Migrating Virtual Machines to OpenShift Virtualization: Real-World Experience

Red Hat Innovate Together

October 6, 2026

Prague

Red Hat

Speaker introduction



Ales Nosek

Principal Architect

Red Hat Consulting

1) M. Sc. in Computer Science, Czech Technical University in Prague
2) Moved to California in 2012
3) Joined Red Hat in 2018
4) Helping customers design and build solutions for the hybrid cloud
5) Been a key contributor to 41 pre-sales engagements and 26 delivery projects

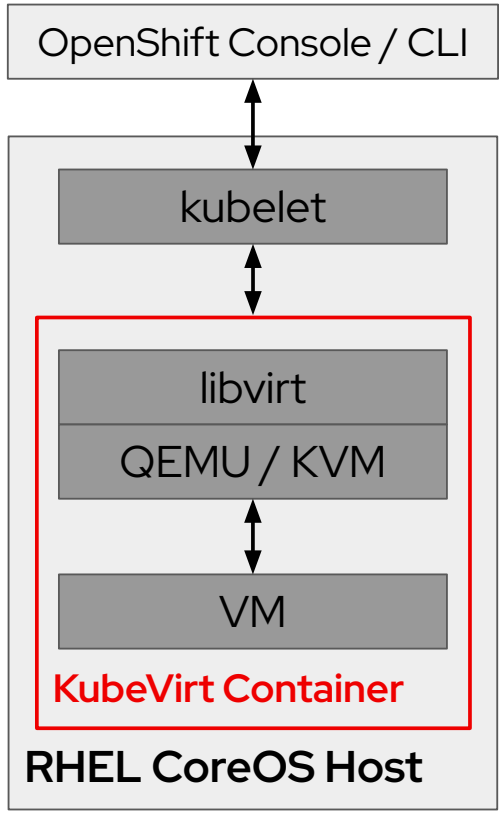
Red Hat

What is OpenShift Virtualization?



Red Hat OpenShift Virtualization

- Red Hat product built on top of the open-source **KubeVirt** project
- KubeVirt runs virtual machines inside a Kubernetes pod
- Battle-tested virtualization stack: QEMU/KVM, libvirt
- Features:
  - CPU and memory overcommit
  - CPU, memory, disk, and NIC hotplug
  - PCI device passthrough
  - Descheduler (equivalent of VMware DRS)
  - VM live-migration (vMotion), storage live-migration (storage vMotion)
  - ...



Migration projects



Migrations from VMware to OpenShift Virtualization

… for which I was the consulting architect

1) **Nationwide financial services provider**
    a) About 2,000 VMs in multiple data centers
    b) Cisco UCS hardware, OpenShift Data Foundation storage, Cohesity backup
2) **Prominent automobile manufacturer**
    a) Over 75,000 VMs in enterprise data centers and plants around the world
    b) Dell hardware, Portworx + Pure Storage FlashArray, PX-Backup
3) **Financial services provider**
    a) Over 12,000 VMs in multiple data centers
    b) HPE and Lenovo hardware, NetApp ONTAP (NFS), Cohesity backup



