

Advanced Technology Group



A Ceph Primer What's the Difference Between IBM Storage Ceph and Ceph in IBM Fusion

Shu Johnbeck (Mookerjee)

ATG Senior Storage Technical Specialist

Shu.Mookerjee@ibm.com



Agenda

- IBM Storage Fusion Overview
- Fusion Data Foundation (FDF)
- Leveraging Ceph from IBM Storage Fusion
- IBM Storage Fusion Walkthrough – A Look at Where Ceph “Lives”
- Standalone Ceph & Fusion Ceph (A Philosophical Discussion)

Goals and Objectives

Objective:

Explain where and how Ceph operates within IBM Storage Fusion

We WILL:

- Review *high-level* architecture of Fusion
- Explain where the Ceph components can be found
- Demonstrate how to utilize Ceph

We WILL Not:

- Cover in-depth deployment of IBM Storage Fusion or the Ceph components
- Provide a deep dive into the IBM Storage Fusion architecture
- Explore all the data service mechanics of Ceph in Fusion

IBM Storage Fusion Overview – The Current Storage Portfolio

Providing solutions around

01 IBM Storage for Hybrid Cloud

Drive innovation and scale application modernization with container-enabled enterprise storage that deploys seamlessly across hybrid infrastructures with a simple and consistent user experience.

Chief Technology Officer | IT Director | VP OpenShift Engineering | Director Open Infrastructure | Cloud Architect | Data Scientist

IBM Storage Fusion

IBM Storage Fusion HCI System

02 IBM Storage for Data and AI

Accelerate business results and innovation and unlock the latent value of unstructured data across the data ecosystem by eliminating data silos, advancing data discovery and classification.

VP of Engineering | VP of Development | Chief Data Officer | Data Architect | HPC Specialist

IBM Storage Scale
IBM Storage Ceph

IBM Storage Scale System

03 IBM Storage for Data Resiliency

Reduce the threat exposure window from days to hours and proactively safeguard data with a multi-faceted and scalable data resiliency approach that defends against cyber vulnerabilities from detection to recovery.

IT Architect | Storage Architect | Chief Information Security Officer | IT Director

IBM Storage Defender

IBM Storage DS8K IBM Storage Tape
IBM Storage FlashSystem

Designed for

Leading with

Delivered on

That run on sustainable infrastructure

Edge – to – Core – to – Cloud

IBM Storage Fusion Overview – HCI and Fusion Software

IBM Fusion HCI



Fully Integrated Appliance

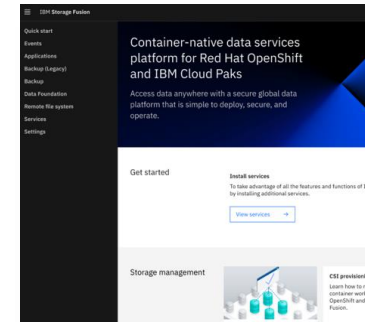
- Fusion software
- Rack
- Servers
- Storage
- Network switches
- Cables
- PDUs

Consistent Data Services

- Data Persistence
- Data Security
- Data Discovery
- Data Mobility
- Data Resilience

