

# IBM Storage for Data & AI: Ceph


Lars Lauber

Subject Matter Expert- Storage for Data & AI

[lauberla@de.ibm.com](mailto:lauberla@de.ibm.com)



# Disclaimer




IBM's statements regarding its plans, directions, and intent are subject to change or withdrawal without notice at IBM's sole discretion. Information regarding potential future products is intended to outline our general product direction and it should not be relied on in making a purchasing decision. The information mentioned regarding potential future products is not a commitment, promise, or legal obligation to deliver any material, code, or functionality. The development, release, and timing of any future features or functionality described for our products remains at our sole discretion.


IBM reserves the right to change product specifications and offerings at any time without notice. This publication could include technical inaccuracies or typographical errors. References herein to IBM products and services do not imply that IBM intends to make them available in all countries.

# IBM Storage Portfolio




**Data and AI** 


IBM Storage Scale  
IBM Storage Ceph




IBM Ceph Ready Nodes




IBM Storage Scale System

**Hybrid Cloud** 


IBM Storage Fusion




IBM Storage Fusion HCI System




Tape Drives - Libraries




IBM Storage TS7770

**Data Resiliency** 


IBM Storage Defender




FlashSystem 9500




IBM Defender Ready Nodes




FlashSystem 7300



FlashSystem 5000




IBM Storage SVC



DS8900F

Tape Drives - Libraries IBM Storage TS7770

# IBM Storage Ceph



## Storage Ceph

### IBM Storage Ceph

An enterprise-grade distributed, universal, software defined storage solution, proven at scale

### Implementation

Flexible, scale-out architecture running on clustered x86 industry-standard hardware

### IBM Storage Ceph provides

A single, efficient, unified storage platform for object, block, and file storage

### IBM Storage Ceph Easy

An IBM Storage Ceph cluster can literally be installed by running one single command

IBM Storage Ceph also features a dashboard UI and APIs for lights out data center operations.

### IBM Storage Ready Nodes

Validated hardware for IBM Storage Ceph, ready for immediate use, with published performance data.

Scale from 4 to 400 nodes  
384TB to 96 PB capacity  
N+1 scalability option

### Enterprise level solution

IBM Storage Ceph is an enterprise-level storage solution, with an open source development model

IBM holds a premier membership seat inside the Ceph Foundation

# IBM Storage Ceph offering options



## IBM Storage Ceph Software only

Install and run IBM Storage Ceph on industry-standard x86 server hardware of a company preferred hardware vendor

IBM Storage Ceph can be bought within a Premium- or Pro licensing option, which defines delivery with- or without inclusion of Red Hat Enterprise Linux

## IBM Storage Ceph on Ready Nodes

IBM Storage Ready Nodes offer a pre-configured, validated, and scalable infrastructure option


A flexible way for organizations to quickly build an on-prem cloud-scale software-defined storage solution with software and hardware support provided by IBM

# IBM Storage for Data and AI

## Our Mission

Unlock unstructured data's potential and propel forward business growth and innovation by accelerating discovery, eliminating data silos, support green and lower cost initiatives and protecting from the unknown

