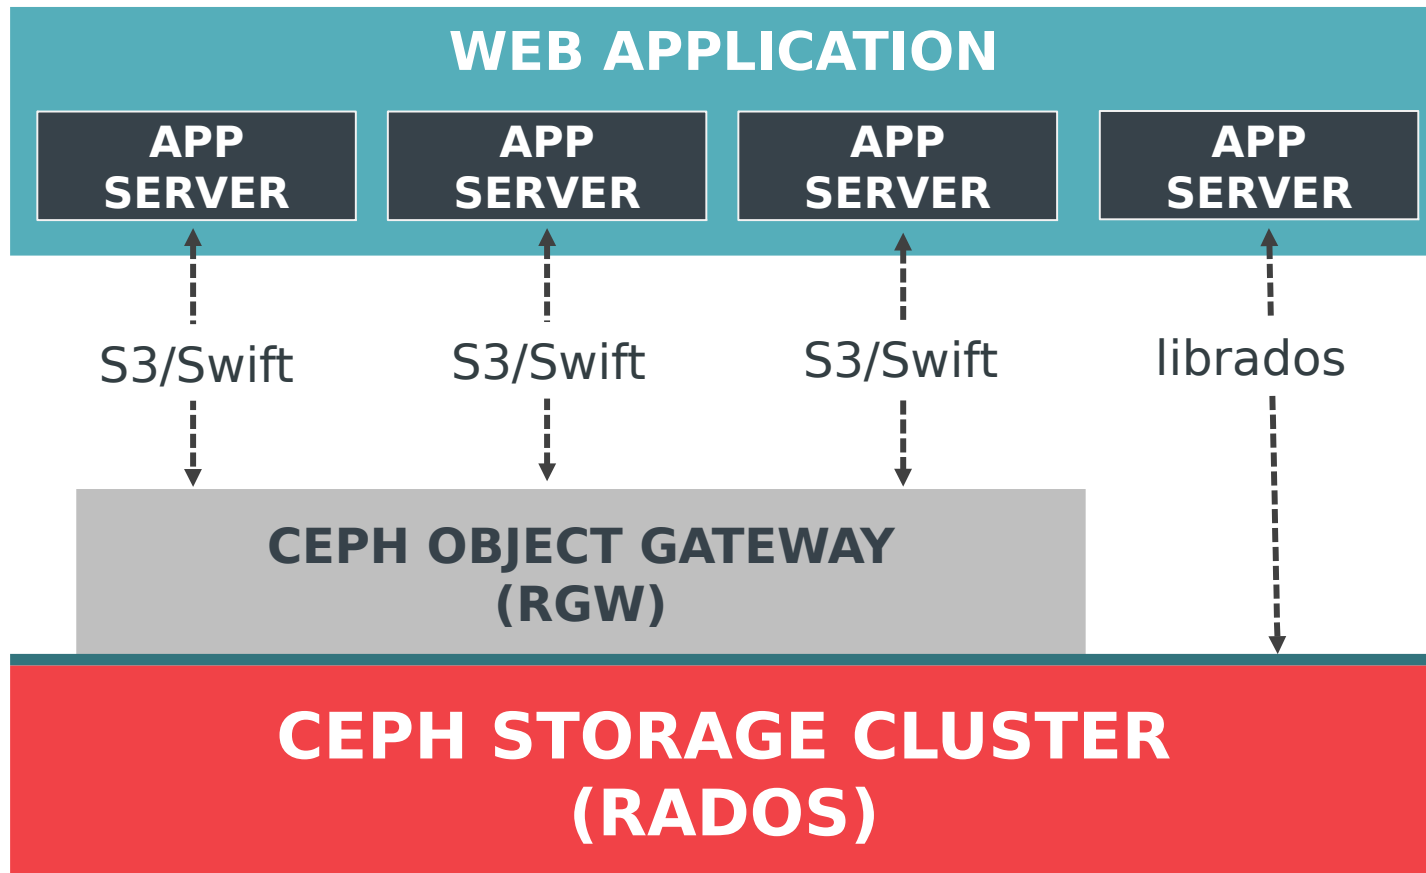
- Pseudo-random placement algorithm
  - Fast calculation, **no lookup**
  - Repeatable, deterministic
- Statistically uniform distribution
- Stable mapping
  - Limited data migration on change
- Rule-based configuration
  - Infrastructure topology aware
  - Adjustable replication
  - Weighting

