



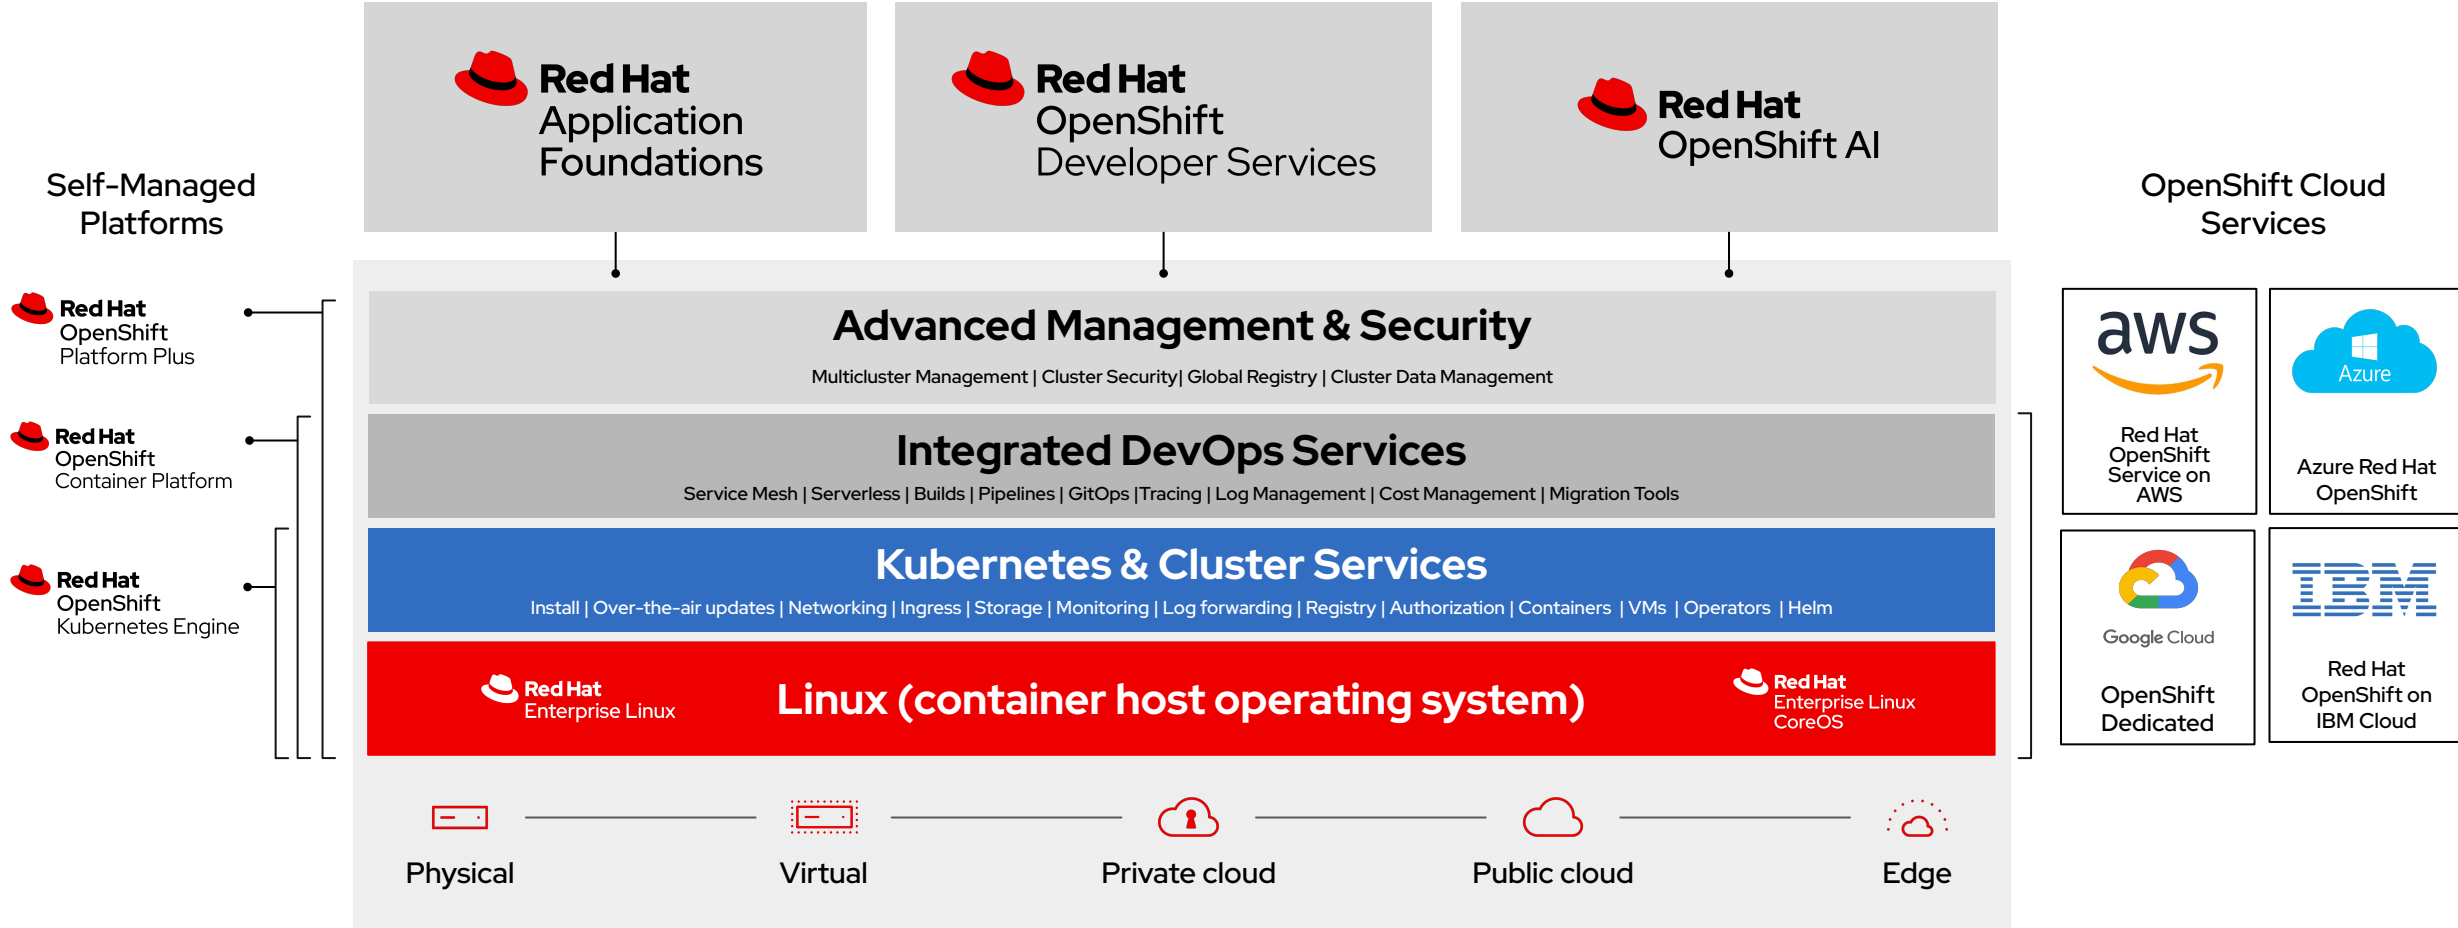
What's New in OpenShift 4.15

February 15th, 2024

OpenShift Product Management

red.ht/whatsnew

Red Hat open hybrid cloud platform



Kubernetes 1.28

'Planternetes'



Major Themes and Features

- ▶ Changes to support skew between control plane and versions
- ▶ Recovery from non-graceful node shutdown (GA)
- ▶ Match conditions for admission webhooks moves to Beta
- ▶ Automatic, retroactive assignment of a default StorageClass graduates to stable
- ▶ Consistent reads from cache
- ▶ Improvements to CustomResourceDefinition Validation Rules
- ▶ ValidatingAdmissionPolicies graduates to Beta
- ▶ Back off limit per index for index jobs
- ▶ Retriable and non-retriable Pod failures to fail faster
- ▶ Support for enabling swap space on Linux moves to Beta


CRI-O
1.28



Kubernetes
1.28



OpenShift
4.15




Notable Top RFEs and Components

Top Requests for Enhancement (RFEs)



- ▶ OVN IPsec support between an OCP cluster and an external provider [N-S] - [RFE-3345](#)
 - ▶ OVN IPsec supports encrypting all data between Red Hat OpenShift and any external provider
- ▶ Grafana dashboard for HAproxy - [RFE-2629](#)
 - ▶ Ingress operator dashboard in the OpenShift Console includes haproxy metrics visualization
- ▶ AWS Wavelength support - [RFE-3045](#)
 - ▶ Deploy compute nodes in AWS Wavelength zones
- ▶ Console improvements
 - ▶ Enable/disable tailing to log viewer [RFE-3560](#) - Choice of first 1000 lines or full pod logs in Console
 - ▶ Show Node Uptime information in the OpenShift Console - [RFE-3790](#)
 - ▶ Show Vertical Pod Autoscaler recommended values on Deployment Details page - [RFE-1068](#)

OpenShift 4.15 Spotlight Features





What's new in Red Hat

OPENSIFT 4.15


CORE

- OVN IPsec between OpenShift cluster and external provider
- Core/infrastructure networking metrics in Console
- Hosted Control Planes support for Virtual Hosts with Agent Provider (Tech Preview)
- Red Hat build of OpenTelemetry
- Power Monitoring (Tech Preview)



EDGE

- AWS Outposts and AWS Wavelength for latency sensitive applications
- Operator Lifecycle Manager on Red Hat Device Edge for simplified operator deployment
- Machine Vision on Arm with Red Hat Device Edge



VIRTUALIZATION

- Instance types for VM provisioning
- Metro DR with OpenShift Data Foundation
- Dynamic reconfiguration of NIC to a running VM







Red Hat Device Edge and MicroShift





Red Hat Device Edge with MicroShift is a Kubernetes distribution derived from OpenShift designed for small form factor devices and edge computing.

Machine Vision on ARM with MicroShift

- MicroShift on RHEL 9.3 for NVIDIA Jetson Orin based devices
- [NVIDIA JetPack 6.0](#) Developer Preview for RHEL
- [NVIDIA DevicePlugin support for RHDE](#)



 **Red Hat**
OpenShift AI  Model Serving

 **Red Hat**
Device Edge  MicroShift V4.15
RHEL9.3 ARM

 Jetson Orin with integrated GPU

Operator Lifecycle Manager with custom catalogs



- Optional component (`dnf install microshift-olm`)
- Build your own (small) catalog with just the operators you need to save resources
- Caveat: check with the operator provider if deployment to MicroShift is supported

MicroShift designed for FIPS

- When installed and running on RHEL in FIPS mode, MicroShift core components use the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-3 Validation on the x86_64 architecture.

Install OpenShift in AWS Edge Locations

Deliver latency sensitive applications closer to end users and on-premises installations



Outposts

Generally Available

- ▶ For **customer managed** OpenShift in AWS
- ▶ Extends **workers** to run in Outposts
- ▶ Deploy post-cluster installation (Day 2)
- ▶ Use **Amazon Elastic Block Store (EBS) gp2** for storage on Outposts




Generally Available

- ▶ For **customer managed** OpenShift in AWS
- ▶ Extends **workers** to Wavelength Zones
- ▶ BYO Virtual Private Cloud (VPC) with Wavelength Zone into existing subnet
- ▶ Deploy using **Installer Provisioned infrastructure (IPI)**
- ▶ Post-cluster installation (Day 2) option

Giant-scale Generative AI with NVIDIA Grace-Hopper

- ▶ Red Hat and NVIDIA collaboration launched after GTC 2021
- ▶ Grace Arm CPU: 2X the performance per watt compared to 2-socket data center systems*
- ▶ OpenShift support for 64k page size kernel
- ▶ CPU+GPU coherent memory model and NVIDIA NVLink Chip-2-Chip interconnect
- ▶ Increase the amount of GPU-accessible memory for large language models
- ▶ NVIDIA GPU Operator enabled for NVIDIA Grace-Hopper systems

* Source NVIDIA: [Hot Chips 2023](#), NVIDIA Chief Scientist, Bill Dally, 5 MW Data Center level projection, NVIDIA Grace Superchip vs x86 2-socket data center systems (AMD Epyc 9654 and Intel Xeon 8480+)



NVIDIA Grace-Hopper Superchip. Source : NVIDIA

```
$ cat /etc/redhat-release
Red Hat Enterprise Linux CoreOS release 4.15

$ lscpu | grep "NUMA node0"
NUMA node0 CPU(s):           0-71

$ uname -r
5.14.0-284.50.1.el9_2.aarch64+64k

$ getconf PAGESIZE
65536
```