## Two products and one mission



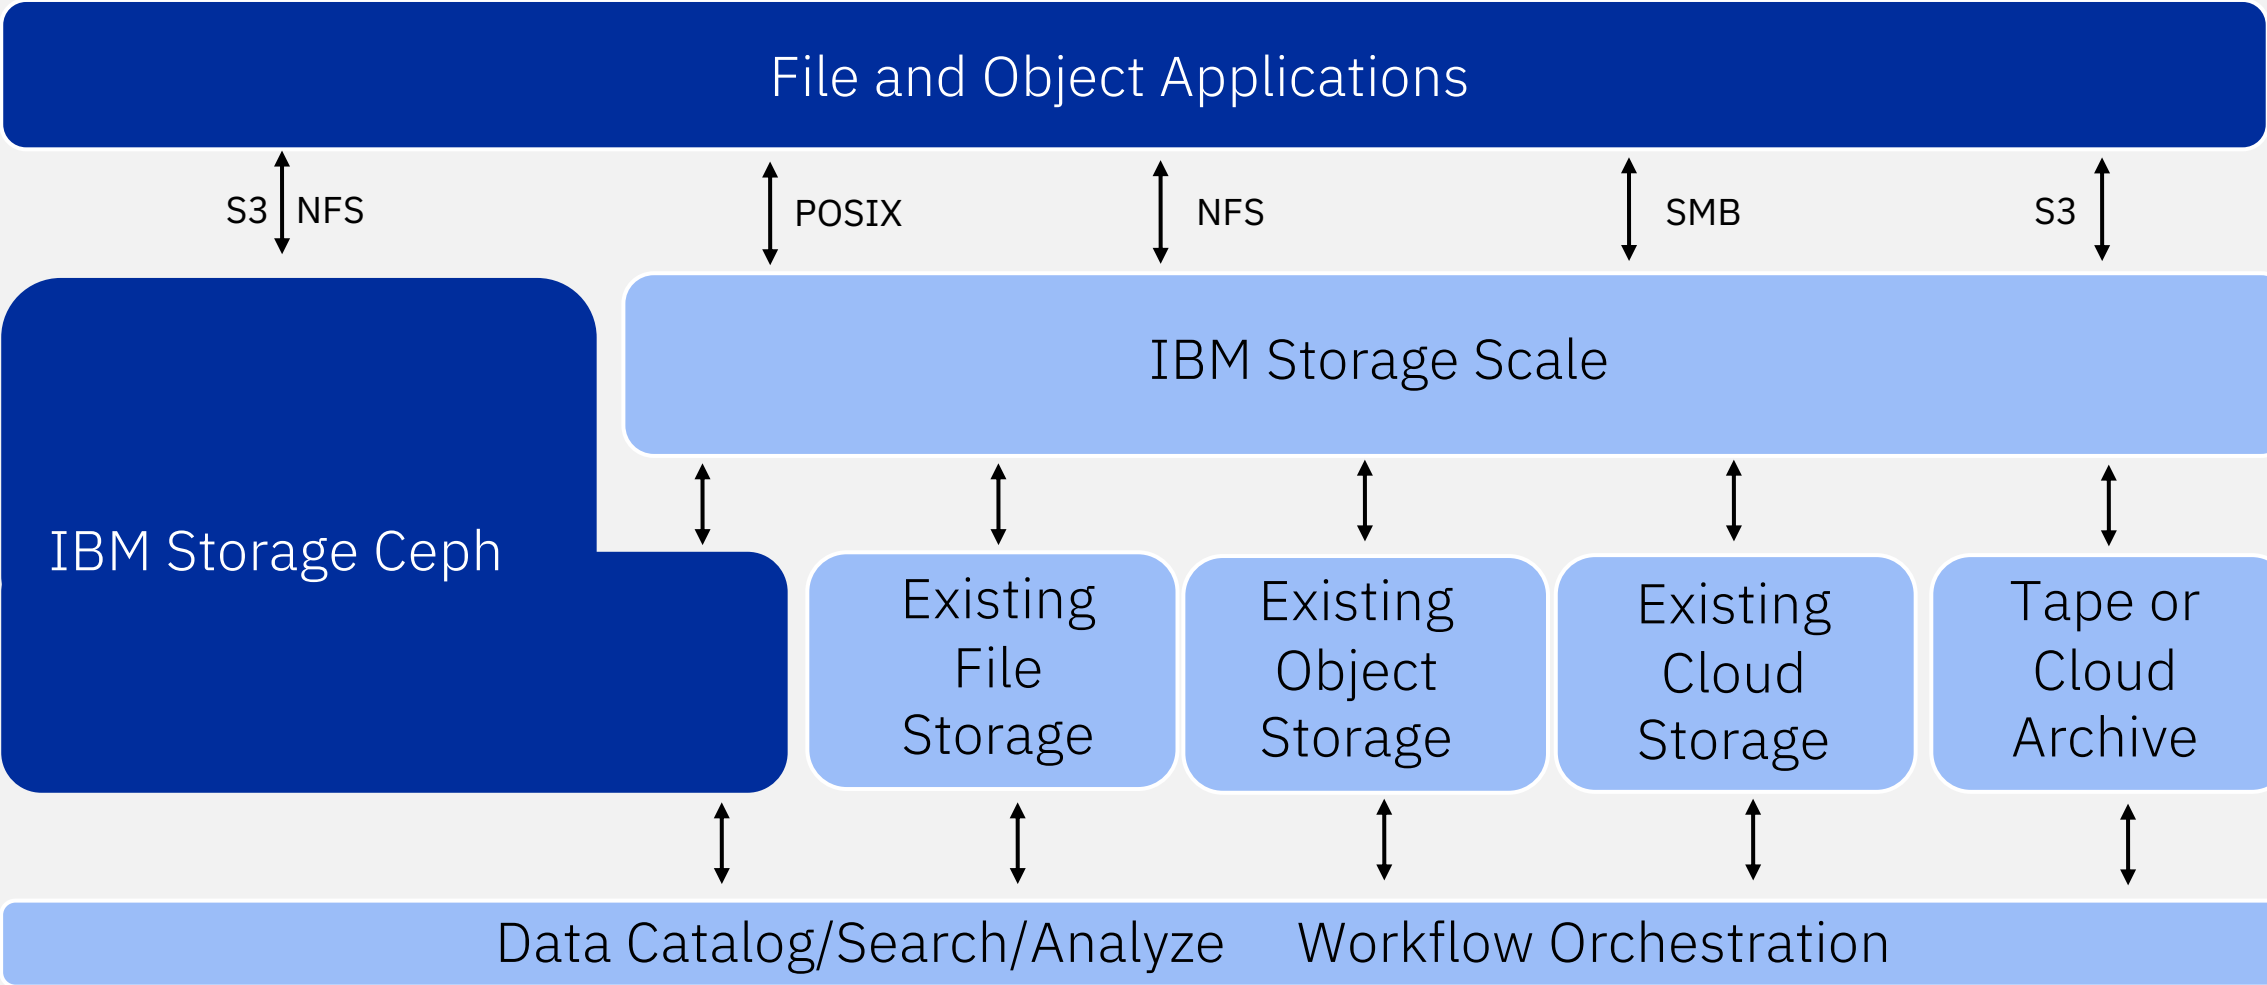
Storage Ceph

Operationalize AI with a data lake designed to improve as it scales




Storage Scale

The foundation for AI innovation and accelerated discovery




# IBM Storage Ceph data services

## Data efficiency




- Multi-protocol
- Hybrid/ Multi-cloud
- Replication
- Erasure coding
- Compression
- Performance-at-scale

## Data resilience



- Snapshots
- Clones
- Backup and recovery
- Migration options
- Business continuity
- Disaster recovery

## Data security



- Authentication and authorization
- Data policies
- At rest encryption
- In flight encryption
- Certified WORM
- Key management service integration



# IBM Storage Ceph

## RBAC support

Role based governance

## Security

Store tables with S3 data encryption

## D3N caching

D3N query acceleration

Datacenter Data Delivery Network

Natively implemented in IBM Storage Ceph RGW

## Parquet support

S3 Select pushdown

## Bucket notification

Notification of bucket events or changes

Can be used in context of serverless runtimes, feeds into messaging systems or similar

## Open-source storage software

For interaction with open-source data analytics applications

Also suitable for any other workload that requires file-, block or object storage

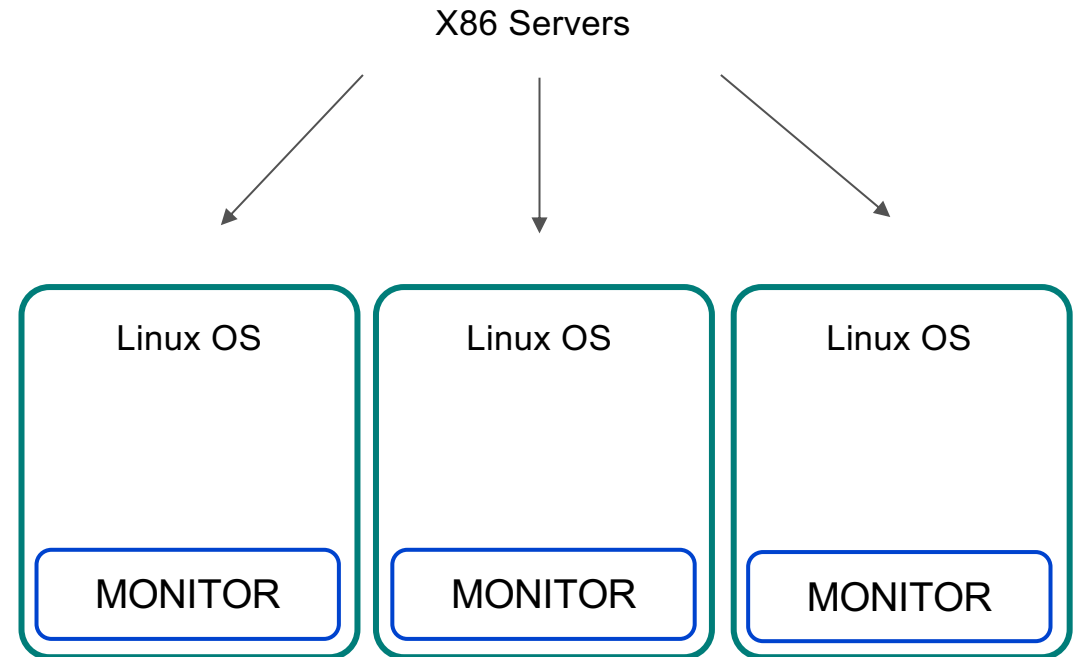
# IBM Storage Ceph basics

# IBM Storage Ceph architecture baseline

MONITOR PROCESS

Ceph uses monitors and object storage daemons

- Monitors maintain the Ceph cluster map
- Decisions are based on consensus: Paxos
- Monitors operate in a small and odd number
- Run as Containerized processes

