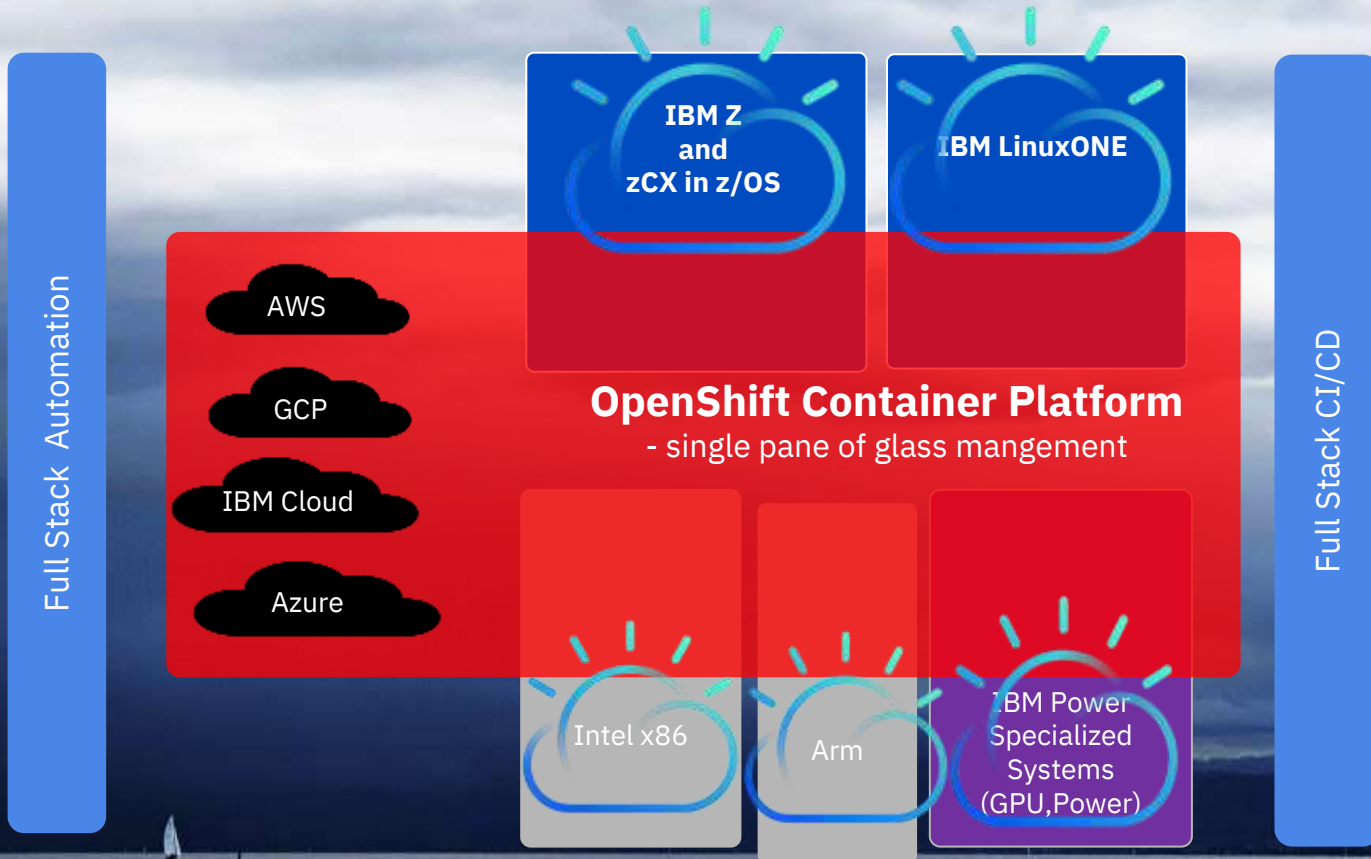


Multi-Architecture Options with Red Hat OpenShift on IBM Z and IBM LinuxONE

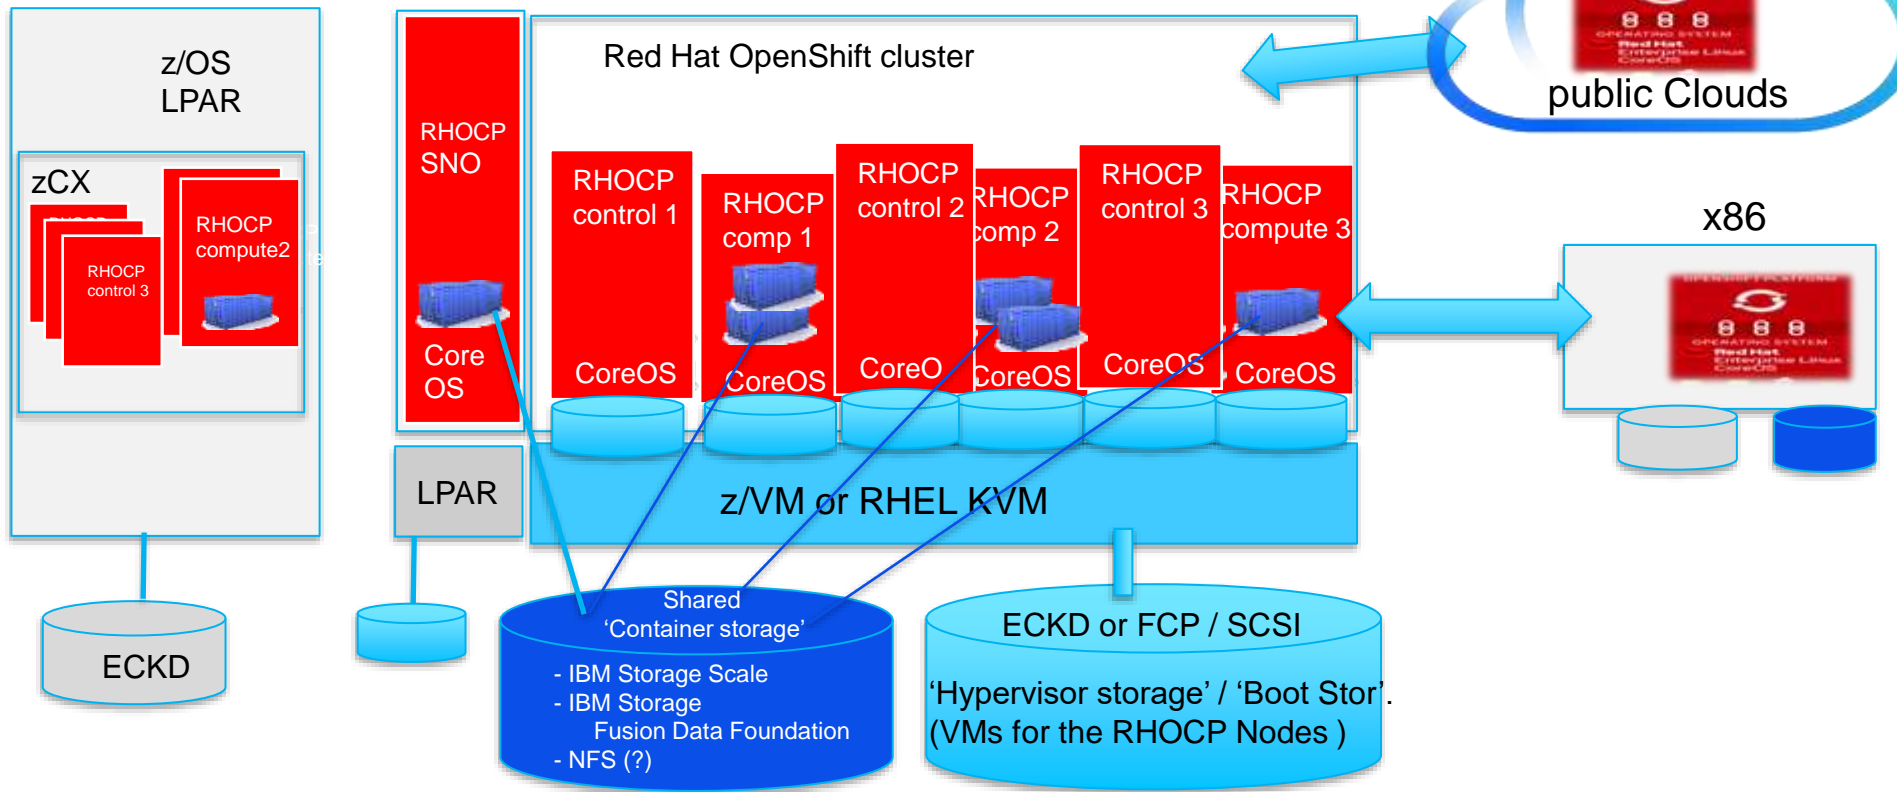
Wilhelm Mild
IBM Executive IT Architect
IBM R & D Lab Germany
wilhelm.mild@de.ibm.com



