Lightning Talks I

IBM Storage Scale Days 2024

March 5-7, 2024 | Stuttgart Marriott Hotel Sindelfingen

Disclaimer

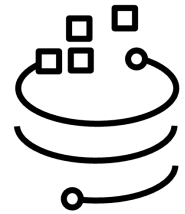

IBM's statements regarding its plans, directions, and intent are subject to change or withdrawal without notice at IBM's sole discretion. Information regarding potential future products is intended to outline our general product direction and it should not be relied on in making a purchasing decision. 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 Scale Days 2024

Lightning Talks I

- Storage Scale (System) Support
- Storage Scale System 6000 – Expert Care Premium
- IBM Technology Expert Labs – A short introduction
- Writeable snapshots with IBM Storage Scale
- Dynamic Infrastructure for Spark in OpenShift
- Data Migration using OpenShift built-in tools
- Multi-Protocol data access in OpenShift

Storage Scale (System) Support

IBM Storage Scale Days 2024

March 5-7, 2024 | Stuttgart Marriott Hotel Sindelfingen

Dennis Kunkel

IBM Storage Scale Days 2024

Support reference guides

- Support Model
- Support Scope
- Contact information
- Escalation handling
- Useful links
- .....

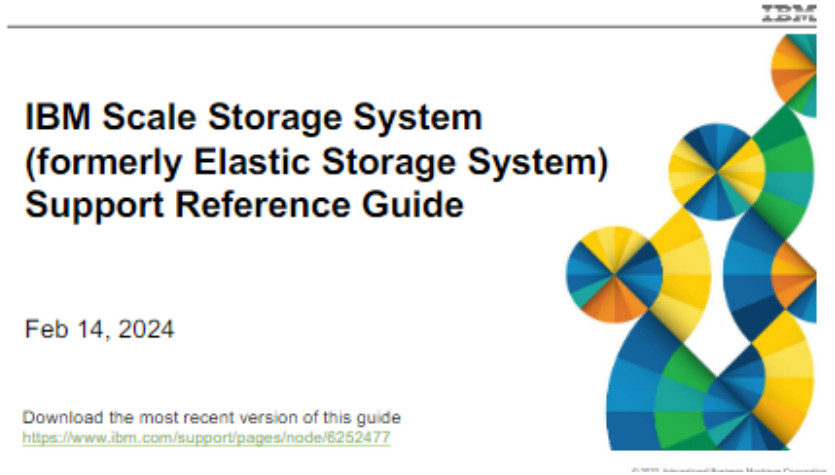

Storage Scale System

- https://www.ibm.com/support/pages/ibm-storage-scale-system-formerly-elastic-storage-system-or-ess-support-reference-guide

Storage Scale

- https://www.ibm.com/support/pages/node/6252403

Questions/Feedback? Don’t hesitate to contact us.

Dennis Kunkel

EMEA Segment Lead for Storage Scale and Storage Scale System

E-Mail: dennis.kunkel@de.ibm.com

Tel: +49-170-3387365

→ LinkedIn

Storage Scale System 6000

Expert Care Premium

IBM Storage Scale Days 2024

March 5-7, 2024 | Stuttgart Marriott Hotel Sindelfingen

Christian Dietrich

IBM Storage Scale Days 2024

Expert Care Premium – General Availability for Storage Scale System 6000 – Jan 23rd 2024

IBM Storage Scale System 6000	Basic	Advanced	Premium
IBM Hardware Maintenance: IBM provides 9x5 - Next Business Day Onsite Repair	√	-	-
IBM Hardware Maintenance IBM provides 24x7 - Same Day Onsite Repair	-	√	√
Predictive Support Predictive Support can be provided only for the Utility Node when it is attached/configured with same generation IBM Storage Scale System hardware.	-	√	√

IBM Storage Scale System 6000	Basic	Advanced	Premium
Dedicate Support from an IBM Expert Technical Account Manager (TAM)	-	-	√
Enhanced 30-minute response time For Severity 1 and Severity 2 Issues	-	-	√
Support Line - Embded with the hardware order of 5149-23E and 5149-F48 Support line for RedHat is outside the Expert Care bundle and is enabled with the hardware machine order of Utility Node (5148-23E) and the Scale System 6000(5149-F48)			

Source: https://www.ibm.com/downloads/cas/XJBDGENA

Christian Dietrich

EMEA Segment Lead for Expert Care
Premium Delivery

E-Mail: christian.dietrich@de.ibm.com

Tel: +49-151-1464 5118

→ LinkedIn

IBM Technology Expert Labs

A short introduction

IBM Storage Scale Days 2024

March 5-7, 2024 | Stuttgart Marriott Hotel Sindelfingen

Fabián Kuhl
Bernd Bäuml

IBM Storage Scale Days 2024

IBM Technology Expert Labs

- ✓ Professional services team that helps clients and business partners to implement and optimize infrastructure.
- ✓ Delivery consultants with deep technical expertise and extensive project experience.

Information about our service:
https://www.ibm.com/services/infrastructure

Contacts

Hilmar Rörig
Manager Storage Services
IBM Technology Expert Labs
Tel.: 0160-9691468
Mail: hroerig@de.ibm.com

Frank Bröde
Manager Storage Services
Brandleader Infrastructure
IBM Technology Expert Labs
Tel.: 0172-7315935
Mail: broede@de.ibm.com

Andre Gaschler
Solution Architect
IBM Technology Expert Labs
Mail: gaschler@de.ibm.com

Wilfried Ochs
Sales Expert Infrastructure Services
IBM Technology Expert Labs
Tel.: 0171-8608694
Mail: wochs@de.ibm.com

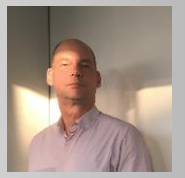

Bernd Baeuml
Storage Consultant
IBM Technology Expert Labs
Mail: bernd.baeuml@de.ibm.com

Jaroslaw Szymczak
Storage Consutant
IBM Technology Expert Labs
Mail: JJS@ibm.com

Fabián Kuhl
Storage Consultant
IBM Technology Expert Labs
Mail: fabikuhl@de.ibm.com

IBM Storage Scale Days 2024

IBM Technology Expert Labs – Service overview

Integration & upgrade services (Deploy/Support)

- DS8000 Logical Configuration Service
- Elastic Storage Server (ESS) Implementation
- IBM Storage Scale Implementation
- IBM Storage Fusion Implementation (HCI & SDS)
- IBM Storage Discover Implementation
- FlashSystem Implementation
- CSM Integration for z/OS or Distributed Platforms
- IBM Storage Protect/Protect Plus Implementation
- IBM Storage Protect Snapshot (FlashCopy Manager)
- IBM Storage Virtualize (SVC) Advanced Integration

Tailored Services (customized)

Customized TAW - Technology Adoption Workshops

- Flexible and customized
- Various levels of skill transfer
- Hands-on

Security Services (Deploy/Optimize)

- Cyber Incident Response Storage Assessment (CIRSA)
- Cyber Vault for IBM Storage Architecture Workshop
- Encryption Implementation (TKLM/SKLM/GKLM)

Health Checks (Optimize)

- IBM Storage Scale Health Check
- Architectural Datacenter Check
- SAN Health Check
- IBM Storage Virtualize / FS9xxx Health Check
- DS8000 Health Check
- IBM Storage Protect Health Check
- TS7700 Health Check Service
- TS7700 Performance Check
- Health Check of TS4500 Libraries and Drives

IBM Storage Scale Days 2024

IBM Technology Expert Labs – Service overview

Migration/Disposal Services (Refresh)

- IBM Storage Scale based migration
- z/OS Data Migraton (TDMF)
- Copy Services based Data Migration
- High End Tape Migration (TS77xx, VTS)
- Enterprise Disk Secure Data Overwrite
- Enterprise Tape Secure Data Overwrite
- FlashSystem Secure Deletion Offering

Consulting Services (Optimize)

- Ansible Automation Service
- IBM Storage Scale Architecture and Consulting
- Elastic Storage Server Consulting Service
- CSM Consulting for z/OS and Distributed
- IBM Storage Protect/IBM Storage Protect Plus
- IBM Storage Protect Plus Consulting
- Cyber Resiliency Consulting Service

IBM Storage Scale Days 2024

IBM Technology Expert Labs – Service offering overview

Standard service offering

- 5-day project units
- Pre-priced
- Standardized content
- Onsite or remote

Expertise Connect

- Subscription-based
- Long-term trusted advisor relationship
- Dedicated expert, brings in highly specialized experts if needed
- Onsite or remote

Time and material

- Project based
- Customized, based on clients requirements
- Flexible - purchase project units at time of sale or during product lifecycle
- Onsite or remote

→ see https://www.ibm.com/support/pages/node/7015883

IBM Storage Scale Days 2024

IBM Technology Expert Labs – Link summary

Information about our service:
https://www.ibm.com/services/infrastructure

Information about our standard offerings:
https://www.ibm.com/support/pages/node/7015883

TEL Mail & Slack:
systems-expert-labs@ibm.com
#ask-systems-expert-labs

TEL Service on Seismic:
https://ibm.seismic.com/Link/Content/DCfQDHQq2qgcMGmBR6bdXW8dD82P

TEL Team (IBM internal):
https://w3.ibm.com/w3publisher/ibm-technology-services/contacts/sales-team

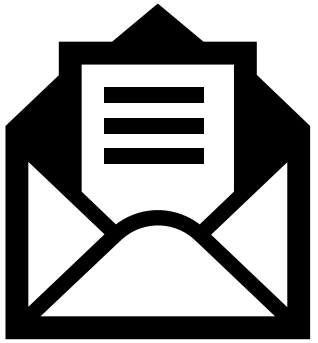

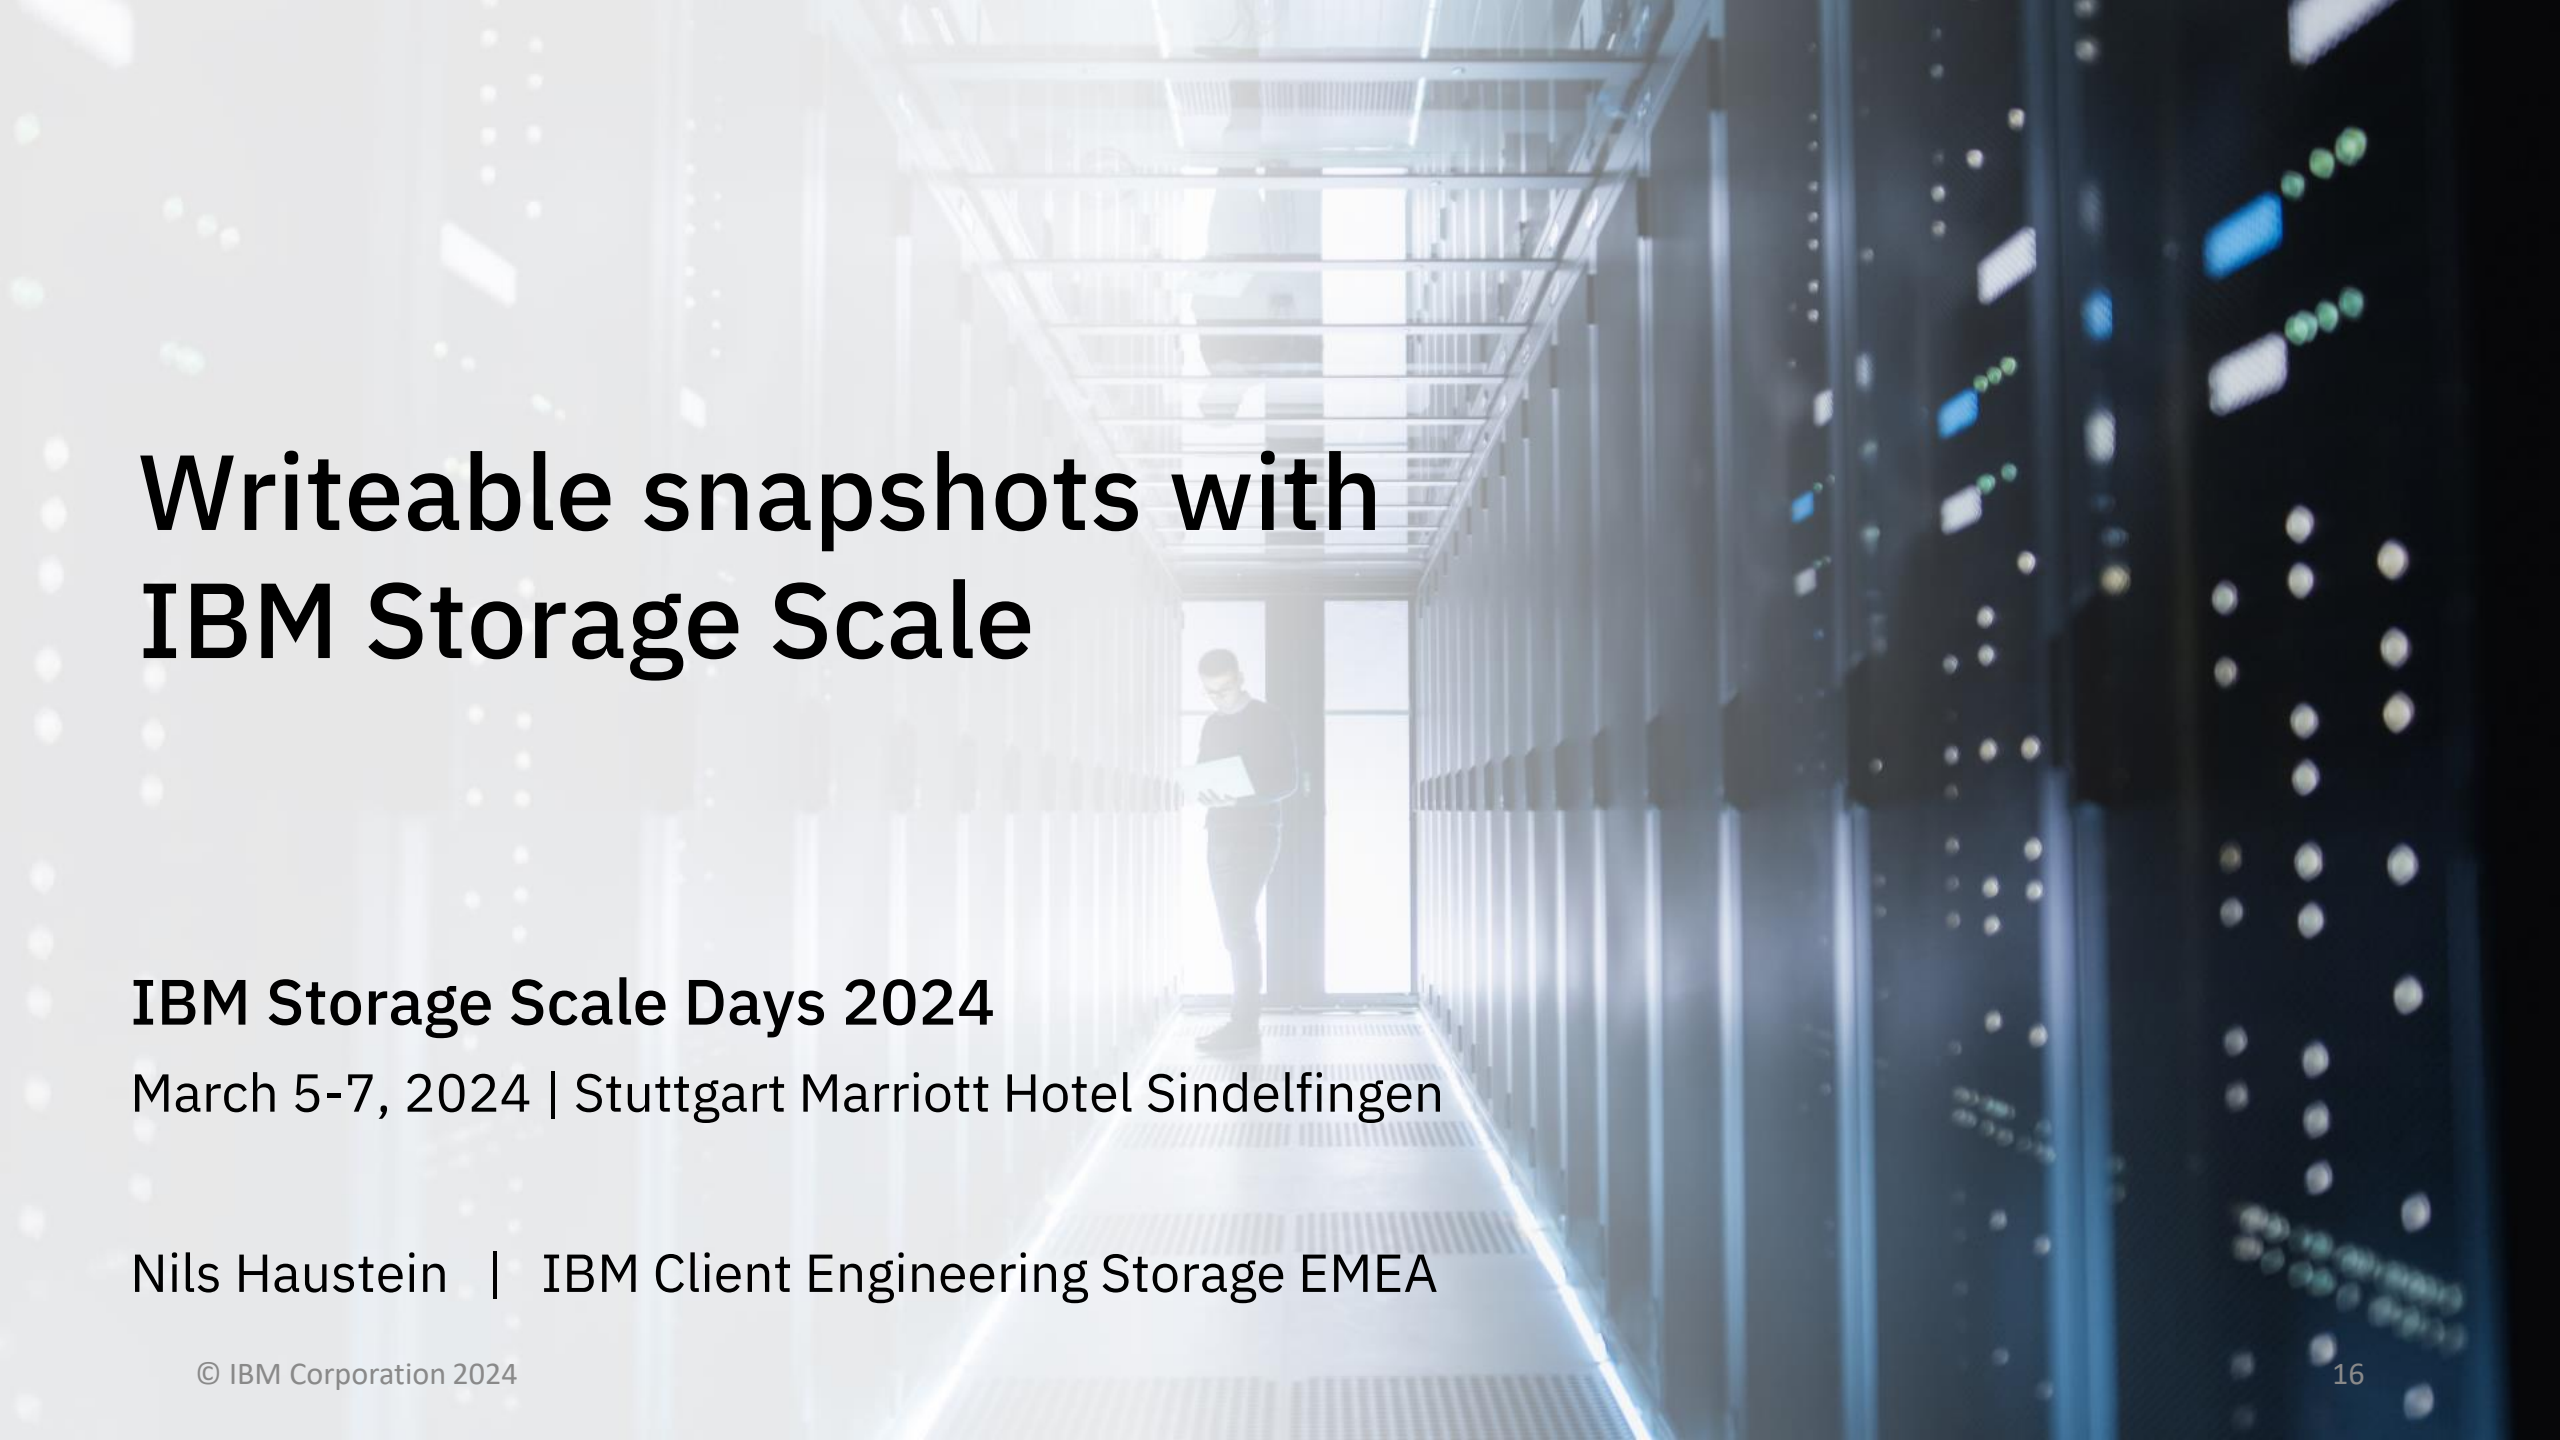

Writeable snapshots with IBM Storage Scale

IBM Storage Scale Days 2024

March 5-7, 2024 | Stuttgart Marriott Hotel Sindelfingen

Nils Haustein | IBM Client Engineering Storage EMEA

© IBM Corporation 2024

Let’s create ↷ **value** together

Client Engineering is an investment by IBM to jointly innovate and rapidly prove solutions to your business opportunities by leveraging IBM hybrid cloud and AI technologies.