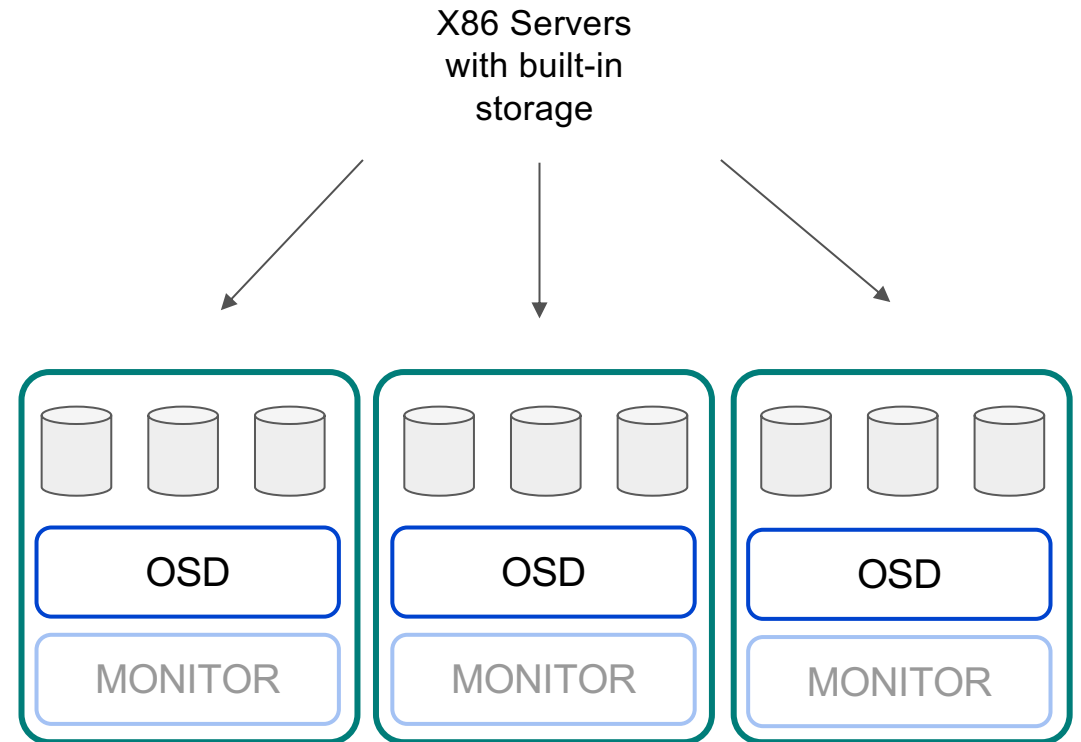


# IBM Storage Ceph architecture baseline

OSD PROCESS

## Ceph object storage daemons

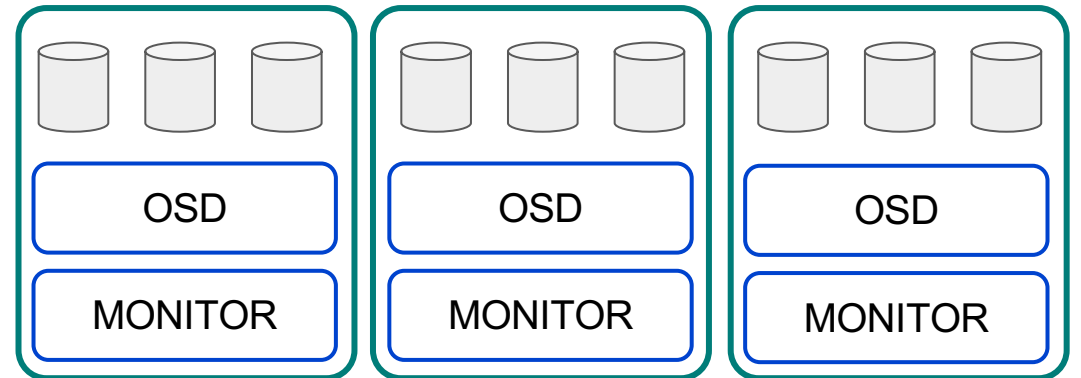
- OSD: Object storage daemon
- Provide direct data access
- Manage layout of data on media
- Peer and coordinate data distribution, integrity checking and recovery



# IBM Storage Ceph architecture baseline

A basic Ceph cluster setup is composed of monitors and object storage daemons (OSDs)

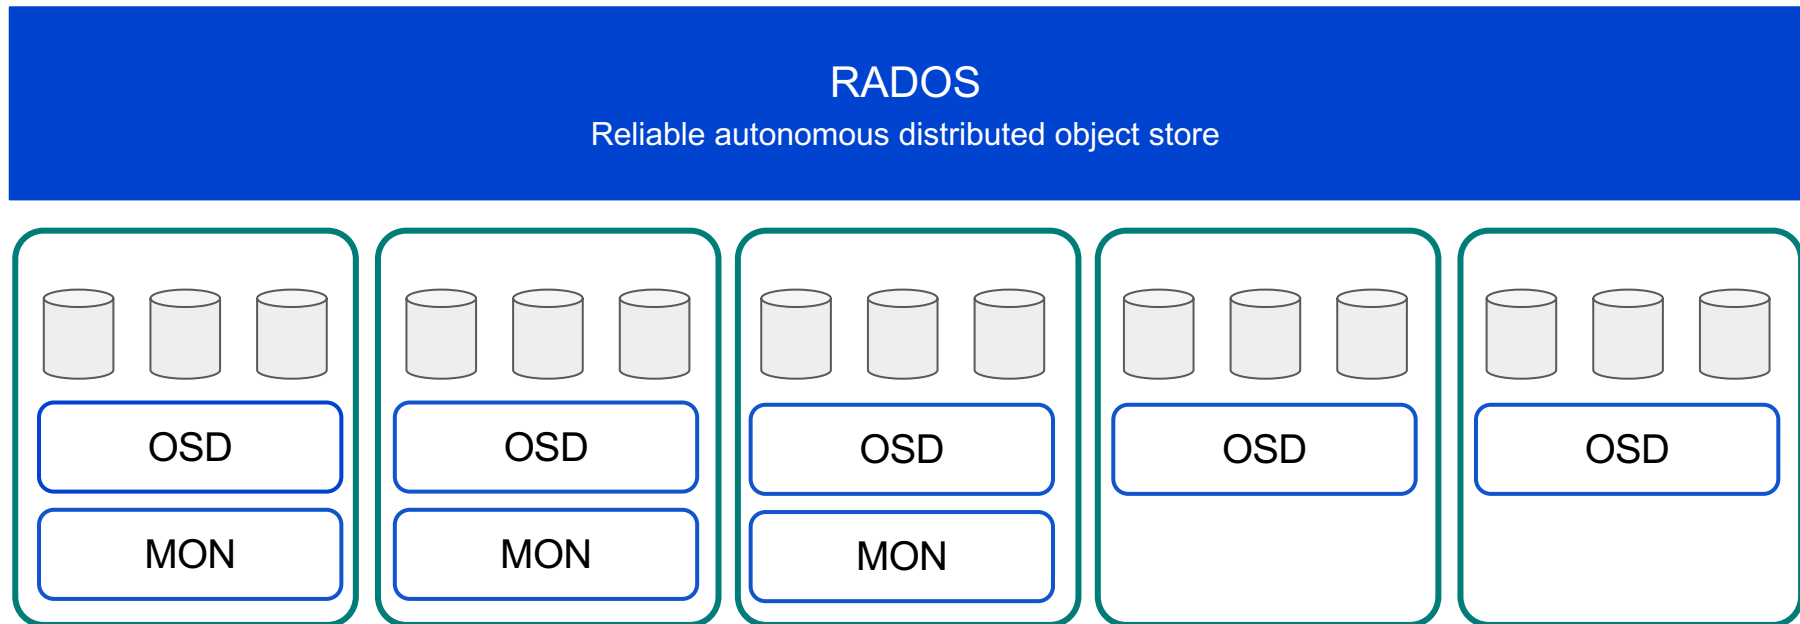
- Easy command-line interface (CLI) and user interface (UI) setup
- A minimal setup contains 4 nodes
- OSDs can scale to 10000s in a cluster
- Tune for performance or capacity