64k page size kernel with Red Hat OpenShift 4.15

Red Hat OpenShift Networking Enhancements



IPSec North-South (Egress-Ingress) Generally Available

Networking Enhancements

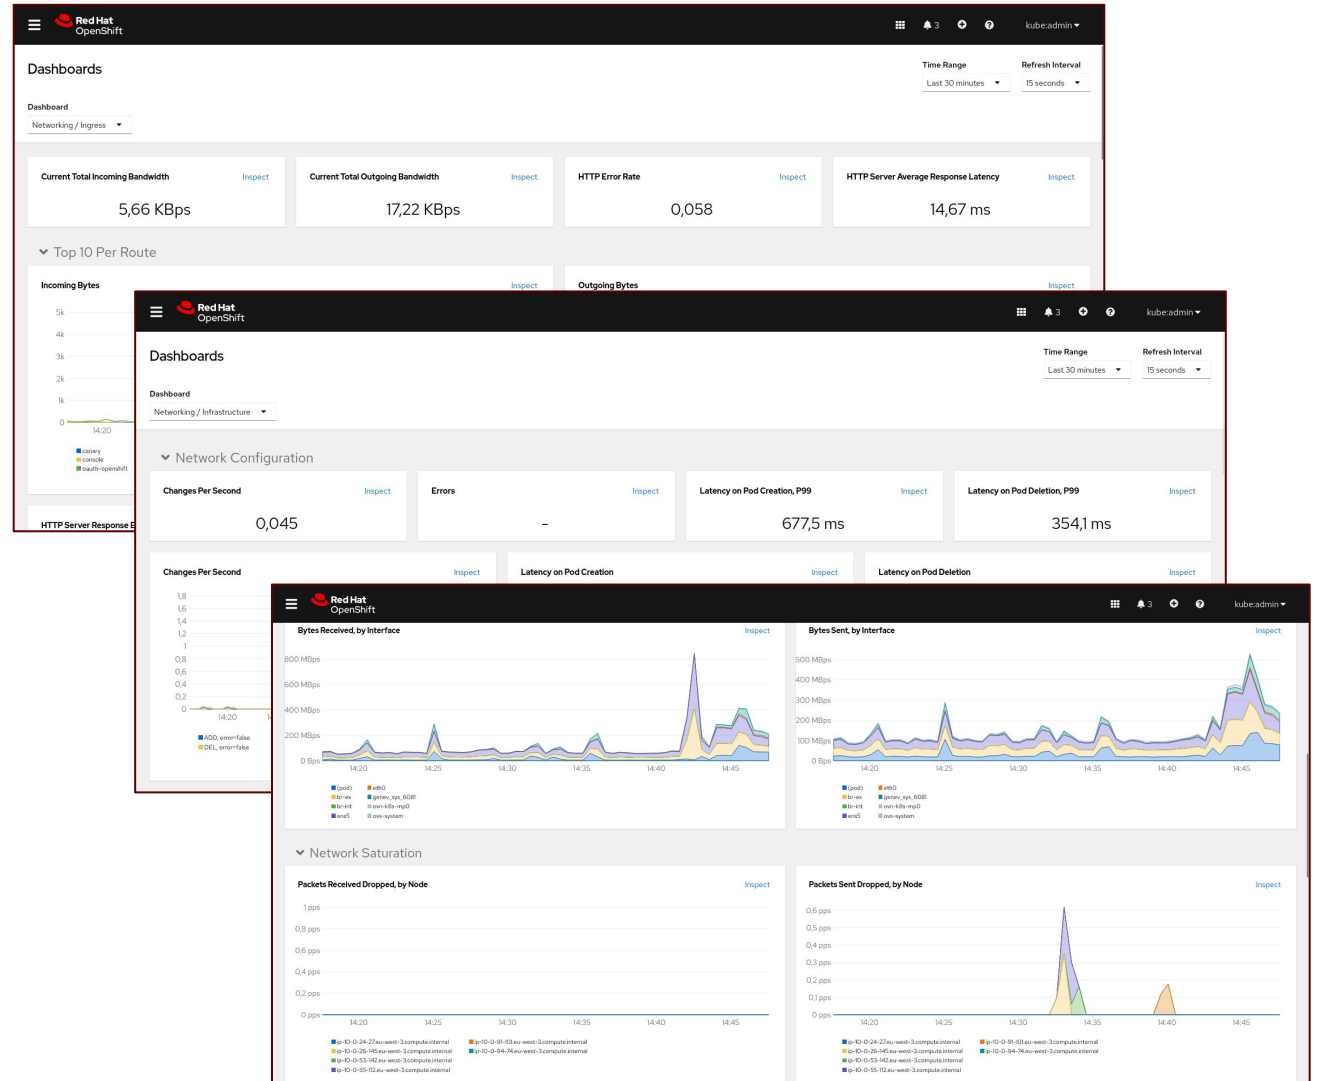
- ▶ OpenShift is adding support for North-South IPsec, and integrating it with existing East-West IPsec capability
- ▶ OVN-Kubernetes, only
- ▶ General Availability at 4.15
- ▶ Mechanics:
 - IPsec East-West: move to Host from cluster pod
 - IPsec North-South: join with E-W on Host
- ▶ Allows for encryption offload
- ▶ Adds telemetry

Enhanced Networking Dashboards

New Networking Dashboards

Observe -> Dashboards

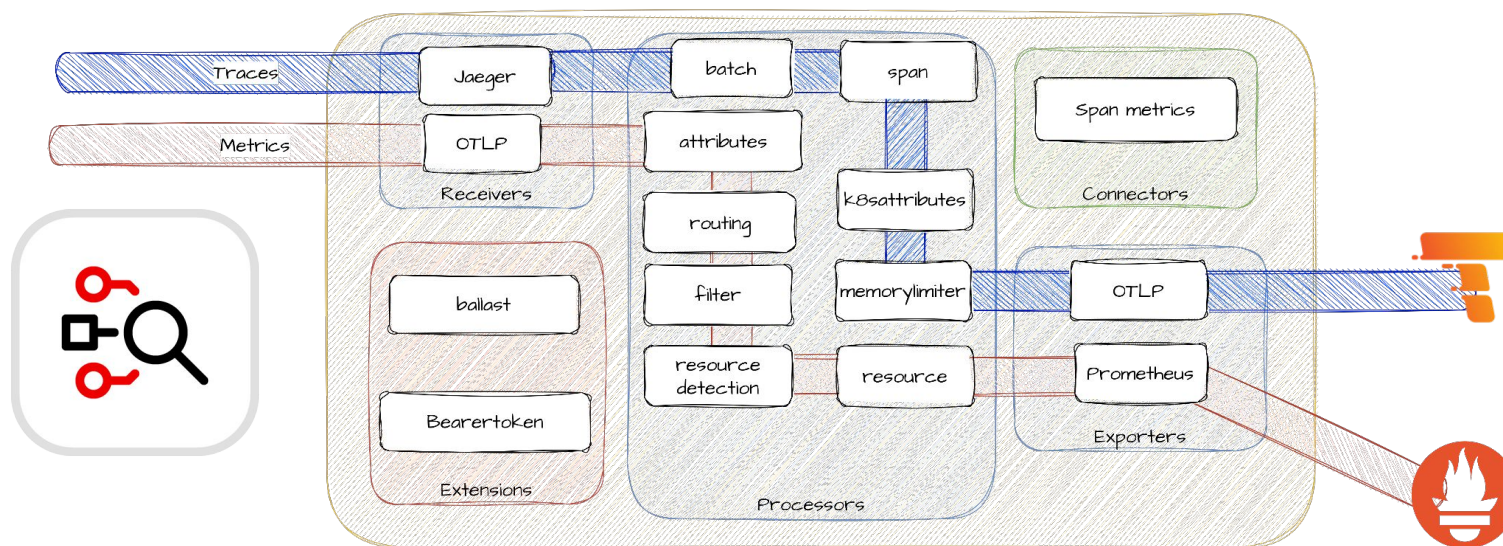
- ▶ **Networking/Ingress**
 - ▶ #routes/shard
 - ▶ HTTP latency and error
- ▶ **Networking/Linux SubSystem Stats**
 - ▶ Network Bandwidth and throughput of various host interfaces
- ▶ **Networking/Infrastructure**
 - ▶ Latency on pod deletion/creation
 - ▶ OVN-K control plane resource usage



Red Hat Build of OpenTelemetry

Open protocol for collecting, storing and exporting data. Avoid vendor lock-in and rely on open standards!

- ▶ The Red Hat build of OpenTelemetry is now **Generally Available** for metrics, logs and traces.
- ▶ **Arm** support added
- ▶ Extract **span metrics** from traces and even create alerts from them.
- ▶ **Automatic Instrumentation** Custom Resource for applications
- ▶ Support for **Prometheus** receiver, **Kafka** receiver and exporter
- ▶ Scale with the **Target Allocator**
- ▶ **filelog** and **journald** receivers (Developer Preview)





Power monitoring for Red Hat OpenShift

Technology Preview

- ▶ Technology Preview of **Power monitoring for Red Hat OpenShift**
- ▶ Monitor total energy consumed in your cluster during last 24 hours
- ▶ Shows breakdown of the **top power consuming namespaces**
- ▶ Exposes the most **power consuming containers and pods**
- ▶ Based on upstream project **Kepler**



Dashboards

Dashboard: API Performance | Apiserver: kube-apiserver | Period: 5m

etcd	etcd-mixin
Kubernetes / Compute Resources / Cluster	kubernetes-mixin
Kubernetes / Compute Resources / Namespace (Pods)	kubernetes-mixin
Kubernetes / Compute Resources / Namespace (Workloads)	kubernetes-mixin
Kubernetes / Compute Resources / Node (Pods)	kubernetes-mixin
Kubernetes / Compute Resources / Pod	kubernetes-mixin
Kubernetes / Compute Resources / Workload	kubernetes-mixin
Kubernetes / Networking / Cluster	kubernetes-mixin
Kubernetes / Networking / Namespace (Pods)	kubernetes-mixin
Kubernetes / Networking / Pod	kubernetes-mixin
Node Cluster	node-cluster-mixin
Node Exporter / USE Method / Cluster	node-exporter-mixin
Node Exporter / USE Method / Node	node-exporter-mixin
Power Monitoring / Overview	kepler-mixin
Power Monitoring / Namespace	kepler-mixin
Prometheus / Overview	prometheus-mixin



OpenShift Virtualization highlights



Modernize your operations with comprehensive lifecycle and infrastructure management

Public cloud experience for VM creation using Instance Types

- Streamlined VM creation: 3-click GUI experience, tuned for multiple purposes
- Simply specify boot source and InstanceType



Compute
Exclusive

CX series



General
Purpose

U series



GPU
NVIDIA

GN series



Memory
Intensive

M series

Ensure continuity of business critical applications.

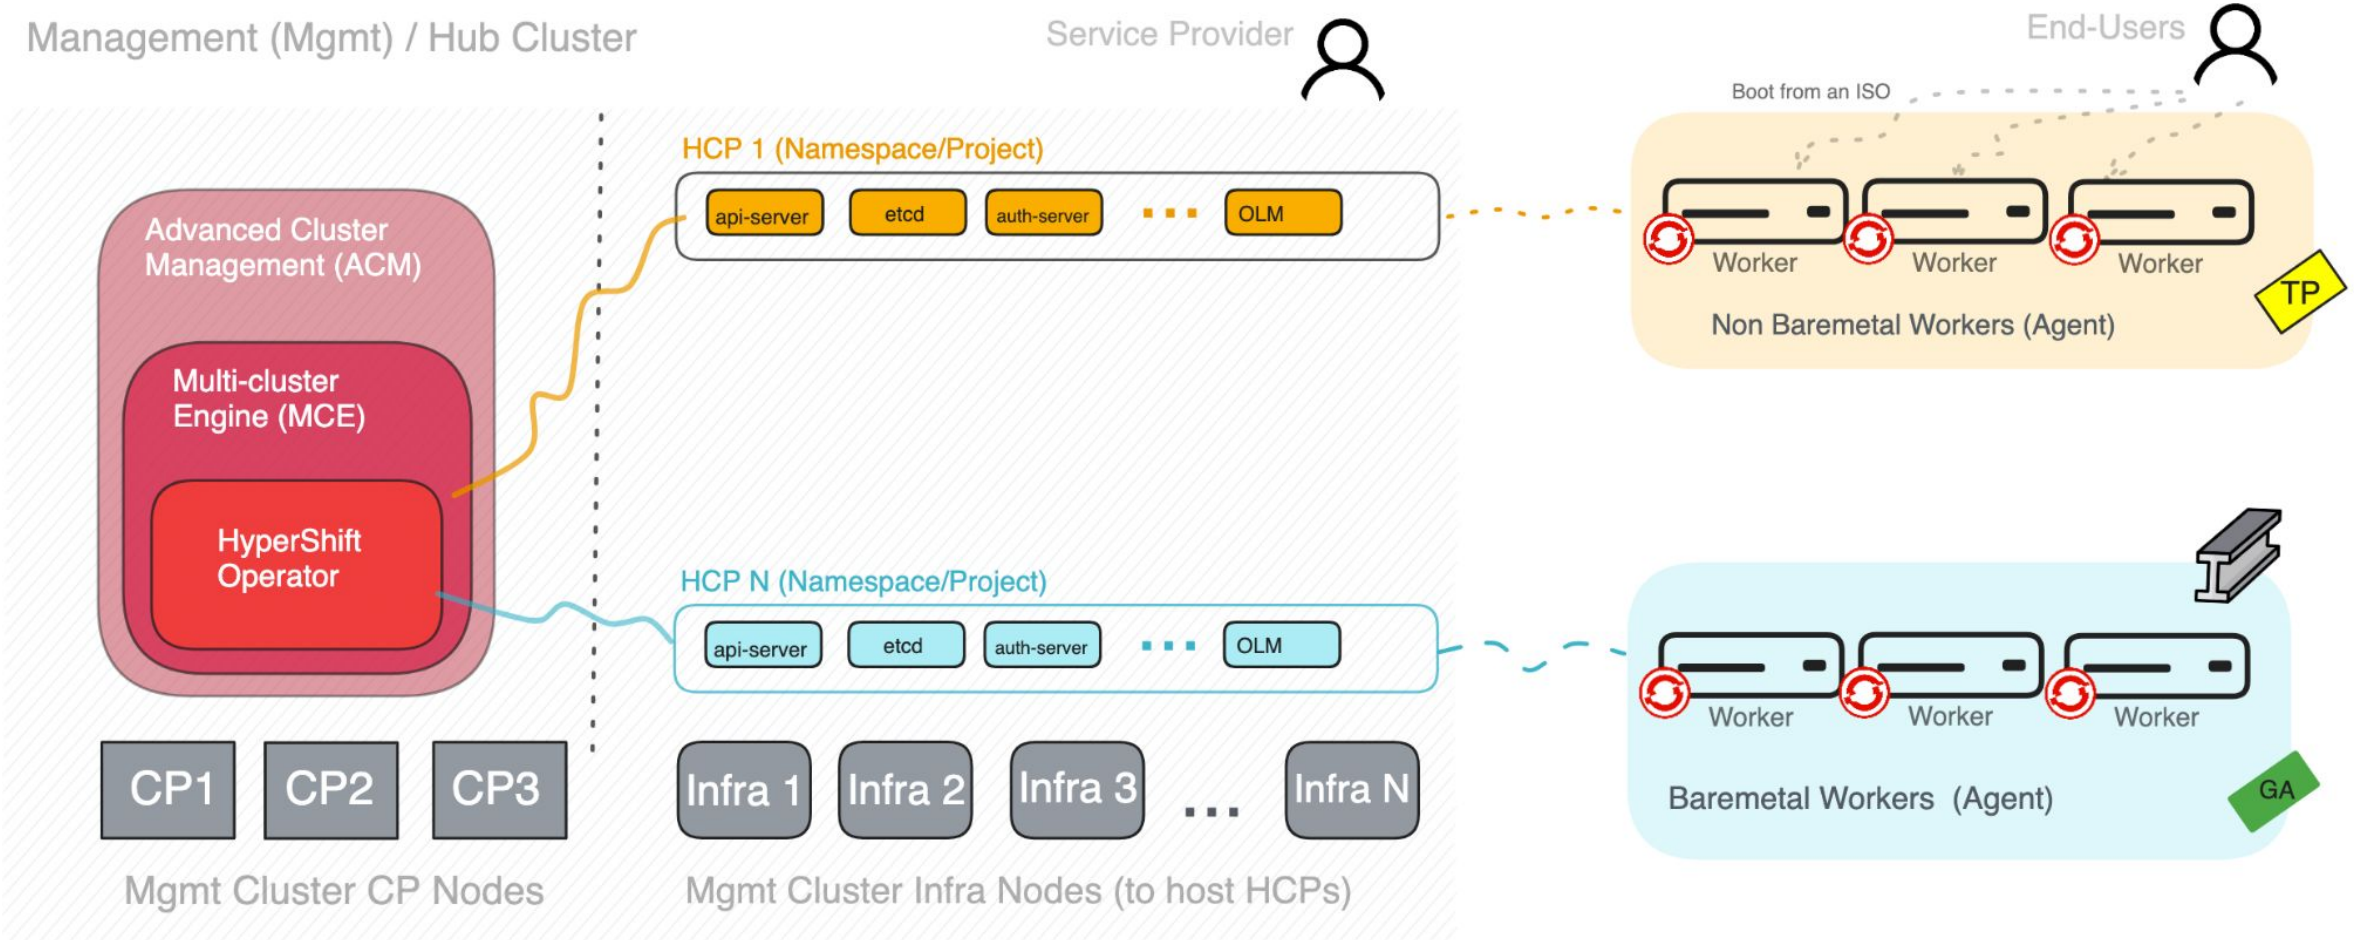
- OpenShift Data Foundation / ACM Metro-DR
 - Support recovery of declarative GitOps virtual machines

Flexibility

- Dynamic reconfiguration - Bridged and SRIOV NIC hotplug
- Micro-segmentation on secondary networks
 - OVN-Kubernetes and ipBlock filtering policies
- Create hosted OpenShift clusters on OpenShift with RHACM.

