

Getting ready for OpenShift migration

Robert Bohne

Sr. Specialist Solution Architect | OpenShift

robert.bohne@redhat.com

Twitter: @Robert Bohne

Sebastian Dehn

Solution Architect Partner Enablement

sdehn@redhat.com

Twitter: @Sifa171

Speaker



Sebastian Dehn
Solution Architect
Partner Enablement
Twitter: @sifa171



Robert Bohne
Sr. Specialist Solution Architect
OpenShift
Twitter: @RobertBohne

Main reason

OpenShift **4.x lacks an in-place upgrade** from existing OpenShift 3.x installs

Customers are required to stand up a new 4.x cluster and bring the workloads they want to the new cluster

Why did we choose this migration strategy?

Reducing risk


A ton of innovation went into OpenShift 4, and an in-place upgrade would have risk of failure in which there is no forwards or backwards remediation. It allows you to skip from 3.7/3.9/3.10/3.11 to 4.x. Skipping the need to install each one.

Useful for 4-to-4 migrations

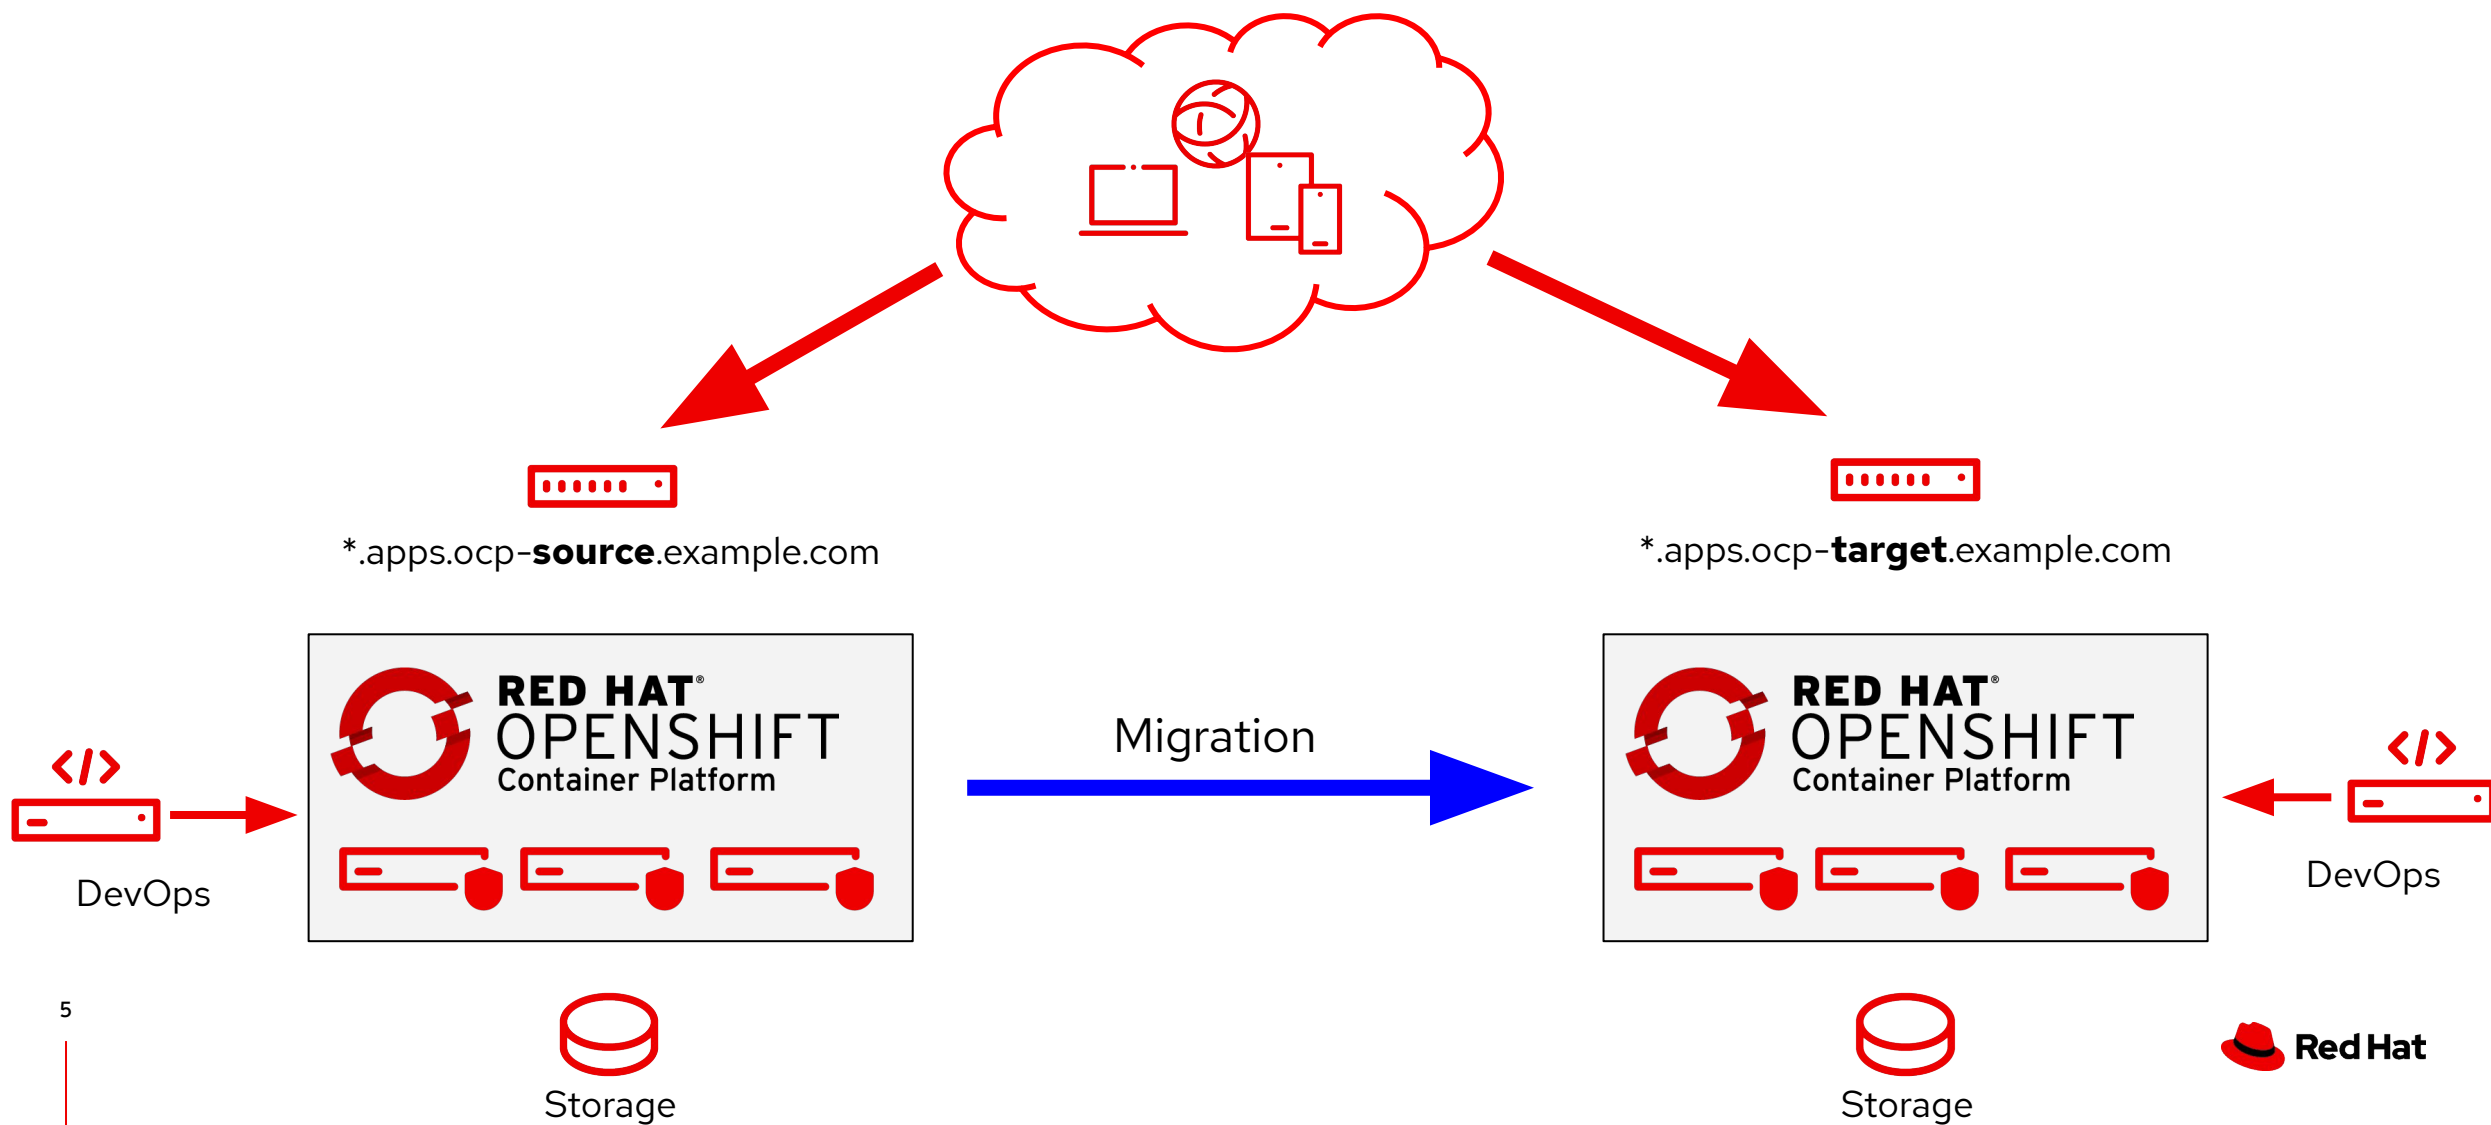
A general migration tool is frequently requested and a better long term investment. Helps you build a foundation towards making your cluster investments less fragile.

Allows for staging

Stage a mock migration before doing it live, on a Project by Project basis. Extremely useful for success.