# IBM Storage Ceph RADOS

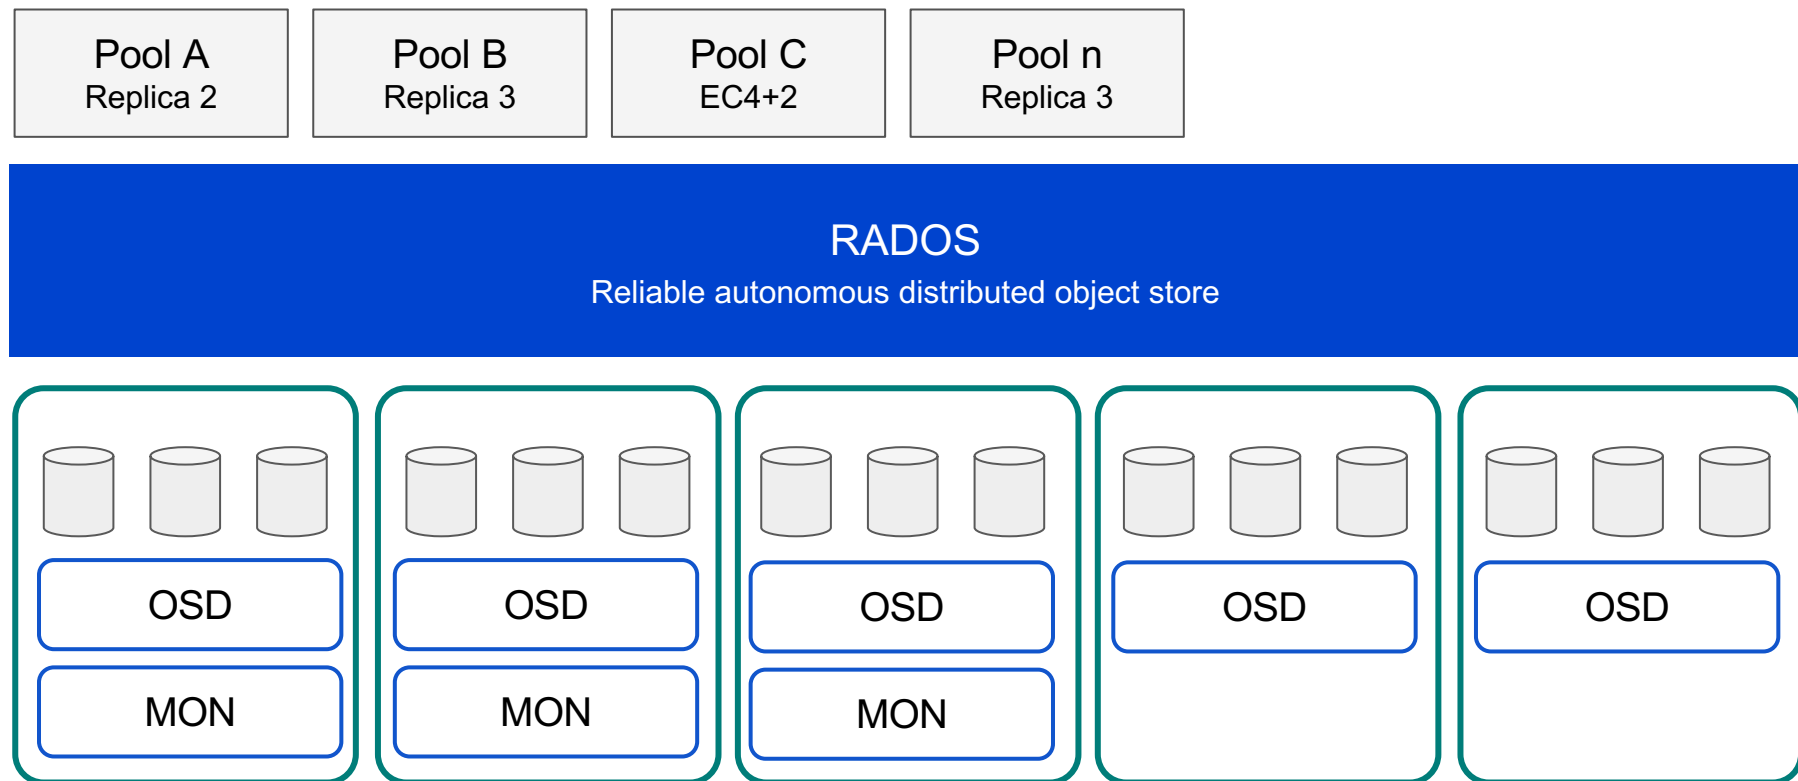
(Reliable Autonomous Distributed Object Store)

- RADOS provides a data access abstraction layer



# IBM Storage Ceph Pools

- Provide a storage access entity for each access method (file, block and object)
- Each pool has an assigned data redundancy scheme (e.g. replicated, erasure-coded)
- Storage pools are usually thin provisioned
- Each data object can only “live” in one Ceph pool

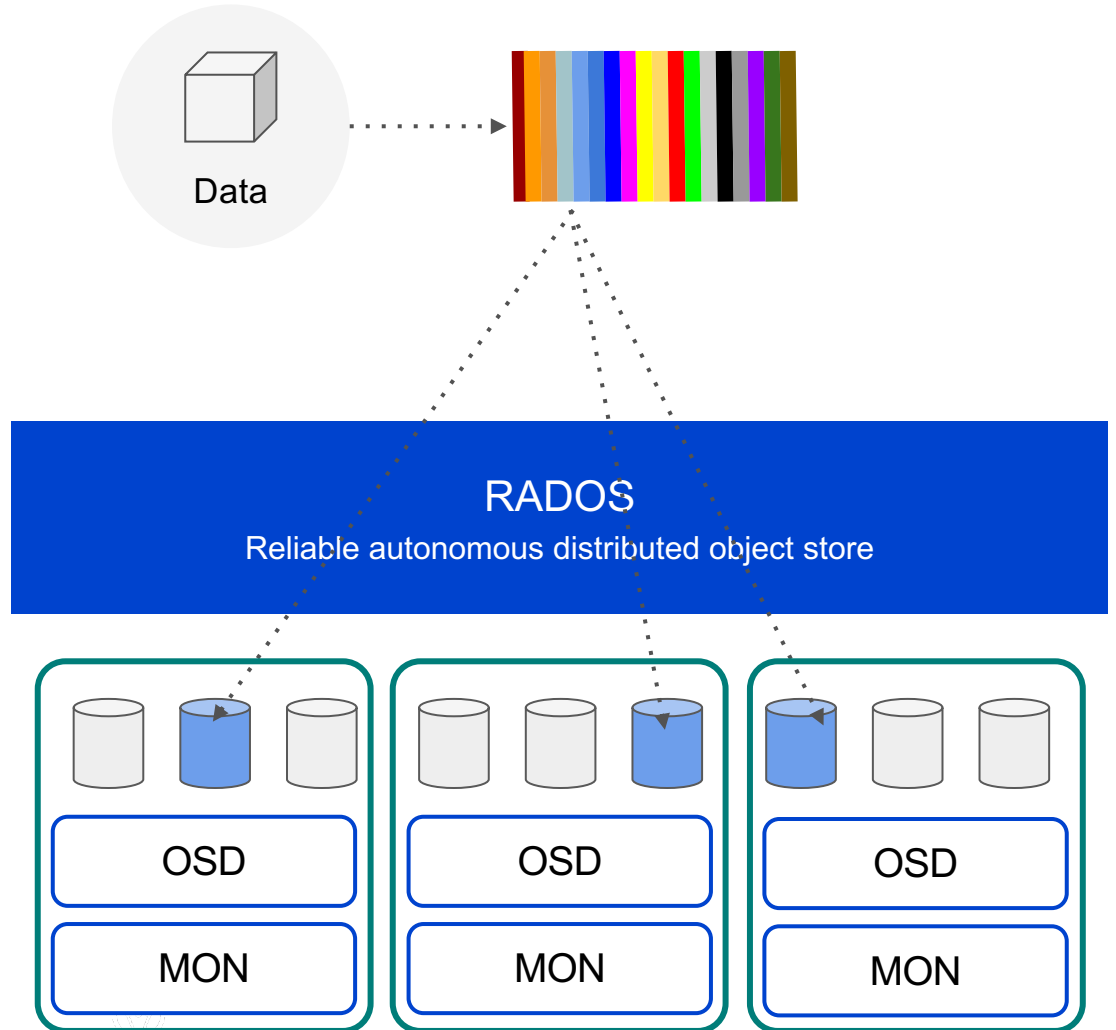


# IBM Storage Ceph data placement via CRUSH

(Controlled Replication Under Scalable Hashing)