Migration project structure

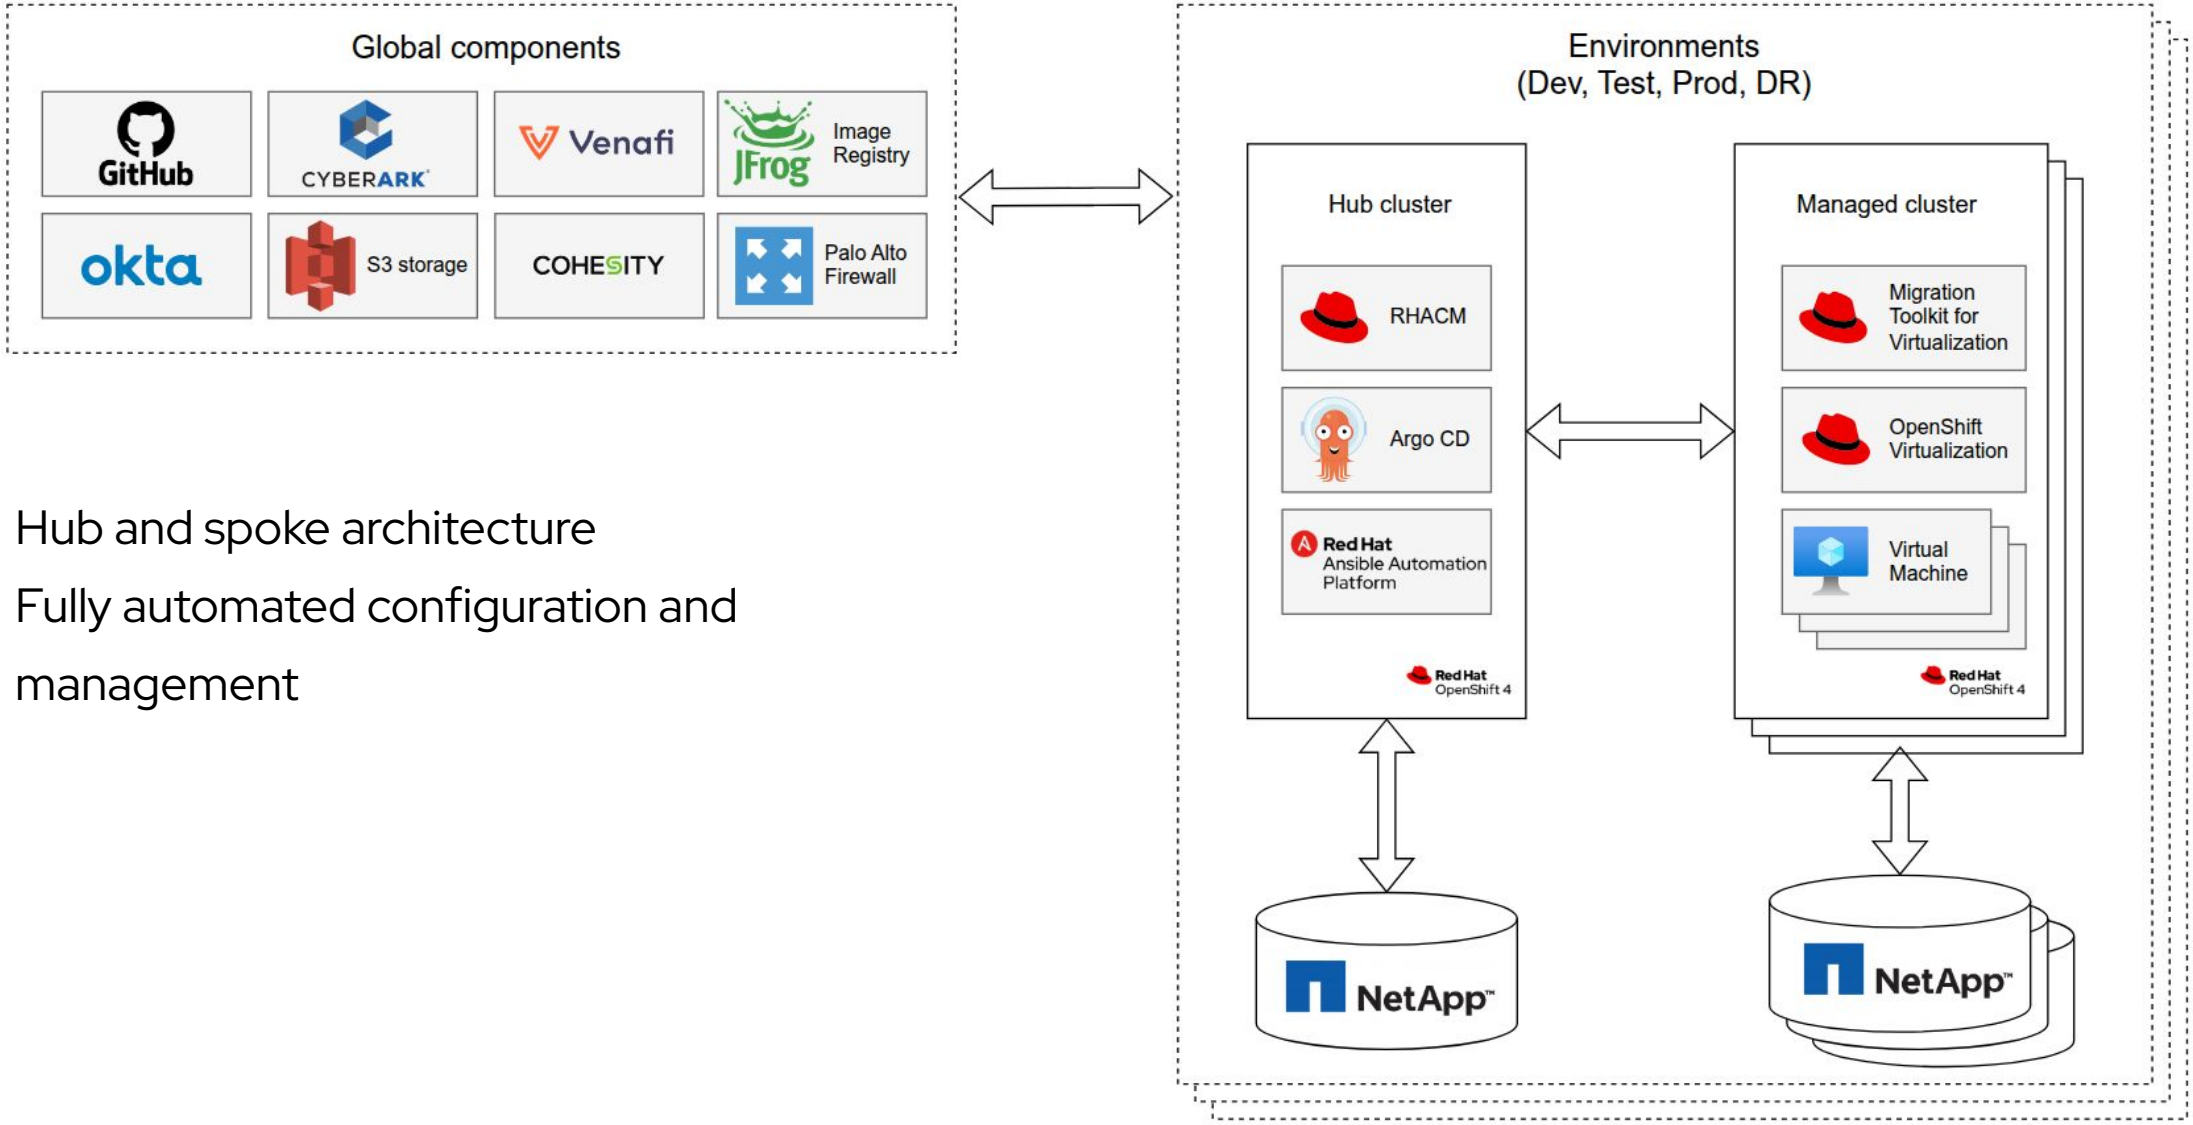
1) **Platform build-out**
    a) Platform architecture
    b) Platform deployment
2) **Virtual machine migration**
    a) Migration planning
    b) Migration execution



