Red Hat OpenShift

CoreOS Layering Updates

Dusty Mabe
Principal Software Engineer

Red Hat

Red Hat OpenShift

Special Thanks

Colin Walters
Software Engineering

Zack Zlotnik
Software Engineering

Mark Russell
Product Management

Red Hat

The OpenShift Journey

OpenShift 3 to OpenShift 4 was a major leap:

- Opinionated openshift-install
- Day 2 Management
- Upgrades and configuration through Machine Config Operator
- Cluster version operator, second level operators
- RHEL CoreOS
  - image-based & purpose built
  - Tested and versioned in concert with OpenShift
  - API-managed

Common Requests

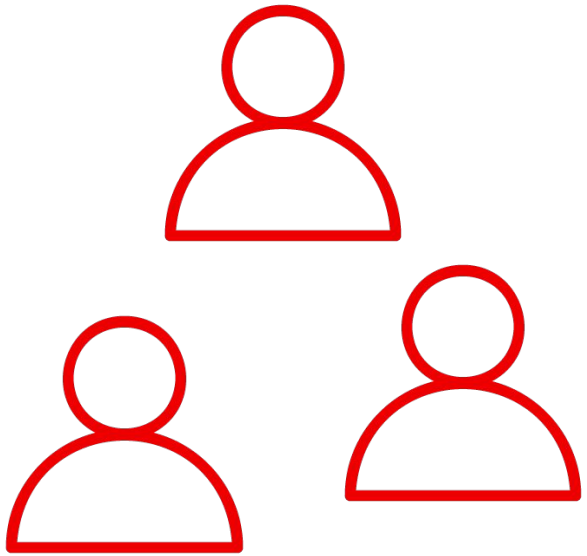

Customer/Partner needs

- Additional RHEL package
- Security Agents
- Kernel Drivers
- Familiar configuration

Evolving RHEL CoreOS

Available today in OCP 4.13+

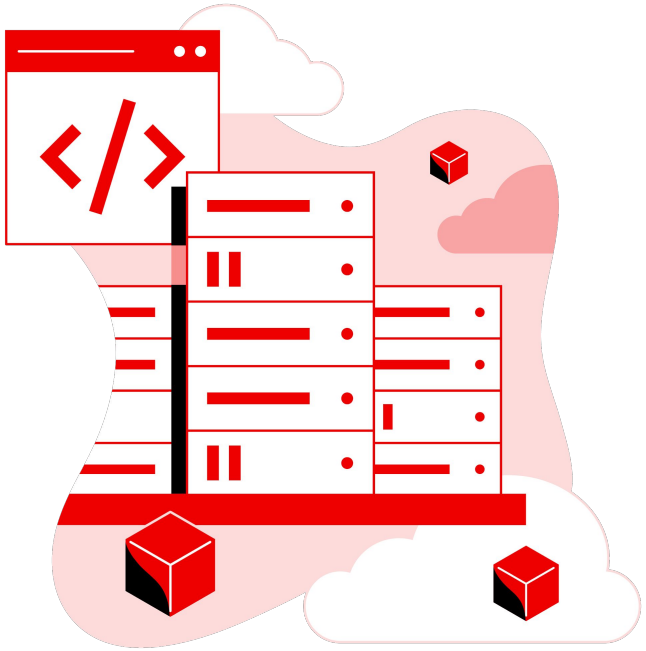

Enabling customization and simpler configuration management in OpenShift 4.13+

- Install 3rd party add-ons including kernel drivers
- Install additional RHEL content
- Simpler configuration file management (optional)
- MachineConfig API is preserved

What is CoreOS Layering?