The scenario



**What we
have to
consider?**



Migration considerations

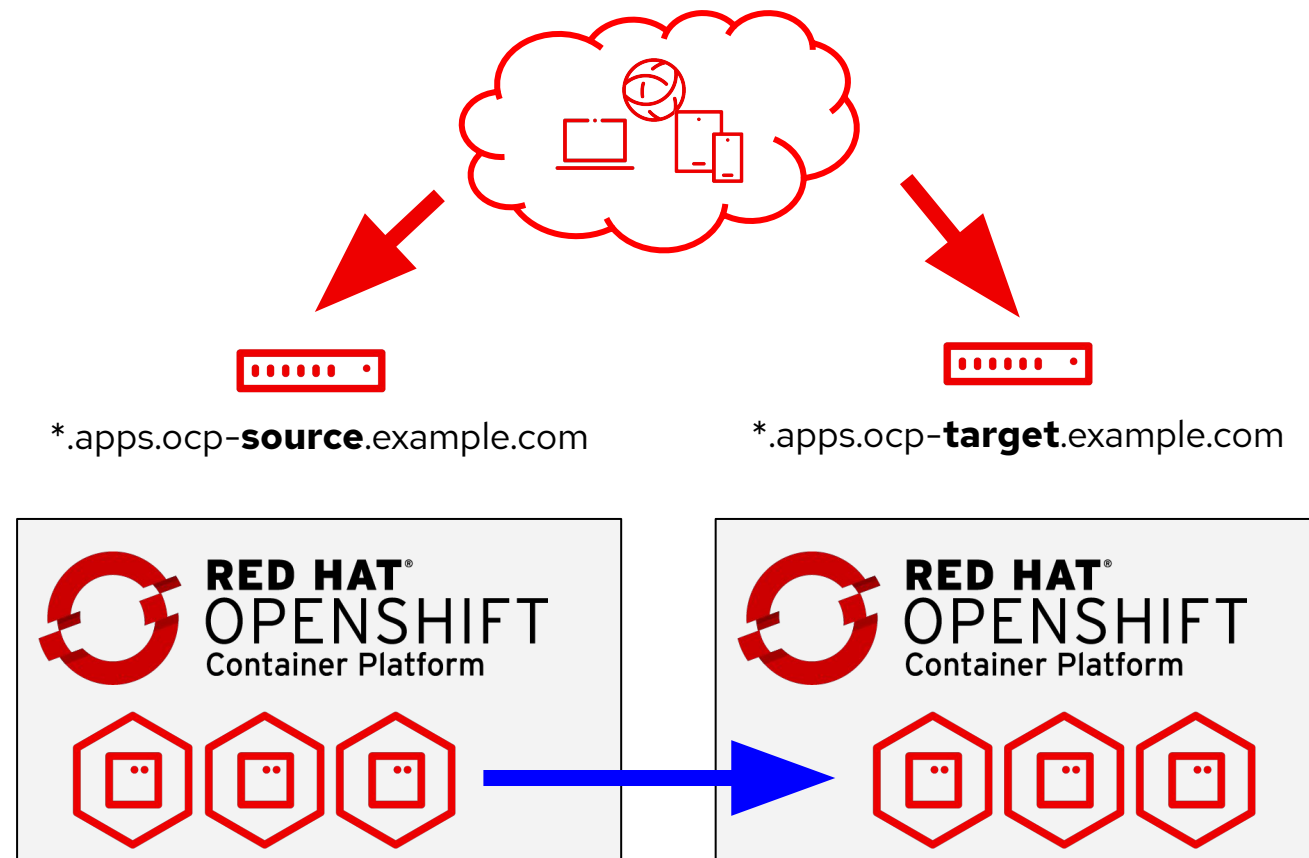
Cluster Migration

- ▶ Storage (Local volume, FlexVolume, CSI, OCS)
- ▶ Networking (multitenant, network policies, Encrypting traffic between hosts)
- ▶ Logging (no log data migration!)
- ▶ Monitoring (DeadMansSwitch vs Watchdog)
- ▶ Operational concept
- ▶ Integration in third party company tools (Ticket system, ServiceNow,...)
- ▶ Same IdM authentication service setup between the two clusters
- ▶ **Amount of migrations: rare!**

Migration considerations


Application Migration

- ▶ ImageStream availability (old versions removed)
- ▶ Storage (Matching storage solutions)
- ▶ **How to handle incoming traffic!**
- ▶ Build (docker vs buildah, rhel build)
"OOMKill because of quota limits"
- ▶ **Amount of migrations: very often!**




Migration considerations

OpenShift v3

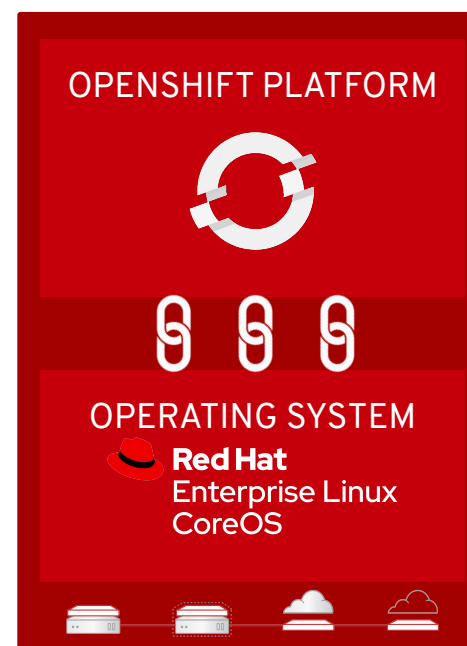


- ▶ Red Hat Enterprise Linux 7
- ▶ Container Runtime: Docker
- ▶ Ansible based installation
- ▶ Container **build** via **Docker** Daemon
external not calculated to container **resources**!



INFRASTRUCTURE

OpenShift v4



- ▶ Red Hat Enterprise Linux CoreOS!
- ▶ Container Runtime: CRI-O, Buildah, Skopeo
- ▶ Operator based installation
- ▶ Container build via **buildah** inside of a container, need **more resource** because **inside of the container**!
- ▶ Build JenkinsPipeline deprecated!

Tools for assistance

Application Workloads

Cluster Application Migration (CAM)

Migrates stateful/stateless applications from source cluster to destination cluster.

Intended initially for addressing OCP 3.7+ to OCP 4.2+ upgrade scenarios

Control Plane Migration Assistance (CPMA)

Helps configure a OCP 4.x cluster to match, **when possible**, settings in an OCP 3.x cluster.

Not all 3.x settings are available in 4.x


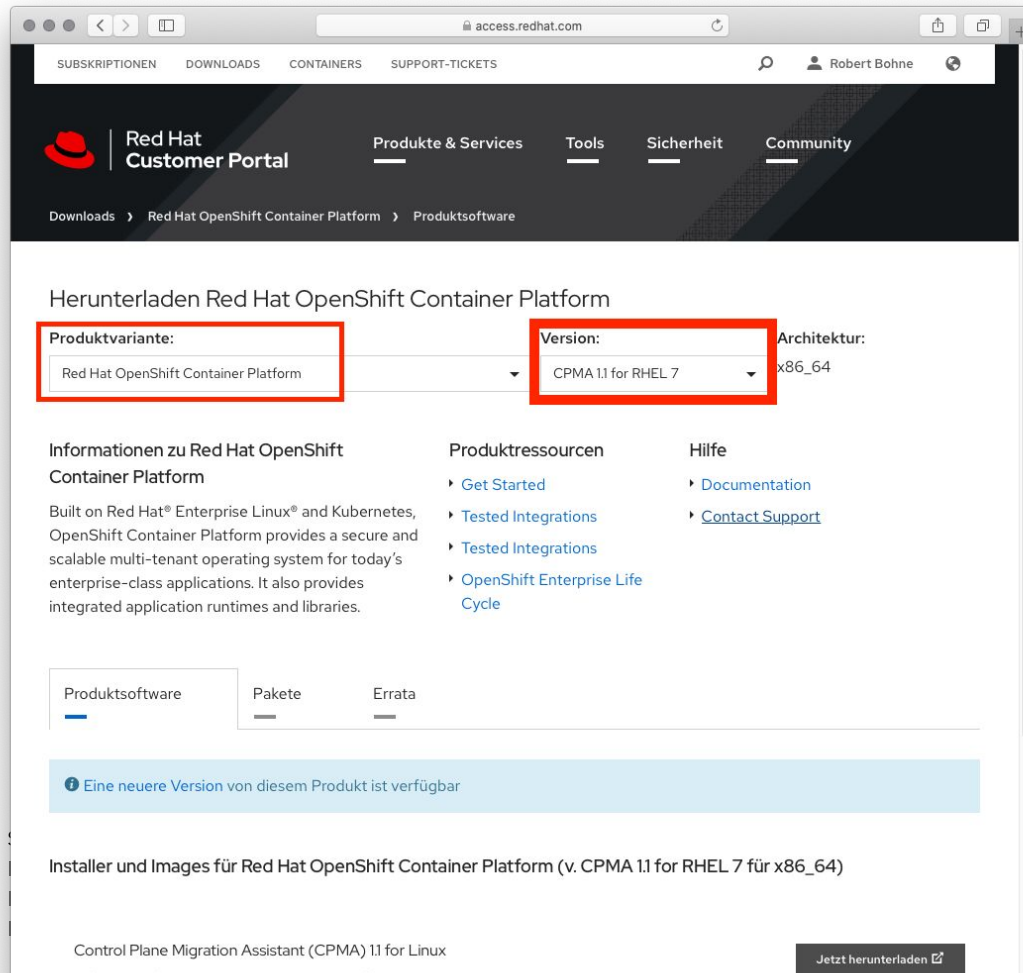
Control Plane Configuration for 4.x

Control Plane Migration Assistance (CPMA)