Investment in your success

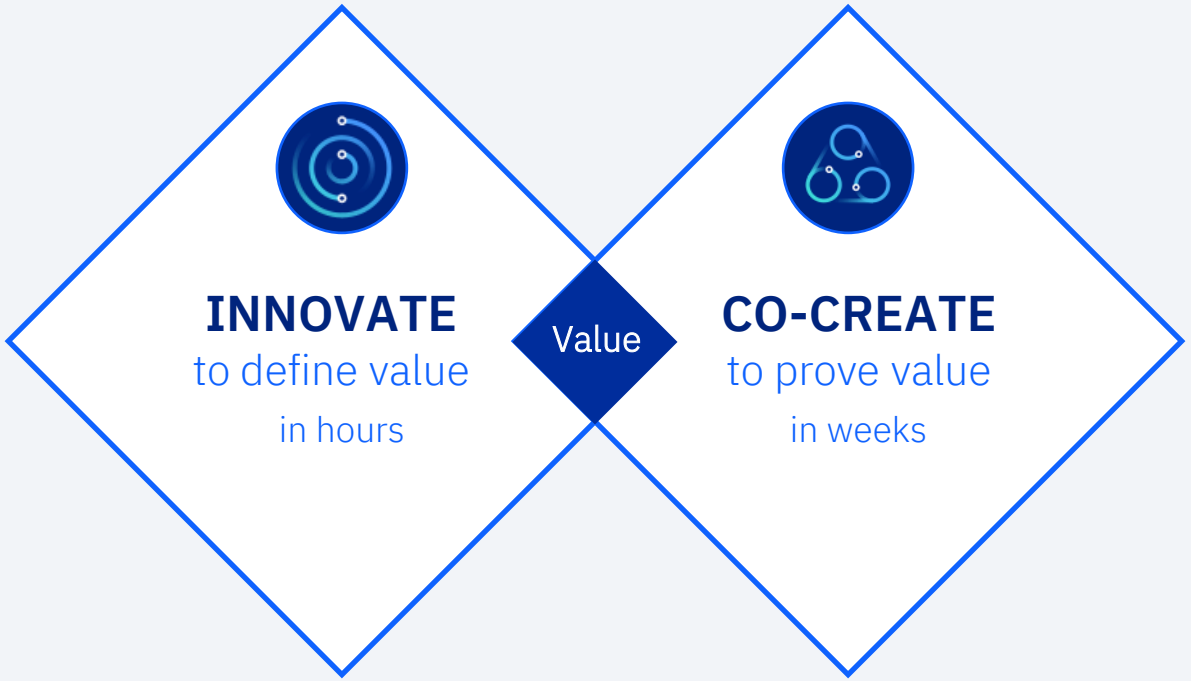

You bring your business and technology context, sponsorship, subject matter experts, and data. We bring a deeply skilled multi-disciplinary squad, technical accelerators, proven method, and a memorable experience.

Snapshots

- Snapshots preserve the content of a file system or independent fileset at a single point in time
  - Snapshots are space-efficient upon creation, only include block pointers
  - Snapshots are read-only
- Writeable snapshots allow reading and writing data from snapshots without changing the original files
- IBM Storage Scale can provide writeable snapshots by leveraging two techniques:
  - Snapshots
  - Active File Management Local update mode

© IBM Corporation 2024

Storage Scale Active File Management LU Mode

- AFM is based on Home – Cache relation
  - Home stores the files,
  - Cache shows files available in home
  - AFM cache mounts home via NFS or GPFS (remote mount)
- After AFM cache fileset is created file metadata is visible in cache
- File data is copied (fetched) from home upon file access or prefetched
- Created, changed and deleted files on cache are not replicated to home

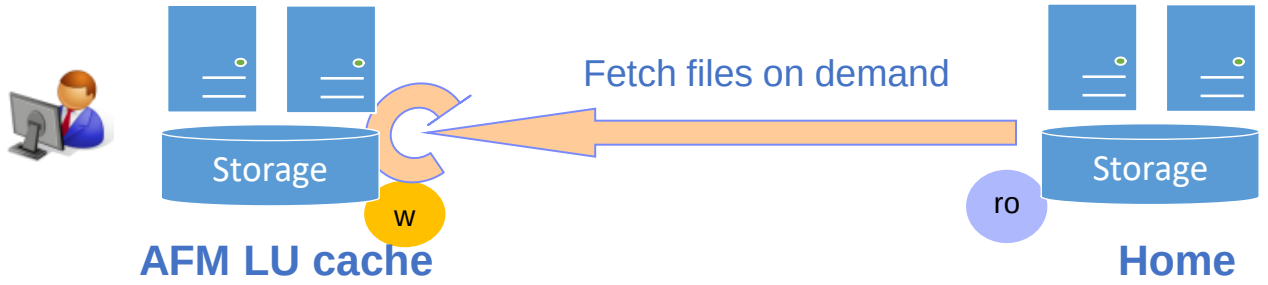

© IBM Corporation 2024

Writable snapshot with IBM Storage Scale

- Home workload uses home file system
  - Snapshot is created in home file system
- AFM LU cache mounts snapshot
  - Files in snapshot are visible in AFM LU fileset
  - Files are fetched from home snapshot to AFM LU fileset on access
  - Files created, changed or deleted in AFM LU fileset remain in AFM LU fileset → writable snapshot
- Cache workload uses AFM LU fileset
  - AFM LU fileset can be in the same home file system or in separate file system (recommended)
  - AFM LU fileset can be in the same home cluster or in a different cluster

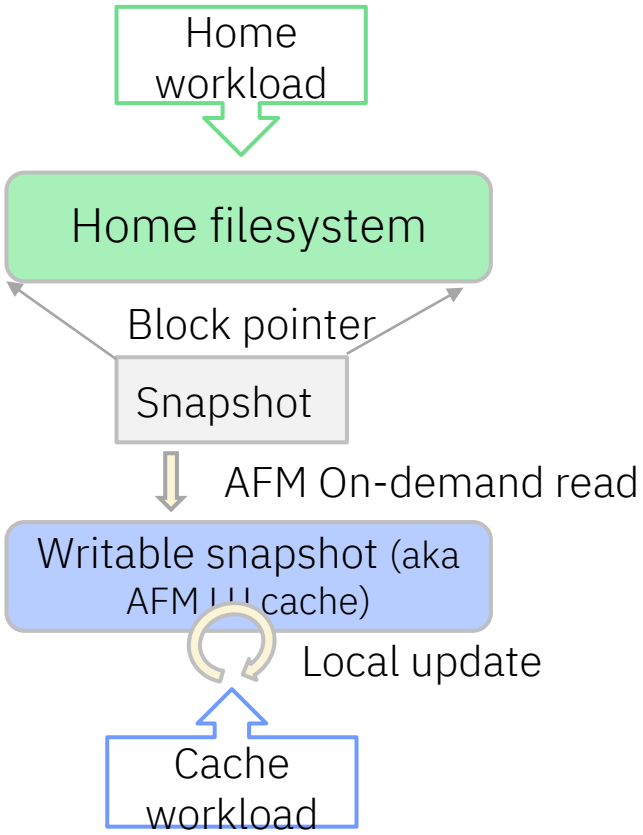

© IBM Corporation 2024

Use case 1: AFM DR testing

- AFM DR allows asynchronous replication from primary to secondary cluster for disaster recovery
- Writable snapshots can be used for DR testing:
  - Writable snapshot is created on secondary system based on peer snapshots
- Characteristics
  - No impact on primary system
  - No impact on data stored on primary or secondary AFM DR fileset
  - Peer snapshots allow application consistence

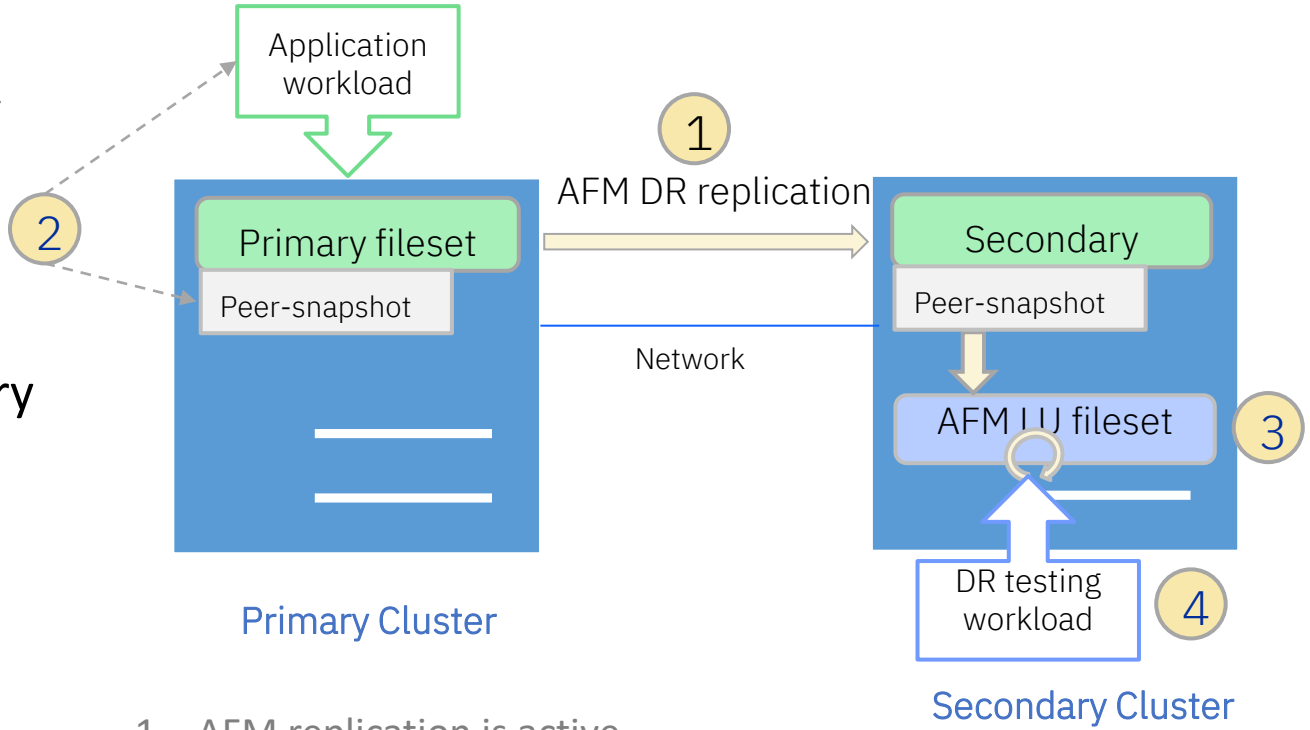

1 – AFM replication is active
2 – Create (consistent) peer snapshot on primary
3 – Create AFM LU fileset on secondary pointing to snapshot
4 – Execute DR testing

© IBM Corporation 2024

Use case 2: Database cloning

- Original Db can be cloned and used as clone Db
- Writable snapshot can be used by clone Db
  - In the same cluster or different cluster
  - Via remote NFS or GPFS mount
- Characteristics
  - No impact on data stored on original Db file system
  - Peer snapshots allow application consistence
  - Some impacts on original Db file system due to reads from snapshot

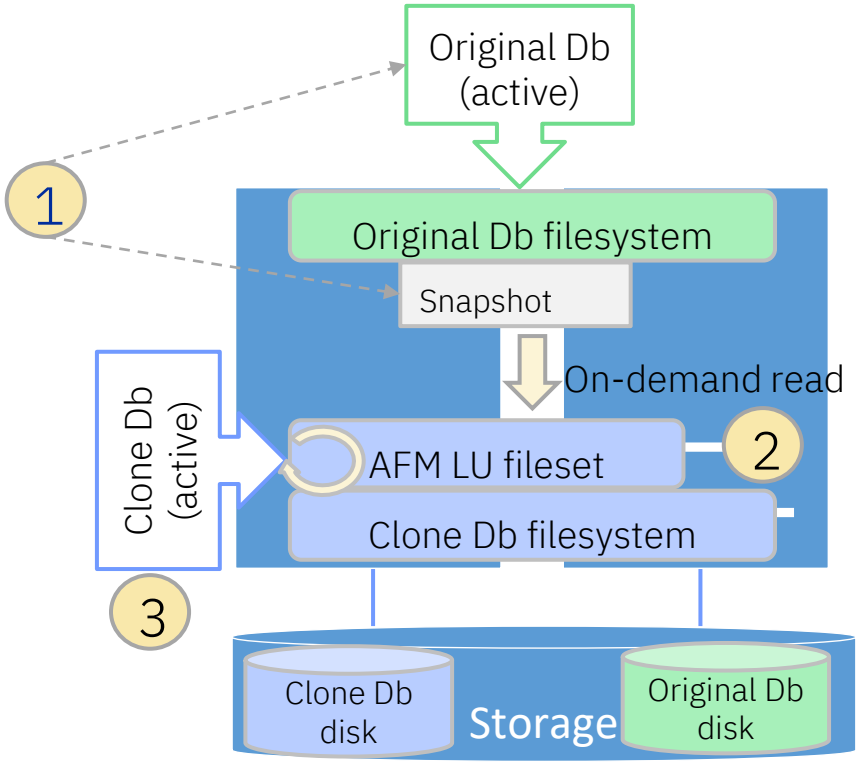

1 – Create (consistent) peer snapshot in original Db file system
2 – Create AFM LU fileset pointing to snapshot
3 – Use clone Db on writable snapshot

© IBM Corporation 2024

Considerations

- After creating snapshots in file system that is used by a database the file system workload increases temporarily
  - Database changes small blocks that must be copied on write
- Reading in AFM LU fileset increases workload in home file system containing the snapshot
  - As more data was read into LU fileset, workload in home file system decreases
- Resource consumption (CPU, memory) increases when reading from snapshot
- Snapshots consume capacity when data is changed in the original file system
  - In addition to this comes the capacity consumed in the AFM LU fileset due to new and changed files