Container images, everywhere

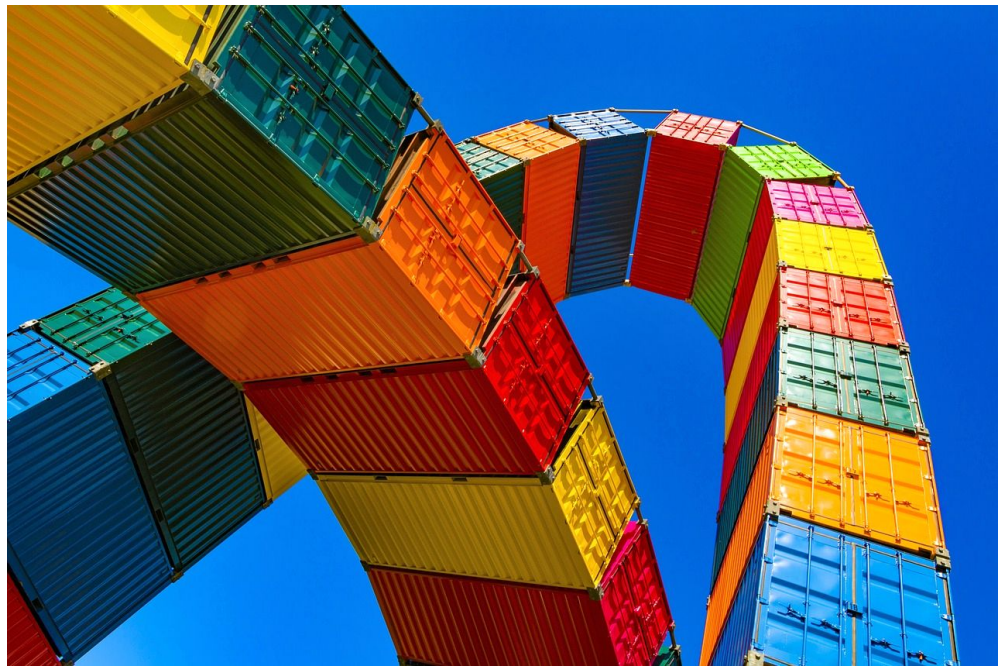

CoreOS Layering is a technology that puts the OS root filesystem in a standard OCI container image.

- Resulting container images are a transport format for updating the operating system root filesystem.
- Container image contents are written to a standard {xfs,ext4} filesystem, OSTree used in the background to manage kernel/bootloader
- Anything built to use standard container images suddenly now interoperates with CoreOS images

Photo by Valdas Miskinis
https://pixabay.com/photos/container-storage-trade-haulage-4203677/

CoreOS Layering Example

```
# Get RHCOS base image of target cluster `oc adm release info --image-for rhel-coreos`
# hadolint ignore=DL3006
FROM quay.io/openshift-release/ocp-release@sha256...

# Install our config file
COPY my-host-to-host.conf /etc/ipsec.d/

# RHEL entitled host is needed here to access RHEL packages
# Install libreswan as extra RHEL package
RUN rpm-ostree install libreswan && \
    systemctl enable ipsec && \
    ostree container commit
```


Source:
https://docs.openshift.com/container-platform/4.13/post_installation_configuration/coreos-layering.html