Red Hat OpenShift - the only Container Platform across these different hardware architectures and clouds



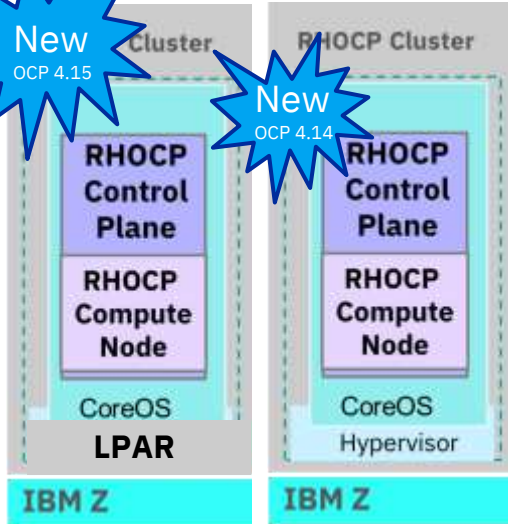
Red Hat OpenShift deployment options




RH OpenShift Single Node (SNO) and HA implementations

Single Node Cluster

New
OCP 4.15

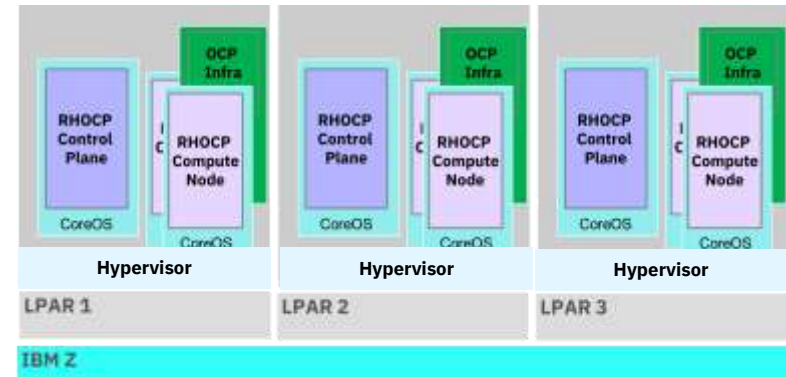


3 Node Cluster



All deployments can be in LPARs
without hypervisor

Multi Node Cluster



Value: Lower resource foot-print due to smaller cluster size

- 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.
- **No high availability**
- **Requires 2 IFLs** (Integrated Facilities for Linux), with SMT-2 enabled **on installation time**
- **Operation** requires **0.5 IFLs on average**; remaining **1.5 IFLs** required for **keeping SNO OCP stable** and for running some workloads

Multi Architecture options with Red Hat OpenShift

What does Multi Architecture mean

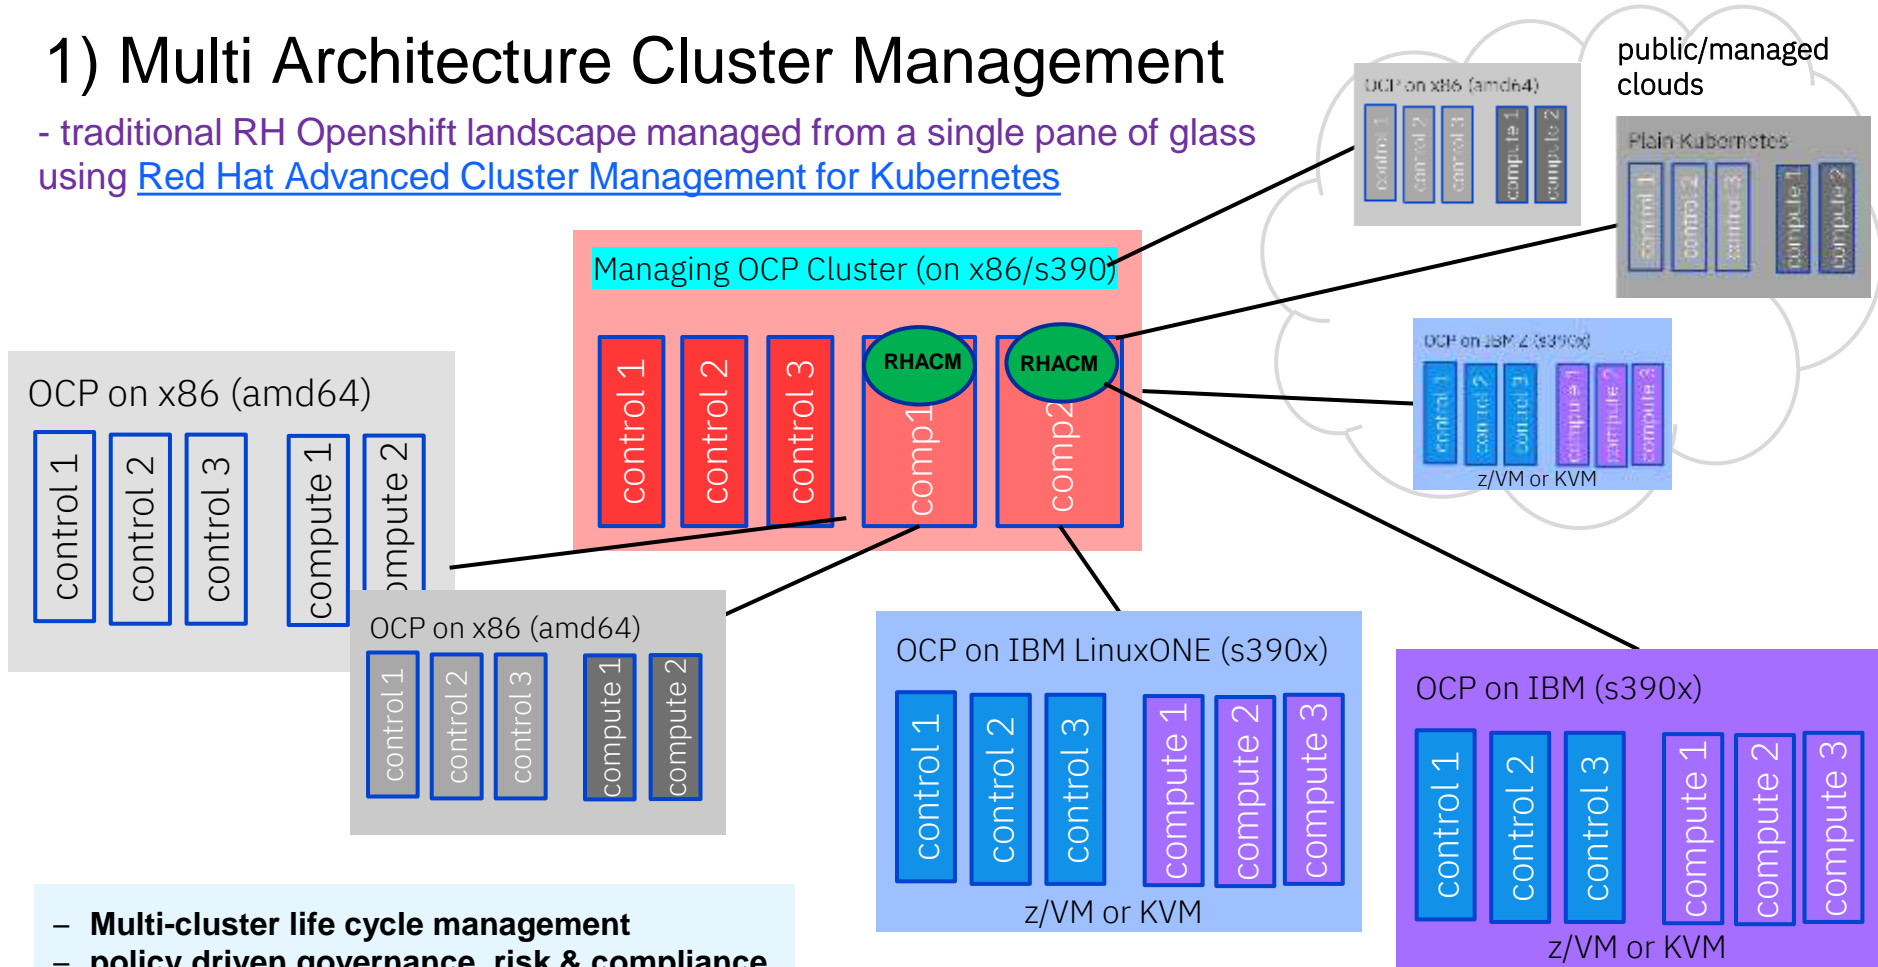
- Solutions spanned across different Cloud environments
- RH Openshift Clusters across different HW Architectures
- Applications can run simultaneous in:
 - different HW Architectures (e.g. x86, IBM Z) and
 - different Platform environments (e.g. LinuxONE, zCX)
- Consistent Management of RH OpenShift environments on different HW Architectures