© IBM Corporation 2024

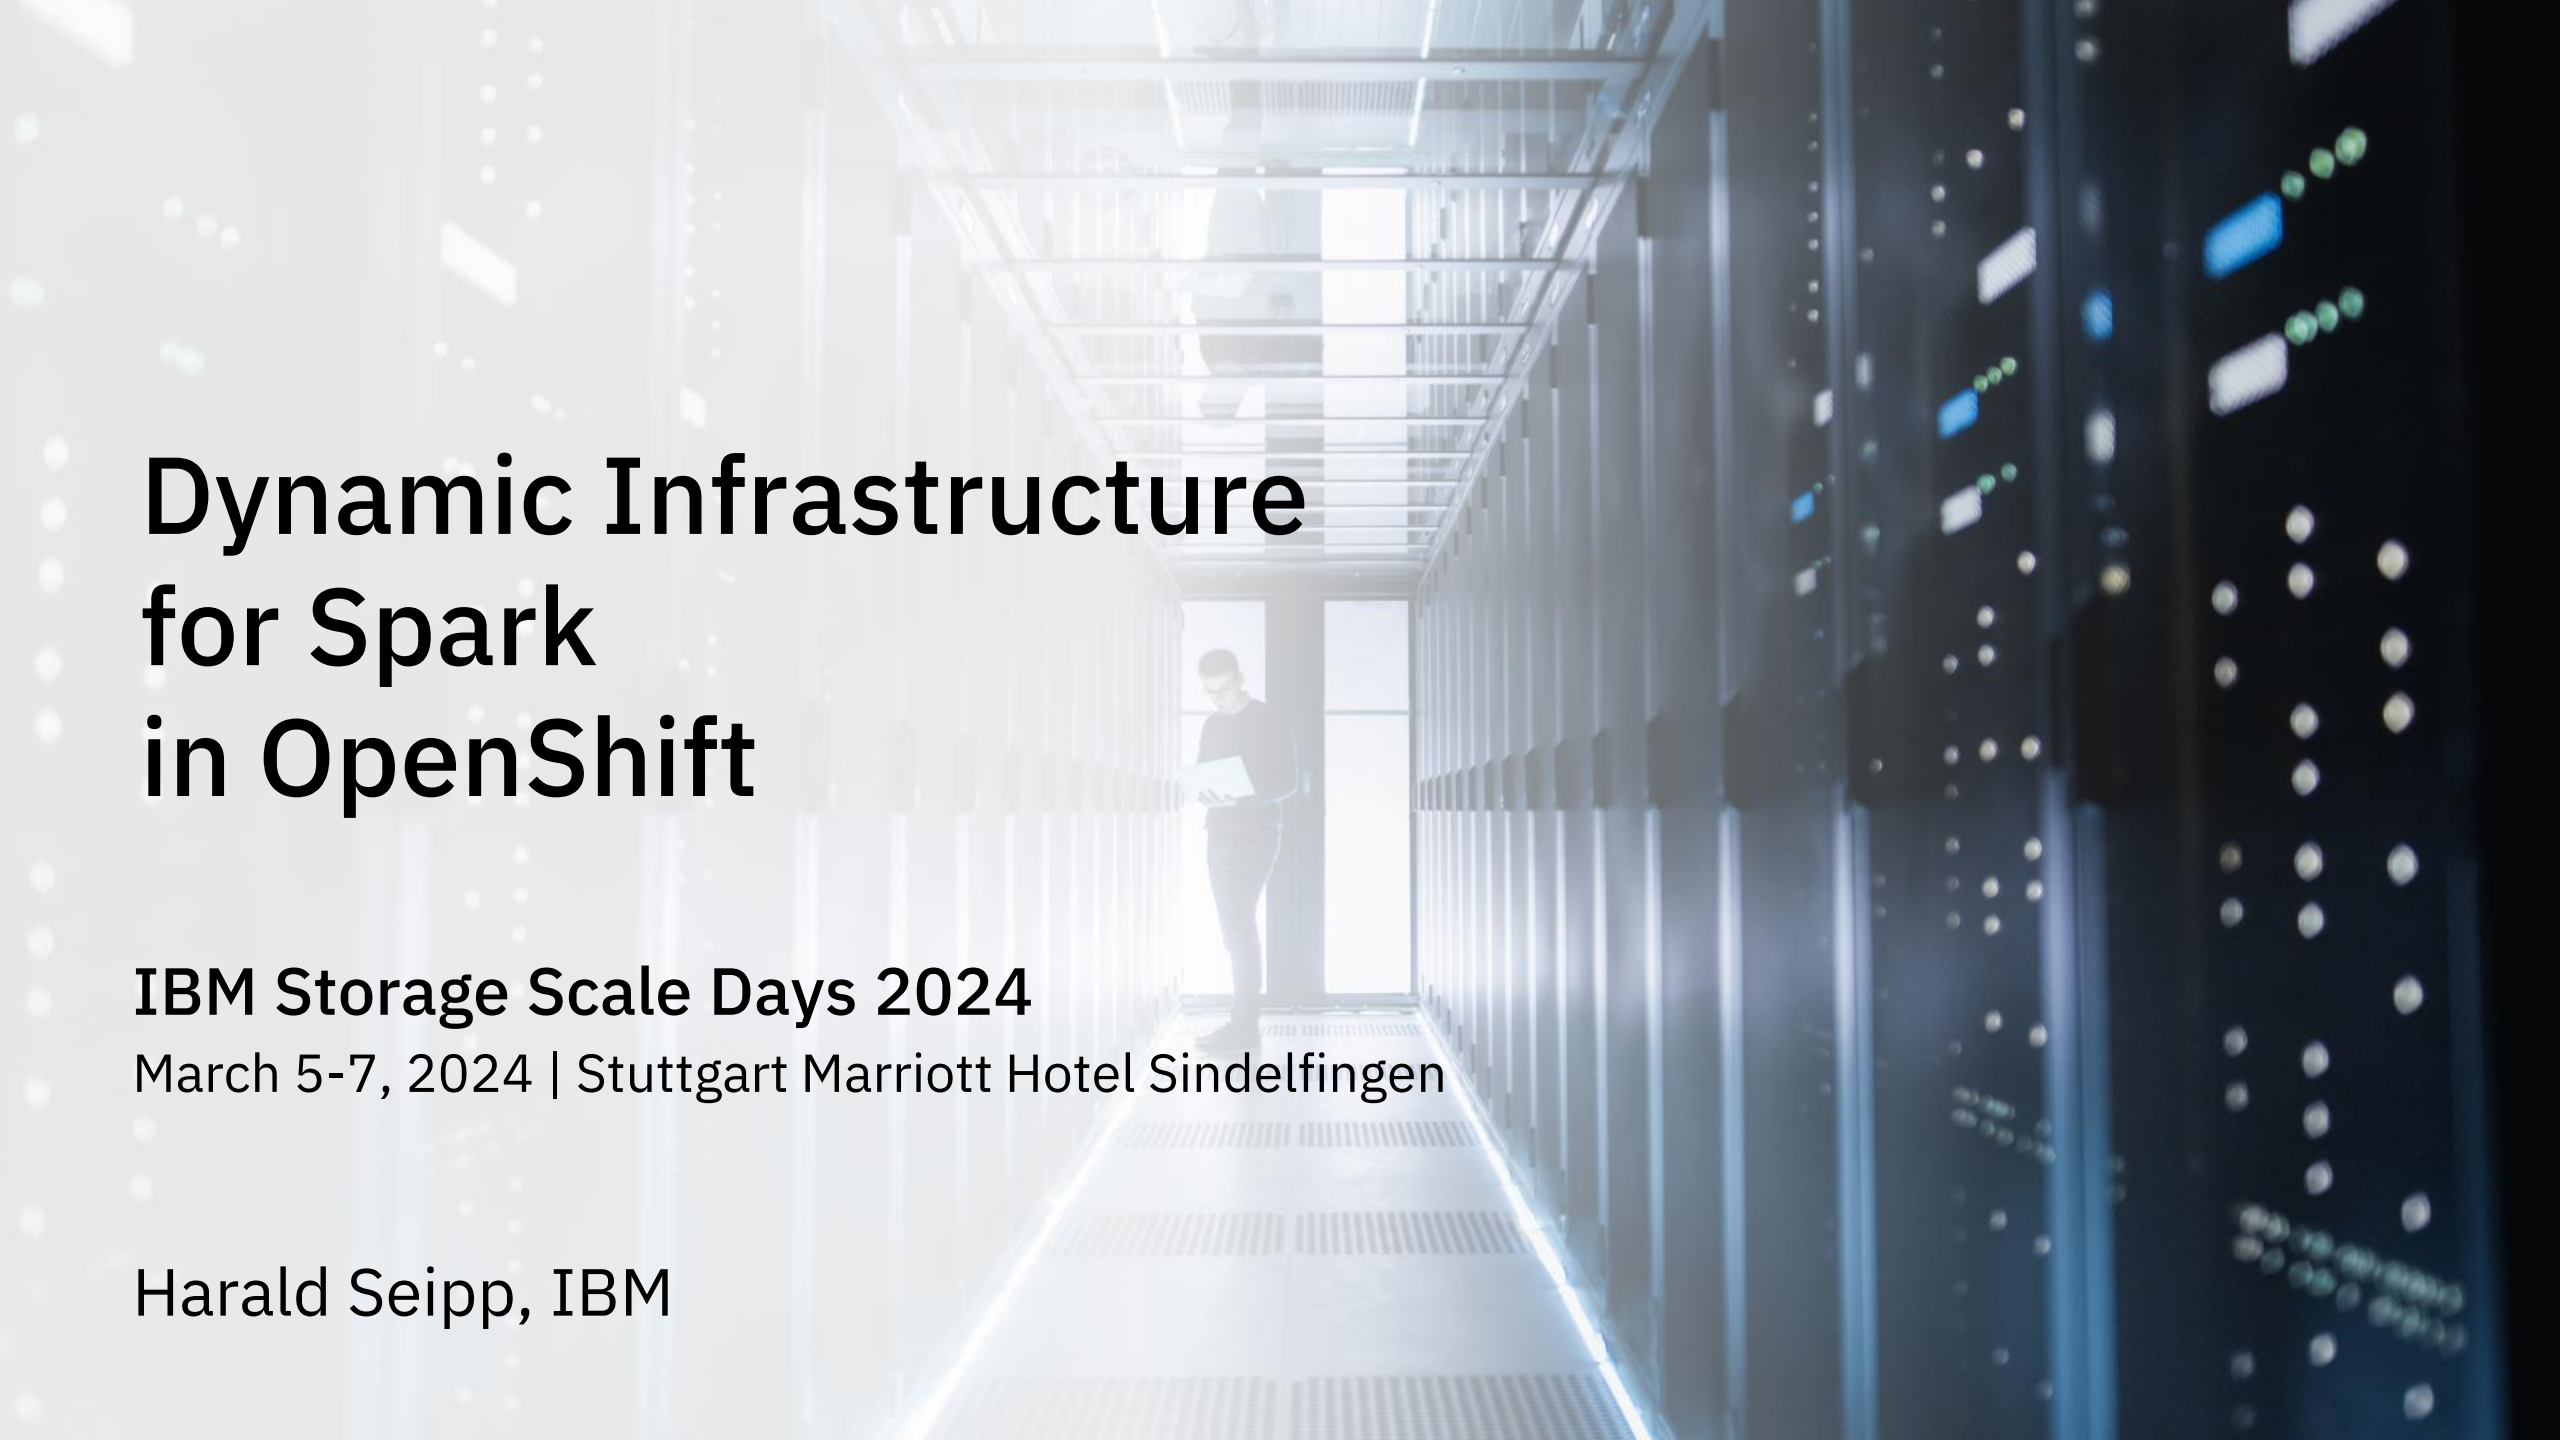

Dynamic Infrastructure for Spark in OpenShift

IBM Storage Scale Days 2024

March 5-7, 2024 | Stuttgart Marriott Hotel Sindelfingen

Harald Seipp, IBM

Key results & Learnings

Operation on OpenShift is feasible and can be realized with comparable hardware equipment as for legacy implementation.

Environments used

- OpenShift Lab Cluster (LAB) with Fusion Data Foundation (FDF)
  Slow performance expected due to not fulfilling OCP/FDF minimum prereqs.
- Hadoop Dev Cluster (HAD)
  Benchmark baseline for all job runs.
- OpenShift Cluster with IBM Storage Scale (POK), IBM Lab
  Executed comparative measurements
- OpenShift Cluster (MOP) with FDF, IBM Lab
  Comparative measurements

Runtime comparison

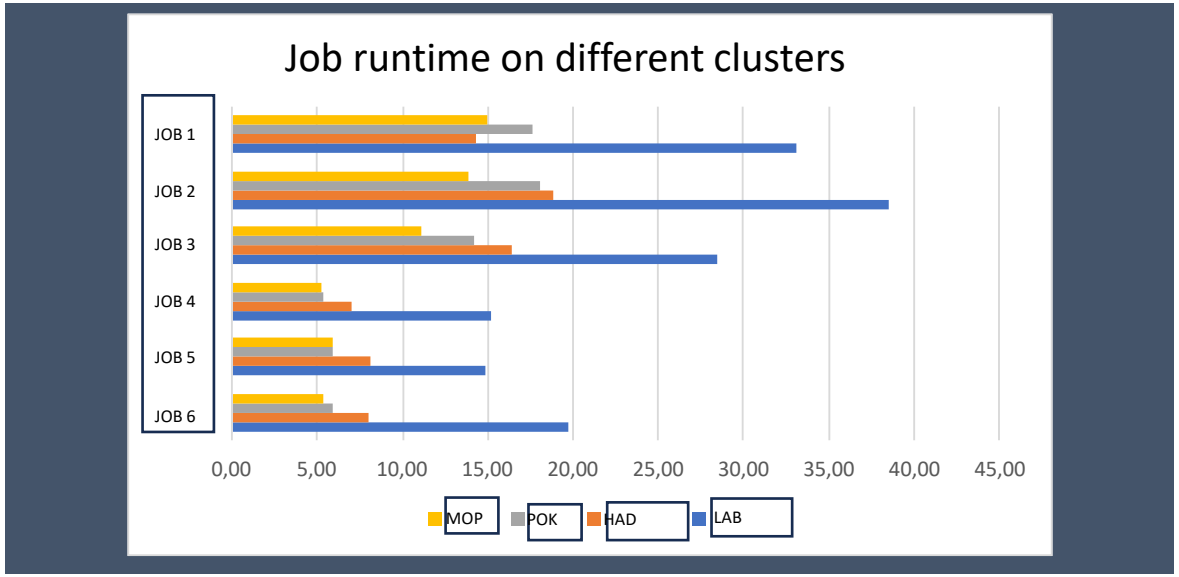

Learnings

- Migration of Spark Jobs could be performed without changing the software; only config changes required.
- FDF/CephFS with OpenShift is a working and performing alternative compared to HDFS with Hadoop.
- A performant network (10 Gbit) for the OCP Cluster is key for quick runtimes. With those, runtimes comparable to the Hadoop Cluster could be achieved.
- The new technology could be quickly adopted by the client team so that they could deploy it successfully during the MVP..

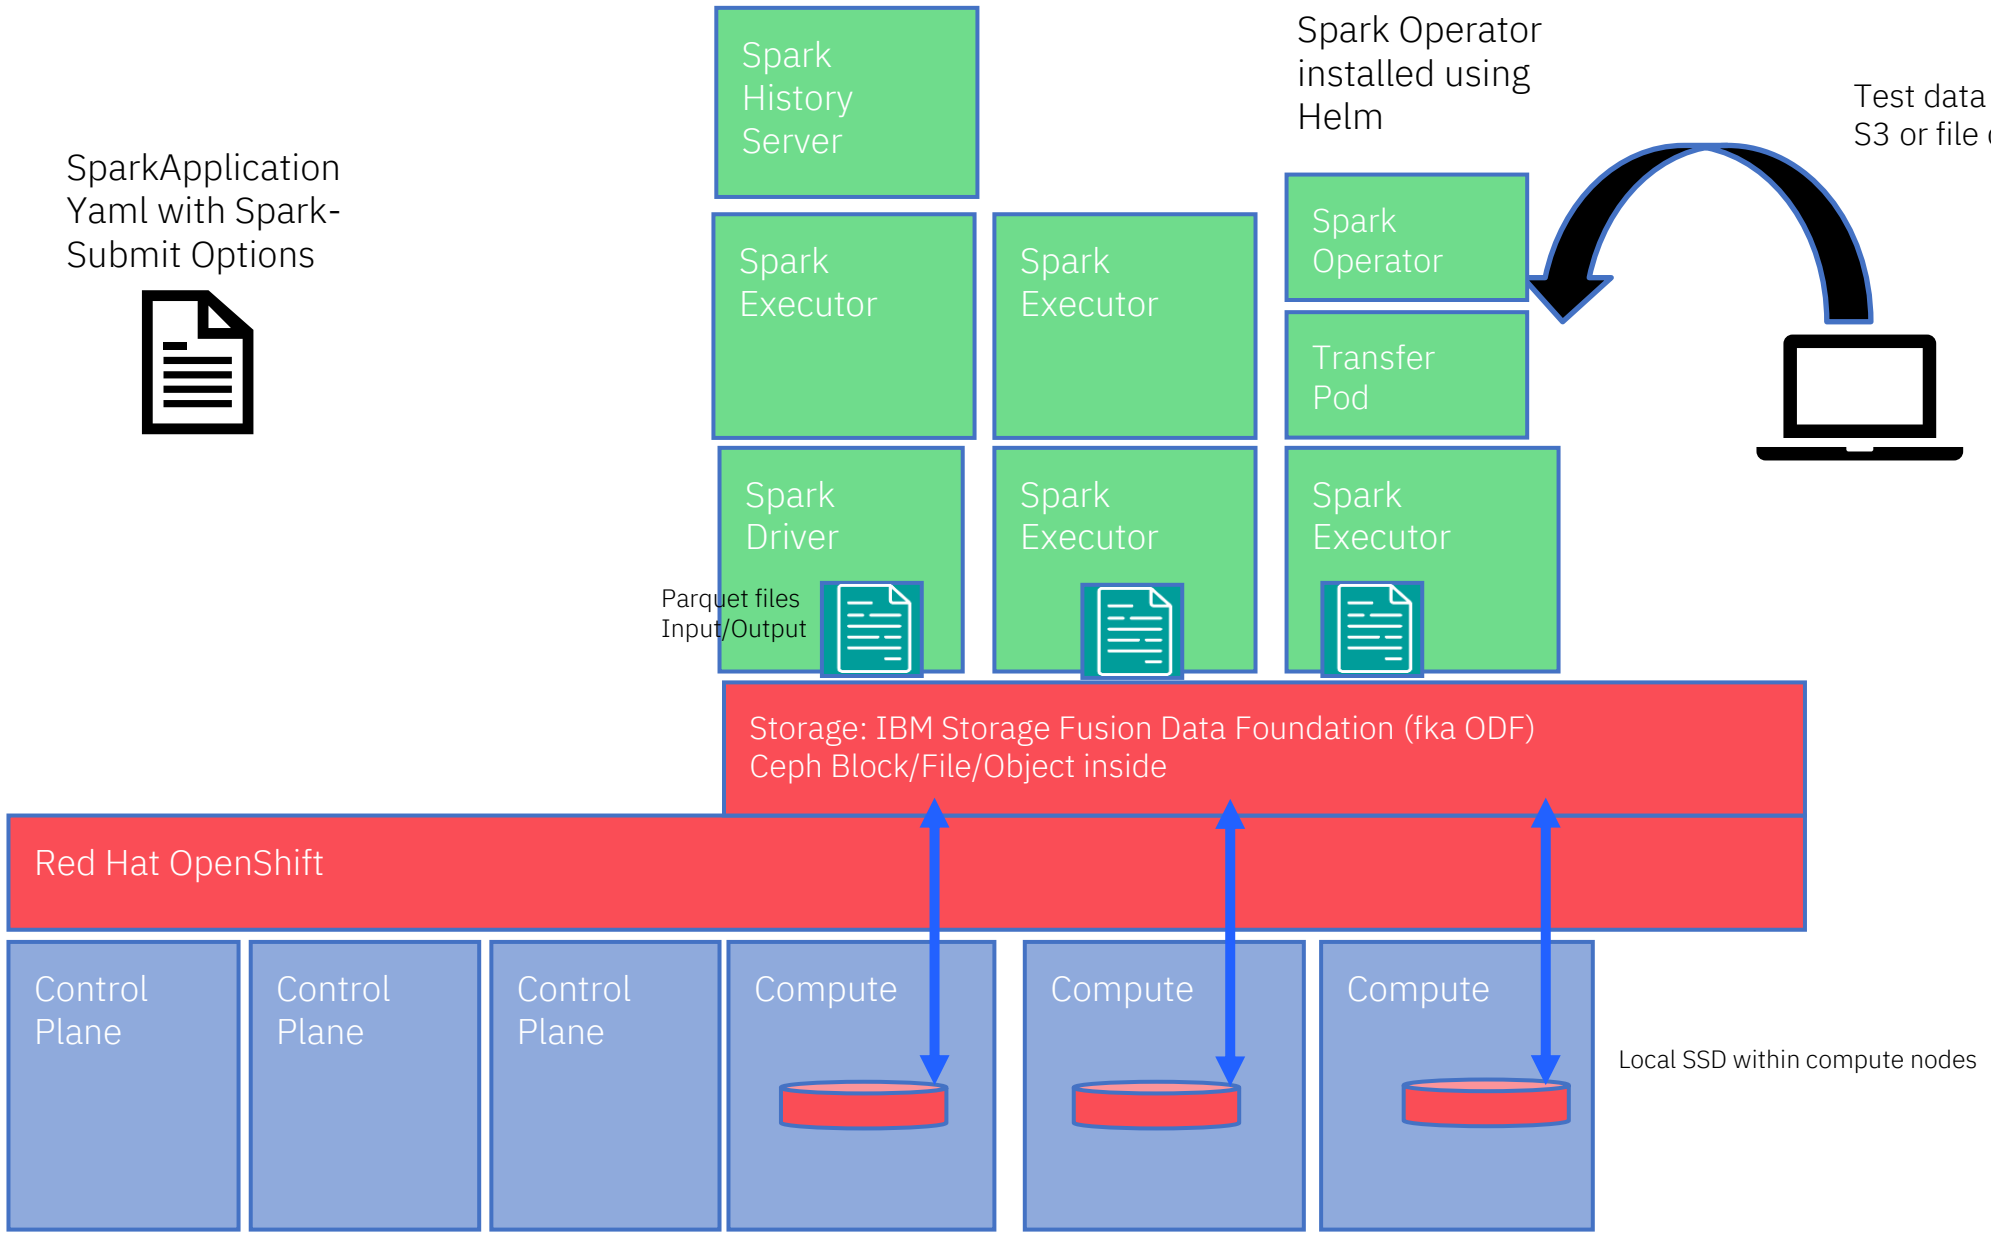
SparkApplication Yaml with Spark-Submit Options
Spark History Server
Spark Operator installed using Helm
Test data provisioning via S3 or file copy
Spark Operator
Transfer Pod
Spark Executor
Spark Executor
Spark Driver
Spark Executor
Spark Executor
Parquet files Input/Output
Storage: IBM Storage Fusion Data Foundation (fka ODF)
Ceph Block/File/Object inside
Red Hat OpenShift
Control Plane
Control Plane
Control Plane
Compute
Compute
Compute
Local SSD within compute nodes

