

Next Generation DataCenter
Mit OpenShift Platform Plus

Storage: MultiCloud, Unified, Converged oder klassisch

Matthias Rettl
Account Solution Architect, Austria
mrettl@redhat.com

9. September 2022

Agenda

1. Red Hat Data Services
2. Red Hat OpenShift Data Foundation - Overview
3. Red Hat OpenShift Data Foundation - Deployment Options
4. Red Hat OpenShift Data Foundation - Benefits in detail
5. Red Hat OpenShift Data Foundation - DEMO
6. Q&A

Stellen Sie Ihre Fragen bitte gerne im Q&A Chat 



Your host for today




Matthias Rettl

Account Solution Architect, based in Austria

Joined Red Hat in December 2019

Previously worked as a Storage Sales Engineer at NetApp

Red Hat Data Services



Why Data Services?



- Data is core to business
- ...bring them **to live** to generate business value
- **The use is key - this gives business value**



Simplified
access

Consistent
experience



Dynamic
scale

Red Hat Data
Services mission:

To make data
accessible to
applications across
the hybrid cloud,
unlocking its power
in new and
impactful ways

Delivering on the Red Hat
OpenShift promise:

Innovation without
limitation





Red Hat Ceph Storage



- 100% Software Defined Storage (SDS) Solution
- Founded 2004 by Sage Weil, acquired 2014 by Red Hat
- The only "Storage product" by Red Hat as of now*

Block, File and Object
Storage

Massively scalable

100's of Petabytes
Billions of Objects

Extremely robust, resilient
and stable

- Product Name: "Red Hat Ceph Storage" (RHCS)
- Community Release Name: "Ceph"
- RHCS runs on RHEL (community on any Linux and even Windows)
- RHCS runs **Containerized**, on Bare-Metal and Virtualized
- RHCS runs on (almost) any Industry-Standard Hardware (Servers and Media-Devices - f.k.a. "Disks")



Red Hat OpenShift Data Foundation



Ceph inside

Easy to consume storage

+ **Rook Operator**

Does the operator's work

+ **NooBaa MCG**

S3 at it's best

- "ODF" is "Ceph on/for OpenShift" and more...
- Contains **Red Hat Ceph Storage** (when used "*internal*")
- + **Rook Operator** (takes care about Ceph)
- + **NooBaa Multicloud Gateway** (distribute S3 Objects everywhere)

- "**Internal Mode**": ODF installs Ceph into the OpenShift Cluster
 - For customers who do not have RH Ceph yet
 - Data can only be accessed by containers inside OCP
- "**External Mode**": An existing Ceph runs outside OpenShift
 - For customers who already have Ceph (mostly)
 - Data can be accessed from everywhere, not only OCP