# CALAMARI

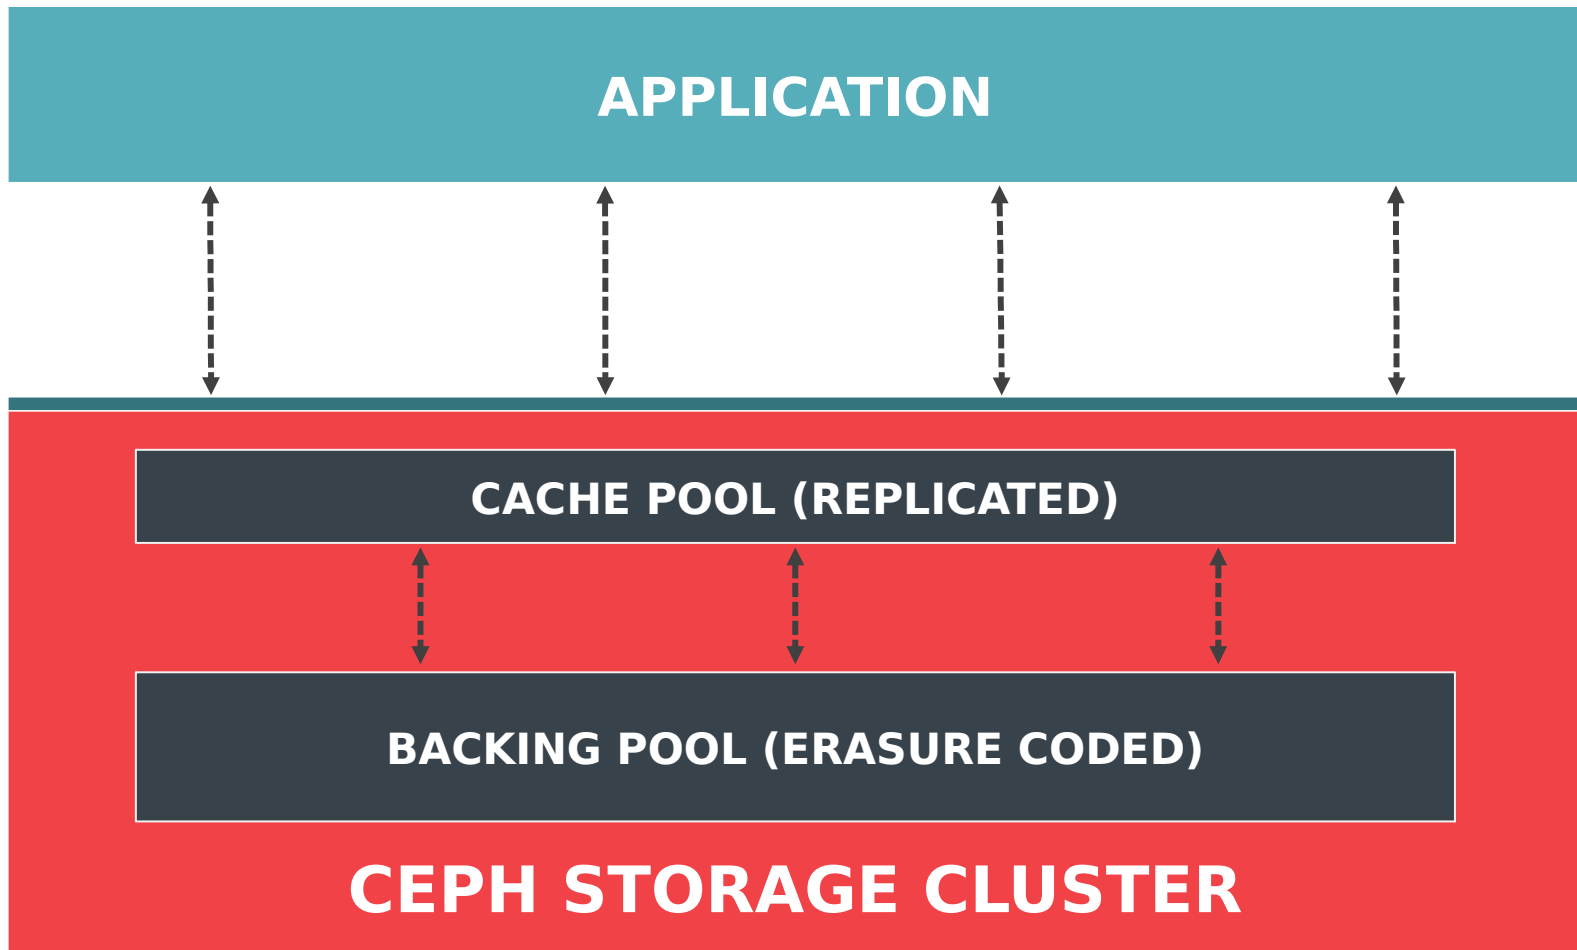
- Graphical management tool for Ceph Enterprise
- Quickly determine cluster health, utilization, etc
- Red Hat open sourced Calamari (only ~30 days after acquisition)