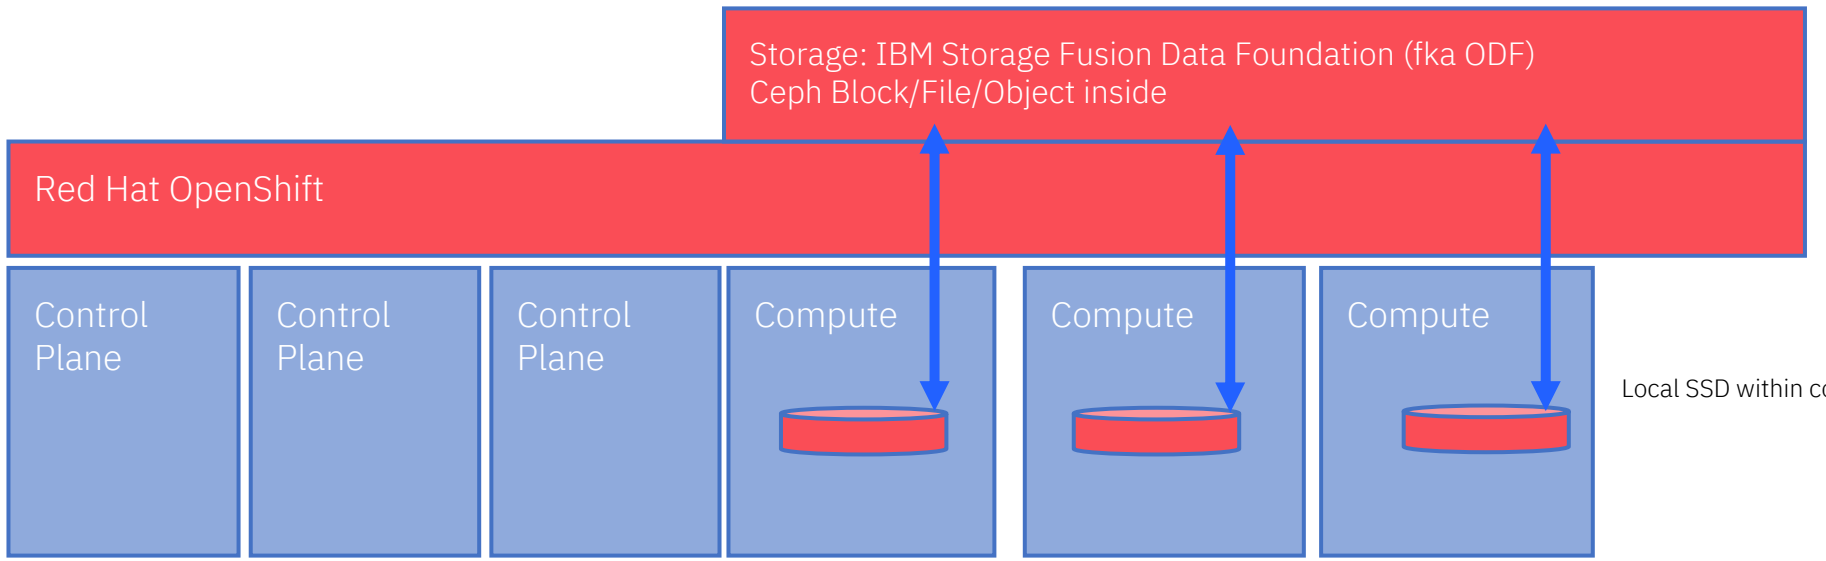

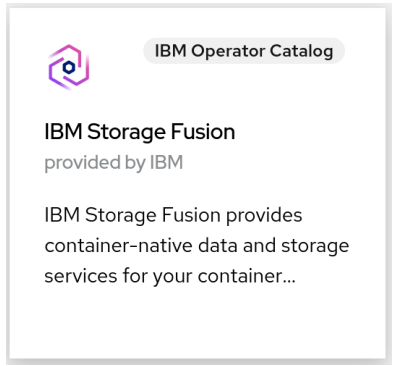

Local SSD within compute nodes

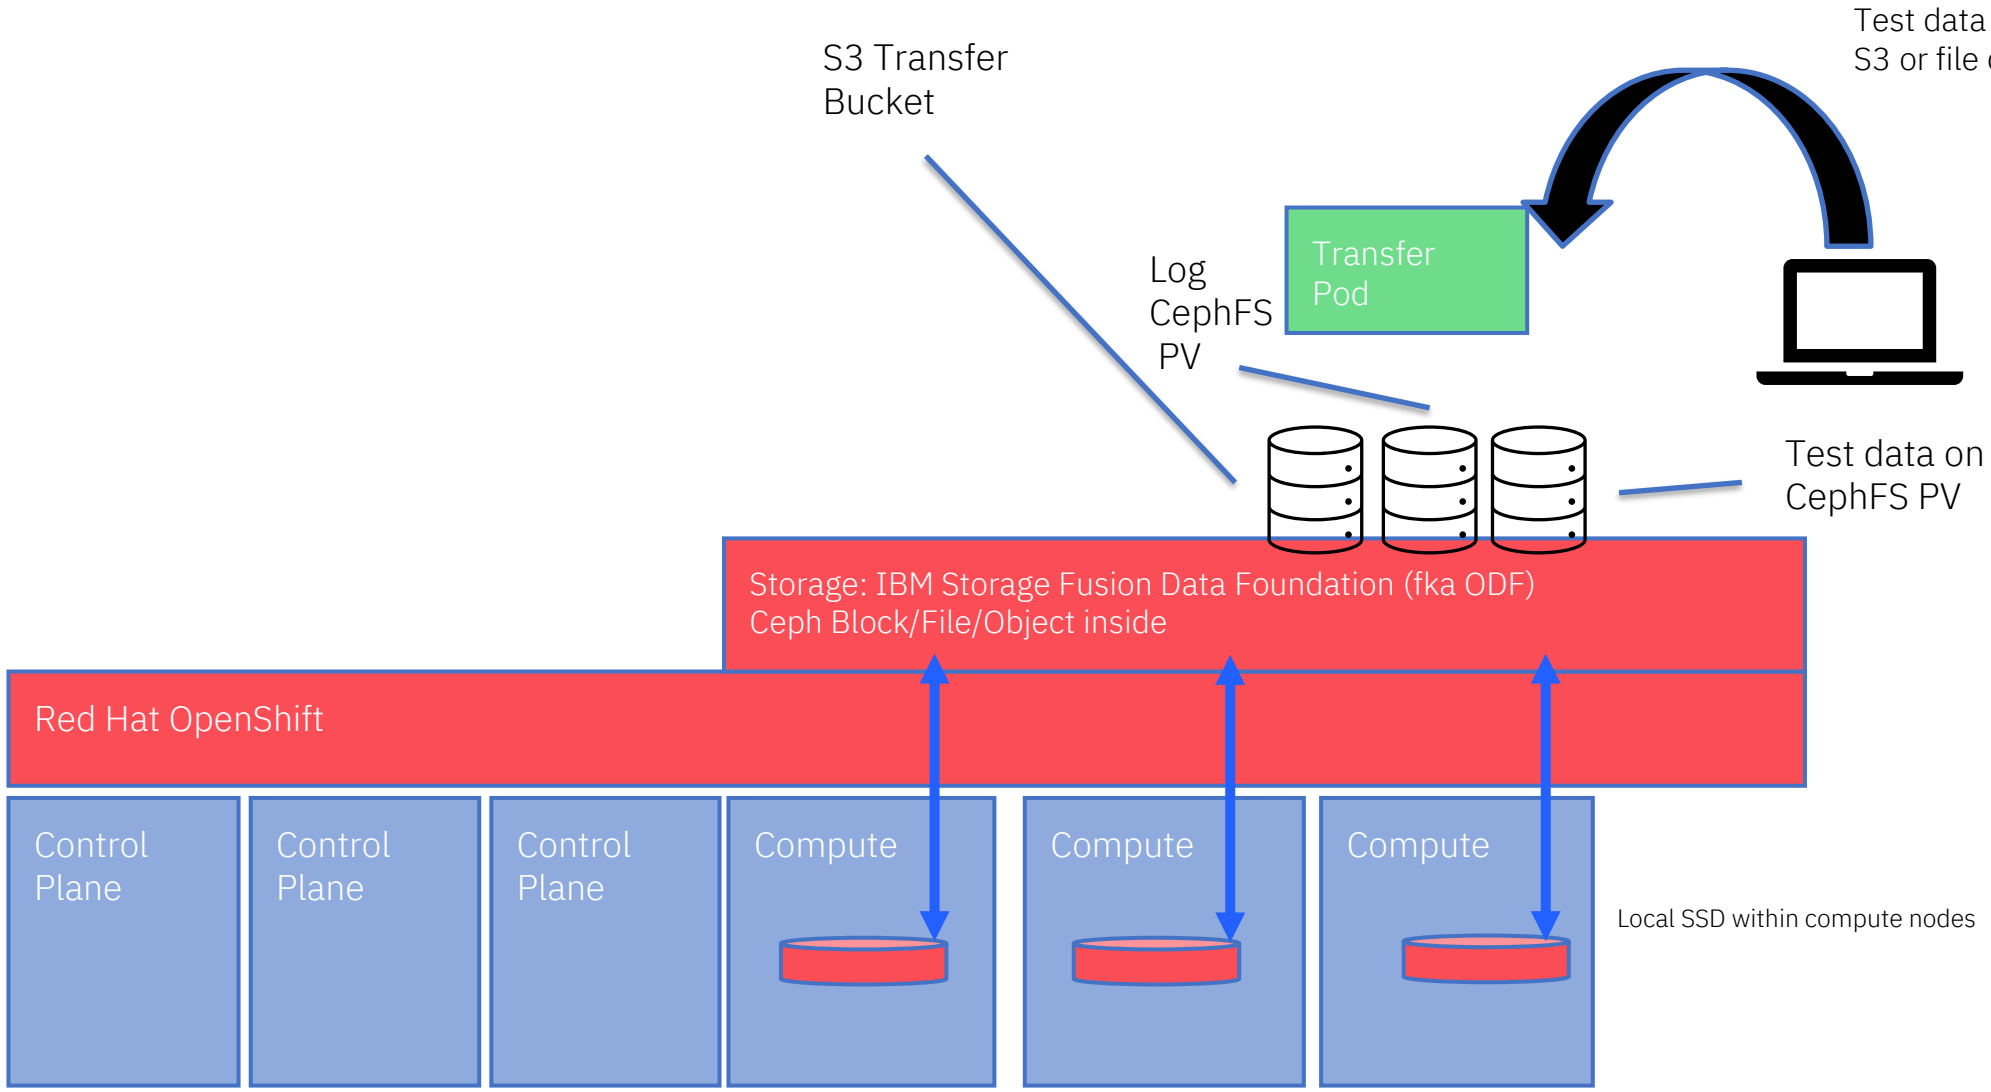
S3 Transfer Bucket
Test data provisioning via S3 or file copy
Log CephFS PV
Transfer Pod
Test data on CephFS PV
Storage: IBM Storage Fusion Data Foundation (fka ODF)
Ceph Block/File/Object inside
Red Hat OpenShift
Control Plane
Control Plane
Control Plane
Compute
Compute
Compute
Local SSD within compute nodes

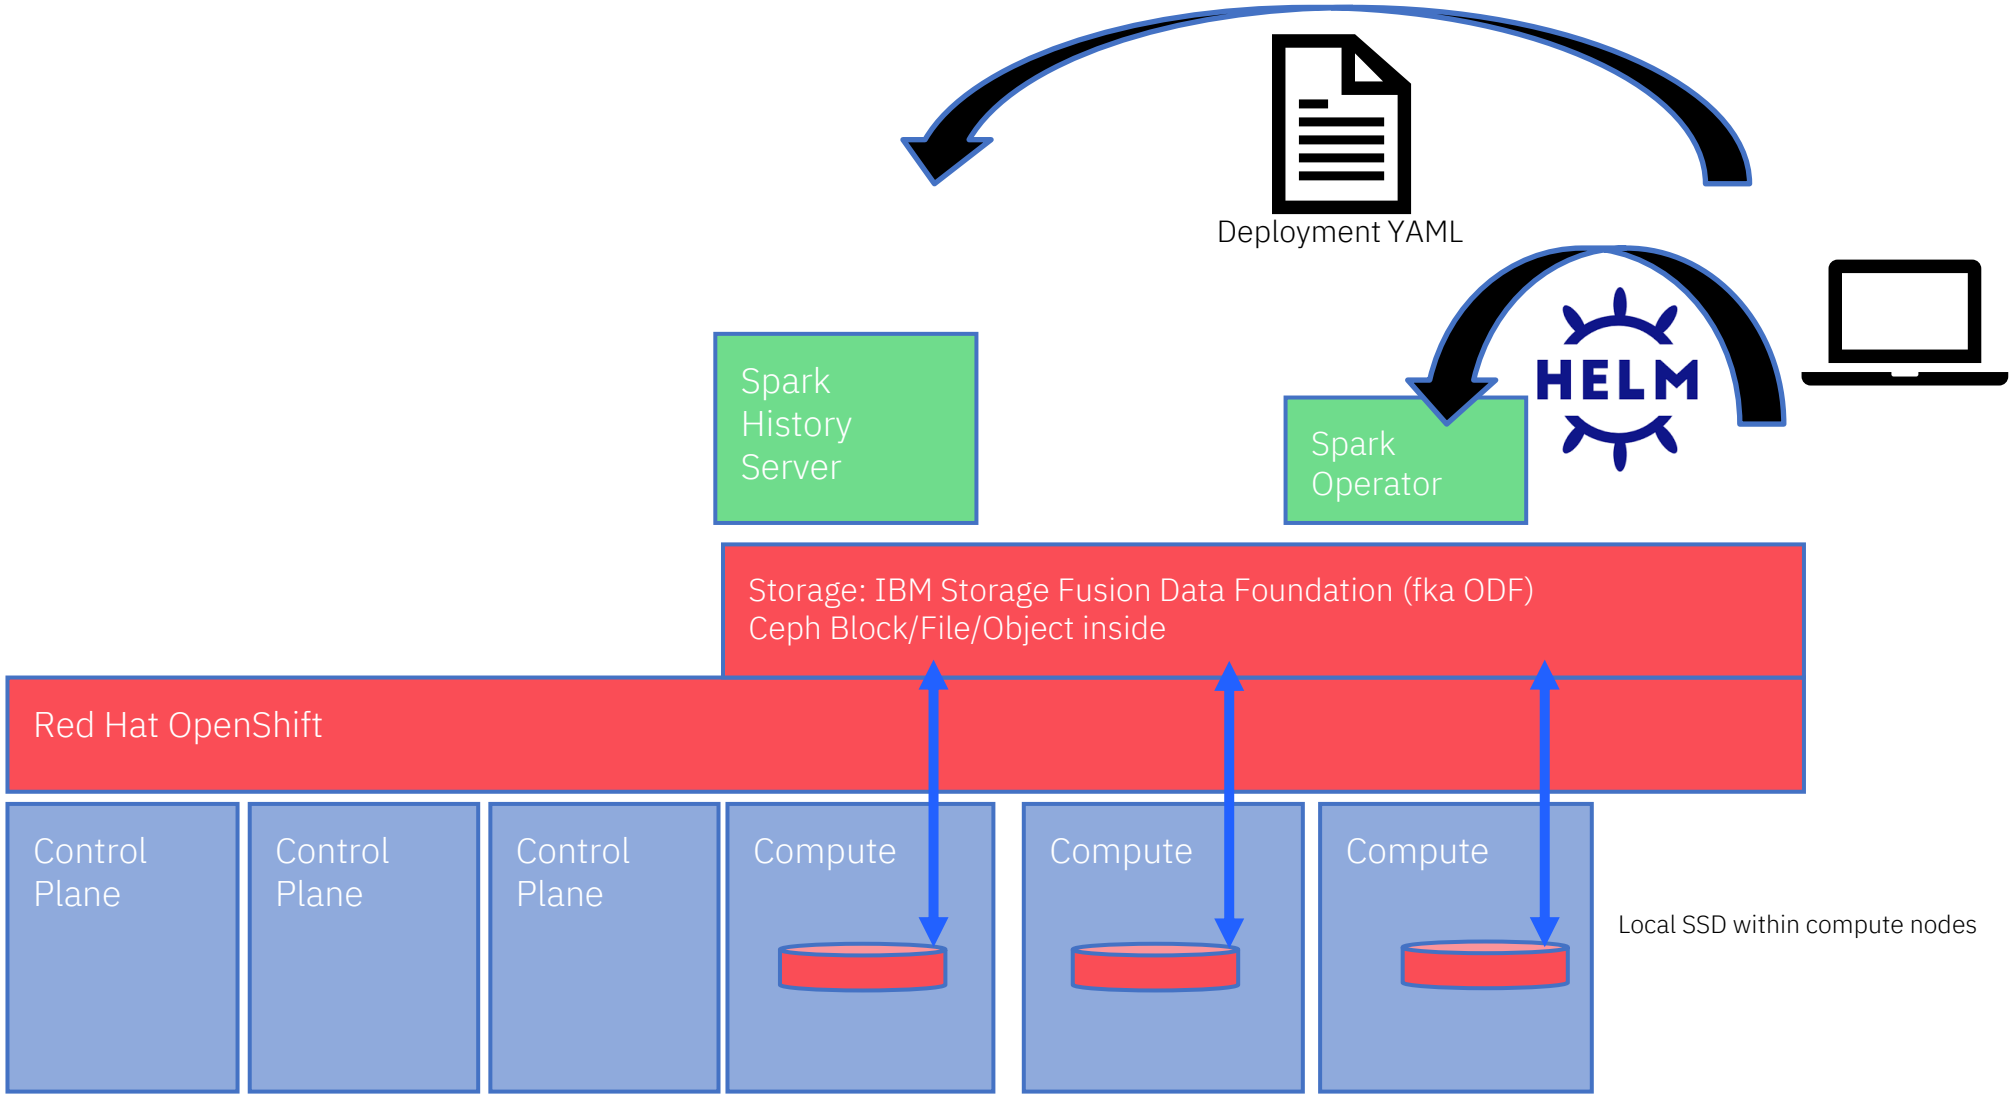
Deployment YAML
HELM
Spark History Server
Spark Operator
Storage: IBM Storage Fusion Data Foundation (fka ODF)
Ceph Block/File/Object inside
Red Hat OpenShift
Control Plane
Control Plane
Control Plane
Compute
Compute
Compute
Local SSD within compute nodes