Red Hat OpenShift Data Science




Sandbox for AI/ML
in the public cloud

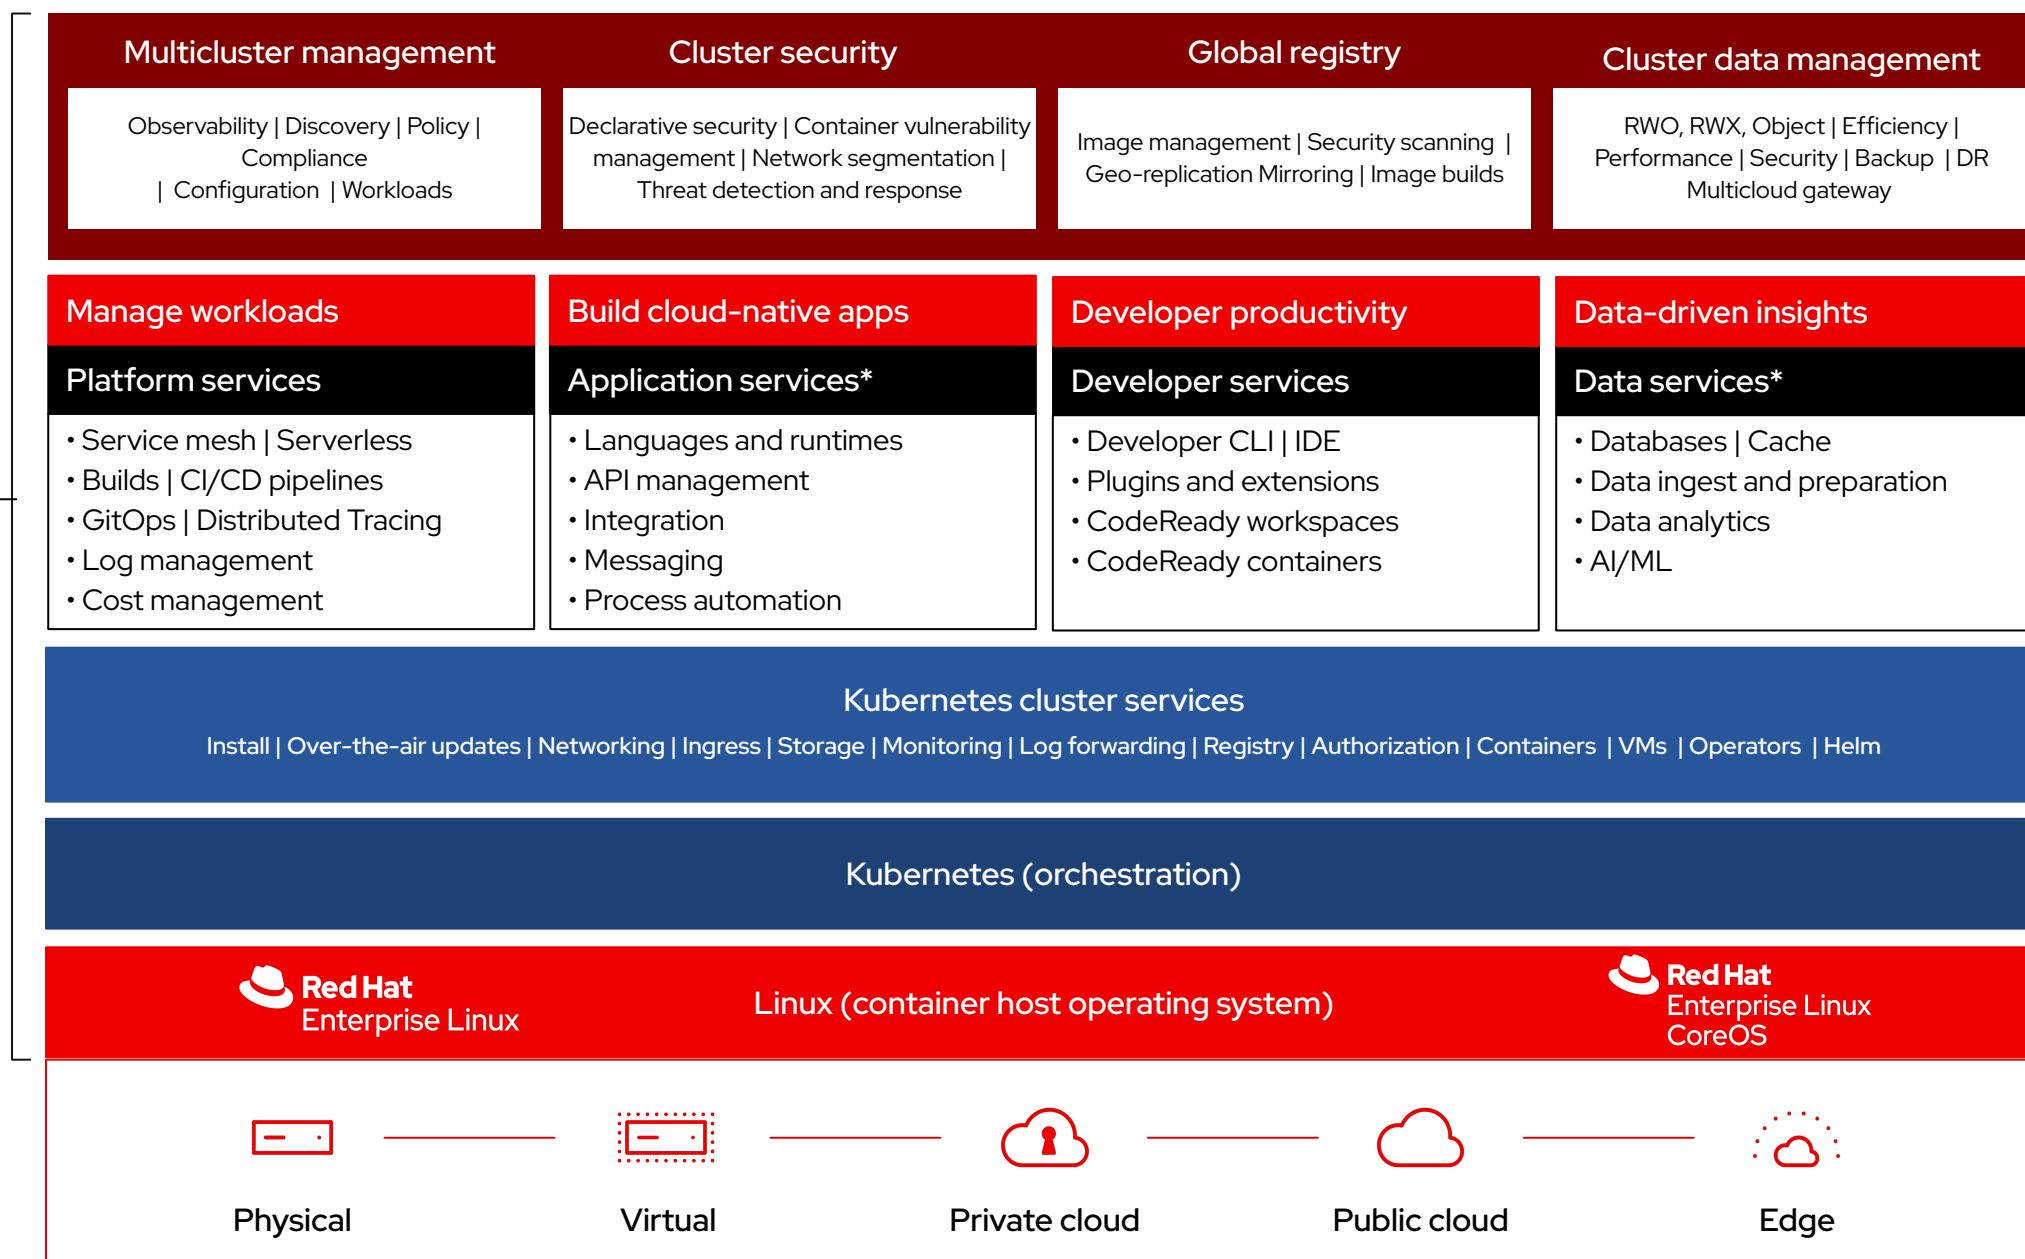
For Data Scientists
to quickly develop
models