Multi-Architecture Options for Red Hat OpenShift Hybrid Cloud environments

- **(1) Multi-Architecture cluster management**
 - **(2) Multi-Architecture hosted clusters**
 - **(3) Multi-Architecture compute**
 - **(4) Multi Architecture Development**
-

1) Multi Architecture Cluster Management

- traditional RH Openshift landscape managed from a single pane of glass using [Red Hat Advanced Cluster Management for Kubernetes](#)



- **Multi-cluster life cycle management**
- **policy driven governance, risk & compliance**
- **Application distribution & life cycle**

Red Hat Advanced Cluster Management Overview (RHACM)



Multicluster lifecycle management



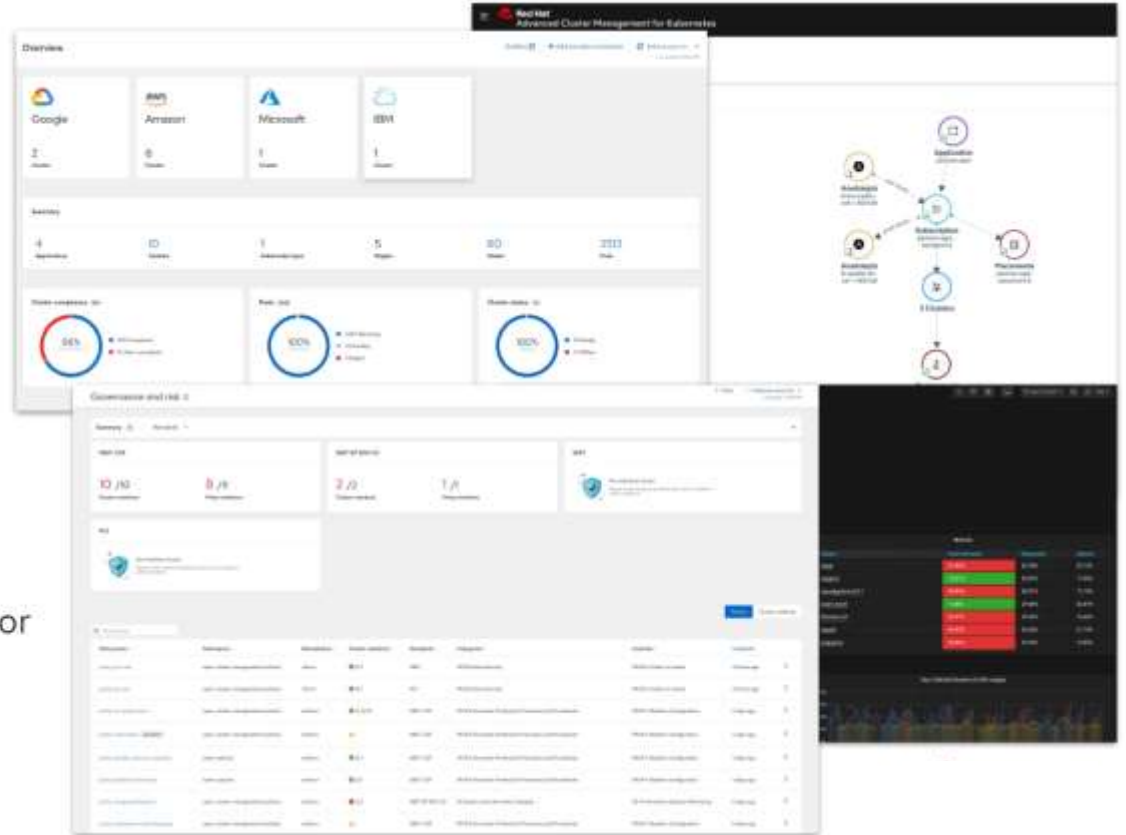
Policy driven governance, risk, and compliance



Advanced application lifecycle management

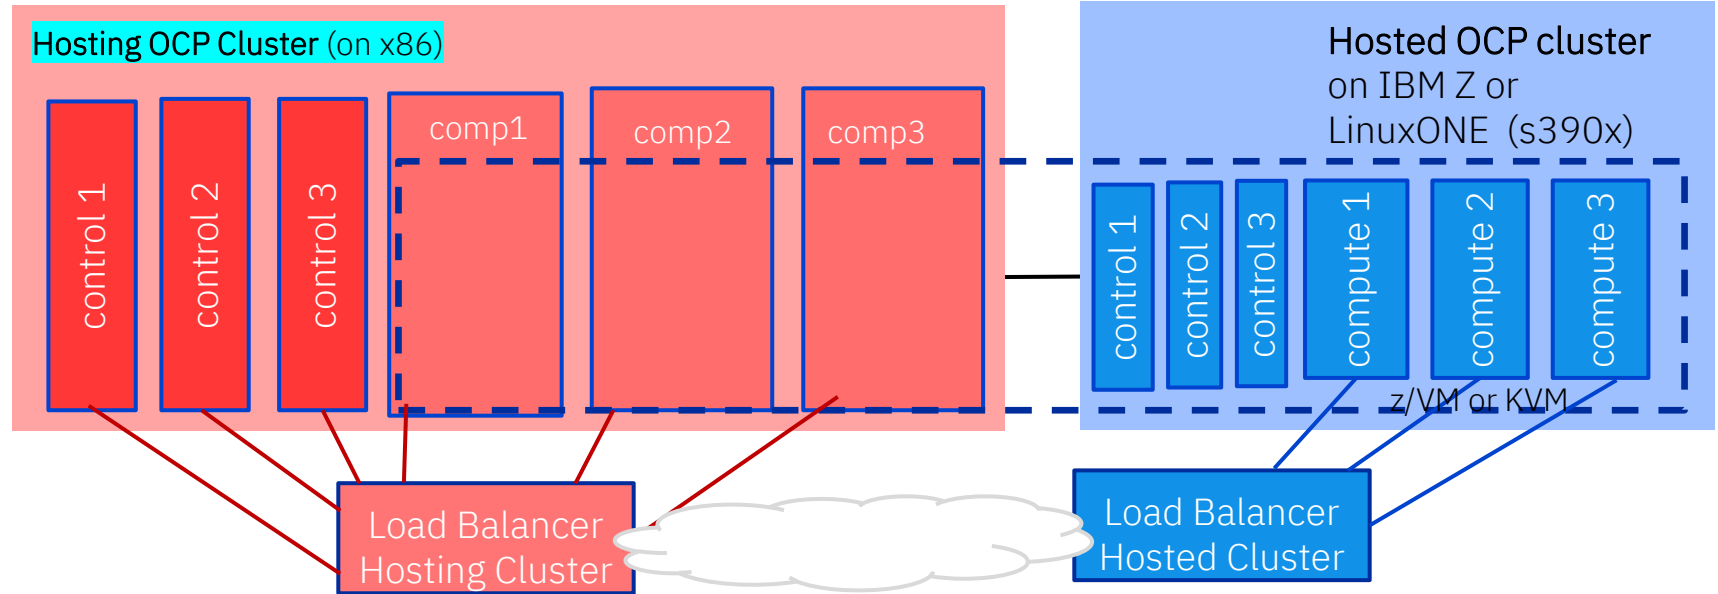


Multicluster observability for health and optimization



2) Multi Architecture hosted cluster management

- RH OpenShift clusters managed via [Hosted Control Planes*](#) former Hypershift



The Hosting Cluster – can host many Control Plane sets

- The Control Planes of Hosted Clusters are Compute Pods in the Hosting Cluster
- that saves resources and simplifies management and control
- isolation of Control and Compute leads to increased flexibility