Platform Build-out

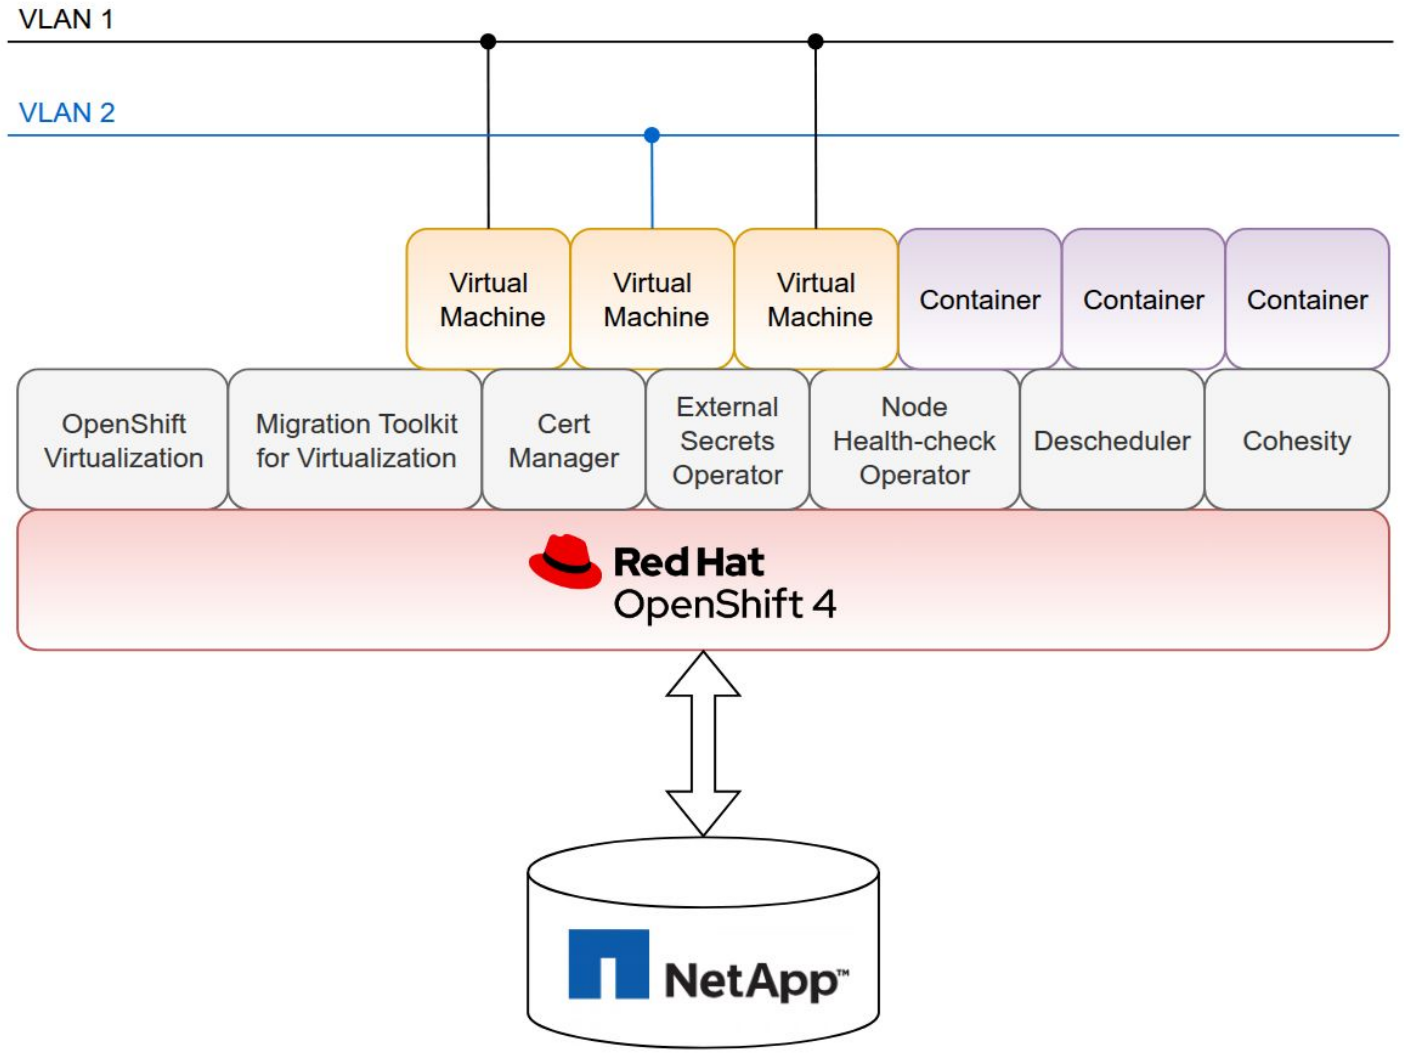


High-level architecture



- Hub and spoke architecture
- Fully automated configuration and management

Managed cluster



Virtual machine migration



Migration planning



Migration planning

What to migrate first and what next?

Goal is to reduce the VMware footprint as fast as possible.

1) VMs are migrated in **batches**. Leverage CMDB.
2) Start migrating **non-production VMs first** (utility servers, test servers, ...)
3) If you plan to reuse current hardware, your migration strategy hinges on how effectively you can **free up those resources** for OpenShift Virtualization.



Migration planning

What to migrate first and what next?

1) Consider migrating RHEL VMs first:

   a) Utilizing Red Hat guest operating systems on top of Red Hat Virtualization ensures support from a single vendor, and it is highly probable that everything will function correctly.

   b) In certain editions of OpenShift, the customer doesn't have to pay for RHEL subscriptions for a VM that runs on OpenShift Virtualization on bare metal.

2) Consider other supported operating systems next:

   a) Article 1351473: Converting virtual machines from other hypervisors to KVM with virt-v2v in RHEL 7, RHEL 8, RHEL 9, and RHEL 10

   b) Article 4234591: Certified Guest Operating Systems in OpenShift Virtualization

Migration backout plan

Example backout plan using Cohesity

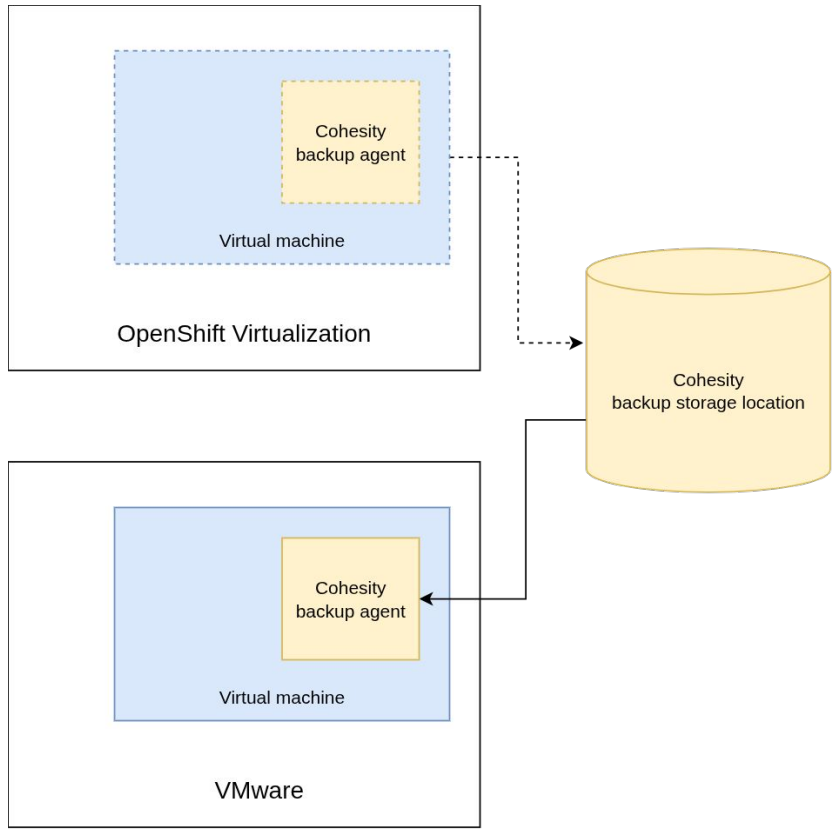
The Cohesity backup agent remains running on the virtual machine after the VM has been migrated into OpenShift Virtualization. If we need to migrate the VM back to VMware:

a. Create Cohesity backup.

b. Shutdown the VM on OpenShift Virtualization

c. Start up the VM on VMware

d. Restore the file system using the Cohesity backup.
This will apply any updates to the file system made while the machine ran in OpenShift Virtualization.



Migration execution



Migration Toolkit for Virtualization

Tool for migrating virtual machines to OpenShift Virtualization

Virtual machine is migrated in several steps:



1) Retrieve virtual machine metadata using vSphere APIs
2) Create virtual machine disks (`PersistentVolumeClaim`s) in OpenShift
3) Copy disk data from VMDK to PVC
4) Convert virtual machine disks using `virt-v2v`:
    a) Uninstall vmware-tools
    b) Adapt network configuration for virtio-net NICs
    c) ...
5) Create `VirtualMachine` resource in OpenShift and start it up



Migration Toolkit for Virtualization

Supported migration types

Migration Toolkit for Virtualization supports:



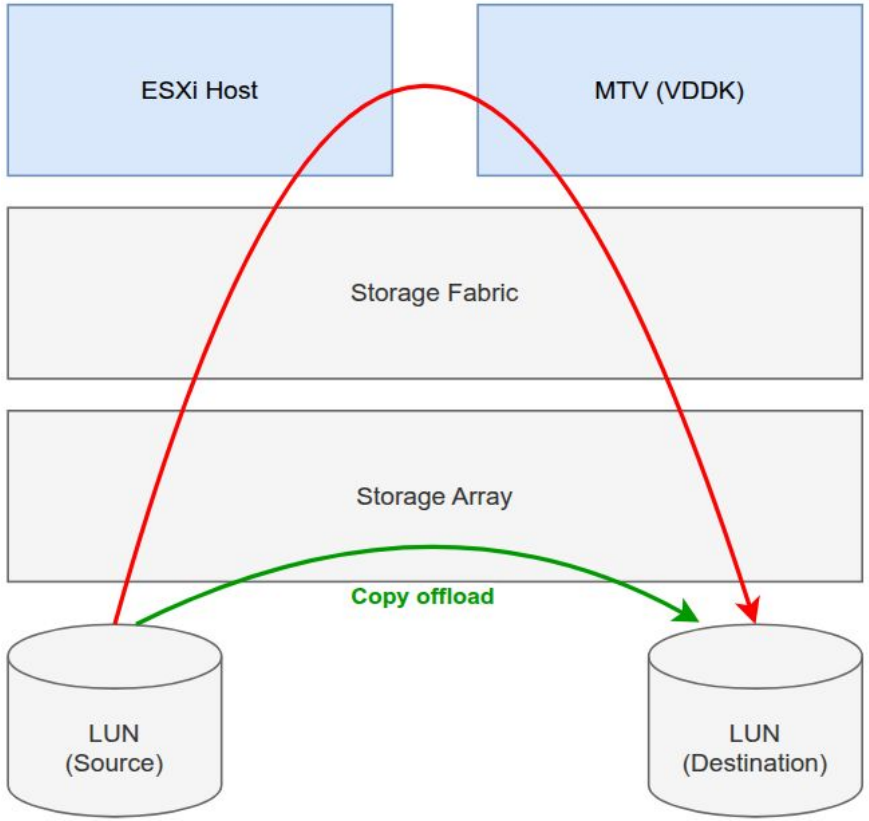
1) **Cold migration**
    a) Shut down the VM first then transfer the disk data
    b) Extended virtual machine downtime
2) **Warm migration**
    a) Warm migration reduces virtual machine downtime in comparison to cold migration
    b) The migration process consists of two stages:
        i) *Precopy stage*: Disk data is copied while VM continues running.
        Requires Changed Block Tracking enabled in VMware
        ii) *Cutover stage*: VM is shut down and the remaining data is copied
3) ~~**Hot migration**~~ **(does NOT exist)**
    a) No VM downtime required during migration



Migration Toolkit for Virtualization

Migration using storage copy offload

With storage copy offload, the VM disks are copied internally within the storage array, rather than being transferred across the storage fabric.