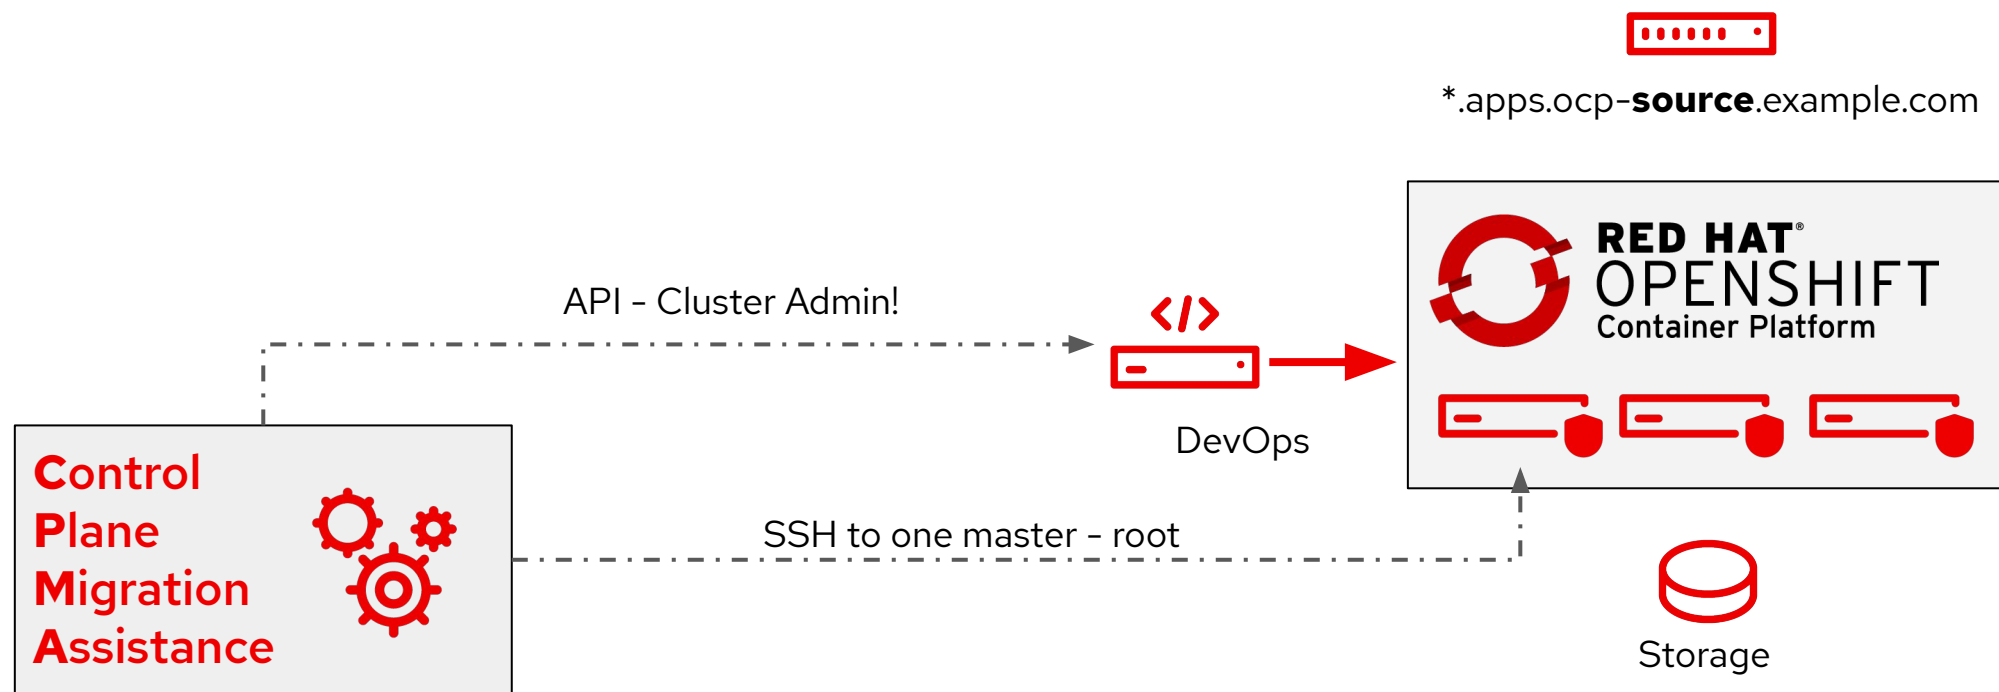
Used version: 1.1.1

Its a command line tool

Download it and run it



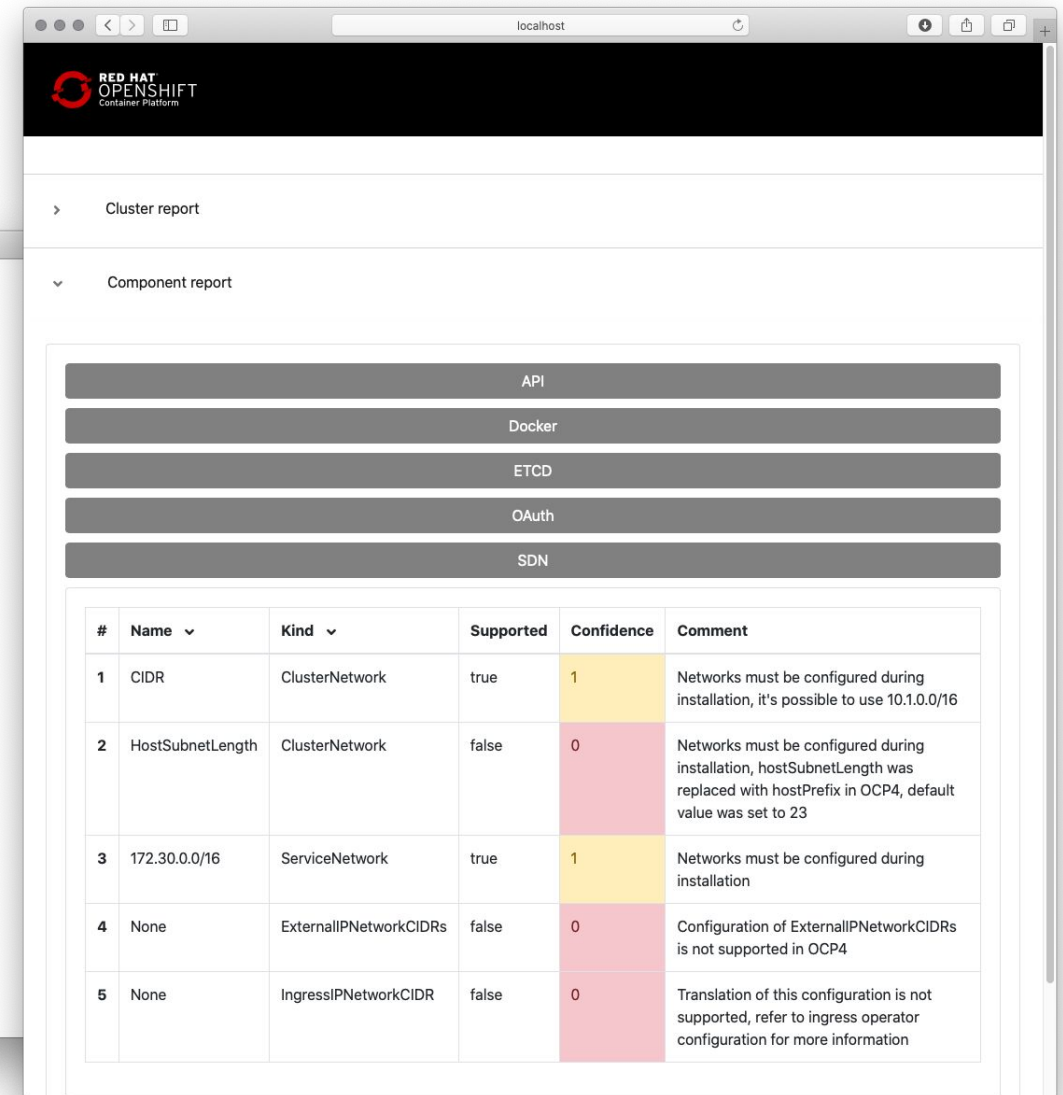
During execution



The output

```
$ find .
.
./report.json
./ocplmaster1.stormshift.coe.muc.redhat.com
./ocplmaster1.stormshift.coe.muc.redhat.com/etc
./ocplmaster1.stormshift.coe.muc.redhat.com/etc/etcd
./ocplmaster1.stormshift.coe.muc.redhat.com/etc/etcd/etcd.conf
./ocplmaster1.stormshift.coe.muc.redhat.com/etc/containers
./ocplmaster1.stormshift.coe.muc.redhat.com/etc/containers/registries.conf
./ocplmaster1.stormshift.coe.muc.redhat.com/etc/origin
./ocplmaster1.stormshift.coe.muc.redhat.com/etc/origin/master
./ocplmaster1.stormshift.coe.muc.redhat.com/etc/origin/master/master.server.key
./ocplmaster1.stormshift.coe.muc.redhat.com/etc/origin/master/master-config.yaml
./ocplmaster1.stormshift.coe.muc.redhat.com/etc/origin/master/htpasswd
./ocplmaster1.stormshift.coe.muc.redhat.com/etc/origin/master/master.server.crt
./cpma.log
./manifests
./manifests/100_CPMA-cluster-config-secret-htpasswd-secret.yaml
./manifests/100_CPMA-cluster-config-sdn.yaml
./manifests/100_CPMA-cluster-config-oauth.yaml
./manifests/100_CPMA-cluster-config-project.yaml
./manifests/100_CPMA-cluster-config-secret-google-secret.yaml
./manifests/100_CPMA-cluster-config-image.yaml
./report.html
$
```

CR(s) to apply



#	Name	Kind	Supported	Confidence	Comment
1	CIDR	ClusterNetwork	true	1	Networks must be configured during installation, it's possible to use 10.1.0.0/16
2	HostSubnetLength	ClusterNetwork	false	0	Networks must be configured during installation, hostSubnetLength was replaced with hostPrefix in OCP4, default value was set to 23
3	172.30.0.0/16	ServiceNetwork	true	1	Networks must be configured during installation
4	None	ExternalIPNetworkCIDRs	false	0	Configuration of ExternalIPNetworkCIDRs is not supported in OCP4
5	None	IngressIPNetworkCIDR	false	0	Translation of this configuration is not supported, refer to ingress operator configuration for more information

Report in HTML & JSON

Things to consider



- ▶ Not all 3.x settings are available in 4.x.
 - Search registries
 - There is no Docker anymore
- ▶ Networks must be configured during installation
- ▶ The API Port for Openshift 4 is 6443 and is non-configurable. Your OCP 3 cluster is currently configured to use port 8443
- ▶ Authentication via manifests/100_CPMA-cluster-config-oauth.yaml
 - Missing ClusterRoleBindings!

Check the report generated by CPMA!

Cluster Application Migration (CAM)

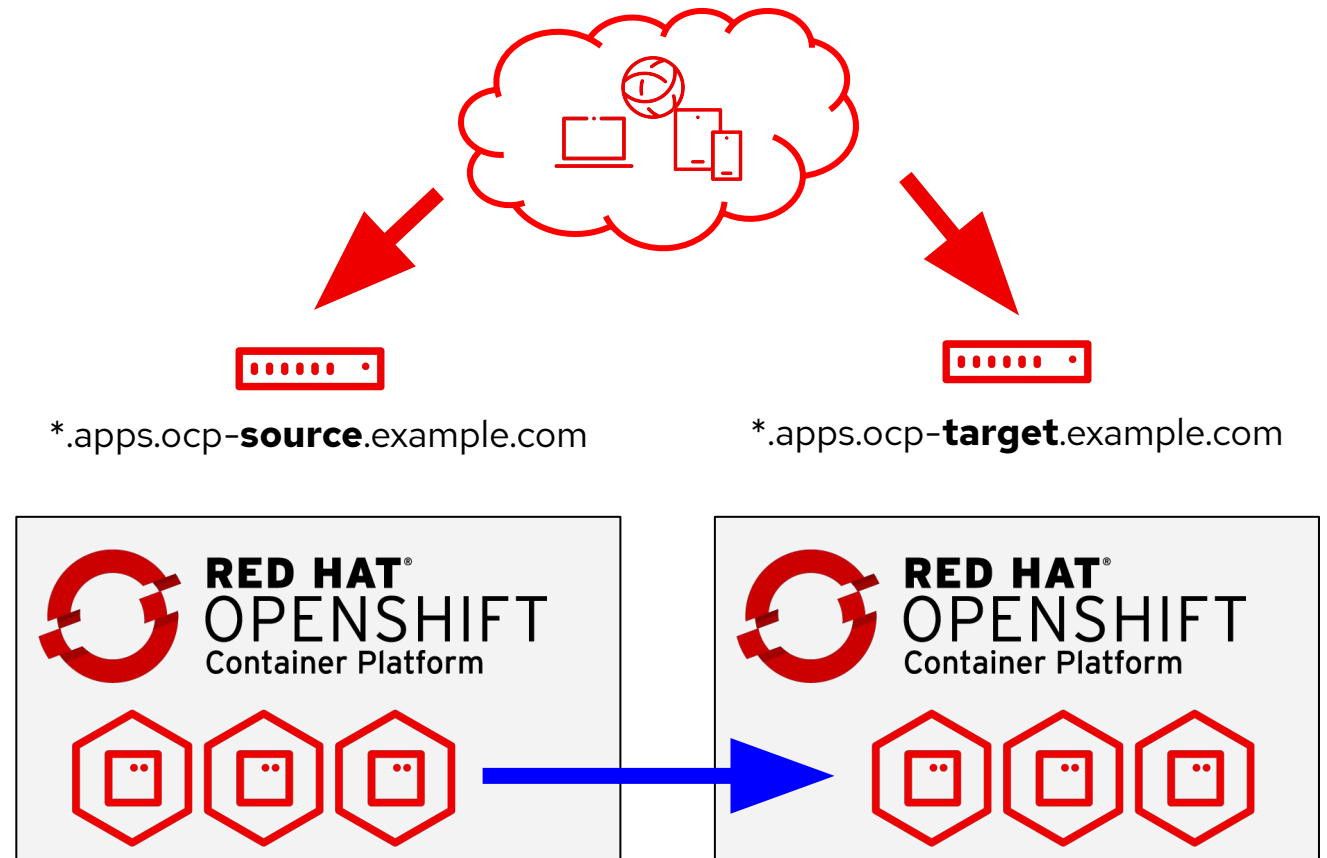
Used version: 1.2

Demo



Migrate workloads between clusters

- ▶ Includes
 - k8s/OCP Object Resources
 - Persistent Volumes
 - Internal Images



Migration Process

Plan

1. Select Source Cluster
2. Select Namespaces
3. Choose Copy or Move for each PV
4. Specify Destination

Stage



Stages the data from Source to Target.