Hosted Cluster consist of Control and Compute Nodes and is an independent entity

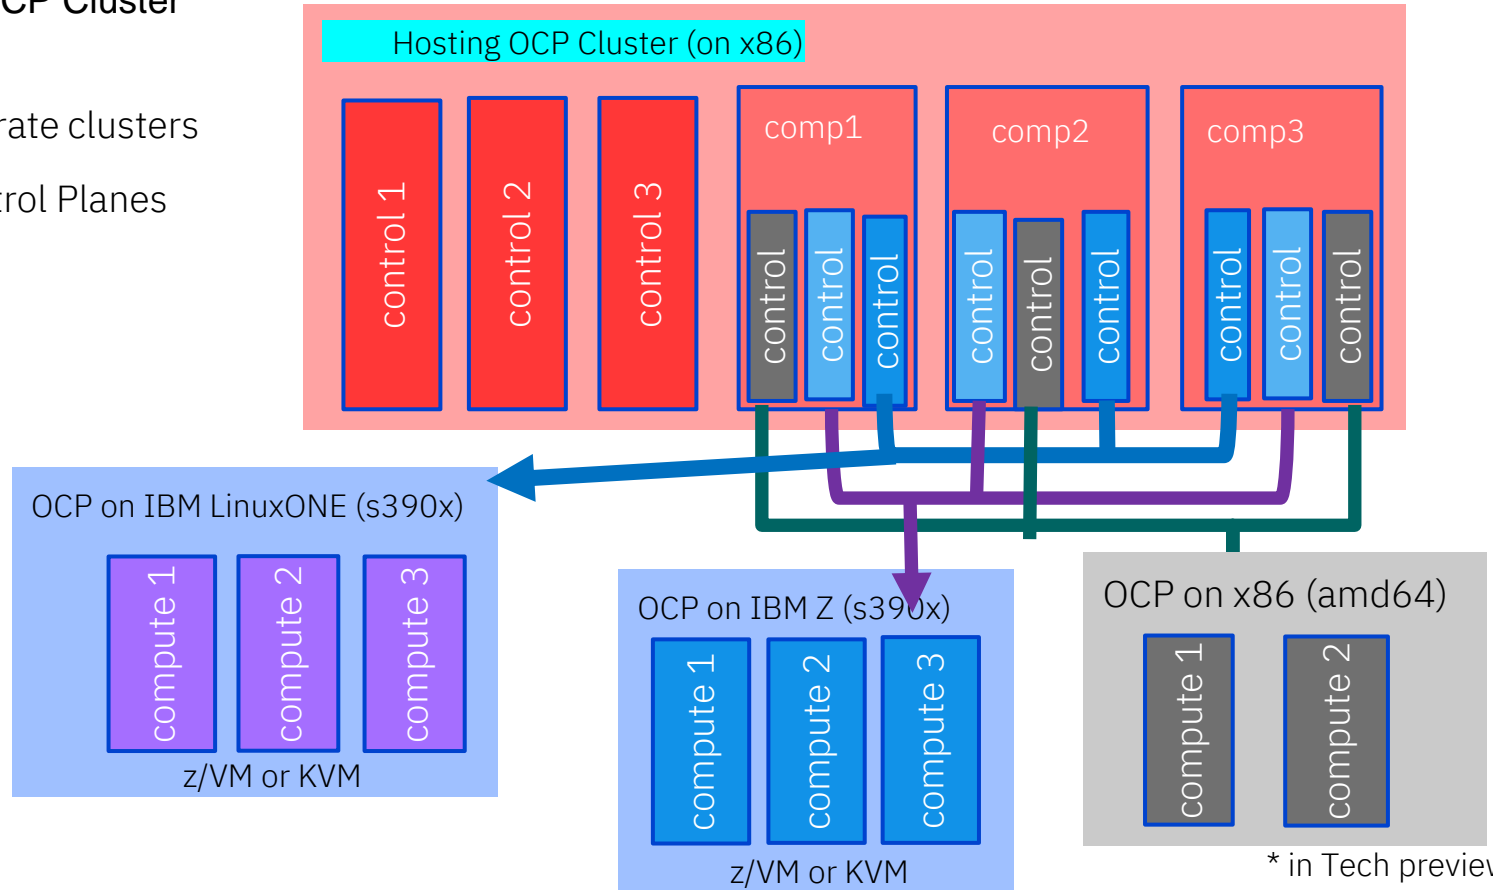
- can have its own rules, security, network setting for tenant isolation
- Hosted Clusters have to be build specifically for being hosted

* in Tech preview

Multi Architecture Hosted Cluster Management

- RH OpenShift environments managed via [Hosted Control Planes*](#) from a single RH OCP Cluster

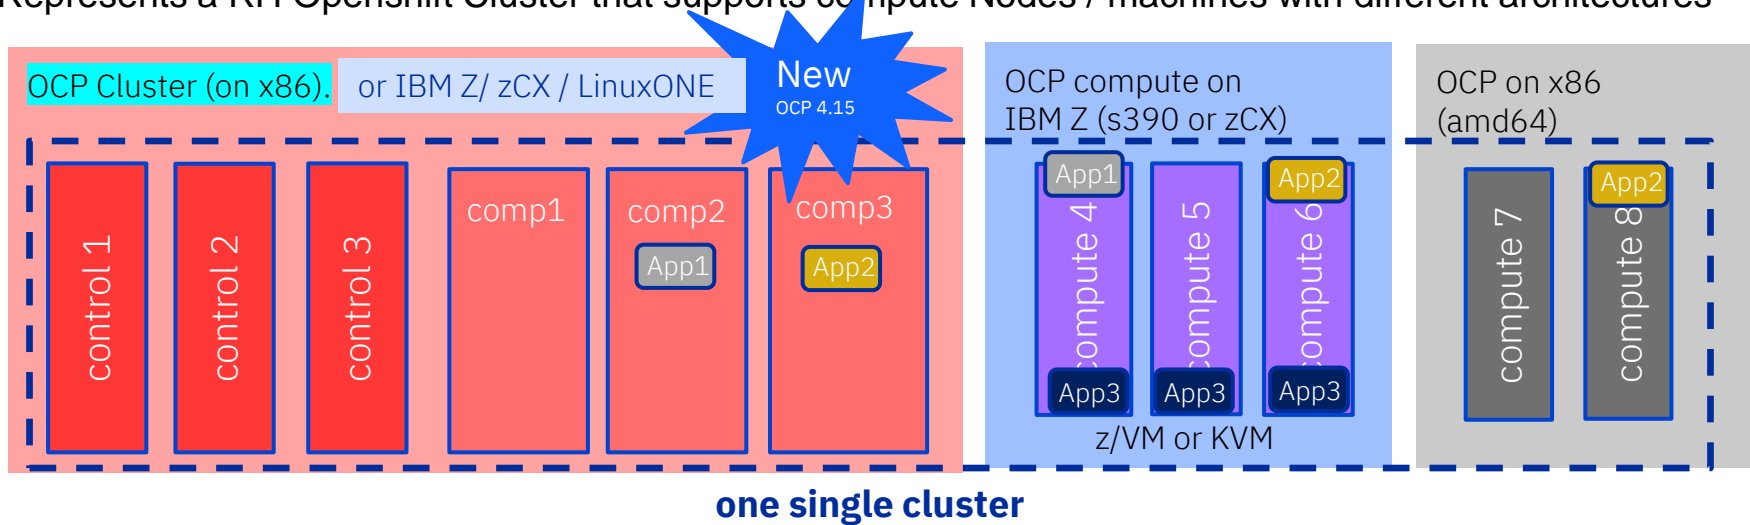
- there are 3 separate clusters
- with hosted Control Planes



* in Tech preview

3) Multi Architecture compute – Generally available

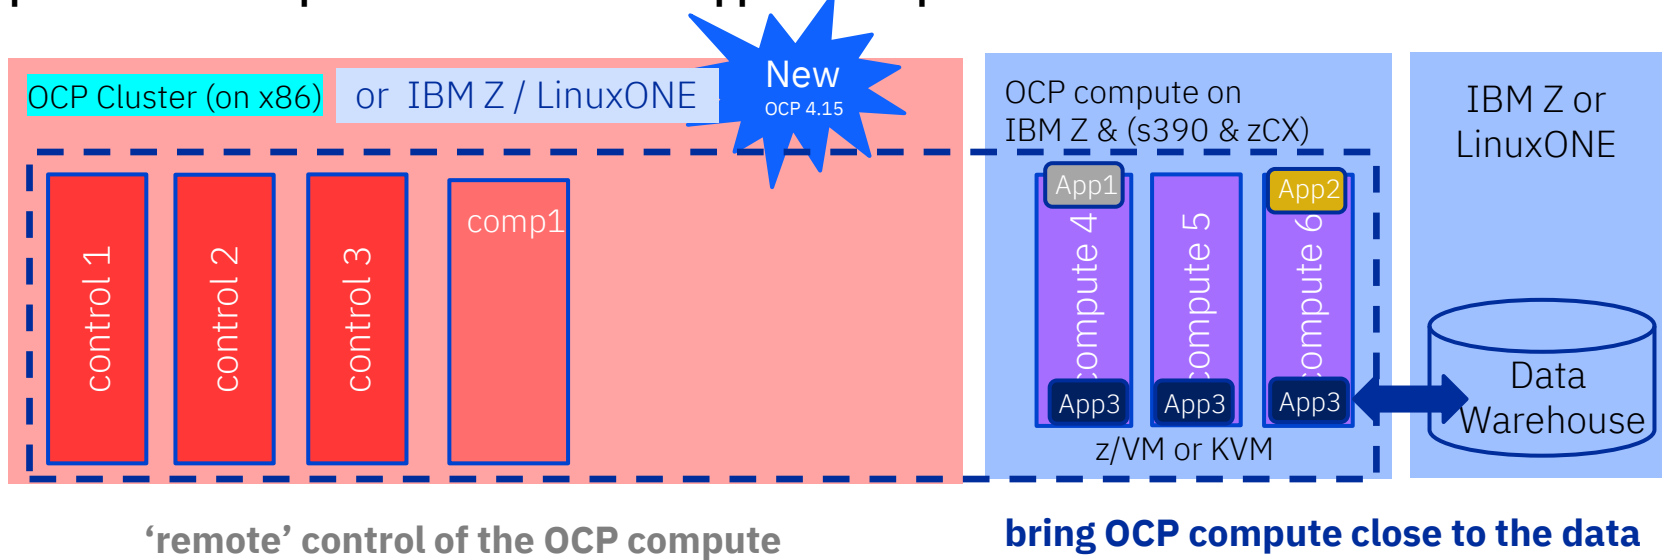
- Represents a RH OpenShift Cluster that supports compute Nodes / machines with different architectures



