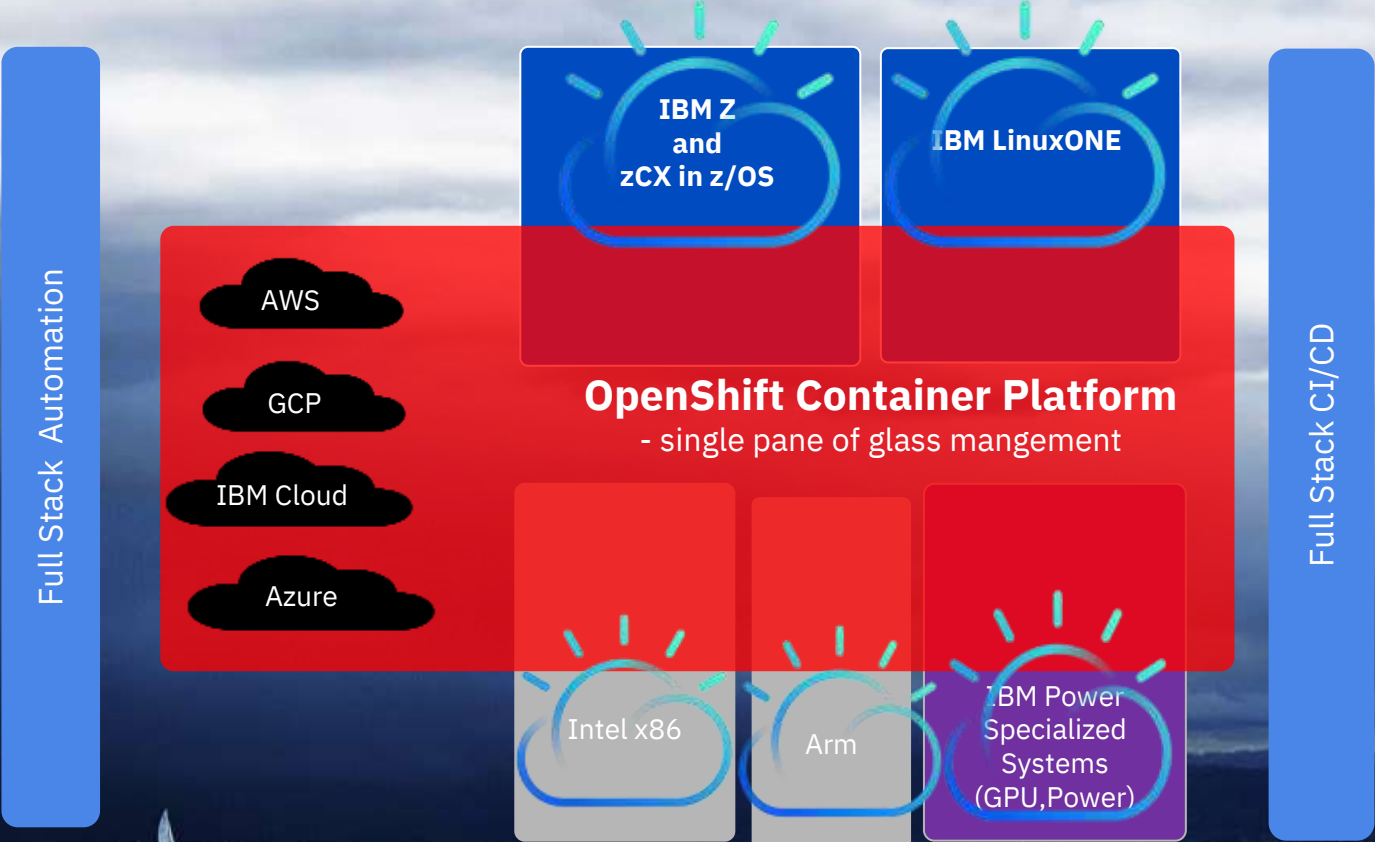


How to Build a Multi-Architecture solution with Red Hat OpenShift on IBM Z, zCX or LinuxONE