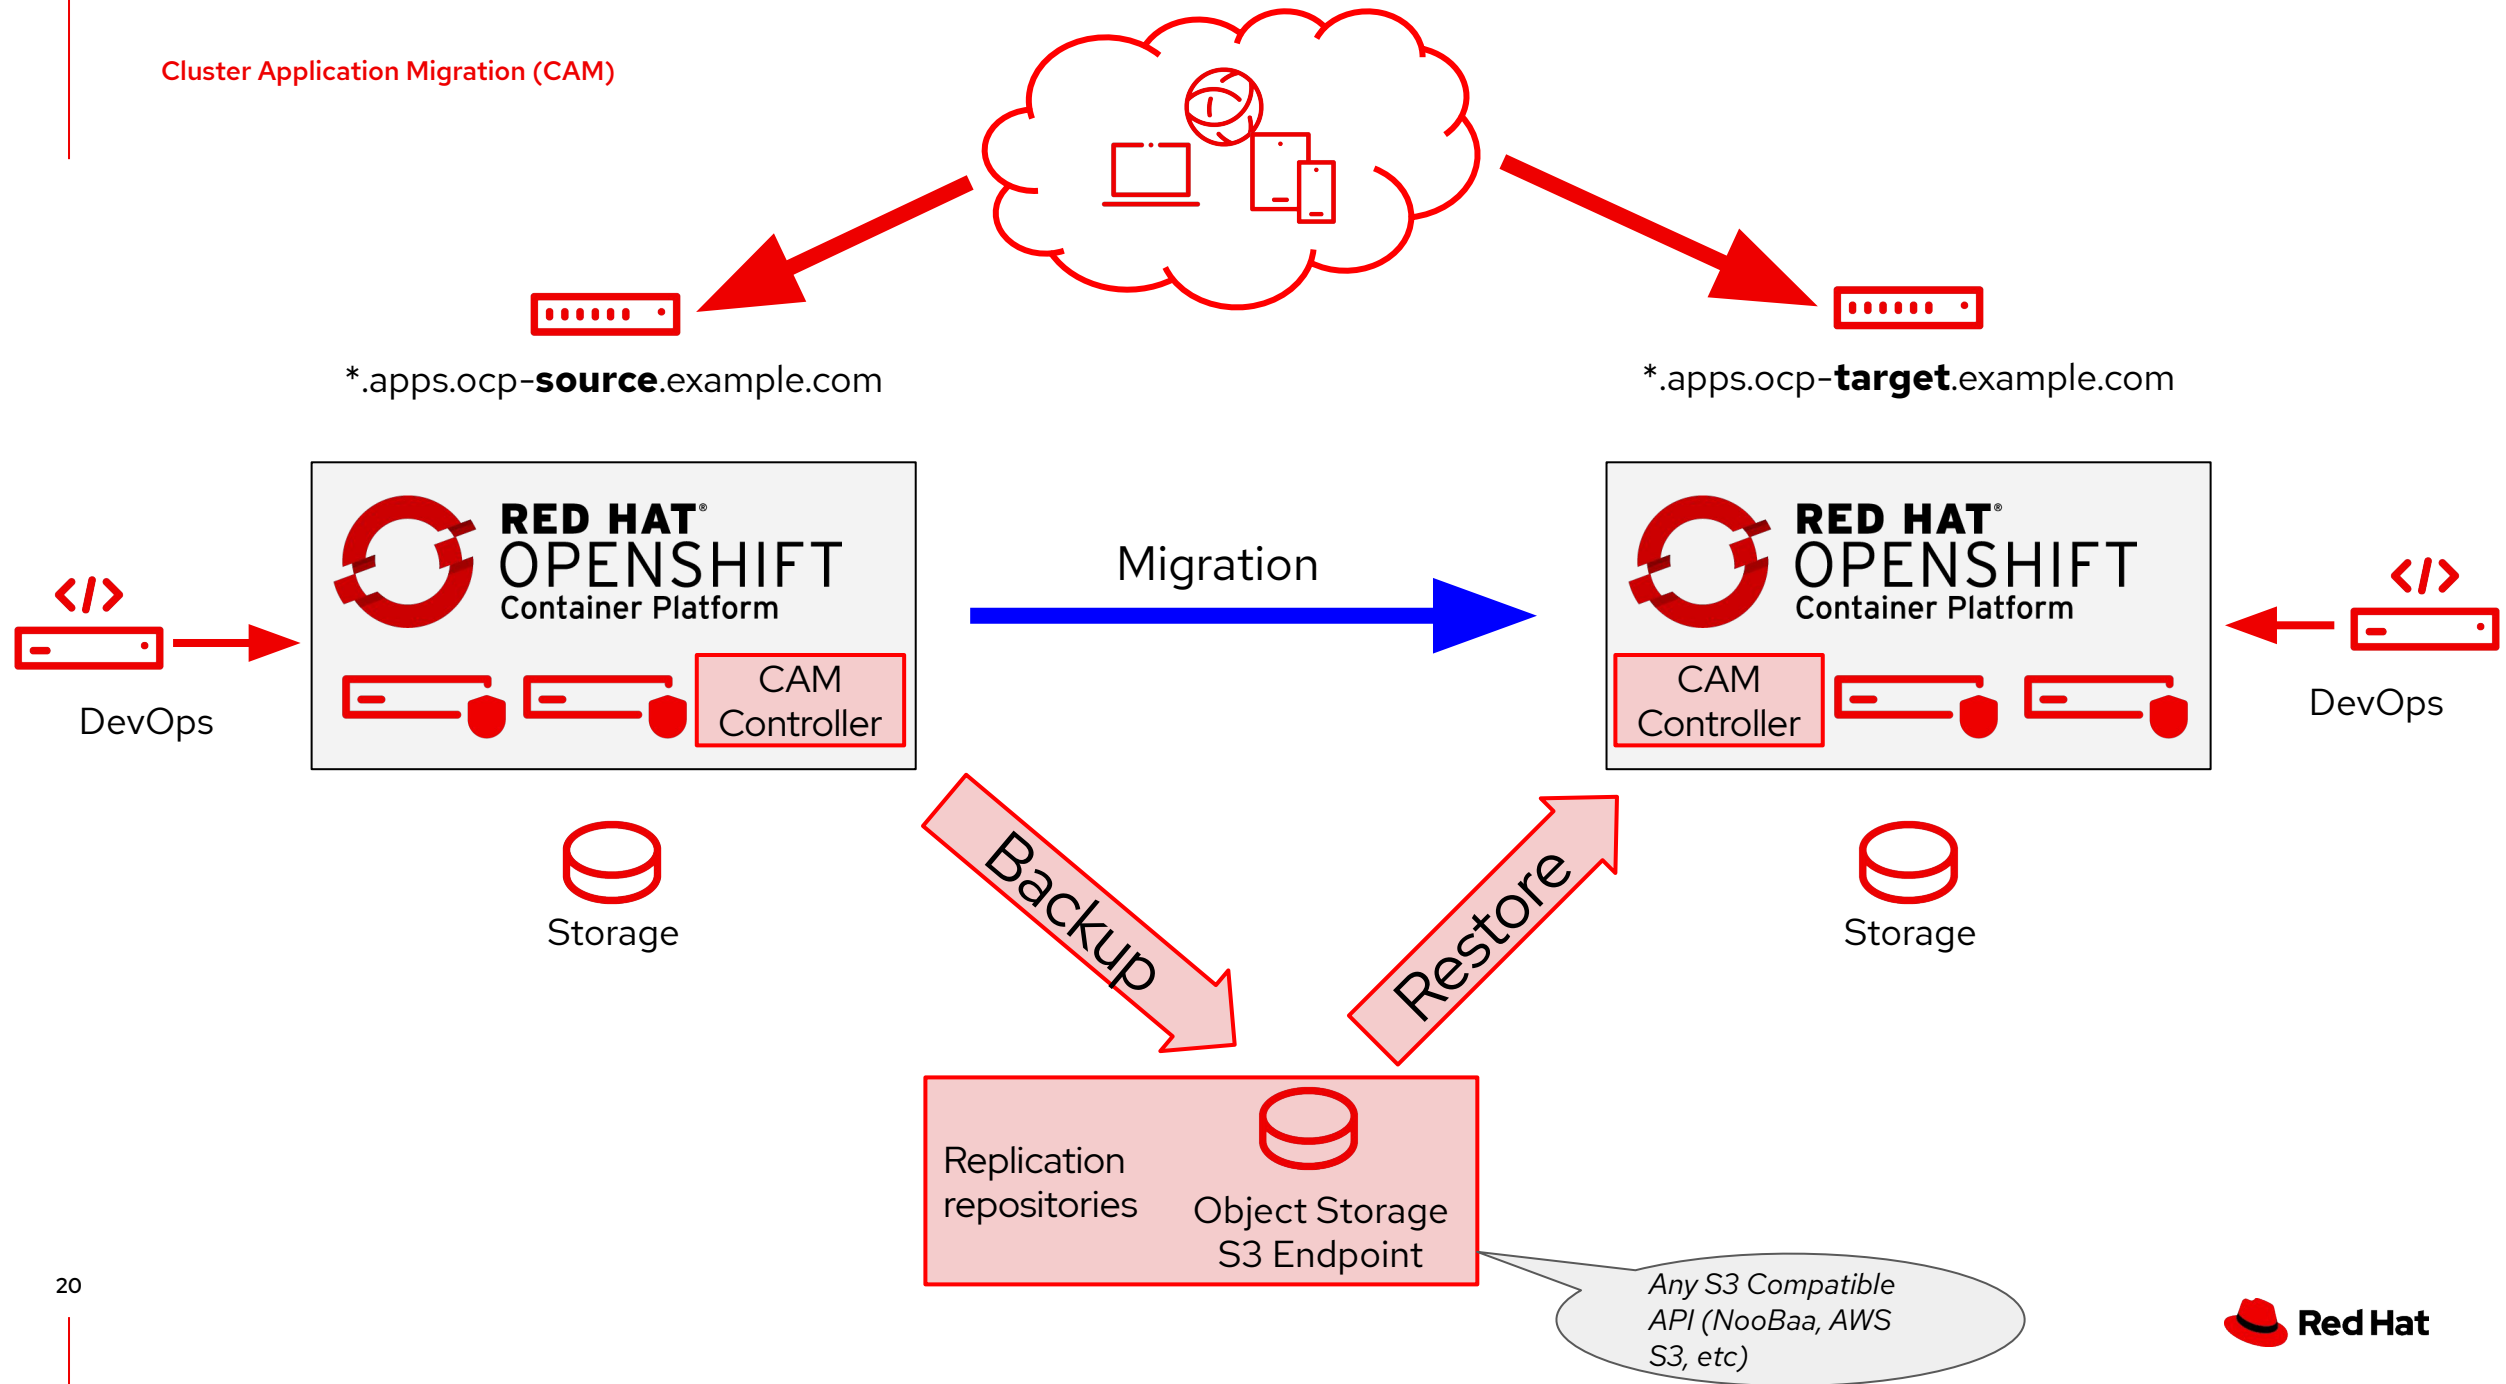
May be run multiple times.