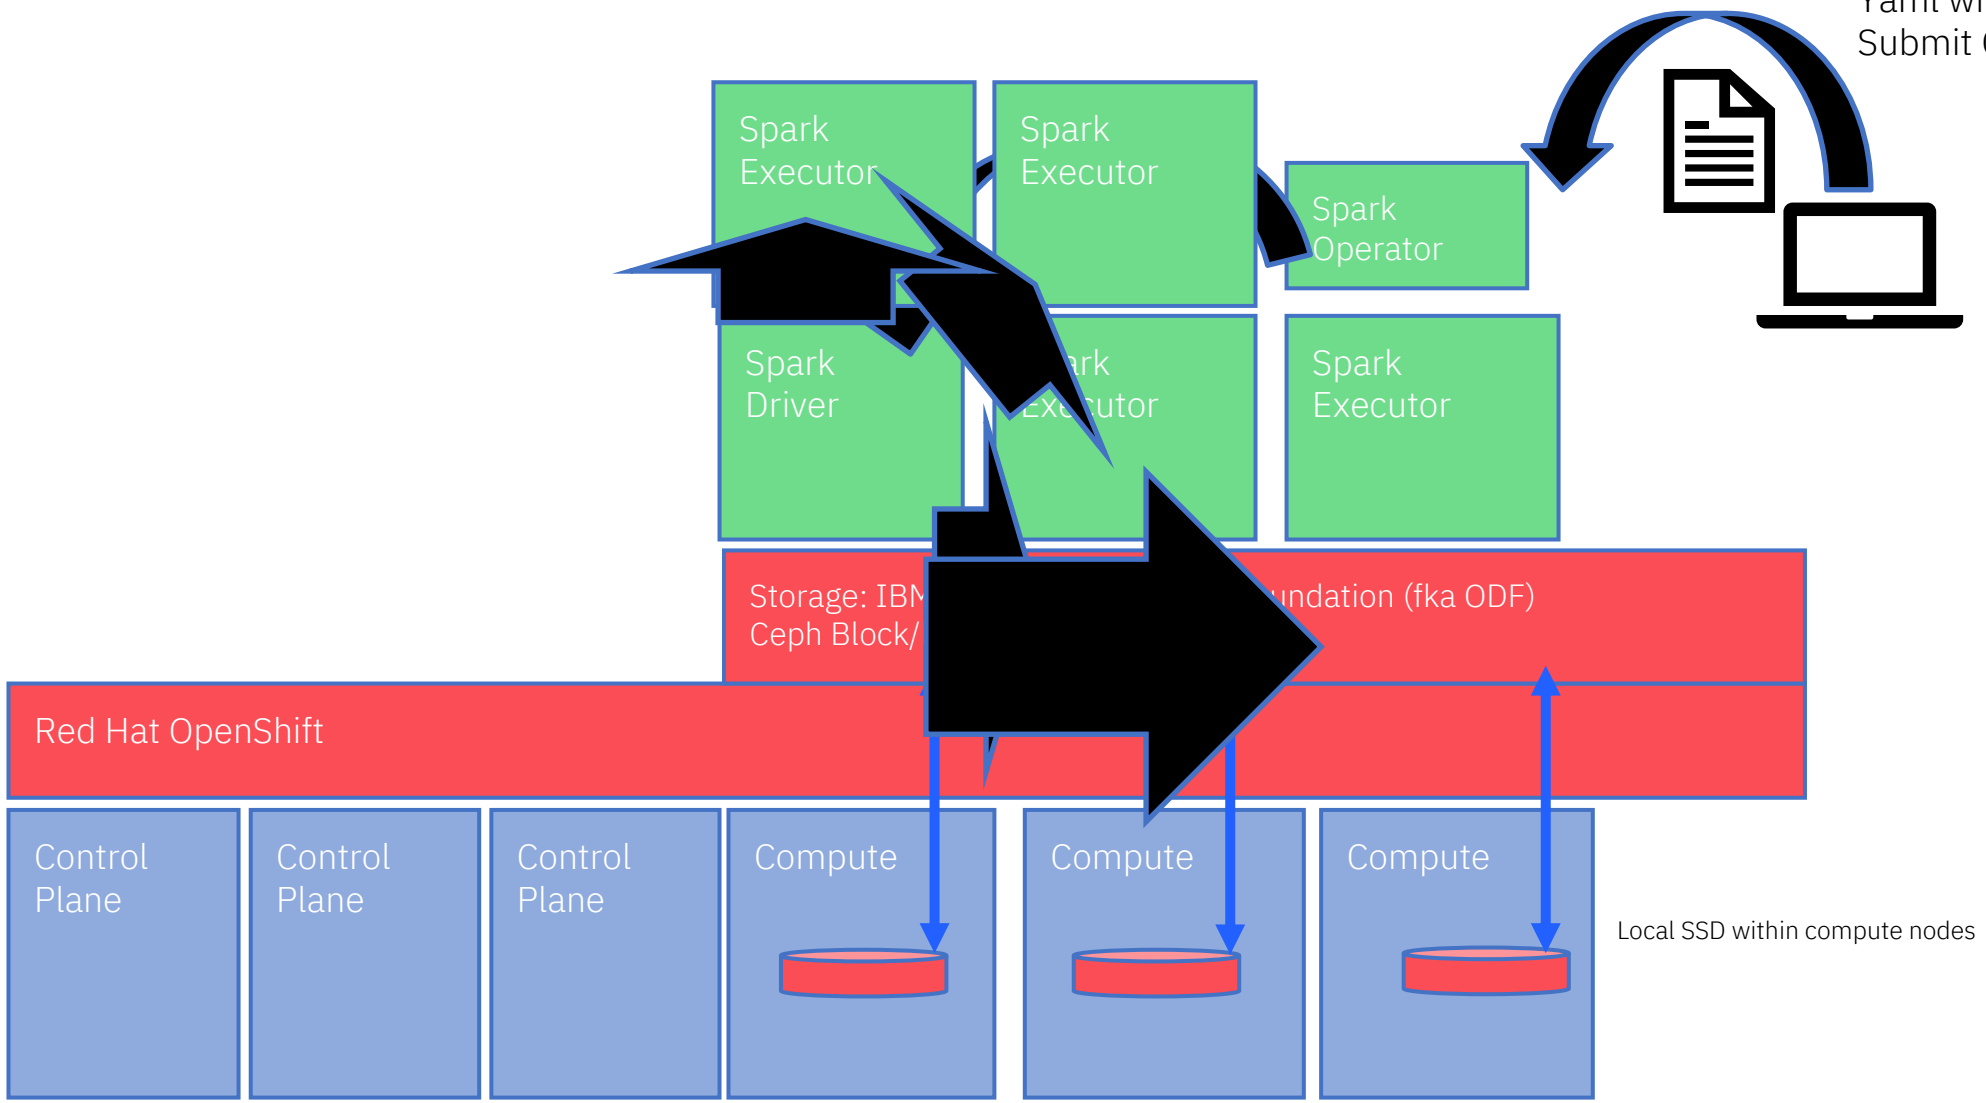
SparkApplication
Yaml with Spark-
Submit Options
Spark
Executor
Spark
Executor
Spark
Operator
Spark
Driver
Spark
Executor
Red Hat OpenShift
Control
Plane
Control
Plane
Control
Plane
Compute
Compute
Compute
Local SSD within compute nodes

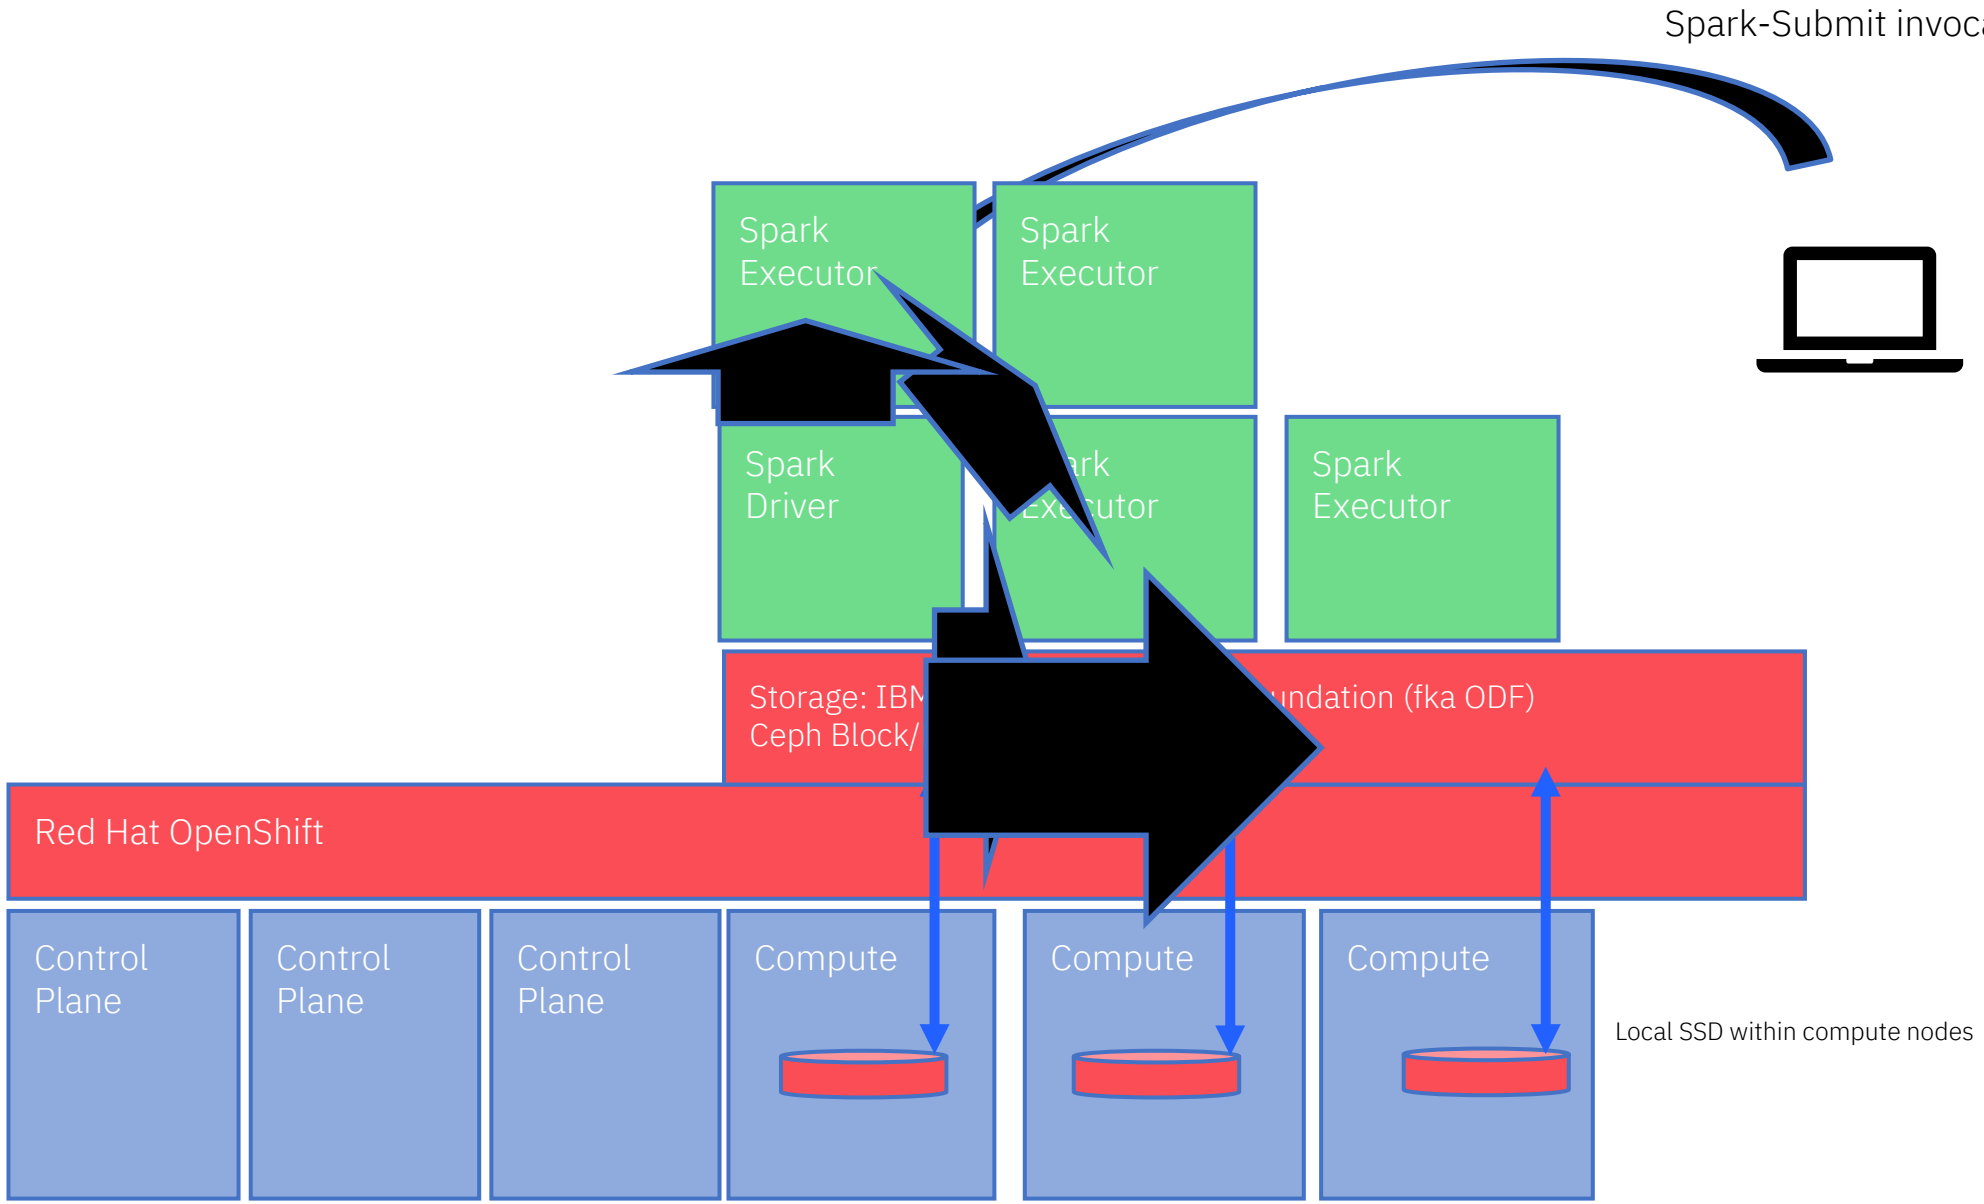
Spark-Submit invocation
Spark Executor
Spark Executor
Spark Driver
Spark Executor
Red Hat OpenShift
Control Plane
Control Plane
Control Plane
Compute
Compute
Compute
Local SSD within compute nodes

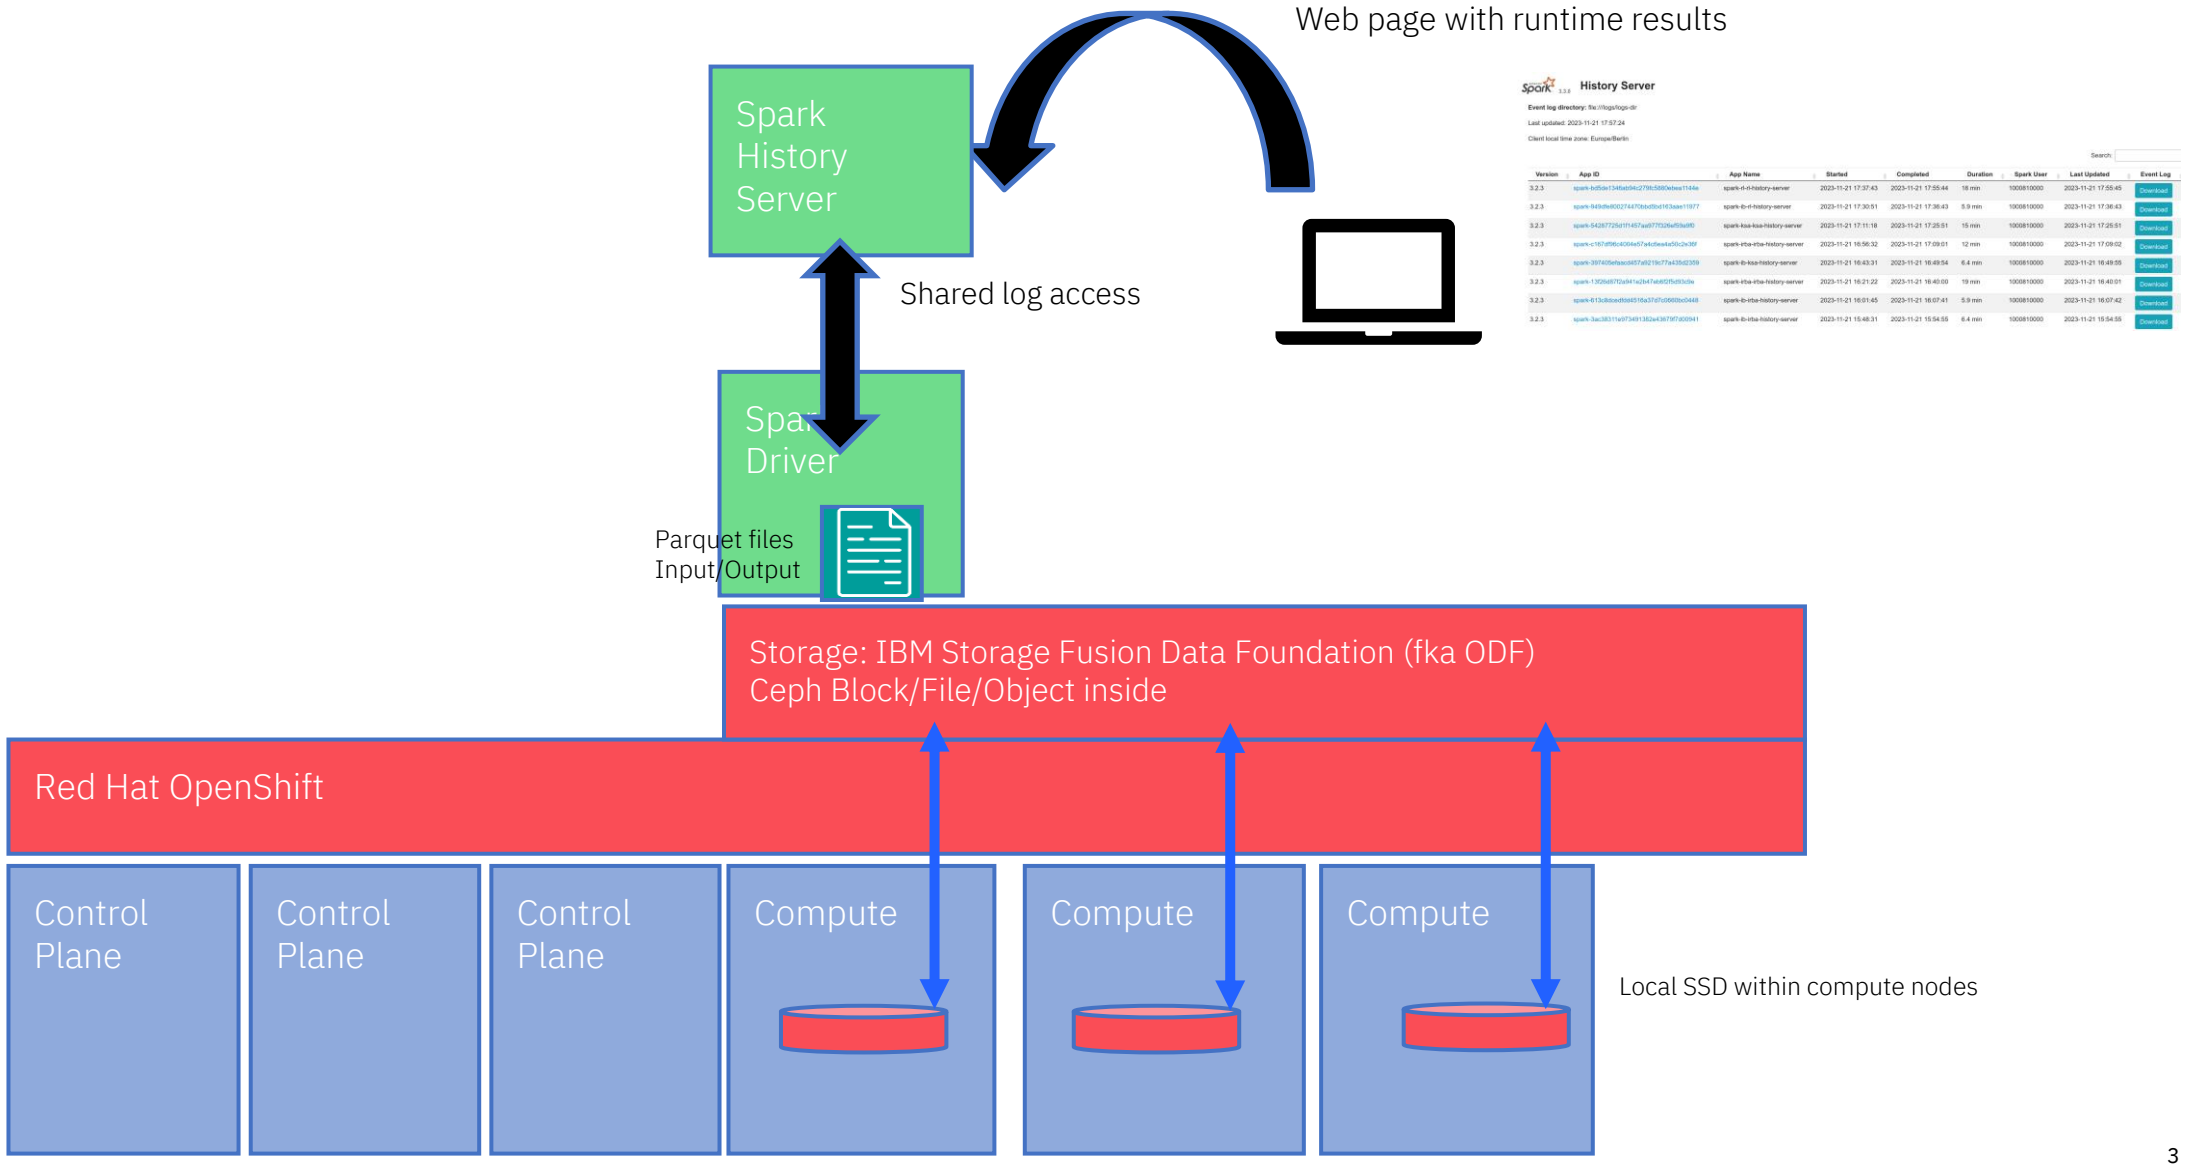
Web page with runtime results
Spark History Server
Shared log access
Spark Driver
Parquet files Input/Output
History Server
Storage: IBM Storage Fusion Data Foundation (fka ODF)
Ceph Block/File/Object inside
Red Hat OpenShift
Control Plane
Control Plane
Control Plane
Compute
Compute
Compute
Local SSD within compute nodes