HCP Non Baremetal Workers (Agent) - TP

Add any node type to your hosted control planes with the Agent provider



Manage at Scale

Red Hat Advanced Cluster Management for Kubernetes

What's New in RHACM 2.10 - **Governance, Risk, and Compliance**

```
apiVersion: policy.open-cluster-management.io/v1beta1
kind: OperatorPolicy
metadata:
  name: quay381policy
spec:
  remediationAction: enforce # or inform
  severity: medium
  complianceType: musthave # or mustnothave
  subscription:
    channel: stable-3.8
    name: project-quay
```

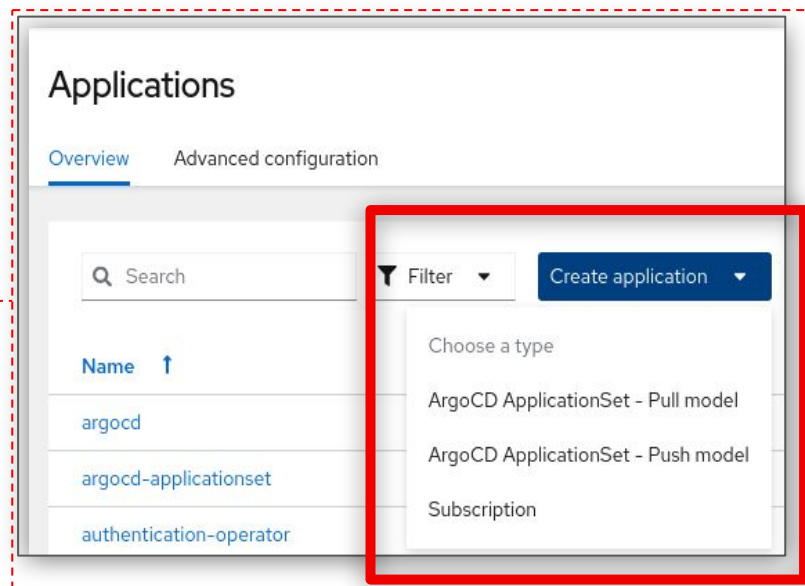
Note: API definition is subject to change upon release

- ▶ Policy compliance history (TP)
 - Track the compliance history for policies across the fleet.
- ▶ Enhanced OLM operator integration (TP)
 - New OperatorPolicy API provides a more native integration for installing and managing OLM operators at scale.
- ▶ Gatekeeper operator uplift to 3.14
 - Alignment with upstream and enhanced configurability of the operator.
- ▶ Improved debugging of policy violations
 - Provide a “diff” of the policy desired state vs actual state to easily understand why a cluster is non-compliant.

Red Hat Advanced Cluster Management for Kubernetes

What's New in RHACM 2.10 - Scale out **application deployments** with OpenShift GitOps; **Deploy and update clusters** with enhanced security; Leverage **fleet observability** for improved operations

- ▶ **Multicluster networking** (submariner) support for bare metal & RHOIC (aka ROKS) (TP)
- ▶ **ApplicationSet pull model** with OpenShift GitOps reaches GA



- ▶ **Cluster Lifecycle** enhancements:
 - RFE: Add authentication for HTTPS osImages content with the Assisted Installer
 - RFE: Allow managed cluster updates to use non-recommended versions
 - RFE: Allow managed OpenShift cluster version to be updated
 - Console support for Hosted Control Planes with OpenShift Virtualization platform
- ▶ **Observability** at scale enhancements:
 - ACM fleet view customization using data from search results
 - Hosted Control Planes hosting cluster capacity monitoring dashboard

Red Hat Quay 3.11

Effective image lifecycle at scale

The screenshot displays the Red Hat Quay 3.11 interface. On the left, a dark sidebar contains navigation options: Overview, Organizations, and Repositories. The main content area shows the 'component-image' repository page, specifically the 'Builds' tab. A 'Build History' table lists recent builds with columns for Build ID, Status, and Triggered by. Overlaid on this is a 'Setup Build Trigger' modal dialog, which is currently on the 'Tagging Options' step. This modal includes a 'Commit-based tagging options' section with checkboxes for 'Tag manifest with the branch or tag name' (checked) and 'Add Latest tag if on default branch'. Below this is the 'Add custom tagging templates' section, which shows 'No tag templates defined' and a field to 'Enter a tag template' with the value '\$[commit_info.short_sha]'. A second modal, 'Enable OIDC Directory Syncing', is also overlaid, providing instructions on how to enter the group name for synchronization and a warning that existing team membership will be revoked.

Repository-level image pruning

Apply policies per image repository to limit storage and artifact growth more nuanced in combination with organization-wide policies.

OIDC team sync


Flexibly map team definitions in Quay to group definitions in OIDC providers to easily manage permissions at scale.

Progress on new UI

Manage container image builds, review audit events and search with expressions in style using dark mode.

ACS 4.4 Highlights

Enhancements and new features



Platform

ACS on ROSA Hosted Control Plane
CO-RE BPF default collection method
BYODB GA



Vulnerability Management

Clair V4 based Scanner v4 GA in ACS



Network Security

Build Time network tools (roxctl): GA

- Generate network policies
- Render connectivity map
- Compare between project versions



Clair v4 based Scanner v4

Consistent and accurate vulnerability reporting across ACS and Quay

[OSV.dev](https://osv.dev) security data

- improves accuracy for language vulnerabilities

Expanded CVE reporting

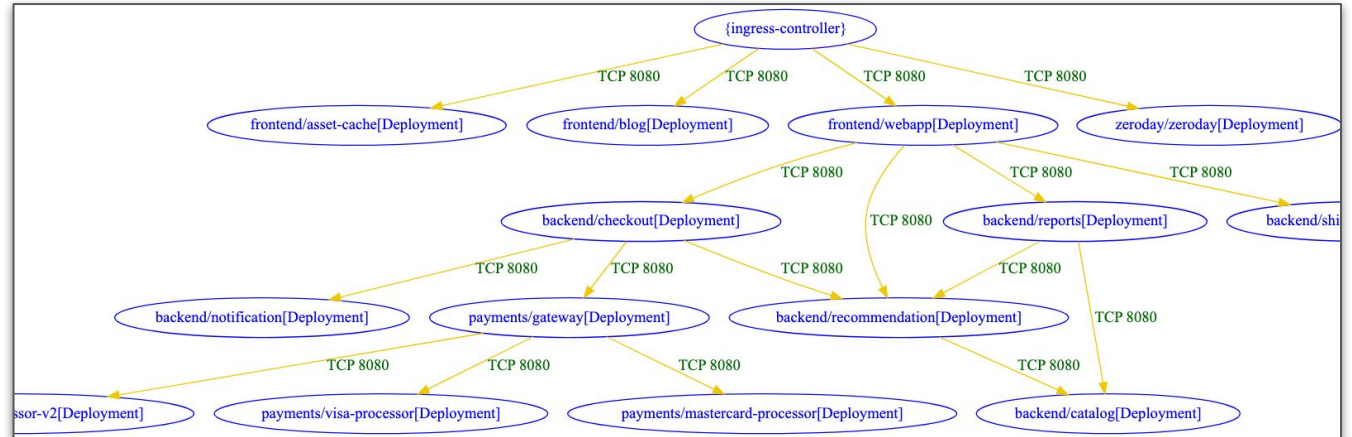
- Include Golang CVEs

Feature	ACS previous scanner (Stackrox)	Clair before consolidation	Clair now and ACS Scanner V4
JavaScript/Node.js (npm - package.json)	✓	✗	✓ (off by default in Clair)
Ruby (Gem)	✓	✗	✓
Golang (binary)	✗	✓	✓
Whiteout files	✓	✗	✓
Oracle Linux	✗ (removed some time ago)	✓	✓
SUSE Linux	✗	✓	✓
Photon Linux	✗	✓	✓
Various other bug fixes/minor support features			

Built Time Network Policy Tools – GA

Allow required network connections, block everything else


- Generate network policies
- Render connectivity map
- Compare between project versions



diff-type	source	destination	dir1	dir2	workloads-diff-info
added	payments/gateway[Deployment]	payments/visa-processor-v2[Deployment]	No Connections	TCP 8080	workload payments/visa-processor-v2[Deployment] added
added	{ingress-controller}	frontend/blog[Deployment]	No Connections	TCP 8080	workload frontend/blog[Deployment] added
added	{ingress-controller}	zeroday/zeroday[Deployment]	No Connections	TCP 8080	workload zeroday/zeroday[Deployment] added

cert-manager Operator 1.13^(*)

Certificate as a Service for Application workloads



API Server & Ingress Controller

Certificates for API Server and Ingress Controller can now be managed through cert-manager .



Multi-Arch Support

ARM64

IBM Z[®] (s390x)

IBM Power[®] (ppc64le)



DNS over HTTPS (DoH)

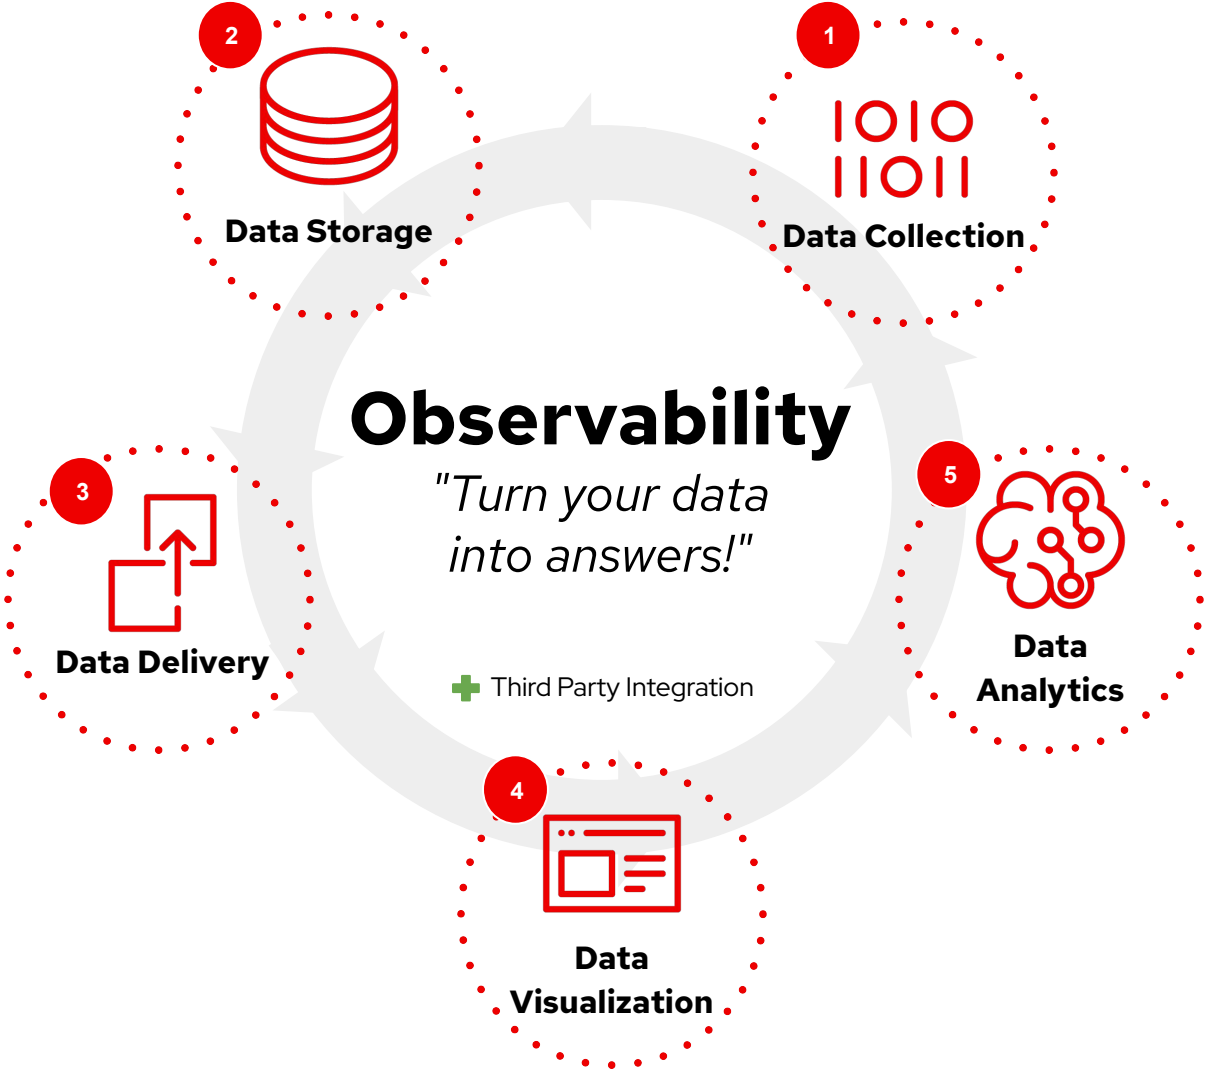
DoH is more secure than plain DNS.