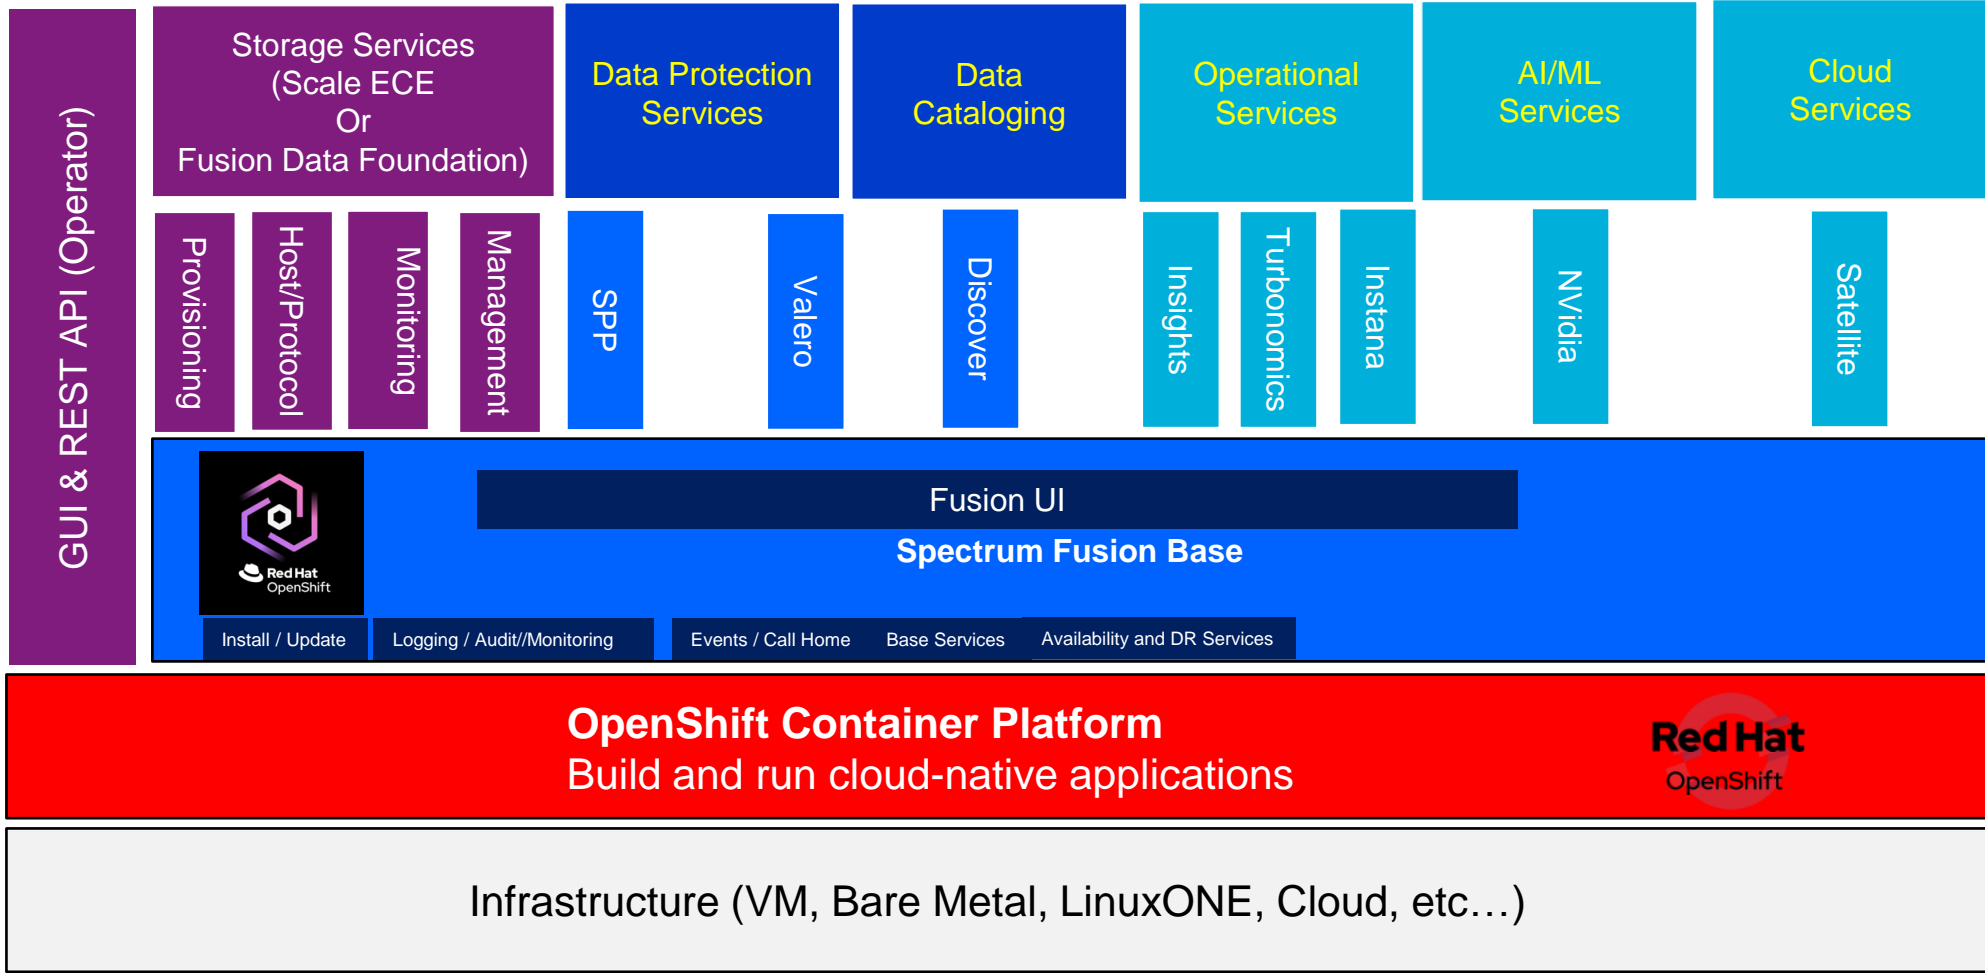


IBM Fusion Software



IBM Storage Fusion Overview – Architecture

Ceph is in here! (If you use Fusion Data Foundation)



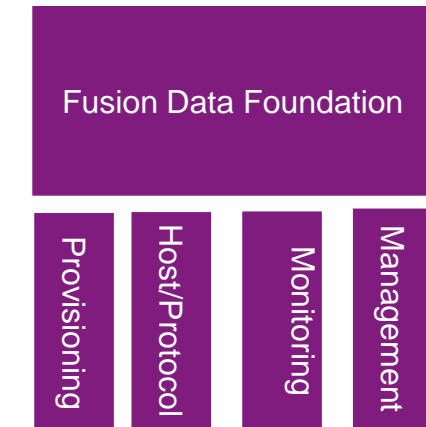
GUI & REST API (Operator)