- **To build compute nodes on other architectures next to a OCP Cluster is a day 2 operation**
- **A RH OpenShift cluster with multi architecture compute nodes has requirements to Nodes & Pods**
 - architecture labels to nodes and labels for nodeSelector/nodeAffinity
 - <https://sdk.operatorframework.io/docs/advanced-topics/multi-arch/#supporting-clusters-with-multi-architecture-compute-nodes>
- **For existing RH OpenShift clusters – there is a path to migrate into a Multi-architecture cluster:**
 - https://docs.openshift.com/container-platform/4.14/updating/updating_a_cluster/migrating-to-multi-payload.html

3) Multi Architecture compute

Represents a RH OpenShift Cluster that supports compute Nodes / machines with different architectures



A RH OpenShift cluster with multi architecture compute nodes has requirements to Nodes and Pods

- **architecture labels** to nodes and labels for **nodeSelector/nodeAffinity**

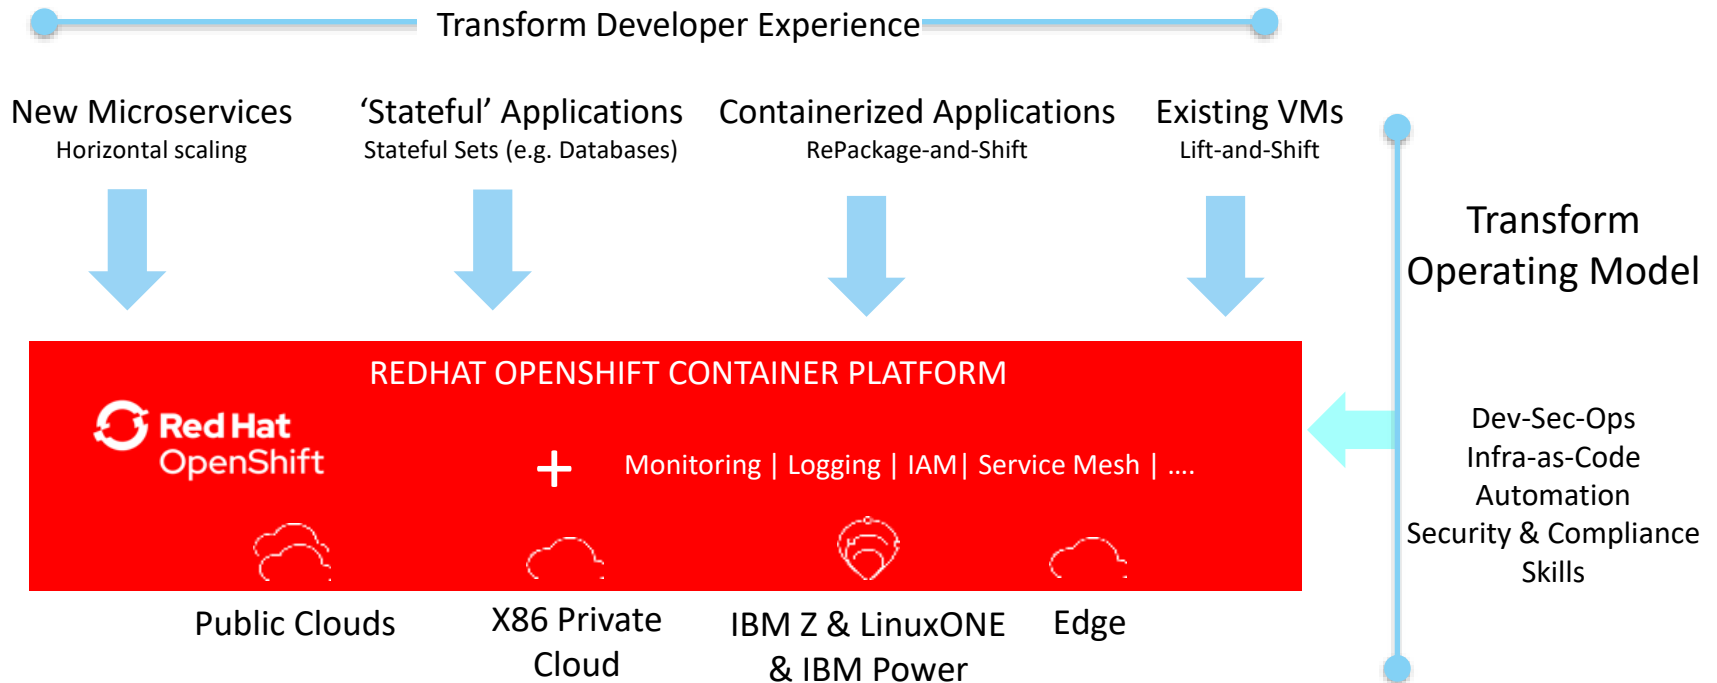
<https://sdk.operatorframework.io/docs/advanced-topics/multi-arch/#supporting-clusters-with-multi-architecture-compute-nodes>

For existing RH OpenShift clusters – there is a path to migrate into a Multi-architecture cluster:

<https://docs.openshift.com/container-platform/4.14/updating/migrating-to-multi-payload.html>


4) Multi Architecture Development – develop once, deploy multiple

Red Hat OpenShift supports a wide range of application types with a consistent developer & ops experience

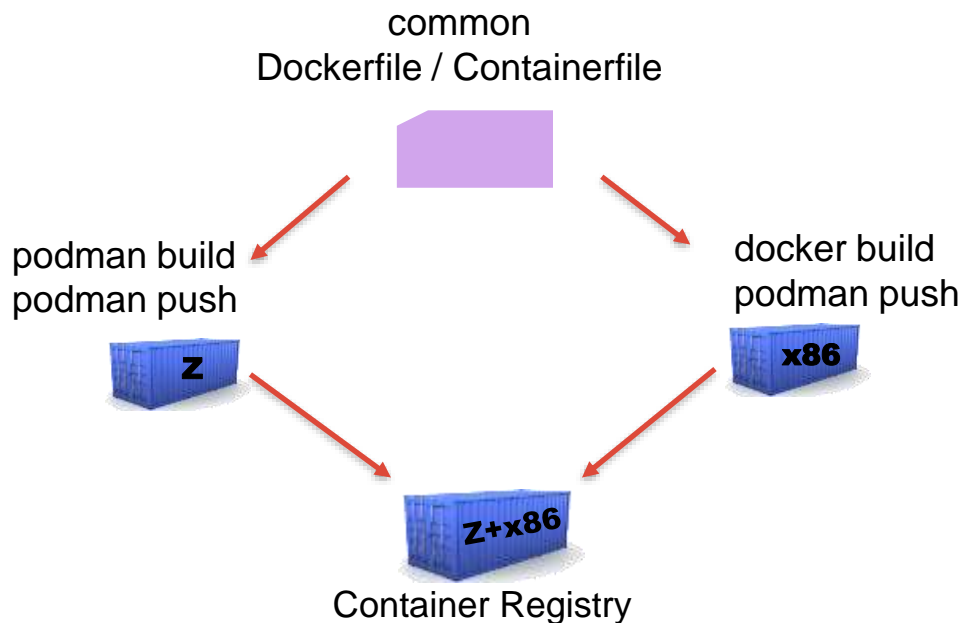


Portability of Container & Multi Architecture registry support

- **Container user experience (CLI, REST API) is identical across platforms**
- **Container images are not portable**, the source code or a s390x binary must be build and available
- Containers are often created with Dockerfiles (build descriptions) containing:
 - Specification of the base image
 - If the same distribution is available on s390x, usually simple
 - Additional steps to modify the image are often platform independent
 - Add packages (needs to match the base image)
 - Download files, Perform build
- Same Dockerfile can be used for multi-platform builds
- Multi-Arch Registry is required and widely available (e.g. Quay, Artifactory,..)
- [Multi-arch tool support available using external tools](#) (i.e. manifest tool)



