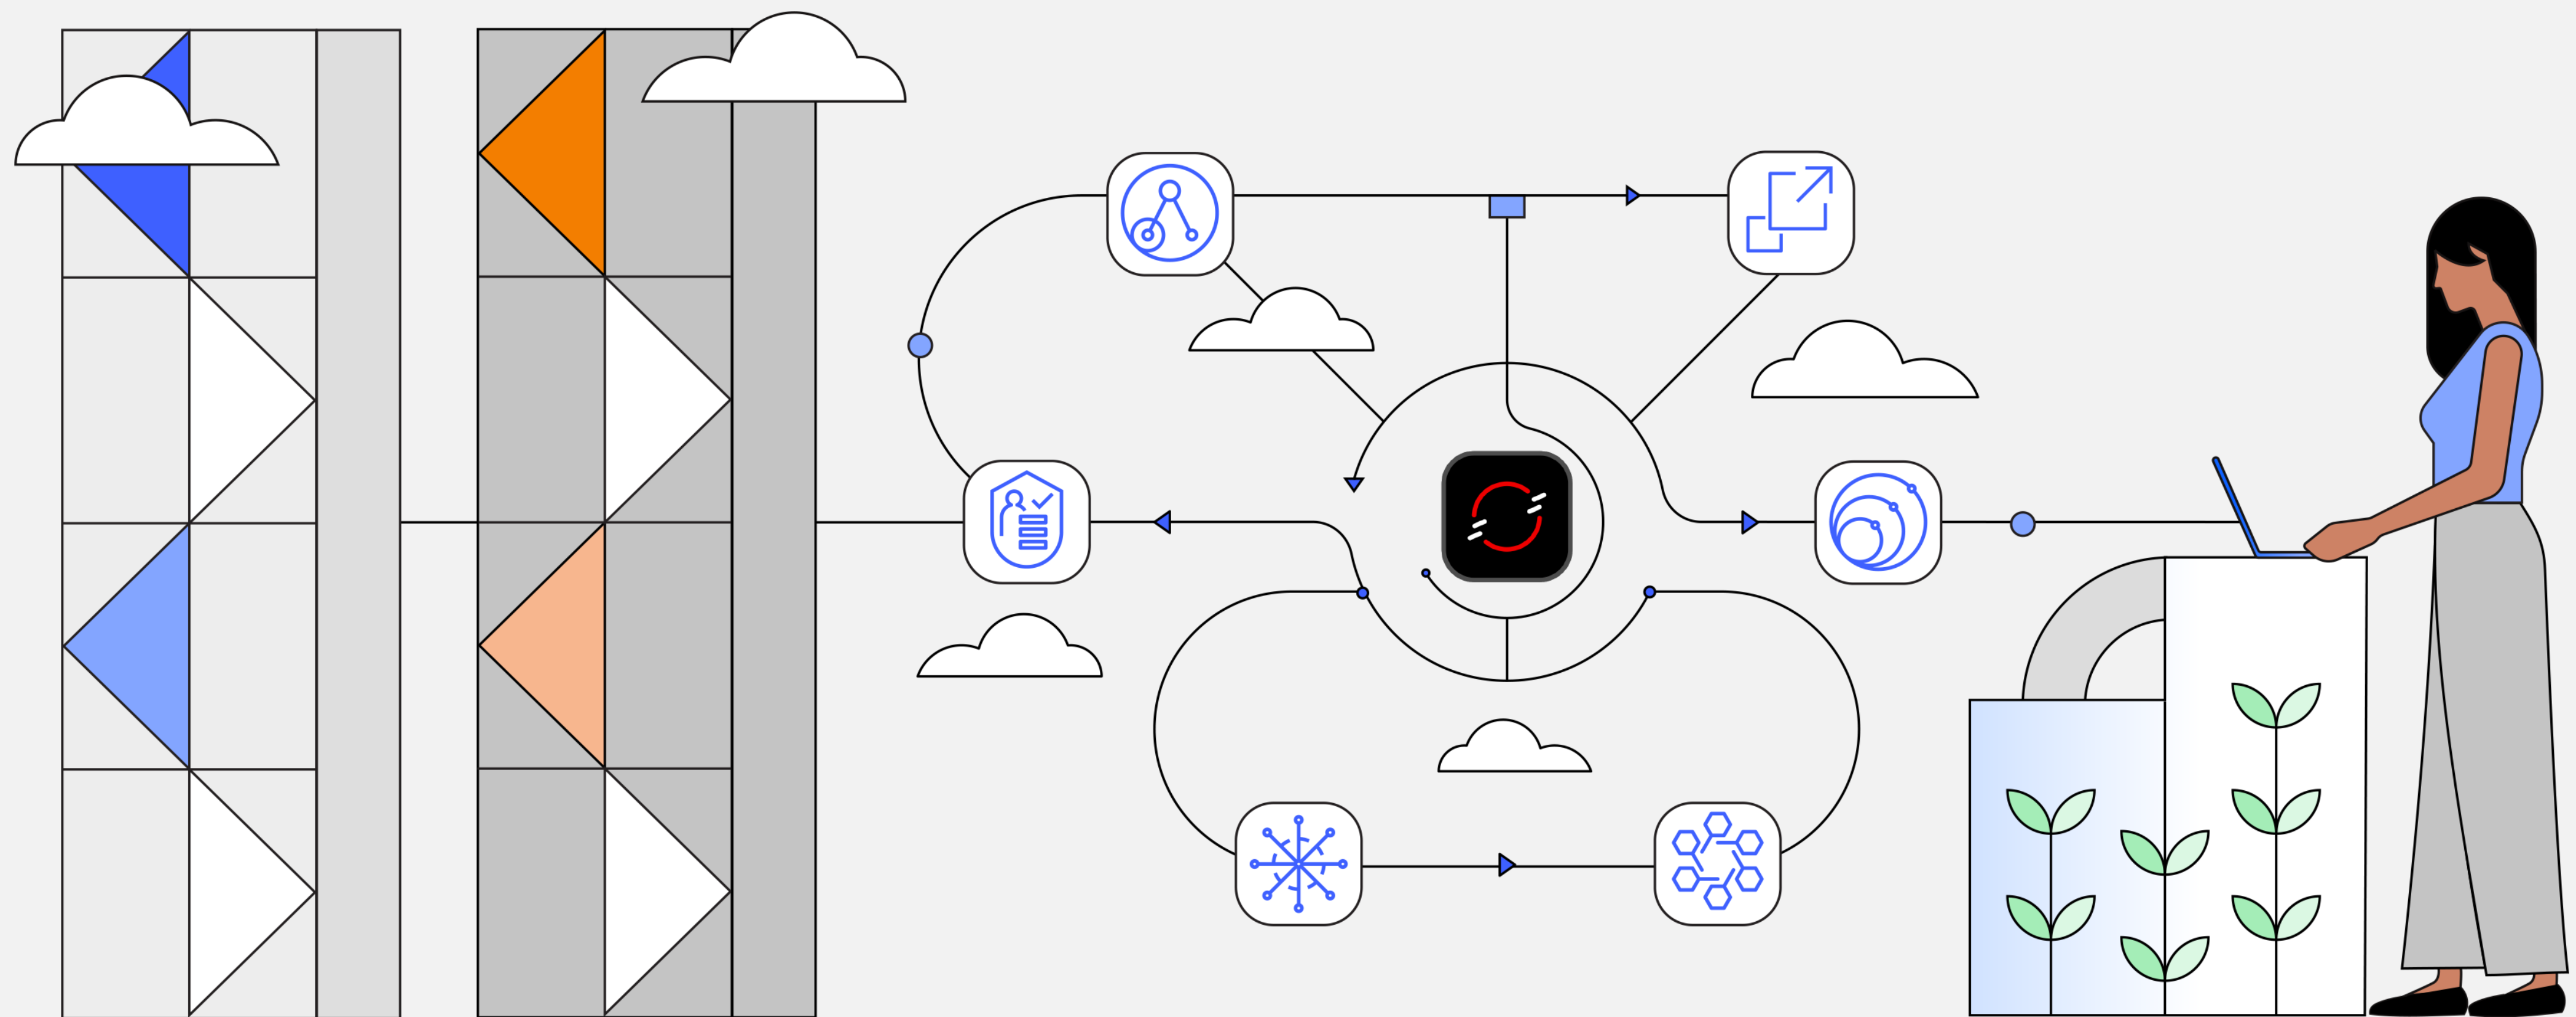


What's New in Red Hat OpenShift Container Platform on IBM Z and IBM LinuxONE



© IBM Corporation 2025

Nourhane Bziouech
Product Owner – Red Hat OpenShift
Virtualization on IBM Z and LinuxONE
Nourhane.Bziouech@de.ibm.com

Dominik Werle, IBM
Product Owner – Base Platform and Performance
of OpenShift on IBM Z and LinuxONE
Dominik.Werle@ibm.com

Holger Wolf, IBM
STSM – Red Hat OpenShift on IBM Z and
LinuxONE
Holger.Wolf@de.ibm.com

Red Hat OpenShift Container Platform and Add-ons



 **Red Hat**
Advanced Cluster Management
for Kubernetes

 **Red Hat**
Advanced Cluster Security
for Kubernetes

 **Red Hat**
Quay

Multicluster management

Observability | Discovery | Policy | Compliance |
Configuration | Workloads

Cluster security

Declarative security | Container vulnerability
management | Network segmentation |
Threat detection and response

Global registry

Image management | Security scanning |
Geo replication Mirroring | Image builds

Manage workloads

Platform services

- Service mesh | Serverless
- Builds | CI/CD pipelines
- GitOps | Distributed Tracing
- Log management
- Cost management

Build cloud native apps

Application services

- Languages and runtimes
- API management
- Integration
- Messaging
- Process automation

Data driven insights

Data services

- Databases | Cache
- Data ingest and prep
- Data analytics | AI/ML
- Data management & resilience

Developer productivity

Developer services

- Developer CLI | IDE
- Plugins and extensions
- CodeReady workspaces
- CodeReady containers

Kubernetes cluster services

Install | Over air updates | Networking | Ingress | Storage | Monitoring | Logging | Registry | Authorization | Containers | VMs | Operators | Helm

Kubernetes (orchestration)

 **Red Hat**
Enterprise Linux

Linux (container host operating system)

 **Red Hat**
Enterprise Linux
CoreOS



Physical



Virtual



Private cloud



Public cloud



Edge

Red Hat OpenShift Add-ons features, Red Hat Cloud Computing, and Red Hat Middleware/Red Hat Application Services products

Red Hat OpenShift add-ons and Red Hat
middleware help to create, deploy, and manage
apps with speed and consistency

- Red Hat OpenShift base Add-ons / Red Hat platforms
 - Red Hat Pipelines (Tekton)
 - Red Hat Serverless
 - Red Hat Service Mesh
 - Red Hat odo
 - Red Hat Dev Spaces
 - Red Hat cert-manager Operator
 - Red Hat Cost Management Metrics Operator
- Red Hat OpenShift separately priced Add-ons / Red Hat Cloud Computing
 - Red Hat Quay
 - Red Hat Advanced Cluster Security
 - Red Hat Advanced Cluster Management
- Red Hat Middleware / Red Hat Application Services
 - Red Hat AMQ Streams
 - Red Hat Fuse
 - Red Hat 3scale API Management
 - Red Hat OpenShift GitOps
 - Red Hat Runtimes

Why Red Hat OpenShift Serverless

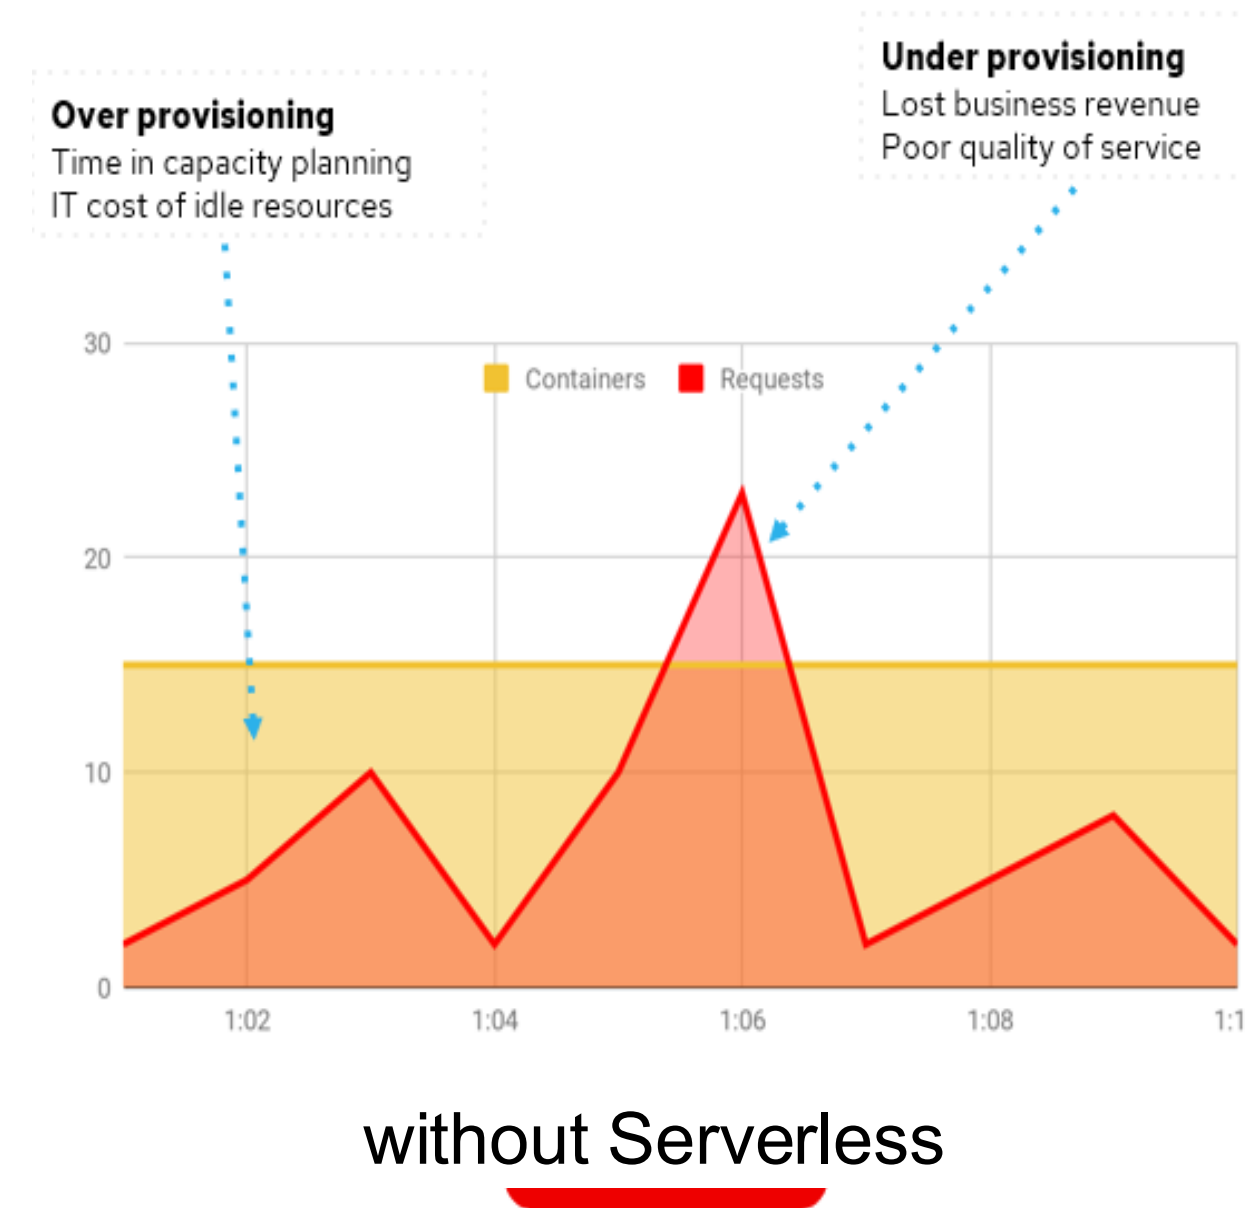
Event-Driven: Scales based on events, adjusting resources dynamically to demand.

Knative Framework: Uses Knative to manage deployment, scaling, and routing of serverless workloads on Kubernetes.

Autoscaling: Scales to zero when idle, reducing costs and automatically scaling up with demand.

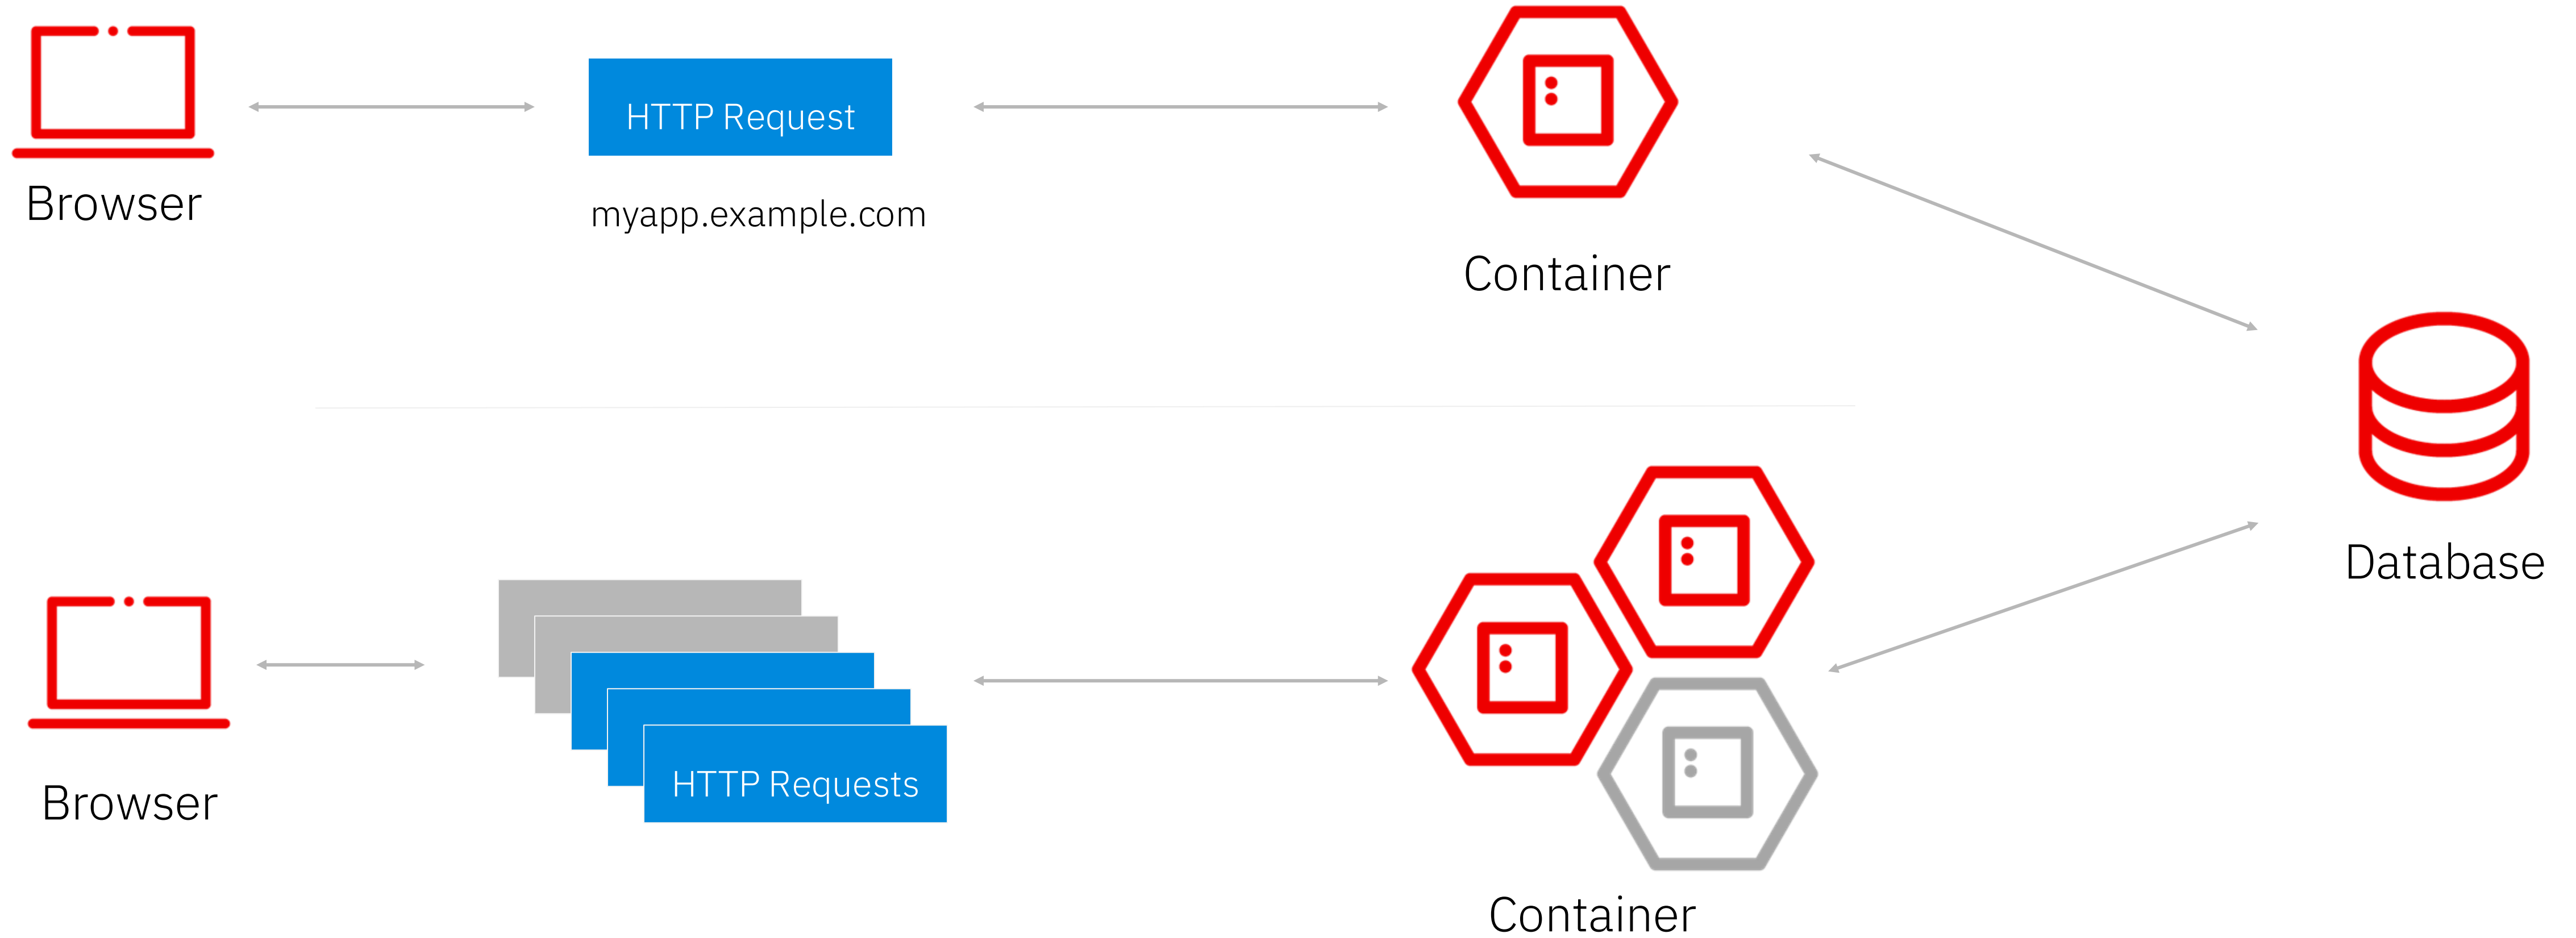
Integration: Works seamlessly with OpenShift services like monitoring, logging, and security.

Language Support: Supports various runtimes such as Quarkus, Node.js, Python, and Go, offering developer flexibility.