Applications are running
no Downtime during step

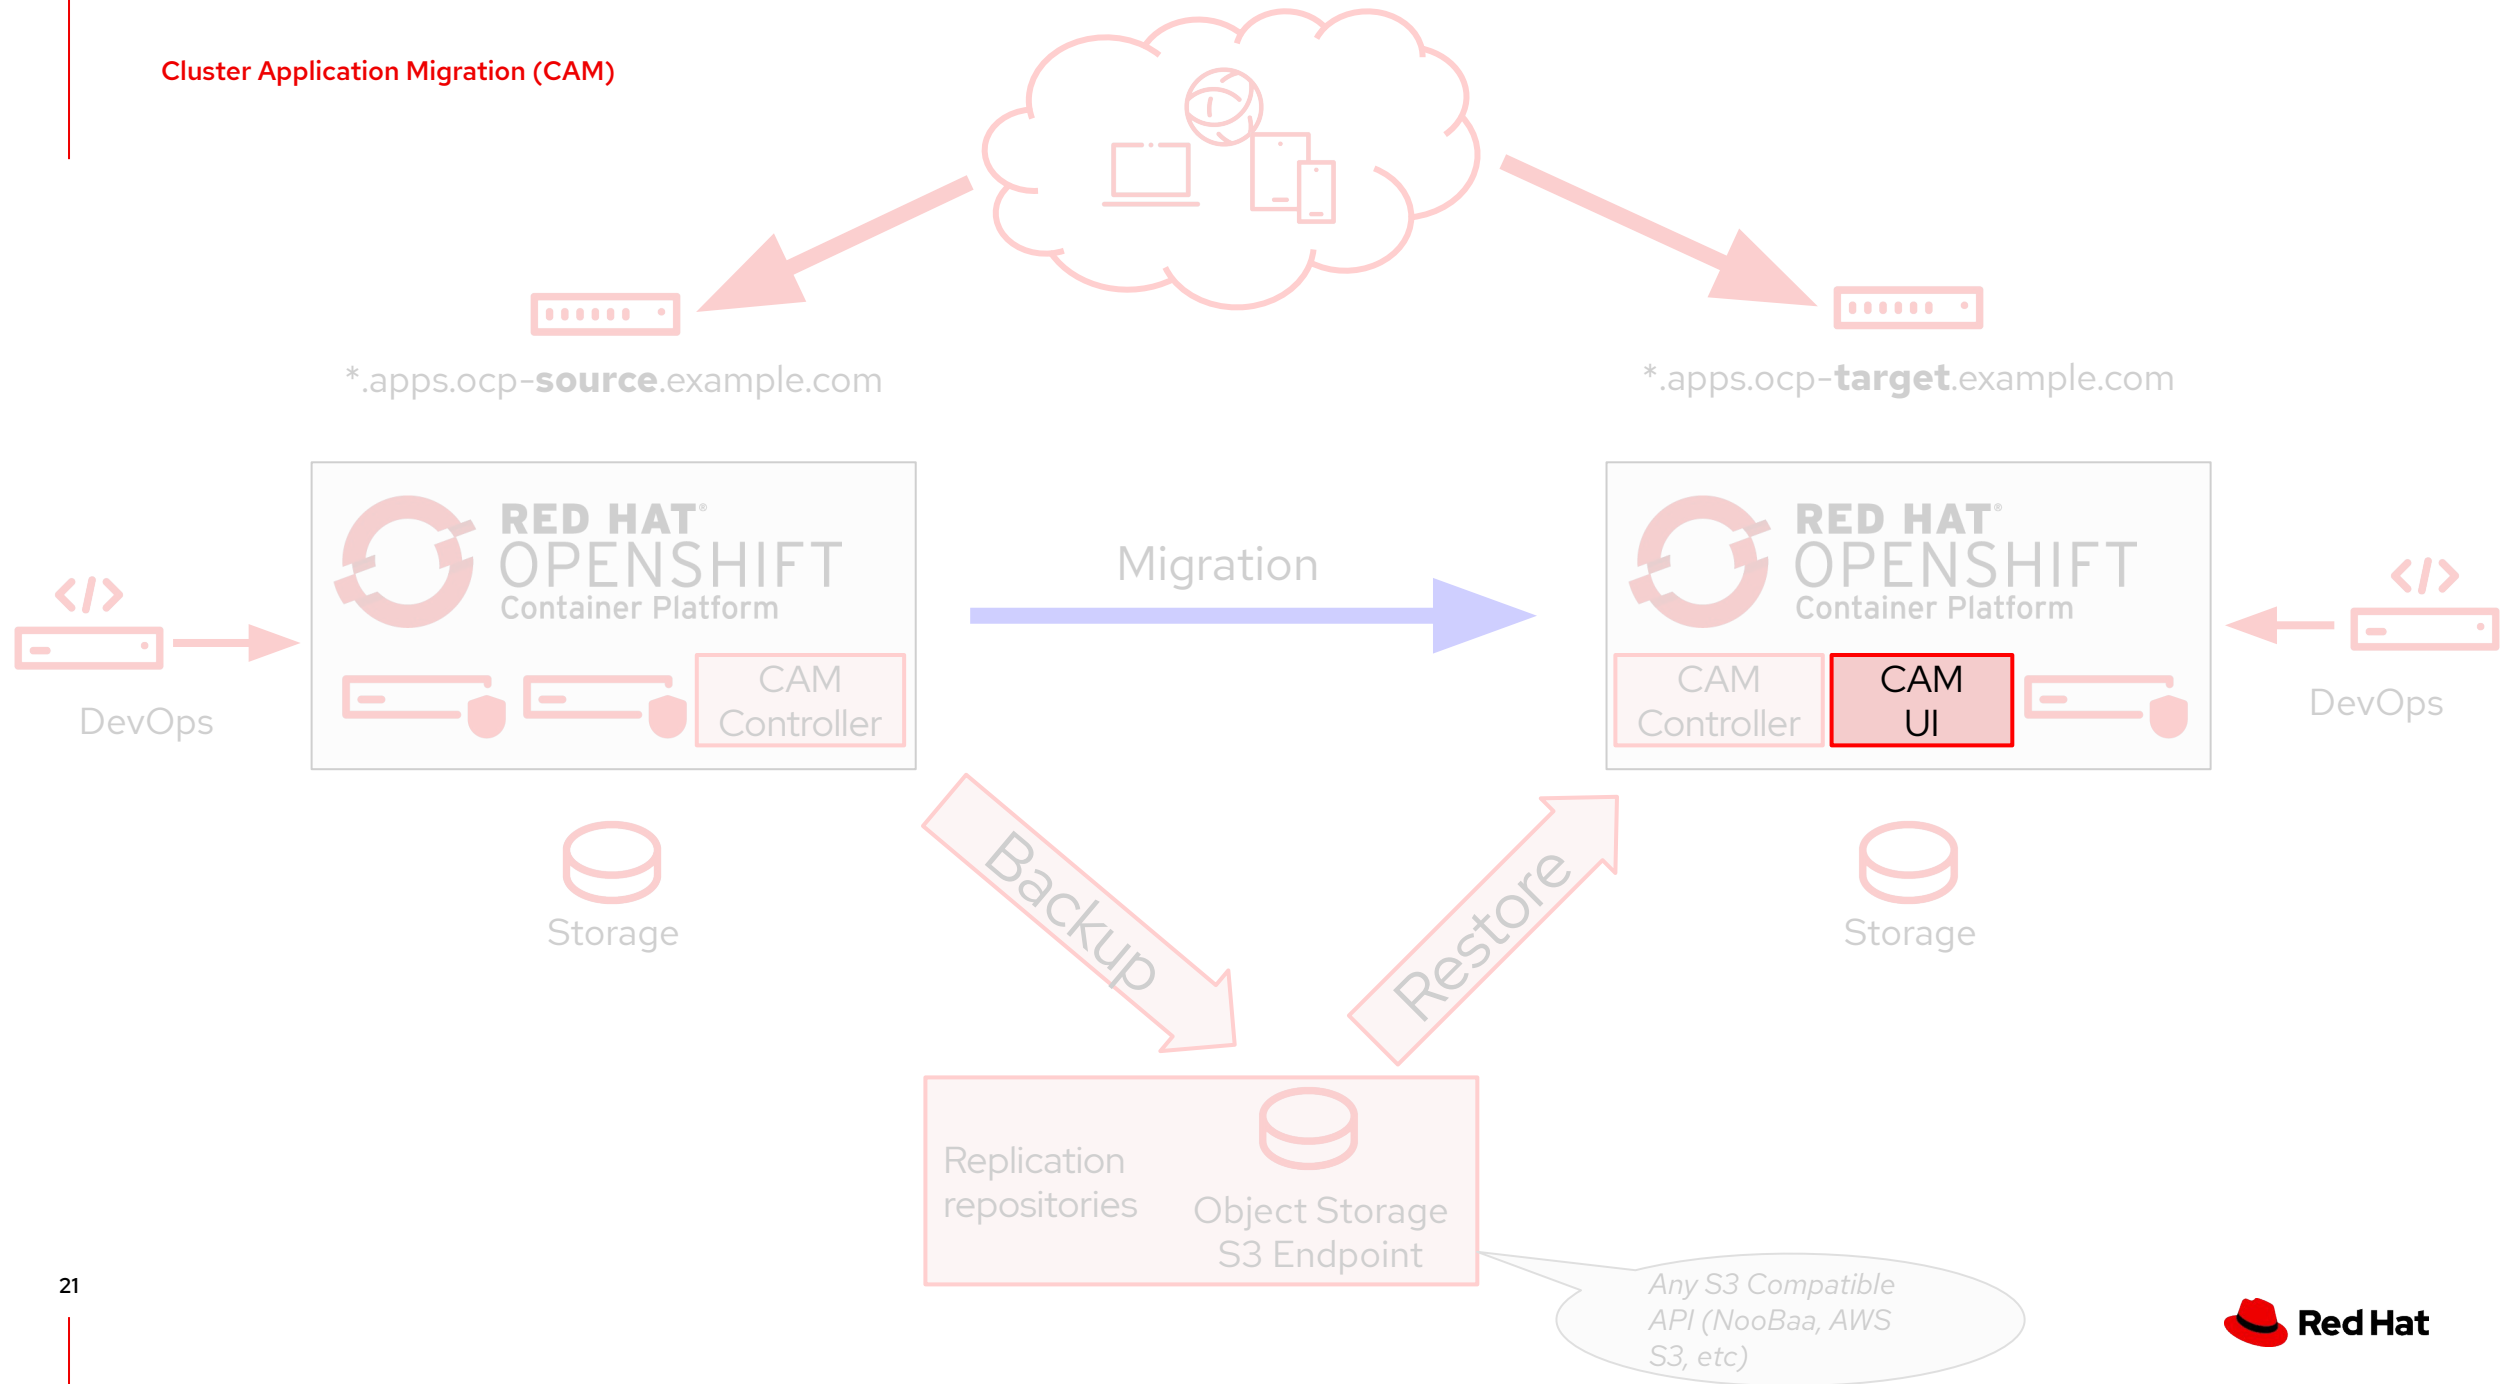
Migrate

Quiesce Application

Migrate any delta bits not
captured in stage.



Cluster Application Migration (CAM)



Supported Use Cases

As a Cluster Admin, I want to ...

1. Migrate my stateless workload
2. Migrate my stateful workload consisting of:
 - a. Source of NFS PVs to Destination of NFS PVs
 - b. Source of EBS PVs to Destination of EBS PVs
 - c. Source of Gluster PVs to Destination of Ceph PVs
 - d. Source of "X" PVs to Destination of "Y" PVs
3. Do the above and also..
 - a. preserve existing internal images in the embedded openshift registry onto the destination cluster



It's OK to mix/match source to destination.

*i.e.
Source (NFS) to Destination (EBS)*

Will allow change of desired Storage Class on Destination, then copy filesystem contents between PVs


Supported Migration Paths:

- ▶ **OCP 3.7+ -> OCP 4.x+ (Current Focus)**
- ▶ OCP 3.7+ -> OCP 3.7+
- ▶ OCP 4.x -> OCP 4.x
- ▶ OCP 4.x -> OCP 3.7+ (*Not recommended, but no technical blocker if desired*)

Persistent Volumes


Move or Swing the PV

- ▶ For a remote disk share (NFS for example) we can recreate the PVC/PV definitions from source cluster to destination cluster



Copy


- ▶ Create a copy of the data from source cluster into the destination cluster.
 - Involves creating a PVC on destination and allowing cluster to find a PV to bind to the claim. We then copy data into the PV.



Persistent Volumes

Snapshot


- ▶ Following the rules for where a snapshot can reach, works for....
 - Azure to Azure Block
 - AWS to AWS EBS
 - GCP to GCP pDisk



Persistent Volumes - Stage and Migrate


Data in a Persistent Volume to be Migrated

Stage



*Stage is run and migrates bulk of data to destination.
No downtime during this phase. For large data sets this may be
expensive and take many hours*

Migrate



*Migrate is run, benefits from incremental backup/restore from last Stage.
Minimize downtime during this phase when Quiescing application*

Downtime?

- ▶ Migrating state can be expensive
 - Downtime is expected for stateful applications
- ▶ Approach to minimize downtime
 - Minimize the time of actual 'migration' by taking advantage of pre-seeding as much data as possible in a prior step.
 - **Stage** - seeds data while leaving application up and running
 - **Migrate** - Quiesces application and migrates deltas from stage
 - Generic approach to work for all application types
 - Later allow more efficient methods of custom Quiesce, read-only modes and the like

Granularity of Namespace(s) & 'cluster-admin' required

- ▶ **Migration is at scope of a Namespace.**
 - *Future will allow selecting resources inside of a Namespace*
- ▶ **Cluster Scoped Resources are not handled**
 - Cluster Role Bindings, SCCs, etc are not handled with migration.
 - Expectation is that cluster admin handles cluster scoped resources ahead of running a Migration.
- ▶ **'cluster-admin' required for initial release targeting OCP 4.2**
 - Future plans to allow end user to migrate what they own post OCP 4.2+

What's under the hood?