Red Hat OpenShift
as the platform

- Not “Storage” in a classical sense
- Fully supported and managed sandbox for data scientists (AI/ML)
- Quickly develop, train, and test containerized ML models in the public cloud*
- OpenShift Data Science is a Subset of *Operate First* delivered as a cloud service on Red Hat OpenShift Managed on Amazon Web Services with optional ISV offerings
- Operate First is a Subset *Open Data Hub* operated at scale for community and university audiences to infuse operational excellence
- Open Data Hub is a Community driven upstream meta-project demonstrating AI/ML platform on Red Hat OpenShift comprised of open source projects

Red Hat OpenShift Data Foundation Overview







Persistent storage for containers

Portability

Enable seamless data placement and access across clouds.

Simplicity

Streamline application development workflows across the hybrid/multicloud.


Scale

Support data-intensive emerging workloads for enterprise Kubernetes customers.

Completeness

Single resource for all data services


What is inside OpenShift Data foundation?



What is Rook?




- Storage Operator for Kubernetes
- turns distributed storage systems into self-managing, self-scaling, self-healing storage services
- automates the tasks of a storage administrator: deployment, bootstrapping, configuration, provisioning, scaling, upgrading, migration, disaster recovery, monitoring, and resource management



What is NooBaa?

- S3 compatible Multicloud Gateway
- aggregate native storage from AWS S3, Azure, Google and even private cloud - like Ceph - into one virtual storage
- quickly connect on-premises resources, with cloud-native storage and let your workload access it seamlessly
- virtualize all your on-premises storage silos from all your data centers into one cloud-like virtual storage

Red Hat OpenShift Data Foundation Deployment Options

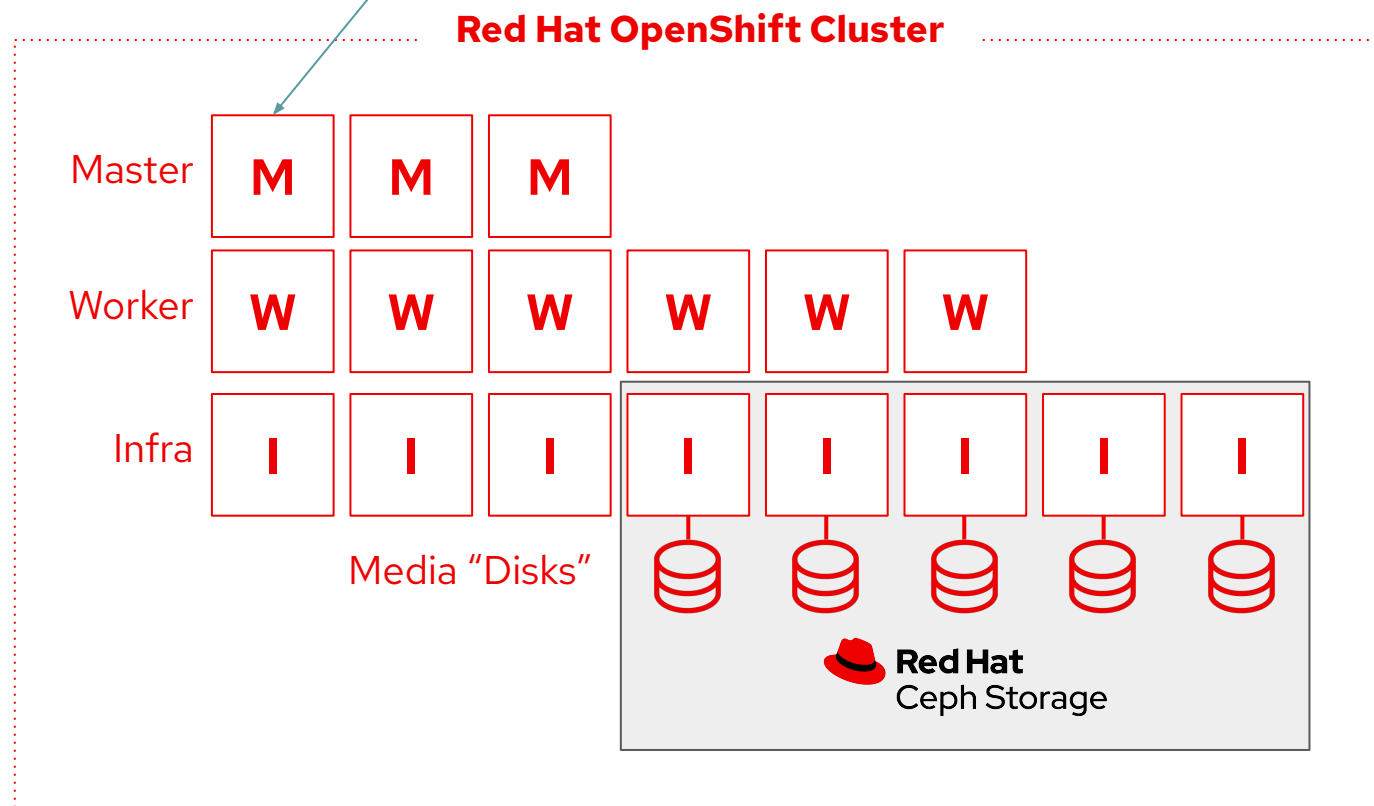


Red Hat OpenShift Data Foundation

ODF “Internal Mode”