The "Serverless Pattern"

A serverless web application



For more information: [About OpenShift Serverless](#)

For IBM Z/IBM® LinuxONE: [Red Hat OpenShift Serverless Supported Configurations](#)

Why Red Hat OpenShift Service Mesh

Red Hat OpenShift Service Mesh provides capabilities to:

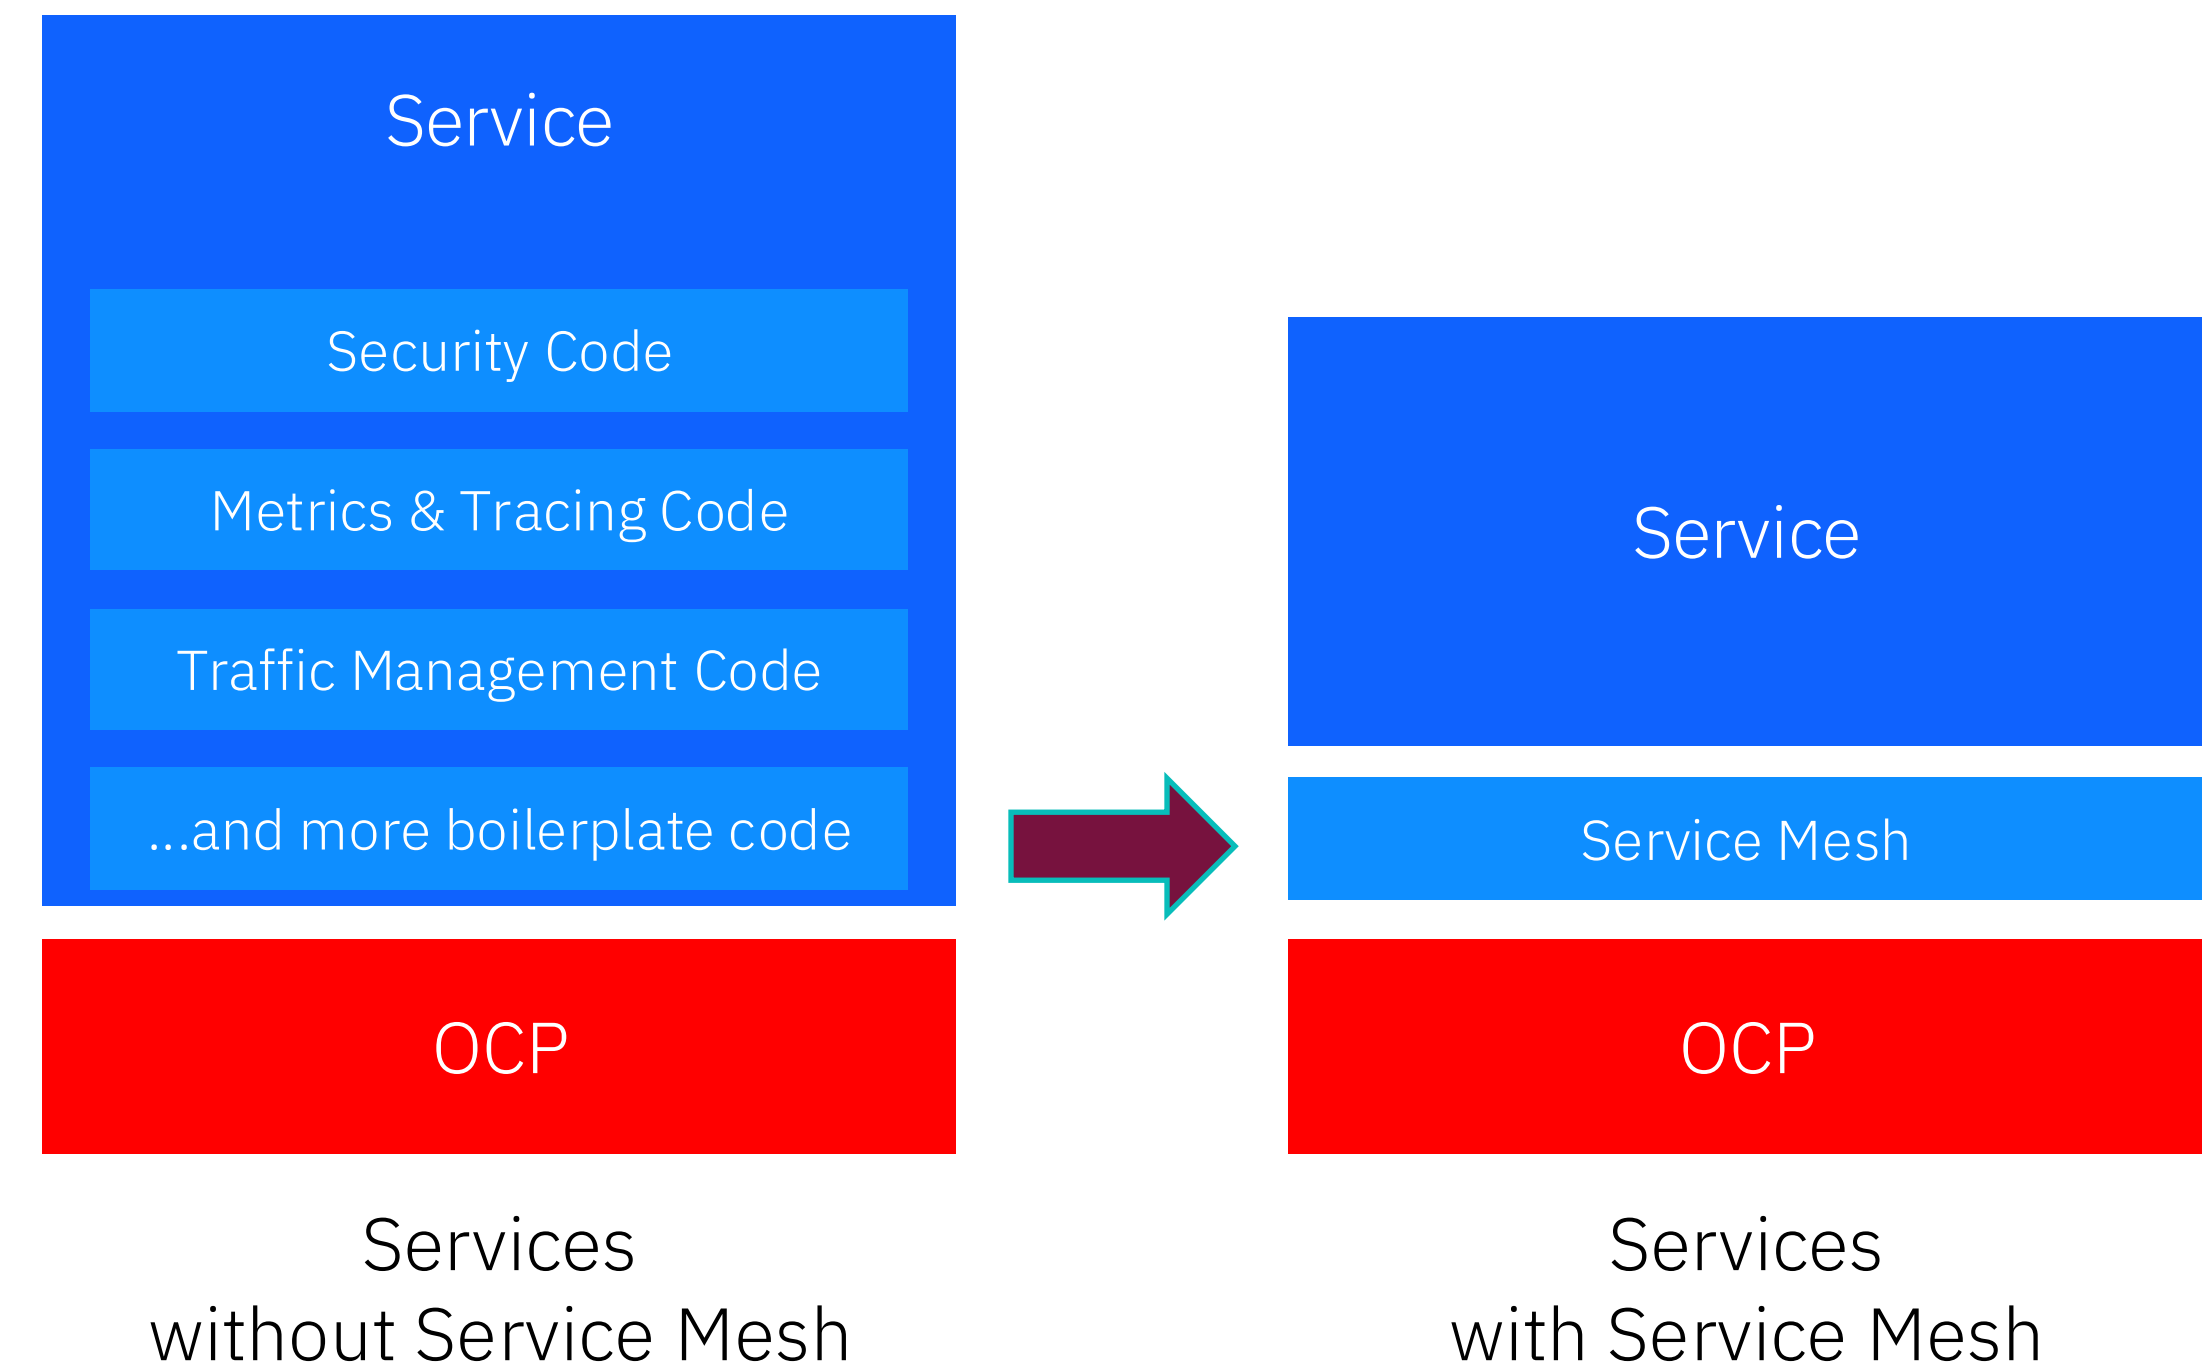
Traffic Management: Controls and monitors traffic flow between microservices, ensuring reliable communication.

Security: Provides mutual TLS encryption, authentication, and authorization for secure service-to-service communication.

Observability: Offers deep insights into service interactions through monitoring, tracing, and logging.

Resilience: Enhances application resilience with features like retries, timeouts, and circuit breaking.

Istio Integration: OpenShift integrates with Istio, a popular service mesh, for advanced traffic management and policy enforcement.



Red Hat OpenShift Service Mesh Architecture

A software infrastructure layer between Red Hat OpenShift and your services for managing communications between the services

Consists of

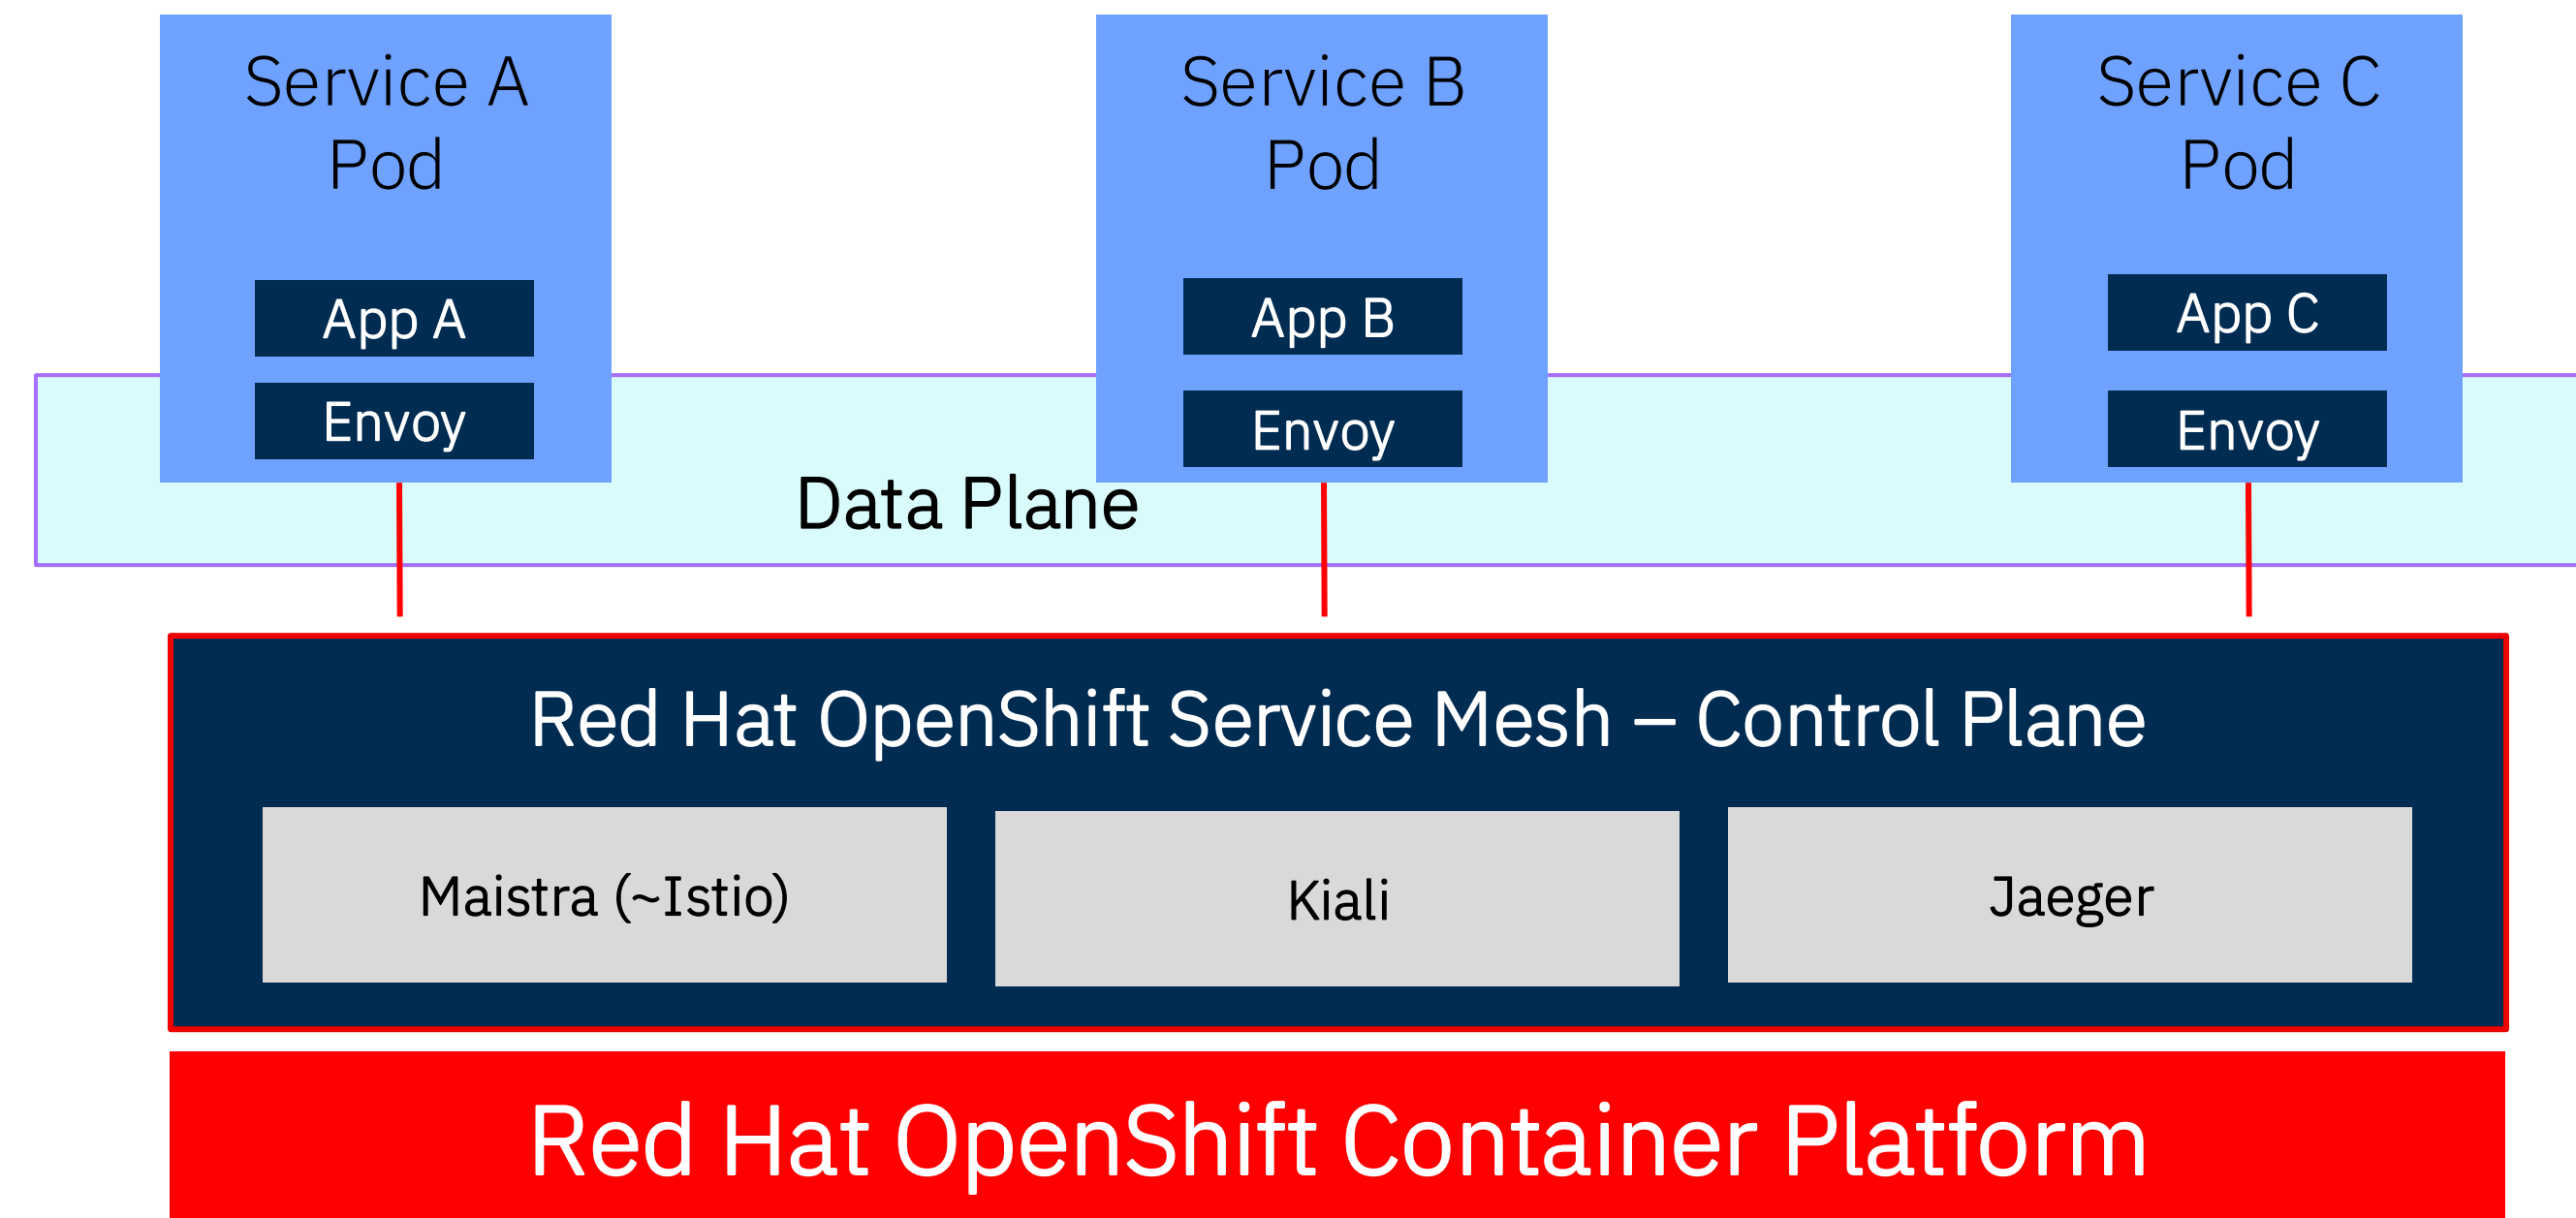
- Data Plane
- Control Plane

For more information:

- [About OpenShift Service Mesh](#)

For IBM Z and IBM® LinuxONE:

- [OpenShift Service Mesh on IBM Z/IBM® LinuxONE Part I: Installation](#)



OpenShift Pipelines: Cloud-native CI/CD built leveraging Tekton on OCP on LinuxONE