Leverages Upstream Projects



VELERO

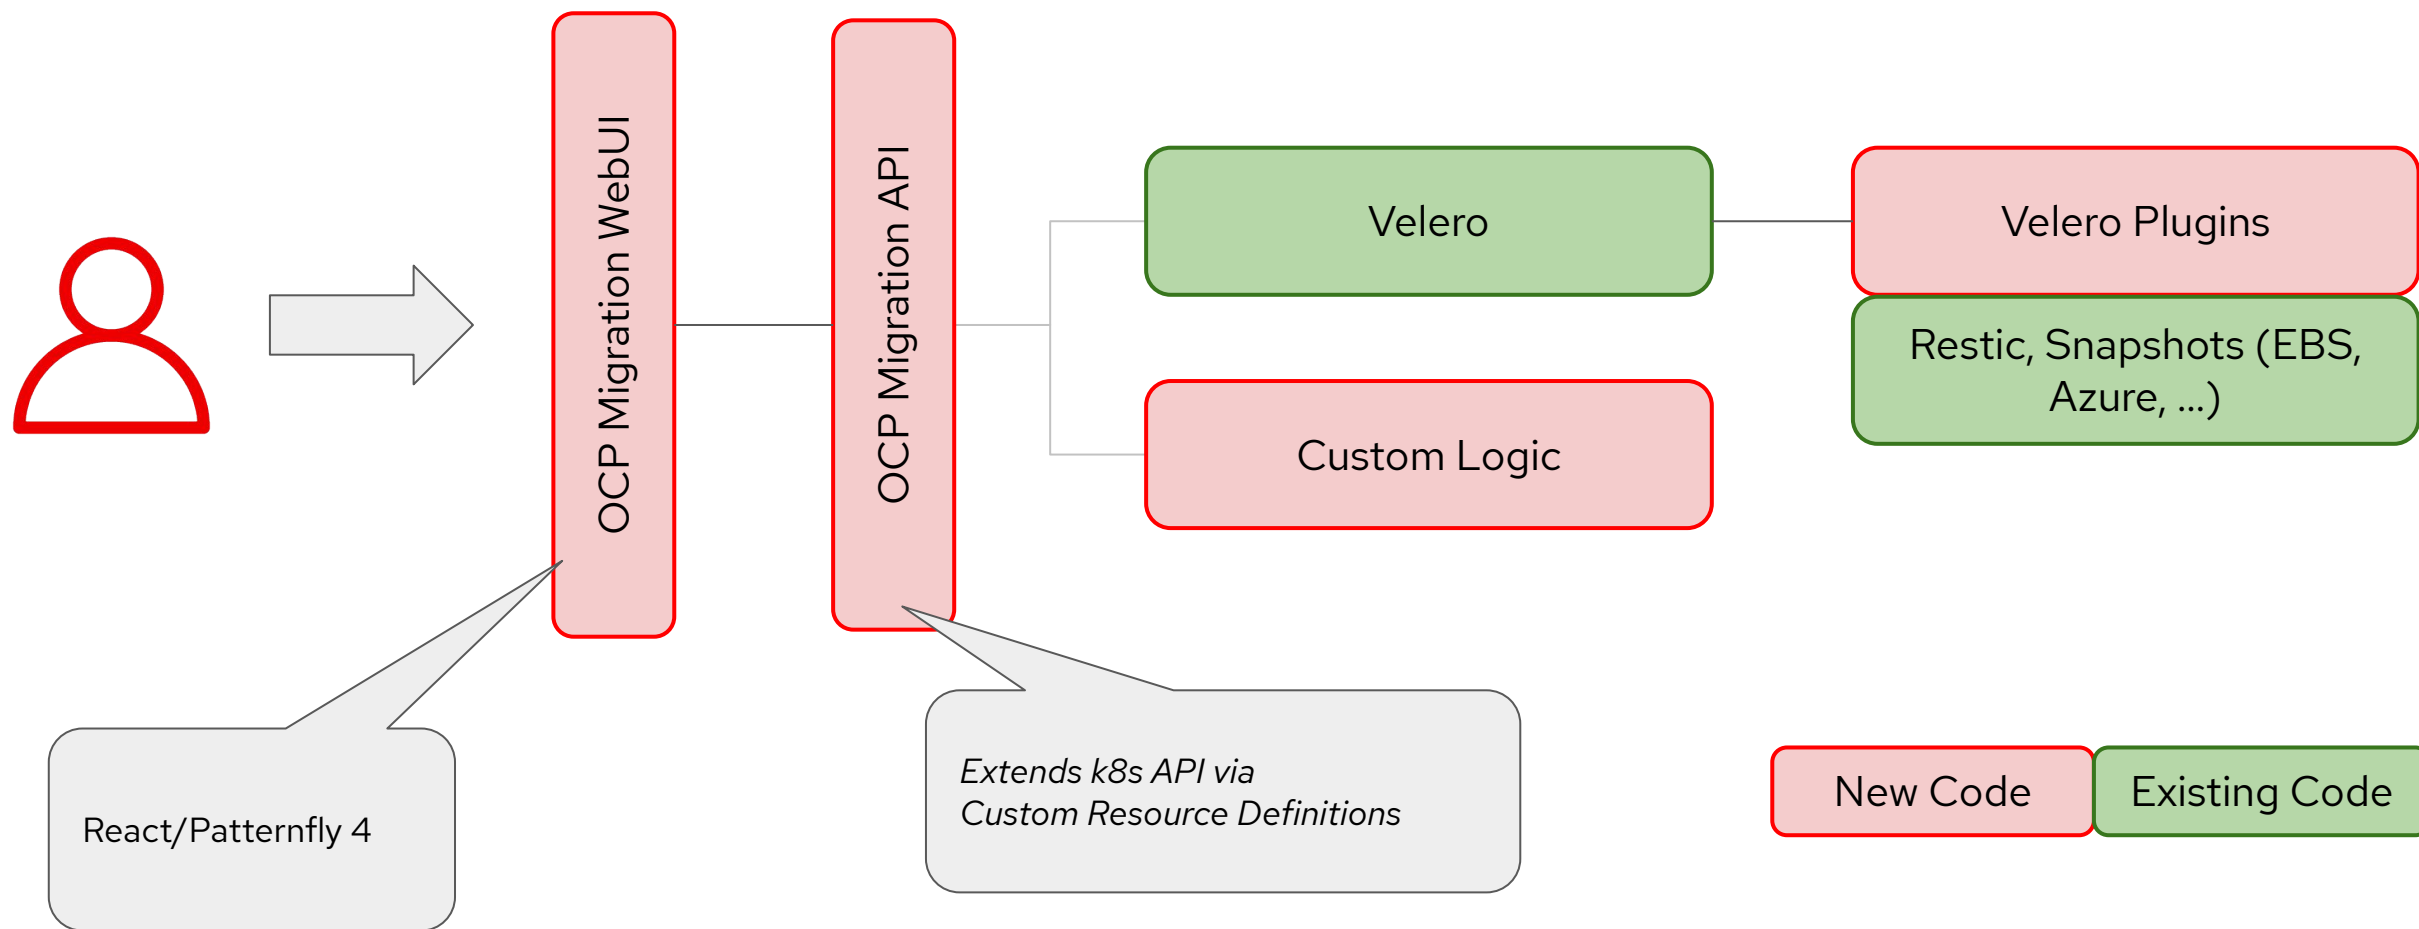
"Backup and migrate Kubernetes applications and their persistent volumes"

Restic

"Fast, secure, efficient backup program"



Architecture & More



Architecture & More

```
$ oc get crd | grep -E '(migration|velero)'
backups.velero.io
backupstoragelocations.velero.io
deletebackuprequests.velero.io
downloadrequests.velero.io
migclusters.migration.openshift.io
migmigrations.migration.openshift.io
migplans.migration.openshift.io
migrationcontrollers.migration.openshift.io
migstorages.migration.openshift.io
podvolumebackups.velero.io
podvolumerestores.velero.io
resticrepositories.velero.io
restores.velero.io
schedules.velero.io
serverstatusrequests.velero.io
volumesnapshotlocations.velero.io
$
```

2020
2020
2020
2020
2020
2020
2020
2020
2020
2020
2020
2020
2020
2020
2020
2020
2020
2020
2020
2020

Cluster Application Migration Tool

3 Clusters

3 Connected
0 Connection failed

View clusters

1 Replication repository

1 Connected
0 Connection failed

View replication

2 Migration plans

0 Not started

Clusters 3

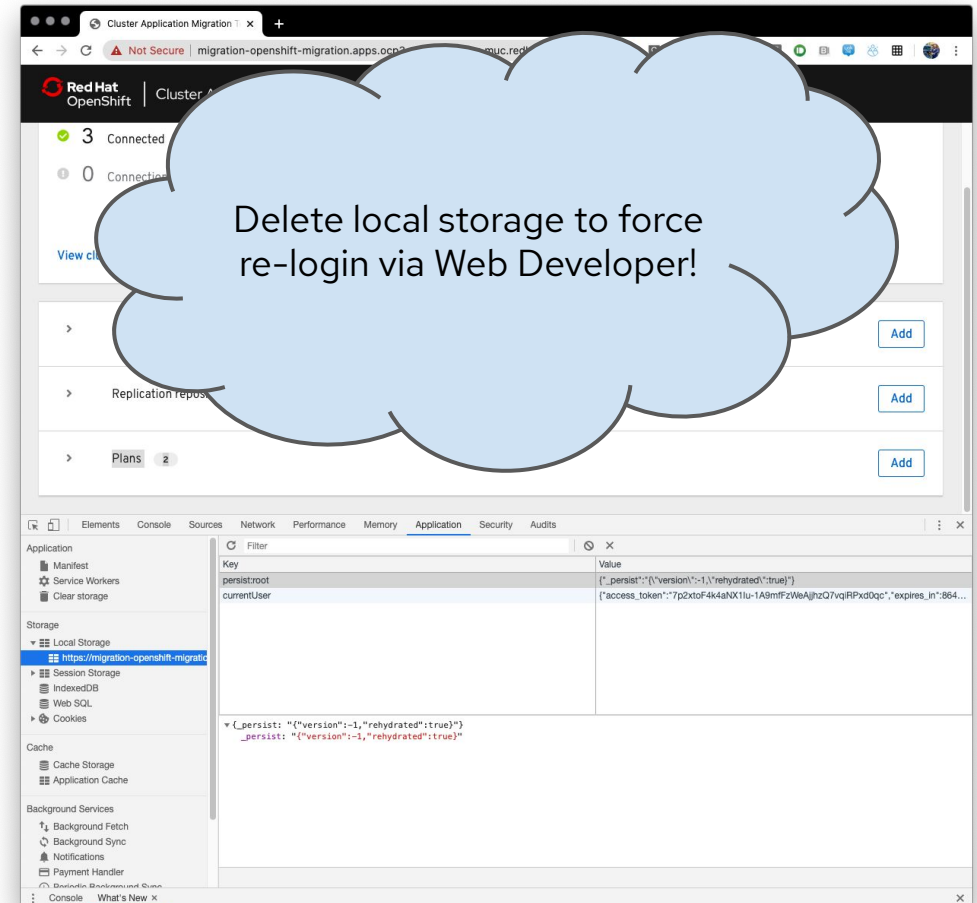
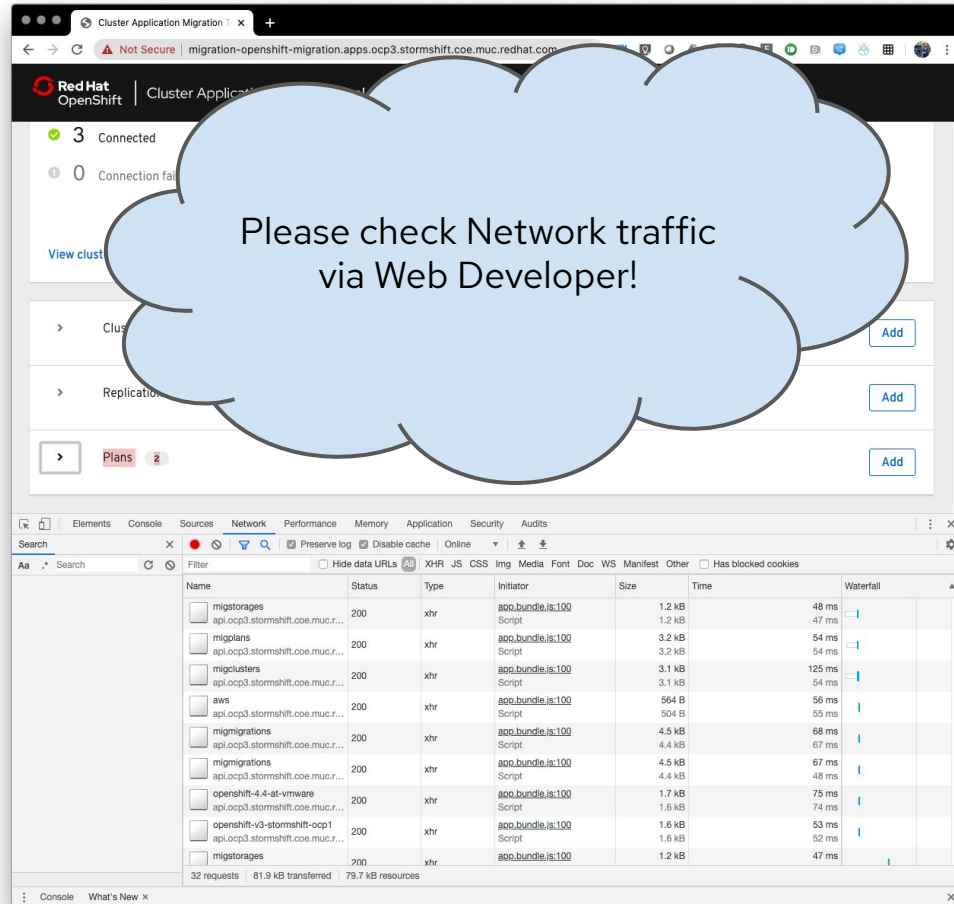
Replication repositories 1