Create Multi-Architecture Images – use the ‘manifest’ tool



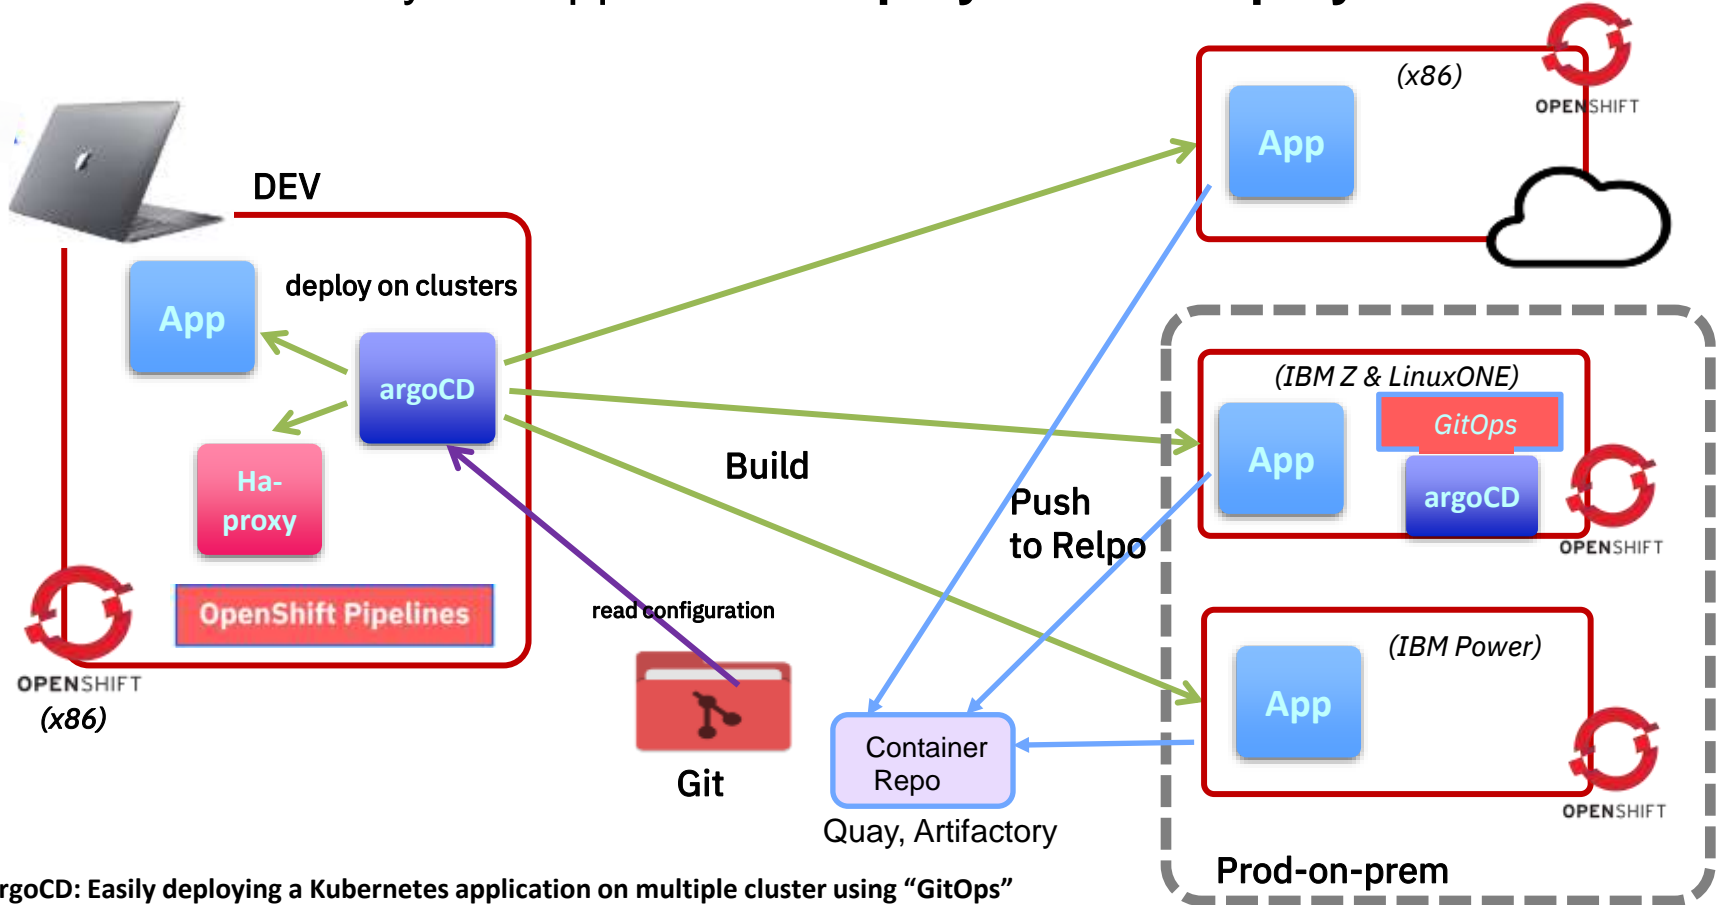
```
image: webapp:latest  
manifests:
```

```
-  
image: webapp-s390x  
platform:  
architecture: s390x  
os: linux
```

```
-  
image: webapp-amd64  
platform:  
architecture: amd64
```

- Many container images on external registries are multi-arch
 - lots of images as s390x versions available
- IBM built trusted image registry: <https://ibm.biz/BdfAdW>

4) Cross Platform Hybrid Application Deployment & Deployment Consistency



ArgoCD: Easily deploying a Kubernetes application on multiple cluster using "GitOps"

GitOps: Deploying and maintaining a Kubernetes application configuration from a Git repository

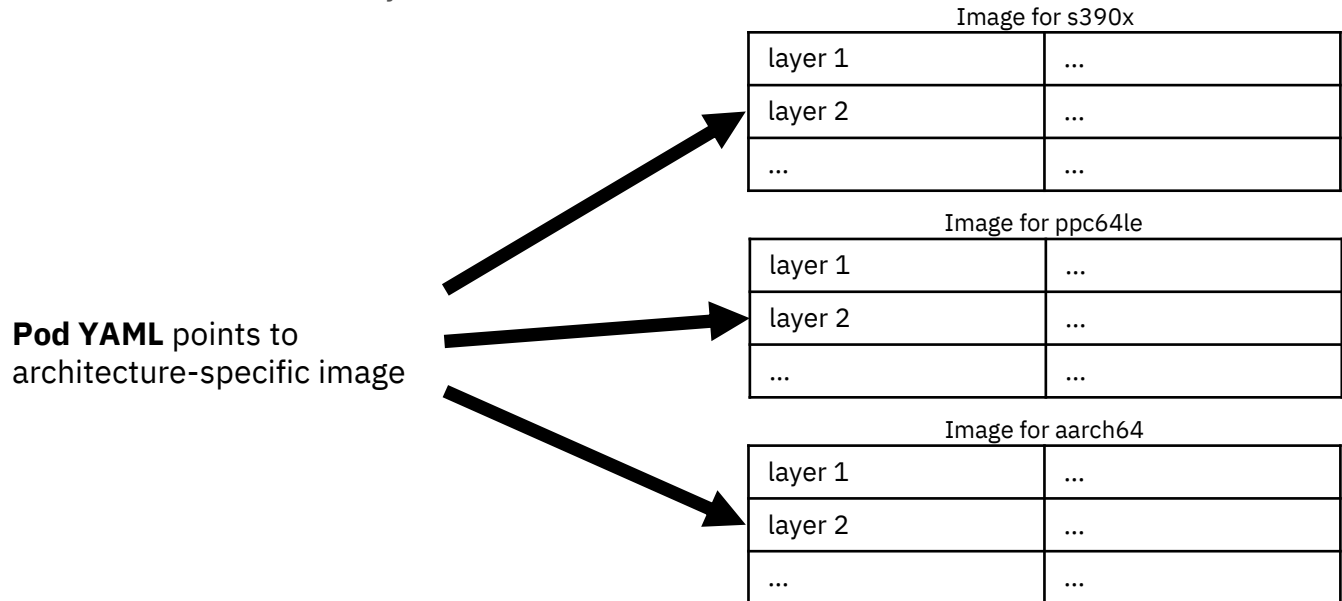
https://mediacenter.ibm.com/media/Multicenter-MultiArch-deployment-with-Red-Hat-OpenShift/1_9qjdhc96