Kubernetes-native Pipelines

- Use Kubernetes resource definitions For Build/Test/Security Checks / Release



Severless - Scale on demand

- Ability to run tasks in parallel
- No central CI/CD engine



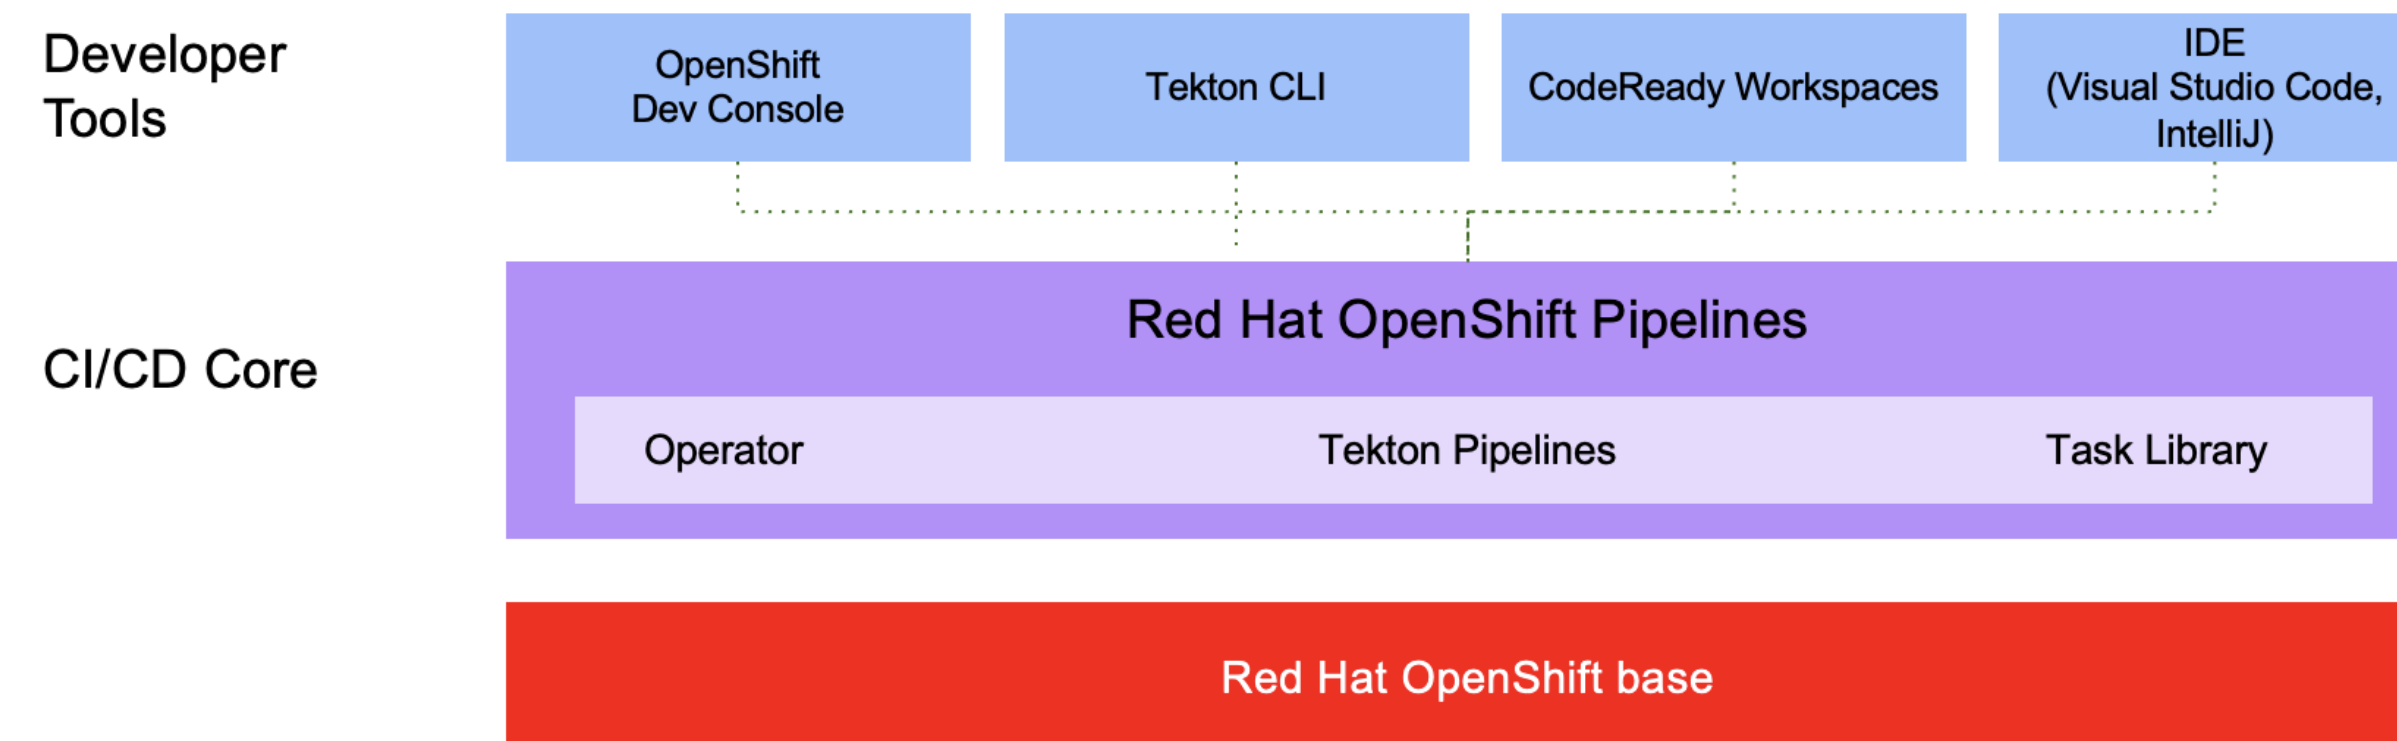
Secured pipeline execution

- Using Kubernetes RBAC and security models



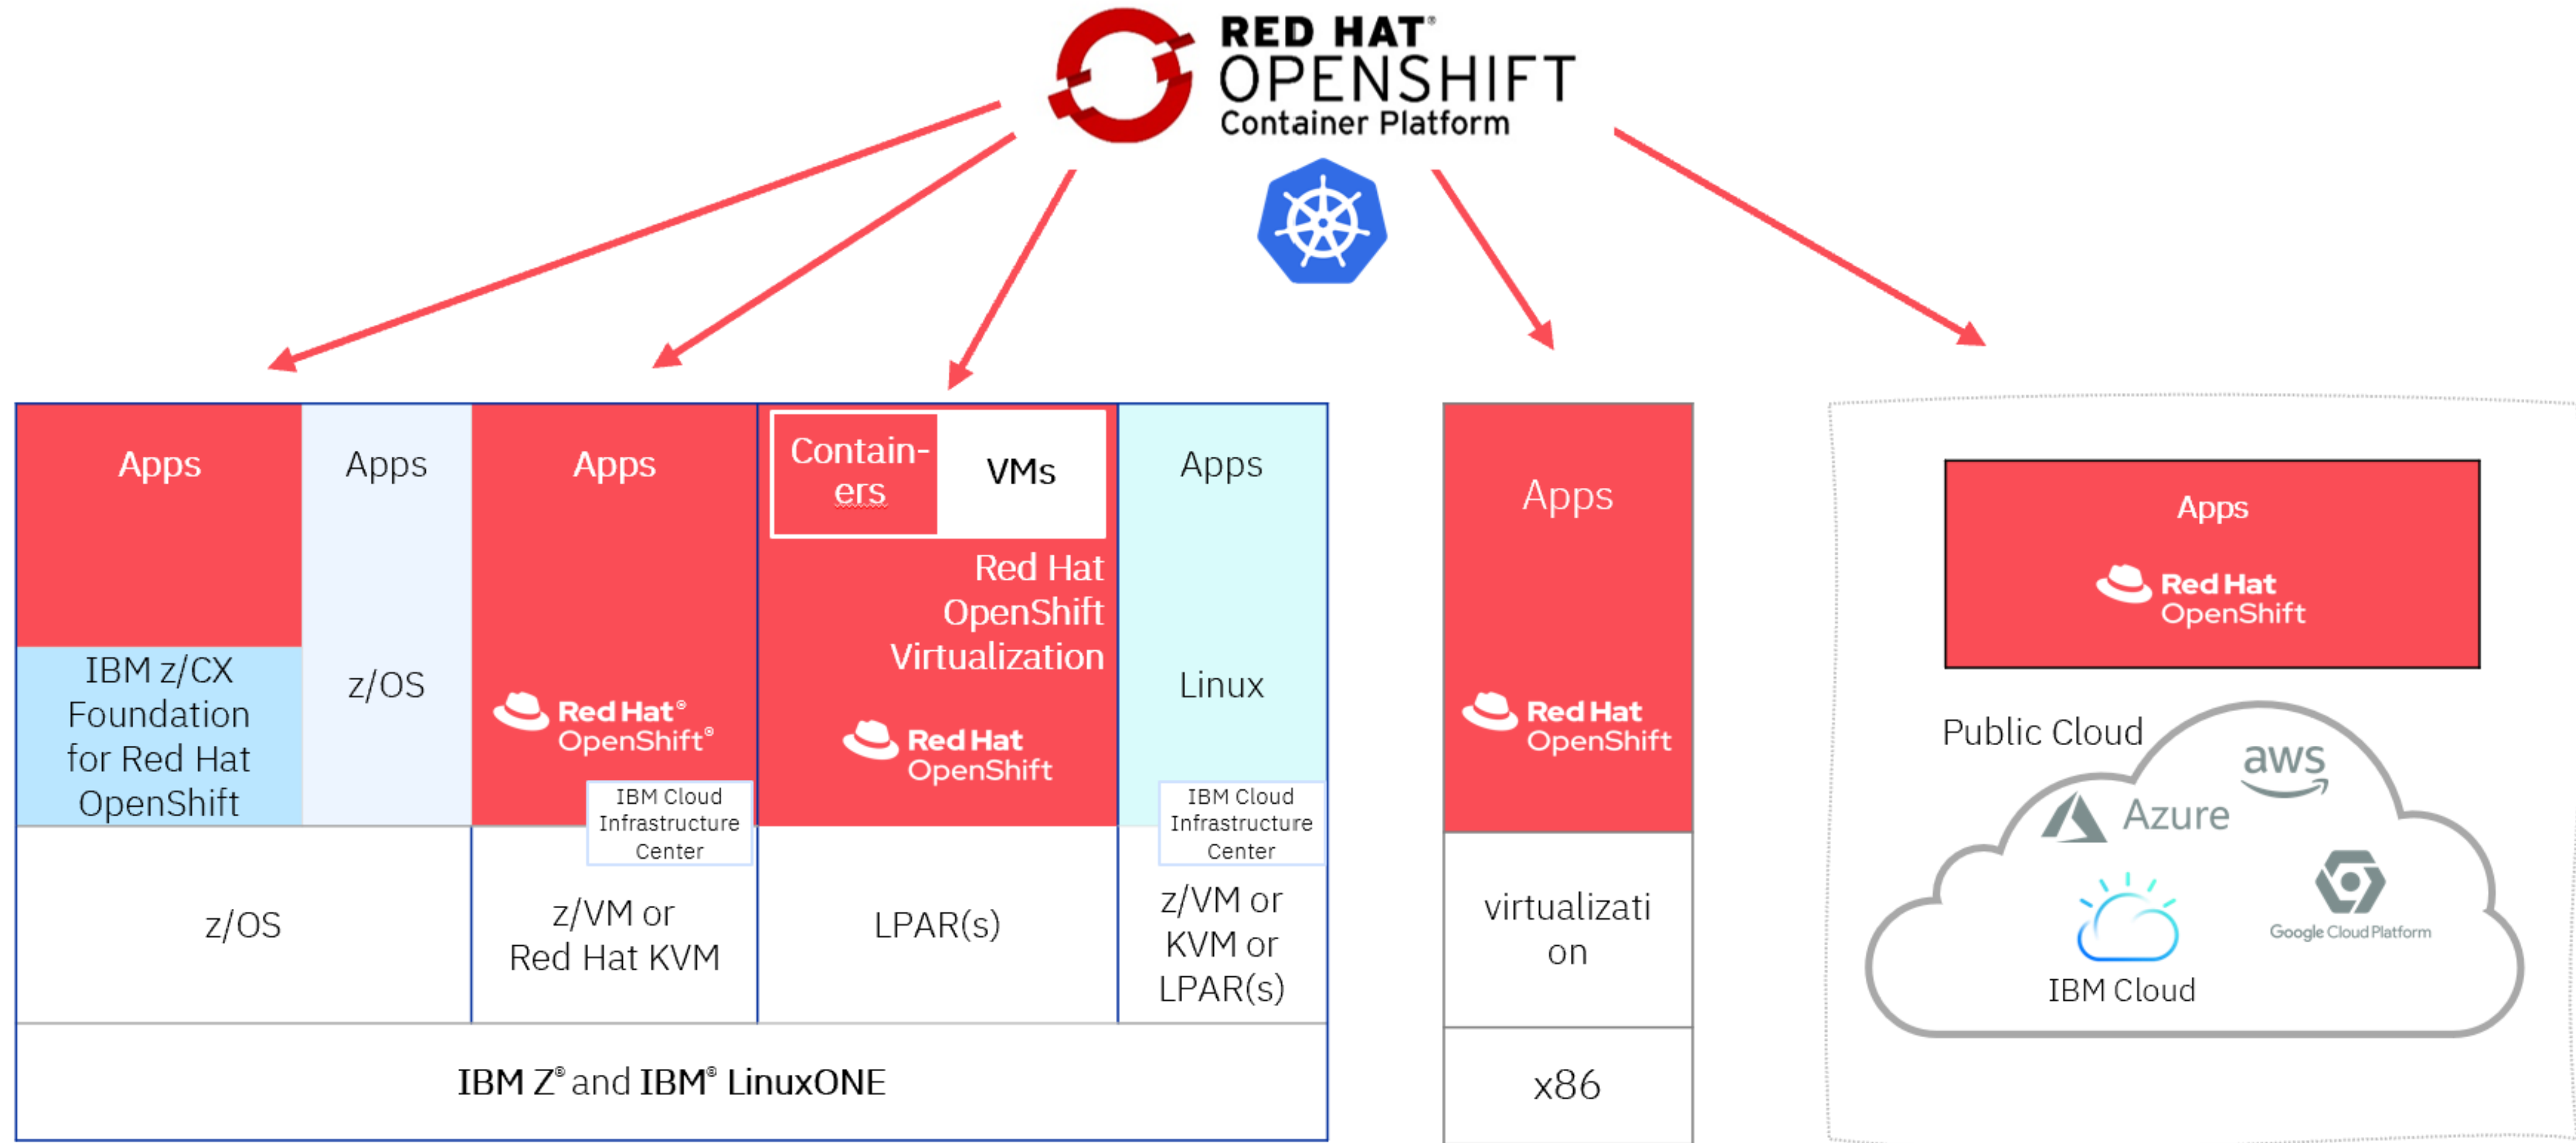
Integrated CI/CD experience

- A streamlined user experience through the OpenShift Console, CLIs, and IDEs.



For more information: <https://tekton.dev>
<https://cloud.redhat.com/blog/introducing-openshift-pipelines>

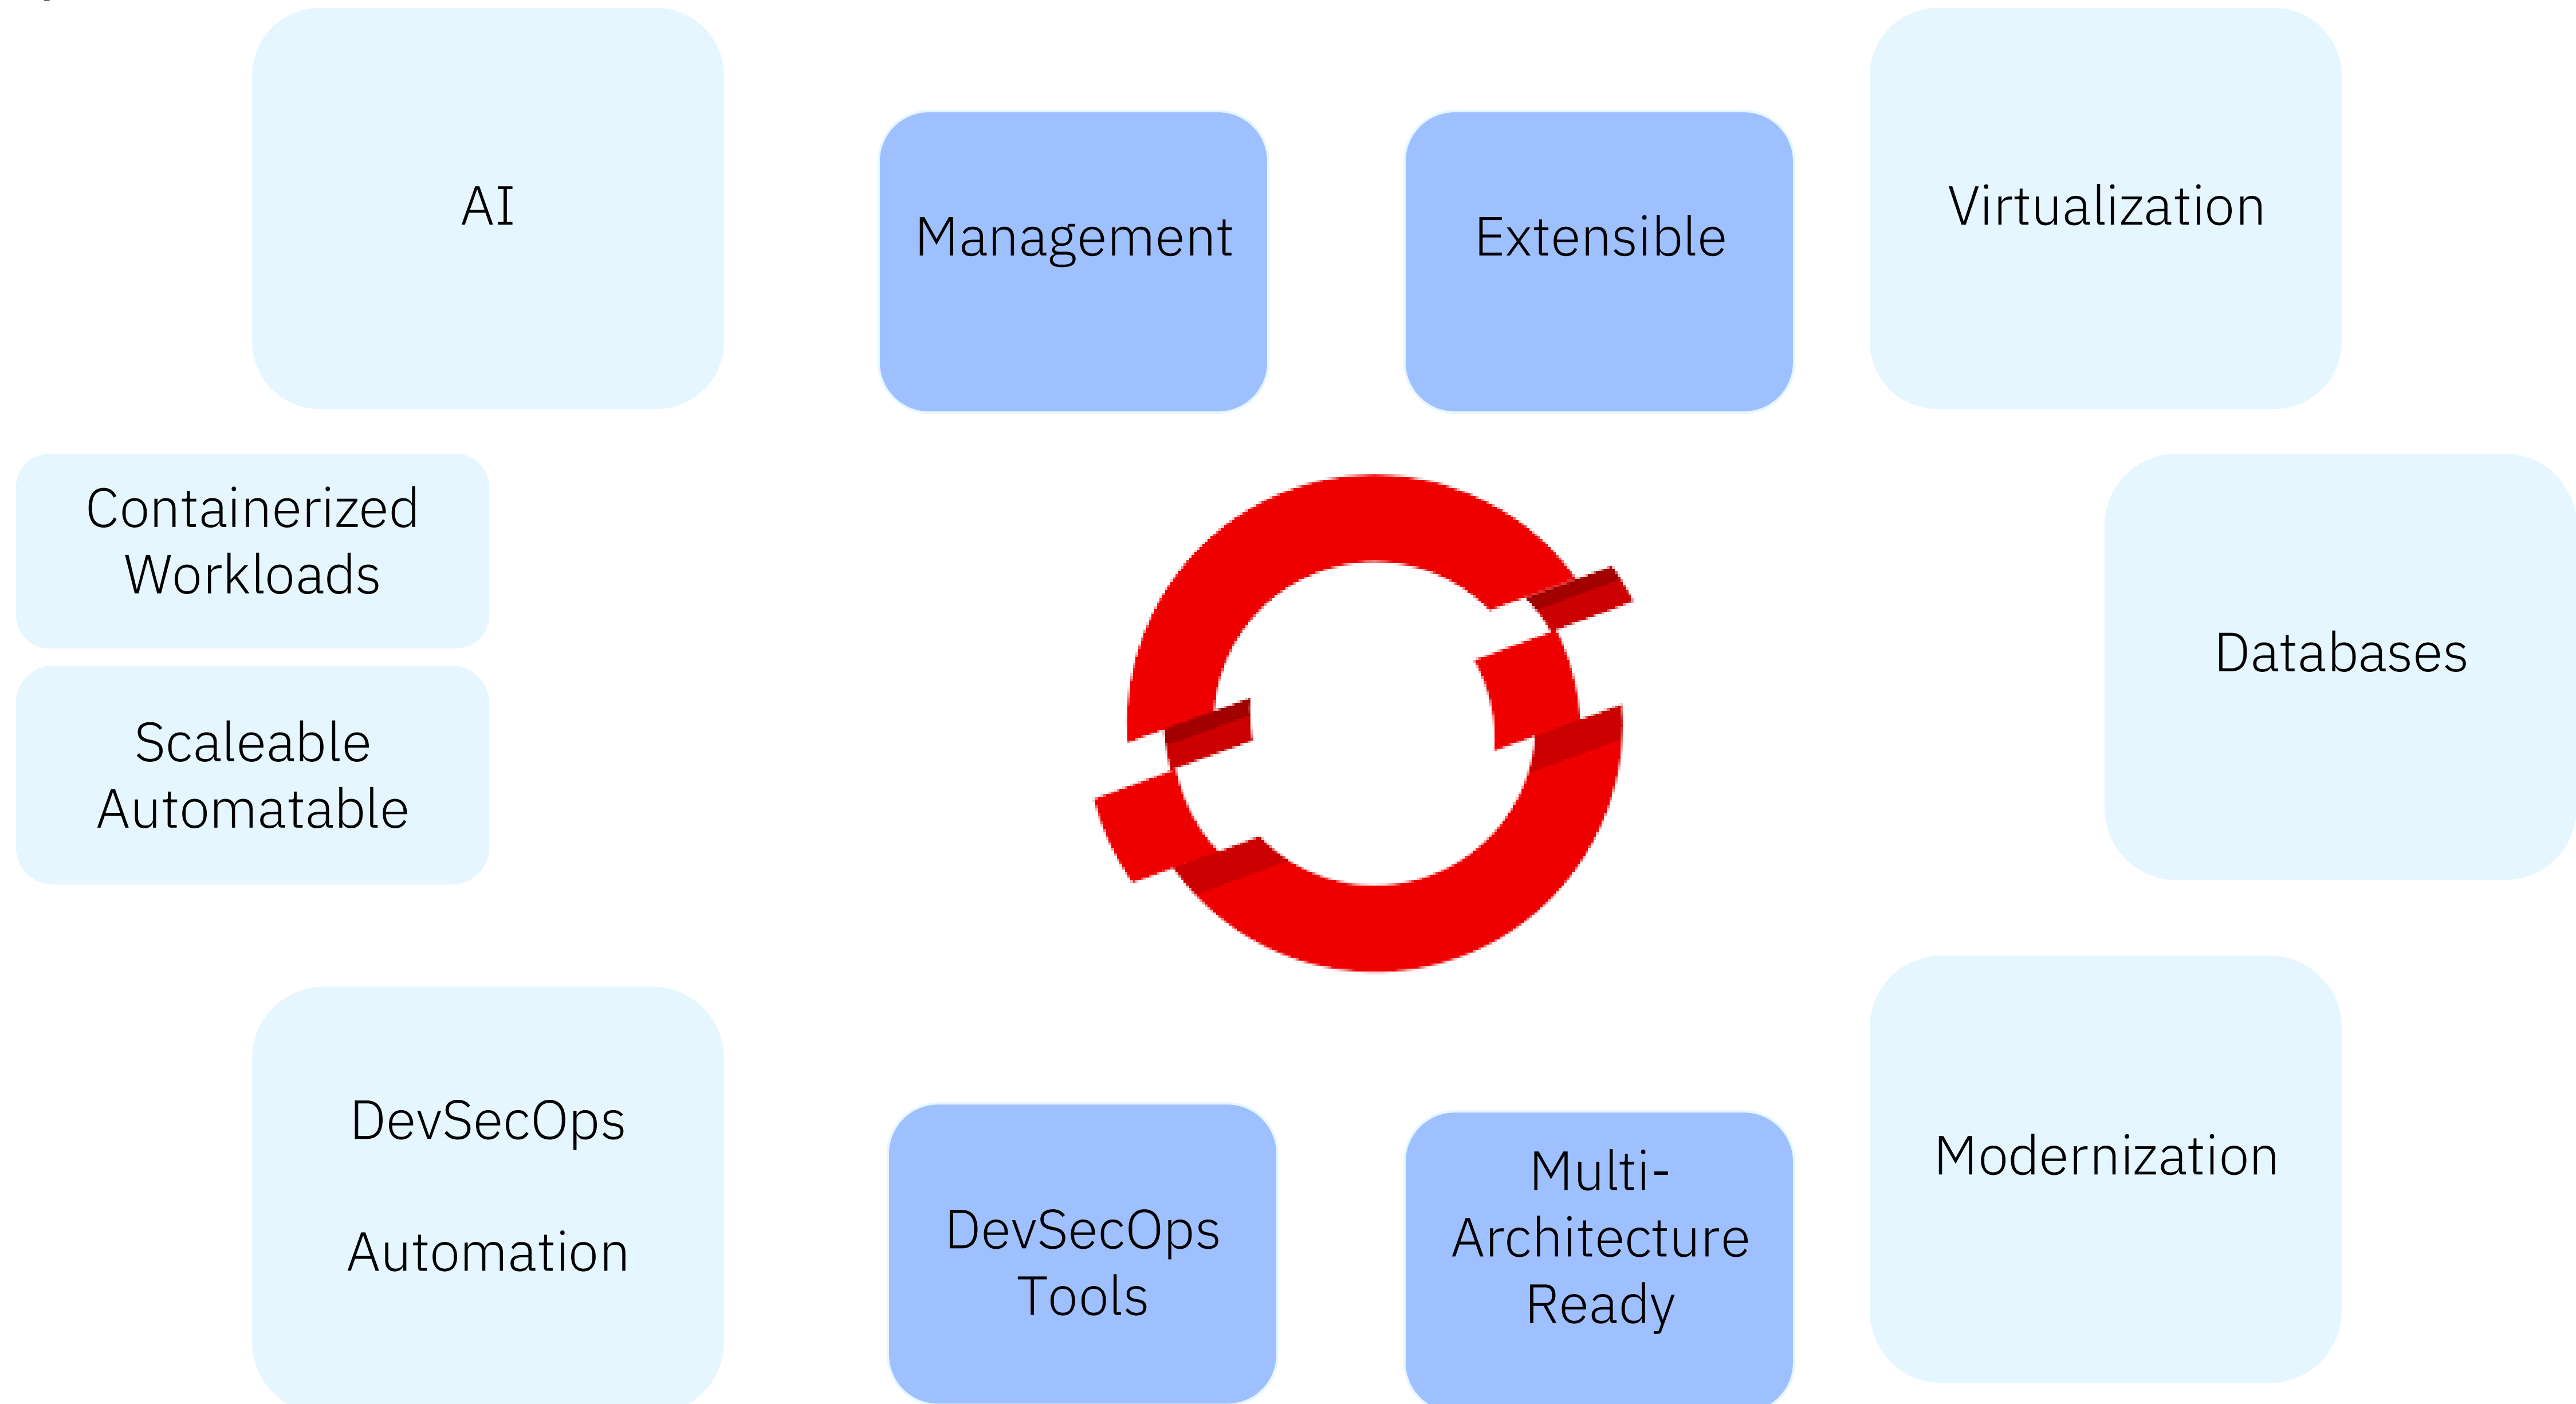
The vision of hybrid and multi-cloud with Red Hat OpenShift



Adapt enterprise DevOps with consistent, cloud native development, and CI/CD tooling.

Hybrid workloads and multiple Red Hat OpenShift clusters can run in parallel on a physical IBM Z or IBM LinuxONE server.