- Ceph runs inside the OCP Cluster
- Rook Operator takes care of almost all Day-2 tasks
- Only Pods of this cluster can access data of Ceph
- No access to Ceph from outside is possible*
- Deploy on Infra-Nodes to save costs (no OCP Subs needed)

One box represents a server, physical (Bare Metal) or virtual (VM)

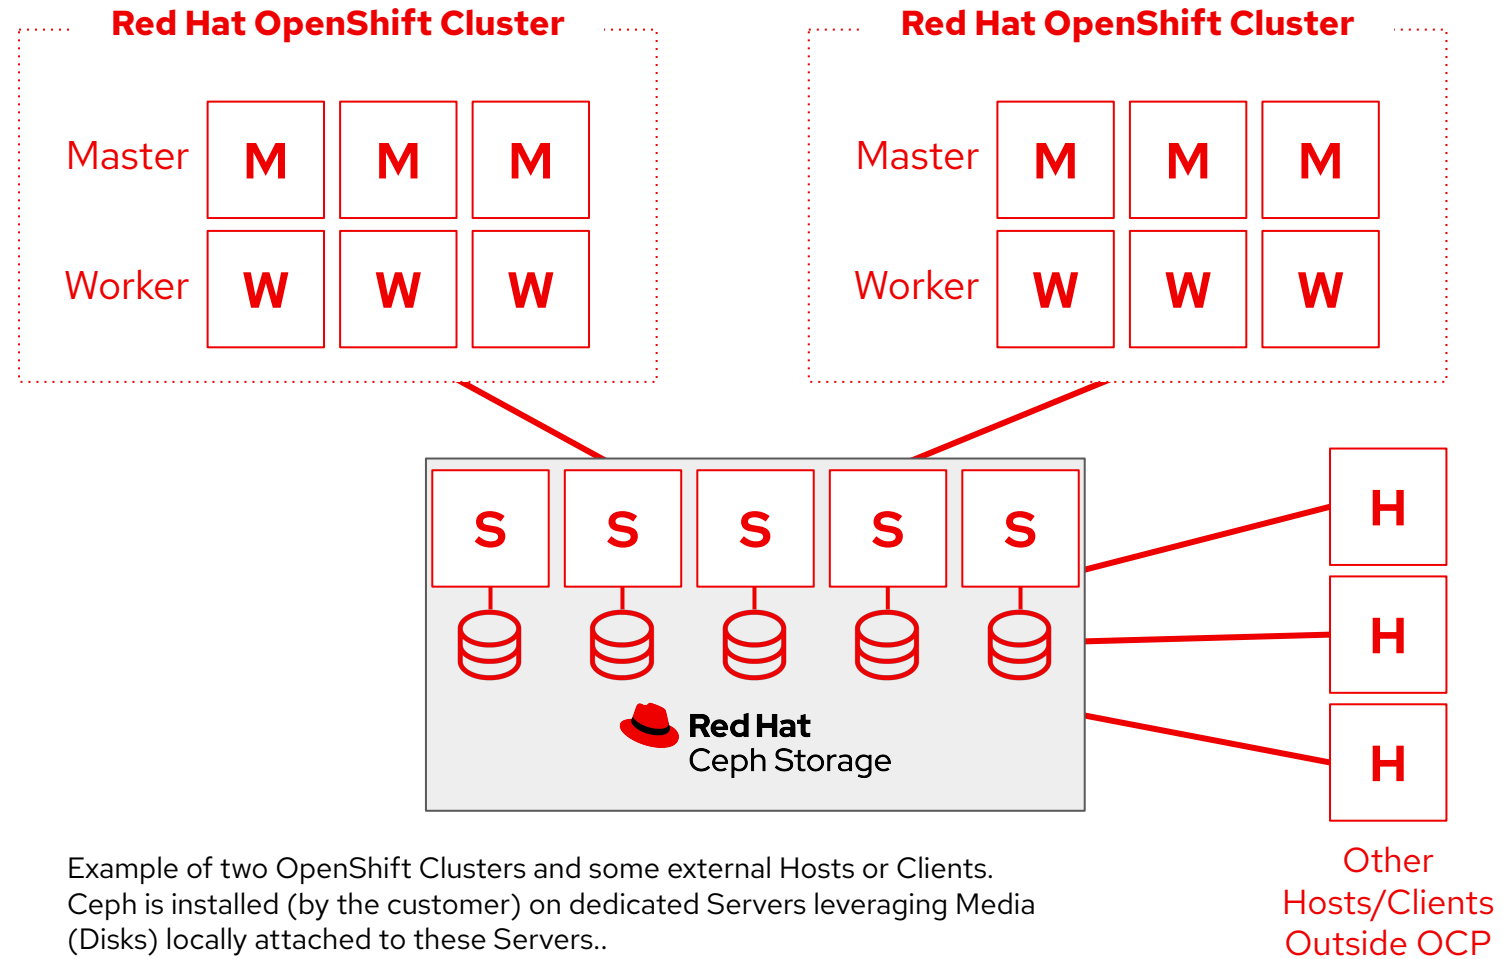


Example of an OpenShift Cluster with 3 Master-, 6 Worker and 8 Infra-Nodes. Ceph is installed (by Rook) on the Infra-Nodes as a 5-Node Cluster, leveraging Media (Disks) locally attached to the Infra-Nodes.



ODF "External Mode"

- Ceph runs outside of the OCP Cluster
- Ceph must be configured and maintained manually
- Pods of every attached OCP Cluster can access data of Ceph
- Ceph access from outside is absolutely possible





Essentials Edition

- Kubernetes RWO (block, file)
- Kubernetes RWX (shared file, shared block)
- Object storage (s3-compatible)
- NooBaa Multi-Cloud Gateway
- Internal mode only (on-cluster)
- Volume snapshots/clones
- Cluster-wide encryption
- Cross-availability zone HA (3 zones)

Included with
Red Hat OpenShift Platform Plus
at no additional cost
up to 256 TB RAW capacity