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



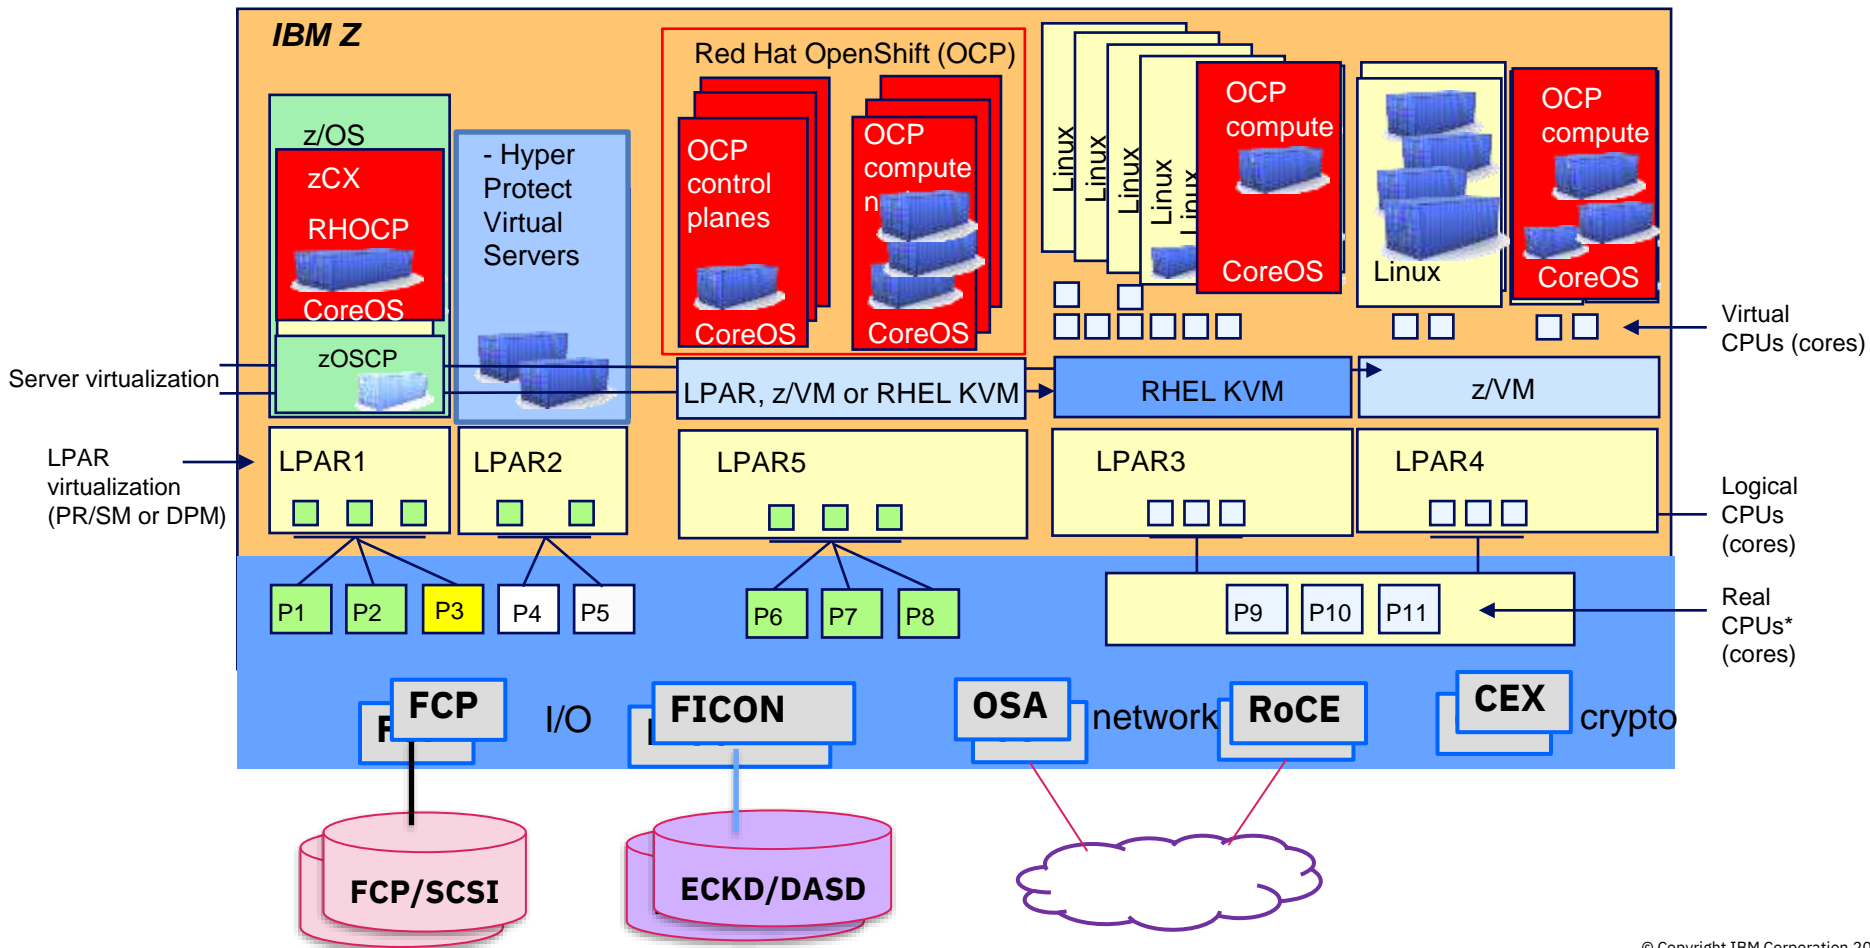
Red Hat OpenShift - manage the different clouds from a single pane of glass



1. Plan your solution regarding Multi Architecture and Multi-cloud

- ❑ Solution spanned across different Cloud environments**
- ❑ Solution spanned across different HW Architectures**
- ❑ The data requirements in the Multi-cloud solution topology**
- ❑ Solution managed from single cluster or from a single pane of glass for all clouds on different HW Architectures**

IBM Z Virtualization and Container options



2. Build the solution Architecture for the cloud services

1. Multi-Architecture compute (MAC) topology

- Cloud Services which run simultaneous in different HW Architectures and environments
- Cloud services which are extended to another HW Architecture
- Cloud services which are restricted to one HW Architecture