Multi-Architecture Compute

Technical features that enable Multi-Architecture Compute:

- Manifest lists (“OCI image index”)^[1]: can point to specific images for architectures
- Kubernetes adds architecture labels to nodes
- Can use labels for nodeSelector/nodeAffinity

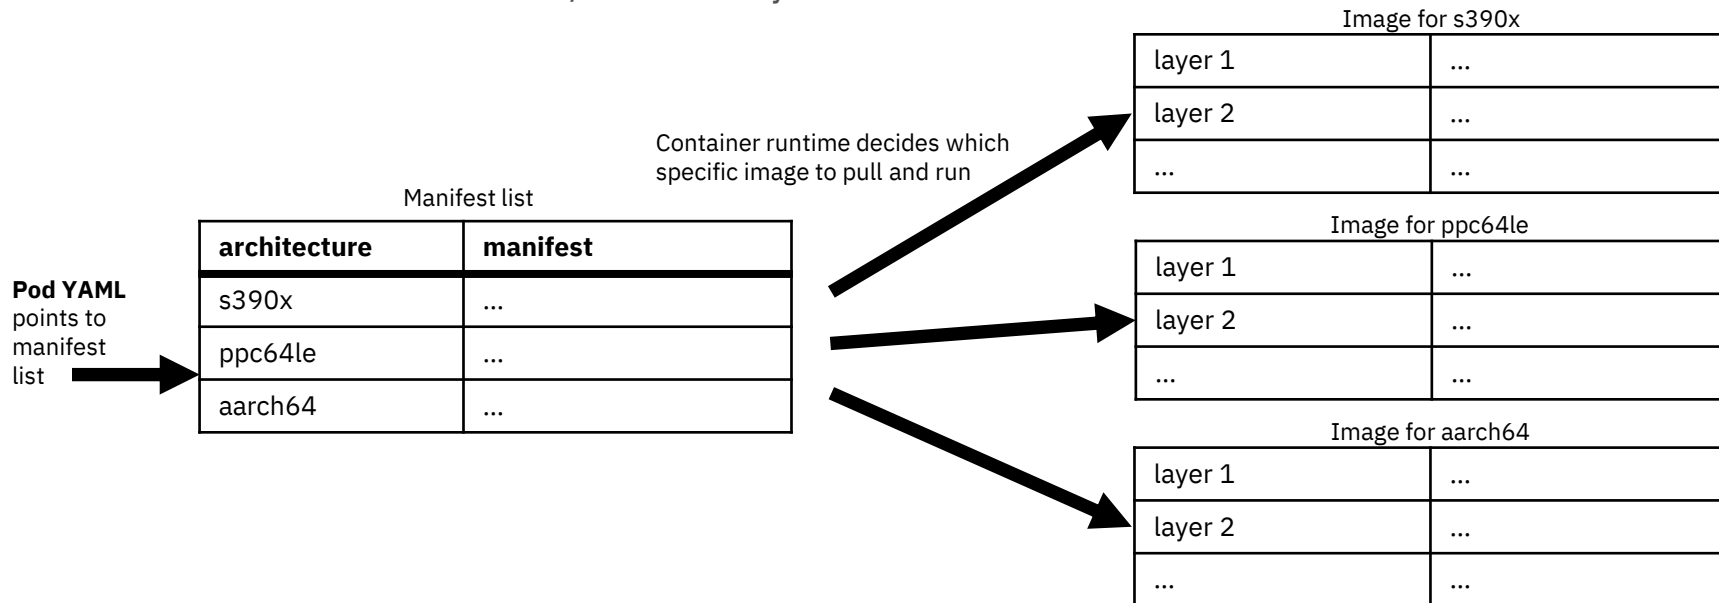


[1] <https://github.com/opencontainers/image-spec/blob/main/image-index.md>

Multi-Architecture Compute

Technical features that enable Multi-Architecture Compute:

- Manifest lists (“OCI image index”)^[1]: can point to specific images for architectures
- Kubernetes adds architecture labels to nodes
- Can use labels for nodeSelector/nodeAffinity



[1] <https://github.com/opencontainers/image-spec/blob/main/image-index.md>

Multi-Architecture Compute

Technical features that enable Multi-Architecture Compute:

- Manifest lists (“OCI image index”)^[1]: can point to specific images for architectures
- [Kubernetes adds architecture labels to nodes](#)
- Can use labels for nodeSelector/nodeAffinity

```
# oc get node compute-3 -oyaml
apiVersion: v1
kind: Node
metadata:
[...]
```

```
  labels:
    beta.kubernetes.io/arch: s390x
    beta.kubernetes.io/os: linux
    kubernetes.io/arch: s390x
    kubernetes.io/hostname: compute-3
    kubernetes.io/os: linux
```

```
[...]
```

[1] <https://github.com/opencontainers/image-spec/blob/main/image-index.md>



Multi-Architecture Compute

Technical features that enable Multi-Architecture Compute:

- Manifest lists (“OCI image index”)^[1]: can point to specific images for architectures
- Kubernetes adds architecture labels to nodes
- Can use labels for `nodeSelector/nodeAffinity`

```
# cat only-on-s390x-nodes.yaml
apiVersion: v1
kind: Pod
metadata:
  name: deploy-this-pod-only-on-s390x-nodes
spec:
[...]
```

```
nodeSelector:
  kubernetes.io/arch: s390x
```

```
# cat on-s390x-or-amd64-nodes.yaml
apiVersion: v1
kind: Pod
metadata:
  name: deploy-this-pod-on-amd64-or-s390x-nodes
spec:
[...]
```

```
affinity:
  nodeAffinity:
    requiredDuringSchedulingIgnoredDuringExecution:
      nodeSelectorTerms:
        - matchExpressions:
            - key: kubernetes.io/arch
              operator: In
              values:
                - amd64
                - s390x
```

[1] <https://github.com/opencontainers/image-spec/blob/main/image-index.md>



Multi-Arch build pipeline example using OpenShift Pipelines

Project: multiarch-build-dev ▾

PipelineRuns > tails

PLR **currencyservice-181c88650f4** Running Actions ▾

[Details](#) [YAML](#) [TaskRuns](#) [Logs](#) [Events](#)

PipelineRun details



Name
currencyservice-181c88650f4

Namespace
NS multiarch-build-dev

Status
Running

Pipeline
PL currencyservice

Reference:

<https://github.com/IBM/ibm-garage-tekton-tasks/tree/main/pipelines>
<https://github.com/IBM/ibm-garage-tekton-tasks/tree/main/tasks>

Multi-Architecture Compute – Setup

- OpenShift docs lists different possible types of setups and compute node combinations in detail [1]
- Steps for adding compute nodes [2]
- Verifying cluster compatibility, possibly update payload
- Adapt environment
- Disable UDP aggregation
- Get ignition file
- Get images
- Boot nodes



[1] https://docs.openshift.com/container-platform/4.15/post_installation_configuration/configuring-multi-arch-compute-machines/multi-arch-compute-machines.html

[2] https://docs.openshift.com/container-platform/4.15/post_installation_configuration/configuring-multi-arch-compute-machines/creating-ibm-z-kvm.html

Operator SDK:

<https://sdk.operatorframework.io/docs/advanced-topics/multi-arch/#supporting-clusters-with-multi-architecture-compute-nodes>

Building an Operator for Multiple Architectures

Kubebuilder explains how you can use docker buildx to build multi-architecture images. Operator SDK leverages KubeBuilder to ensure that

Manifest lists

The most straightforward way of building operators is to leverage manifest lists, specified by `ImageManifestList`. A manifest list points to specific image manifests for

For convenience tools like `buildah` allow to cross-build images for multiple architectures on a single host. For instance with `buildah`:

```
for a in amd64 arm64 ppc64le s390x; do \
  buildah bud --manifest registry/username/manifest-list:$a
done
```

This creates the manifest list, builds each image, and pushes them to the registry.

The result can then be pushed to the desired registry.