Also useful in proxy environments where traditional DNS resolution


is not available

(*) introduced in OCP 4.14

Observability



What's new is OpenShift Monitoring 4.15?




New Features

- ▶ **Cluster Observability Operator** - enable Red Hat monitoring stack with initial set of features (Tech Preview)
- ▶ Switch to **metrics server** (Tech Preview)
- ▶ Kubelet staleness handling
- ▶ Support **sendExemplars** via UWM remote write
- ▶ Tolerate scrape timestamp jitter



Improvements

- ▶ Improved **query alerts** for User Workload Monitoring (UWM)
- ▶ UserWorkLoad components failures won't degrade core monitoring
- ▶ Alert for **PTP-Operator time** synchronization (Telco)
- ▶ **externalLabels** defined in "cluster monitoring config" to be visible in Alerts triggered in OCP web console



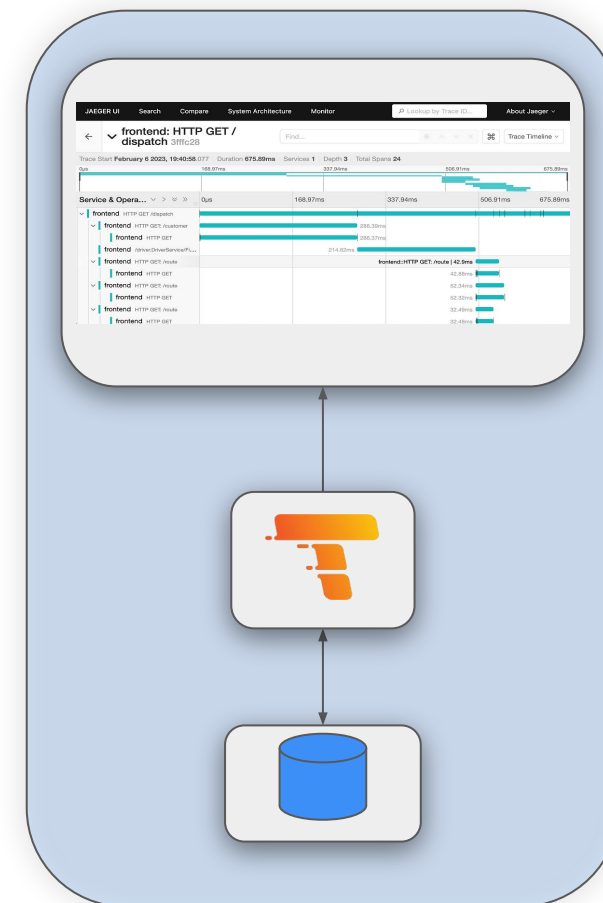
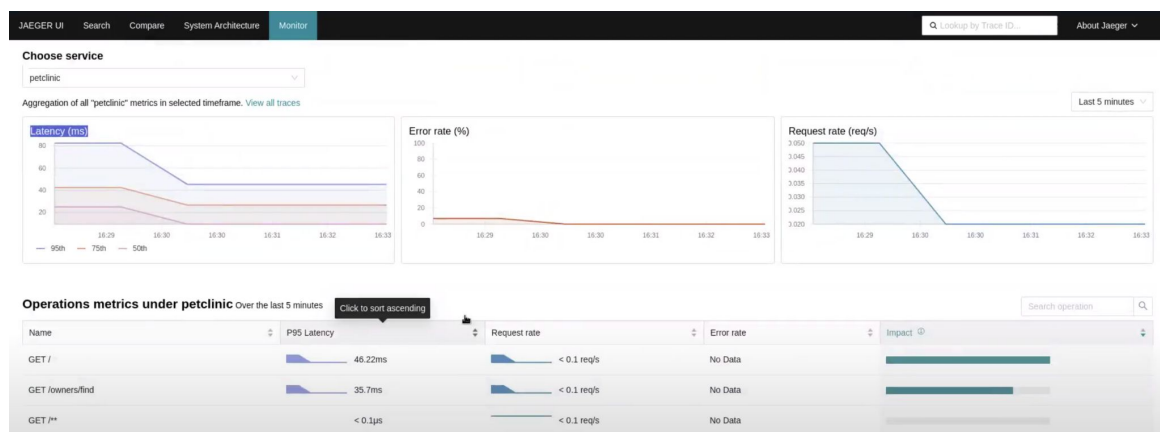
What's new in Logging 5.9?

Logs

- ▶ **OpenTelemetry Data Model** for Vector and Loki (Tech Preview)
- ▶ **Log Forwarding** Integration with Azure
- ▶ AWS and Azure Object Storage Identity Federation with Loki
- ▶ Vector can receive logs from rsyslog
 - For Red Hat Enterprise Linux and Red Hat OpenStack support
- ▶ Display **Log Metrics** in Logs UI - OCP web console
- ▶ **Search across multiple namespaces** in Logs UI
 - OCP web console (Developer Perspective -> Observability UI)

What's new for Distributed Tracing?

- ▶ Tempo operator is now **Generally Available** – easy, cost-efficient, and scalable alternative to Jaeger
- ▶ **Arm** support for Jaeger and Tempo
- ▶ **Jaeger** deprecated in favor of Tempo
- ▶ **Monitor tab** enabled in the Jaeger console
- ▶ **Visualize** Request, Error and Duration (**RED**) metrics
- ▶ **Monolithic deployment** in developer preview



Deployment Validation with Insights

- ▶ **Deployment validation** for on-premise clusters: operational overview of deployment configuration issues.

Suggest to follow best practices for k8s deployments - Apps with no resource limits, wrong pod disruption budget definitions, containers with allowed privilege escalation, network policy violations and more.

- ▶ **Conditional data gathering:** send data relevant only to debugging of an issue. If no known issue is detected, bare bone data collection only. Reduced footprint.
- ▶ **Fleet Insights** in Red Hat ACM: Display summary view of most important information about clusters

The screenshot shows the Red Hat Hybrid Cloud Console interface. The main content area displays the 'Advisor clusters' page for a cluster named 'rosa-g4vz5'. The UI includes a search bar, a 'Filter by description' dropdown, and a table of recommendations. The table has columns for 'Description', 'Modified', 'First impacted', and 'Total risk'. A single recommendation is shown: 'Workloads are prone to security breaches when no NetworkPolicy matches them', with a 'Moderate' risk level. Below the table, there is a 'Detected issues' section with a detailed description of the problem and a 'Steps to resolve' section listing three affected namespaces with their respective IDs.

Description	Modified	First impacted	Total risk
Workloads are prone to security breaches when no NetworkPolicy matches them	6 months ago	3 months ago	Moderate

Detected issues

The Deployment Validation Operator has detected workloads that are not selected by any NetworkPolicy. Network access to such workloads is unrestricted which makes them prone to security breaches. It is a best practice to create NetworkPolicy resources to explicitly define the minimal required network access for each workload.

Kube-linter check: [non-isolated-pod](#)

Steps to resolve

Red Hat recommends that you [create network policies](#) for the affected workloads.

A total of 3 objects in 2 namespaces are affected:

- Namespace UID: 0621ba24-070c-407e-b8d4-a6da8c9e6fc9, Kind: Deployment, UID: a32b52ec-a193-44d9-b971-86cf45a7f00f
- Namespace UID: 0621ba24-070c-407e-b8d4-a6da8c9e6fc9, Kind: Pod, UID: 2f785cce-e07a-4bb2-97ba-83c88e22a924
- Namespace UID: 8664baf1-2052-47ae-be59-f2248df58cdd, Kind: Deployment, UID: 57f63721-467d-41c2-ac6f-596759ba3ac3

Note: Red Hat avoids gathering and processing namespace and resource names as these may reveal confidential information. Namespaces and resources are identified by their UIDs instead. You can use in-cluster commands like the ones below to translate UIDs of affected resources to their names.

Console

Dynamic Plugin Framework

OCP 4.15 Dynamic Plugin Enhancements

- ▶ New DetailPage Extension
- ▶ CronTab Examples Added:
 - Annotation Modal
 - Label Modal
 - Delete Modal
- ▶ Support for both PF4 & PF5 design library



The screenshot illustrates the dynamic plugin framework in the OpenShift console. The main interface shows the 'CronTab details' page for a CronTab named 'my-new-cron-object' in the 'default' namespace. The page includes fields for Name, Namespace, Labels, CronSpec, Image, Replicas, and Created at. Three modal windows are overlaid on the page, demonstrating the framework's capabilities:

- Edit labels:** A modal window for editing labels for the CronTab. It shows a list of labels for 'my-new-cron-object' and provides a 'Save' button.
- Edit annotations:** A modal window for editing annotations. It displays a table with 'Key' and 'Value' columns, each containing a 'Text box' input. It includes an 'Add more' button and 'Save' and 'Cancel' buttons.
- Delete CronTab?:** A warning modal window asking for confirmation to delete the CronTab 'my-new-cron-object' in the 'default' namespace. It features a red 'Delete' button and a blue 'Cancel' button.

Developer Tools Update

Developer Tools Update

Check out:

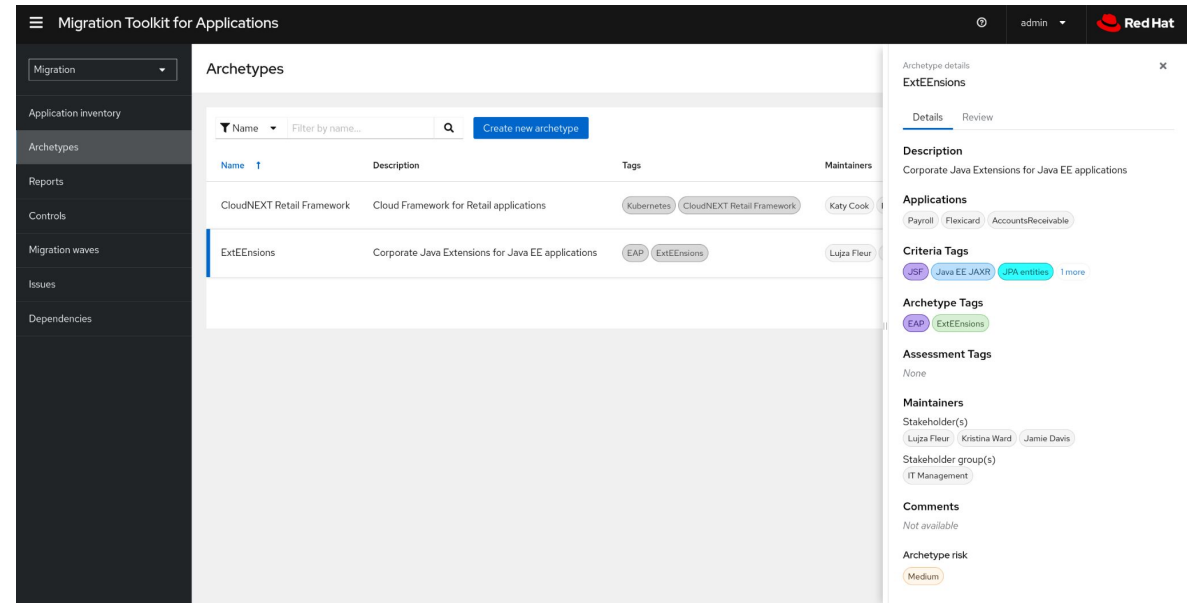
- ▶ The **Developer Perspective** in **OpenShift Console** includes new dynamic plugin-based dashboard for OpenShift Pipelines, access to additional information in OpenShift Pipelines for Trusted Software Supply Chain, enhancements for build strategy and creating Serverless functions.
- ▶ **Podman Desktop 1.7**, includes ways to create local clusters with the OpenShift Local extension, UI to manage Pods, Services, Deployment and Routes and manage local kubernetes contexts. The entire onboarding experience has been improved, making it easier for developers to setup their environments locally.
- ▶ **OpenShift Toolkit IDE** extension, includes a new UI for **Helm** charts with Helm repo management. Allows users to do remote container development for OpenShift and Kubernetes application resource management. It supports **OpenShift Serverless 1.32** with remote deployment using tekton and on-cluster builds. This extension is available for **Visual Studio Code** and **IntelliJ**.
- ▶ **Developer Hub v1.0 GA** offers software templates and plugins for OpenShift deployments, monitoring, accessing pipeline runs, Quay container images, and viewing clusters from OCM.