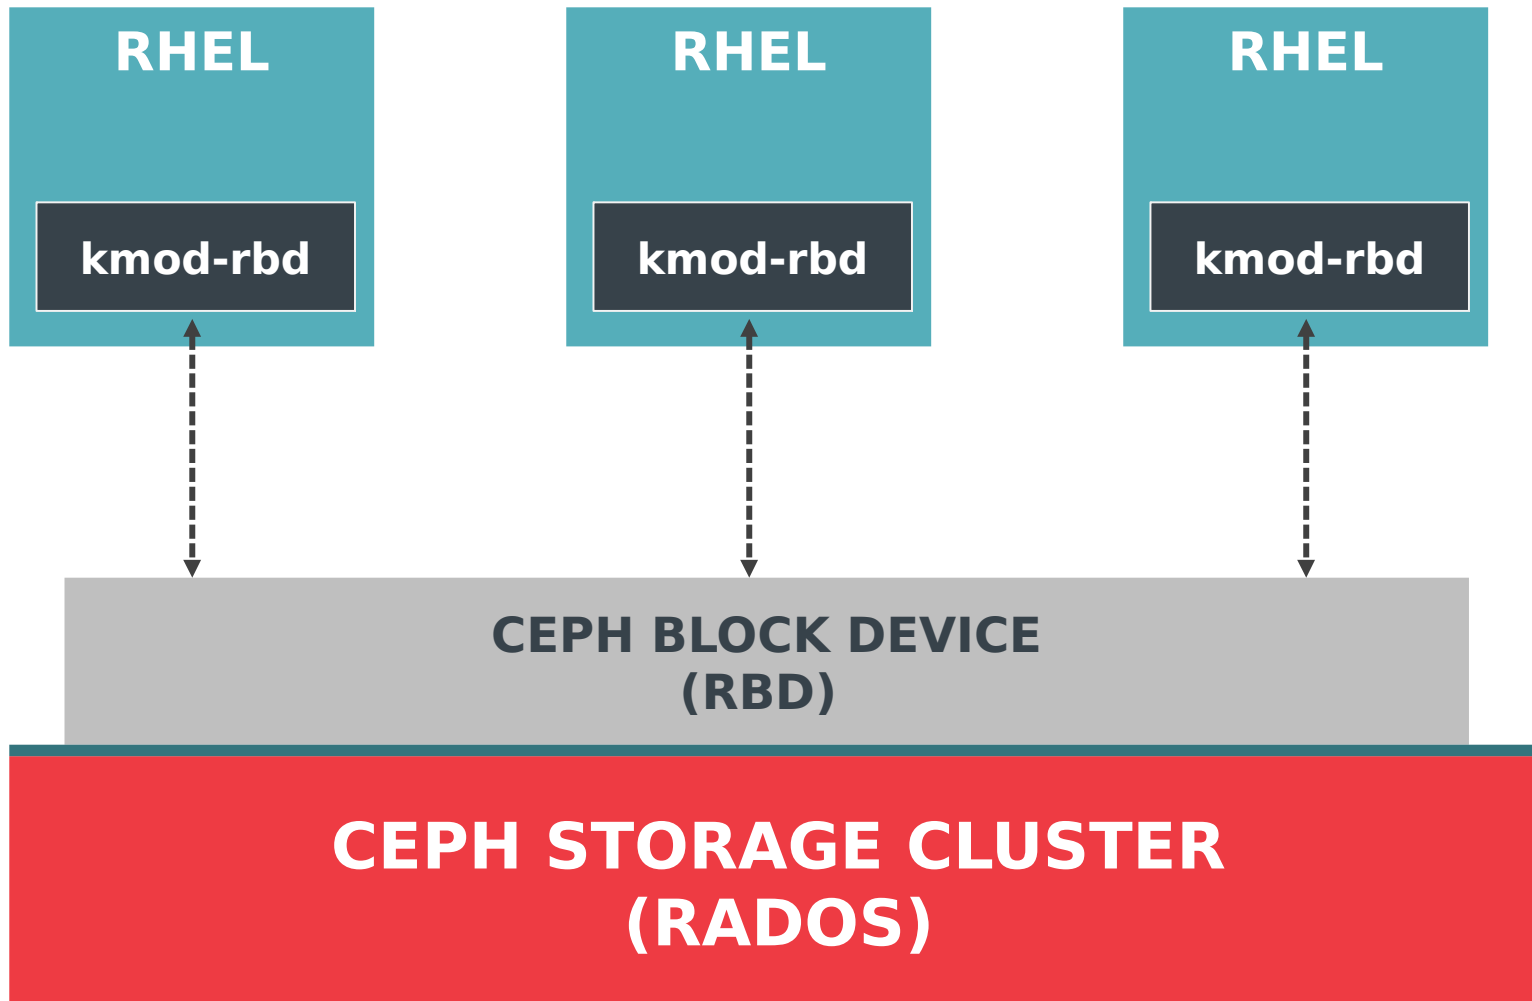
# USE CASE: WEB APP STORAGE