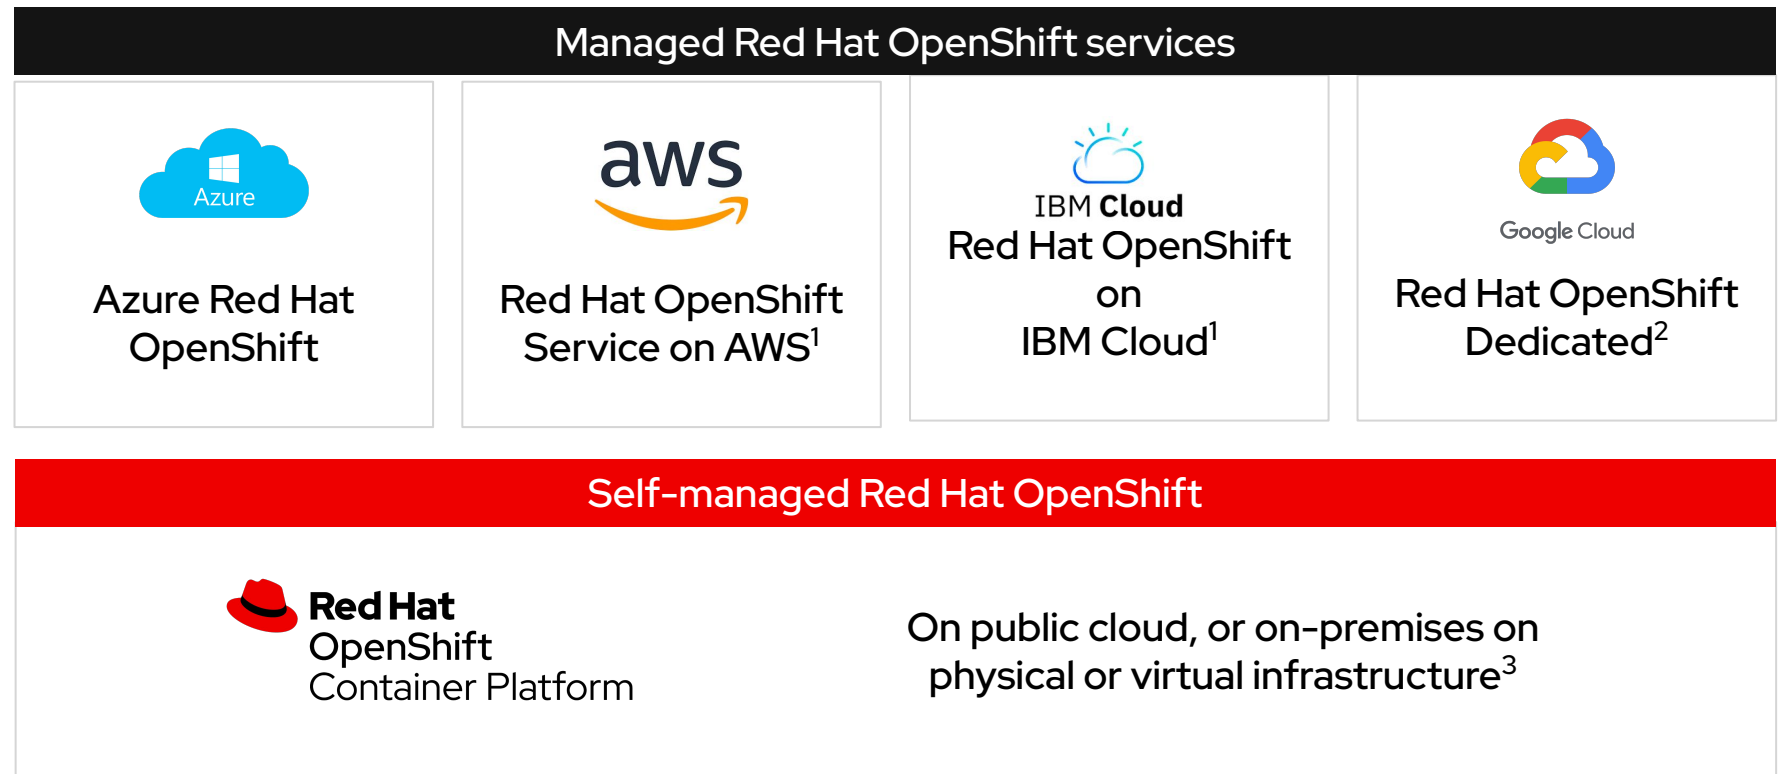
Advanced Edition

- All Essentials-Features plus:
- **External mode** (shared cluster)
- **Granular encryption** (Volume-level)
- **External Key Management System** (BYOK)
- **Metro-DR** (Tech Preview)
- **Regional-DR** (Tech Preview)

Supporting hybrid usage and buying patterns

A consistent platform no matter how or where you run

Start quickly, we manage it for you




You manage it, for control and flexibility

¹ In preview as of 1/1/2021. Also available as Red Hat OpenShift Dedicated managed service running on user-supplied AWS infrastructure.

² Red Hat managed service running on user-supplied GCP infrastructure

³ See docs.openshift.com for supported infrastructure options and configurations

Red Hat OpenShift Data Foundation Benefits in detail



Why Red Hat OpenShift Data Foundation ?



Abstraction

no learning curve,
no specialities




Economical Control

all small to huge PVs at
max performance




Flexibility

any number
of PVs per node



Performance

no performance variation
based on capacity




Self-Service

fully automated -
no special access
credentials needed



Fully Secured

complete data
encryption



ACM integration

application + data
mobility + DR




Data Availability

Failure domain
bridging



Freedom of Data Location

aggregation /
distribution of objects



Full feature support

no dependency on k8s
upstream readiness

Other important functionality of Red Hat OpenShift Data Foundation ?



True block device for RWO and RWX

local filesystem for single reader/writer, live migration for VMs



True POSIX file system for RWX (no NFS)

no impact for applications that do not support NFS



Object access acceleration

performance at scale



Backup / Restore

OADP,
established Backup S/W vendors



No data locality

no difference for workload with location of data



True object storage with full security

S3 objects unlimited* w/o filesystem dependency



Self-service for object storage

no need to wait for external provisioning and credentials



Full object storage location and flavor abstraction

true hybrid cloud, cross-cloud redundancy



No additional infrastructure

runs internal to and is maintained by OCP and OCP networking



No incompatibility:

consumption of PVs is native to OCP nodes, unlike external storage

Why Red Hat OpenShift Data Foundation for **storage teams** ?