Watch out for a separate DEVELOPER EDITION presentation coming the next weeks!!
developers.redhat.com

Runtimes

Cloud Native Runtimes

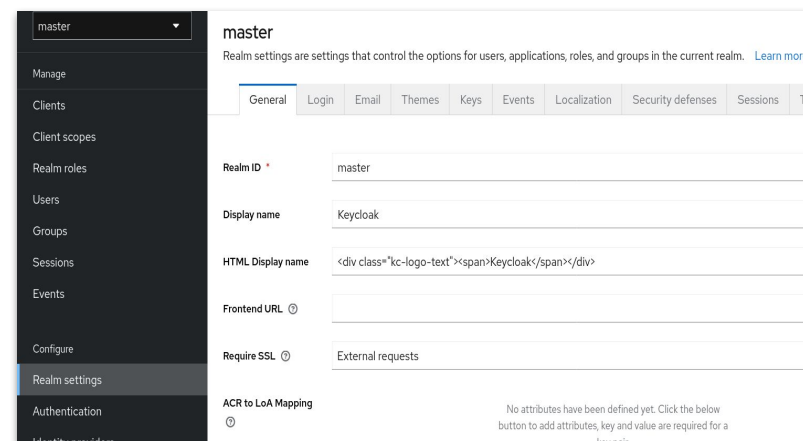
- ▶ **Migration Toolkit for Applications 7.0**
 - ▶ Multi-language, new rules syntax, automated classification, dynamic reports
- ▶ **Quarkus 3.8** is coming soon:
 - ▶ New OpenSearch Dev Service
 - ▶ Redis 7.2 support
 - ▶ Java 21 support - including Virtual Threads
 - ▶ ARM native build support
- ▶ **Node.js 20** container images for OpenShift now available
- ▶ **Java 21** builder & runtime container images for OpenShift now available
- ▶ **Spring Boot 3.1.x, 3.2.x** tested & verified runtimes on OpenShift



MTA: Automated classification via Archetypes

Red Hat build of Keycloak

- ▶ **Red Hat build of Keycloak 22: Cloud-friendly Identity Access Management solution**
 - ▶ Built on Quarkus: Kube-native, faster, reduced resource consumption
 - ▶ Focus on usability, better UX
 - ▶ Seamless User Experience, Login, Logout, Self-registration, User Account Management
- ▶ **Enterprise single sign-on capabilities**
 - ▶ Strong Authentication, MFA, Passwordless authentication
 - ▶ Enhanced security, FIPS compliance (critical to NAPS, FedRamp)
 - ▶ Identity Brokering, authenticating with external OpenID Connect or SAML Identity Providers
- ▶ **Container images and zip distros available**
- ▶ **Migration guide & tooling for RH-SSO users**



New Administrator UI

```

Keycloak - Open Source Identity and Access Management
Find more information at: https://www.keycloak.org/docs/latest

Usage:
kc.sh [OPTIONS] [COMMAND]

Use this command-line tool to manage your Keycloak cluster.
Make sure the command is available on your "PATH" or prefix it with "./" (e.g.:
"./kc.sh") to execute from the current folder.

Options:
-cf, --config-file <file>
    Set the path to a configuration file. By default, configuration properties are
    read from the "keycloak.conf" file in the "conf" directory.
-h, --help
    This help message.
-v, --verbose
    Print out error details when running this command.
-V, --version
    Show version information

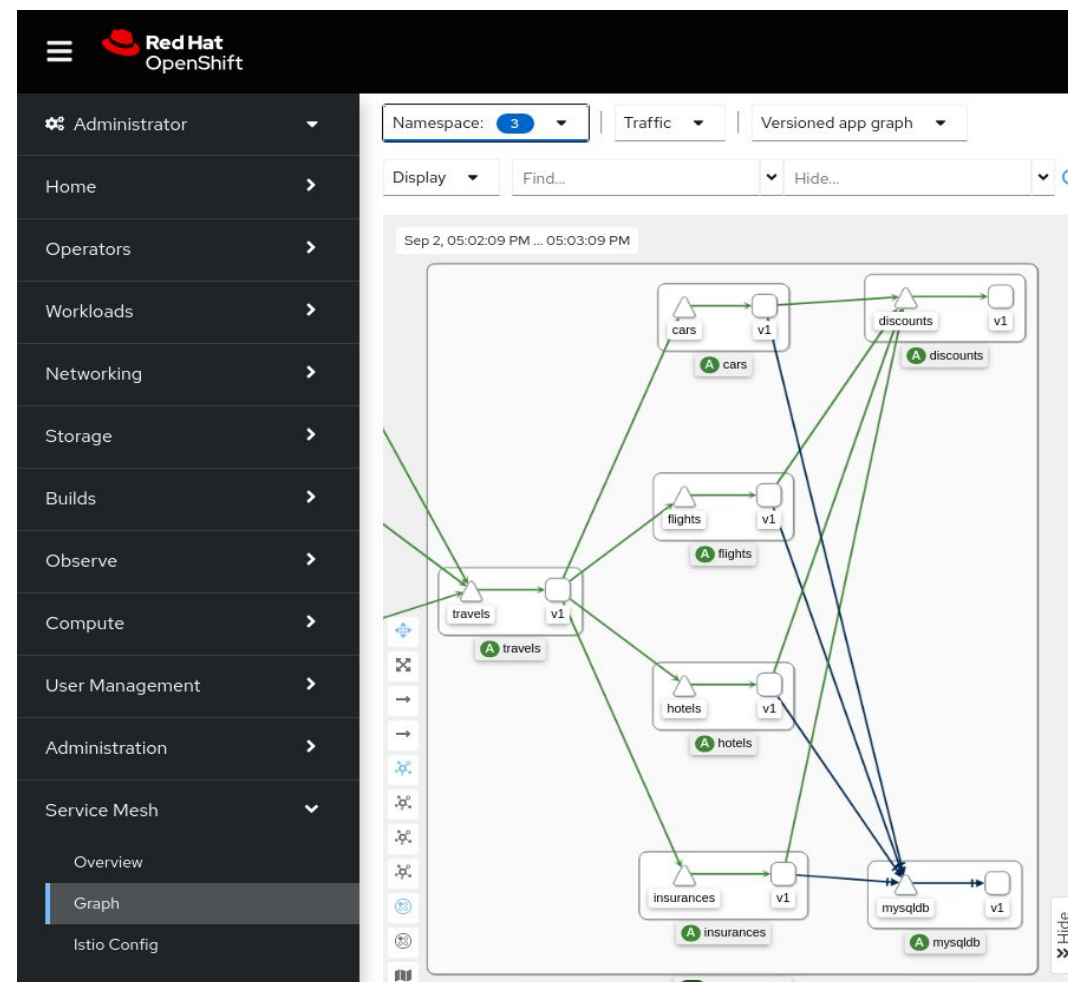
Commands:
build
    Creates a new and optimized server image.
start
    Start the server.
start-dev
    Start the server in development mode.
export
    Export data from realms to a file or directory.
import
    Import data from a directory or a file.
show-config
    Print out the current configuration.
tools
    Utilities for use and interaction with the server.
completion
    Generate bash/zsh completion script for kc.sh.
  
```

User-friendly CLI

Platform Services

OpenShift Service Mesh

- ▶ OpenShift **Service Mesh 2.5**:
 - ▶ Based on **Istio 1.18** and **Kiali 1.73**
 - ▶ GA Support for Service Mesh on **Arm** clusters
 - ▶ GA of OpenShift Service Mesh Console plugin
 - ▶ Certificate Revocation Lists (CRLs) for gateways
 - ▶ GA of zipkin, opentelemetry and envoyOtelAls extension providers
 - ▶ Support for tracing with the **Tempo** operator
 - ▶ Developer preview of IPv4/IPv6 Dual-Stack
- ▶ [Kiali on RH Developer Hub Dev Preview](#)
- ▶ “Sail Operator” - Updated Developer Preview of OpenShift **Service Mesh 3**:
 - ▶ See [update blog post](#) - Dec 2023



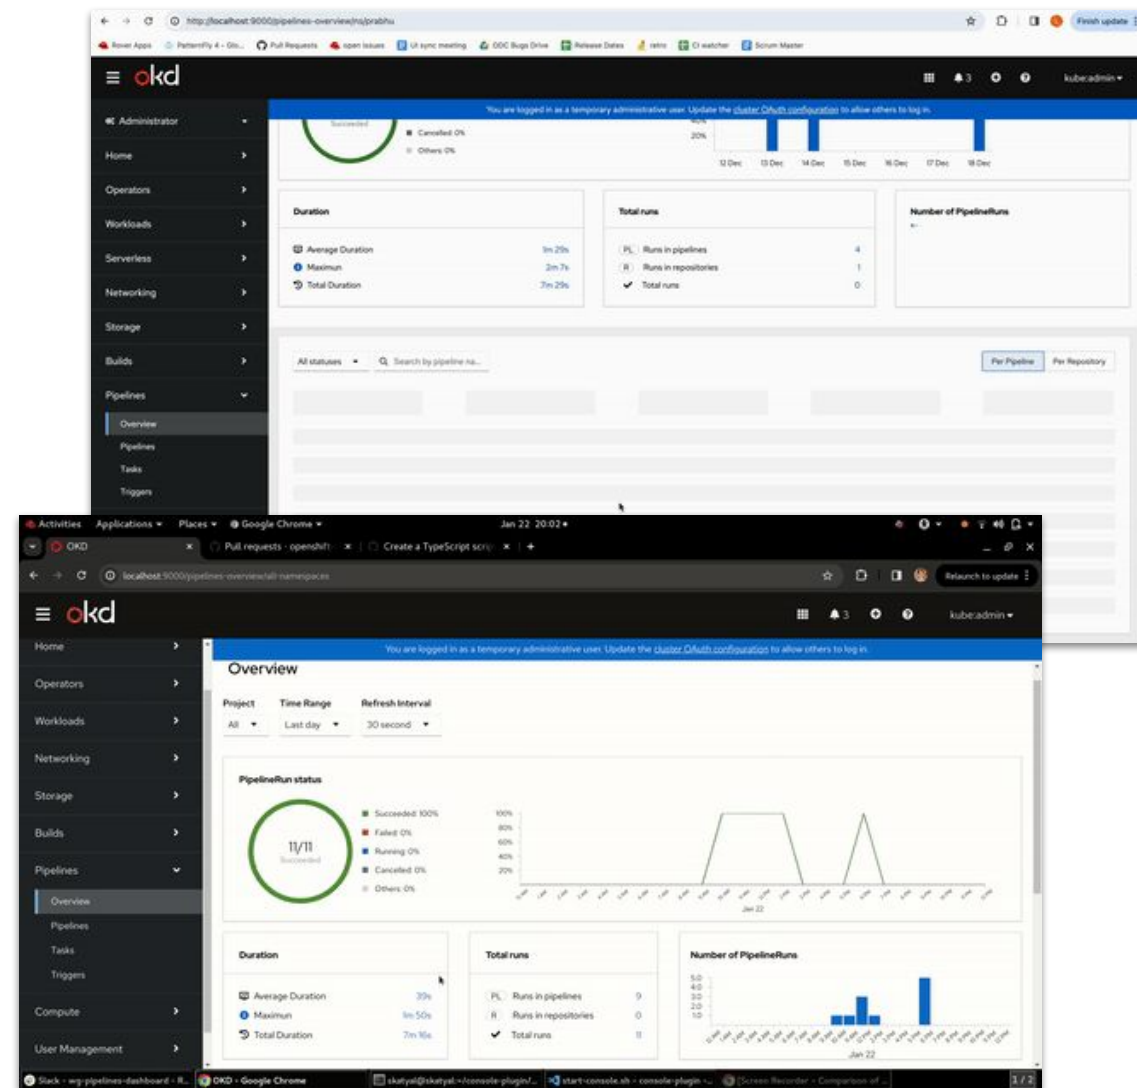
OpenShift GitOps

- ▶ OpenShift GitOps **1.12** coming in March
 - ▶ Includes **Argo CD 2.10**
 - ▶ Small footprint GitOps + MicroShift support - TP
 - ▶ Argo CD CLI support - TP
 - ▶ Notifications goes GA
 - ▶ OpenShift Routes support in Rollouts - TP
- ▶ OpenShift GitOps **1.11** released December includes
 - ▶ Includes **Argo CD 2.9**
 - ▶ Dynamic shard rebalancing - TP
 - ▶ Gitlab SCM provider now supports self-signed certs

OpenShift Pipelines

OpenShift Pipelines 1.13 released, 1.14 coming in February

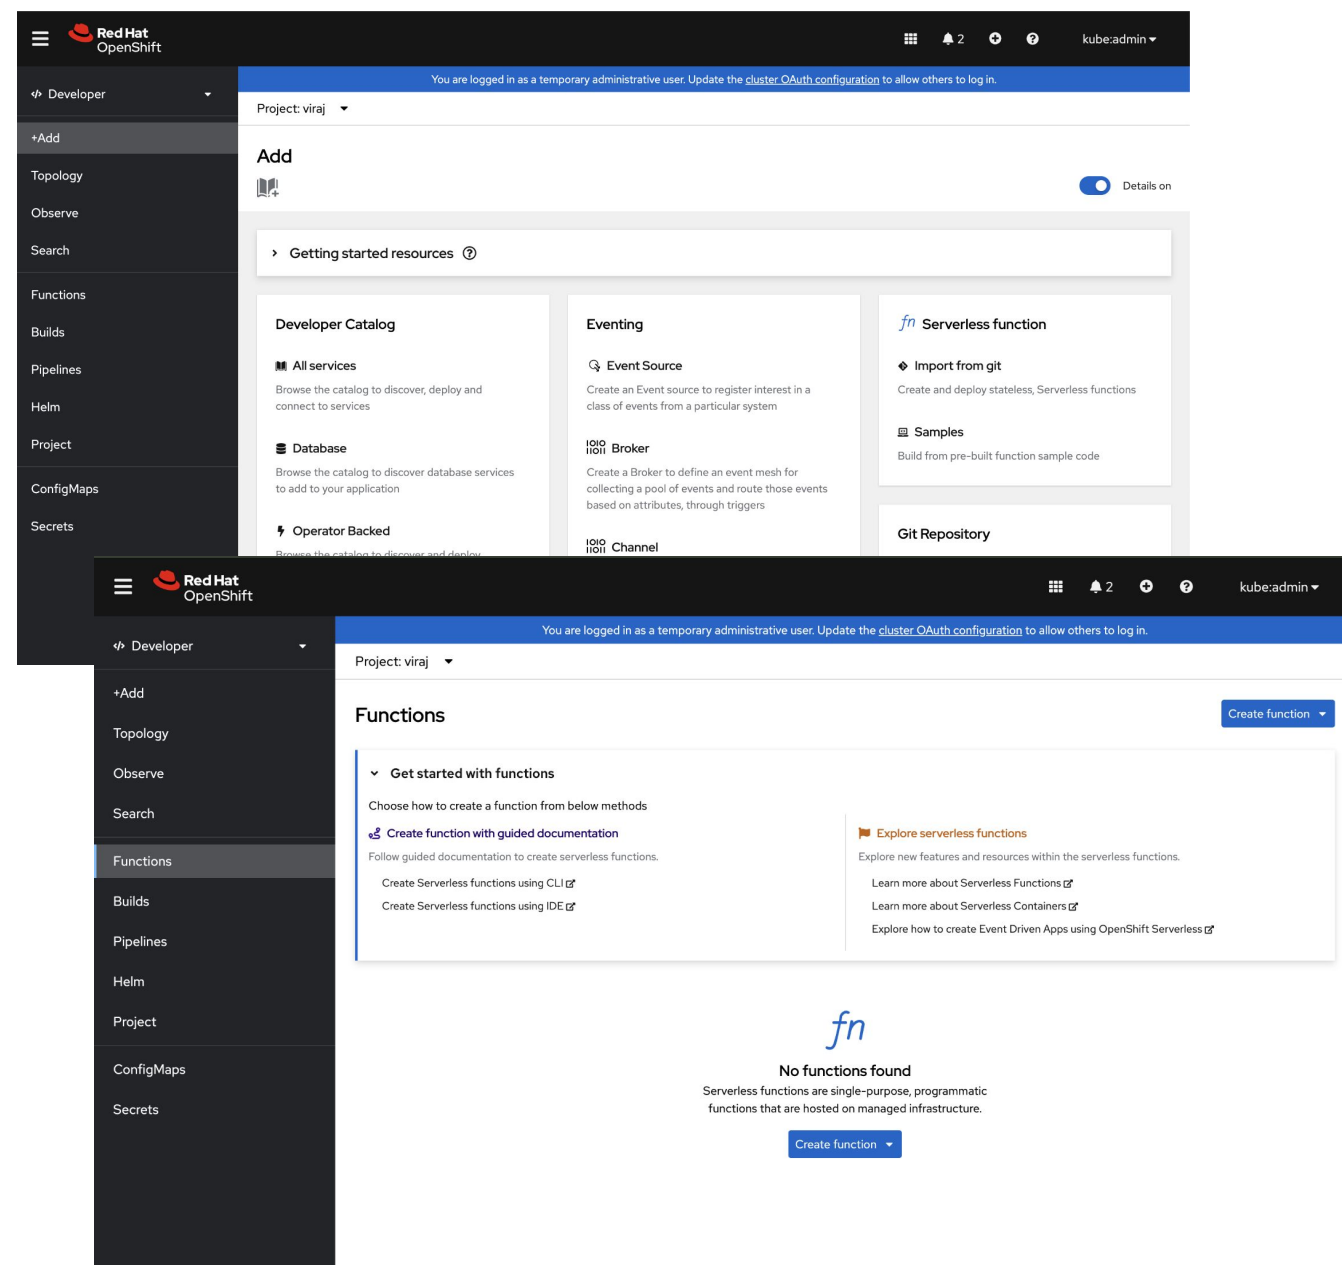
- ▶ Tekton Results released in **Tech Preview**
 - ▶ Includes support for external DB and storage
 - ▶ New API for summary of logs with various filter
- ▶ **Pipelines As Code**
 - ▶ Multiple GitHub Apps support
 - ▶ Remote pipeline support in PAC resolver
- ▶ Validation of secrets store CSI driver and Tekton integration for making RHEL entitlements available in buildah pods
- ▶ **Tekton Controller performance** testing and recommendation for enabling HA of controllers for performance improvements
- ▶ **Console Improvements**
 - ▶ Tekton Results integration with OCP console, Pipelines dynamic plugin for a CI centric dashboard
 - ▶ Vulnerability flags and signed PR indicators