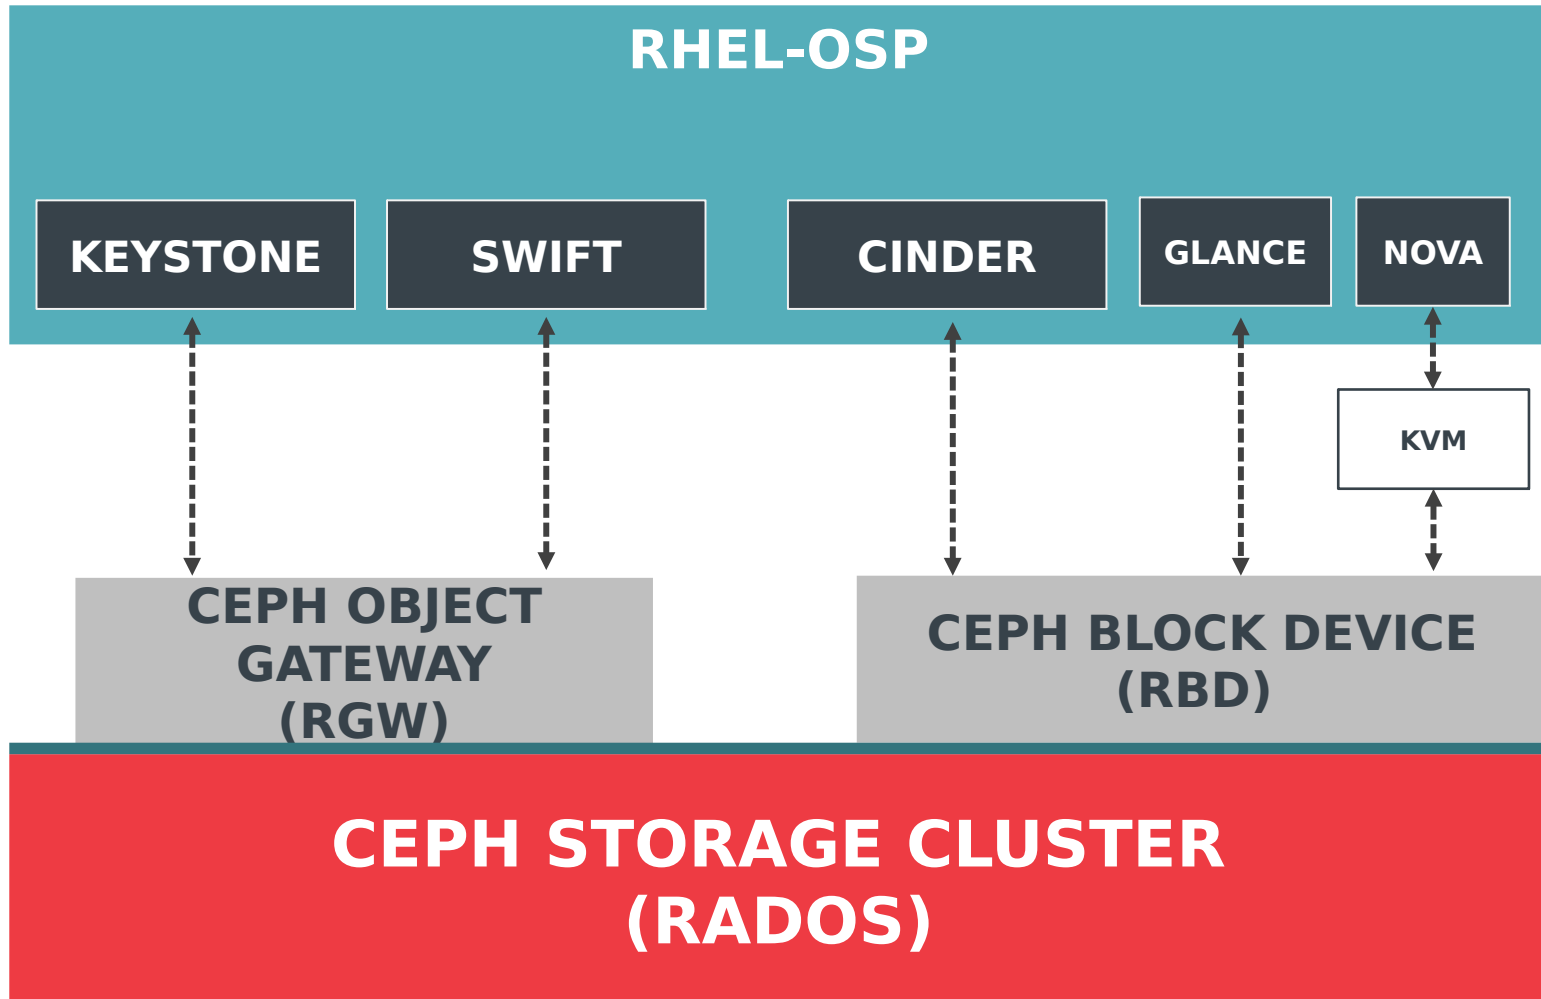
# USE CASE: COLD STORAGE



# USE CASE: RHEL BLOCK STORAGE



# USE CASE: OPENSTACK



**“...A DISRUPTIVE AND  
UNSTOPPABLE FORCE.”**

–IDC REPORT

**THANK YOU**