- Pseudo-random placement algorithm
- Fast calculation, no lookup, no gateways
- Repeatable and deterministic
- Stable mapping
- Unlimited scalability
- Rule-based configuration
- Adjustable replication
- Weighting
- Rack / Room and DC-aware

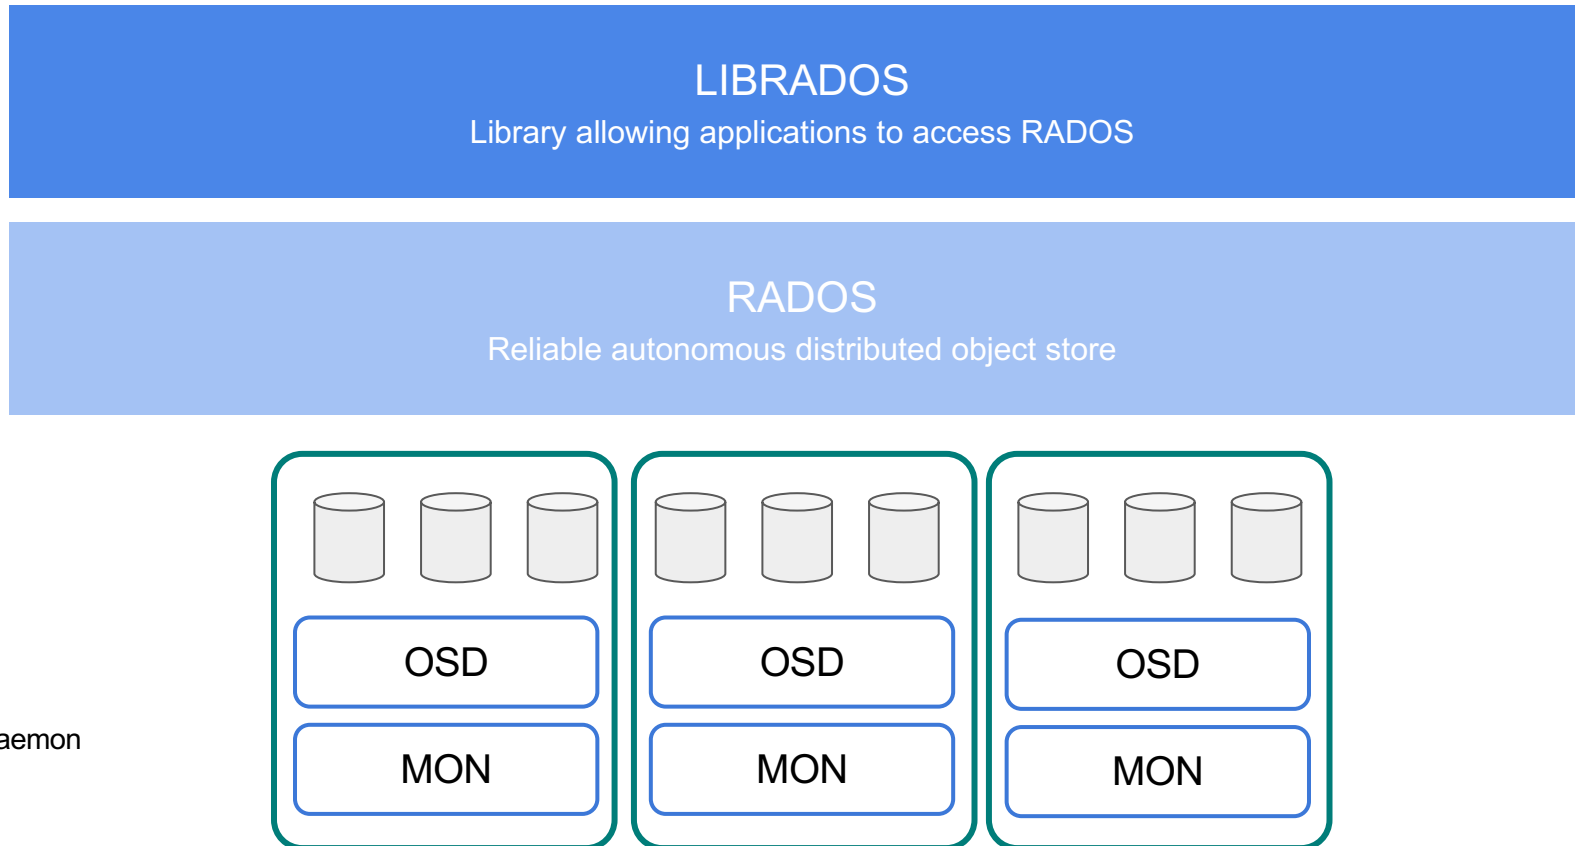
MON	Monitor
OSD	Object storage daemon



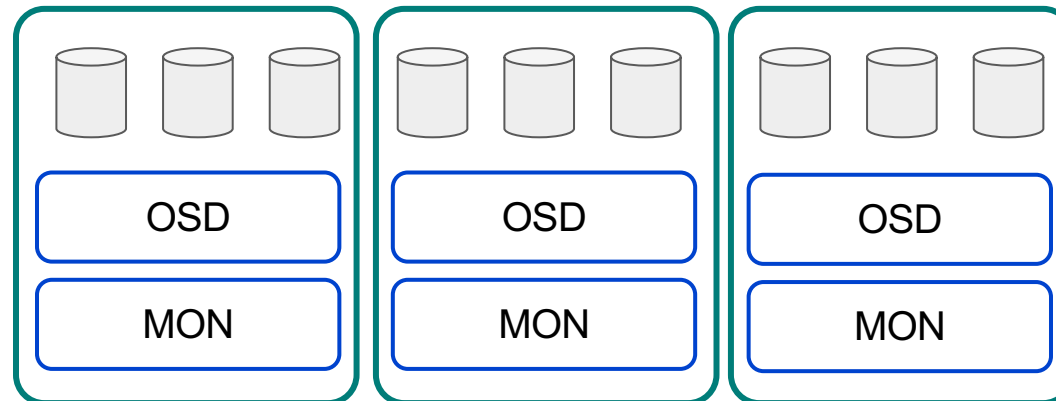



# IBM Storage Ceph Data Access

# IBM Storage Ceph LIBRADOS



# IBM Storage Ceph LIBRADOS API data access



# IBM Storage Ceph Object Access

 S3, Swift API



RADOS gateway  
(RGW)

LIBRADOS

Library allowing applications to access RADOS

RADOS

Reliable autonomous distributed object store



# IBM Storage Ceph Block Access

 S3, Swift API

LibRBD, kRBD

RADOS gateway  
(RGW)

RADOS block device  
(RBD)

LIBRADOS

Library allowing applications to access RADOS

RADOS


Reliable autonomous distributed object store

MON

Monitor

OSD

Object storage daemon



# IBM Storage Ceph File Access

 S3, Swift API

LibRBD,kRBD

FUSE, kernel

RADOS gateway  
(RGW)

RADOS block device  
(RBD)