OpenShift Serverless

Key Features & Updates

- ▶ Serverless 1.32 : Update to Knative 1.11
- ▶ Platform Agnostic (Tier 2) support
- ▶ Serverless functions
 - ▶ Configuration of PVC
 - ▶ Dev console presence
- ▶ Multi Tenancy with ServiceMesh - TP
 - ▶ Serving and Eventing
- ▶ Single Node OpenShift support
- ▶ Enhanced Security and Performance
 - ▶ More configuration option



Installer Flexibility

OpenShift 4.15 Supported Providers


Installation Experiences



Outposts
Wavelength
Local Zones



(Tech Preview)




IBM Z and
IBM LinuxONE



Bare Metal



Cloud (Tech Preview)



Automated

Installer Provisioned Infrastructure

- Auto-provisions infrastructure
- *KS like
- Enables self-service



Full Control

User Provisioned Infrastructure

- Bring your own hosts
- You choose infrastructure automation
- Full flexibility
- Integrate ISV solutions



Interactive - Connected

Assisted Installer

- Hosted web-based guided experience
- Agnostic, bare metal, vSphere and Nutanix
- ISO driven



Local - Disconnected

Agent-based Installer

- Disconnected / air-gapped
- Automatable installations via CLI
- Bare metal, vSphere, SNO
- ISO driven

Installation Highlights for Cloud Providers



- ▶ Custom MTU at install time for AWS
- ▶ Support for AWS Wavelength
- ▶ AWS Outposts graduates to General Availability
- ▶ Support for Tel Aviv AWS Region



- ▶ User managed DNS support for GCP
- ▶ Out-of-tree cloud controller Manager (CCM) for GCP graduates to General Availability



- ▶ OpenShift on Oracle Cloud Infrastructure with Virtual Machines is now Technology Preview



- ▶ Installation on restricted networks for IBM Cloud VPC
- ▶ User managed encryption key for IBM Cloud VPC



- ▶ User managed encryption key for Azure Storage Account

Agent-Based Installer

Day-1 Bare metal hosts' BMC config

No day-2 BMC config, add BareMetalHosts like you do in IPI for MAPI integration at install time

Improved bare metal compatibility

Add root device hints, host network, and other config directly to the install-config.yaml

Configure vSphere credentials on day-1

No day-2 config needed for vSphere, add your vCenter credentials to install-config.yaml

Platform External Support

Allow easier new provider integrations following the platform external model.

<pre>platform: baremetal: hosts: bmc: username:</pre>	The username for the BMC.
<pre>platform: baremetal: hosts: bmc: password:</pre>	Password for the BMC.
<pre>platform: baremetal: hosts: bmc: address:</pre>	The URL for communicating with the host's BMC controller. The address configuration setting specifies the protocol. For example, <code>redfish+http://10.10.10.1:8000/redfish/v1/Systems/1234</code> enables Redfish. For more information, see "BMC addressing" in the "Deploying installer-provisioned clusters on bare metal" section.

More at the [Agent-based Installer documentation](#)

OpenShift on vSphere

Updated minimum privileges on vSphere

- ▶ Set granular permissions while staying secure and functional in IPI and UPI installations

ControlPlaneMachineSets (Tech Preview)

- ▶ Simplify management of your cluster and improve its reliability with ControlPlaneMachineSets
- ▶ Available for vSphere in Tech Preview in OpenShift 4.15

q

- ▼ Installing on vSphere
 - Preparing to install on vSphere
 - ▶ [Installer-provisioned infrastructure](#)
 - ▼ User-provisioned infrastructure
 - vSphere installation requirements
 - Installing a cluster
 - Installing a cluster with network customizations
 - Installing a cluster in a restricted network
 - Assisted Installer
 - Agent-based Installer
 - Installing a three-node cluster
 - Uninstalling a cluster
 - Using the vSphere Problem Detector Operator
 - Installation configuration parameters
 - ▶ Installing on any platform

Required vCenter account privileges

To install an OpenShift Container Platform cluster in a vCenter, your vSphere account must include privileges for reading and creating the required resources. Using an account that has global administrative privileges is the simplest way to access all of the necessary permissions.

⤴ Collapse all

▼ Roles and privileges required for installation in vSphere API

vSphere object for role	When required	Required privileges in vSphere API
vSphere vCenter	Always	Cns.Searchable InventoryService.Tagging.AttachTag InventoryService.Tagging.CreateCategory InventoryService.Tagging.CreateTag InventoryService.Tagging.DeleteCategory InventoryService.Tagging.DeleteTag InventoryService.Tagging.EditCategory InventoryService.Tagging.EditTag Sessions.ValidateSession StorageProfile.Update StorageProfile.View
vSphere vCenter Cluster	If VMs will be created in the cluster root	Host.Config.StorageResource.AssignVMToPool VApp.AssignResourcePool VApp.Import VirtualMachine.Config.AddNewDisk

OpenShift on Nutanix

Fault tolerant deployments using multiple Nutanix Prism Elements (Nutanix clusters)

- ▶ Control plane and compute nodes can be distributed across multiple Nutanix clusters for high availability.
- ▶ A failure domain represents a Prism Element instance that is available to OpenShift machine pools during and after installation.
- ▶ Available in IPI-deployed clusters

```
platform:
  nutanix:
    failureDomains:
      - name: <failure_domain_name>
        prismElement:
          name: <prism_element_name>
          uuid: <prism_element_uuid>
        subnetUUIDs:
          - <network_uuid>
    controlPlane:
      nutanix:
        failureDomains:
          - failure-domain-1
          - failure-domain-2
          - failure-domain-3
    compute:
      nutanix:
        failureDomains:
          - failure-domain-1
          - failure-domain-2
```

OpenShift on Bare Metal

Configure hardware RAID for Dell nodes via Redfish

- ▶ You can now configure hardware RAID on Dell hosts from the OpenShift IPI installer.
- ▶ This adds support for Dell hardware, along with existing support for Fujitsu.

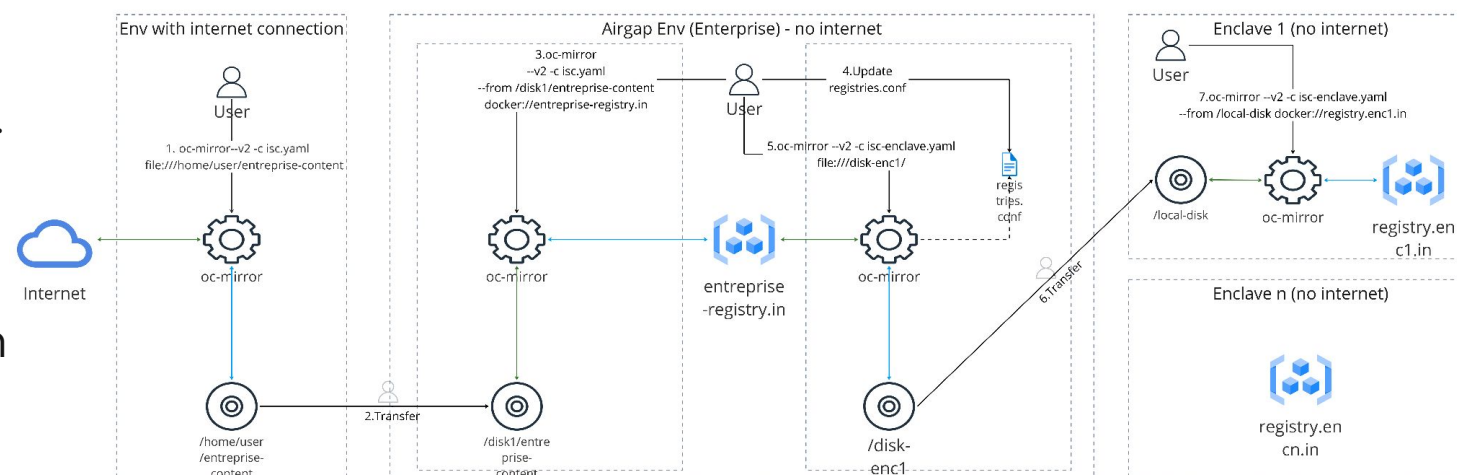
```
$ vim clusterconfigs/openshift/99_openshift-cluster-api_hosts-*.yaml

spec:
  raid:
    hardwareRAIDVolumes:
      - level: "0"
        name: "sda"
        numberOfPhysicalDisks: 1
        rotational: true
        sizeGibibytes: 0
```


OpenShift oc-mirror plugin (Developer Preview)

oc-mirror enclaves Developer Preview

- ▶ Mirror images to and from disconnected environments (enclaves).
- ▶ Save time, effort and bandwidth by mirroring images centrally and only transferring the necessary ones to each enclave.
- ▶ Introduced as Developer Preview in OpenShift 4.15 for testing it.
- ▶ Tech Preview planned for OpenShift 4.16.



Developer Preview documentation: https://github.com/openshift/oc-mirror/blob/main/docs/enclave_support.md

OpenShift On OpenStack 4.15 Update

▶ Dual Stack Support GA

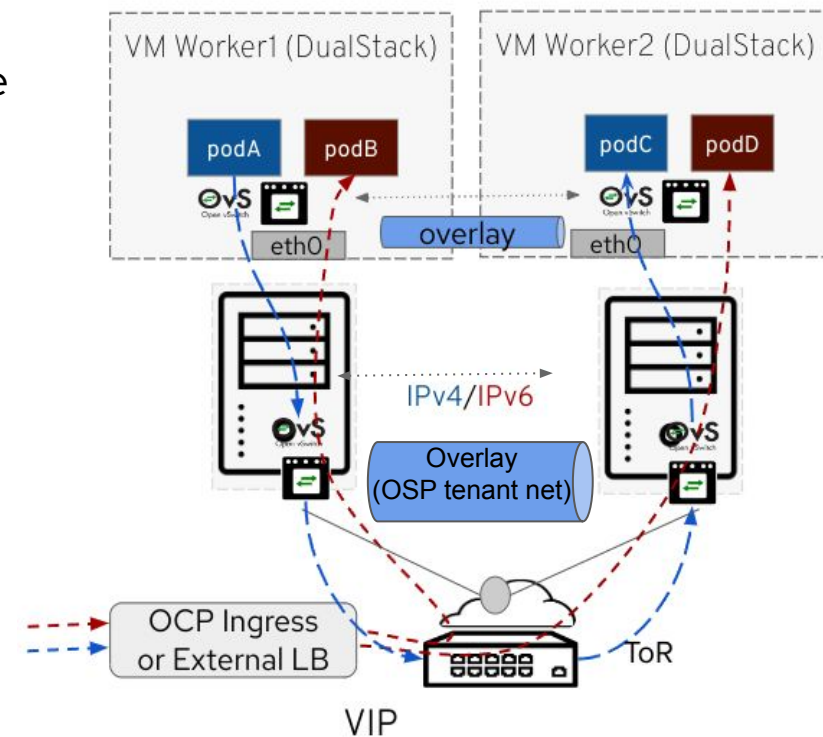
- Supported in Both UPI and IPI deployed clusters
- Requires underlying Stack to be deployed with Dual Stack as a prerequisite
- Supporting both IPv4 and IPv6 as primary stacks

▶ Custom MachineSets for Fast Datapath nodes Tech Preview

- Automates the creation of Fast datapath instances via MachineSets for:
 - SRIOV-DPDK enabled flavor
 - OVS-DPDK (VirtIO) enabled flavor
- Highly sought after by Red Hat Openstack Telco customers

▶ Kuryr CNI EOL

- No new installation are supported with Kuryr
- A migration from kuryr to OVN-K is captured in the [Documentation](#)



CoreOS Updates

RHEL CoreOS & MCO

- ▶ ARM 64k page kernel extension
- ▶ **Tech Preview** support for primary disk using the iSCSI Boot Firmware Table driver (`iscsi_ibft`)
- ▶ **Dev Preview** custom first boot images (RAW disk format)
- ▶ Improvement to custom machineconfigpool config merging logic
- ▶ **Tech Preview** enhanced MCO state reporting

Control Plane Updates

Enabling /dev/fuse

With annotation `io.kubernetes.cri-o.Devices: "/dev/fuse"`

The Need for /dev/fuse in Containers

Within a container, access to certain host devices is restricted for security and isolation reasons. However, there are scenarios where a container might need to interact with specific host devices. One such device is /dev/fuse, used for FUSE (Filesystem in Userspace) operations.