GA in 4.13

Off-cluster Builds

Build, Test, Deploy

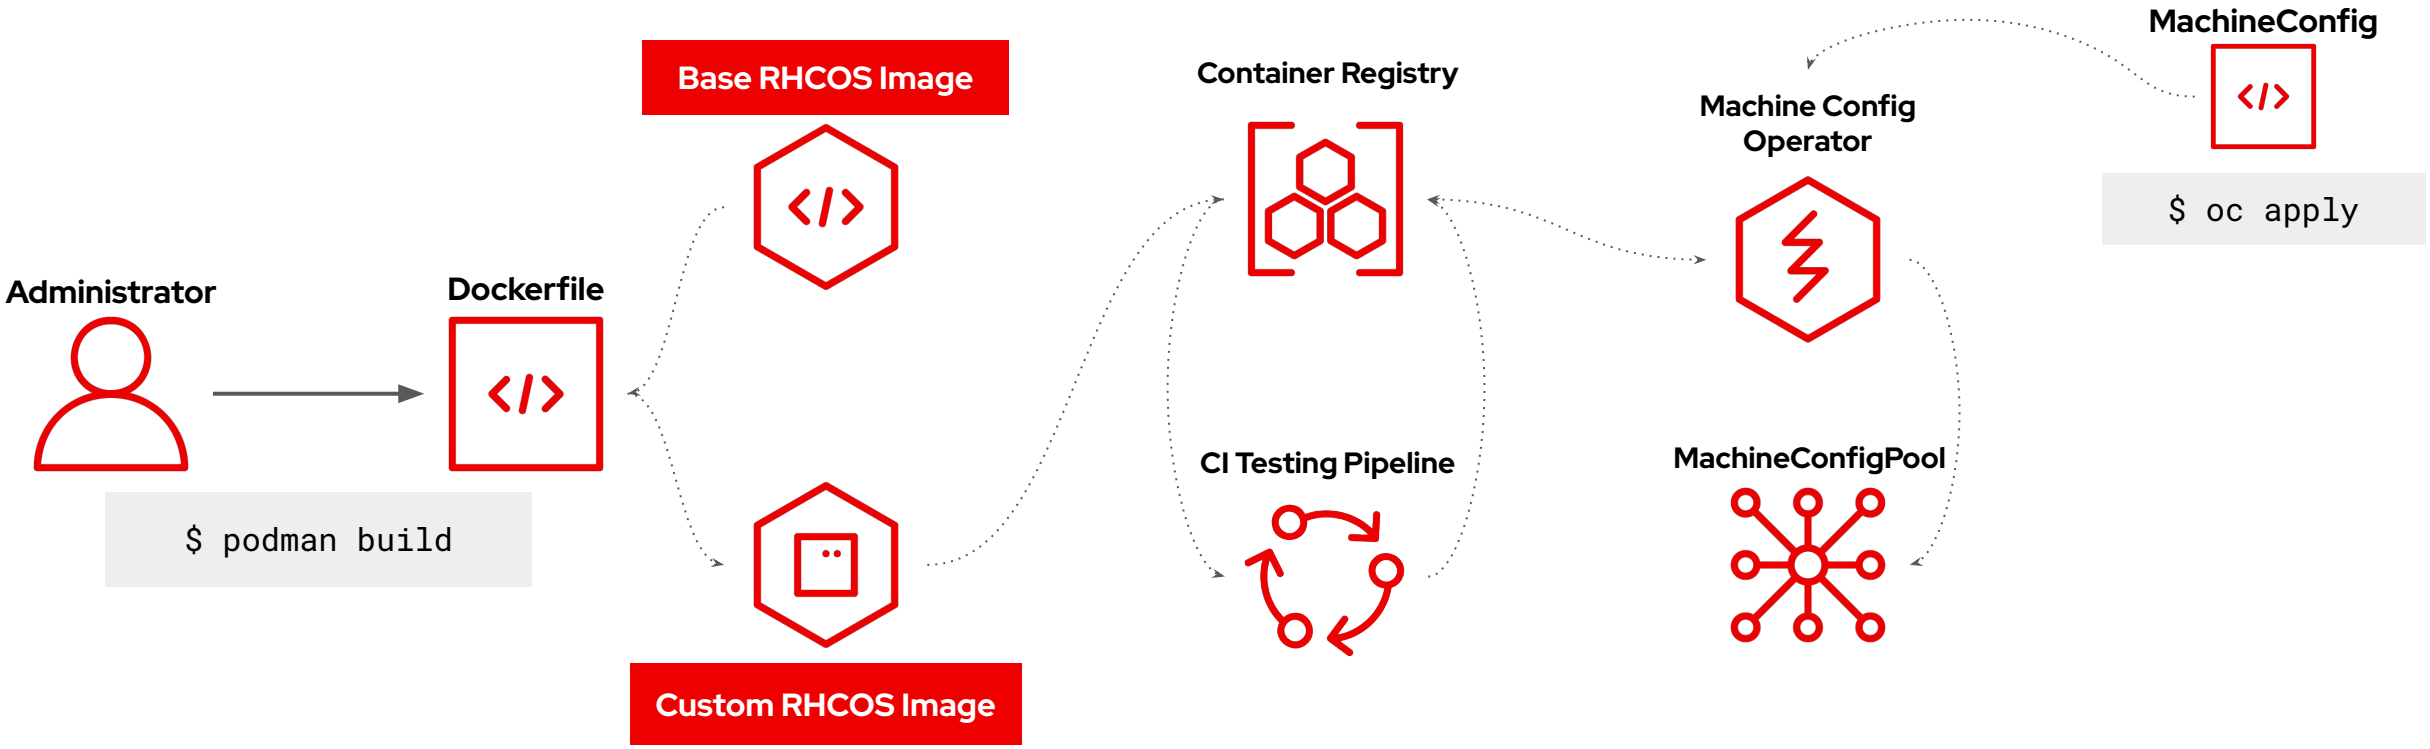

Red Hat

Future Feature

On-cluster Builds

Make it so!