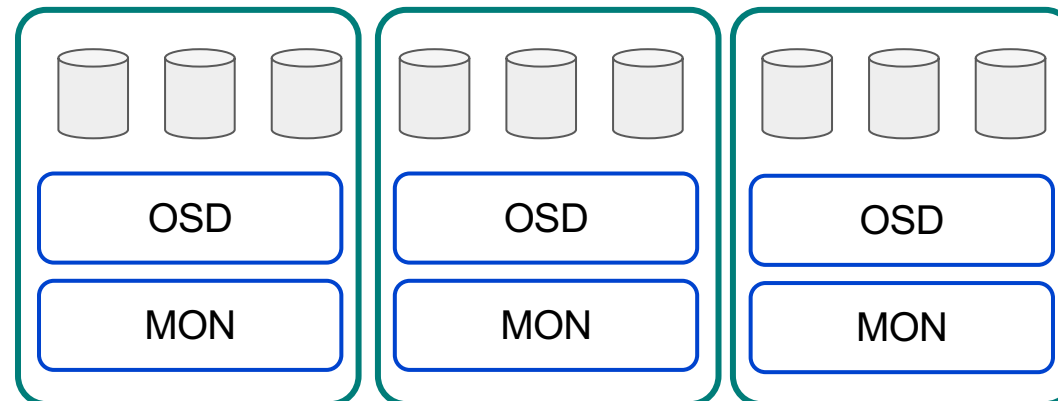

Ceph filesystem  
(CephFS)