Migration Toolkit for Virtualization

Summary of migration methods

1) **Warm migration** reduces virtual machine downtime in comparison to cold migration.
2) With **storage copy offload**, the VM disks are copied internally within the storage array.

	MTV storage copy	Storage copy offload
Cold migration	Fully supported	Fully supported since MTV 2.11.0
Warm migration	Fully supported	Fully supported since MTV 2.12.0



Migration Toolkit for Virtualization

Virtual machine networking

1) The virtual machines being migrated are "pets". Their IP addresses are well-known and referenced in other parts of the environment. **The IP address cannot be changed**.



2) To preserve the networking configuration, we use a network attachment provided by Multus to **attach the virtual machine's NIC to the same VLAN as on VMware**. On this VLAN, the virtual machine is directly accessible using its configured IP address.

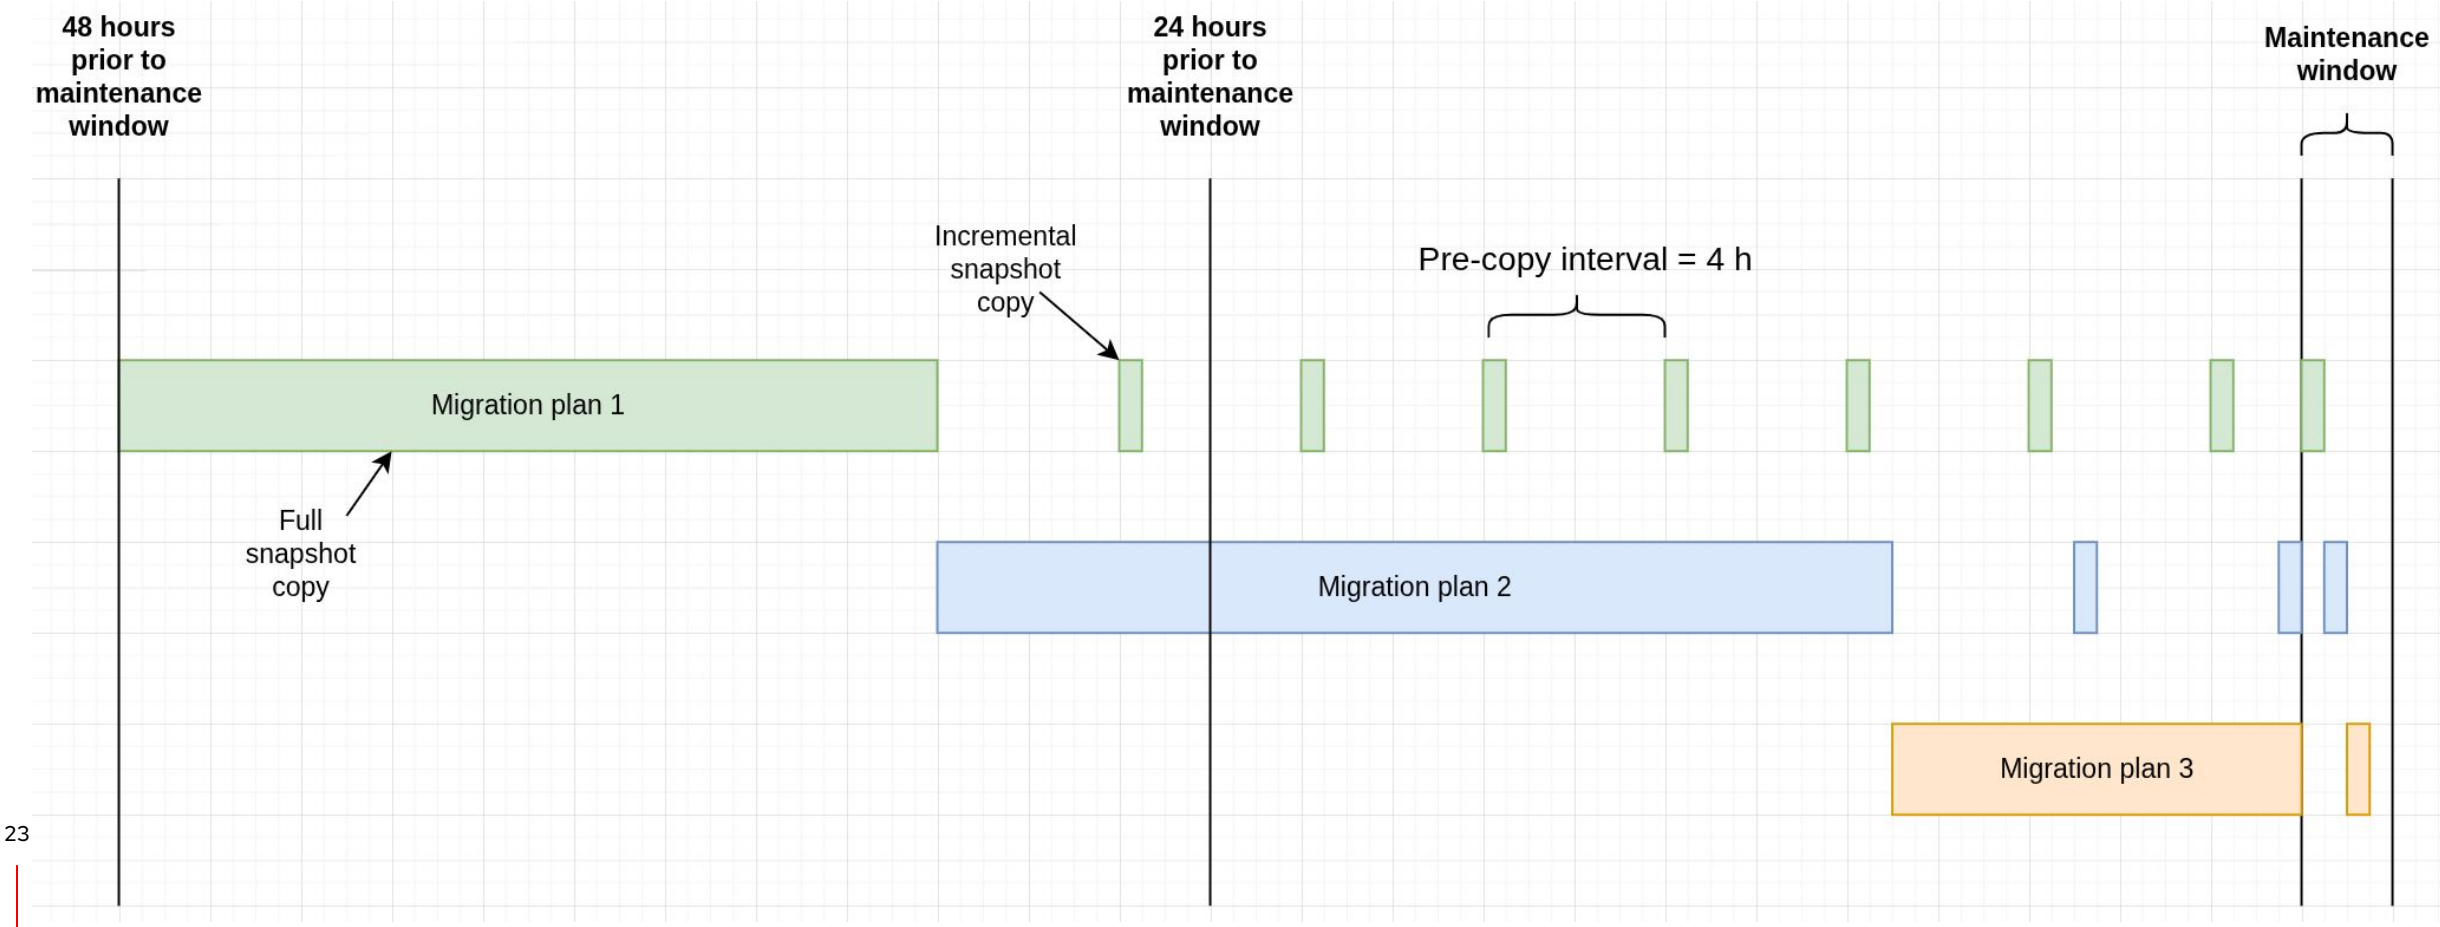


Migration examples

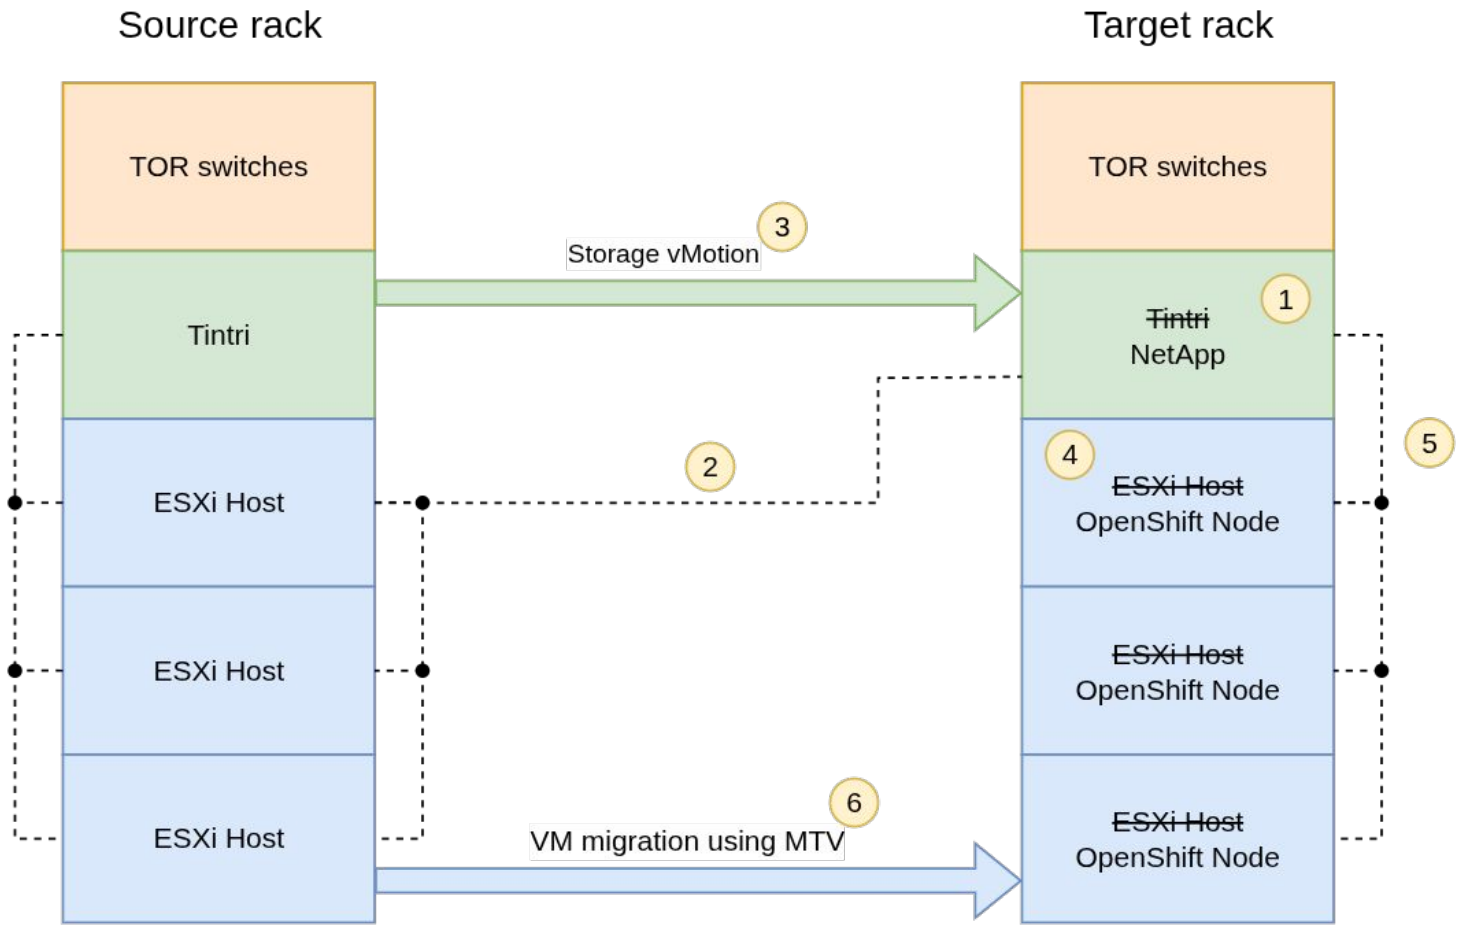


Avoiding storage overload

Scenario: warm migration with no storage offload



Example rack-by-rack migration using storage copy offload



Migrating clustered applications

*Goal: Migrate clustered applications that utilize **shared disks**.*

1) Example applications:
    a) Windows Server Failover Cluster (WSFC) with shared disk
    b) Oracle Real Application Cluster (RAC)
2) In VMware, shared disk access is enabled via *SCSI Bus Sharing* or *Multi-writer*.
3) VMware is unable to create snapshots of shared disks and hence **warm migration is not possible**.



Once the VM has been migrated ...

The following actions should be taken after the virtual machine has been successfully migrated to OpenShift Virtualization:

1) **Rename the VM** in vCenter so that we know that it was migrated.
2) **Disconnect** the virtual machine's **network interfaces** to avoid any potential MAC address conflicts in the event that the machine is inadvertently started in vCenter.
3) **Remove the backup tags** to deactivate VM backups on VMware.
4) **Configure VM placement** (anti-affinities) in OpenShift as necessary.



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/@redhat



facebook.com/redhat



x.com/RedHat

Red Hat