OpenShift the Infrastructure Foundation



OpenShift Exploits Technology for Growth

AI

Management


Extensible

Virtualization

OpenShift
Virtualization

Containerized
Workloads

Scaleable
Automatable



Databases

Modernization

Exploit
Installation
and
automation

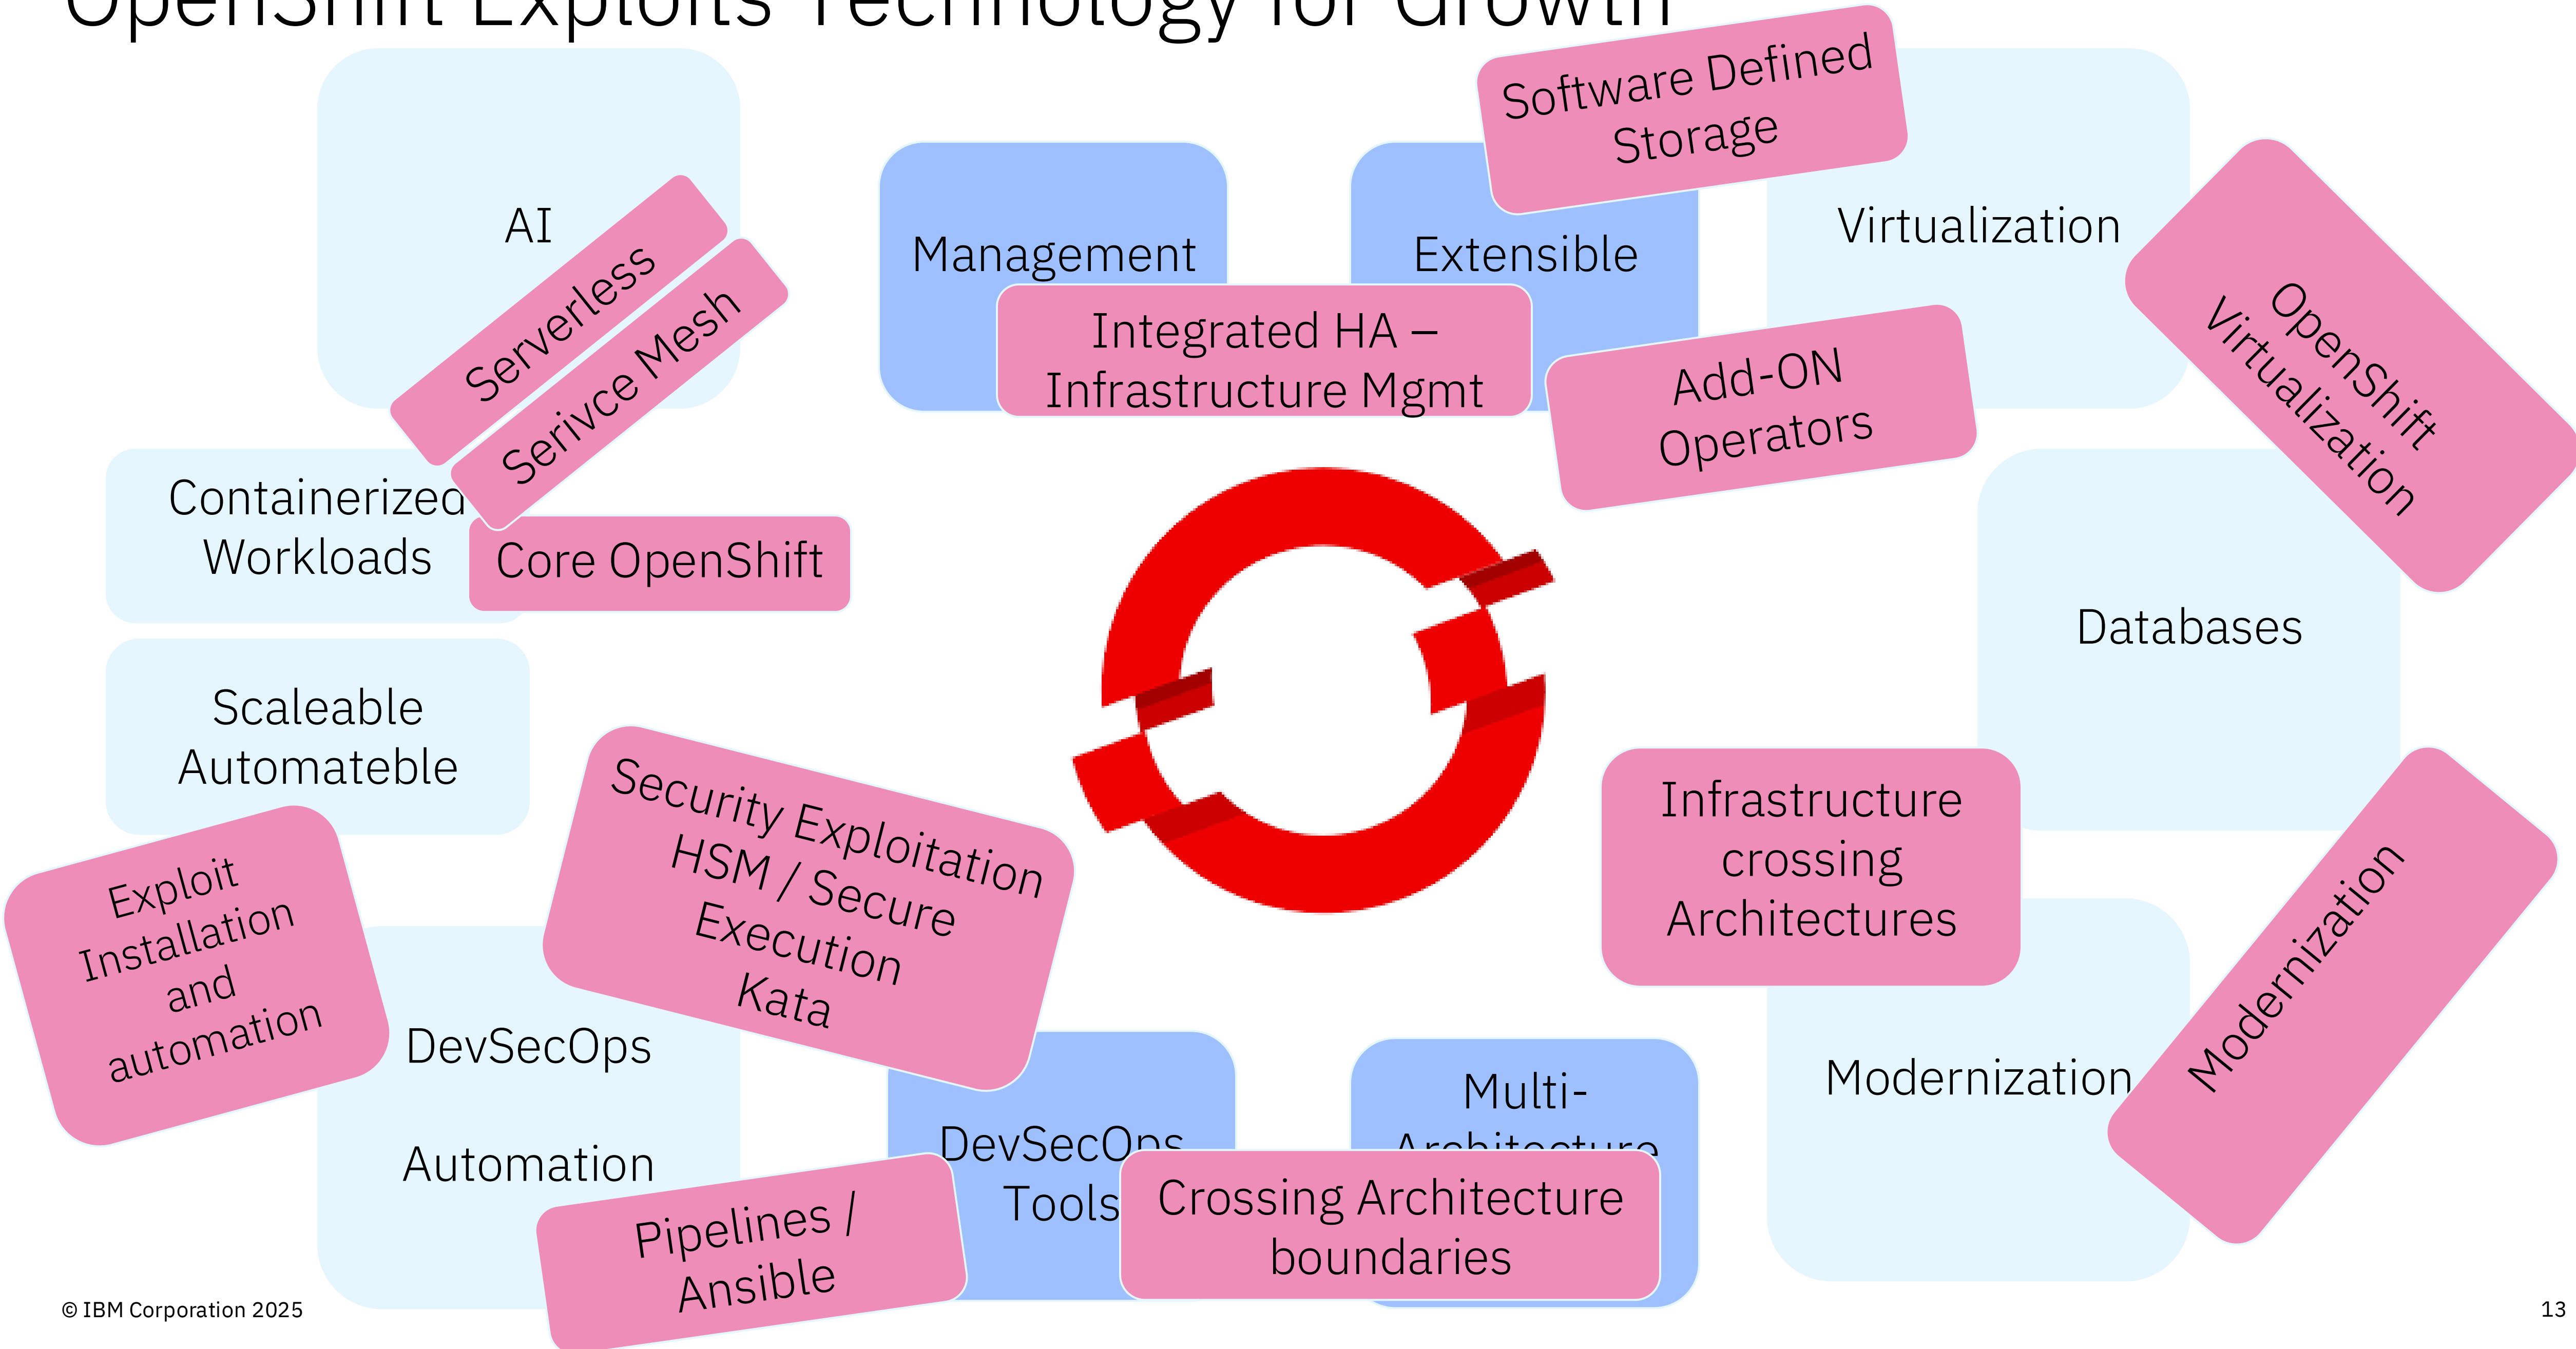
DevSecOps
Automation

DevSecOps
Tools

Multi-
Architecture
Ready

Modernization

OpenShift Exploits Technology for Growth



Overview items delivered with RHOCP 4.17 / 4.18

Management

Hosted Control Planes (GA)

- Enable infrastructure to spin-up clusters by hosting the control plane on a hosted cluster
- Enable LPAR installation
- Support: LPAR, z/VM, and KVM

CPU Manager

- Allows to provide CPU topology to be used for scheduling considerations

Secondary Scheduler

- Support customer owned scheduler implementation

OpenShift Virtualization (Tech Preview)

Platform

Hosted Control Plane (GA)

- Allows to spin up IBM Z / IBM® LinuxONE clusters by hosting control components on x86.

NVMe

- Support for NVMe local storage for boot nodes or use as local disk
- Support: LPAR (boot and data), z/VM (data)

Install

Hosted Control Planes (GA)

- Based on assisted service Support of LPAR added

Assisted Installer

- Adding Day 2 node install and NVMe
- Support: LPAR, z/VM, and KVM

Agent-based Installer Day 2 Installation

[Release Notes Red Hat OpenShift 4.17](#)

[Release Notes Red Hat OpenShift 4.18](#)

New capabilities of Red Hat OpenShift

Red Hat OpenShift 4.16

- Ingress Node Firewall Operator
- Assisted Installer LPAR support
- Agent-based Installer ISO Support
- Secure Boot for zVM and LPAR
- Compliance Operator PCI profile

Red Hat OpenShift 4.15

- Multi-Arch compute
 - IBM Z / IBM® LinuxONE (s390x) control/compute and x86 compute
- LPAR support
- Agent-based Installer
- Assisted Installer for KVM and z/VM
 - New: Discovery feature for kernel arguments
- cert-manager Operator (Add-on)
- Cost Management Metrics Operator (Add-ON)

Red Hat OpenShift 4.14

- single-node OpenShift
- Multi-Arch compute
 - x86 control/compute and IBM Z / IBM® LinuxONE (s390x) compute
- Hosted Control Planes (Tech Preview)
 - x86 hosting cluster and IBM Z / IBM® LinuxONE (s390x) compute
- NIST Profile supported in Compliance Operator
- oc-mirror plugin
- Assisted Installer for KVM and z/VM

Red Hat OpenShift 4.13

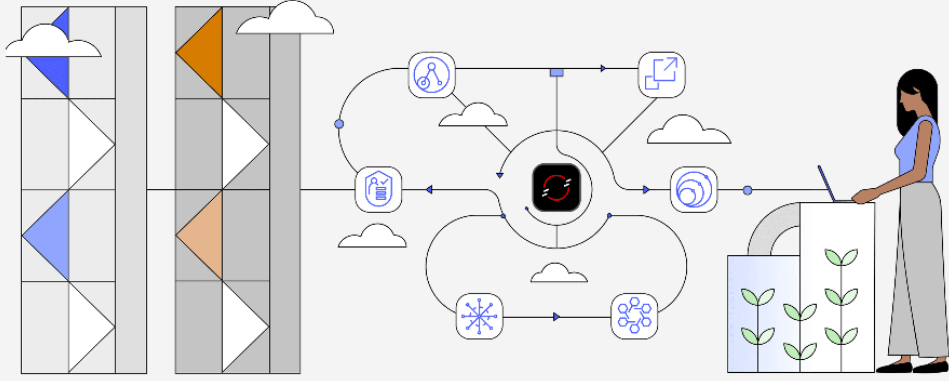
- FIPS as an installation option
- Deploy Red Hat OpenShift with Secure Execution
- NBDE network-bound disk encryption
- Expose services with MetalLB
- Control outgoing traffic with Egress IP
- Cluster Resource Override Operator
- Extension of Logging Operator
- Support of UBI9 and RHCOS based on RHEL9

IBM's webpage for Red Hat OpenShift on IBM Z and LinuxONE

- “One-stop shop” for everything related to Red Hat OpenShift on IBM Z and IBM LinuxONE
 - ibm.com/docs/en/rhocp-ibm-z
- Content
 - Reference architectures and whitepapers
 - Value of RHOCP on IBM Z
 - Performance insights
 - Links to official Red Hat documentation
 - Links to relevant IBM resources
 - List of blogs written by IBM's RHOCP team
 - Kubernetes device plug-in for CEX cards documentation

[All products](#) / [Red Hat OpenShift on IBM Z and IBM LinuxONE](#) /

Red Hat OpenShift on IBM Z and IBM LinuxONE



Recommended topics

Value of Red Hat OpenShift on IBM Z and IBM® LinuxONE →	Red Hat OpenShift, add-ons, layered products support →	Red Hat OpenShift Container Platform on IBM Z and IBM® LinuxONE Reference Architecture →
IBM Storage Fusion and IBM Storage Fusion Data Foundation Reference Architecture →	IBM resources →	Blogs →

Your choice

–

Deployment options

New deployment options

Single-node OpenShift (SNO)

- Reduces the prerequisites for a Red Hat OpenShift deployment down to a minimum.
- Requires 2 IFLs with SMT-2 enabled on installation time
- Can be used for production workloads that are not mission-critical

LPAR deployment

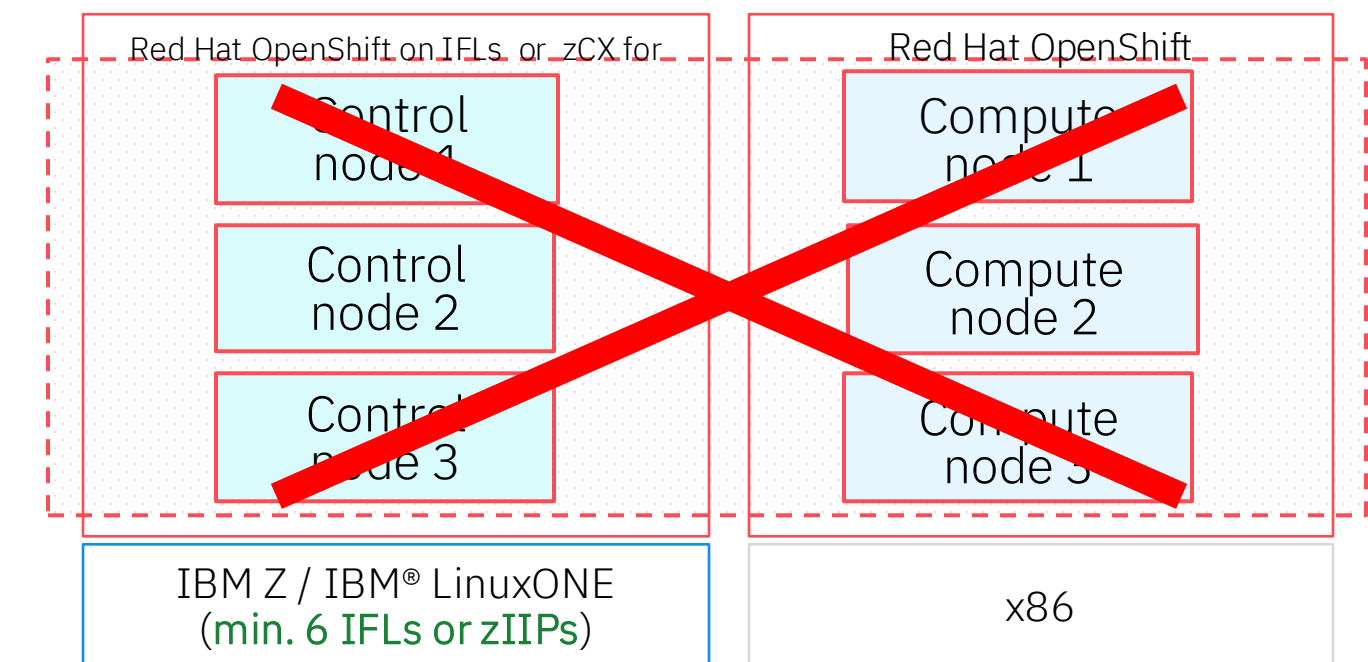
- According to the desired setup, one to several LPARs are required.
- No hypervisor required; therefore an improved performance can be expected for most workloads.
- LPAR installation does not provide the flexibility of a hypervisor-based installation regarding the placement of guests, memory, etc., and the tools to manage a dynamic guest creation.

Multi-architecture compute

- Provides the full function of Red Hat OpenShift cluster on IBM Z / IBM® LinuxONE and integrates x86 with just the small compute node footprint – and vice versa.
- Enables to build optimized composite solutions

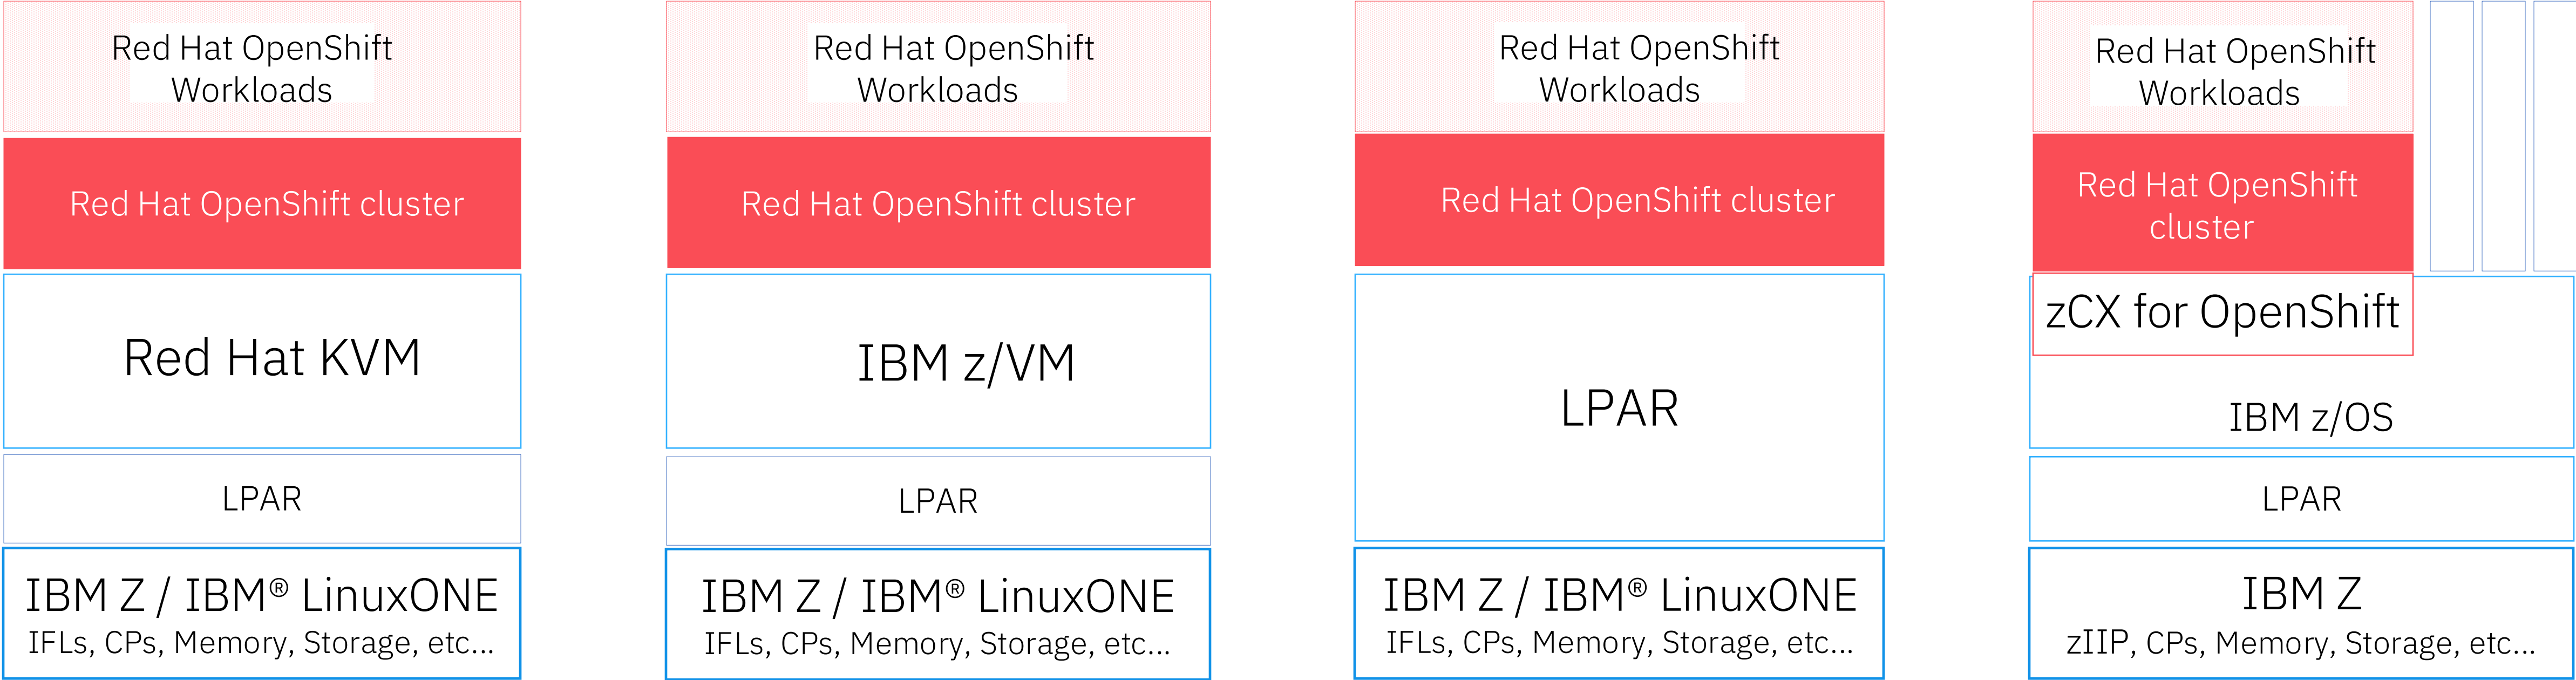
Be aware that:

- It is not supported to only install the control nodes on one architecture and the compute nodes on another architecture



Red Hat OpenShift deployment options

Red Hat OpenShift on IBM Z and IBM® LinuxONE



Control Plane Topology Options

SNO


Built from single node that consolidates control and compute nodes

Pro:

- Reduced footprint, can start with 2 IFLs/zIIPs
- Compute can be added
- Kubernetes function available
- Good for Dev/Test

Con:

- No infrastructure High Availability
- Updates result in service interruptions



3 Node


Initially 3 converged control and compute nodes

Pro:

- Reduced IFL footprint
- Can grow by adding compute nodes
- Cluster is high availability

Con:

- Cluster size might be limited for applications and grow cluster might be required



Production (HA)