2. Hosted Control Planes topology

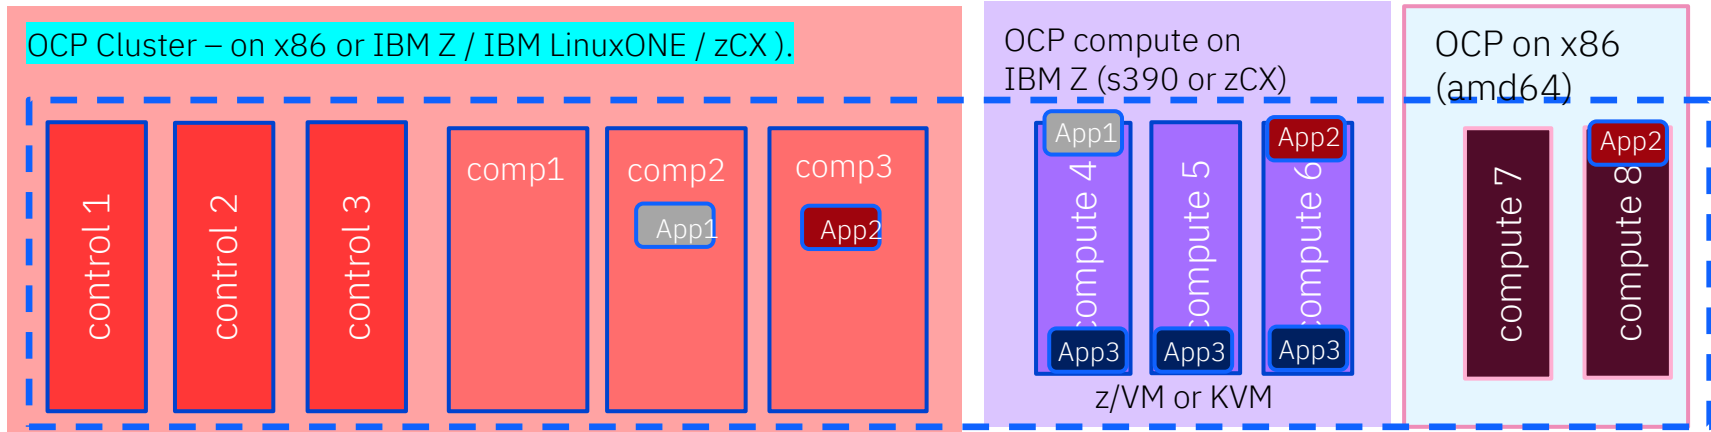
- Manage Multi-Architecture clusters from one central OpenShift cluster

Consider the Operational aspects

- Who will manage and operate the infrastructure and Container Applications

1) Multi Architecture compute (MAC)

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

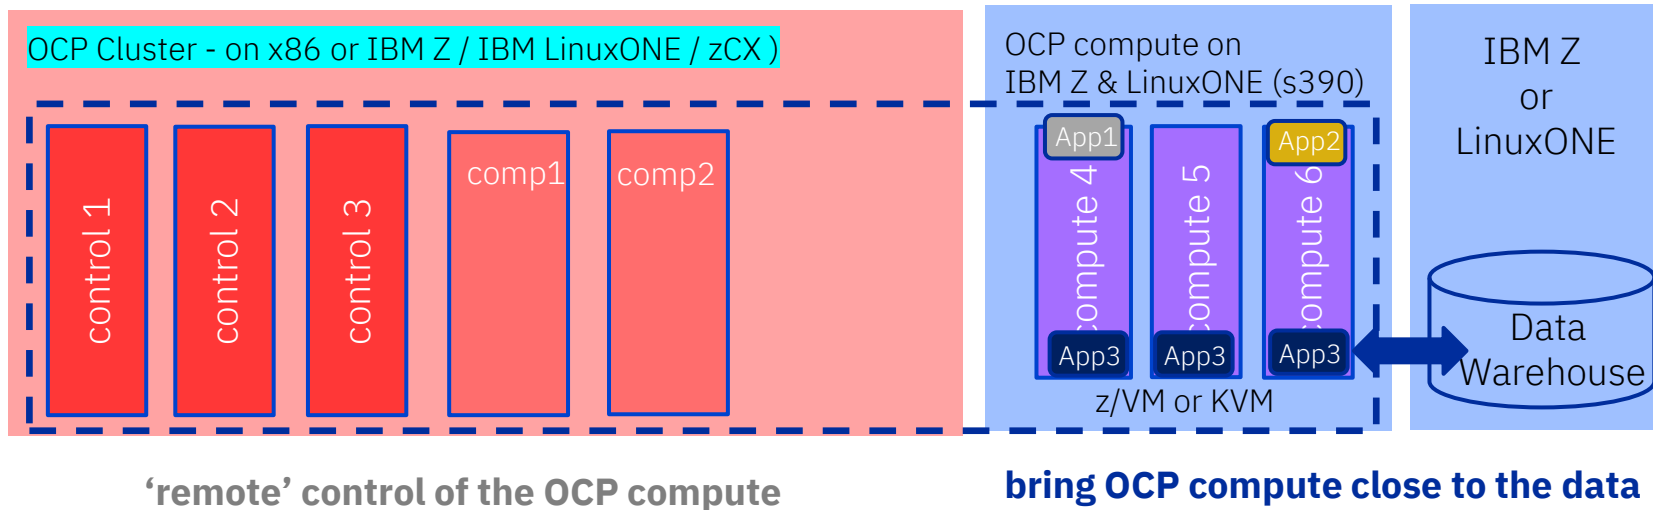


one single cluster

- **Note: A base cluster on x86 is required to be installed with ‘platform: NONE’**
- **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/4.17/updating/updating_a_cluster/migrating-to-multi-payload.html