Plans 2

Name	Migrations	Source	Target	Replication repository	PVs	Last state
migration-demo-1588071435	2	openshift-v3-stormshift-ocp1	host	aws	1	Stage Succeeded
rbohne-mig	2	openshift-v3-stormshift-ocp1	host	aws	1	Closed

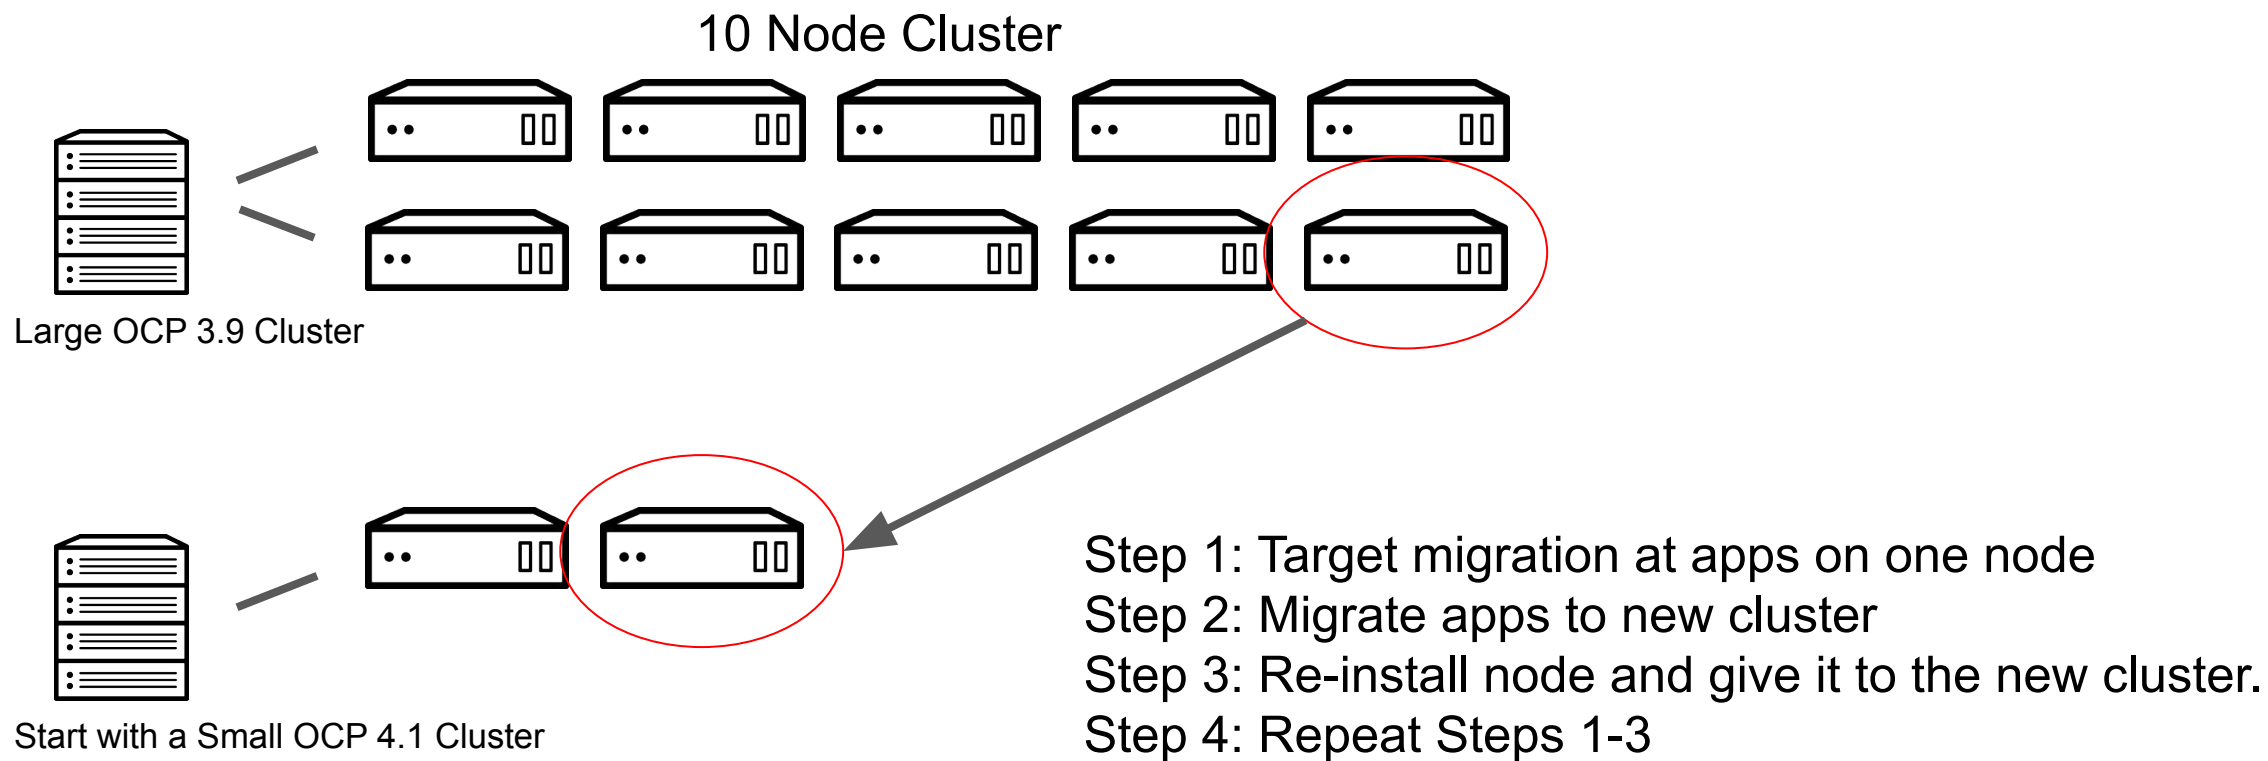
```
$ oc get migplan
NAME                                AGE
migration-demo-1588071435          4h
rbohne-mig                          31h
$
```


In case you have trouble with the WebUI



Customer Journey From OCP 3.x to OCP 4.x


How Do I Find Equipment for This?



If you are using a cloud provider enabled IaaS, we can automatically add nodes in OCP4 via worker machineSets based on load!

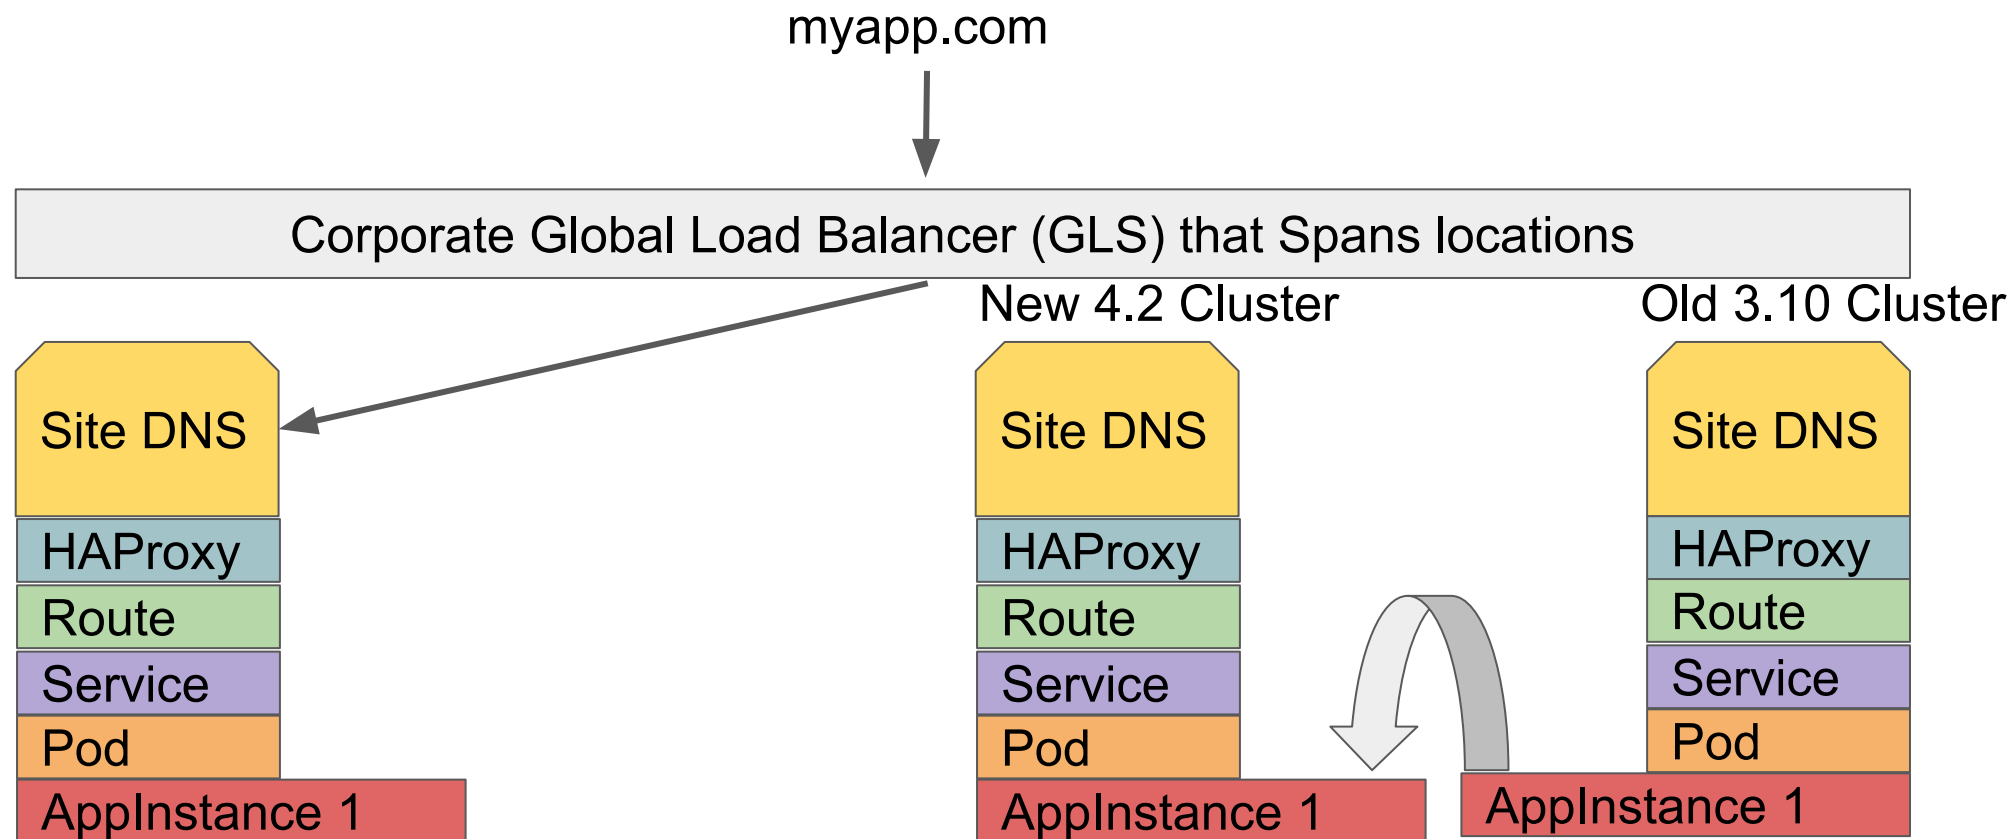
Talk to Me About Application Downtime

If you don't already have an install that offers no application downtime if a cluster is lost, maybe that is the first step. If you have that already, we can leverage it. Common setup is the following.