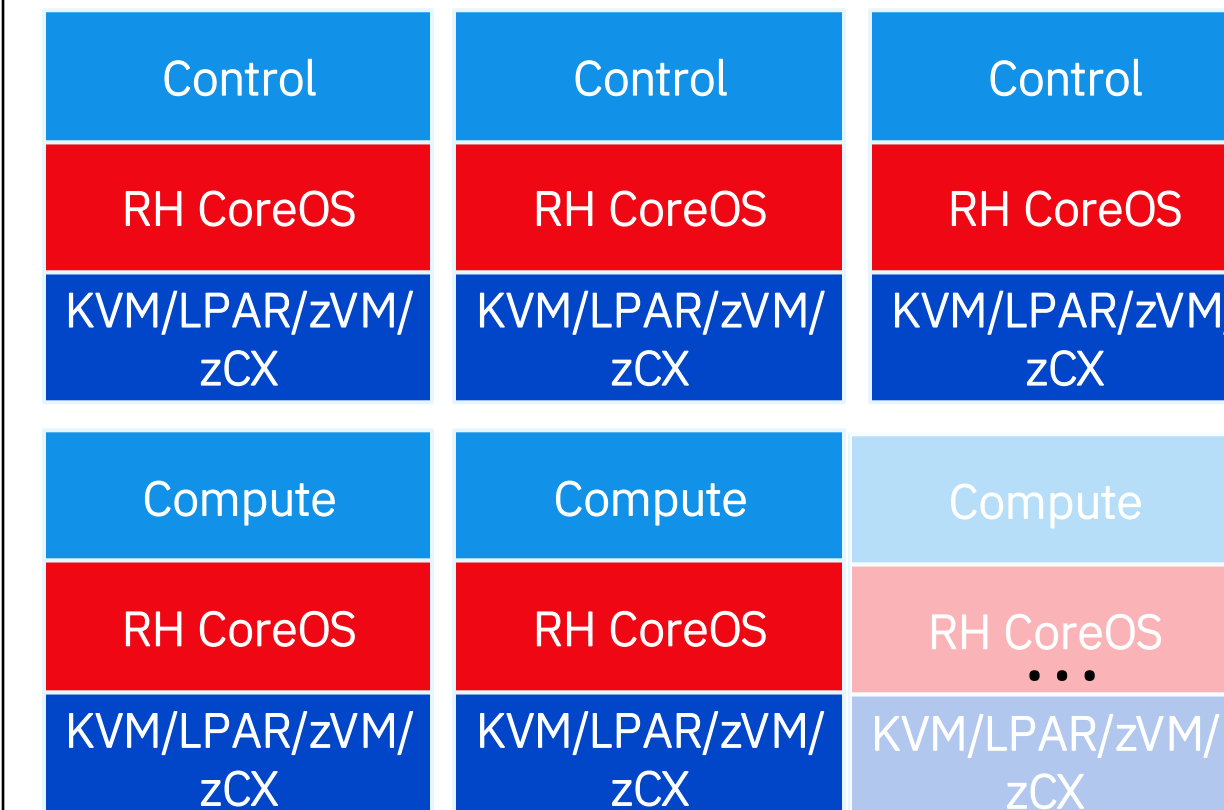
Minimum 3 control nodes and 2 compute nodes. Dedicate control nodes, for cluster growth

Pro:

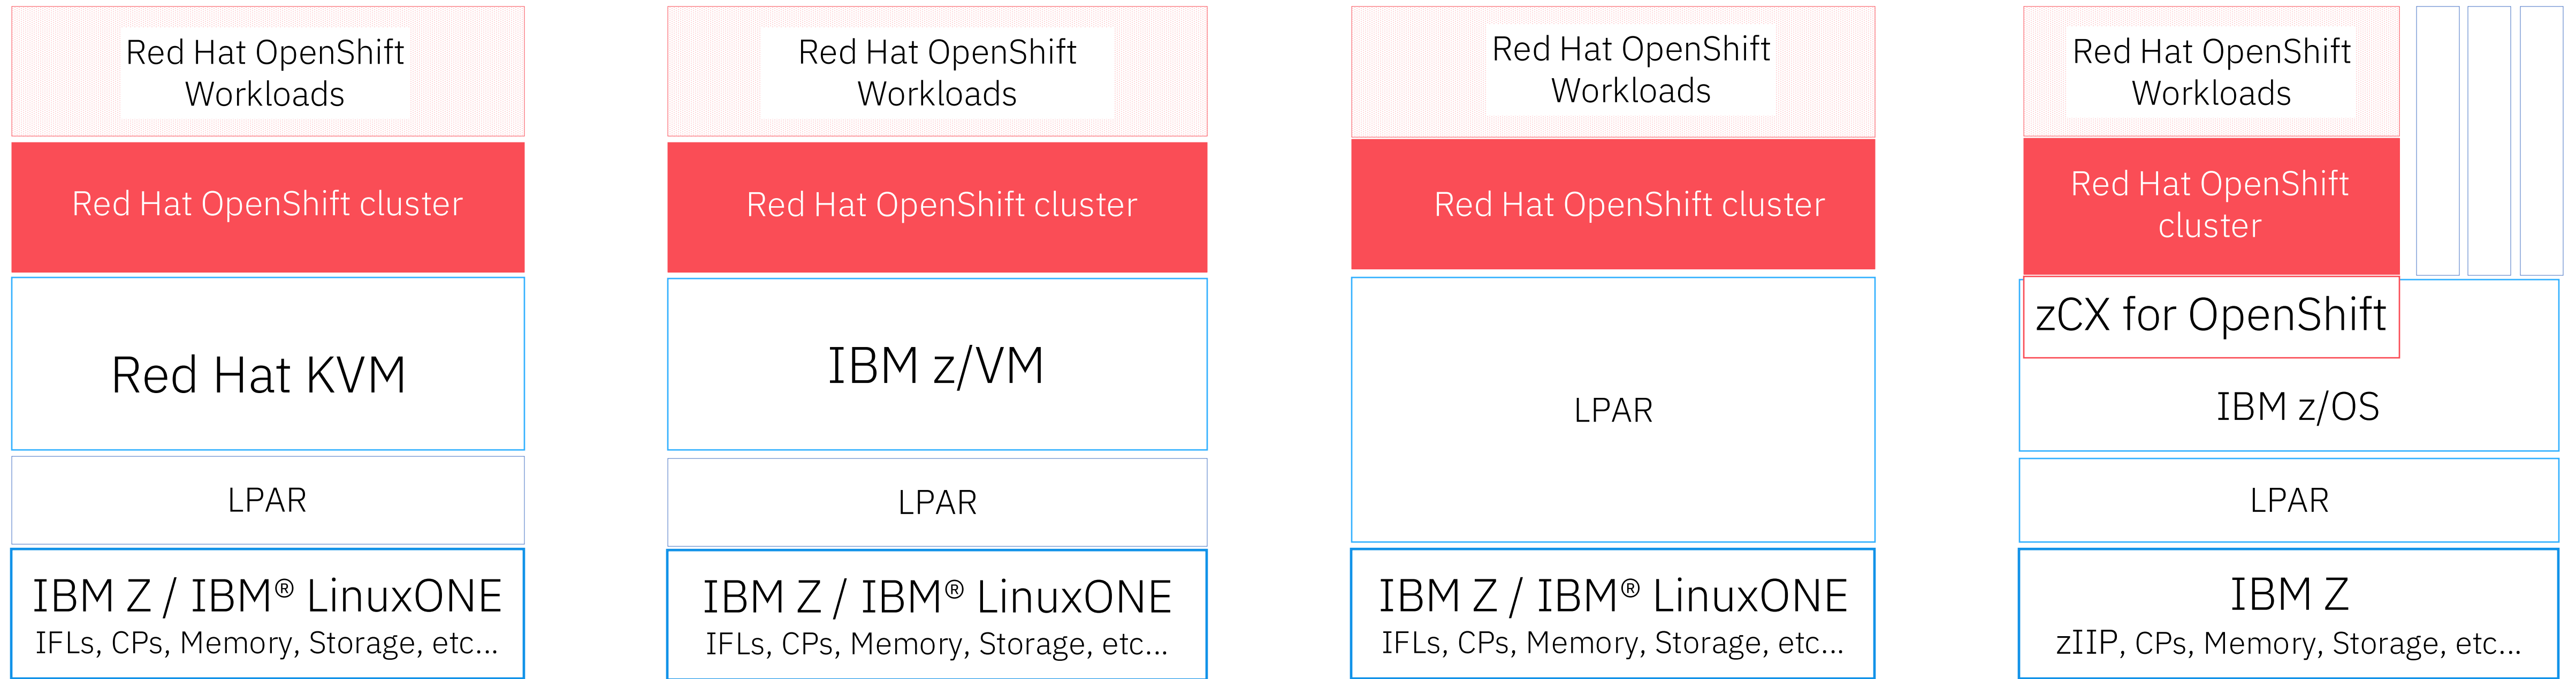
- Full cluster availability
- Full flexibility for applications and growth
- Dedicated control nodes for High Availability

Con:

- Highest cluster IFL capacity footprint (6 IFLs)
- Setup for Dev/Test env potentially oversized



Red Hat OpenShift deployment & configuration options on IBM Z and IBM® LinuxONE



- Standard 6 node cluster
- High availability cluster
- 3 node cluster
- SNO
- Multi-arch compute

Reduced cluster
footprint –

Single-node OpenShift
- Use cases

Single-node OpenShift (SNO)

- Single-node OpenShift offers both control and compute node capabilities in a single server.
- Single-node OpenShift also allows you to add additional compute nodes to your cluster. Compute node Scaling is possible
- No cluster high availability
- Provides Container availability and monitoring
- Requires 2 IFLs (Integrated Facilities for Linux), with SMT-2 enabled on installation time (for single node)
- Operation requires 0.5 IFLs on average; remaining 1.5 IFLs required for keeping SNO OCP stable and for running some workloads

Use cases:

- Test and Dev clusters - deploy and delete clusters and applications
- Clusters for apps which need a reliable runtime environment with comprehensive tools for production environments.

Apps can have short downtimes (in case SNO is updated and no dedicated compute node is added)

Example: zSCC

Cluster on speed

—

LPAR

Use cases

Provide an environment for higher performance needs and without need to have a hypervisor installed

Needs and Goals:

- Higher performance and by direct access to devices performance is increased
- Hypervisor management not required

Outcome:

- Cluster can be installed in LPAR additional to z/VM, KVM, zCX environments.
- UPI installation is supported
- Multi-Arch environments will be supported

Benefits:

Workload Optimization:

- Providing higher performance and capabilities to provide efficient throughput.

Cost Efficiency:

- Enhanced capacity available and exceed available options, which provides higher consolidation in throughput.
- Save effort for hypervisor

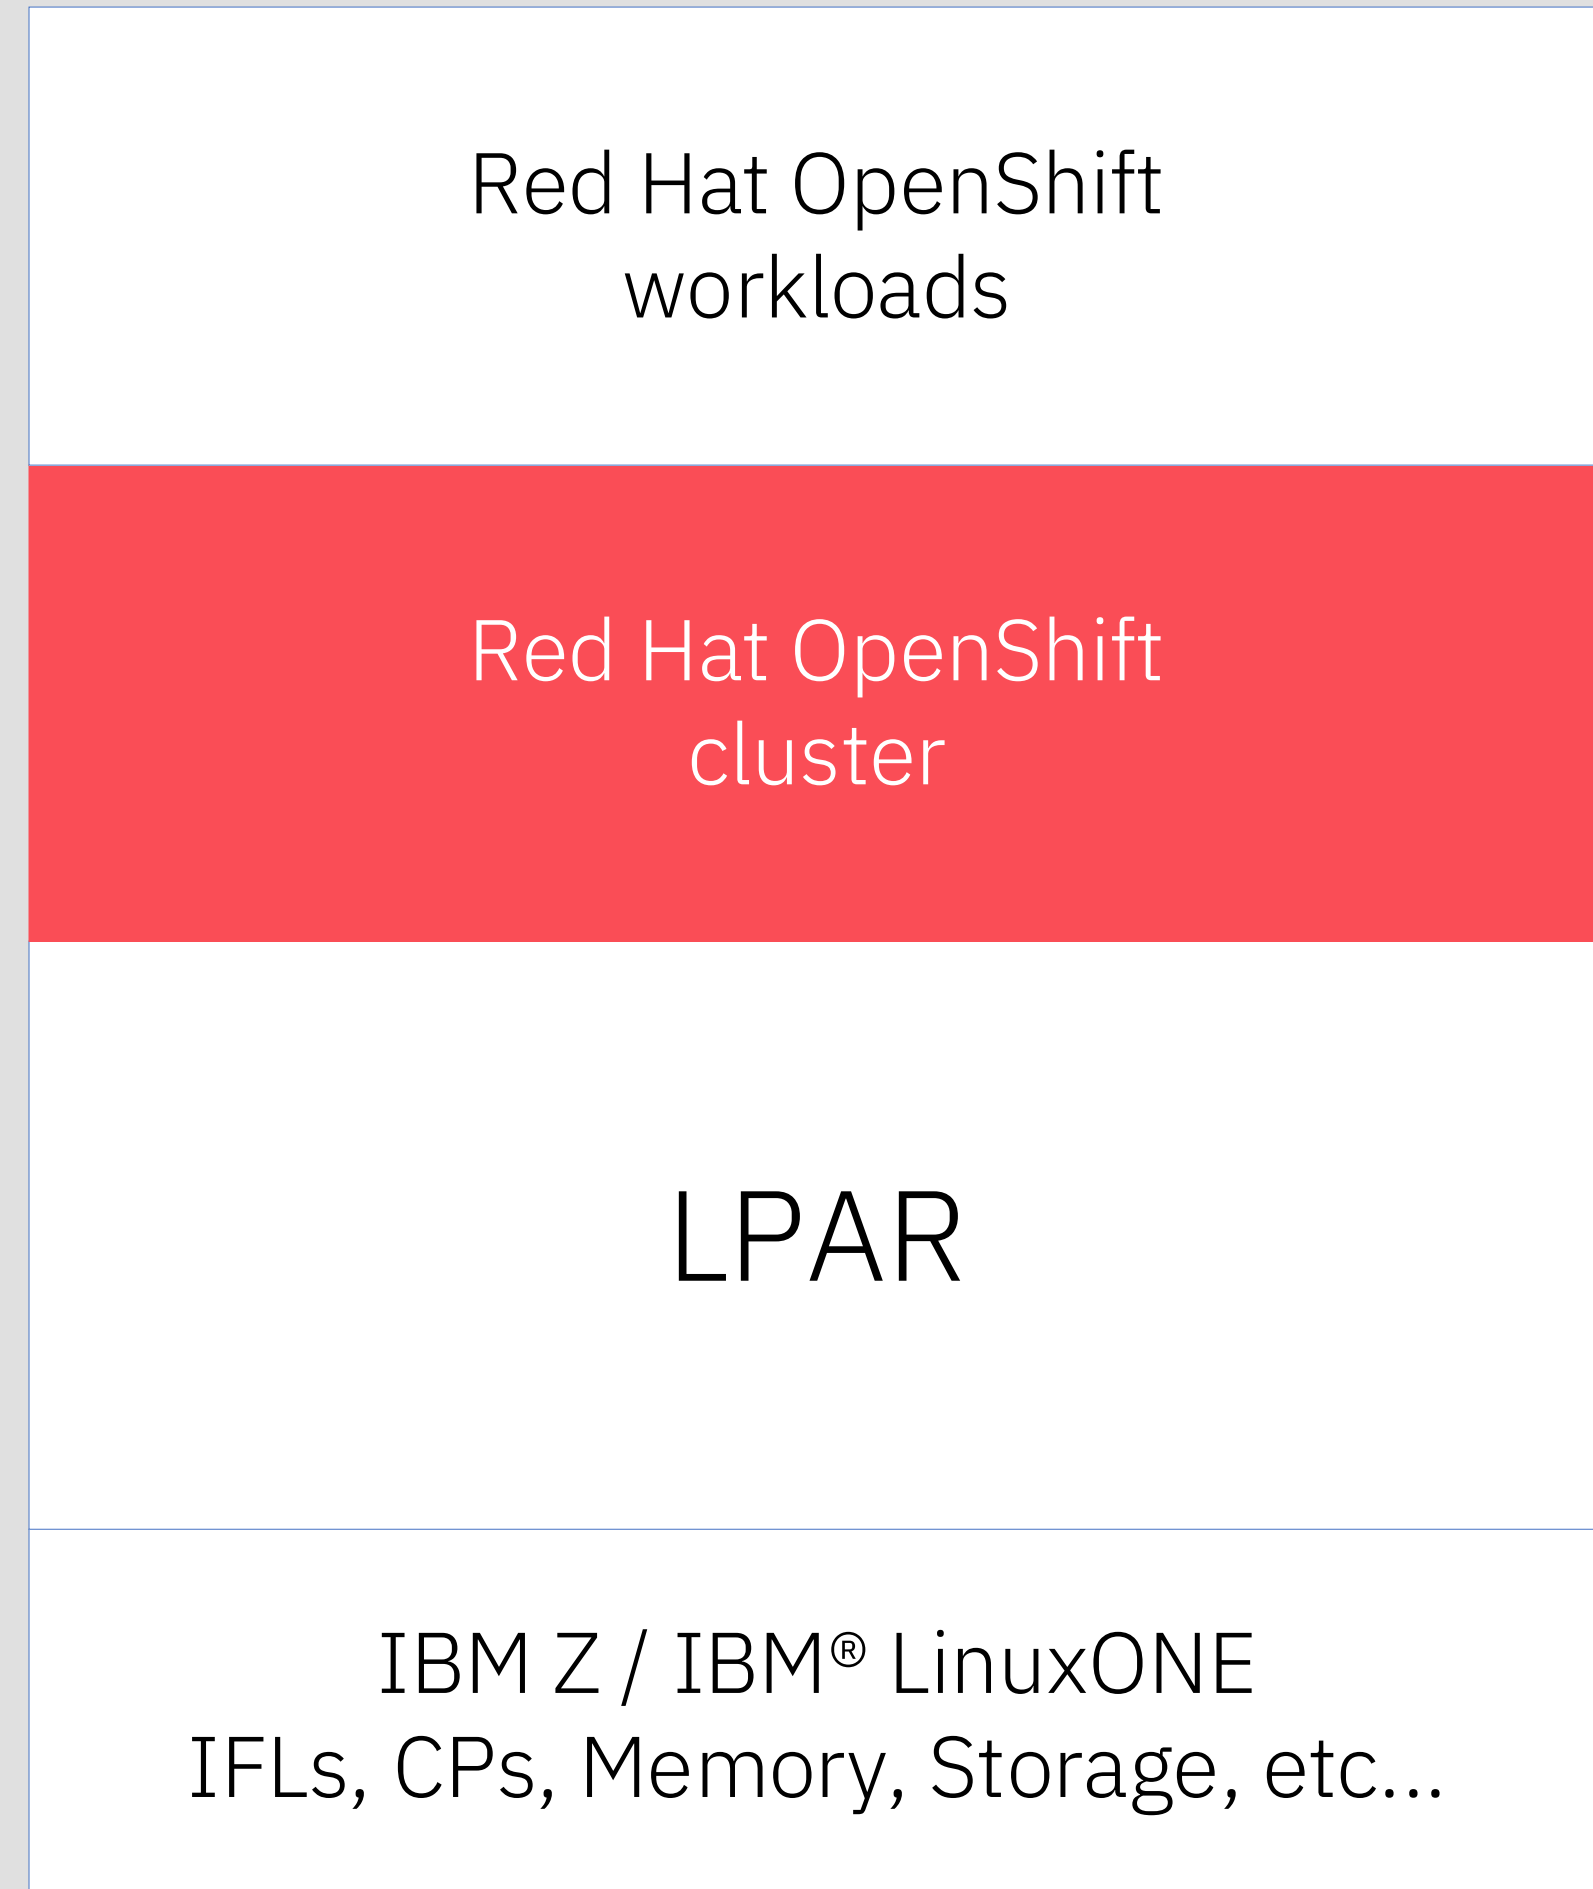
Restrictions:

- Less flexibility how to manage guests. Each OCP Node is a single LPAR. Full cluster would require 5 Nodes == 5 LPARs.

LPAR creation is recommend with DPM due to its more flexible management capabilities.

LPAR support

- *Each node is running in a separate LPAR.*
- No hypervisor required
- Depending on the setup, single-node OpenShift, 3 or more nodes cluster can be set up
- Benefits
 - Workload optimization
 - Cost efficiency



RHOCP on LPAR

Red Hat OpenShift provides in LPAR the same capabilities as on virtualized environments.

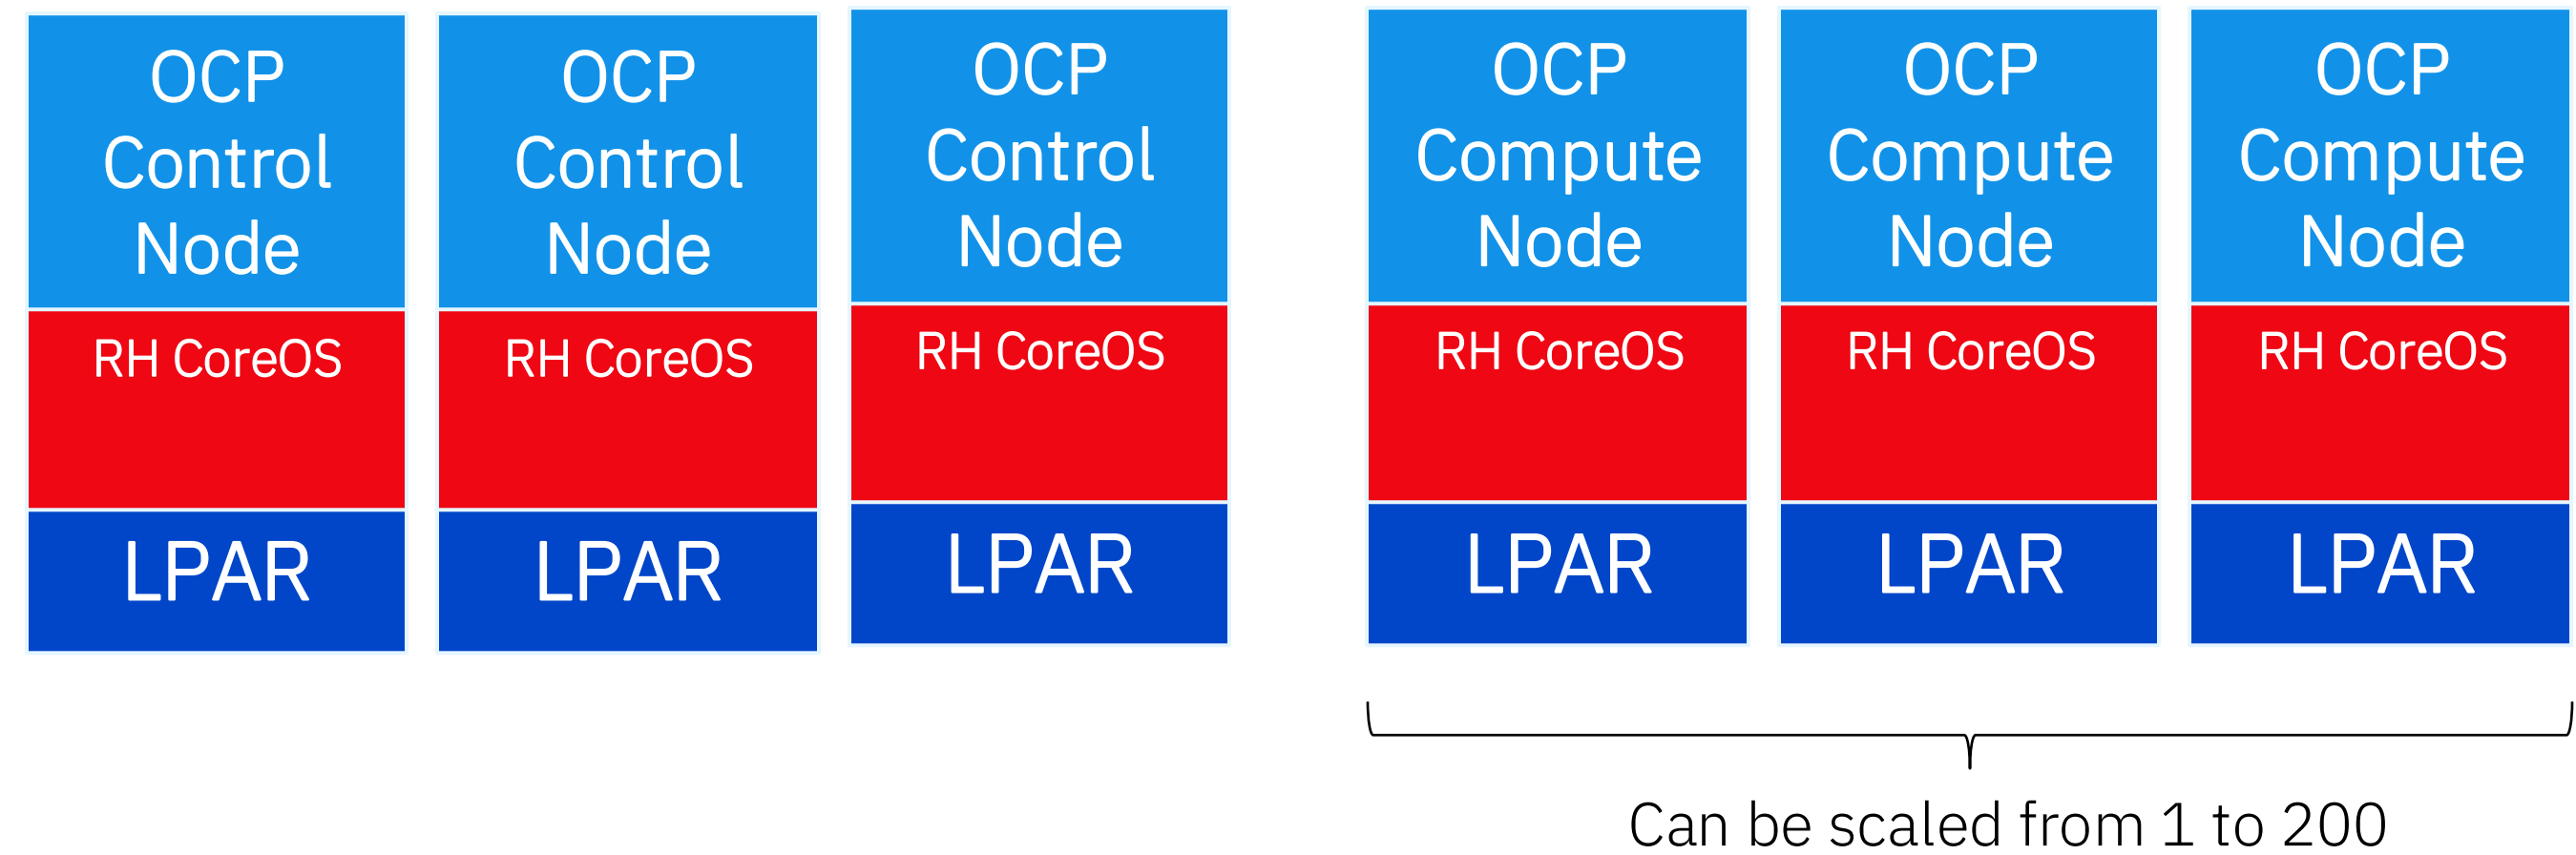
- Operational model
- Scalable runtime infrastructure
- High Availability capabilities
- Network and storage operation