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>

IBM zCX Foundation for Red Hat OpenShift environment

IBM pre-packaged OpenShift Environment

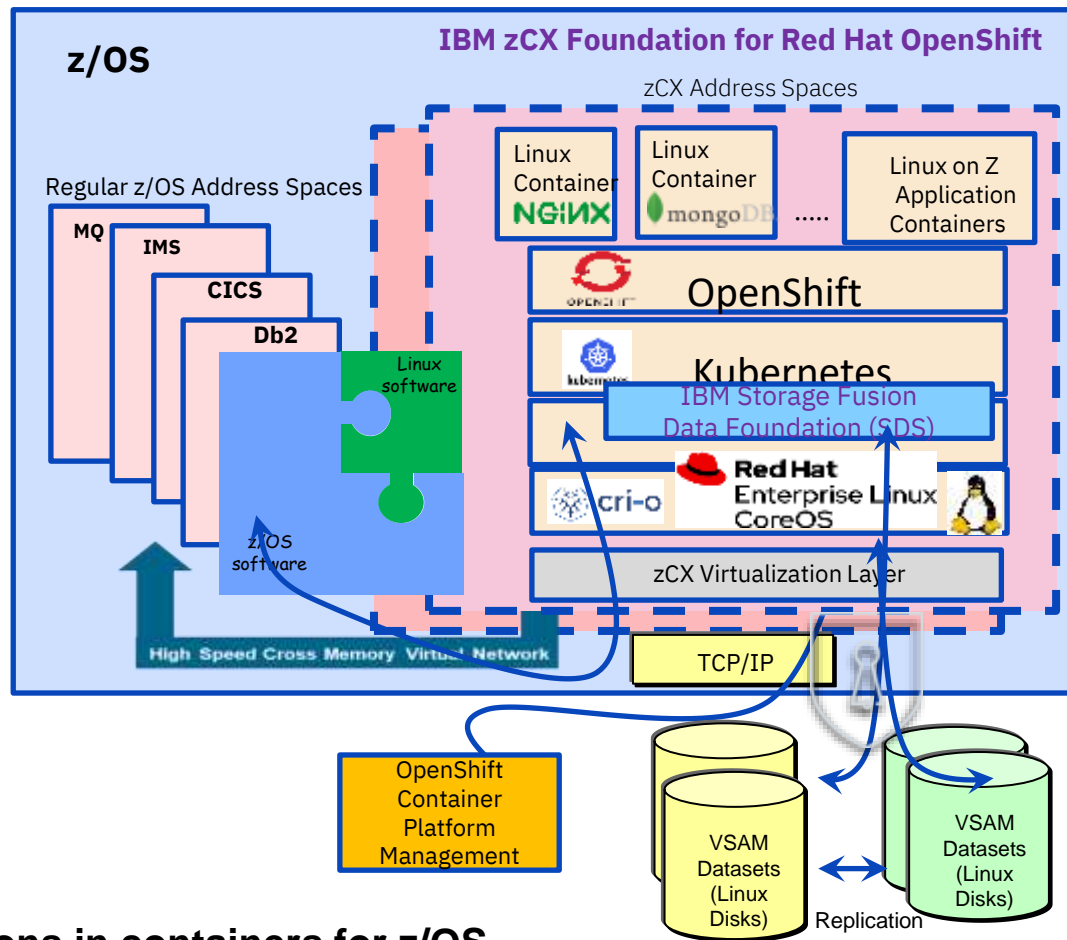
- Includes full stack OpenShift (CoreOS + K8S + Openshift components)
- Supported directly by IBM
- Can include clustering and registry capabilities
- Competitive price/performance (Exploits zIIPs)

Application developers

- can deploy software using OpenShift interface
- Any software available as a Container image (s390x) - continuously growing ecosystem
- Any home-grown Linux on zSystems container images
- Using open standard interfaces

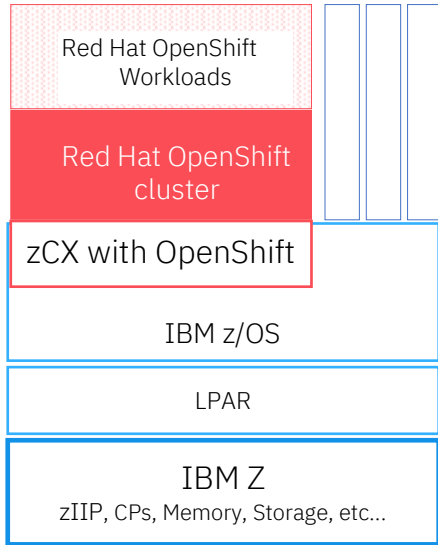
Access to underlying environment via RHOCP APIs

- No root access
- Administrative tasks via RHOCP and z/OS
- Secure virtual network – SAMEHOST
- Requires packaging of software as Container images
- Same code for RH OpenShift in IBM zSystems
- WLM Mgmt – like any other z/OS Address Space



Provides various solution options in containers for z/OS

Point of Value for OpenShift on zCX



Enable cloud native applications in z/OS

Containerize existing or new z/OS applications to use cloud services with OpenShift in zCX that were previously unavailable.



Design solutions with z/OS

Include z/OS in the design of new Multi-Architecture solutions that exploit a containerized Linux software deployment and orchestrate it with Red Hat OpenShift.



Expand your hybrid cloud strategy

Leverage Kubernetes orchestration on z/OS using zCX. Enable z/OS system programmers to integrate with zCX into Kubernetes clusters managed by Red Hat OpenShift and expand your hybrid cloud strategy.