LIBRADOS

Library allowing applications to access RADOS


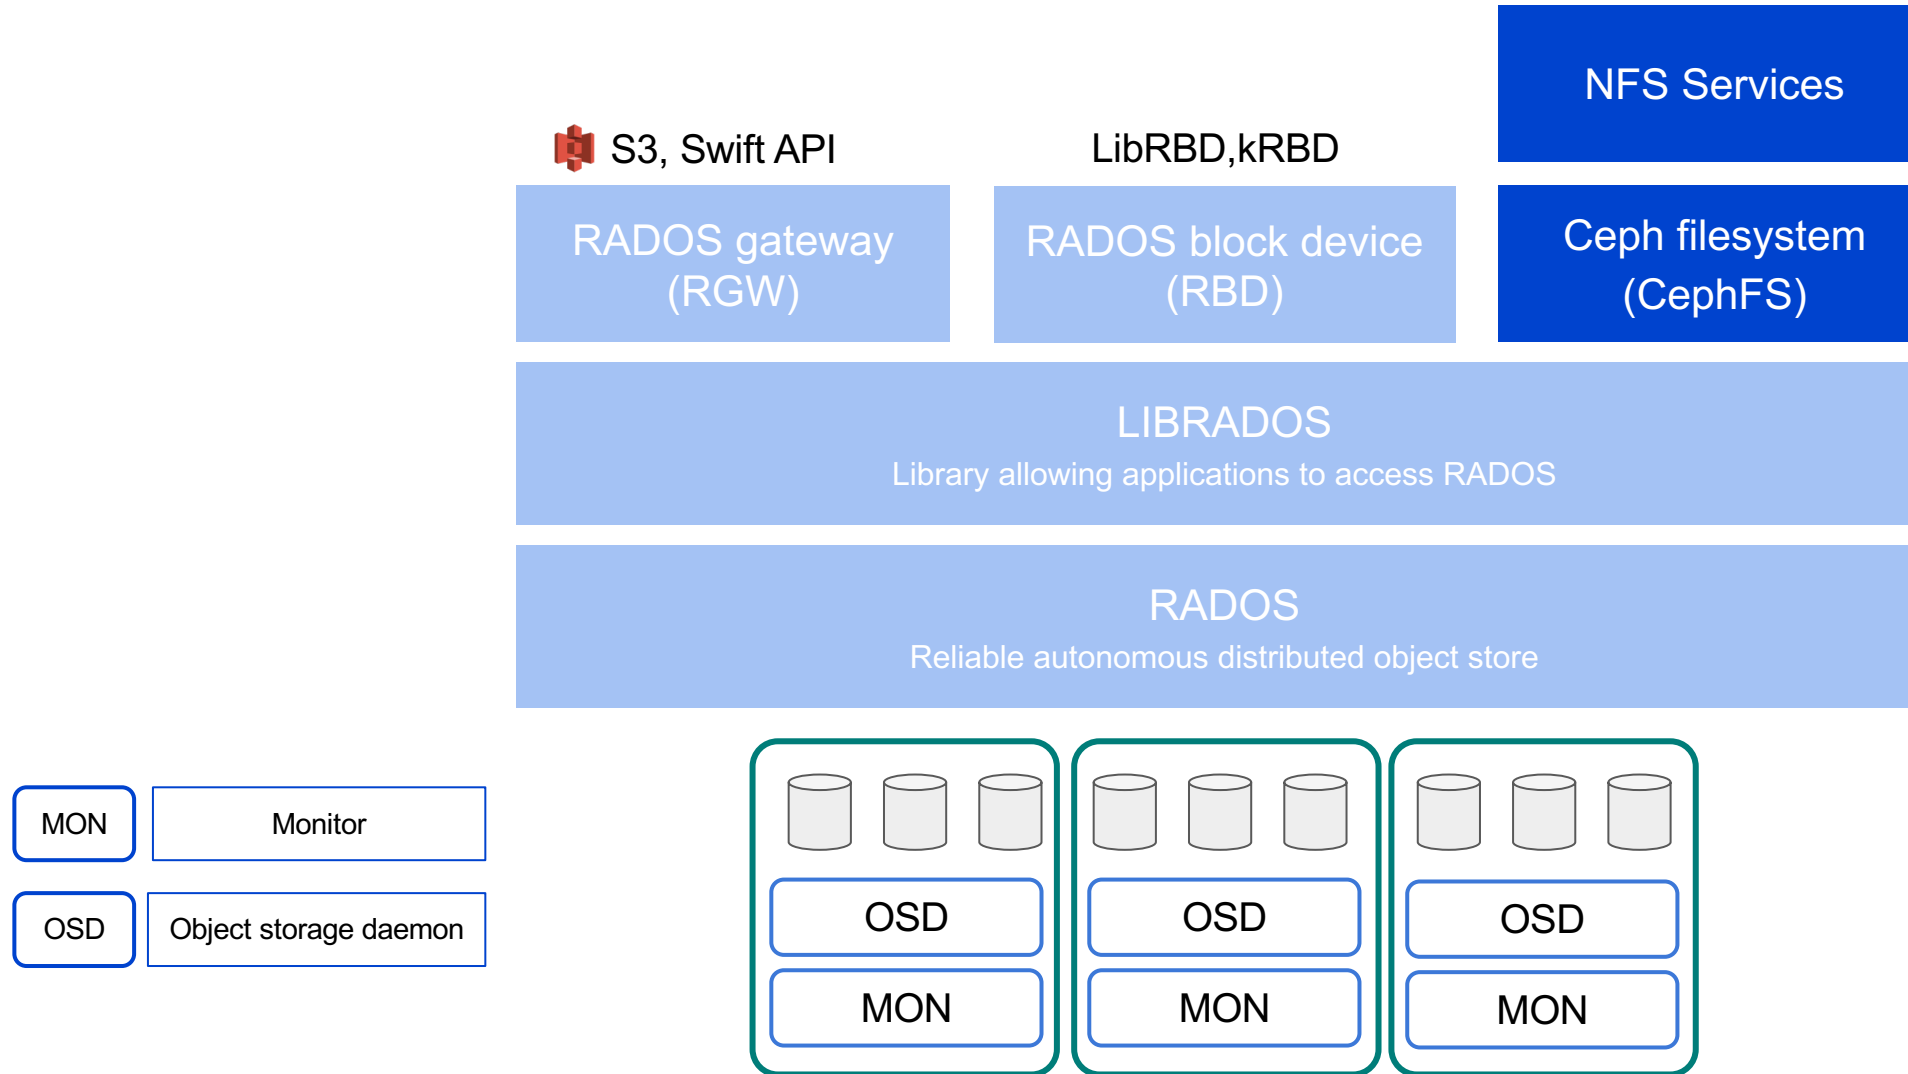
RADOS

Reliable autonomous distributed object store

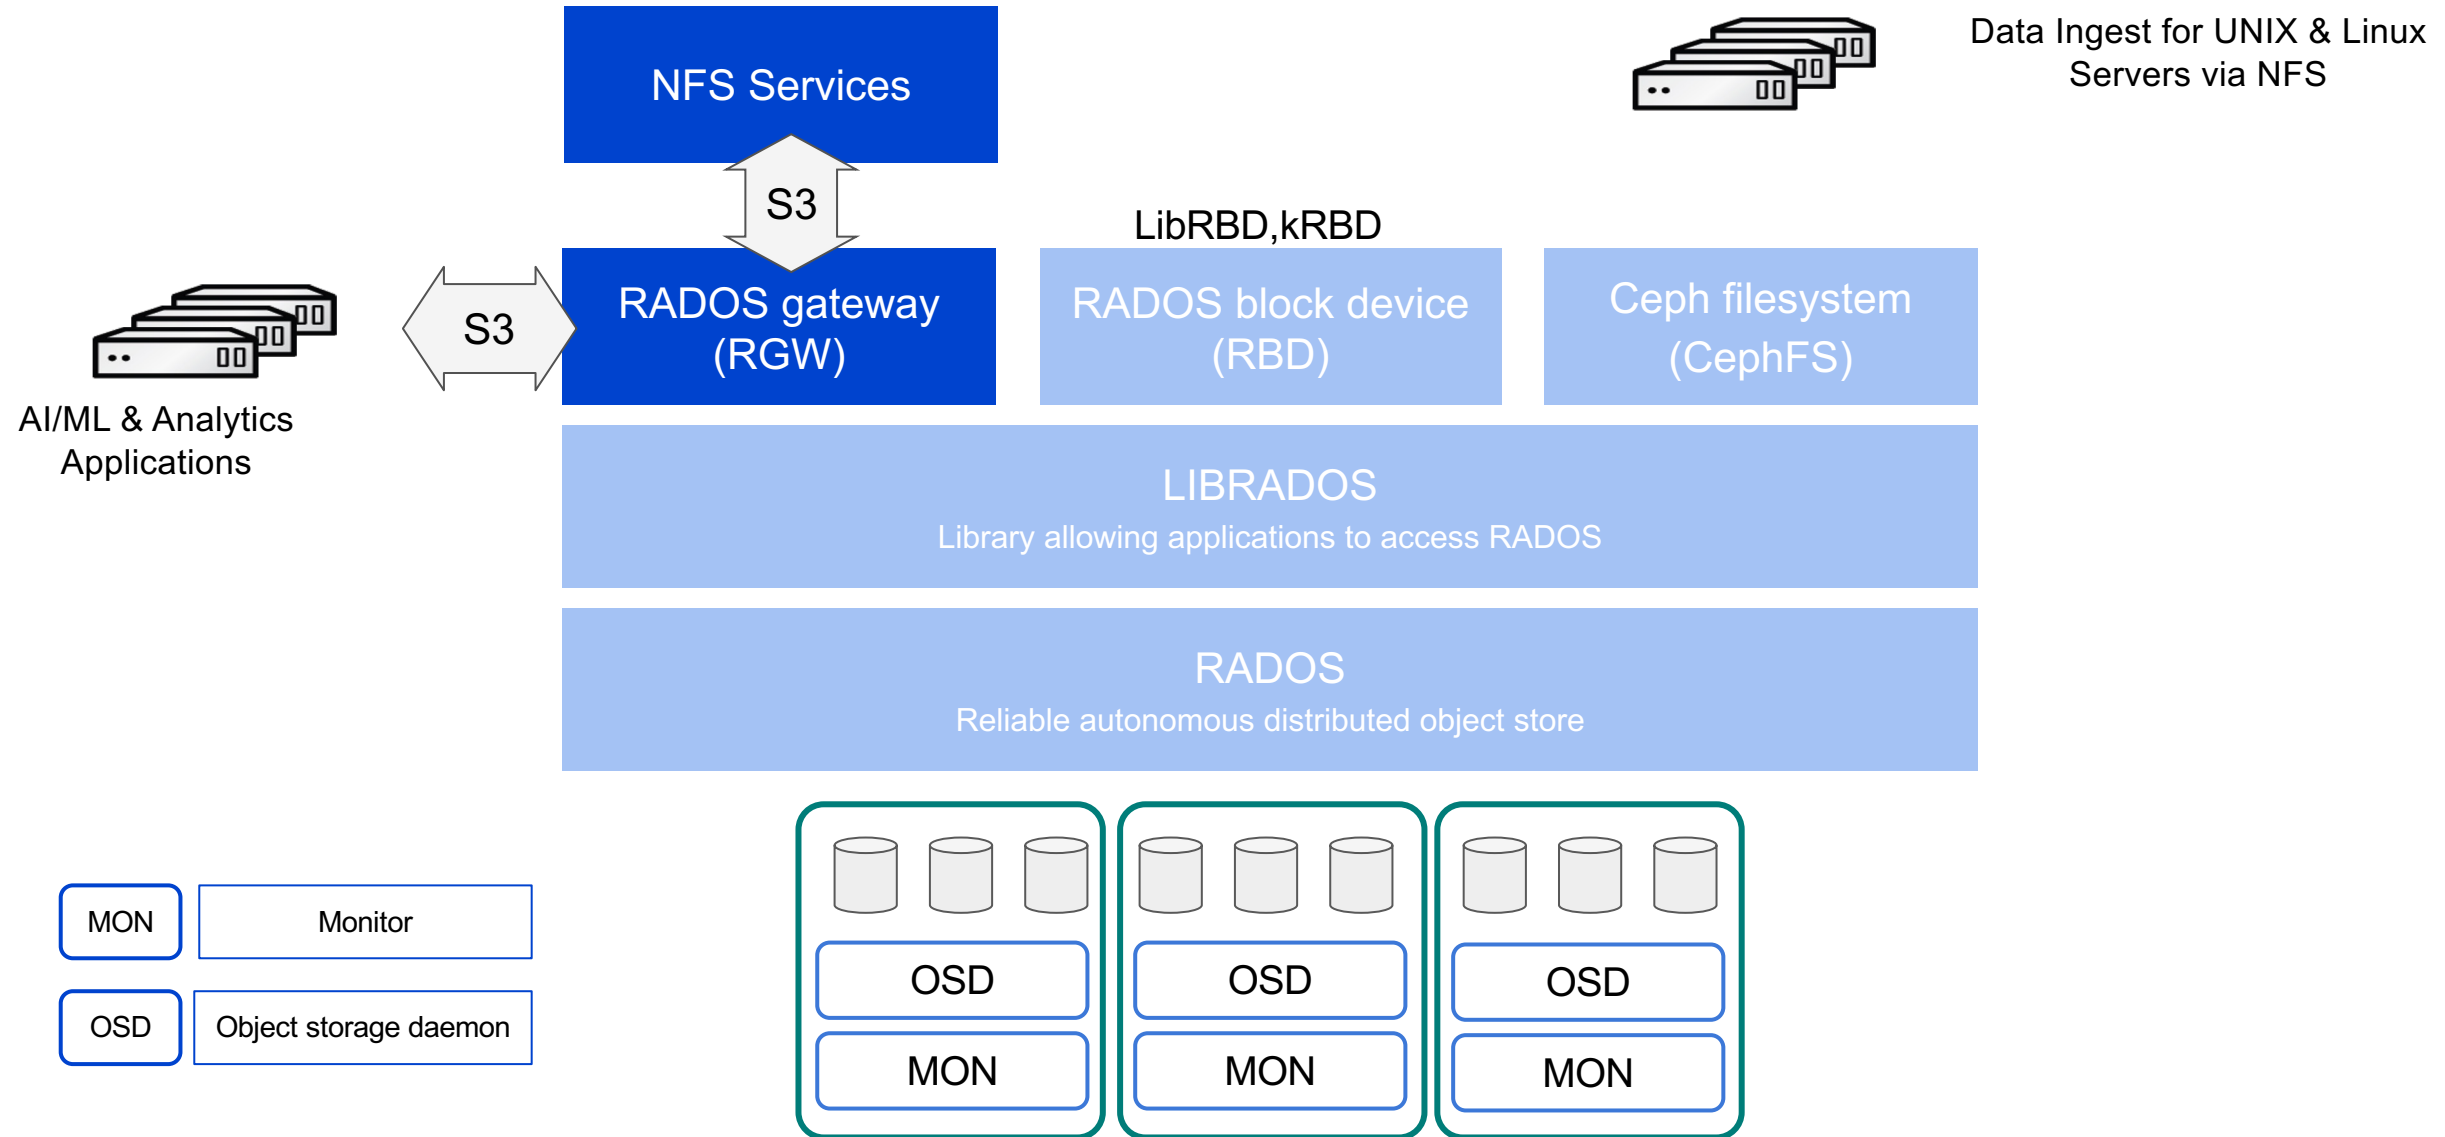


# IBM Storage Ceph File Access (through NFS)

File Shares for  
UNIX & Linux  
Servers

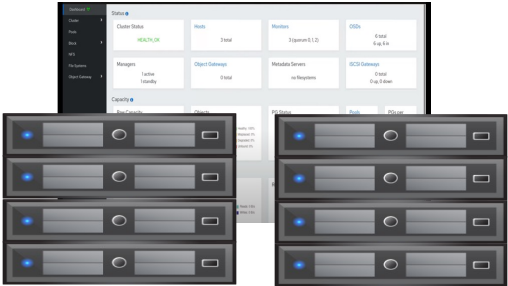
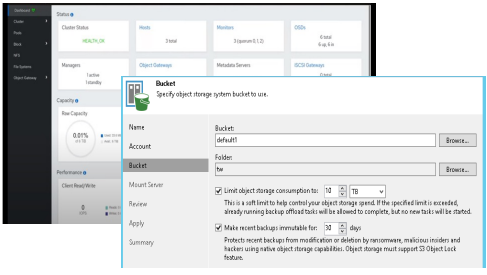
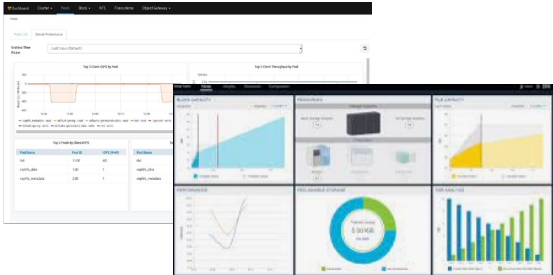


# IBM Storage Ceph File Access (to the Object Store)





# IBM Storage Ceph Easy start process



1

Install your industry-standard x86 servers with Red Hat Enterprise Linux

Start with 4 server nodes

2

Acquire and install IBM Storage Ceph with IBM Storage Insights