UPI Installation support:

- DPM: boot with ftp, PXE
- Classic: boot with ftp

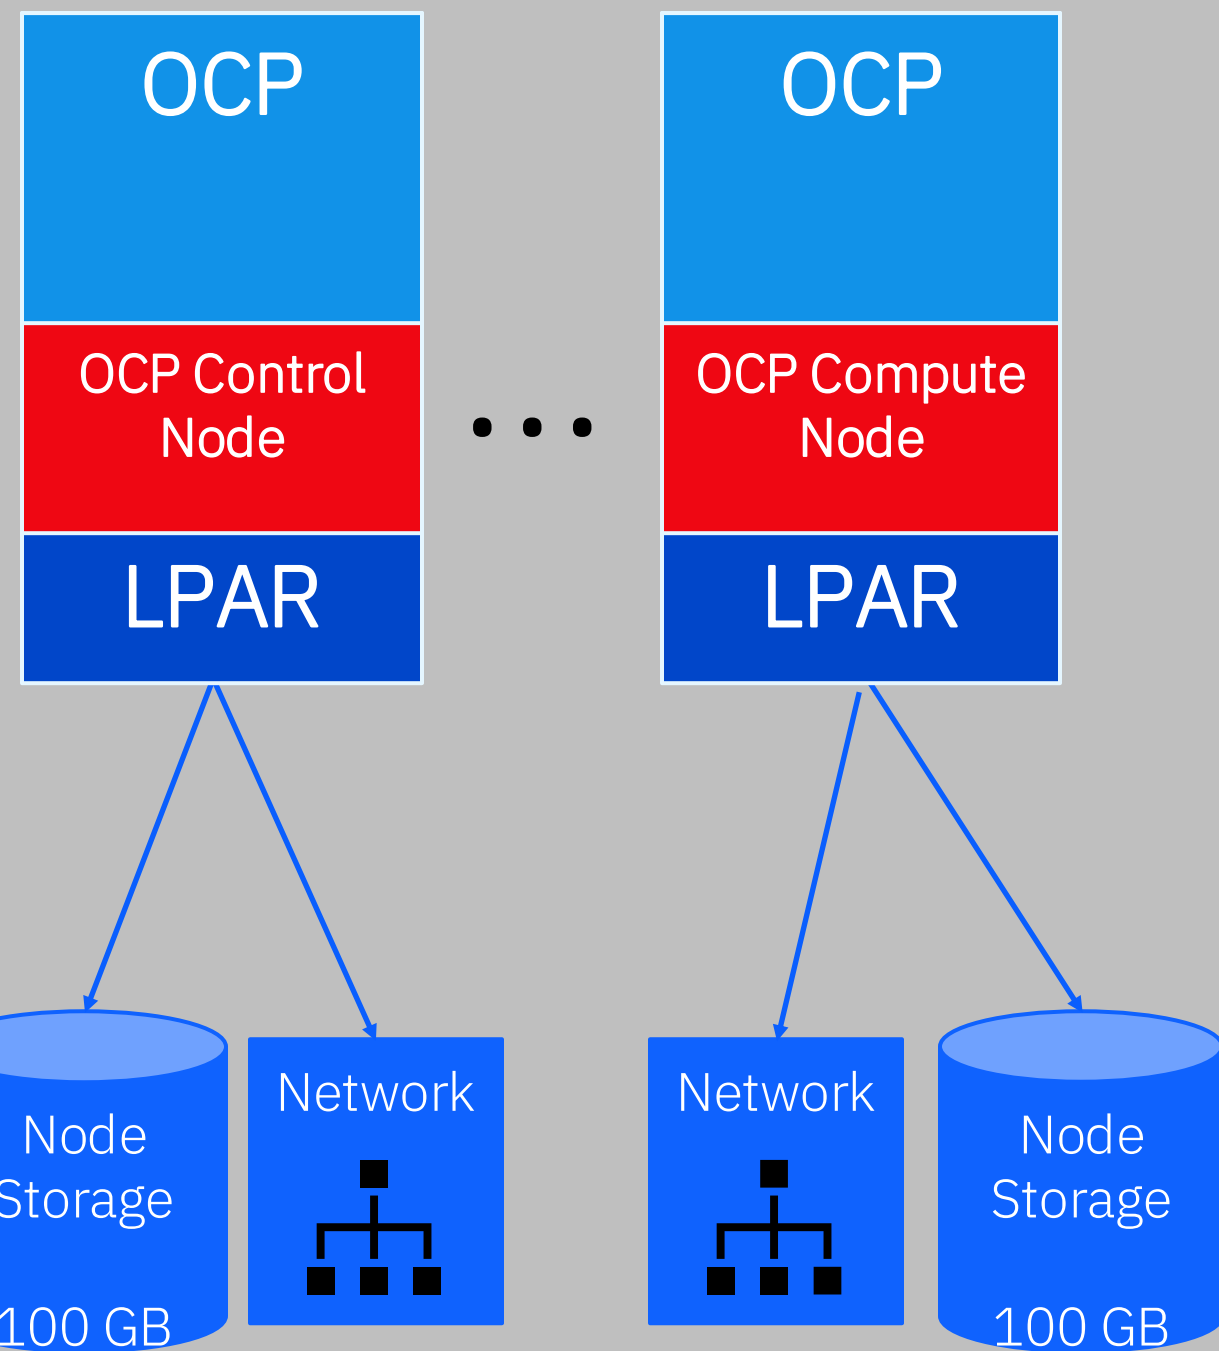
Red Hat OpenShift is by default covered cross nodes

Control nodes contain operator artifacts and configuration data



LPAR requirements

IBM Z / LinuxONE Machine



Storage

1 Disk per node for RHCOS (holding components for control or compute) as root device

Storage Types supported:

ECKD
FCP

Network

1 Network adapter for connection of the cluster network

Supported Network types:

OSA
ROCE
Hipersocket (only internal network external connection required)

CPU / Memory

Control node minimum:
2 shared IFLs
16 GB RAM

Compute node minimum:
1 shared IFL
8 GB RAM

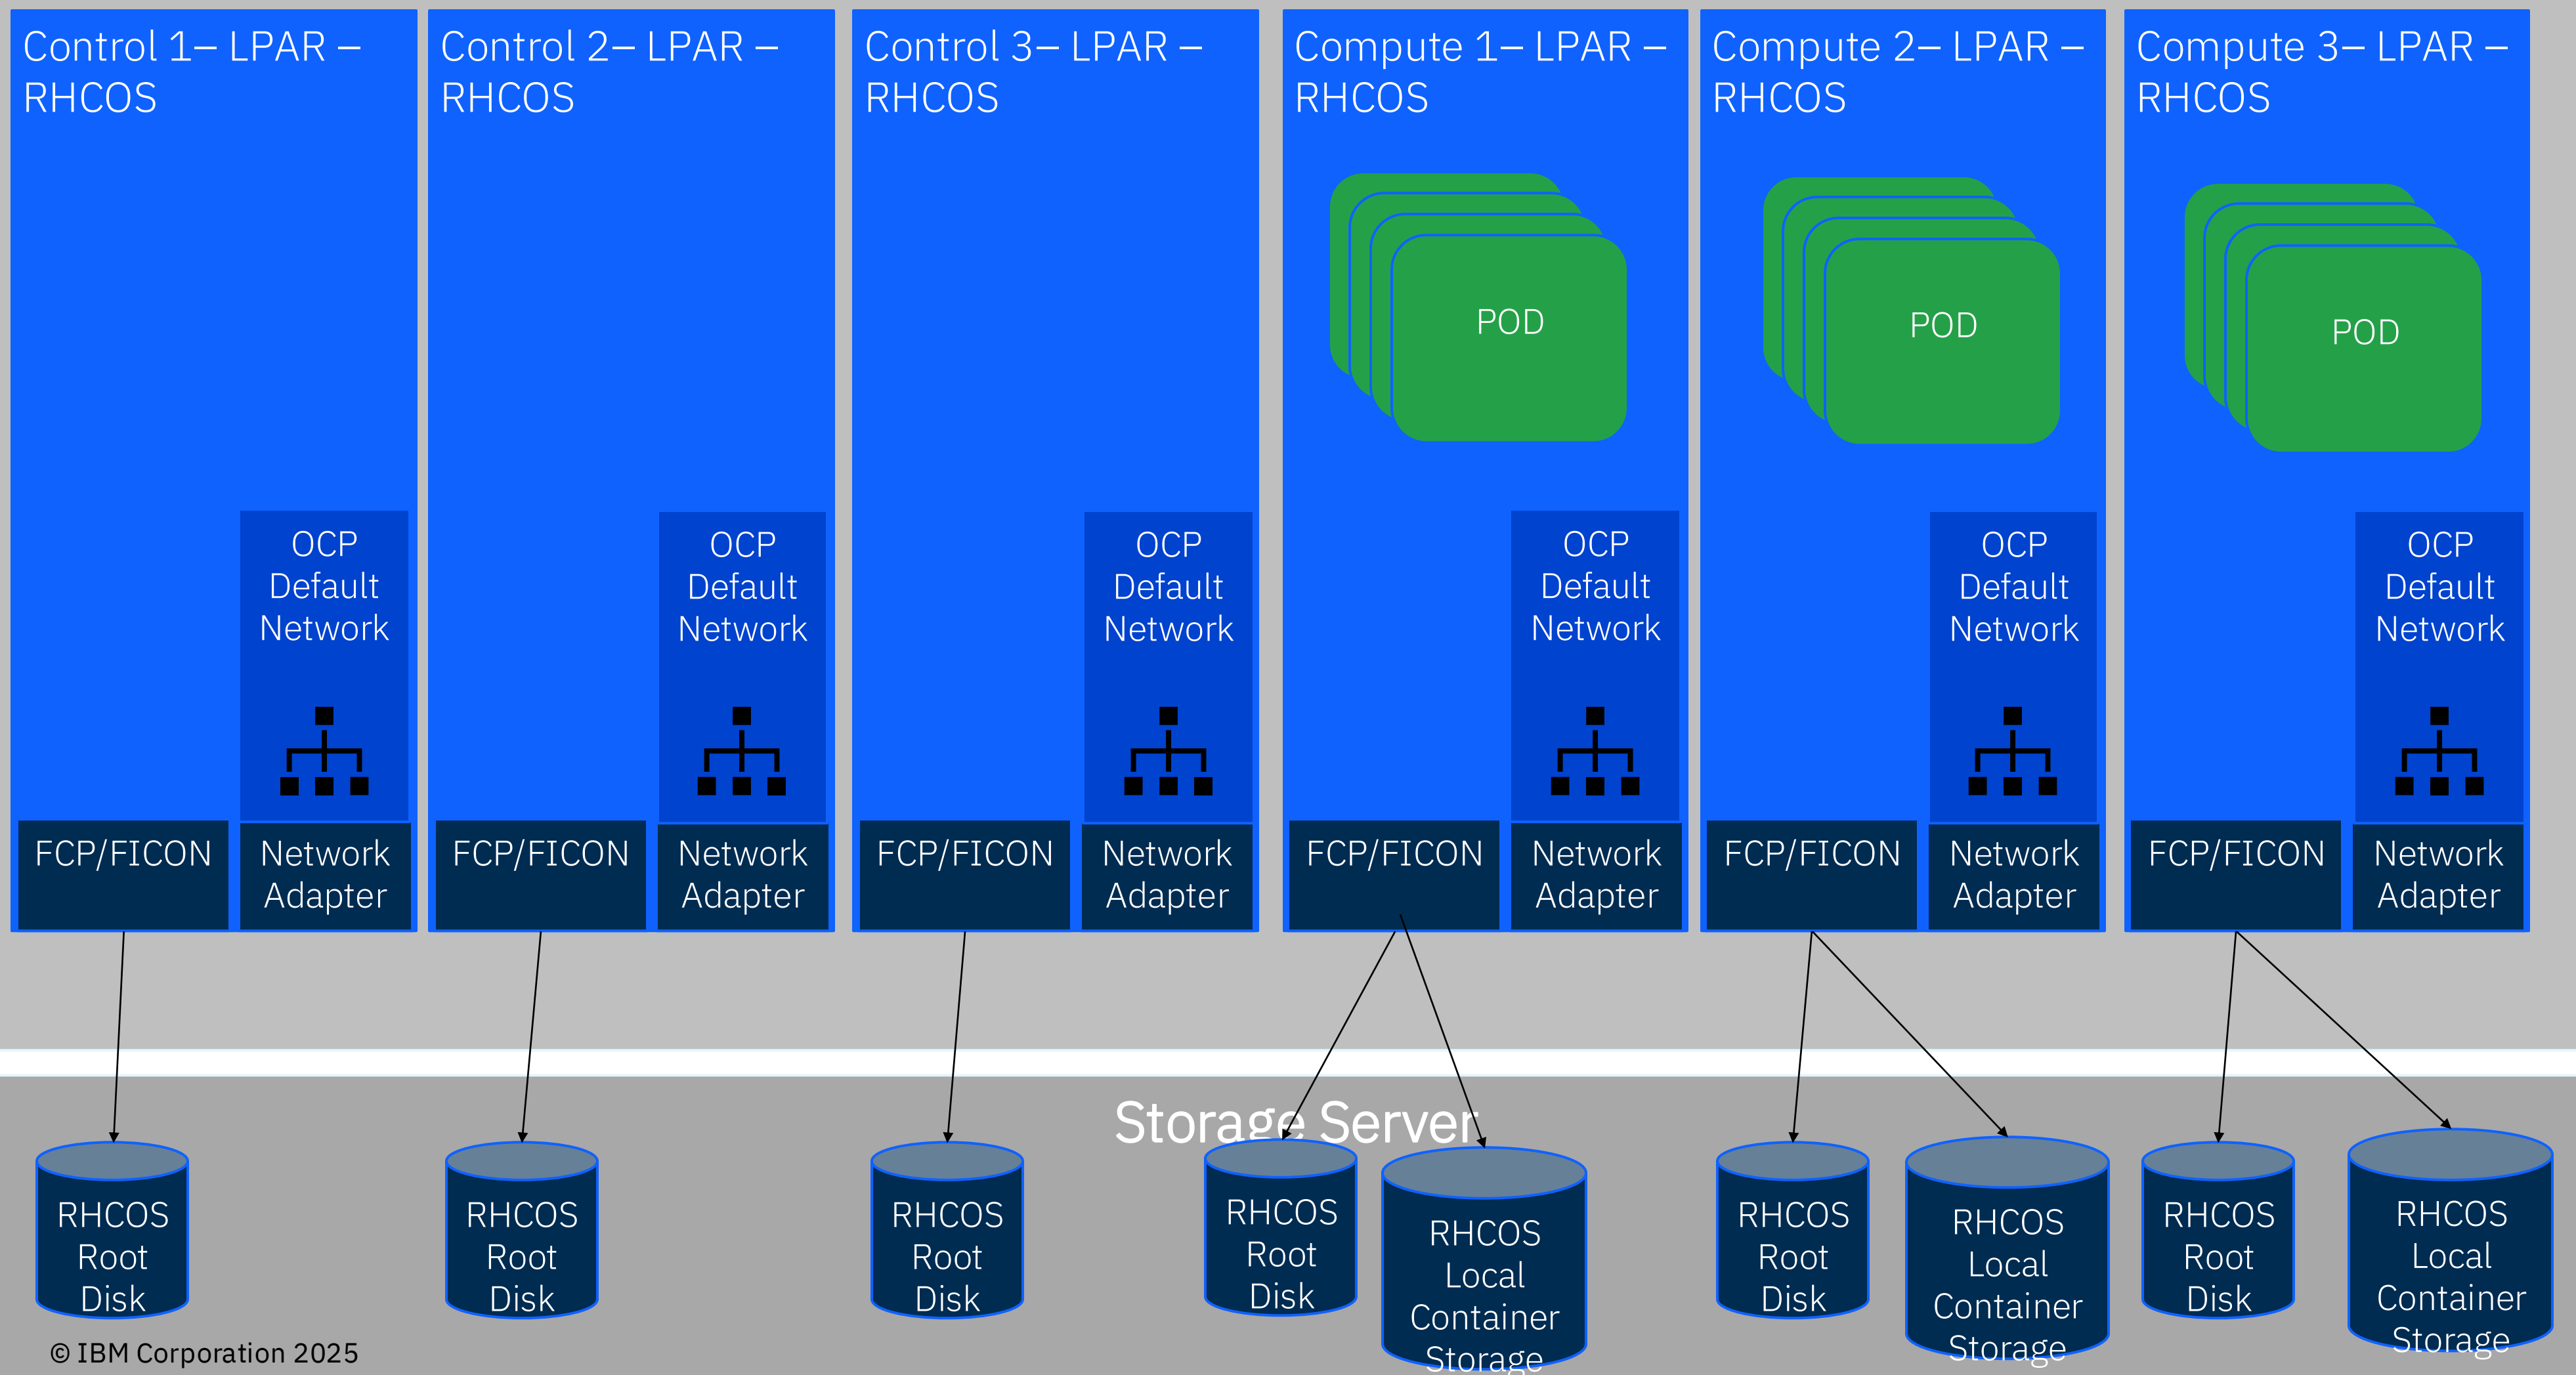
Production Cluster minimum contains
3 control
3 compute

6 IFLs can be shared among all nodes.

Resources need to be adjusted to the real need!

RHOCP setup – node setup

IBM Z / IBM LinuxONE Machine



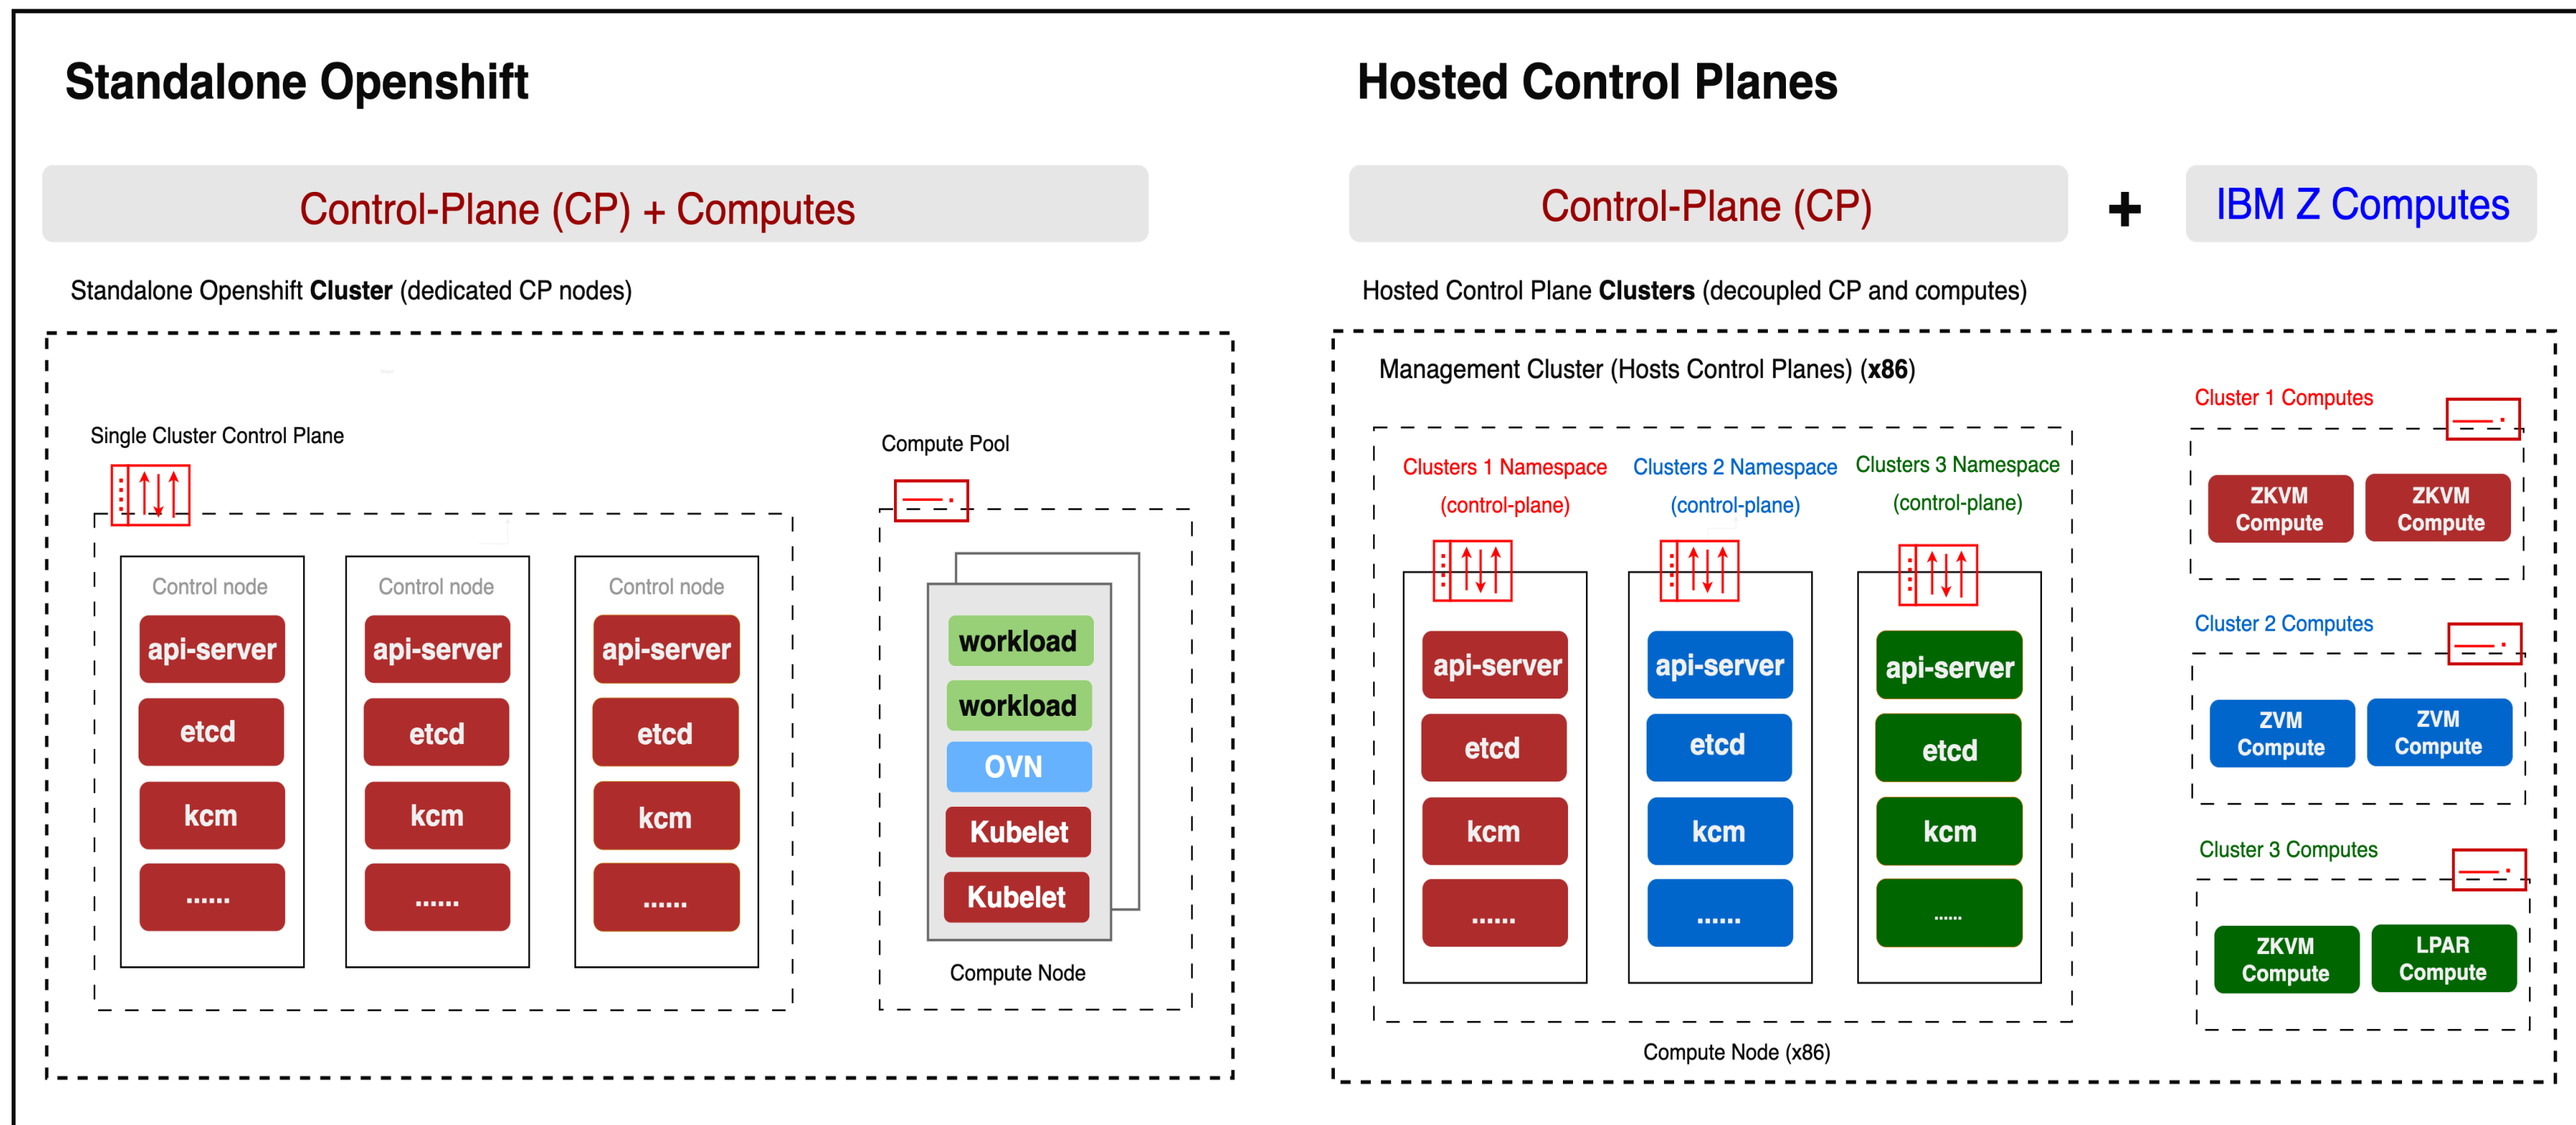
- Compute + Control require minimum 1 RHCOS root disk
- Pod persistent storage is either logically (for Spectrum Scale) or physically for Local Storage Operator / ODF.
- There is minimum 1 network used for cluster communication and egress / ingress from/to pods
- Nodes are reachable via routed network and can “see” (DNS-resolveable and able to ping) each other.
- Cluster can be installed on a single or crossing multiple machines