Home Build Documentation Releases

Documentation

Overview

Installation

Building Operators

Testing Operators

Upgrade SDK

Version

Commands

OLM Integration

Advanced Topics

Custom Bundle

Validation

Multiple

Architectures

Multiple Service

Accounts

Best Practices

Contribution Guide

FAQ

Supporting Clusters with Multi-Architecture Compute Nodes

The [Fundamentals](#) above aim to guide authors on the key steps to building and distributing operators that can run on multiple architectures. These instructions are sufficient when your cluster's compute nodes share the same architecture. However, operator authors should also understand the implications of running their operators in a cluster with multi-architecture compute nodes since it is not always guaranteed that the architectures of the compute nodes will match the architectures supported by the operator.

Safe Scheduling Using Node Affinity

Node affinity is a mechanism exposed in a Kubernetes pod template that allows a `PodSpec` author to instruct the scheduler to restrict a pod to run only on (or with a preference for) nodes that meet specific criteria. To ensure that pods are always scheduled to nodes of compatible architecture, it is a best practice for authors to set node affinity requirements to ensure their operators and operands will only schedule to the nodes with architectures available to the images in the pod. If you don't do this, a container scheduled to an incompatible node will immediately crash with an `exec format error`, which will ultimately lead to an `ImagePullBackoff` event as the pod is restarted only to crash again with the same error.

Determining the Architectures Supported by an Image

For a given container image, you can check which architectures are supported by listing them by inspecting the manifest. Piping the output to the `python json.tool` module enables pretty-printed JSON output.

IBM Collaboration with cloud providers on a hybrid cloud approach

Hyperscaler Collaboration

IBM
Cloud

- [Announcement](#) (Feb '22)
- [z/OS Dev Test available on IBM Cloud](#) (Jun '22)
- [Hybrid Cloud Blog](#)

Amazon Web
Services
(AWS)

- [Announcement Blog](#)
- IBM Z and Cloud Modernization Stack on [AWS Marketplace](#)

Microsoft
Azure

[Announcement Blog](#)

Google Cloud
Platform
(GCP)

[Secure data pathing](#)

Hybrid Application Modernization Patterns



Secure applications and data in the cloud



Hybrid storage with cloud storage



Cloud-native developer experience for z/OS applications



Simplify access from digital channels with APIs



Enterprise automation across z/OS and Cloud



Share near real-time information between z/OS apps and Cloud

IBM zSystems and LinuxONE Container Registry



foundational distros*



languages




registries, scanners



dashboards, UIs



shell, utilities



CI/CD infrastructure



messaging




* Red Hat base image available at Red Hat Marketplace

Many images, multiple versions, chosen based on customer input



databases/datastores



web serving



Elastic stack

Top 10 Containers

- Kafka
- Development-workspace
- Elasticsearch
- Logstash
- Ubuntu
- Grafana
- OpenJDK
- Alpine
- Portainer
- Tensorflow

Images in the pipeline



Official interface made available end of October, 2021

<https://ibm.biz/BdfAdW>

Red Hat catalog with container images

<https://catalog.redhat.com/software/containers/search?p=1&architecture=s390x>



Red Hat Ecosystem Catalog

Containers Search Ecosystem Catalog

Contact us Help Resources All Red Hat

[Home](#) > [Containers](#) > [Browse Containers](#)

Certified s390x container images

Container images offer lightweight and self-contained software to enable deployment at scale.

Filters

Find specific filters

Architecture **s390x** Clear filters

Provider

- CA Technologies
- EnterpriseDB
- Fujitsu Limited
- IBM Corporation
- IBM Japan, Ltd.
- Instana
- LLC BPC Banking Technologies
- NGINX, Inc.

[See more](#)

Category

1,201 Results found

Sort by: Relevance



rhel7
Red Hat Enterprise Linux 7

By Red Hat

Provides the latest release of Red Hat Enterprise Linux 7 in a fully featured and supported base image.



redhat-openjdk-18/openj...
Java Applications

By Red Hat

OpenJDK 8 image for Java Applications



rhel7/etcd
Etcd

By Red Hat

A highly-available key value store for shared configuration



rhel7-atomic
RHEL Atomic Base Image

By Red Hat

Provides the latest release of Red Hat Enterprise Linux 7 in a fully supported base image which is optimized for minimal size.

Documentation

OpenShift:

- https://64355--docspreview.netlify.app/openshift-enterprise/latest/operators/operator_sdk/osdk-multi-arch-support

Operator Lifecycle Management (OLM):

- <https://olm.operatorframework.io/docs/advanced-tasks/ship-operator-supporting-multiarch/>

Multi-Arch Manager Operator:

- <https://github.com/openshift/enhancements/blob/master/enhancements/multi-arch/multiarch-manager-operator.md>

Useful links for Linux and RH Openshift on IBM Z & LinuxONE

- **Technical Linux on zSystems customer webinars:** <http://ibm.biz/LinuxonZandLinuxONEwebcasts>
- **[IBM Knowledge Center](#) for Linux on Z and LinuxONE**
- [Blog: Linux and Mainframe](#)
- [News and tips for running Linux on IBM Z and LinuxONE](#)
- [OpenShift on IBM Z](#)
- **The Reference Architecture:** <https://lnkd.in/dpdpz8V>
Redbook: <https://www.redbooks.ibm.com/redbooks/pdfs/sg248515.pdf>
- The Reference Architecture for [IBM Spectrum Scale Container Native Storage Access \(CNSA\)](#)
- The Reference Architecture for [IBM Storage Fusion Data Foundation \(former OpenShift Data Foundation\)](#)
- Blog: <https://www.openshift.com/blog/installing-ocp-in-a-mainframe-z-series>
- **[Virtualization on IBM Z & LinuxONE](#)**
- [z/VM resources](#)
- [KVM on Z](#) blog
- **[Containers on IBM Z](#)**
- **Trusted IBM Container Image Registry -Sign up now:** <https://ibm.biz/zregeap>
- IBM Z container blog: [Linux on Z and Containers](#)
- zCX, Containers in z/OS : <https://www.ibm.com/support/z-content-solutions/container-extensions/>

Questions?



Wilhelm Mild
IBM Executive IT Architect

*IBM Deutschland Research
 & Development GmbH
 Schönaicher Strasse 220
 71032 Böblingen, Germany*



*Office: +49 (0)7031-16-3796
 wilhelm.mild@de.ibm.com*

Twitter: @mildw



Notices and disclaimers

- © 2019 International Business Machines Corporation. No part of this document may be reproduced or transmitted in any form without written permission from IBM.
- **U.S. Government Users Restricted Rights — use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM.**
- Information in these presentations (including information relating to products that have not yet been announced by IBM) has been reviewed for accuracy as of the date of initial publication and could include unintentional technical or typographical errors. IBM shall have no responsibility to update this information. **This document is distributed “as is” without any warranty, either express or implied. In no event, shall IBM be liable for any damage arising from the use of this information, including but not limited to, loss of data, business interruption, loss of profit or loss of opportunity.** IBM products and services are warranted per the terms and conditions of the agreements under which they are provided.
- IBM products are manufactured from new parts or new and used parts.
In some cases, a product may not be new and may have been previously installed. Regardless, our warranty terms apply.”
- **Any statements regarding IBM's future direction, intent or product plans are subject to change or withdrawal without notice.**
- Performance data contained herein was generally obtained in a controlled, isolated environments. Customer examples are presented as illustrations of how those
- customers have used IBM products and the results they may have achieved. Actual performance, cost, savings or other results in other operating environments may vary.
- References in this document to IBM products, programs, or services does not imply that IBM intends to make such products, programs or services available in all countries in which IBM operates or does business.
- Workshops, sessions and associated materials may have been prepared by independent session speakers, and do not necessarily reflect the views of IBM. All materials and discussions are provided for informational purposes only, and are neither intended to, nor shall constitute legal or other guidance or advice to any individual participant or their specific situation.
- It is the customer’s responsibility to insure its own compliance with legal requirements and to obtain advice of competent legal counsel as to the identification and interpretation of any relevant laws and regulatory requirements that may affect the customer’s business and any actions the customer may need to take to comply with such laws. IBM does not provide legal advice or represent or warrant that its services or products will ensure that the customer follows any law.

Notices and disclaimers continued

- Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products about this publication and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products. IBM does not warrant the quality of any third-party products, or the ability of any such third-party products to interoperate with IBM's products. **IBM expressly disclaims all warranties, expressed or implied, including but not limited to, the implied warranties of merchantability and fitness for a purpose.**
- The provision of the information contained herein is not intended to, and does not, grant any right or license under any IBM patents, copyrights, trademarks or other intellectual property right.
- IBM, the IBM logo, ibm.com and [names of other referenced IBM products and services used in the presentation] are trademarks of International Business Machines Corporation, registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at: www.ibm.com/legal/copytrade.shtml