IBM Storage Insights is included for use with IBM Storage Ceph

3

Configure object storage buckets and access keys and parameters

Configure your applications with relevant bucket access information

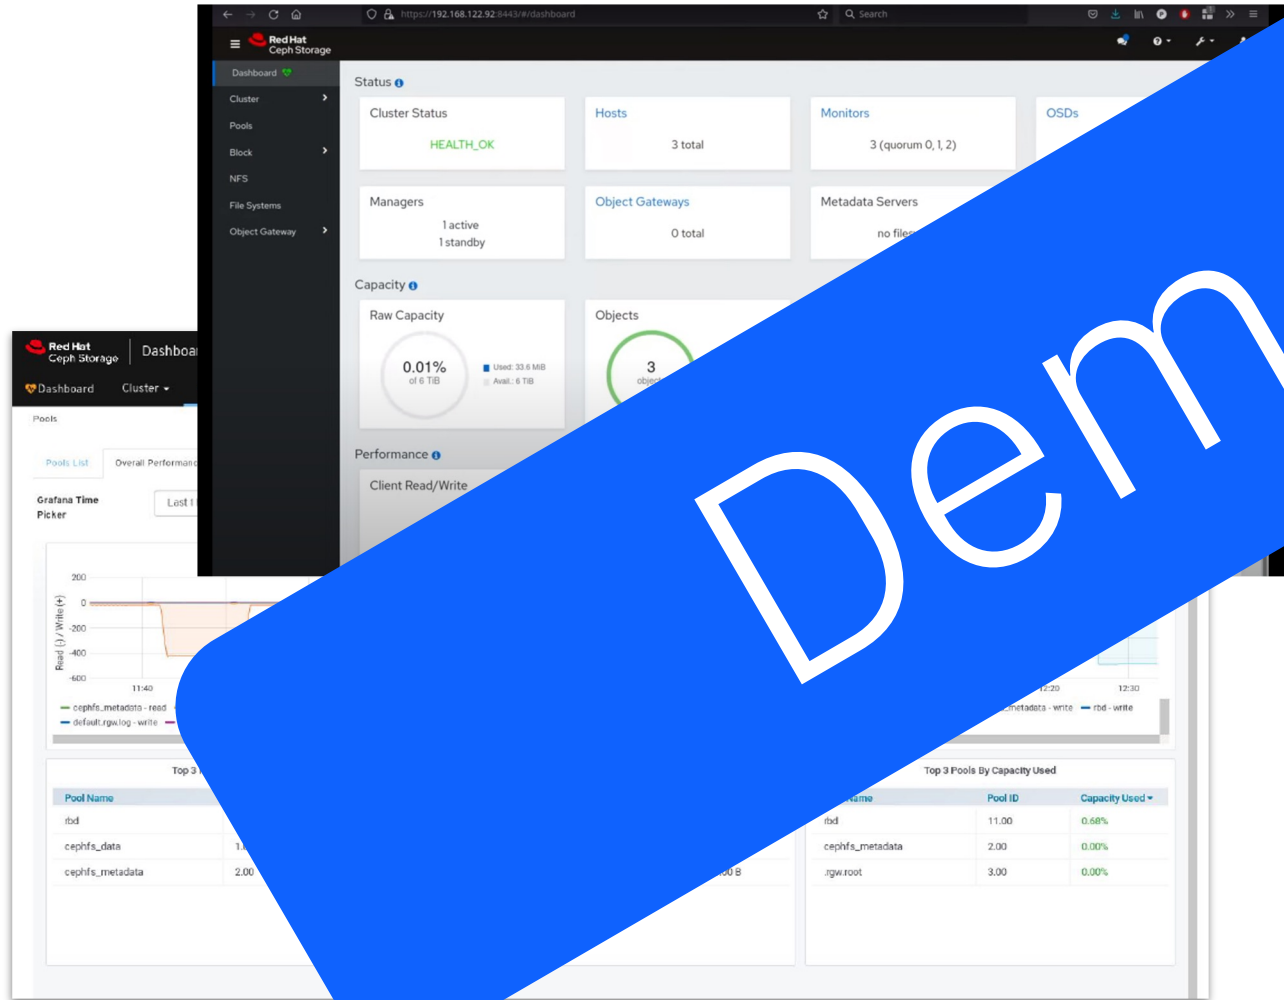
4

Add additional resources and capacity according to needs

Perform day-2 operations, like in example, setup remote replica cluster locations

# IBM Storage Ceph Dashboard

For improved monitoring and management



Demo

- Expansion
- Level overview
- Information of Ceph infrastructure
- SERVICES
- Device and Performance Monitoring
- Simple Day-2 Operations
- Role-based access control (RBAC)



Backup

# Positioning with IBM File/Object Storage