Data Migration using OpenShift built-in tools

IBM Storage Scale Days 2024

March 5-7, 2024 | Stuttgart Marriott Hotel Sindelfingen

Harald Seipp, IBM

Application Stack in OpenShift

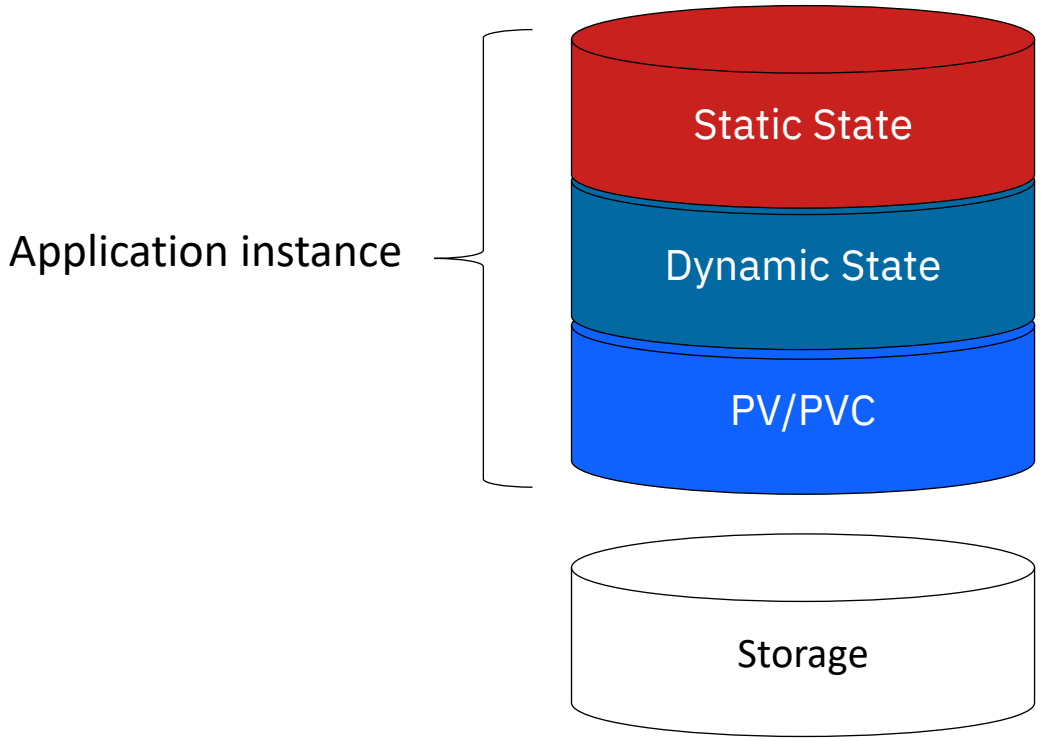

- All kube objects created once
- In source control

- All kube objects created during install and runtime
- Not in source control

- PV/PVC YAML

- Storage in the data plane

Data Migration Options

- MTC - Red Hat Migration Toolkit for Containers
  - https://docs.openshift.com/container-platform/4.14/migration_toolkit_for_containers/about-mtc.html
- Container based RSYNC
- Not covered here:
  - AFM - Storage Scale Active File Manager
    - https://www.ibm.com/docs/en/sfhs/2.7.x?topic=data-active-file-management-afm-fileset-creation
  - Fusion Data Foundation MCG - Multi-Cloud Object Gateway Services – Namespace buckets
    - https://www.ibm.com/docs/en/storage-fusion-software/2.7.x?topic=resource-managing-namespace-buckets
  - IBM Storage Fusion Backup and Restore
    - https://www.ibm.com/docs/en/sfhs/2.7.x?topic=protection-backup-restore
  - MTV – Red Hat Migration Toolkit for Virtualization
    - https://access.redhat.com/documentation/en-us/migration_toolkit_for_virtualization/2.5

MTC Migration

Before Migration

OpenShift, Filebrowser Namespace

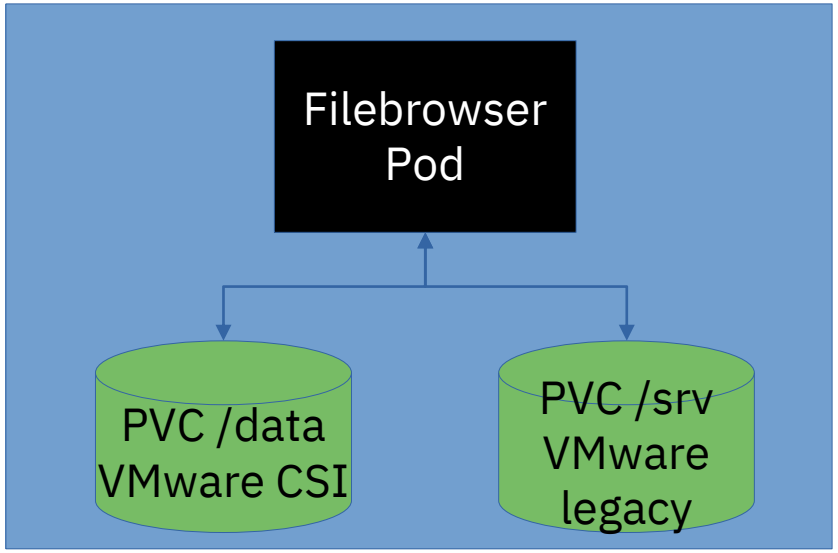

VMware vSphere VMFS

After Migration

OpenShift, Filebrowser Namespace

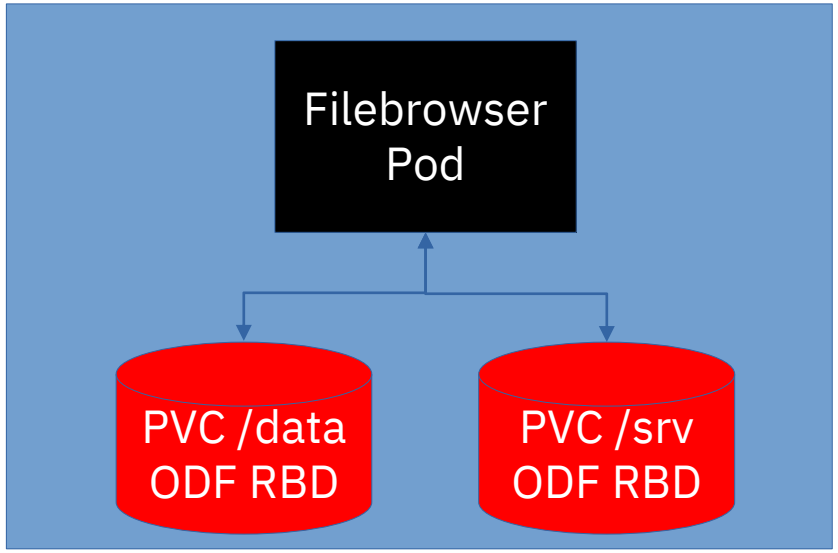

Red Hat OpenShift Data Foundation

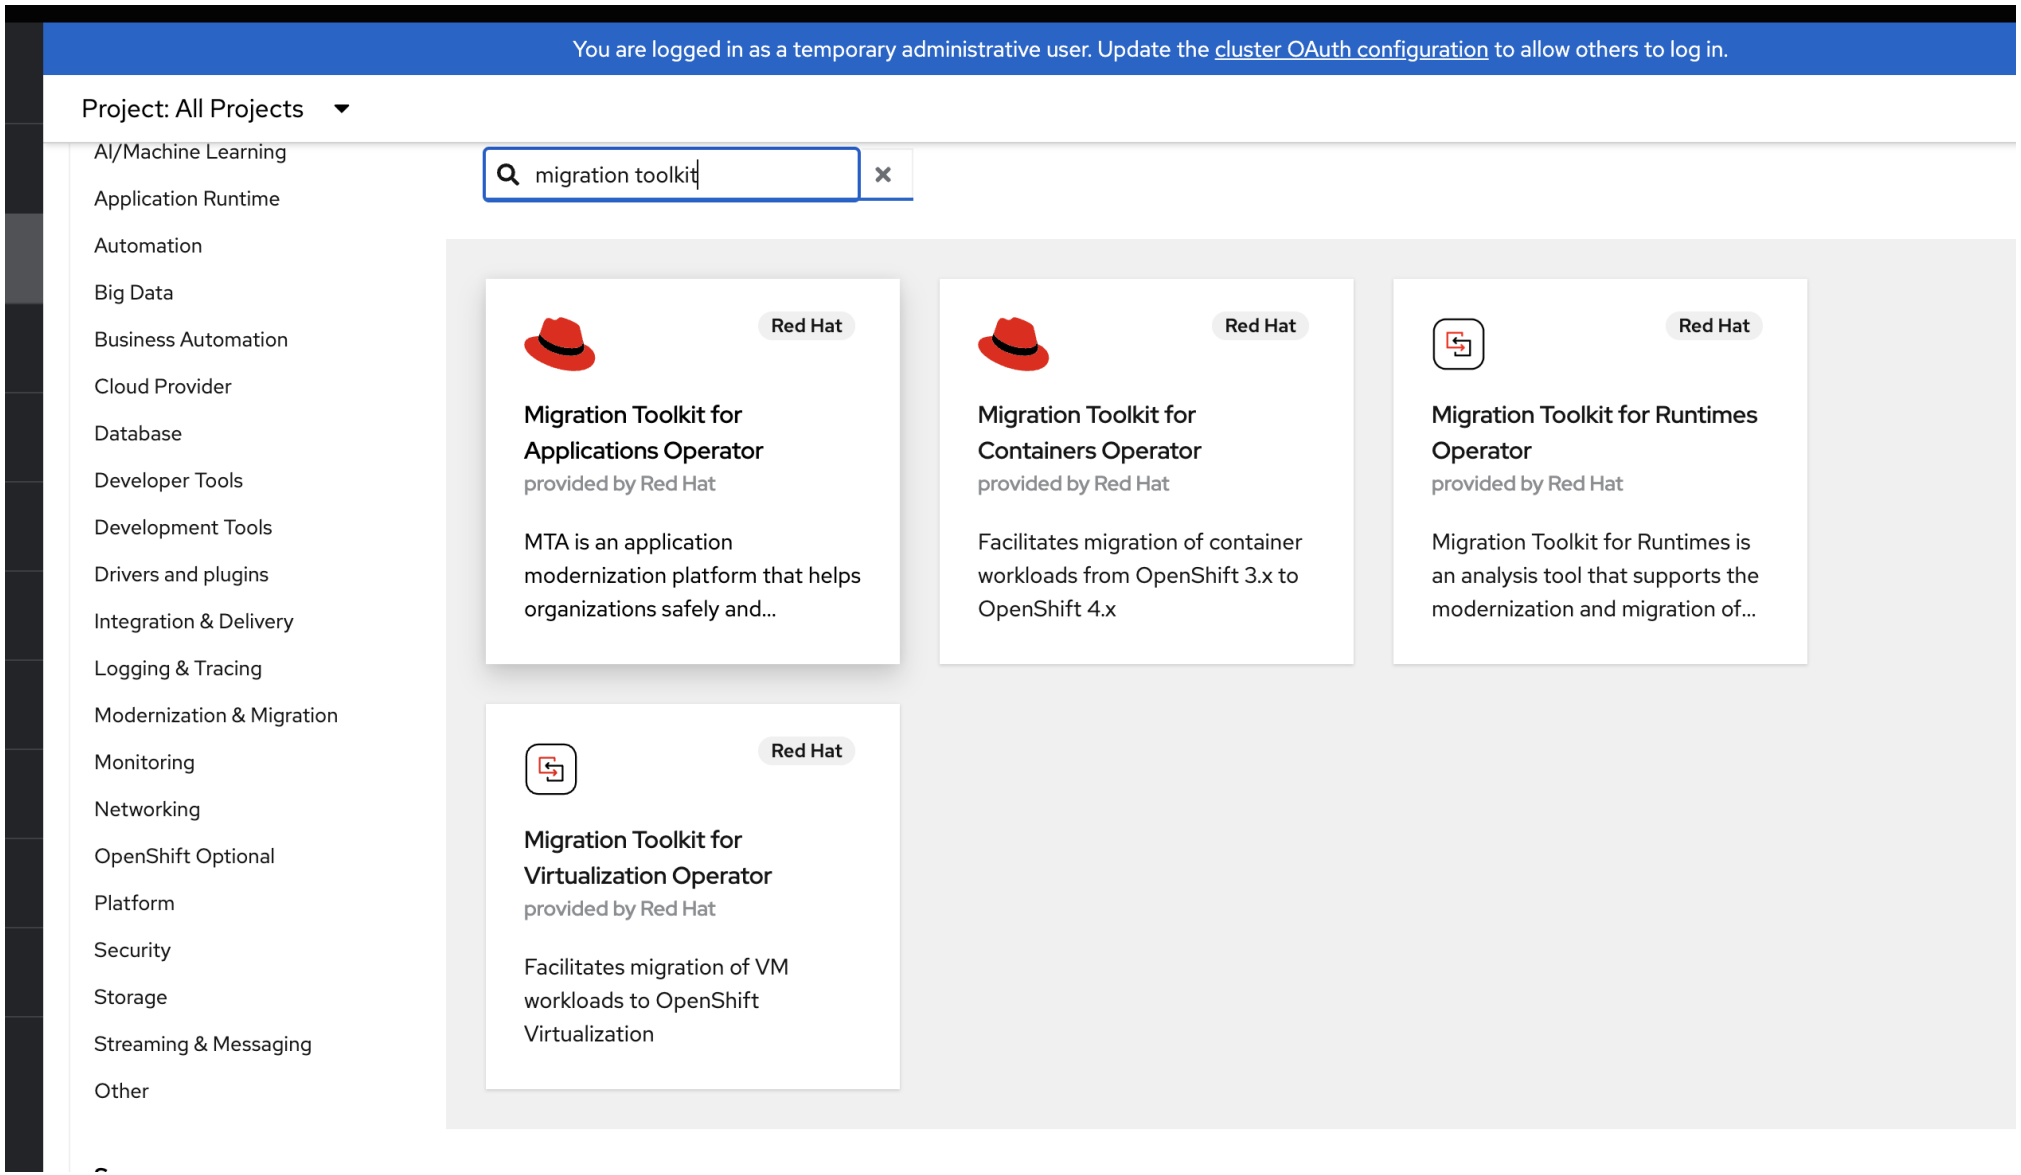
You are logged in as a temporary administrative user. Update the cluster OAuth configuration to allow others to log in.
Project: All Projects
AI/Machine Learning
Application Runtime
Automation
Big Data
Business Automation
Cloud Provider
Database
Developer Tools
Development Tools
Drivers and plugins
Integration & Delivery
Logging & Tracing
Modernization & Migration
Monitoring
Networking
OpenShift Optional
Platform
Security
Storage
Streaming & Messaging
Other
migration toolkit
Red Hat
Migration Toolkit for Applications Operator
provided by Red Hat
MTA is an application modernization platform that helps organizations safely and...
Red Hat
Migration Toolkit for Containers Operator
provided by Red Hat
Facilitates migration of container workloads from OpenShift 3.x to OpenShift 4.x
Red Hat
Migration Toolkit for Runtimes Operator
provided by Red Hat
Migration Toolkit for Runtimes is an analysis tool that supports the modernization and migration of...
Red Hat
Migration Toolkit for Virtualization Operator
provided by Red Hat
Facilitates migration of VM workloads to OpenShift Virtualization

Migration Toolkit for Containers Operator

1.8.2 provided by Red Hat

Install

Latest version

1.8.2

Capability level

- Basic Install
- Seamless Upgrades
- Full Lifecycle
- Deep Insights
- Auto Pilot

Source

Red Hat

Provider

Red Hat

The Migration Toolkit for Containers Operator enables installation of the application migration tool components.