Manage consolidated cluster

—

Hosted Control Planes

Hosted Control Planes - overview



What are Hosted Control Planes

- A way to host OpenShift control planes on top of an existing OpenShift cluster.
- Decouples the control plane from the compute nodes
- Provides the flexibility to focus on workloads and reduce management overhead.
- Control planes no longer required dedicated VMs

Low CAPEX and OPEX costs
(bundling of CPs + CP as pods)



Central Management of CPs
(easy operation & maintenance)



Multi-arch support
(e.g. CP x86, workers ARM)



Network & Trust
segmentation



Mixed IaaS For CP and
Workers



Fast cluster bootstrapping
(CP as Pods)



As of now, HCP with s390x compute nodes on x86 arch management OpenShift clusters are supported.

Hosted Control Planes (HCP) - Use cases

Less infrastructure overhead, reduces costs

- Control planes no longer required as dedicated VMs

Easy to Scale

- Easy to create multiple clusters

Provide a hosted RHOCP on Premise

- Implementation of self service cloud instances for dev/test

Consolidation of control plane on x86. Compute on IBM Z and IBM® LinuxONE


- Control and data planes can run on different hardware

Fast to Provision (Time-To-Value)

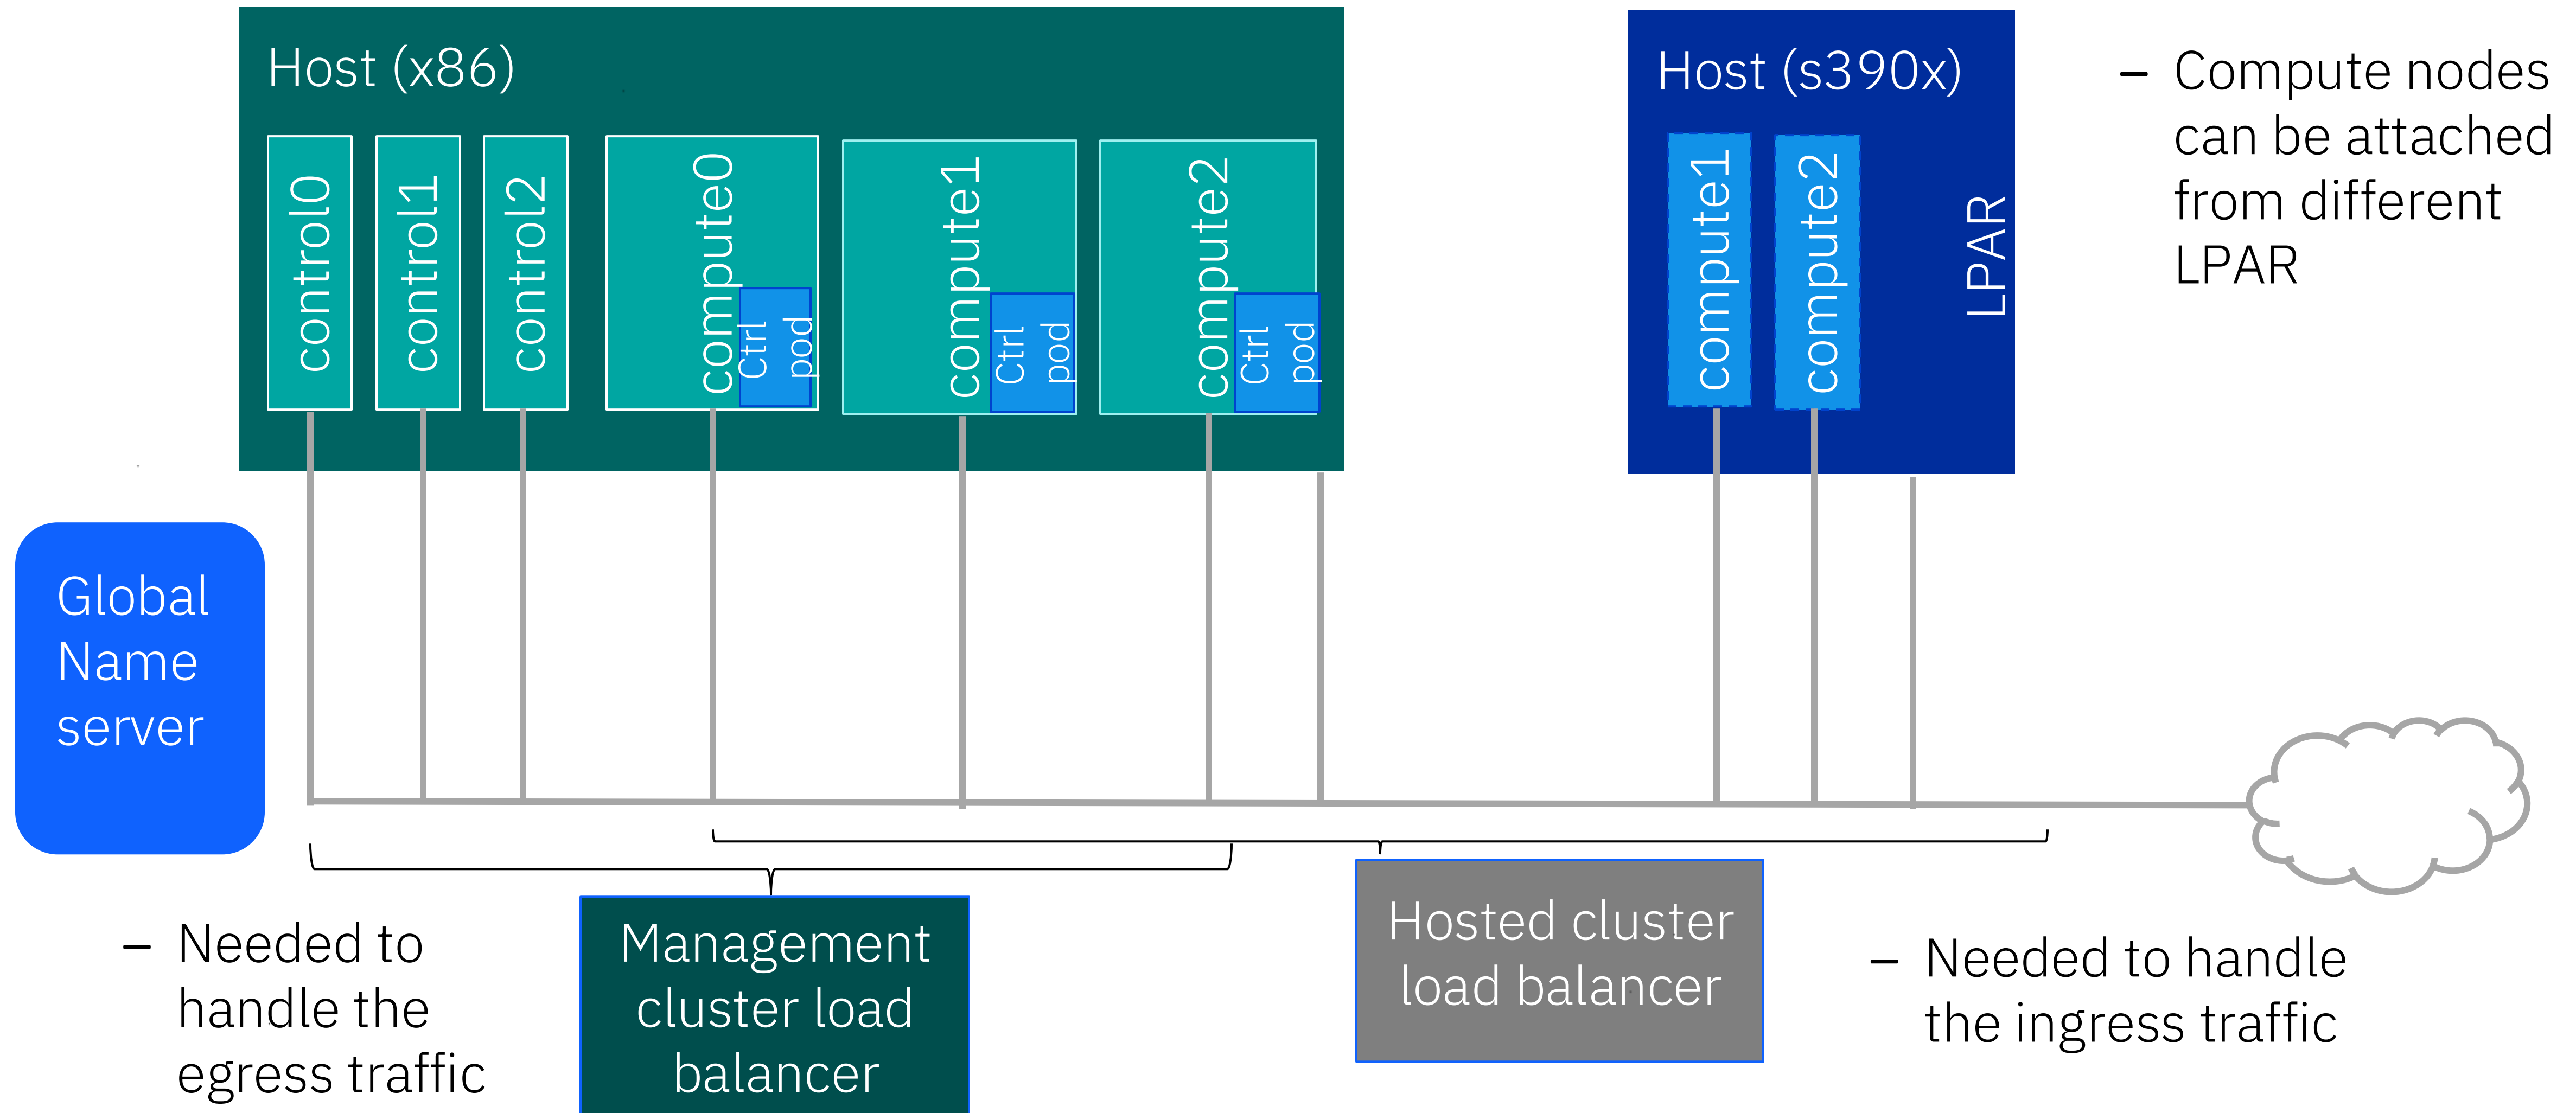
- Control plane is deployed as pods in a namespace

Centralized Management of cluster creation

- Single management of all clusters



Hoster Control Planes: Cluster setup with x86 HCP and s390x compute nodes



Hosted Control Planes – features

Platform (Computes)	Network	Boot Storage (Disk Type)
KVM on IBM Z	MacVTap	QCOW DASD
z/VM	RoCE OSA vSwitch HiperSocket	FCP DASD
LPAR on IBM Z	OSA RoCE HiperSocket	FCP DASD NVMe

Important Features:

- Supports disconnected installation of hosted control plane and data plane (4.17)
- Supports Isolated upgrades of control-plane and data-plane
- Upgrade to data-plane should be done until control-plane version
- Ease of autoscaling the compute nodes via Agents for KVM on IBM Z.

Delivery Notes:

- Hosted Control Plane is delivered as part of Multi Cluster Engine Operator (MCE).
- Hosted Control Planes with Agent-based Installation method is supported.

Building one cluster
across architectures

—

Multi-Architecture
compute

Use cases

User with an existing OpenShift cluster seeks to *expand workload options by integrating nodes based on another CPU architecture.*

Needs and Goals:

- Introduce workload types optimized for specific architecture (e.g., mainframe applications s390x or x86 based container).

Co-location

- Preserve the flexibility and efficiency of the existing infrastructure and add a second architecture to close gaps.

Outcome:

- Your OpenShift cluster seamlessly combines the capabilities of both x86 and s390x architectures while ensuring compatibility with your existing applications and services.

Benefits:

Workload Optimization:

- Incorporate s390x/x86-specific workloads, capitalizing on the unique strengths of this architecture.

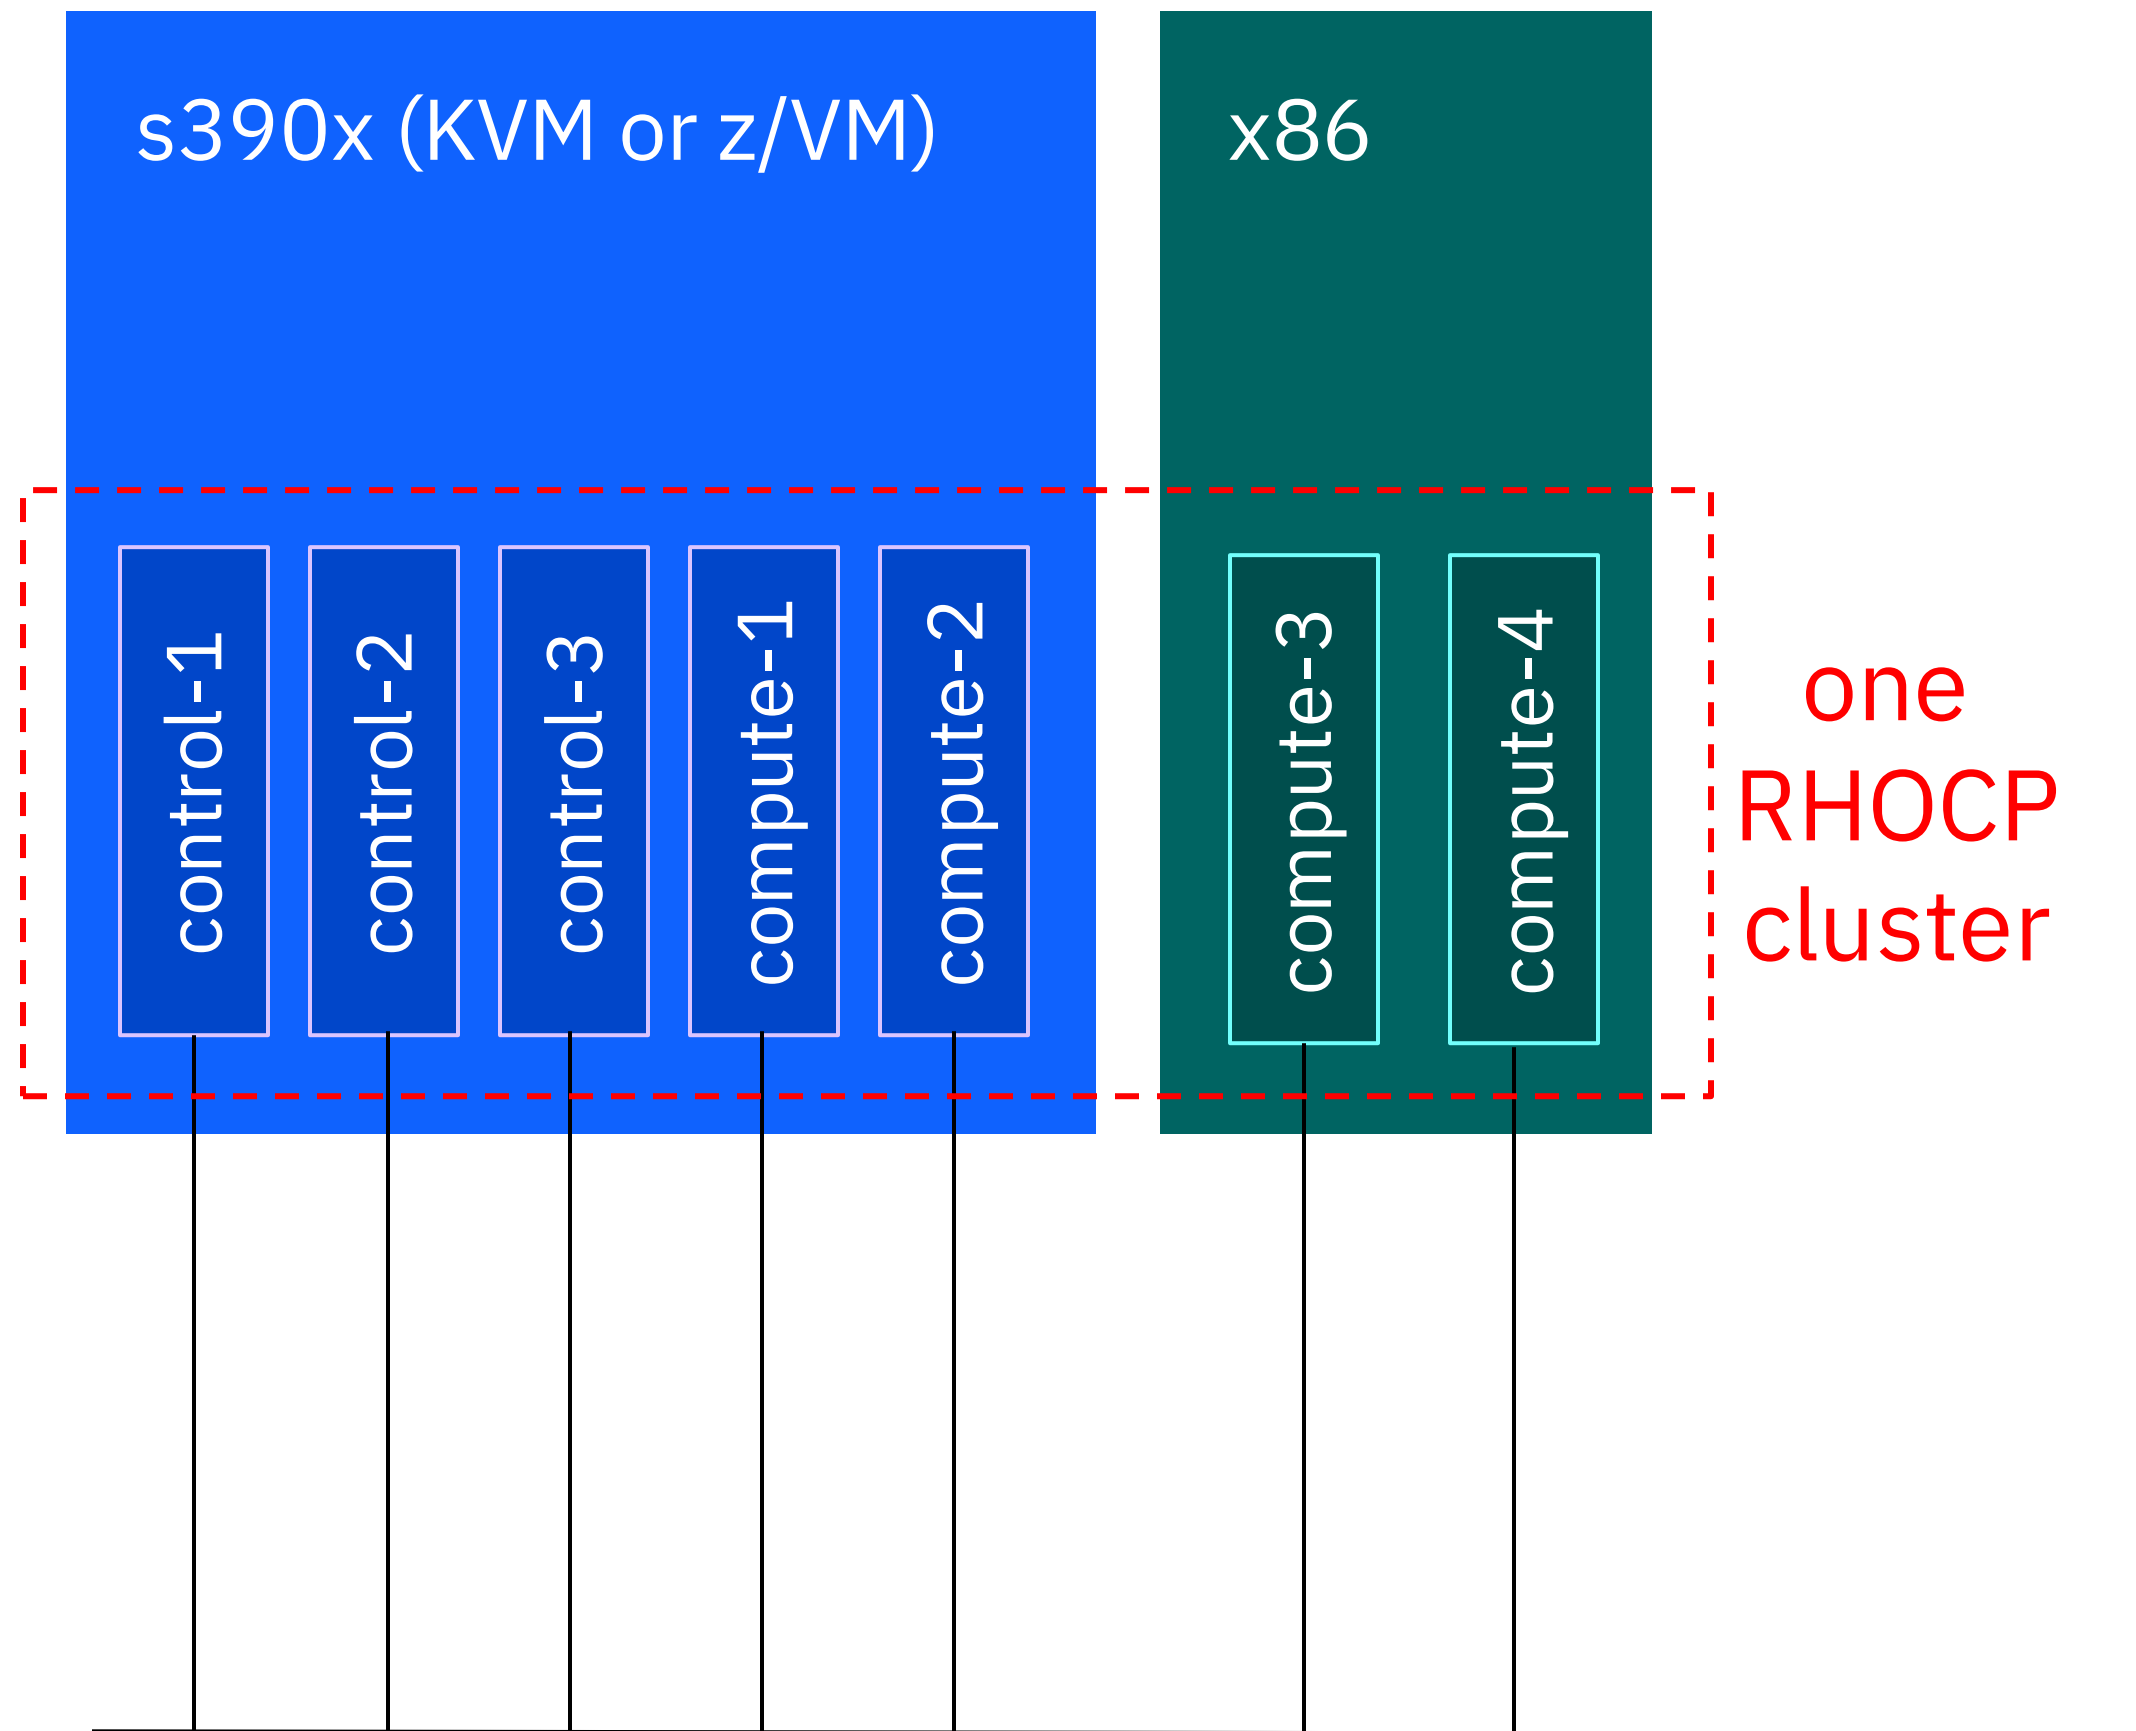
Cost Efficiency:

- Avoid the need for a separate infrastructure solely for s390x or x86 workloads. By integrating s390x/x86 compute nodes, you can leverage architectural diversity without additional operational expenses.