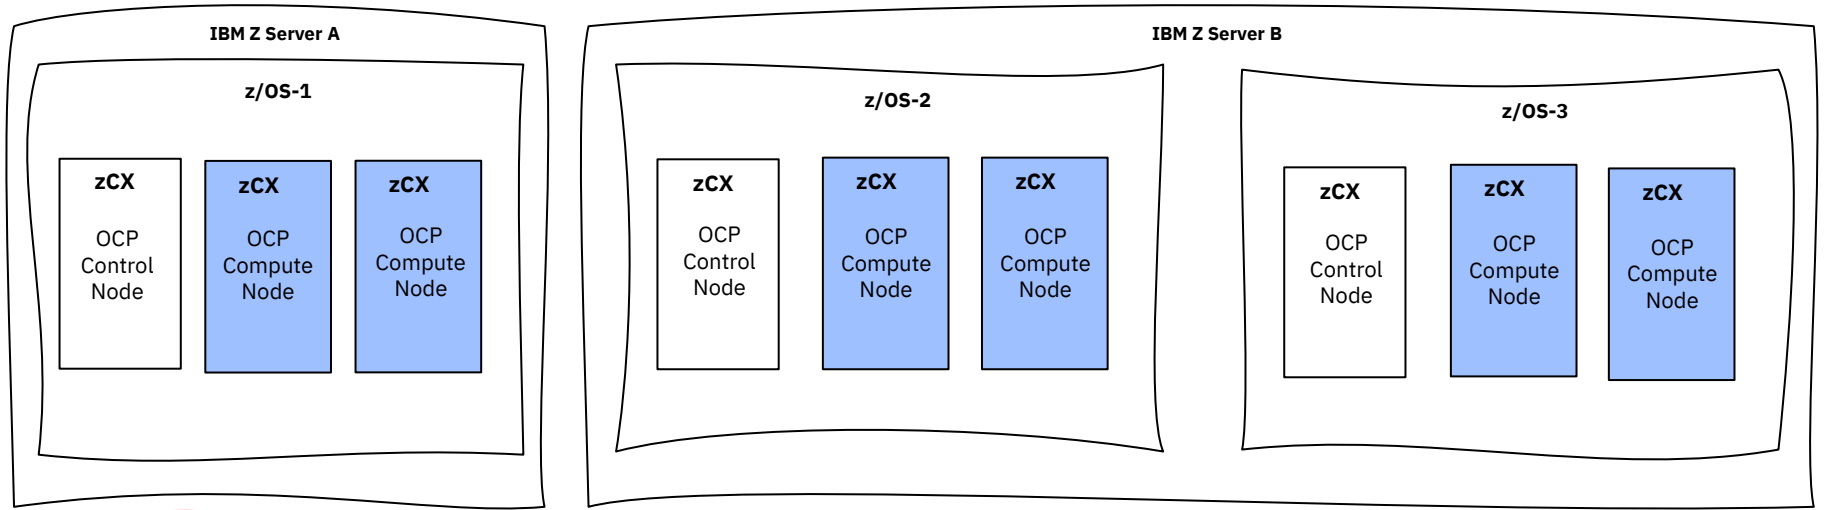


No specialized z/OS skills or expertise required

Develop and deploy containerized software in support of and inside z/OS using standard Red Hat OpenShift interfaces, processes, and tooling, with Cloud native skills.

IBM zCX Foundation for OpenShift


Red Hat OpenShift Cluster deployed in an HA configuration on z/OS zCX Address Spaces



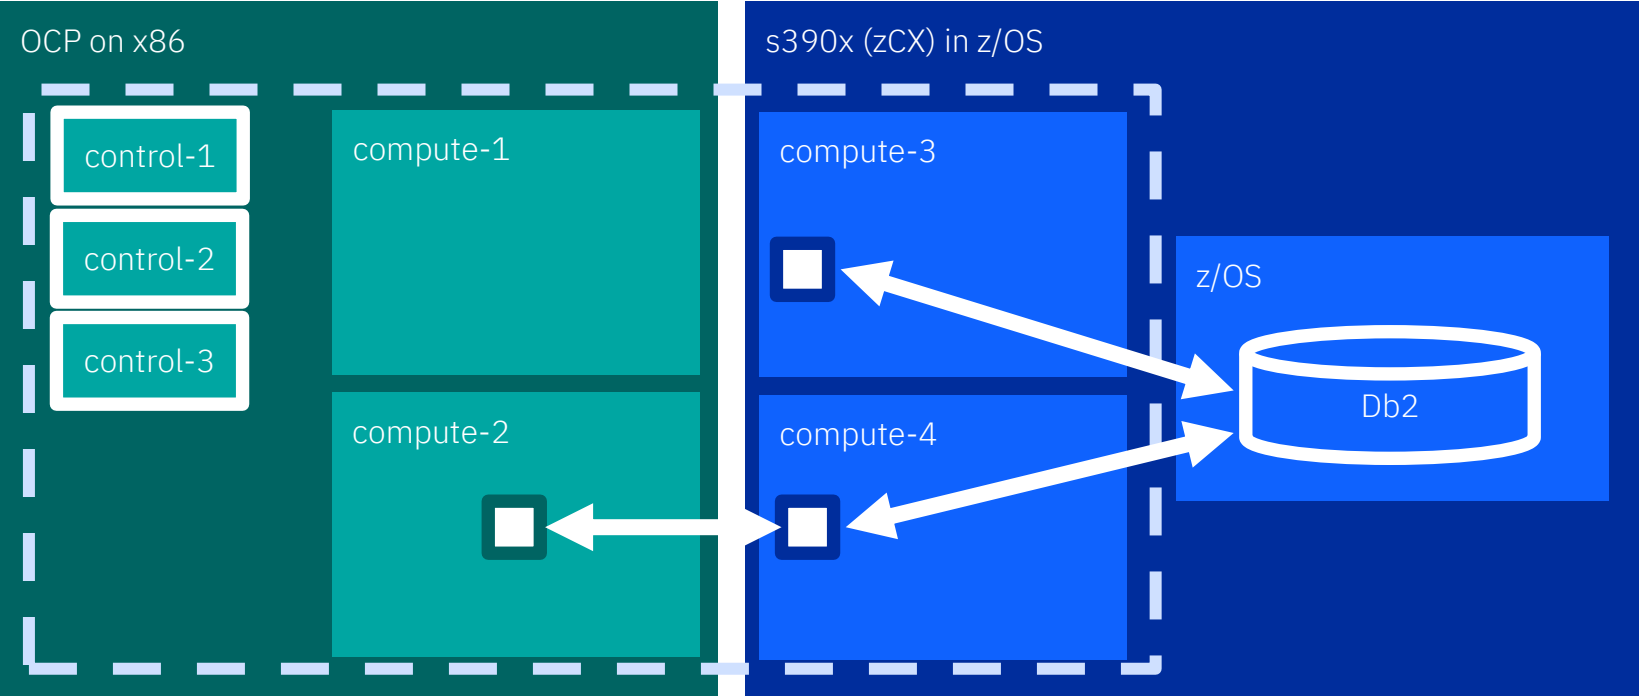
Best Practices HA Configuration


A single OpenShift cluster that manages applications on and across different architectures.