Abstraction:
no learning curve,
no specialities



Secure by default:
full isolation from
providing storage



Access from nodes only:
no need to integrate
workloads with storage



Scale easily:
More Capacity ?
Provide more volumes.



**Fast and Simple
resources:**
no special features
needed, no intelligence in
storage systems



Cross location redundancy:
transparent replication built-in




No complexity:
no dependency between
PVs and volumes on
storage



Speed of management:
once installed, all
provisioning done in ODF

Red Hat OpenShift Data Foundation DEMO



Overview - Red Hat OpenShift

https://console-openshift-console.apps.cluster-glkvb.glkvb.sandbox1154.opentlc.com/dashboards

120%

opentlc-mgr

Red Hat OpenShift Container Platform

Administrator

Home

Overview

Projects

Search

API Explorer

Events

Operators

Workloads

Networking

Storage

Builds

Observe

Compute

Overview

Cluster

Details

View settings

Cluster API address

https://api.cluster-glkvb.glkvb.sandbox1154.opentlc.com:6443

Cluster ID

667dbaeb-0e0c-4e7c-b158-62c68b3fad20

OpenShift Cluster Manager

Provider

AWS

OpenShift version

4.10.26

Update cluster

Service Level Agreement (SLA)

Self-support, 60 day trial

59 days remaining

Manage subscription settings

Update channel

stable-4.10

Cluster inventory

Status

View alerts

Cluster

Control Plane

Operators

Insights

1 issue found

A cluster version update is available

Update cluster

2 Sep 2022, 10:11

Alerts are not configured to be sent to a notification system, meaning that you may not be notified in a timely fashion when important failures occur. Check the OpenShift documentation to learn how to configure notifications with Alertmanager.

Configure

Cluster utilization

Filter by Node type

1 hour

Resource

Usage

CPU

2,58

57,42 available of 60

Memory

26,73 GiB

205,7 GiB available of 232,4 GiB

Activity

View events

Ongoing

There are no ongoing activities.

Recent events

Pause

readyz=true

All pre-shutdown h...

All pending reques...

Received signal to ...

The minimal shutd...

readyz=true

Received signal to ...


All pre-shutdown h...

All pre-shutdown h...

All pending reques...

The minimal shutd...


All pre-shutdown ...



Red Hat OpenShift Data Foundation

What did we learn today

- OpenShift Data Foundation = Ceph+Rook+NooBaa
- 100% Software Defined Storage für die Hybrid-Cloud
- OnPrem, Bare-Metal, Virtualisiert, Edge & Cloud
- Maximale Abstraktion, Sicherheit und Robustheit
- Self-Service ohne Storage Know-How
- start small – GROW BIG
- Scale-Out ohne Limits
- Integration in Advanced Cluster Manager (ACM)



Red Hat OpenShift Data Foundation Q&A

Overview: Das Next Generation Datacenter mit Red Hat gestalten

Franz Theisen

19.8.2022, 11.00 - 12.00 CEST

Compute: Virtualisierung und Container auf einer Plattform

Domenico Piol

26.8.2022, 11.00 - 12.00 CEST

Management

Robert Baumgartner

2.9.2022, 11:00-12.00 CEST

Storage: MultiCloud, Unified, Converged oder klassisch

Matthias Rettl

9.9.2022, 11:00-12.00 CEST


Networking & Security

Robert Bohne

16.9.2022, 11:00-12.00 CEST

Anmeldung & Recordings:

<https://red.ht/3STjQg3>



Thank you

Red Hat is the world's leading provider of enterprise open source software solutions. Award-winning support, training, and consulting services make Red Hat a trusted adviser to the Fortune 500.

 linkedin.com/company/red-hat

 youtube.com/user/RedHatVideos

 facebook.com/redhatinc

 twitter.com/RedHat