IBM Fusion Data Foundation - Overview

Fusion Data Foundation (FDF):

- Software-defined storage that's optimized for container environments.
- Runs as an operator on OpenShift Container Platform to provide highly integrated and simplified persistent storage management for containers
- Efficiently attaches and detaches persistent data volumes
- Provides:
 - Block storage for databases
 - Shared file storage for continuous integration, messaging, and data aggregation
 - Object storage for cloud-first development, archival, backup, and media storage
- Scales data exponentially to accommodate applications
- Enables stretched clusters across multiple data-centers or availability zones
- Establishes a comprehensive application container registry

In other words, it provides most of the day-to-day storage services for IBM Fusion



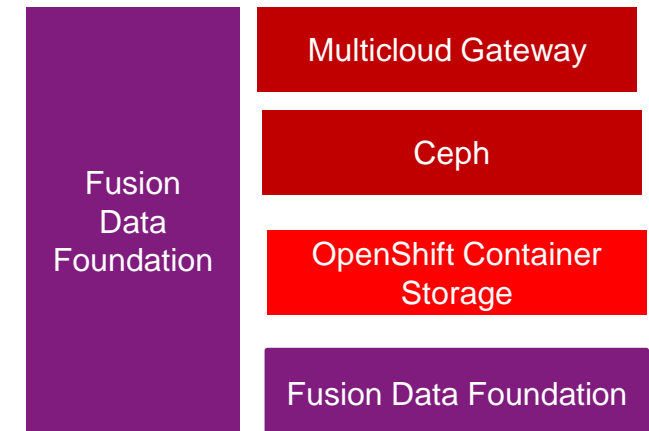
IBM Fusion Data Foundation - Architecture

FDf provides storage services largely by leveraging the following software packages:

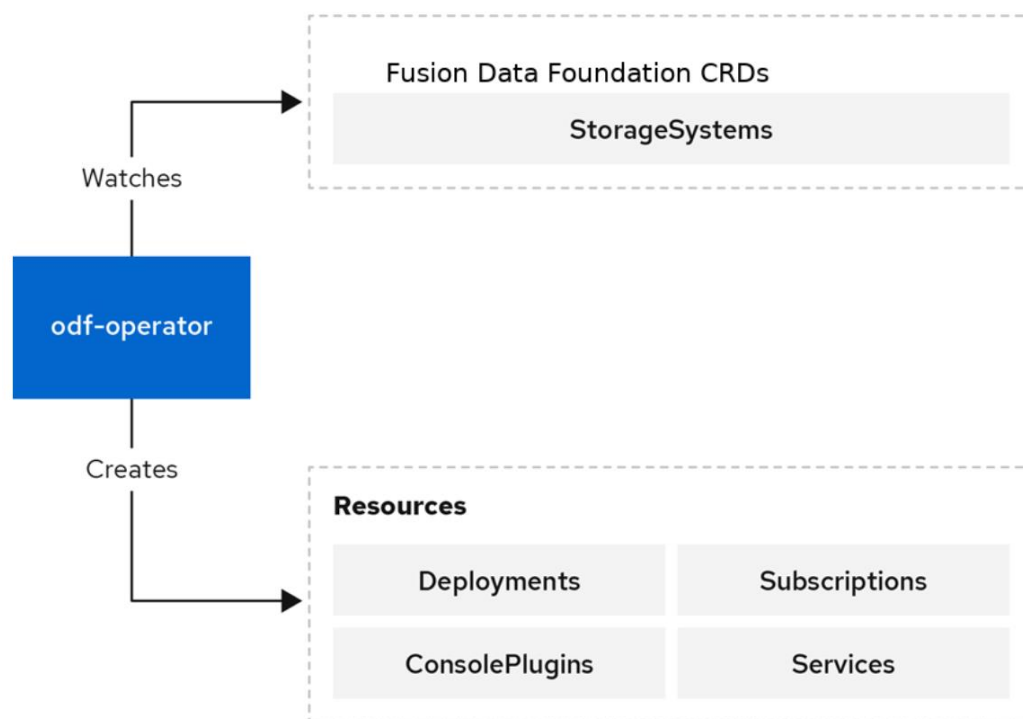
- Ceph – Provides block storage, a shared and distributed file system, and on-premises object storage
- Ceph CSI – Manage provisioning and lifecycle of persistent volumes and claims
- NooBaa - Provide a Multicloud Object Gateway (MCG)
- Fusion Data Foundation Operators –Initialize and manage Fusion Data Foundation services

FDf is comprised of three (3) Operator Lifecycle Manager (OLM) bundles that deploy four (4) operators:

- **Fusion Data Foundation**
 - Data Foundation (odf-operator)
- **Container Storage**
 - OpenShift Container Storage (ocs-operator)
 - Ceph (rook-ceph-operator)
- **Multicloud Gateway**
 - Multicloud Gateway (mcg-operator)

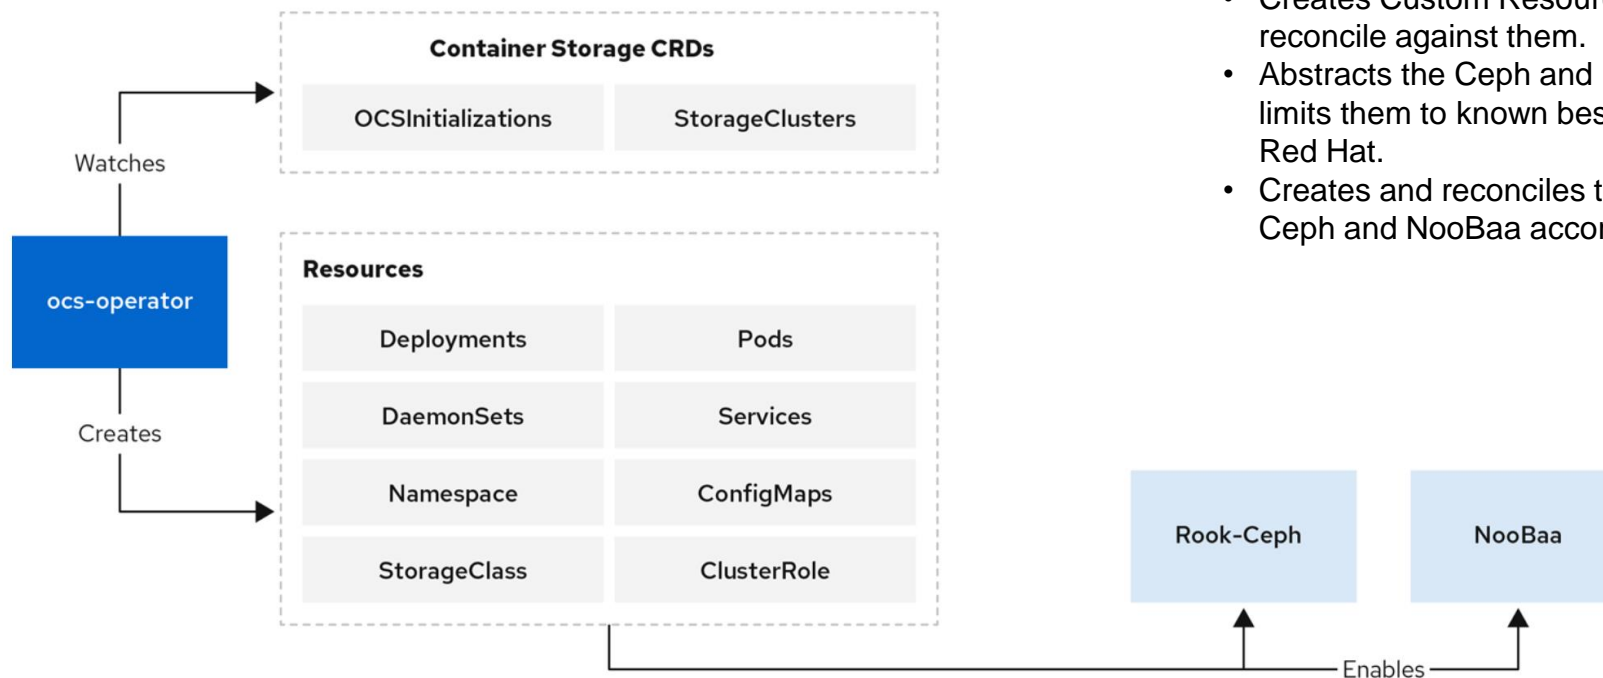


IBM Fusion Data Foundation – Operators - odf-operator



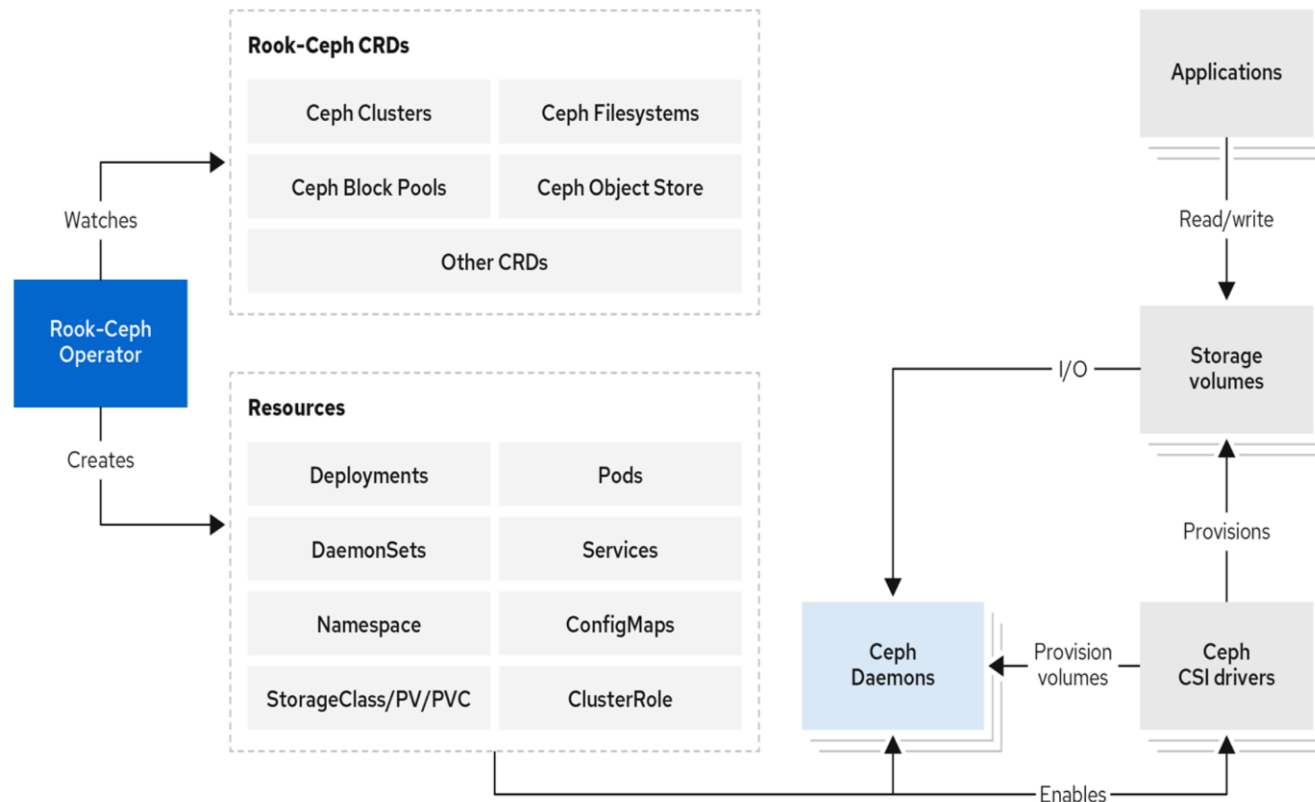
- Deployed from the Fusion UI
- Install the Fusion packages and Data Services on the OpenShift cluster
- "Meta" operator designed to influence other operators
- Provides the following services:
 - Enforces the configuration and versioning of the other operators that comprise Fusion Data Foundation through operator dependencies and subscription management.
 - Provides the Fusion Data Foundation external plugin for the OpenShift Console.
 - Provides an API to integrate storage solutions with the OpenShift Console
- Has a dependency on the ocs-operator

IBM Fusion Data Foundation – Operators - ocs-operator



- "Meta" operator that serves as a configuration gateway
- Has the following primary functions:
 - Creates Custom Resources (CRs) that trigger the other operators to reconcile against them.
 - Abstracts the Ceph and Multicloud Object Gateway configurations and limits them to known best practices that are validated and supported by Red Hat.
 - Creates and reconciles the resources required to deploy containerized Ceph and NooBaa according to the support policies.

IBM Fusion Data Foundation – Operators – rook-ceph-operator



Manages the core Data Services components as part of the deployment:

- **Ceph-CSI Driver** - Creates and updates the CSI driver, including a provisioner for the RADOS block device (RBD) and Ceph filesystem (CephFS) as well as the volume plugin daemonset for each of the two drivers.
- **Ceph daemons**
 - MON (Monitors) - Provide the core metadata store for Ceph.
 - OSD (Object Storage Daemons) - Store the data on underlying devices.
 - MGR (Manager) - Collects metrics and provides other internal functions for Ceph.
 - RGW (RADOS Gateway) - Provides the S3 endpoint to the object store
 - MDS (Metadata Server) - Provides CephFS shared volumes.