12

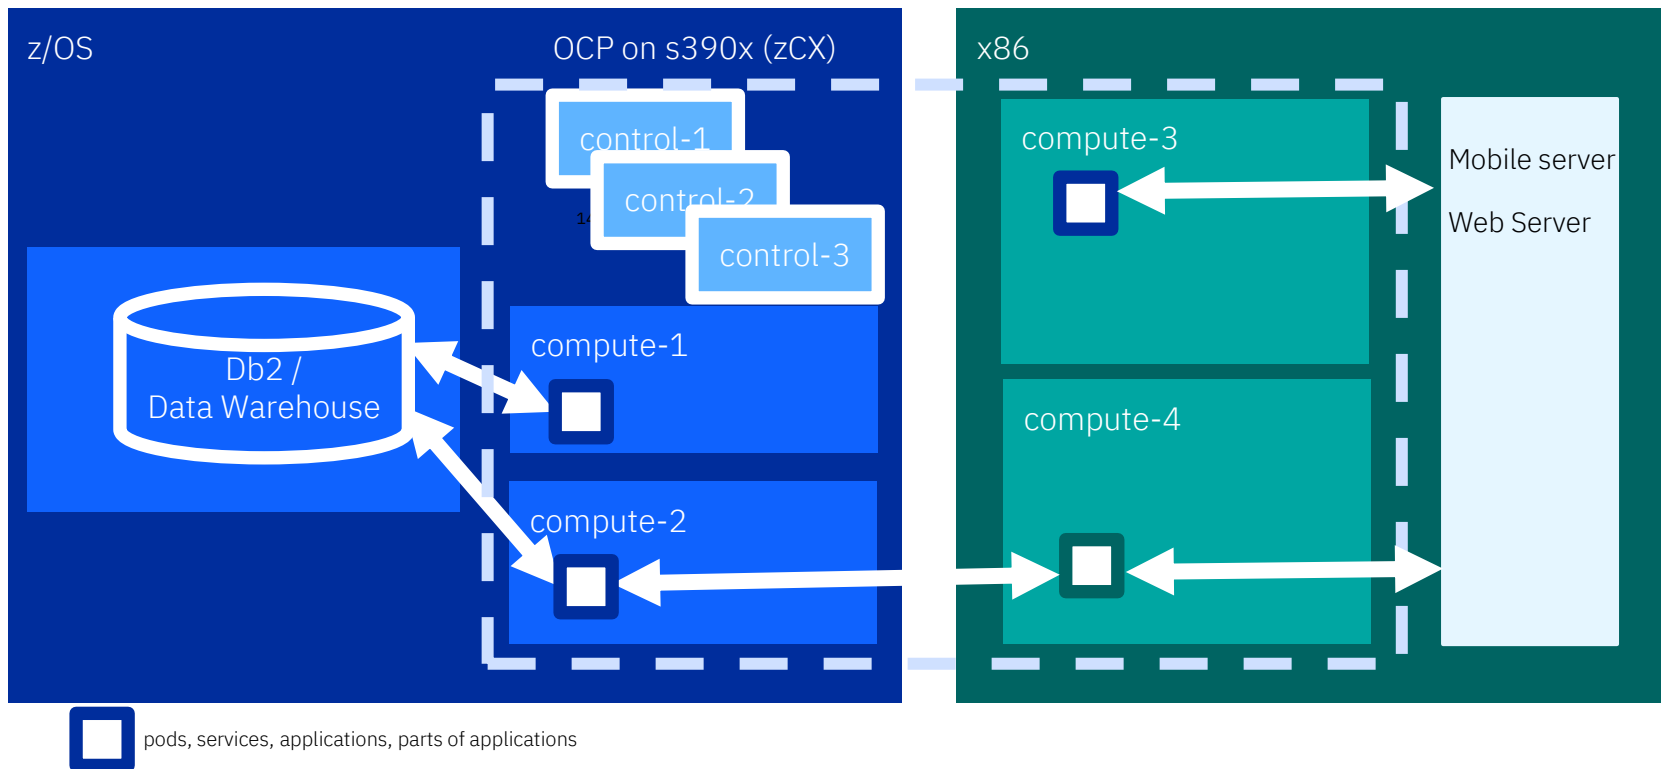


Strategic Workload Placement with Control Planes on x86



 pods, services, applications, parts of applications

Strategic Workload Placement with Control Planes on s390



Configuring multi-architecture compute machines on an OpenShift Container Platform cluster

https://docs.redhat.com/en/documentation/openshift_container_platform/4.17/html/postinstallation_configuration/configuring-multi-architecture-compute-machines-on-an-openshift-cluster#post-install-multi-architecture-configuration

Verifying cluster compatibility

Before you can start adding compute nodes of different architectures to your cluster, you must verify that your cluster is multi-architecture compatible.