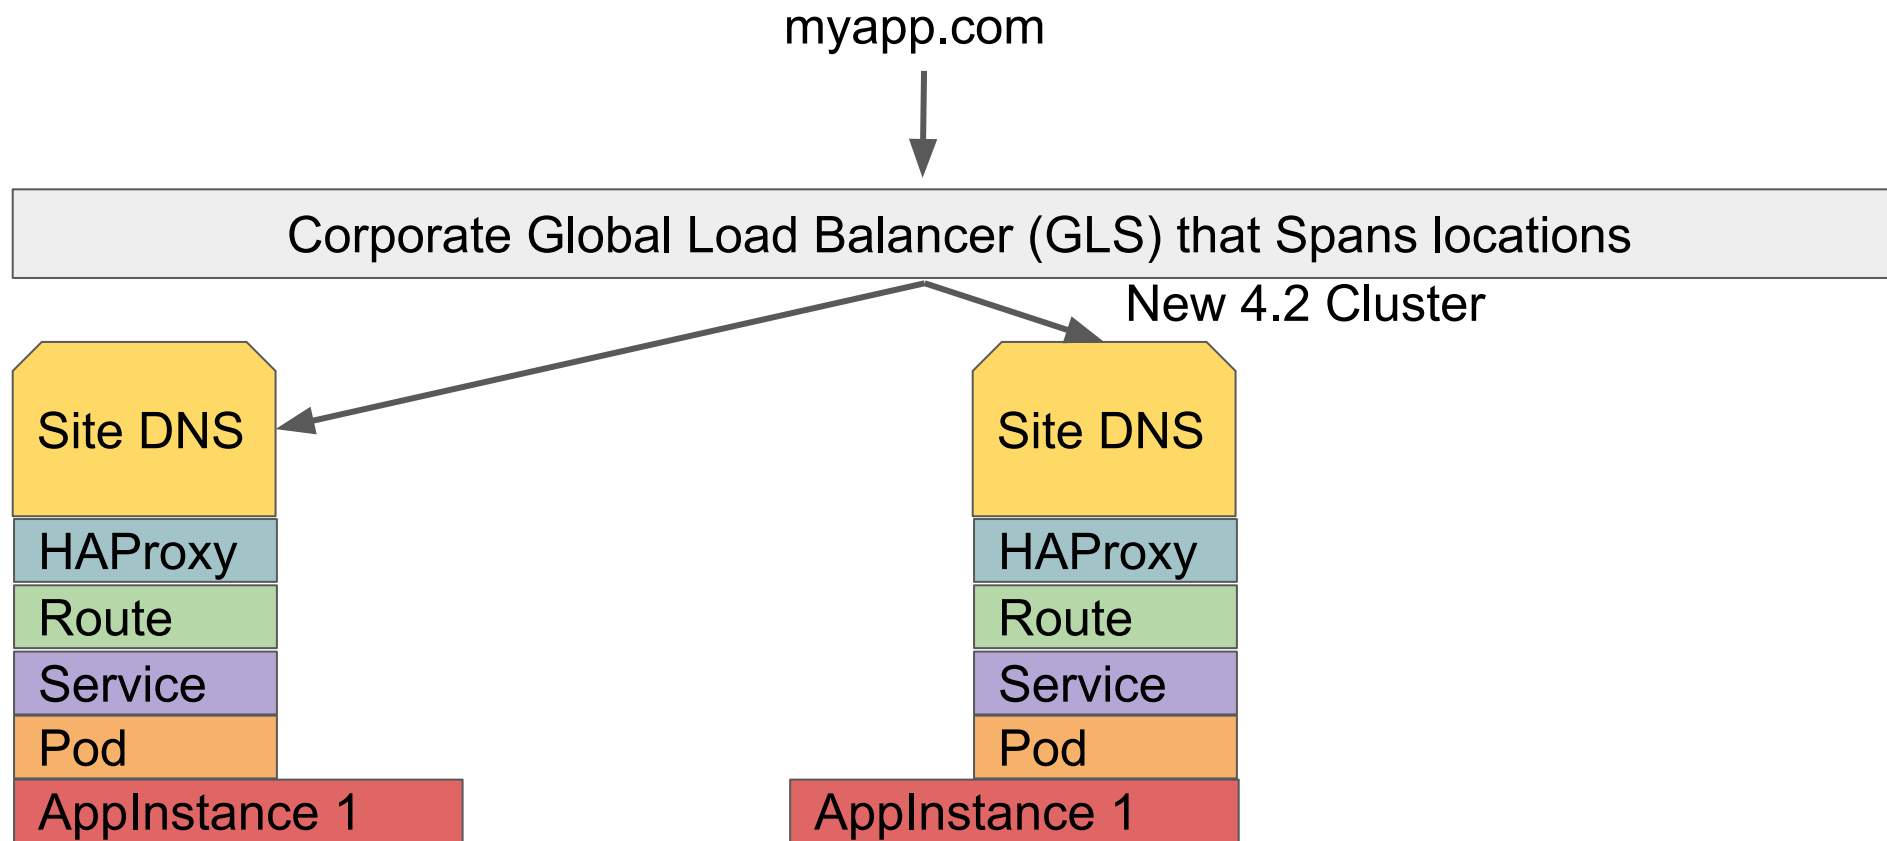
No Application Downtime Continued...

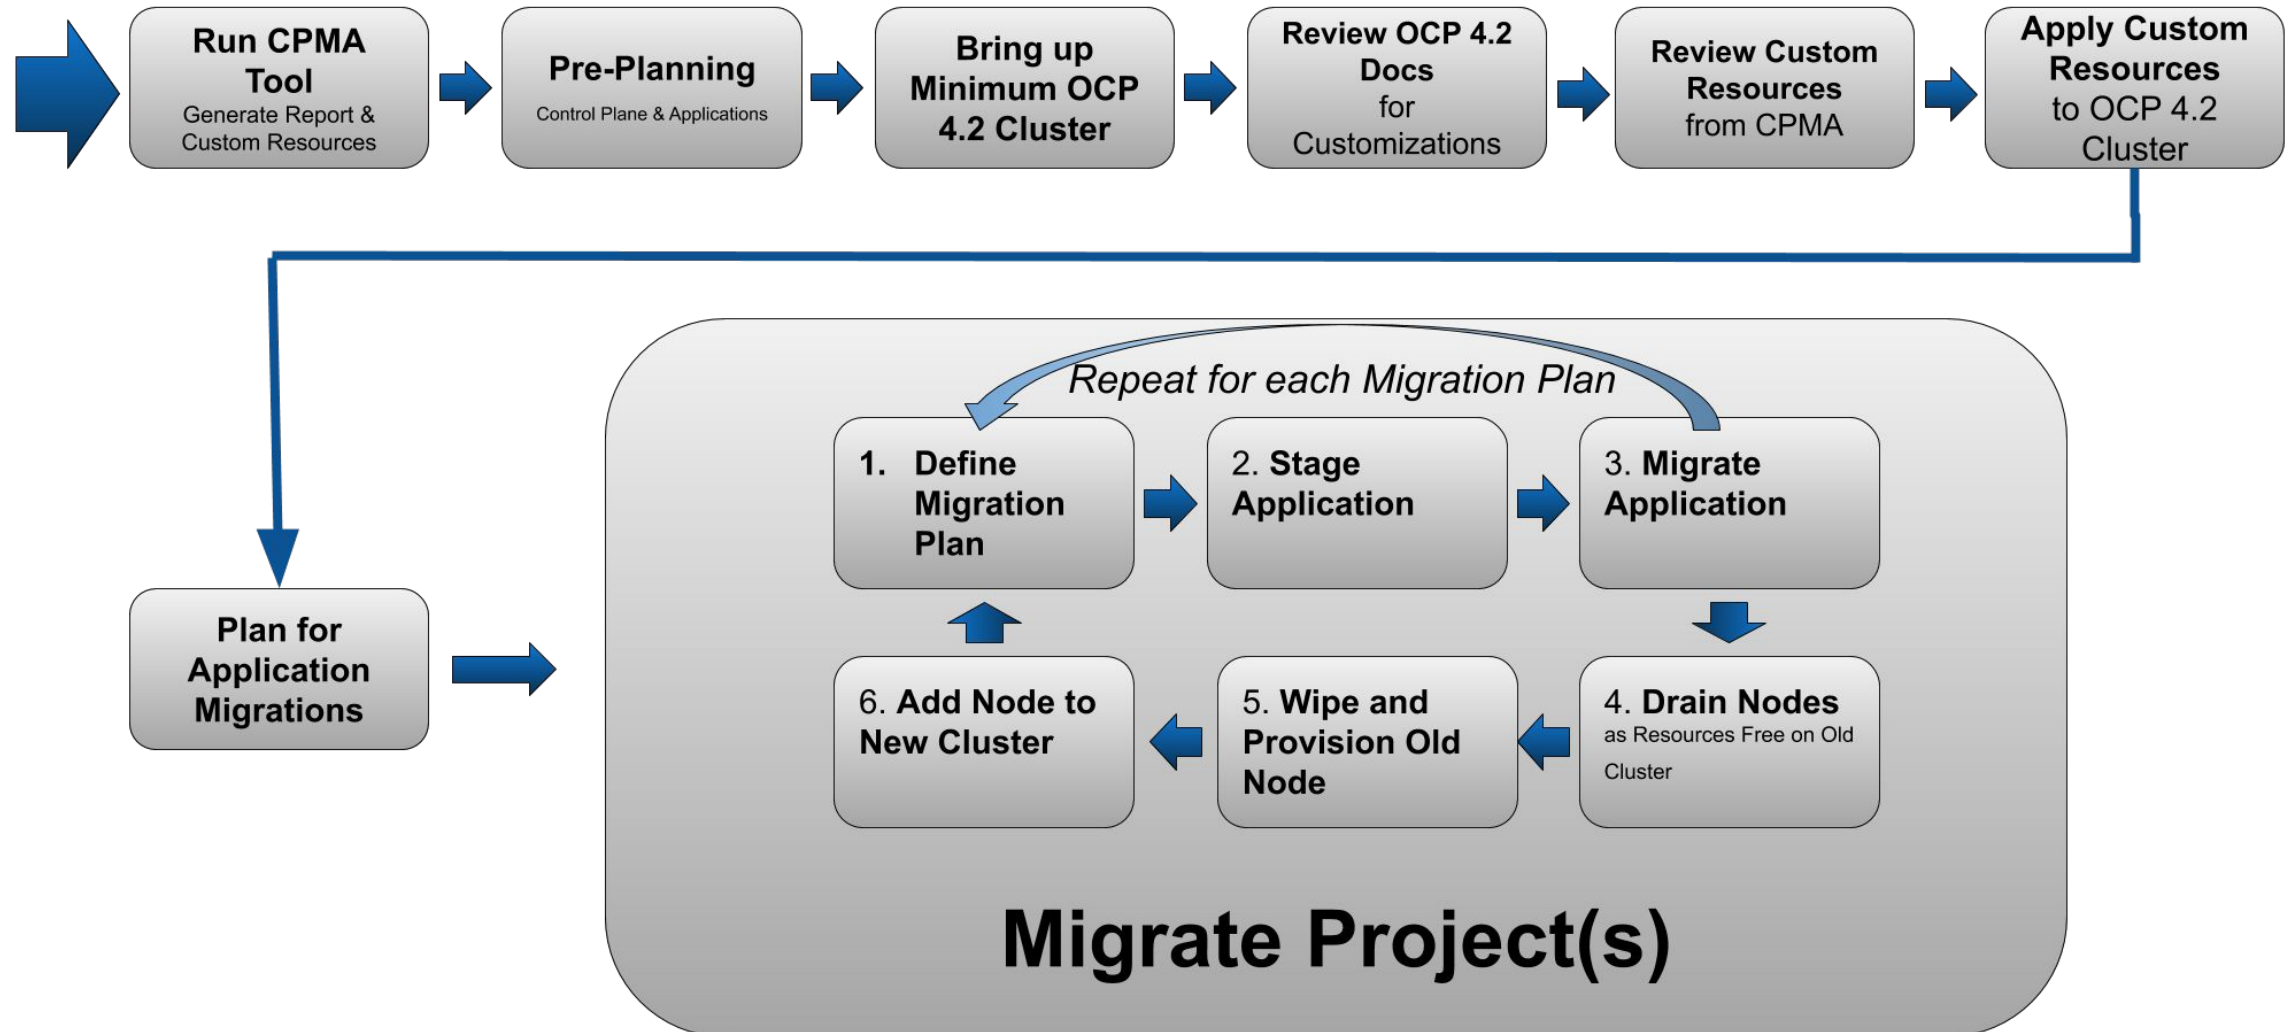
Direct traffic to a single cluster. Bring up a OpenShift 4 cluster. Use the OpenShift Migration Console to replicate and then migrate the application from the old 3.x cluster to the new 4.x cluster.



No Application Downtime Continued...

Add the new cluster's site DNS into the Global Load Balancer and allow traffic.





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