What is new in 4.15

Customer can use annotation `io.kubernetes.cri-o.Devices: "/dev/fuse"` to grant access to the /dev/fuse device on the host.

Benefit

Customers are now able to run podman or buildah with fuse-overlayfs instead of vfs resulting in faster build in pods.

Deprecating ICSP

*We are not removing ICSP support but encouraging customer to use IDMS instead

ImageContentSourcePolicy (ICSP) and Image Digests Mirror Service (IDMS) in OpenShift are used to manage and control the sources and integrity of container images in an OpenShift environment. They play crucial roles in ensuring that the right images are used in the right places, especially in restricted or highly-controlled environments.

- ▶ We are trying to deprecate the use of ICSP and encourage customer to use IDMS instead
- ▶ ICSP and IDMS will be both supported in a cluster together
- ▶ Migration steps from ICPS to IDMS are available in OpenShift documentation

Prevent must-gather from filling up master node

The Need for limiting size of must-gather logs

Must-gather is a tool to collect system configuration that can be sent to Red Hat for further analysis . Must-gather runs on a control plane node and based on the how big is the cluster (number of resource , configurations) it might fill up the storage space of Master node due to the size of log it collects.

What's new in 4.15

In Openshift 4.15, we have added a customer configured limit to the size of must-gather logs which is set to default value to 30% of the total volume size out of box.

Benefit

This will prevent must-gather logs from filling up the master node.

Selective Workload Monitoring with Vertical Pod Autoscaler (VPA)

For Efficient Resource Management in Large OpenShift Clusters

The Need for selective workload monitoring in VPA

In an OpenShift cluster with lot of workloads . When user deploys VPA to scale up few selective workloads. The VPA recommender by default watches all workloads in that cluster. This causes VPA recommender to use lot of memory and might stop working.

What's new in 4.15

Customers can configure the VPA Operator to monitor only those workloads that are being managed by a VPA CR.

Benefit

By configuring the Operator to monitor only selected workloads with a VPA CR, customers can save on CPU and memory resources.

Networking & Routing

Red Hat OpenShift Networking Enhancements

- **Removal of openshift-sdn CNI option for all newly-installed clusters at 4.15+ ***

- The openshift-sdn CNI plug-in will no longer be an install-time option for newly installed 4.15+ clusters across installation options.
- Note that customer clusters currently using openshift-sdn that upgrade to 4.15 or 4.16 with openshift-sdn will remain fully supported.

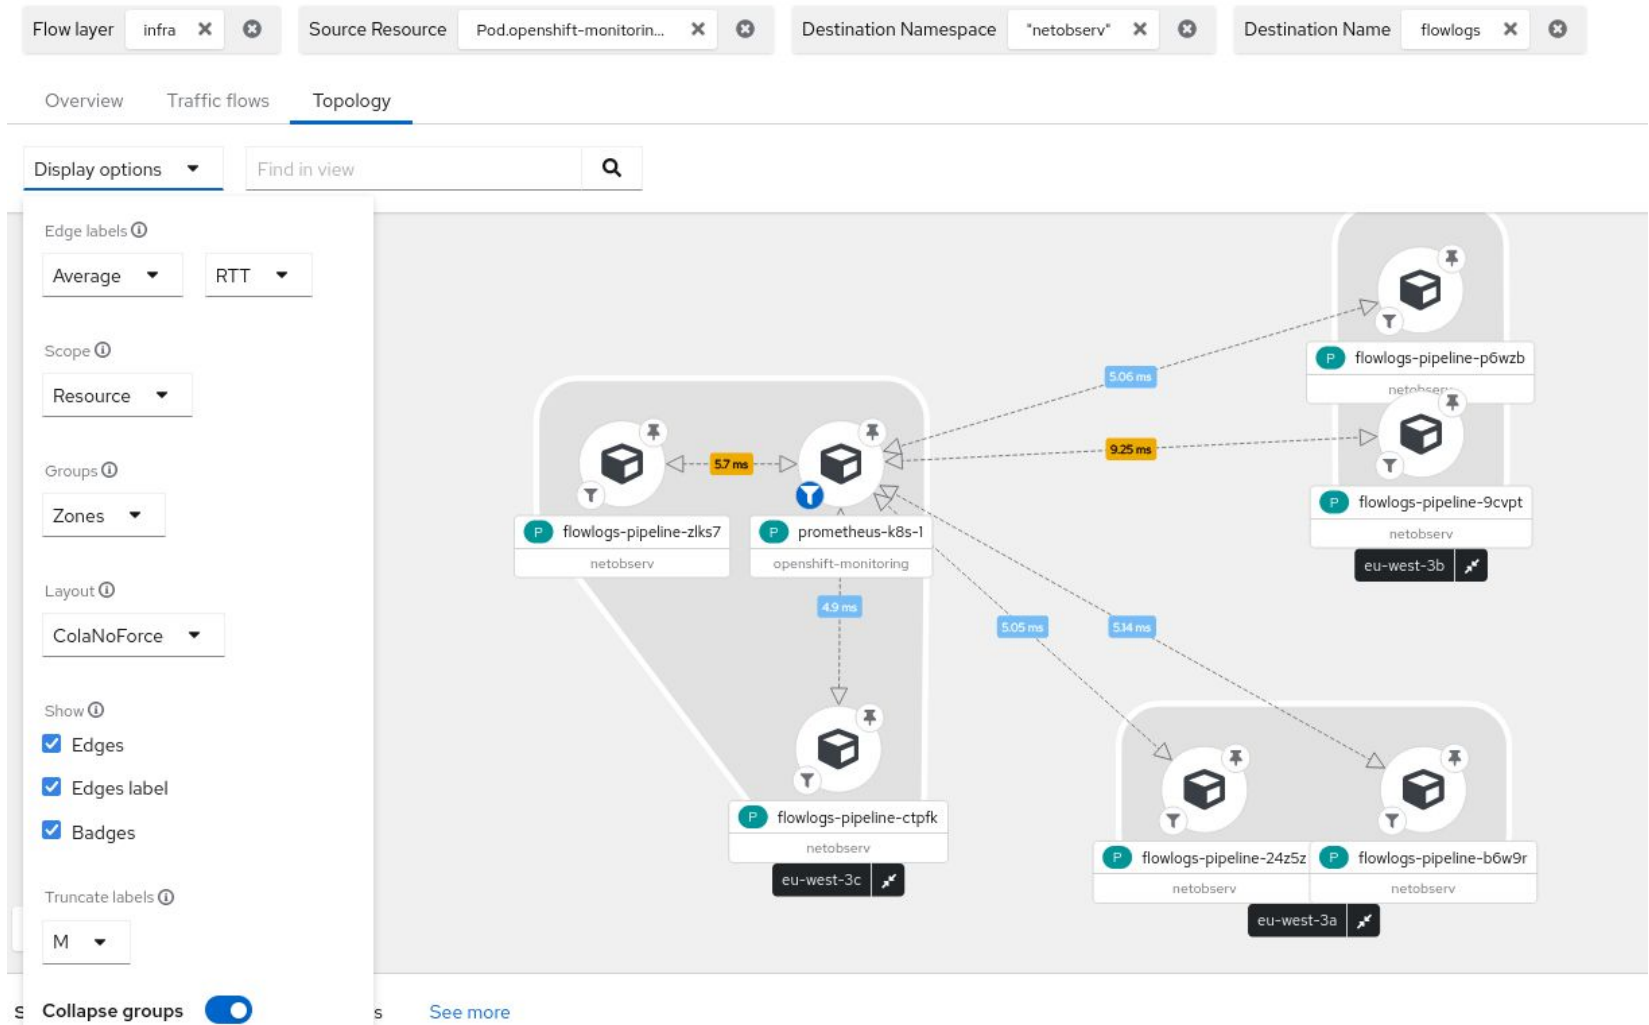
*IBM POWER platforms are exempt until 4.16.

- **Support Kubernetes MultiNetwork Policy [GA]**

- Multi-Network in Kubernetes deployed via Multus provides
 - Enhanced Tenant Isolation
 - Regulatory Compliance
 - Support advanced Network topologies
 - IPv4 and IPv6 (dual stack)
 - SR-IOV kernel CNI
 - macvlan CNI
- Supported via multi-network policy upstream project which helps enhances security for secondary networks

Network Observability Operator v1.5

- **Cluster and Zone Aware**
 - Report traffic on per-cluster basis
 - Traffic per zones
- Reporting **Round Trip Time [RTT]** per flow basis for latency analysis
- Now reporting **Differentiated Services Code Point [DSCP]** field
- API updates and UI improvements



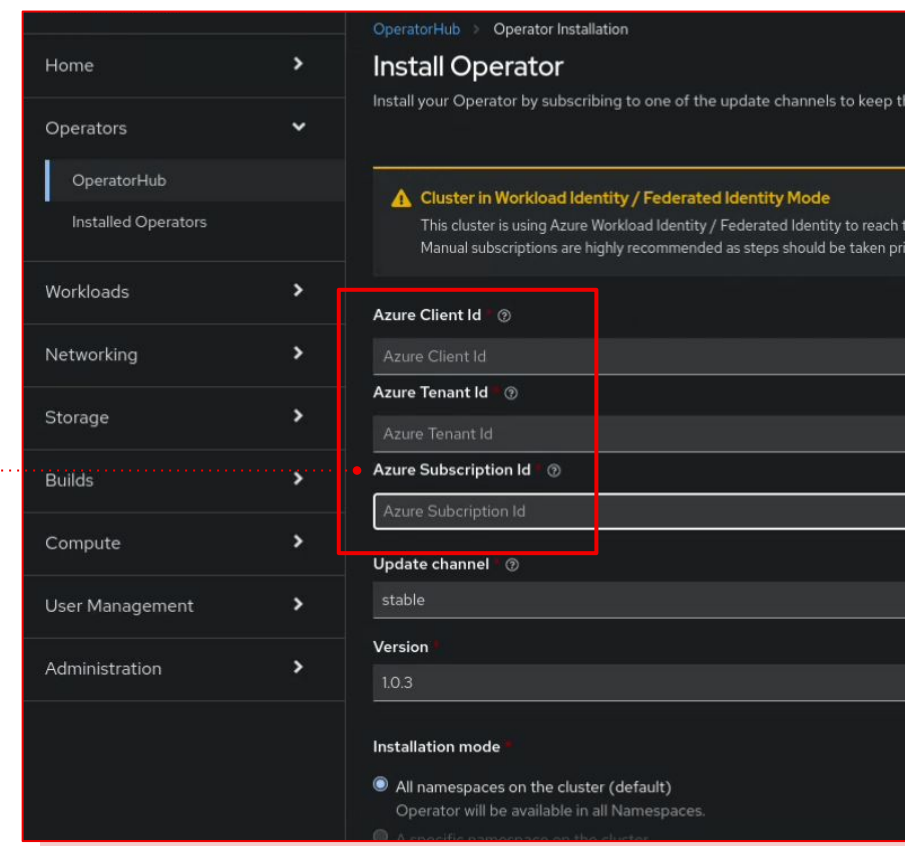
Operator Framework

OperatorHub: Install operators with tokenized cloud auth



Securely access cloud resources using a short-lived token with Azure Workload Identity

Strong security posture by managing operators talking with cloud provider API with short-lived authentication tokens

- ▶ OLM managed operators will be enabled to support this standardized tokenized cloud authentication flow:
 - **Discoverable security:** The console will show which operators support short-lived token authentication and their IAM requirements.
 - **Guided setup:** OperatorHub will guide users to fill in client ID, tenant ID, and subscription ID during operator installation.
 - **Cloud access on tap:** The CloudCredentialOperator will configure a secret which contains credentials for API access in the cloud accounts.
- ▶ AWS STS has been supported since OCP 4.14 release and we expanded to cover Azure Identity in OCP 4.15 release.



Operator Framework




Easy access to Deprecation Information

See if an installed operator is deprecated entirely, currently subscribed to a deprecated channel, or stays in a deprecated version, and know how to stay within the support boundary.



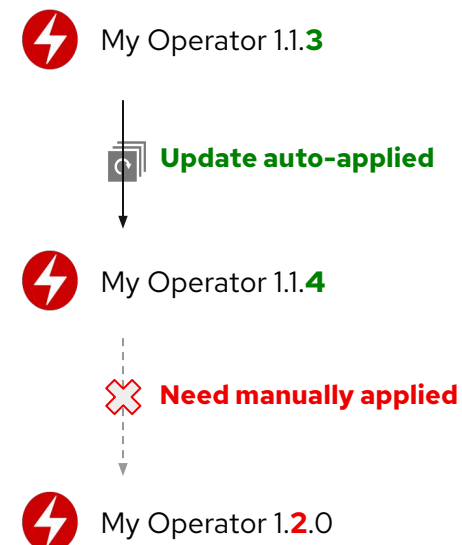
Z-stream only automatic updates (OLM 1.0 Tech Preview)

"Operator" API: all the security/CVEs patches can be auto-applied without human interactions, and no worries about introducing any breaking changes by the auto-updates.



Performance improvements in Catalog API (OLM 1.0 Tech Preview)

"Catalog" API: replaces the original CRD-based approach with a new HTTP Service to serve catalog content to reduce the burden on the Kube API server.