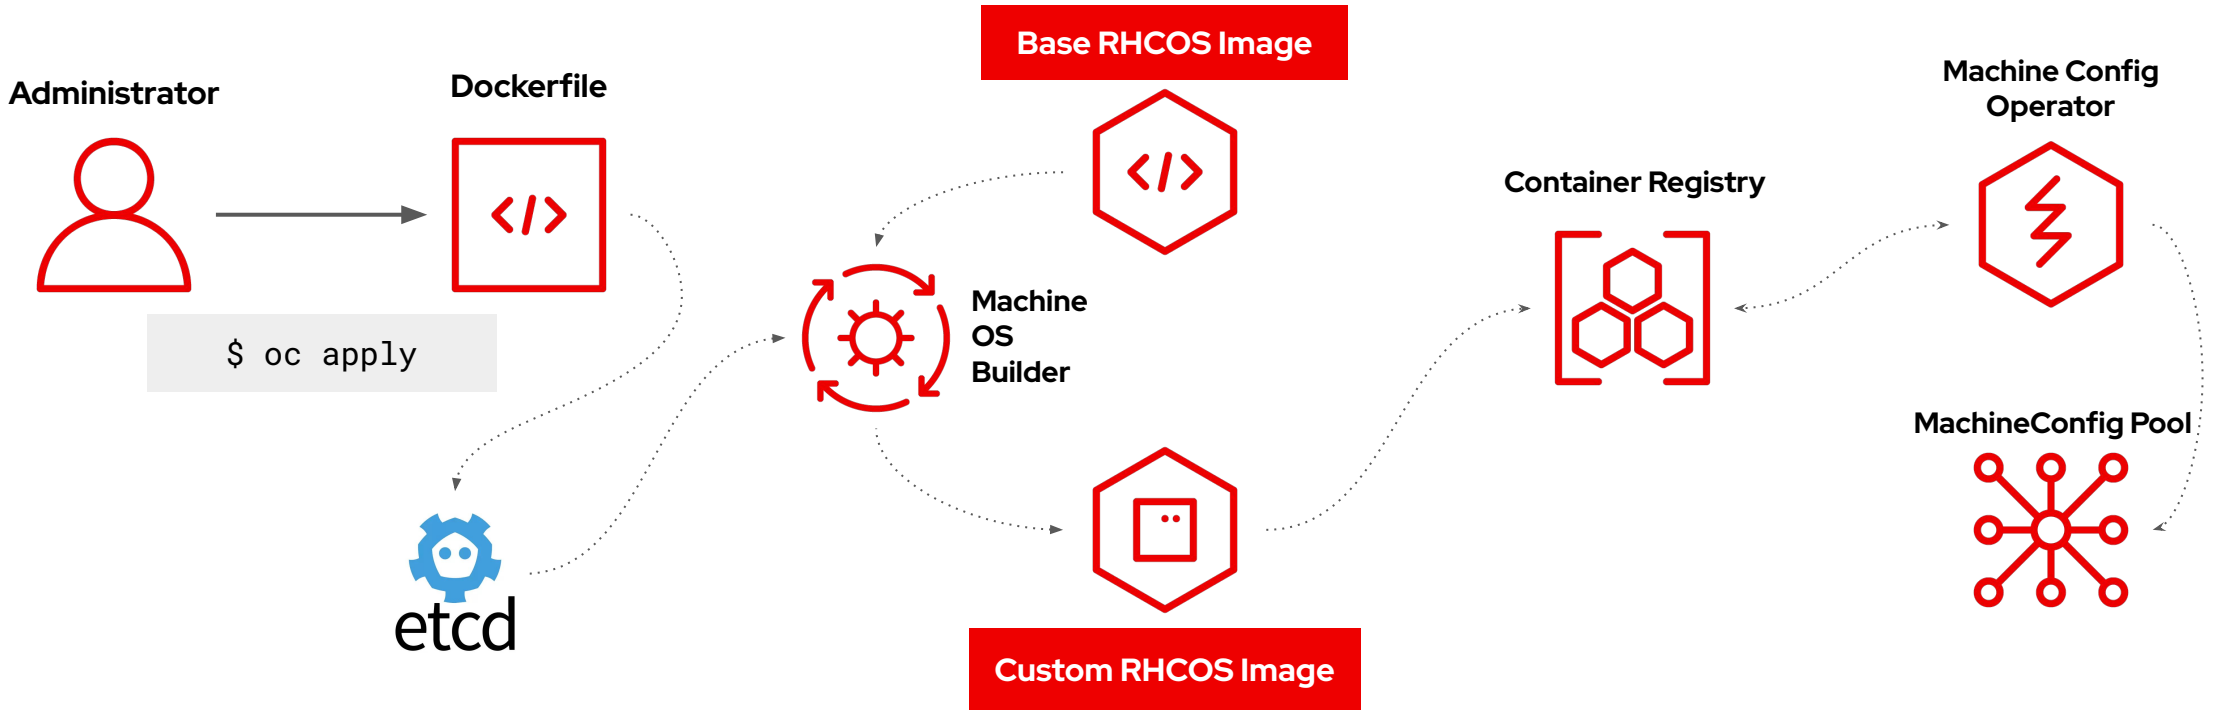

Red Hat

Demo 1

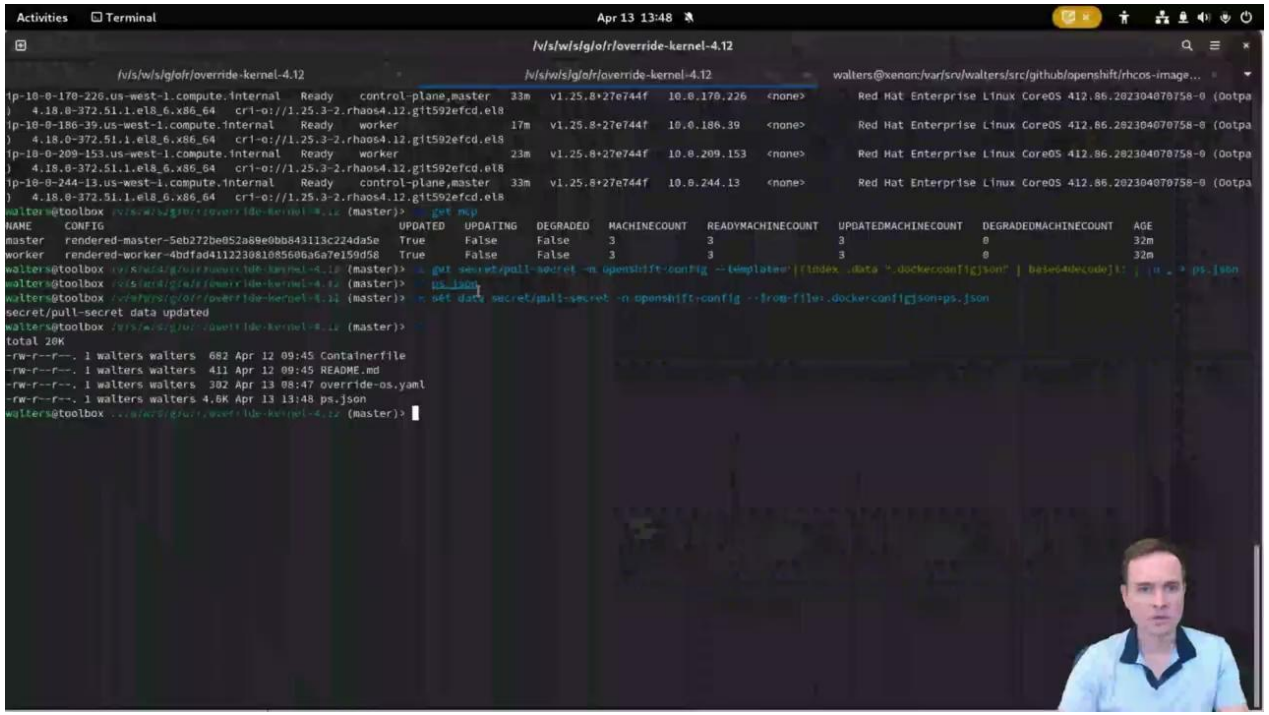

Red Hat

Demo 2

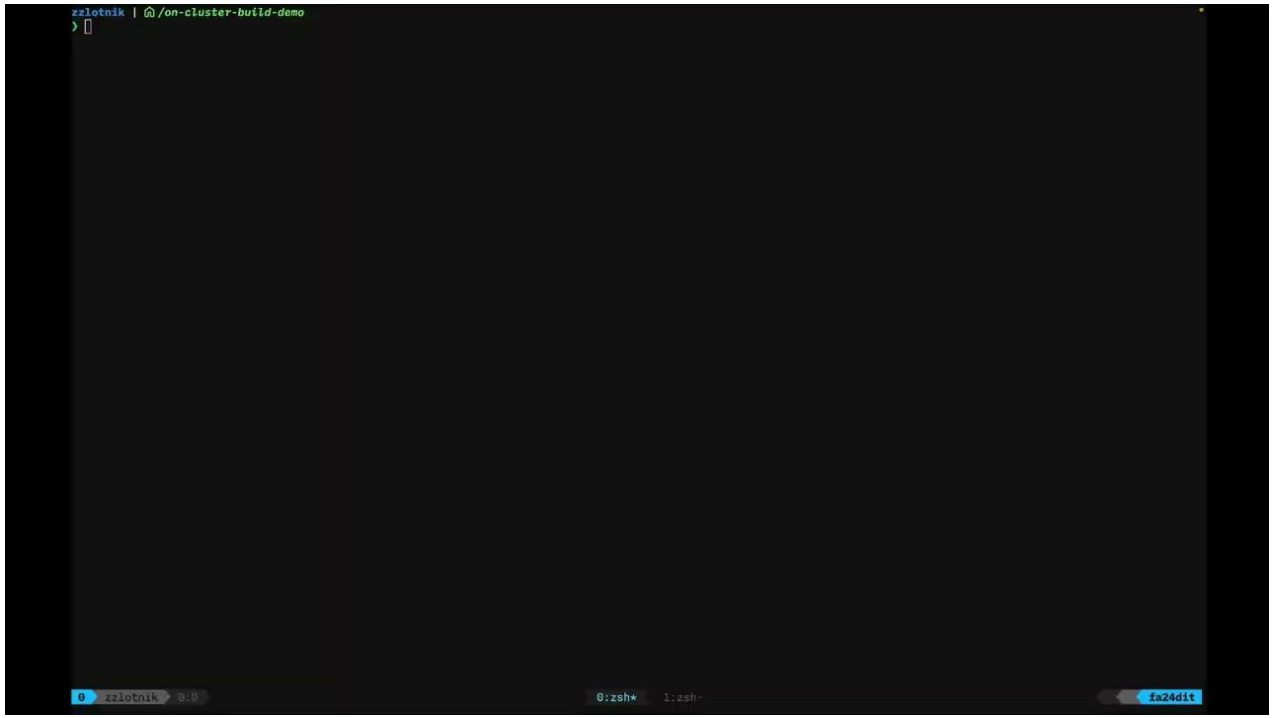

Red Hat

Off-cluster vs. On-cluster

Manage the pipeline OR let OpenShift do it for you

Off-cluster building is great for

- ✓ Production support today
- ✓ Taking ownership & responsibility for OS updates