Prerequisites

- You installed the OpenShift CLI (oc)

Procedure

- You can check that your cluster uses the multi architecture payload by running the following command:

```
$ oc adm release info -o json | jq .metadata.metadata
```

Verification

1. If you see the following output, then your cluster is using the multi-architecture payload:

```
$ "release.openshift.io/architecture": "multi"
```

You can then begin adding multi-arch compute nodes to your cluster.

2. Your cluster is not using the multi-architecture payload with the response:


```
$ null
```

You can migrate your cluster to one that supports multi-architecture compute machines !

Steps are in: "[Migrating to a cluster with multi-architecture compute machines](#)".

2) Multi Architecture Hosted Control Planes cluster management

- Red Hat OpenShift clusters managed via [Hosted Control Planes](#)



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

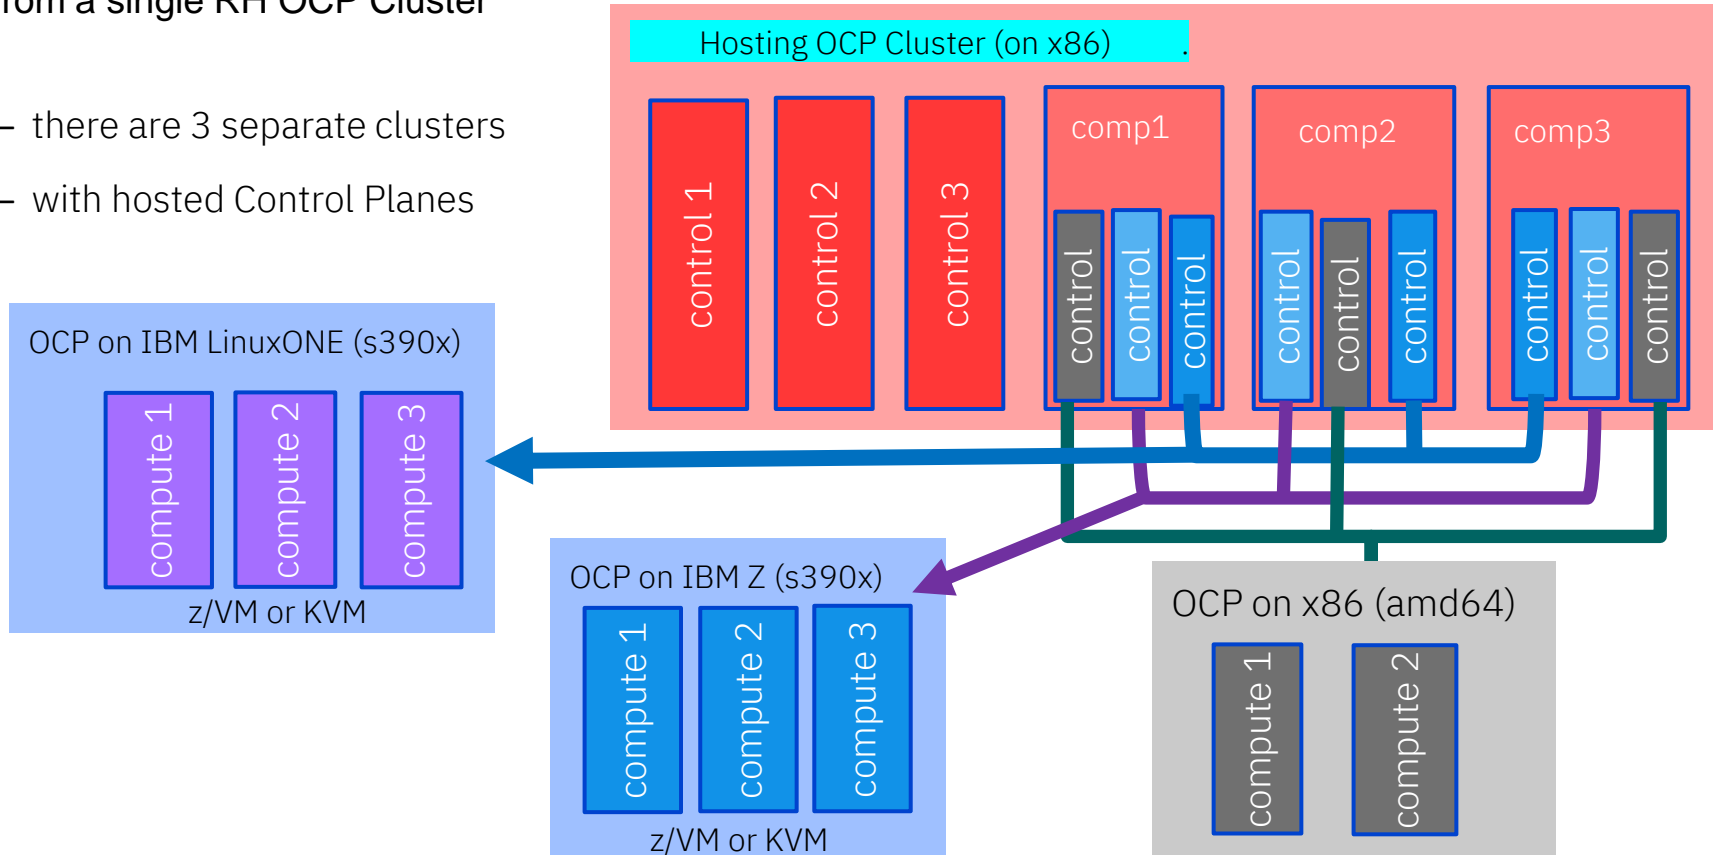
Hosted Cluster consist of Control and Compute Nodes and is an independent entity

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

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



Implementation decisions & steps

- Solution Architecture design**

- Do you want to extend a single OCP cluster (Multi Arch Compute)
- or keep the environments separated (Hosted Control Planes)