After you have installed the Operator, you can install components by creating a MigrationController CR.

By default, the Operator installs the following components on an OpenShift Container Platform 4 target cluster:

- velero and restic, to back up data on the source cluster and restore it on the target cluster
- mig-controller, to coordinate the backup and restore processes
- mig-ui, the web console for managing migrations

If you experience any issues or have feature requests, please file an issue

For more information, see the OpenShift Container Platform Migration Guide.

Project: openshift-migration

Installed Operators > Operator details

Migration Toolkit for Containers Operator

1.8.2 provided by Red Hat

Actions

Details | YAML | Subscription | Events | All instances | DirectVolumeMigration | DirectVolumeMigrationProgress | DirectImageMigration | DirectImageStre

MigrationControllers

Create MigrationController

No operands found

Operands are declarative components used to define the behavior of the application.

You are logged in as a temporary administrative user. Update the cluster OAuth configuration to allow others to log in.

Project: openshift-migration

Installed Operators > Operator details

Migration Toolkit for Containers Operator
1.8.2 provided by Red Hat

Details | YAML | Subscription | Events | All instances | DirectVolumeMigration | DirectVolumeMigrationProgress | DirectImageMigrat

MigrationControllers

Name | Search by name... /

Name	Kind	Status	Labels	Last updated
MC migration-controller	MigrationController	Condition: Running	No labels	Jan 23, 2024, 3:38 PM

Connect Clusters via the Migration Toolkit for Containers

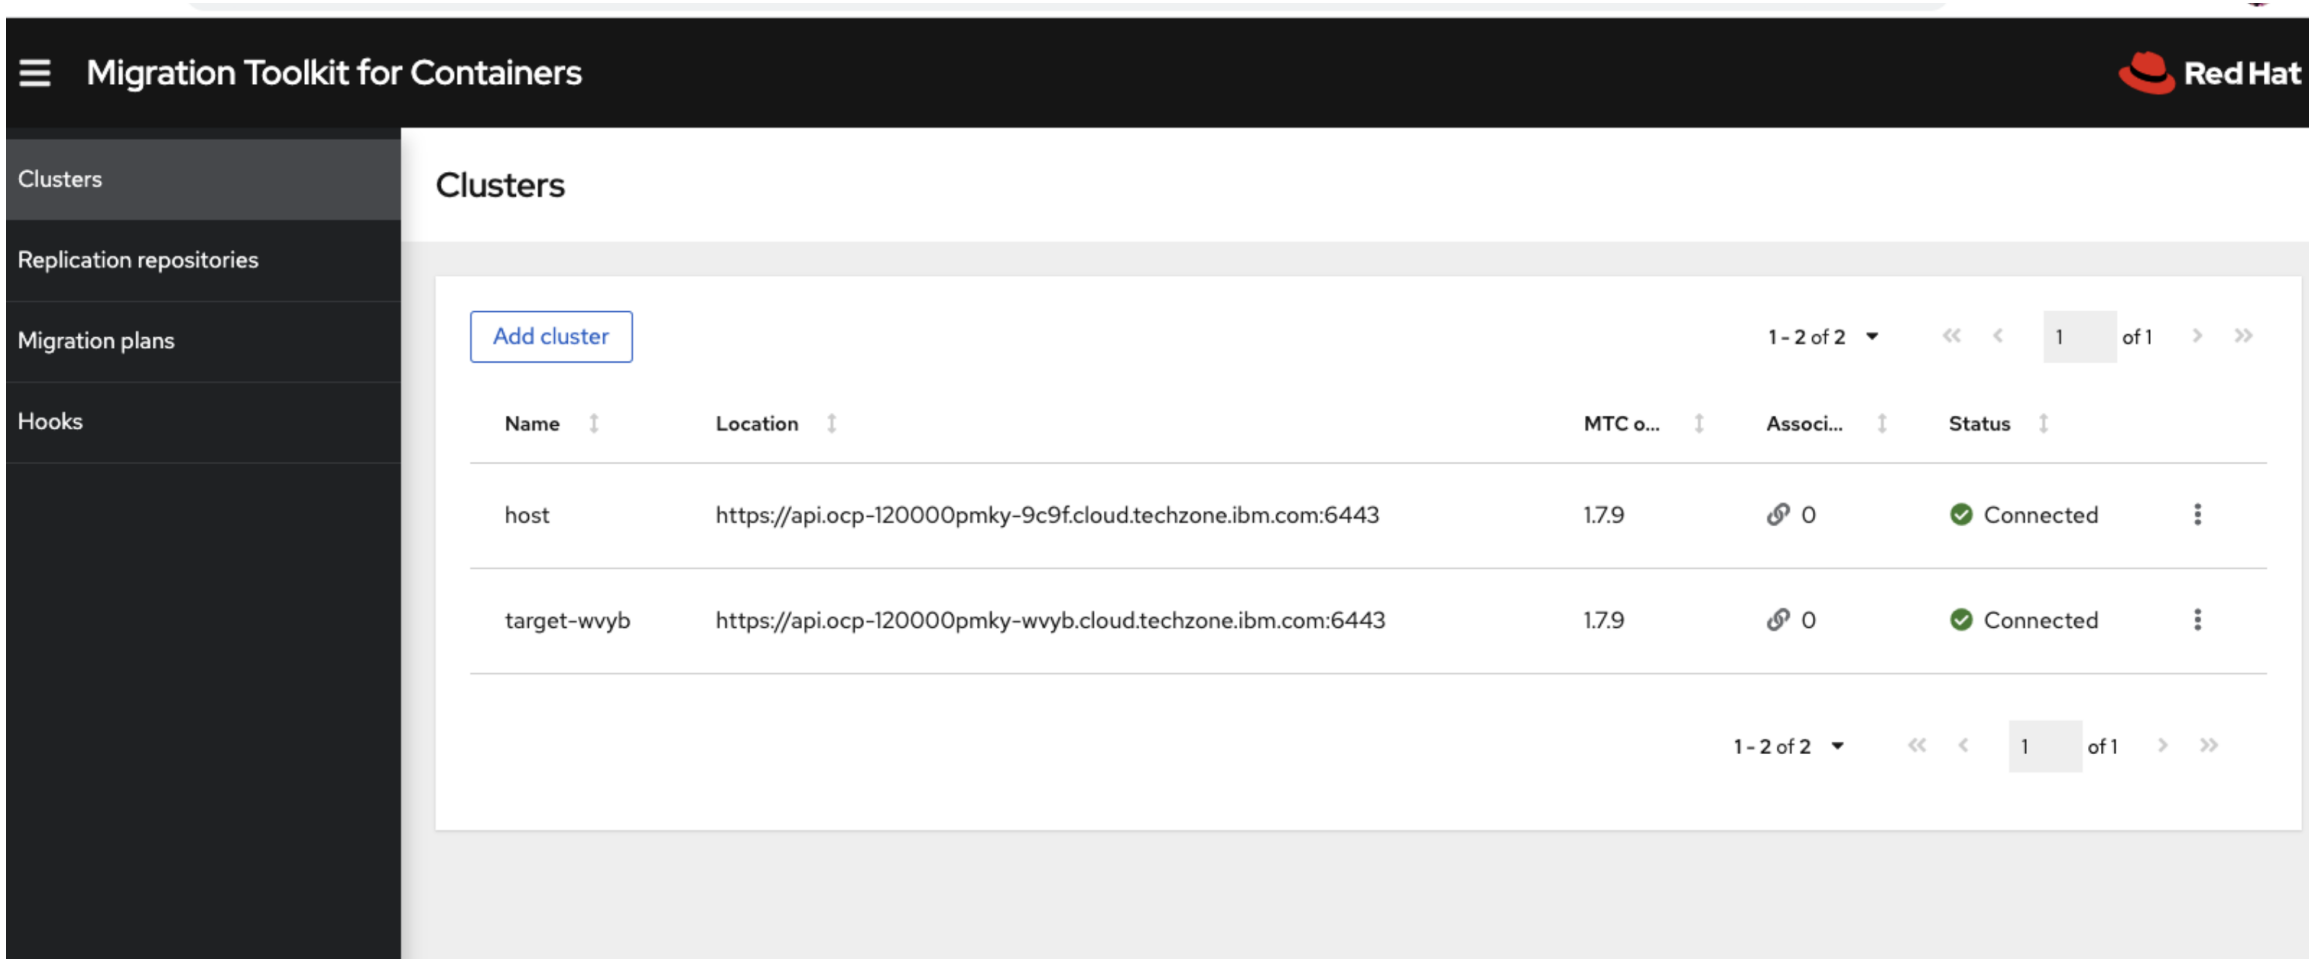

Create a Migration Plan

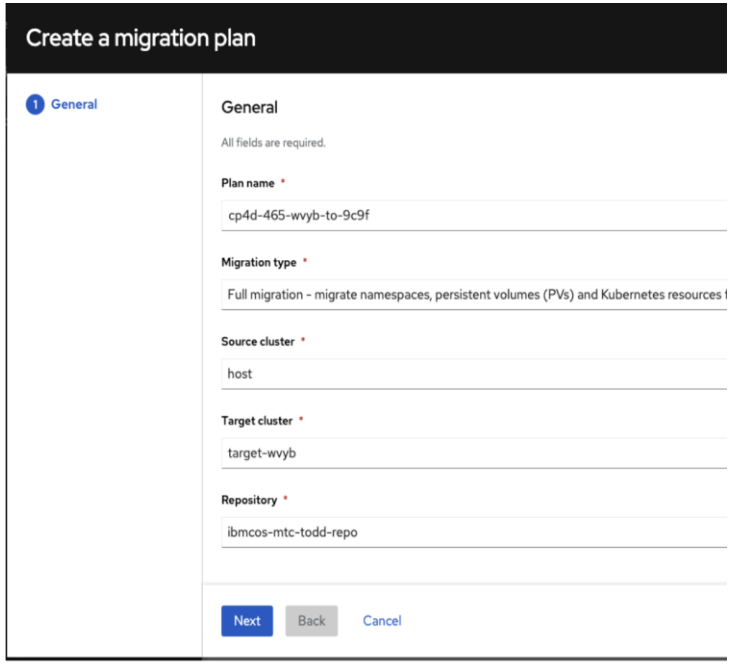

1. Define Namespaces

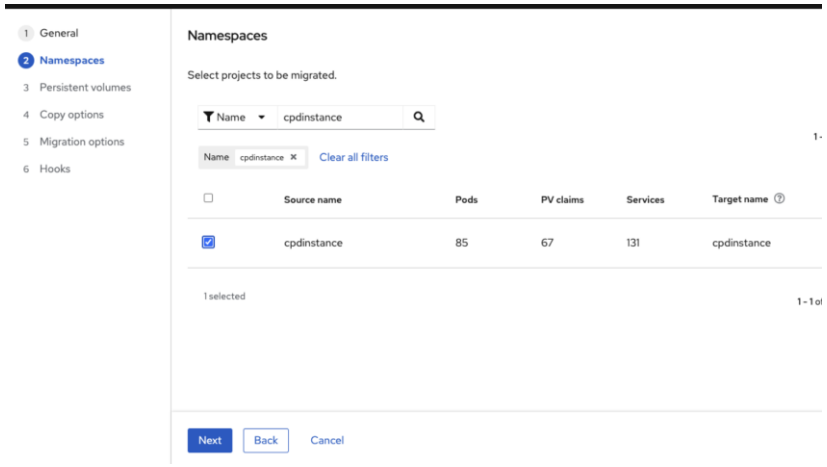

2. Select Volumes

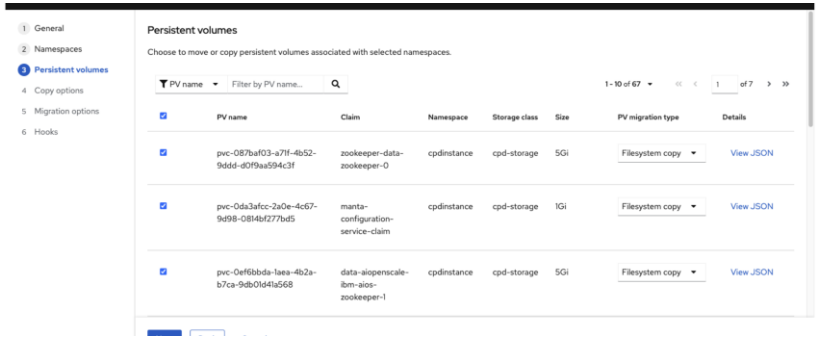

3. Define Copy Actions

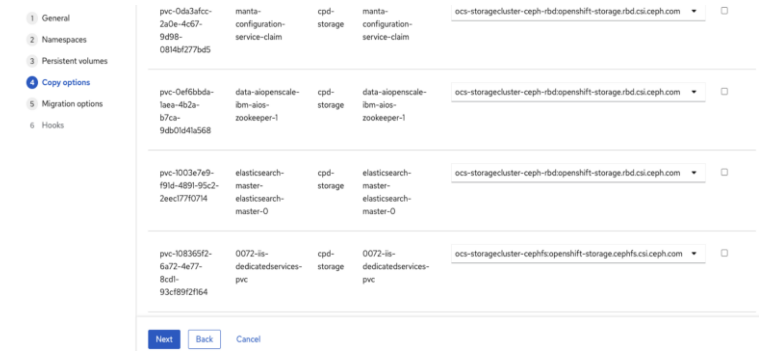

RSYNC Migration – when MTC is not available*

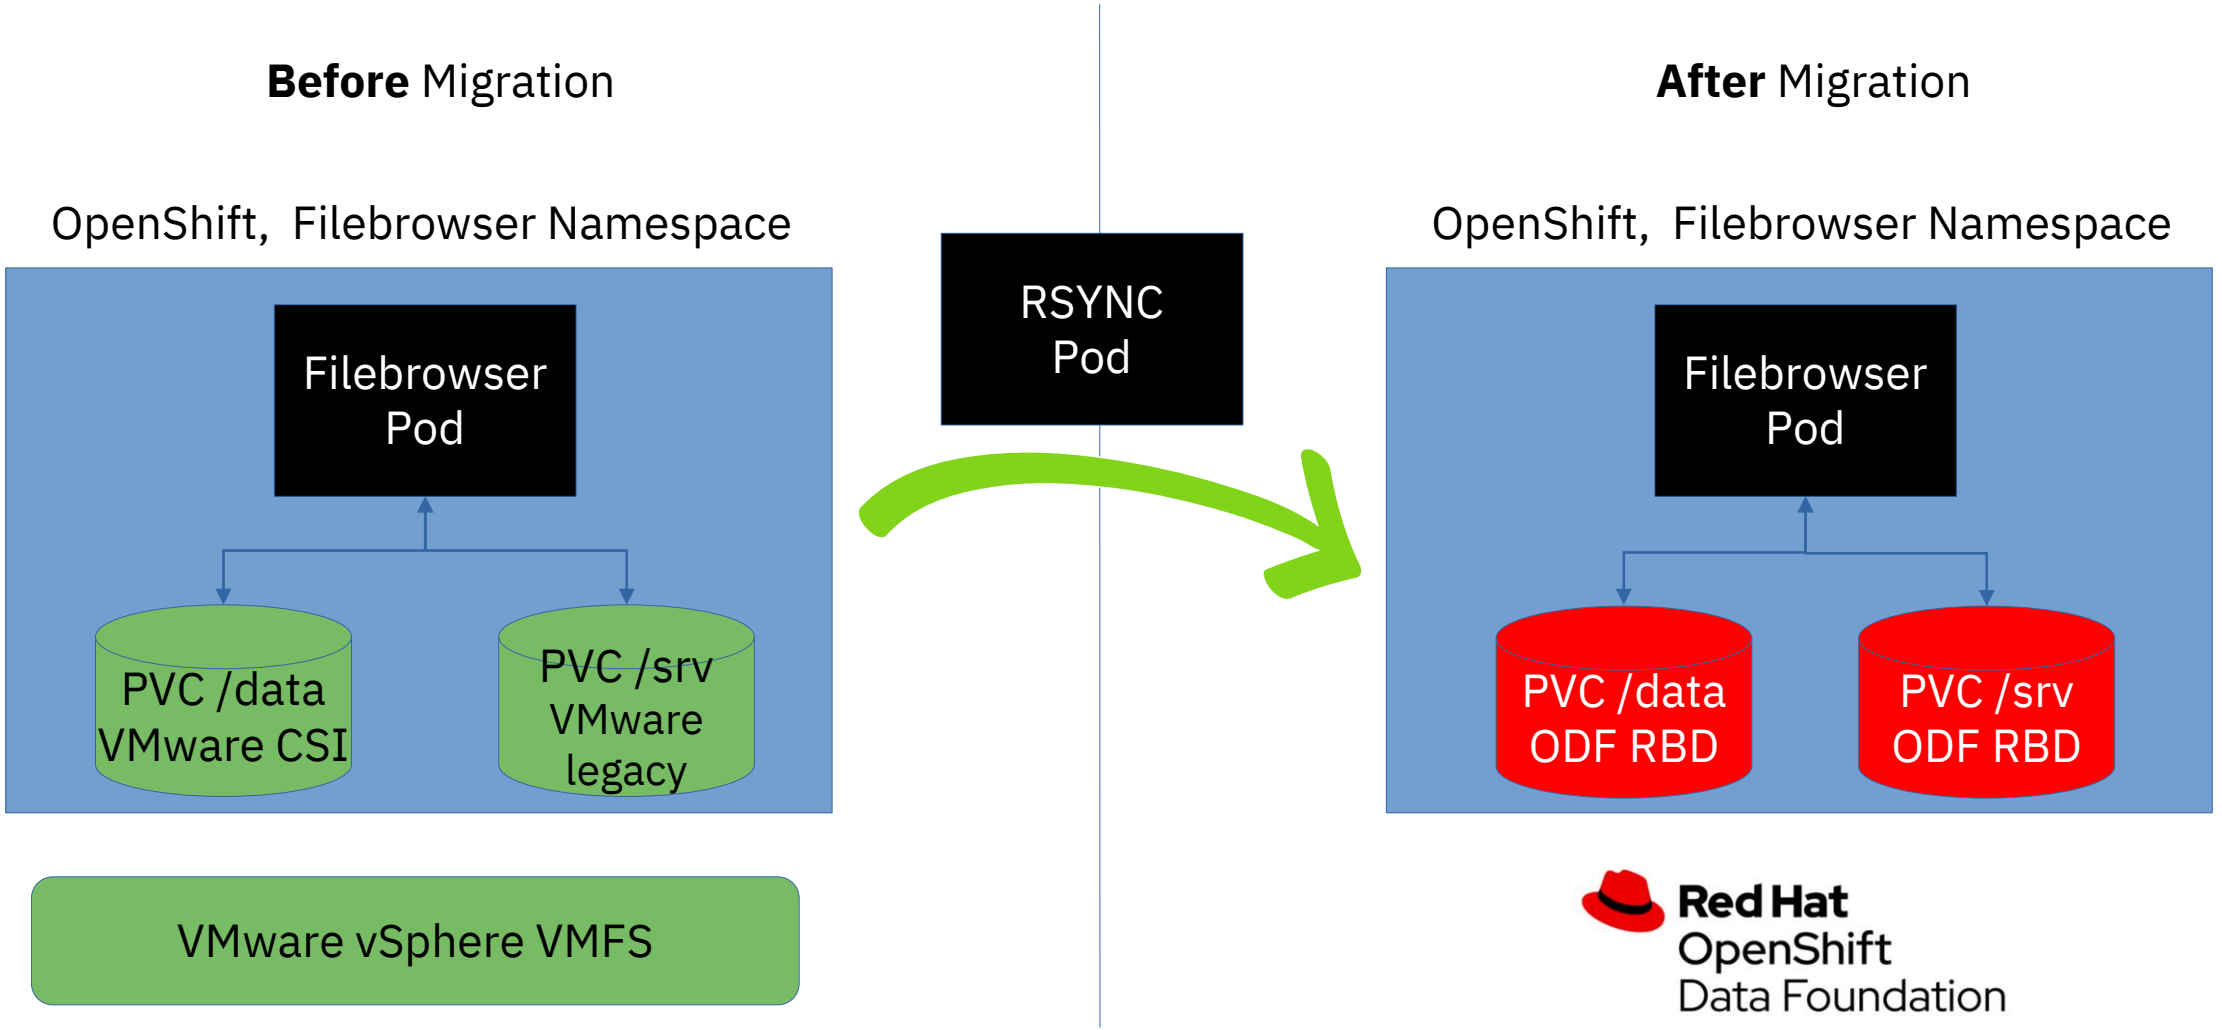

* MTC is only available on x86

RSYNC Migration – Steps

Get Current PVC

```
$ oc get pvc
NAME                          STATUS   VOLUME
CAPACITY   ACCESS MODES   STORAGECLASS                AGE
filebrowser-storage           Bound    pvc-1a53de42-7dae-4713-8820-30cd7dcec2a0   1Gi
RWO            nfs-client                  27d
```