- ✓ Creating custom test pipelines
- ✓ Centralizing image builds for many clusters
- ✓ Integrating with existing CI/CD and build systems

On-cluster building is great for

- ✓ Easy, automatic builds
- ✓ Ensuring that OpenShift upgrades automatically merge with your custom content
- ✓ Temporary requirements, e.g. a test or hotfix package

Important notes!

- Enabling off-cluster builds means taking **responsibility** for OS image updates
- Staying up to date is a virtue
- Standard Red Hat support policies do apply
- Currently Day-2 **only**

https://access.redhat.com/third-party-software-support
https://docs.openshift.com/container-platform/4.13/post_installation_configuration/coreos-layering.html

What's next?

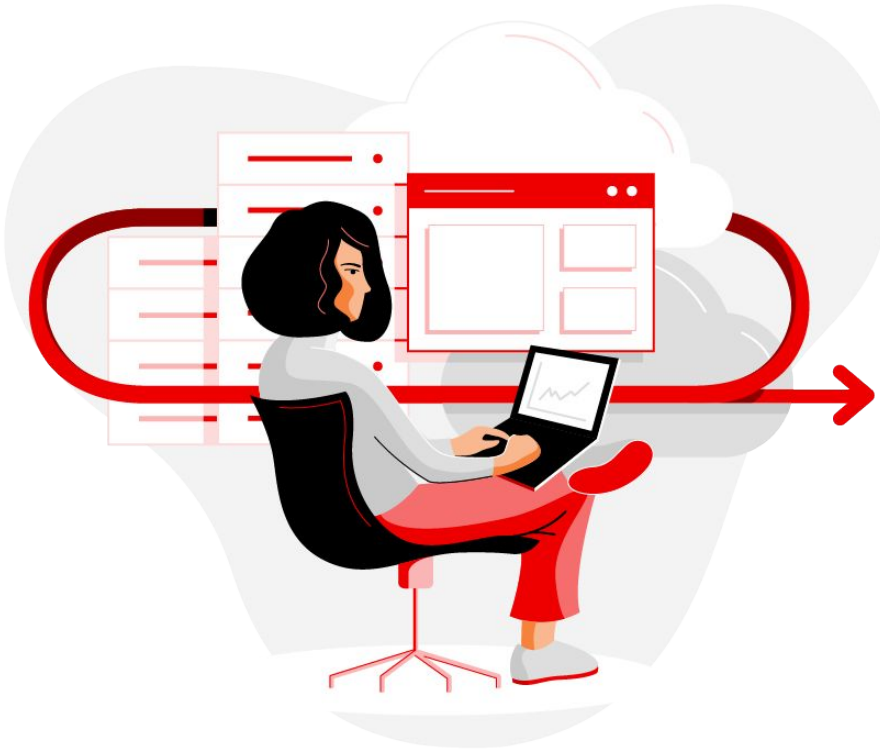

Development priorities:

- Day-0 custom install media for new machines
- On Cluster Build support to GA
- Console integration

Red Hat
Enterprise Linux
CoreOS

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/OpenShift


facebook.com/redhatinc


twitter.com/OpenShift

Red Hat