Leveraging Ceph from IBM Storage Fusion

```
[root@bootnode ~]# oc get pod -n openshift-storage
NAME                                                    READY   STATUS    RESTARTS   AGE
csi-addons-controller-manager-68b84b665c-thjtc         2/2     Running   0           11d
csi-cephfspplugin-c6t2f                                2/2     Running   2           4d16h
csi-cephfspplugin-h6hn4                                2/2     Running   0           4d16h
csi-cephfspplugin-j9zjh                                2/2     Running   0           4d16h
csi-cephfspplugin-pjkq2                                2/2     Running   0           11d
csi-cephfspplugin-provisioner-77bdf4cbc9-2qgwp        5/5     Running   0           11d
csi-cephfspplugin-provisioner-77bdf4cbc9-z7ckj        5/5     Running   0           11d
csi-cephfspplugin-q777p                                2/2     Running   0           11d
csi-cephfspplugin-wv525                                2/2     Running   0           11d
csi-rbdplugin-gc7d5                                    3/3     Running   0           11d
csi-rbdplugin-l9wvv                                    3/3     Running   0           11d
csi-rbdplugin-provisioner-57f748bb78-4w2t2            6/6     Running   0           11d
csi-rbdplugin-provisioner-57f748bb78-js85w            6/6     Running   0           11d
csi-rbdplugin-psnkq                                    3/3     Running   0           4d16h
csi-rbdplugin-svvgg7                                   3/3     Running   0           4d16h
csi-rbdplugin-tkltb                                    3/3     Running   3           4d16h
csi-rbdplugin-xztt9                                    3/3     Running   0           11d
noobaa-core-0                                          1/1     Running   0           11d
noobaa-db-pg-0                                         1/1     Running   0           11d
noobaa-endpoint-c9d6fbf5f-gsqmk                       1/1     Running   0           11d
noobaa-operator-6cc6d548f5-z4mdg                      1/1     Running   10 (23h ago)  11d
ocs-metrics-exporter-5db796c5b4-b8n62                 1/1     Running   0           11d
ocs-operator-78b5974554-2mx7r                         1/1     Running   2 (23h ago)   11d
odf-console-6775954ff4-jrpfw                          1/1     Running   0           11d
odf-operator-controller-manager-5dccc79d7-9rj49       2/2     Running   2 (23h ago)   11d
rook-ceph-crashcollector-demoocp80-1777b-worker-5v95f-8495922h2 1/1     Running   0           11d
rook-ceph-crashcollector-demoocp80-1777b-worker-6p994-74f6mmww4 1/1     Running   0           11d
rook-ceph-crashcollector-demoocp80-1777b-worker-zqvps-5cdcrqgn7 1/1     Running   0           11d
rook-ceph-mds-ocs-storagecluster-cephfilesystem-a-5fd8b68415vlf 2/2     Running   0           11d
rook-ceph-mds-ocs-storagecluster-cephfilesystem-b-dc7d5d69sgblt 2/2     Running   0           11d
rook-ceph-mgr-a-89dfc7c79-2512h                      2/2     Running   0           11d
rook-ceph-mon-a-6b5f89cdb7-sgp46                     2/2     Running   0           11d
rook-ceph-mon-b-5b5b476df7-p289q                    2/2     Running   0           11d
rook-ceph-mon-c-74b845d689-77lrv                    2/2     Running   0           11d
rook-ceph-operator-79cb79795b-dmv7s                  1/1     Running   0           11d
rook-ceph-osd-0-598b979f89-2c7lz                     2/2     Running   0           11d
rook-ceph-osd-1-64799dbf8f-p2jgn                     2/2     Running   0           11d
rook-ceph-osd-2-7fd665c46b-9zmbc                    2/2     Running   0           11d
rook-ceph-osd-prepare-fusion-storage-0-data-0xfmcj-rfcs5 0/1     Completed 0           39d
rook-ceph-osd-prepare-fusion-storage-1-data-0dmkvm-gw249 0/1     Completed 0           39d
rook-ceph-rgw-ocs-storagecluster-cephobjectstore-a-599bbd54cjnt 2/2     Running   0           11d
[root@bootnode ~]#
```

After FDF has been deployed, a quick validation check will reveal all installed operators:

OpenShift Container Storage operator

OpenShift Data Foundation operator

Ceph operator

Leveraging Ceph from IBM Storage Fusion

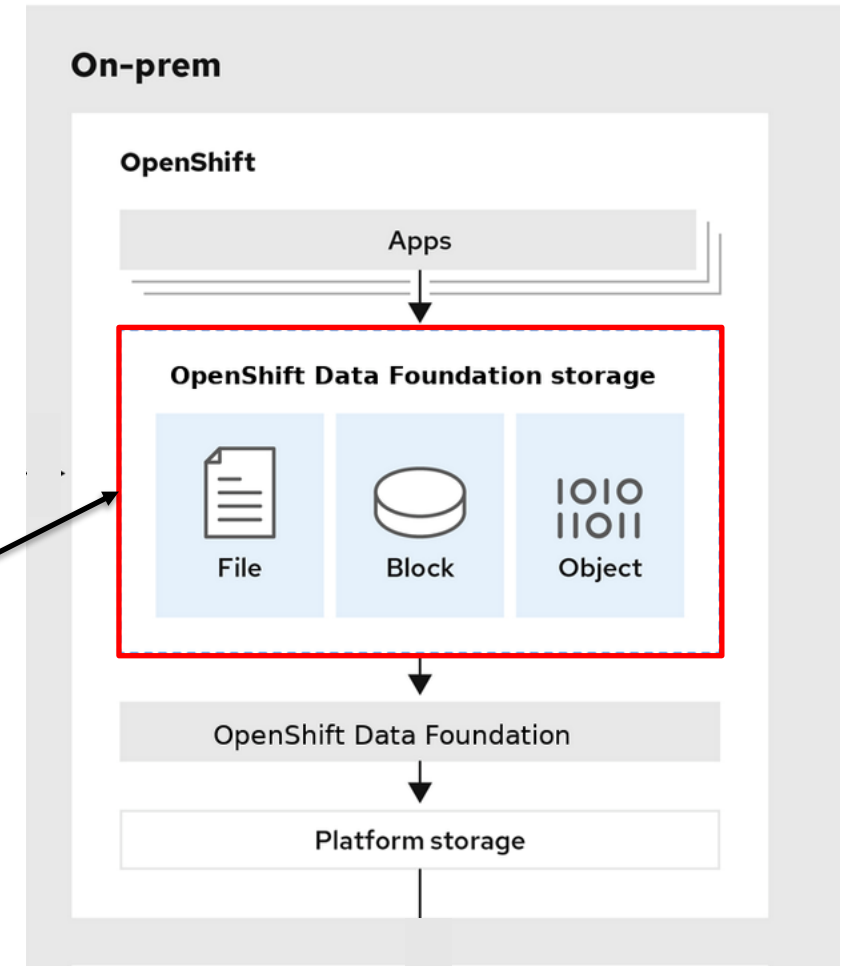
Once the Fusion Data Foundation operator has been deployed, Ceph will run “under the covers” to:

- Expose physical storage to applications that are deployed on Red Hat OpenShift
- Abstract the storage using container storage interfaces (CSI) and Storage Classes
- Deliver consistent volume types, resilience, and tiering wherever the storage layer is provisioned

So...where's Ceph?! It's in there!

But don't take MY word for it. Let's take a look:

Ceph



So Let's See Where Ceph Lives...

It's Walkthrough Time!

Standalone Ceph & Fusion Ceph – Available Offerings

IBM Storage Ceph and Red Hat Ceph packaging



IBM Storage Ceph

On-prem S3 storage at scale and performance

- Object storage
- Block storage
- File storage
- Presence at the on-prem object market at 10-Petabyte+ scale
- S3 compatibility with AWS



Red Hat
OpenStack
Platform



Red Hat
Ceph Storage

Ceph for OpenStack

1 in OpenStack storage

- Cinder block storage
- Nova ephemeral storage
- Glance image storage
- Swift object store
- Manila file storage
- Advanced integration
- Unified management
- Hyperconverged and Edge capabilities