Storage

Journey to CSI



Operators & Drivers

- GPC Filestore
 - Now supports Shared VPC deployments
- IBM Cloud
 - Support for BYOK
- LSO
 - Wipe partition table before provisioning



CSI Migration

- All CSI migrations are enabled!
- Recommended to move CSI SC as default



Misc

- New behavior for PVCs pending due to no default SC
 - Retroactively assigns futur default SC

CSI Operators

Operator	Migration	Driver
AliCloud Disk	n/a	GA
AWS EBS	GA	GA
AWS EFS	n/a	GA
Azure Disk	GA	GA
Azure File	GA	GA
Azure Stack Hub	n/a	GA
GCE Disk	GA	GA
GCP Filestore	n/a	GA
IBM Cloud	n/a	GA
RH-OSP Cinder	GA	GA
RH-OSP Manila	n/a	GA
vSphere	GA	GA
SecretStore	n/a	TP

Non Graceful node shutdown

(GA)

- **Release CSI volume attachments** when the node's shutdown is not detected by Kubernetes.
- **Volumes can be reattached** on other nodes
- Taint the node with
 - `out-of-service=nodeshutdown:NoExecute`
- Remove the taint once the node is back online
- Can be **automated** with the Self Node Remediation Operator

```
# Ensure the node is down

# Taint the node
$ oc adm taint node <node-name> \
node.kubernetes.io/out-of-service=nodeshutdown:NoExecute

# Start the node and ensure it is online

# Untaint the node
$ oc adm taint node <node-name> \
node.kubernetes.io/out-of-service=nodeshutdown:NoExecute
-
```

Improve SELinux for RWOP PVs

(Tech Preview)

- Apply SELinux context at **mount time**
 - With a `-o context=`
- **Replaces** the default recursive chcon approach
 - Addresses pod's startup timeouts
- Applies to **RWOP PVs** only for now
 - Active RWO/RWX work upstream
- CSI Drivers must explicitly **expose** support
 - `CSIDriver.SELinuxMountSupported: true`
- Currently **enabled** by default in
 - AWS EBS
 - Azure Disk
 - GCP PD
 - IBM VPC Block
 - Openstack Cinder
 - VMware vsphere
 - ODF RBD & CephFS

```

apiVersion: v1
kind: Pod
metadata:
  name: pod-example
spec:
  containers:
  # Specs

  volumes:
  # Specs

  securityContext:
    fsGroup: 1234
    supplementalGroups: [5678]
    selinuxOptions:
      level: "s0:c12,c34"

# Volume is mounted with
# -o context=system_u:object_r:container_file_t:s0:c12,c34

```

LVM Storage

What is it?

CSI driver for node local storage backed by RHEL's logical volume manager, i.e. for each PVC a logical volume is dynamically created

New Features:

Designed for FIPS:

- When installed and running on OpenShift / RHCOS in FIPS mode, LVMS uses the RHEL cryptographic libraries that have been submitted to NIST for FIPS 140-3 Validation on the x86_64 architecture.

Support for on multi node clusters:

- Previously only Single Node OpenShift was supported
- Caveat: it's still node **local** storage. There's no distribution/replication across nodes. Workload has to ensure replication (e.g. PSQL active/passive) to avoid single point of failures.

Support for software RAID:

- Leverage LVM / mdadm software RAID capabilities to protect against single disk outages

Wipe local volumes before first use:

- Can optionally wipe disks to help with automated testing on re-used devices.

OpenShift Data Foundation 4.15 updates

- Data Resiliency
 - RDR support for existing customers
 - Non resilient storage class (replica 1)
- Performance profiles
 - Lean, balanced, performance
- Side by Side Internal and External mode
 - The ability to scale internal mode deployment with external mode
 - Multiple storage tiers

Out of the box support

Block, File, Object, NFS

Platforms

AWS/Azure

Google Cloud (GA)

OpenShift Virtualization

OSP (Tech Preview)

Bare metal/IBM Z/Power

VMWare Thin/Thick IPI/UPI

ARO - Self managed ODF

IBM ROKS & Satellite - Managed ODF (GA)

Any platform using agnostic deployment mode for self managed OpenShift deployments.

Deployment modes

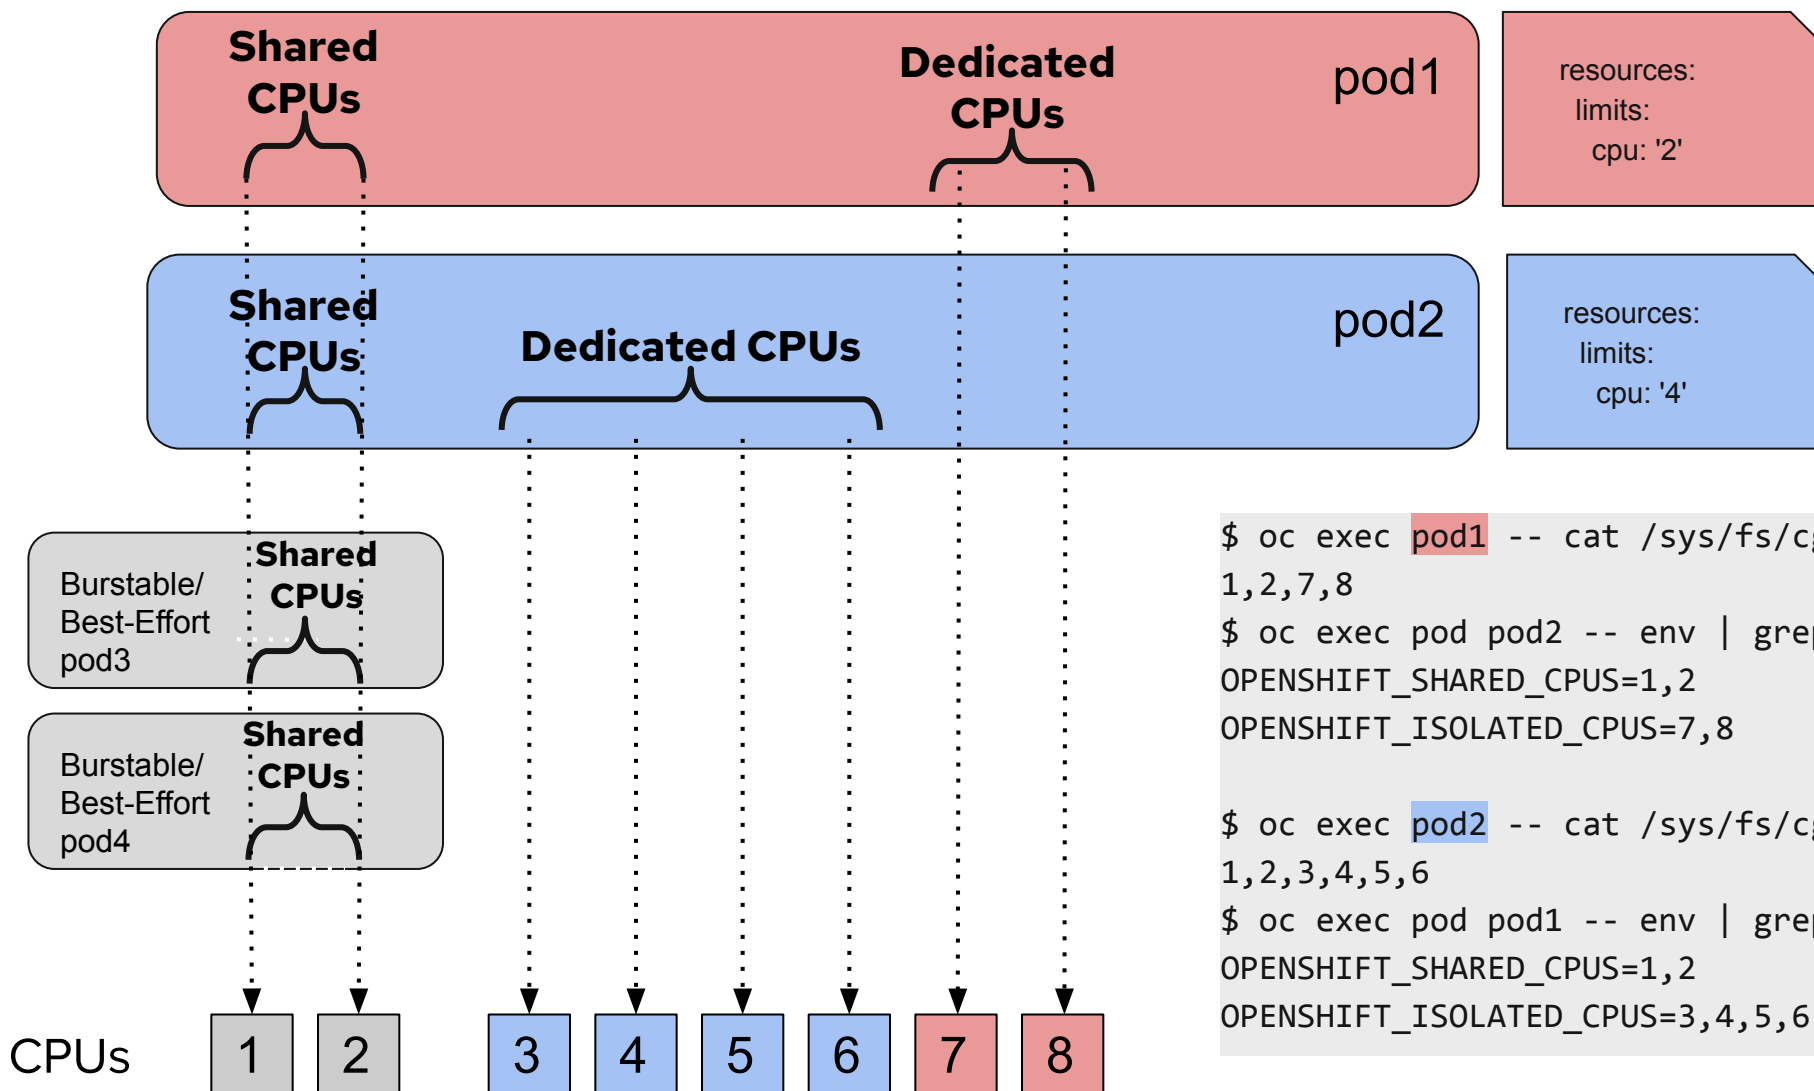
Disconnected environment and Proxied environments



Telco 5G

OPEX, CAPEX & sustainability: Efficient CPU allocation

Before: all CPUs are either dedicated, either shared - Developer Preview in 4.15 ([specifications](#))



```
resources:
limits:
cpu: '2'
```

```
resources:
limits:
cpu: '4'
```

```
$ oc exec pod1 -- cat /sys/fs/cgroup/cpuset/cpuset.cpus
1,2,7,8
$ oc exec pod pod2 -- env | grep CPUS
OPENSIFT_SHARED_CPUS=1,2
OPENSIFT_ISOLATED_CPUS=7,8

$ oc exec pod2 -- cat /sys/fs/cgroup/cpuset/cpuset.cpus
1,2,3,4,5,6
$ oc exec pod pod1 -- env | grep CPUS
OPENSIFT_SHARED_CPUS=1,2
OPENSIFT_ISOLATED_CPUS=3,4,5,6
```

Dev Preview: Accelerate RAN vDU Upgrade on Single Node OpenShift

Goals:

- Reduce upgrade time and service downtime for DU-configured OpenShift deployments

What we plan to do:

- Replace the existing upgrade procedure with **Image Based Upgrade** procedure

Steps to upgrade a DU-configured Single Node OpenShift using Image Based Upgrades (IBU)




STEP 1
seed-image generated from DU-configured Single Node OpenShift installation



STEP 2
seed-image uploaded to image registry




STEP 3
seed-image downloaded to Far Edge server and Lifecycle Agent Operator prepares node for pivot to new version



STEP 4
Backup CNF kubernetes artefacts and reboot to updated OpenShift version



STEP 5
Lifecycle Agent Operator finalizes install by applying site-specific configuration. CNF can now be re-instantiated.



IF NEEDED
Rollback to the previously working OpenShift version.

Major Benefits

- Significantly faster upgrade time
- Upgrade from n to n+2 (EUS to EUS) not n to n+1 then n+1 to n+2

Grand Master Clock (T-GM)

Precision Timing Protocol

Single card connectivity to an external GNSS.

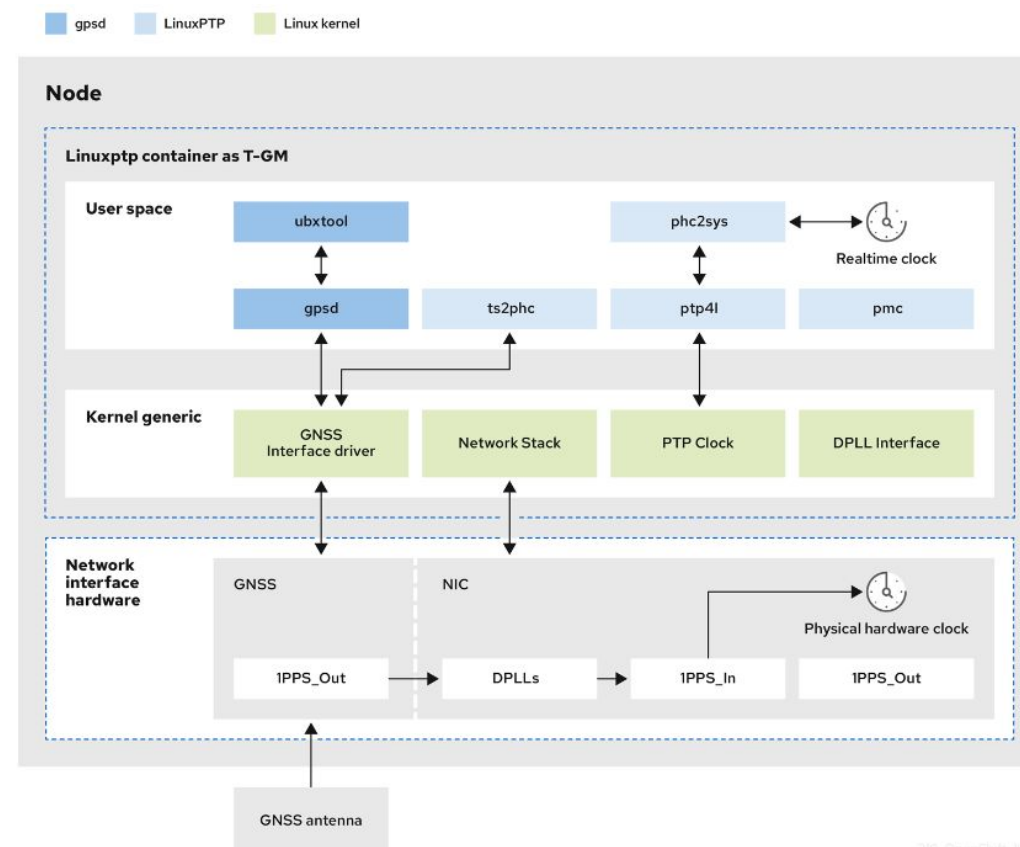
Inter-card connectivity to share timing across NICs, with physical interconnect.

PTP events and metrics added for grandmaster.

Validated with Intel E810-XXVDA4T (West Port Channel / WPC) NIC.

Backports available

- Single NIC GMC 4.14.6 (already available)
- Dual NIC GMC 4.14.14 (mid February)



399_OpenShift_1023

Thank you for joining!

Guided demos of
new features
on a real cluster

learn.openshift.com

OpenShift info,
documentation
and more

try.openshift.com

OpenShift Commons:
Where users, partners,
and contributors
come together

commons.openshift.org