Resource Flexibility:

- With the addition of s390x/x86 nodes, resource allocation becomes more flexible. Customize resource allocation to meet the specific requirements of different workloads, ultimately improving the overall performance of your cluster.

Multi-Architecture compute

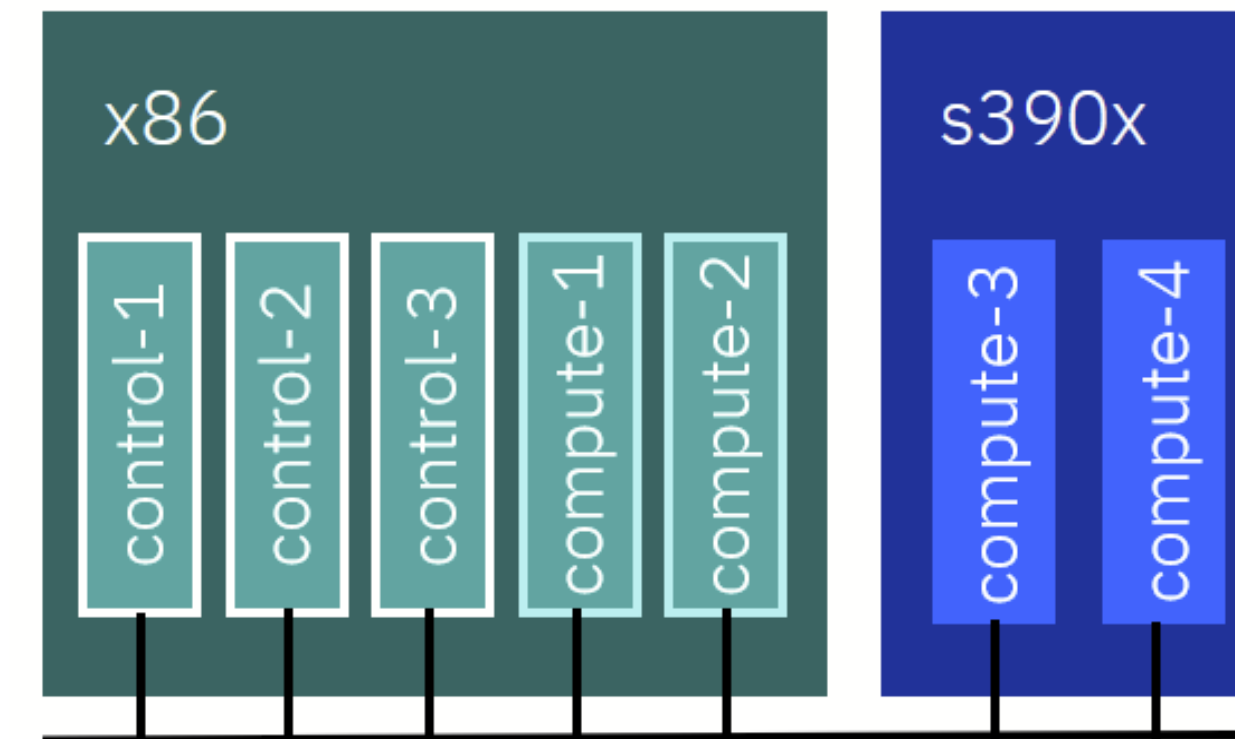


- Add additional x86 compute nodes to an existing OpenShift cluster running on z/VM or KVM on IBM Z.
- Now, both s390x control planes and x86 control planes with mixed s390x and x86 compute nodes are supported.
- OpenShift Multi-Architecture Compute clusters look and work like single-architecture clusters.
- Update path for moving from single-architecture clusters to multi-architecture version supported.
- Current restrictions:
 - Not all Add-ons work out of the box on multi-architecture compute
 - Understanding of the cluster is required for schedule applications

Multi-Architecture compute

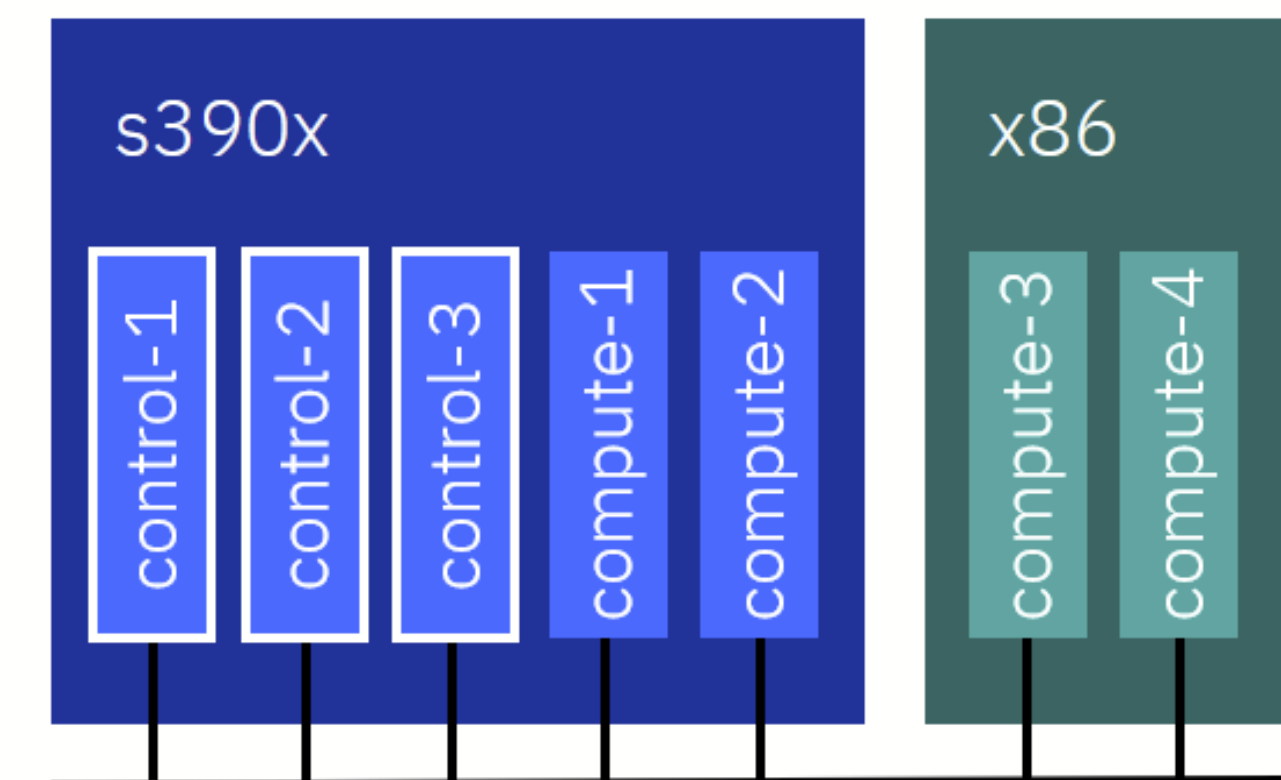
4.14

- A cluster with control planes and compute nodes on x86 architecture.
- With additional s390x compute nodes.



4.15 (new)

- A cluster with control planes and compute nodes on s390x architecture.
- With additional x86 compute nodes.



Multi-Architecture deployments – considerations

Multi-Arch Compute

Single cluster contains nodes out of different architectures

Use Case:

- Solution deployments which require different architecture as application dependency
- Pods should run close to a Database on s390x or frontend be on x86 but backend on S390x

Application container can start on each node. Scheduling rules might be required.

Applications require multi-arch enabled container images as manifest listed containers.

Cluster must be installed as platform “none”

Existing x86 cluster based on other platforms e.g. vsphere or baremetal require re-install to “none”

Hosted Control Plane

Multi-Arch by hosting control nodes on x86 but compute on s390x.

Use Case:

- Consolidate Control Nodes
- Simplified deployment
- Can be used on existing x86 cluster to deploy IBM Z / LinuxONE cluster nodes

Hosted Cluster can execute only on single architecture. Compute nodes cannot be multi-arch nodes.

Only hosted control nodes are on x86.

Multi-architecture compute

Possible combinations of cluster and additional compute nodes

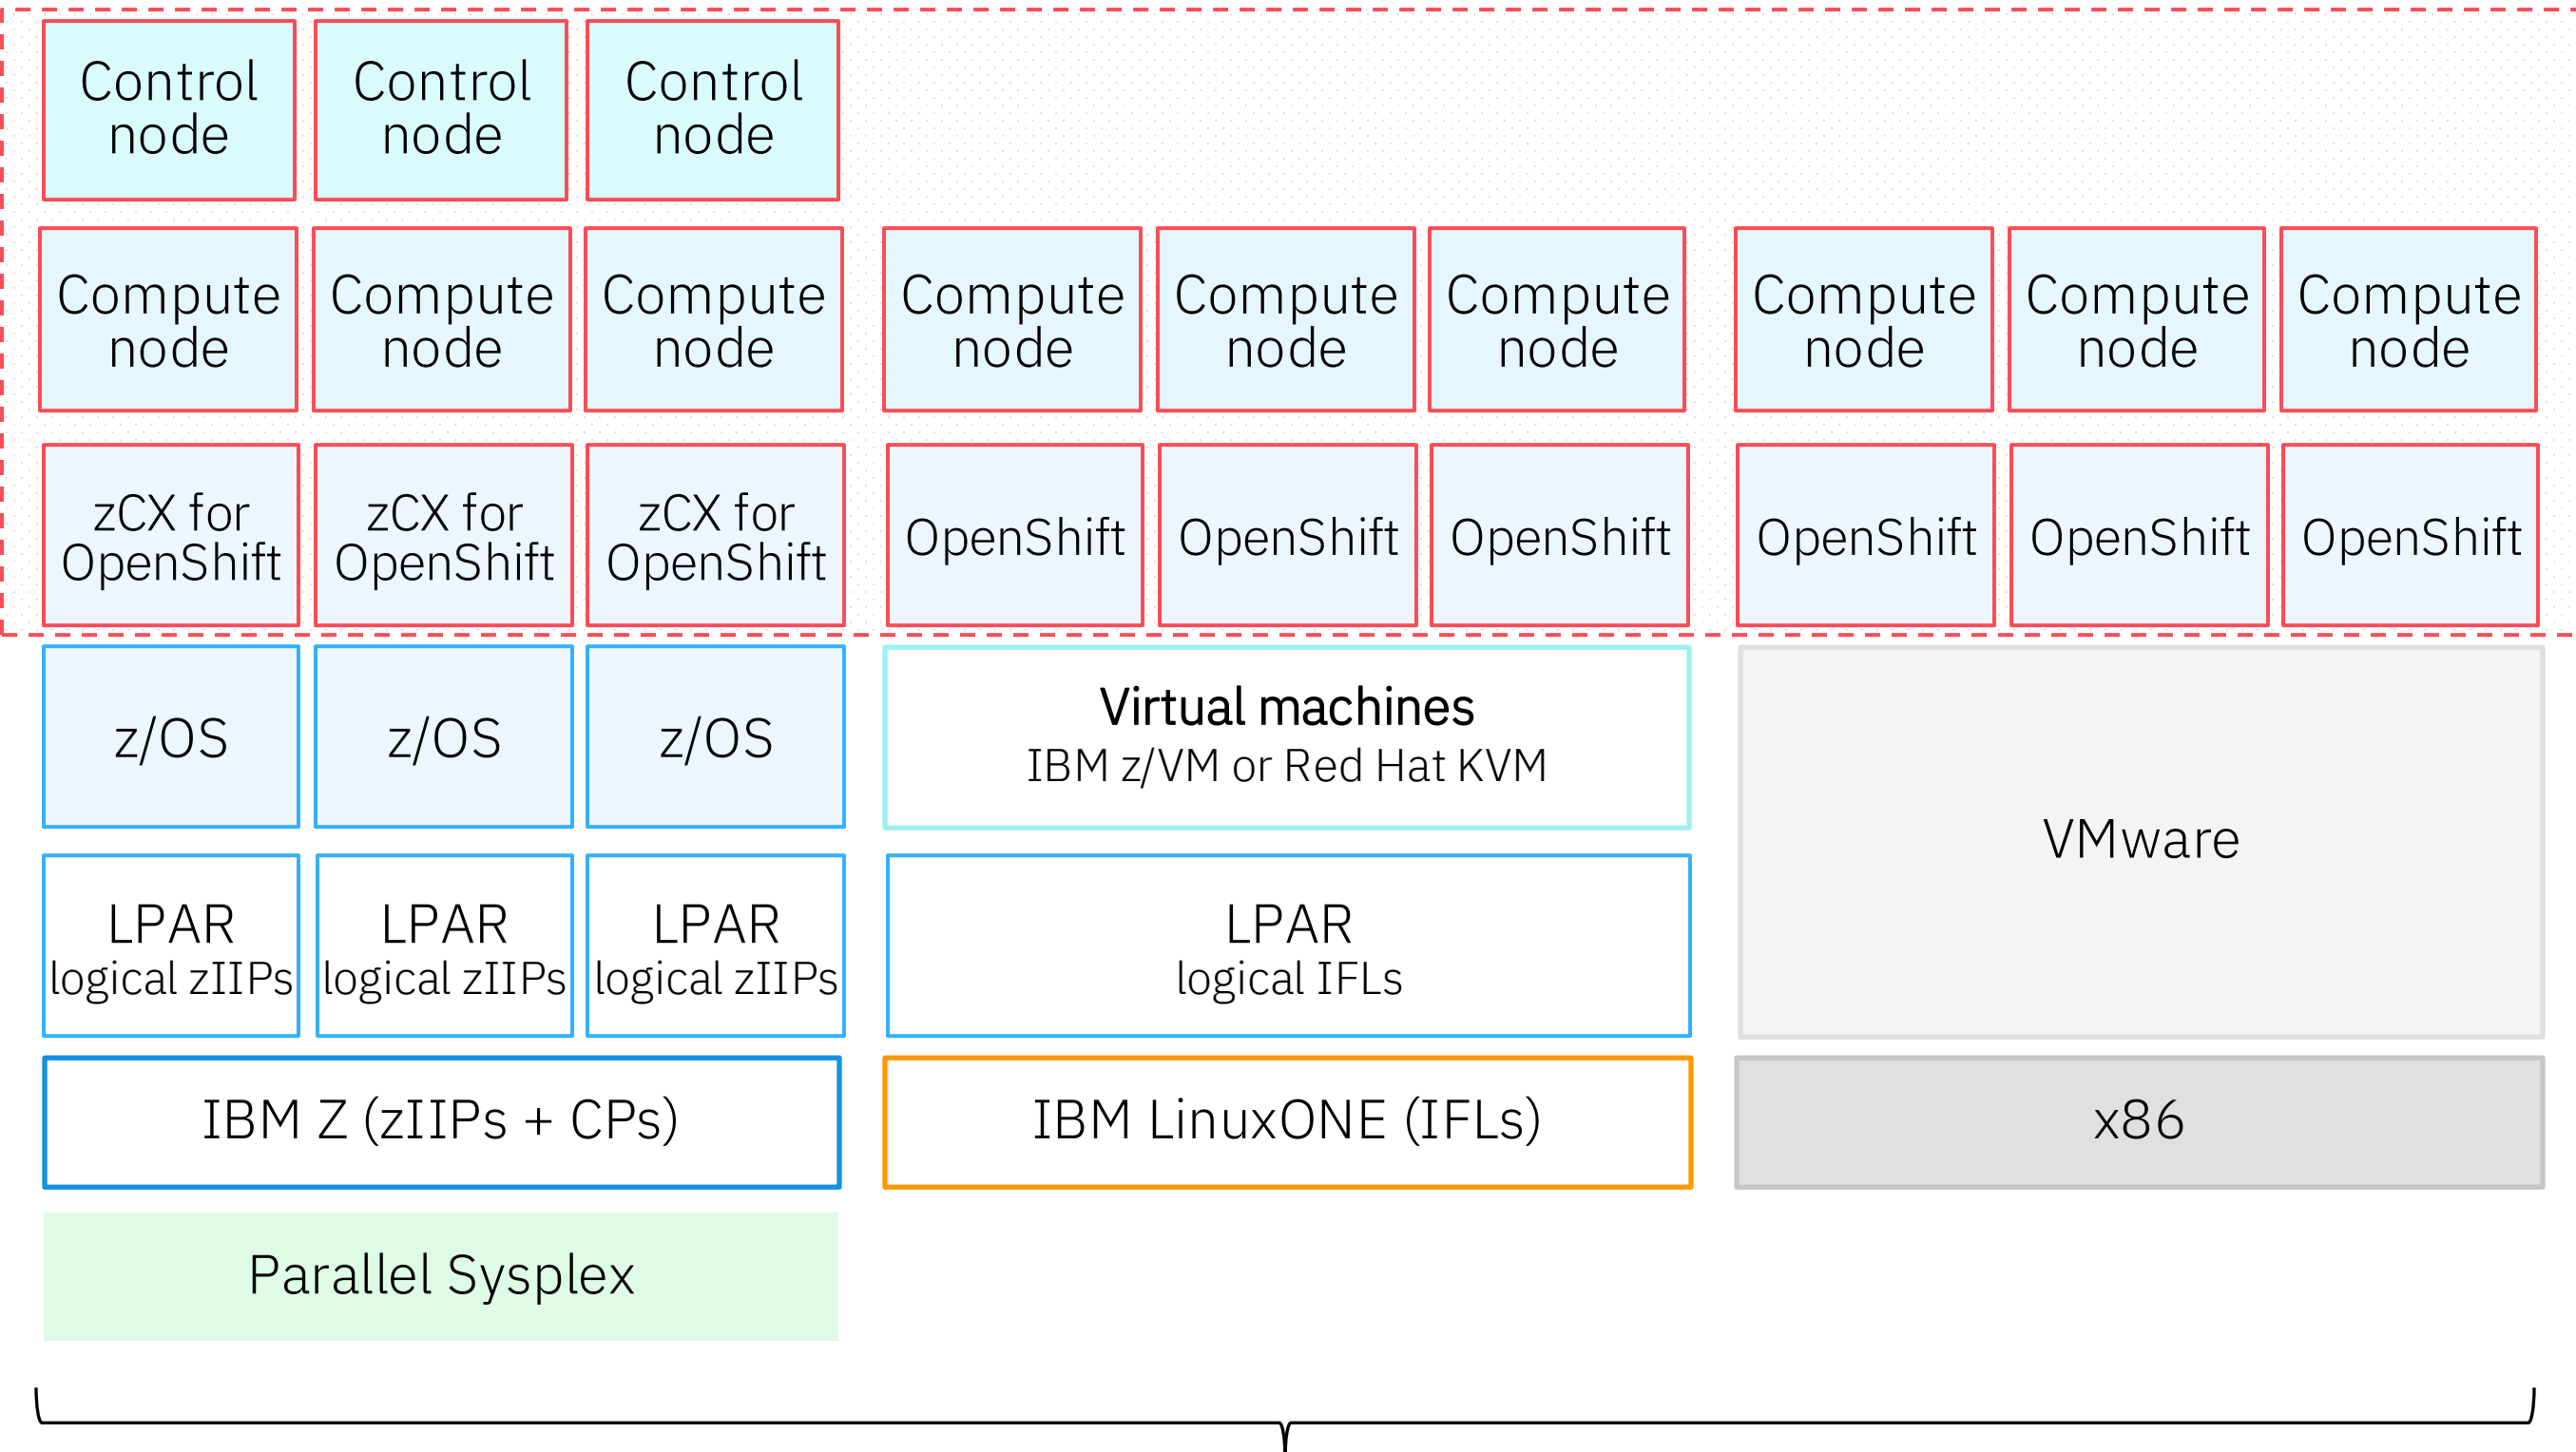
Cluster		additional compute nodes		
Cluster platform	Deployment option	IBM Z	IBM® LinuxONE	x86
IBM Z	Virtual machines*	<ul style="list-style-type: none"> LPARs zCX for OpenShift 	<ul style="list-style-type: none"> Virtual machines* LPARs 	<ul style="list-style-type: none"> x86
	LPARs	<ul style="list-style-type: none"> Virtual machines* zCX for OpenShift 	<ul style="list-style-type: none"> Virtual machines* LPARs 	<ul style="list-style-type: none"> x86
	zCX for OpenShift	<ul style="list-style-type: none"> Virtual machines* LPARs 	<ul style="list-style-type: none"> Virtual machines* LPARs 	<ul style="list-style-type: none"> x86
IBM LinuxONE	Virtual machines*	<ul style="list-style-type: none"> LPARs zCX for OpenShift 	<ul style="list-style-type: none"> LPARs 	<ul style="list-style-type: none"> x86
	LPARs	<ul style="list-style-type: none"> Virtual machines* zCX for OpenShift 	<ul style="list-style-type: none"> Virtual machines* 	<ul style="list-style-type: none"> x86
x86	...	<ul style="list-style-type: none"> Virtual machines* LPARs zCX for OpenShift 	<ul style="list-style-type: none"> Virtual machines* LPARs 	<ul style="list-style-type: none"> ...

* Virtual machines based on IBM z/VM or Red Hat KVM

Multi-architecture compute – example

One Red Hat OpenShift cluster spans across

- zCX for OpenShift on **IBM Z** – as part of a Parallel Sysplex,
- virtual machines on **IBM LinuxONE**,
- and virtual machines on **x86**



Planning hint:

When IBM Z / LinuxONE is included in the multi-architecture deployment, then platform "none" must be set during the installation of all architectures included in the 'multi-architecture compute' deployment

IBM