Create the target PVC (simple RWX PVC using Scale StorageClass)

```
$ oc apply -f datacopy-target-pvc.yaml
```

Start the deployment of the "copy container“ (see →)

```
$ oc apply -f datacopy-deployment.yaml
```

Using that container, we can perform the rsync operation to copy the data

```
$ oc exec -it $(oc get pods -l app=datacopy -o name) \
 -- rsync -av /mnt/src/ /mnt/tgt/
```

Patch the deployment to use the target volume

```
$ oc patch deployment filebrowser -p '{"spec":\
{"template": {"spec": {"volumes": [{"name":\
"filebrowser-storage", "persistentVolumeClaim":\
{"claimName": "filebrowser-storage-target"}}]}}}}'
```

```
apiVersion: apps/v1
kind: Deployment
metadata:
  name: datacopy
spec:
  selector:
    matchLabels:
      app: datacopy
  replicas: 1
  template:
    metadata:
      labels:
        app: datacopy
    spec:
      volumes:
        - name: source
          persistentVolumeClaim:
            claimName: filebrowser-storage
        - name: target
          persistentVolumeClaim:
            claimName: filebrowser-storage-
target
      containers:
        - name: datacopy
image: registry.redhat.io/rhmtc/openshift-
migration-rsync-transfer-
rhel8@sha256:4376360c9c5d9106dd690d4a0cdda9a095
eb868c92fa21d7e70a56fd6a353bba
          command: [ "/bin/sh", "-c", "--" ]
          args: [ "while true; do sleep 30;
done;" ]
          securityContext:
            allowPrivilegeEscalation: false
            runAsUser: null
            capabilities:
              drop:
              - ALL
            runAsNonRoot: true
            seccompProfile:
              type: RuntimeDefault
          volumeMounts:
            - name: source
              mountPath: /mnt/src
            - name: target
              mountPath: /mnt/tgt
```

Migrate existing Storage Classes to OpenShift Data Foundation

https://community.ibm.com/community/user/blogs/harald-seipp/2023/01/12/datamigration2odf

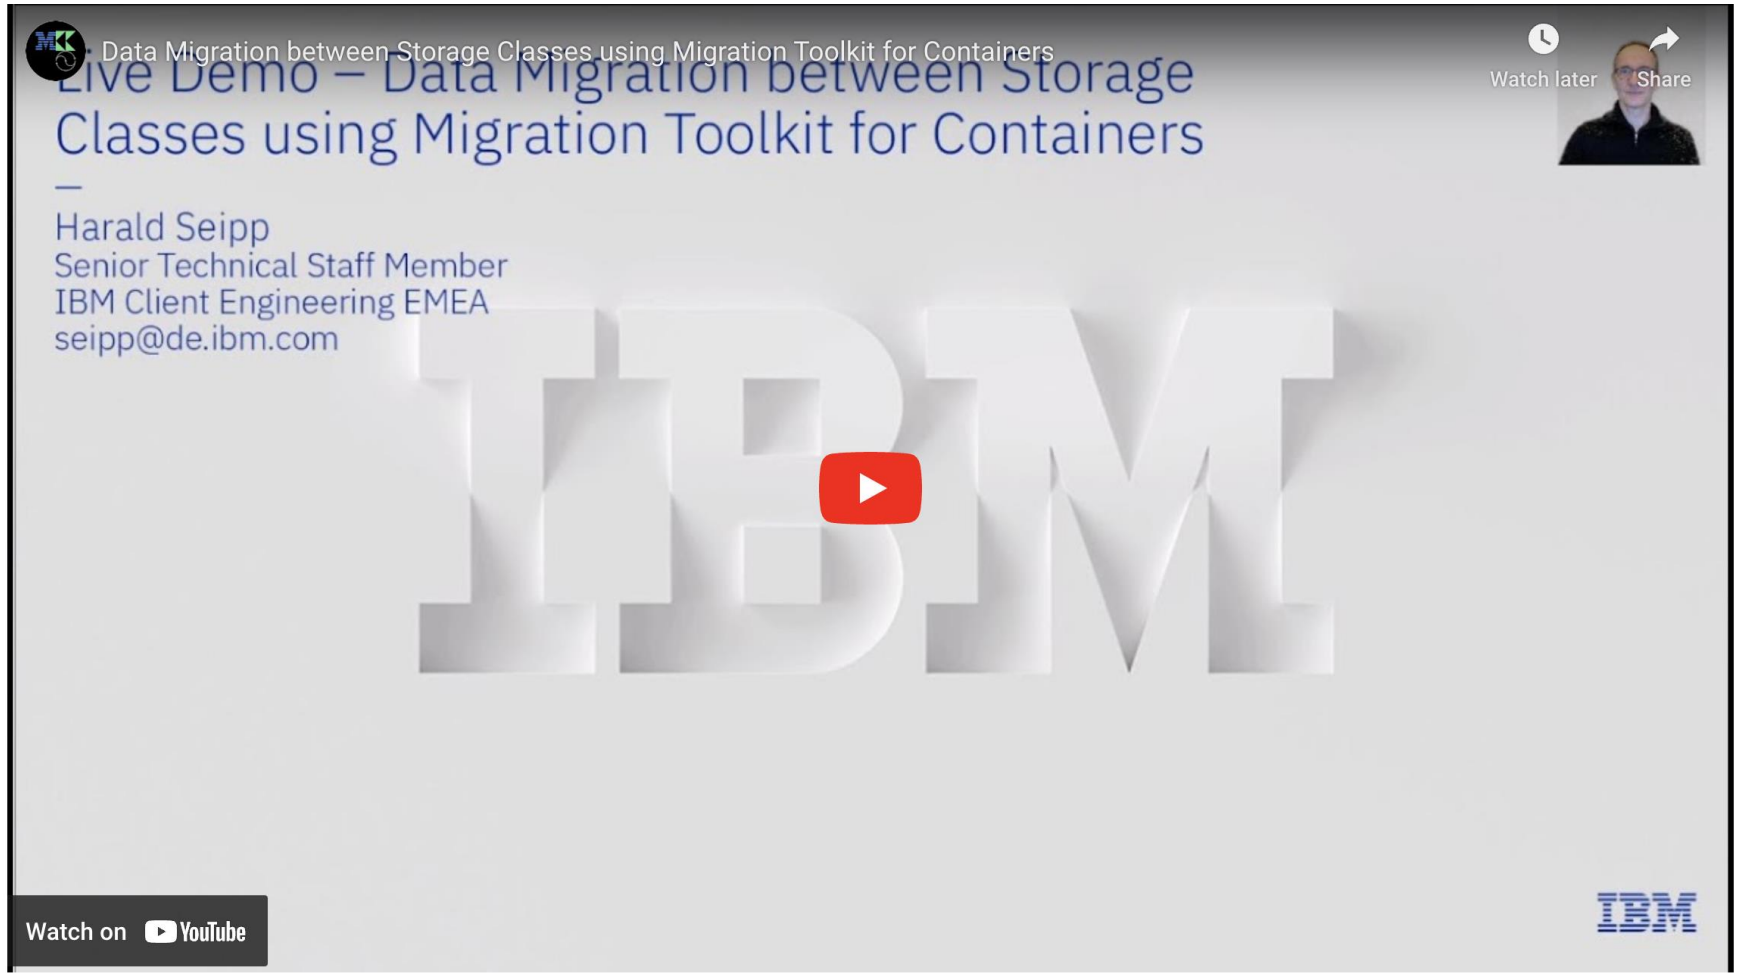

Multi-protocol data access in OpenShift

IBM Storage Scale Days 2024

March 5-7, 2024 | Stuttgart Marriott Hotel Sindelfingen

Harald Seipp, IBM

Overview Diagram

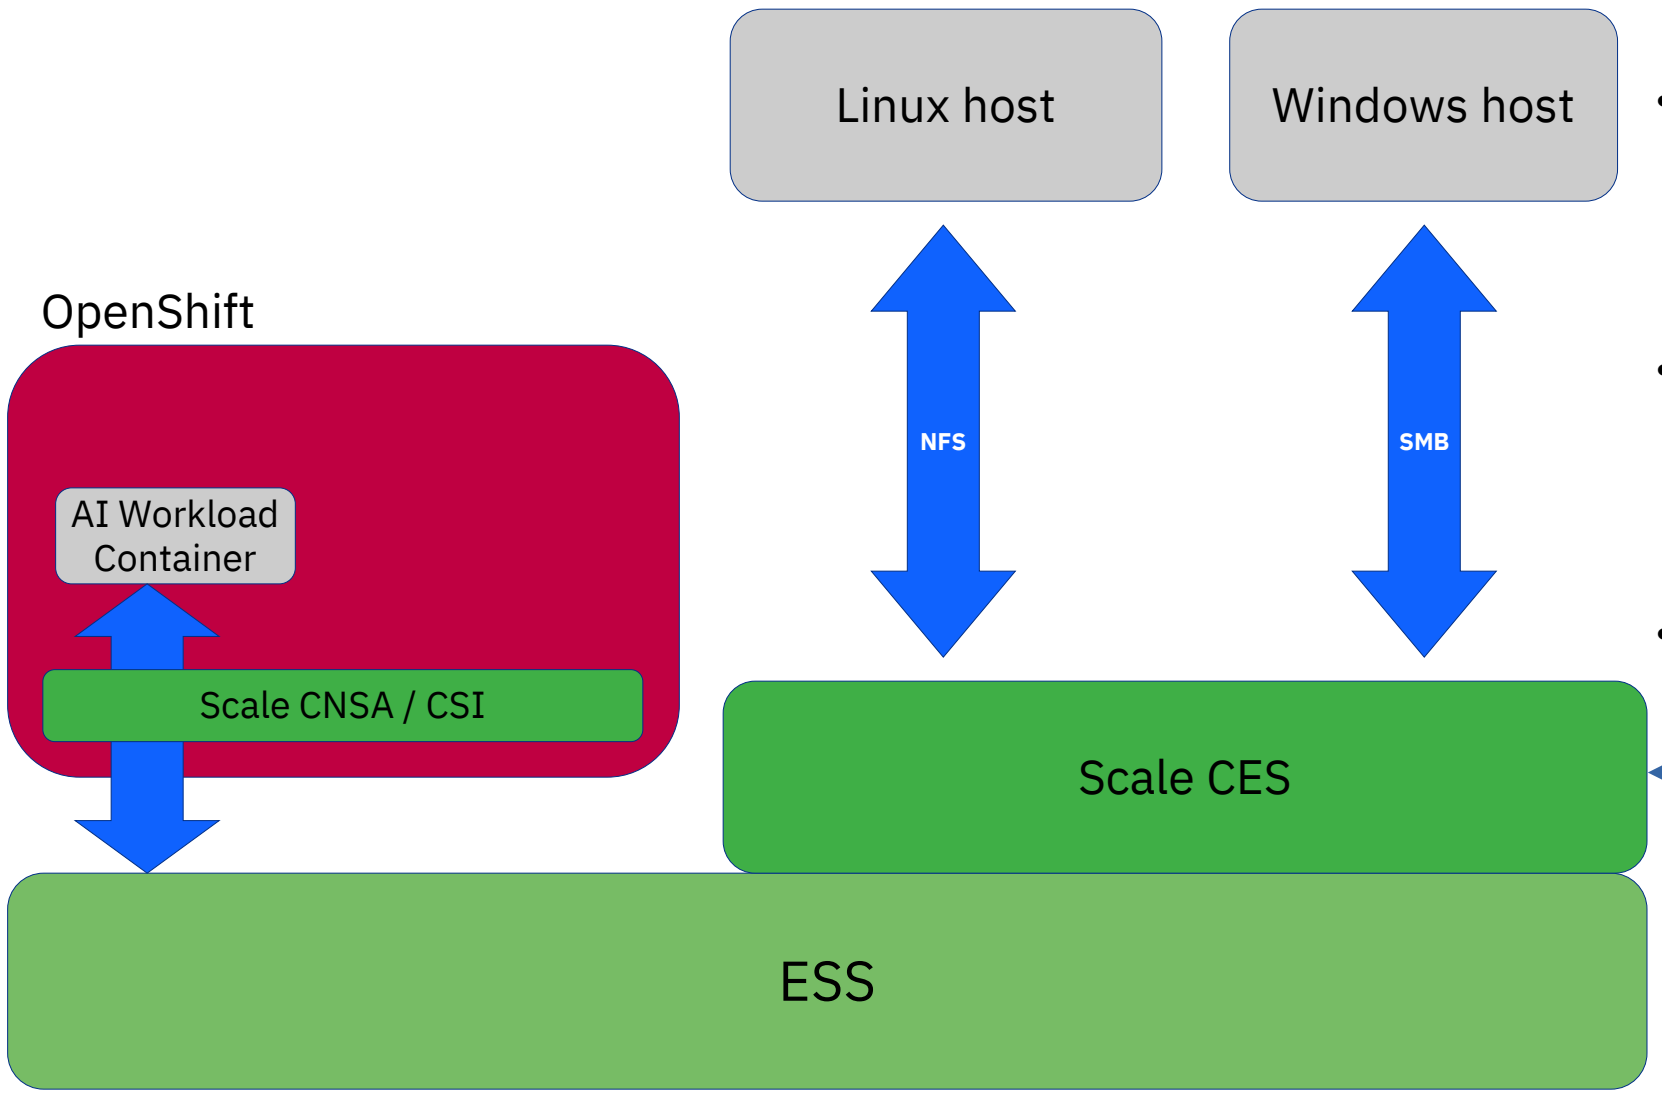

- CES exported SMB+NFS
  - Multiple filesets
- Data is ingested from Linux and Windows hosts
  - Access control via NFSv4 ACLs
  - SID-to-ACL NFSv4 mapping with SMB
  - Inheritance mode
- Data is processed in AI workload container
  - Non-root user within container
  - Results are written back to the Scale file system
- SELinux on ESS/Scale Storage cluster is set to permissive mode

Active Directory

© 2024, IBM. All rights reserved.

Problem Statement

- Define best practice / guidance for container application data access to Storage Scale filesets that contain data ingested by NFS and SMB and are access-controlled through NFSv4 ACL
  - Security / access control is supposed to be as strict as possible
    - Setting an ACL to everyone:write (proven to be working) is not an option
  - The container application must not invalidate the ACLs
    - AD managed using Microsoft Windows is the imperative
- There is no prior art for the given problem, at least not discoverable on the Internet
- The kubelet is not aware of NFSv4 ACL

© 2024, IBM. All rights reserved.

Assumptions

- IBM Storage Scale CNSA 5.1.7+ and IBM Storage Scale CSI 2.9+ are used
  - SELinux issues do no longer apply as the file system is mounted with a container permissive SELinux context with the mount context of the file system set to system_u:object_r:container_file_t:s0 (Documentation)
- PVs are provisioned statically to allow access to existing NFS and SMB filesets

© 2024, IBM. All rights reserved.

Initial attempt on recommendations

- Add access for a specific (service) GID to the ACLs
- Add the `fsGroup` Parameter with the service GID to the accessing Pod (Reference)
- Ensure that the container process(es) are members of the service GID group
- Double-check that the container UID dynamically created by OCP does not overlap with the AD UID range
  - If an overlap is detected, modify the OCP namespace `UidRange` settings (Reference)

© 2024, IBM. All rights reserved.

Why initial recommendations did not work

- fsGroup doesn’t work with statically provisioned PVs
- fsGroup can cause file system tree ownership recursive change (chown)
  - might destroy NFSv4 ACLs
  - fsGroupChangePolicy can reduce the impact, but does not help when ownership change further down the tree is required

© 2024, IBM. All rights reserved.

Solution

- Add access for a specific (service) GID to the ACLs for the given filesets
- Use the `runAsGroup` parameter for the Pods to ensure that the service GID is the primary group (Documentation)
- Ensure that the uid:gid for the process inside the container is not messing with the imperative OCP uid:gid handling
  - Simplest solution: use uid:gid of root:root for the container process
  - This is not a security issue, OCP will enforce non-root container context
- Double-check that the container UID dynamically created by OCP does not overlap with the AD UID range
  - If an overlap is detected, modify the OCP namespace `UidRange` settings (Reference)

© 2024, IBM. All rights reserved.

References

- Advanced Static Volume Provisioning with IBM Storage Scale for OpenShift: https://community.ibm.com/community/user/storage/blogs/gero-schmidt1/2022/04/01/advanced-static-volume-provisioning-on-ocp
- IBM Storage Scale CSI Driver documentation: https://www.ibm.com/docs/en/spectrum-scale-csi?topic=spectrum-scale-container-storage-interface-driver-29
- OpenShift documentation for SCC: https://docs.openshift.com/container-platform/4.12/authentication/managing-security-context-constraints.html
- OpenShift documentation for fsGroup: https://docs.openshift.com/container-platform/4.12/storage/understanding-persistent-storage.html#using_fsGroup_understanding-persistent-storage

© 2024, IBM. All rights reserved.

Thank you for using

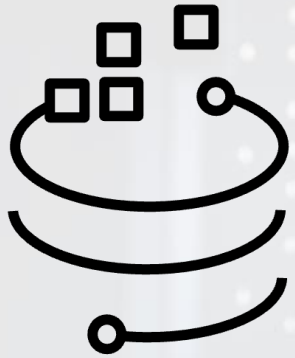

Storage Scale

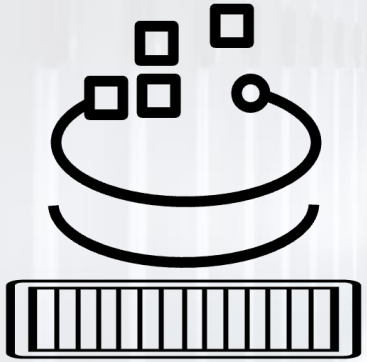

Storage Scale System