IBM Storage Fusion

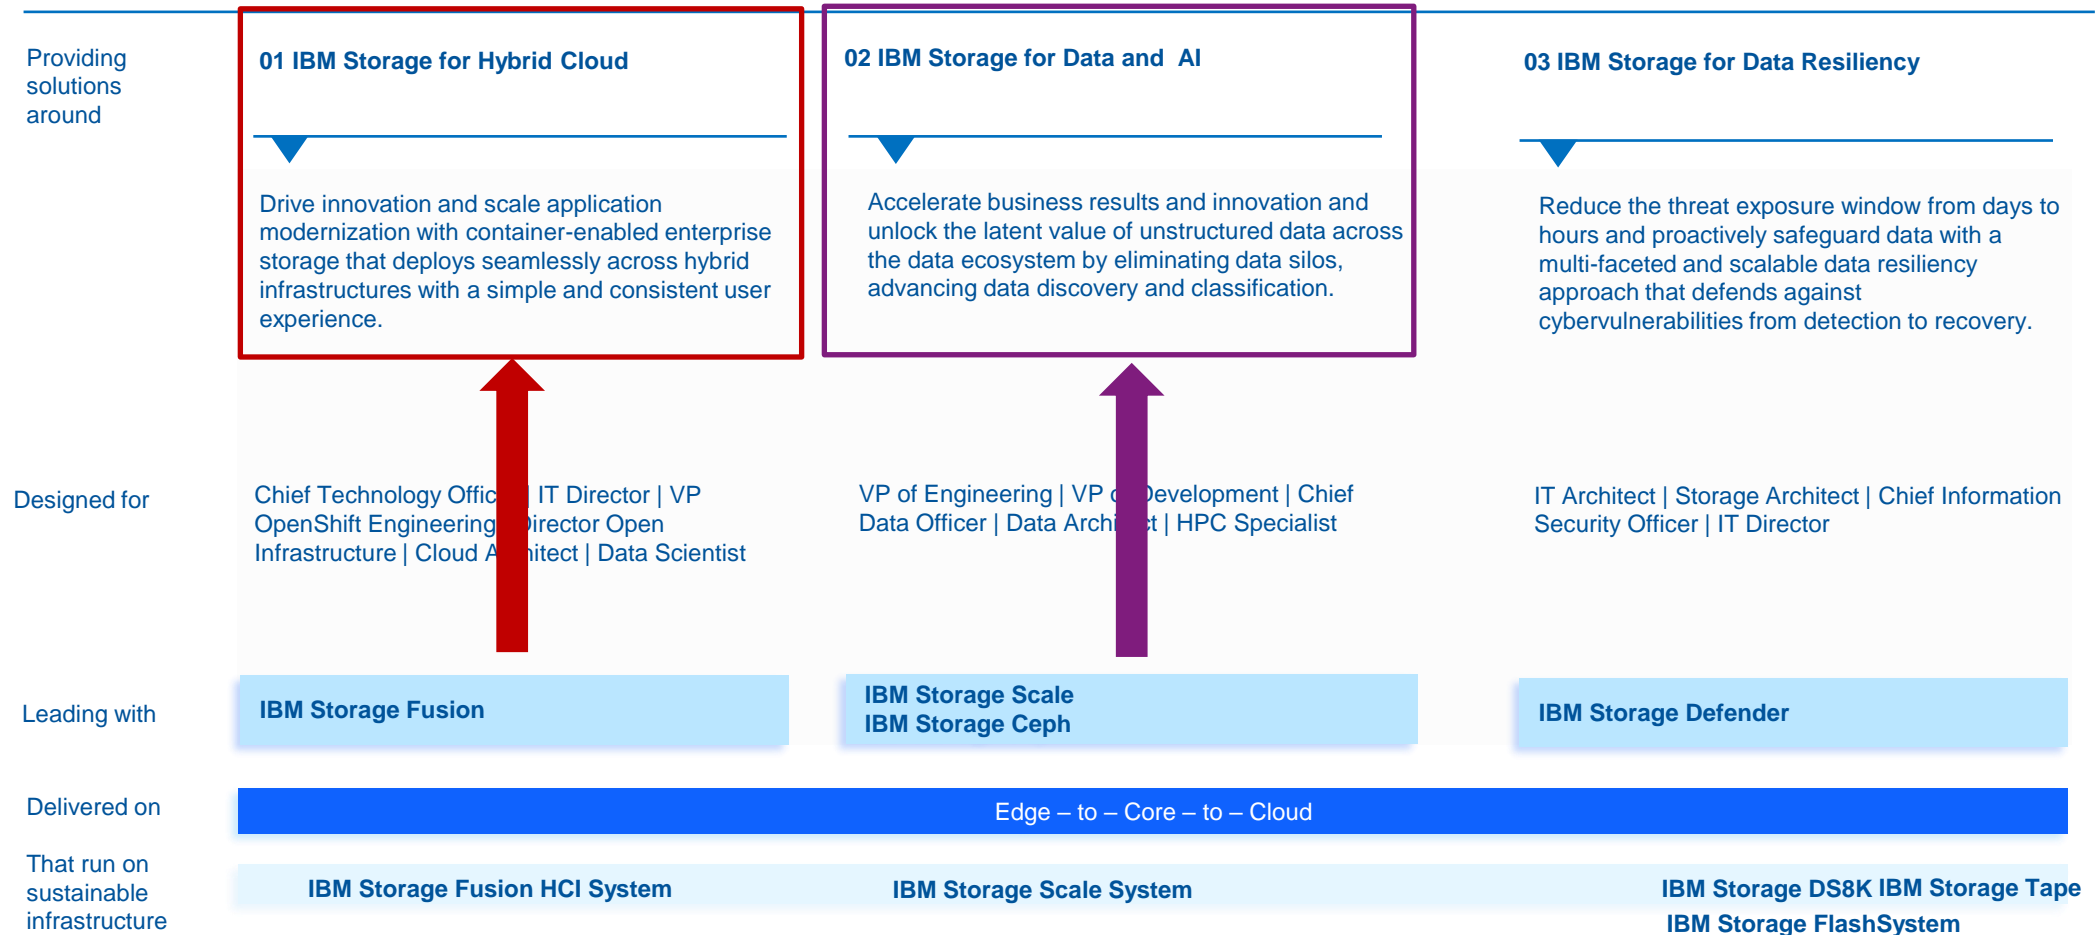
Ceph for OpenShift

- Self-managing storage powered by Red Hat Ceph Storage
- Automated by Rook and completed with Multicloud object gateway (MCG)
- Advanced integration, automation, ease of use
- Persistent storage for OpenShift stateful workloads

Standalone Ceph & Fusion Ceph - A Philosophical Discussion...

So when should I use IBM Storage Ceph vs. IBM Fusion Ceph?

- IBM Storage Fusion – Hybrid cloud, Application Modernization, OpenShift, etc...
- IBM Storage Ceph – Big Data, AI, Data Lakehouse, etc...



Thank you!

Accelerate with ATG Technical Webinar Series - Survey

Please take a moment to share your feedback with our team!

You can access this 6-question survey via [Menti.com](https://www.menti.com) with code 2243 3599 or

Direct link <https://www.menti.com/albneqj15g57>

Or

QR Code