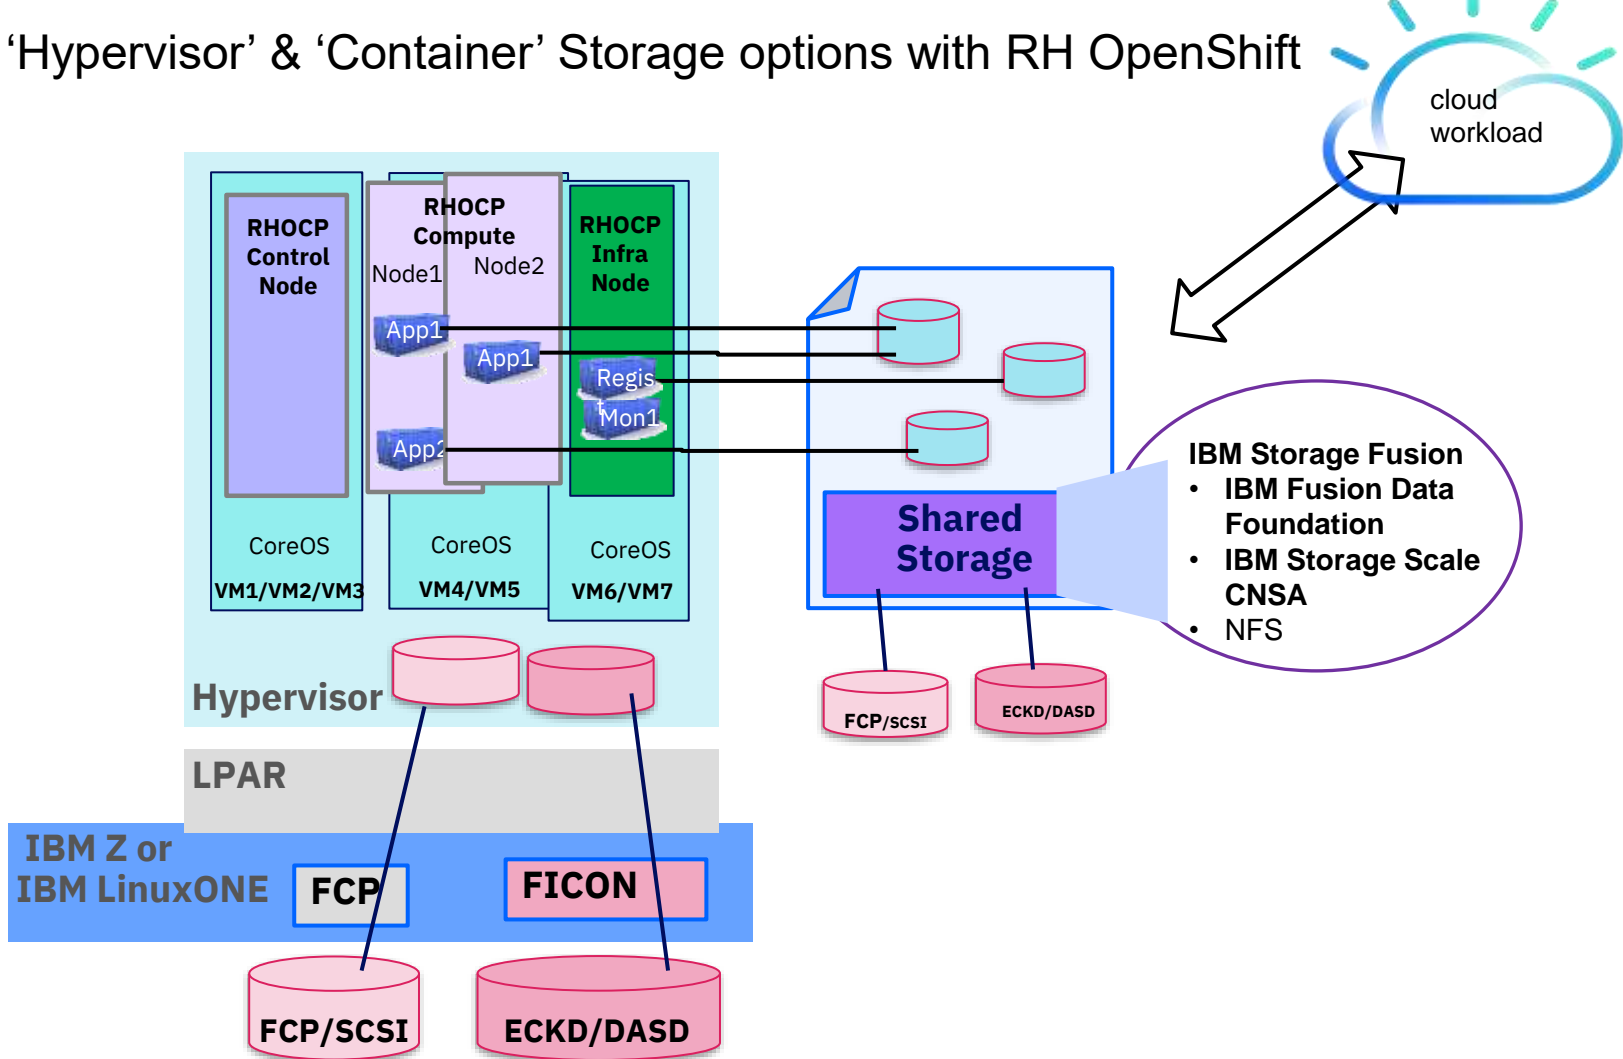
- Where is the base / main cluster from which you want to extend workloads:**

- on IBM Z / LinuxONE
- or on x86

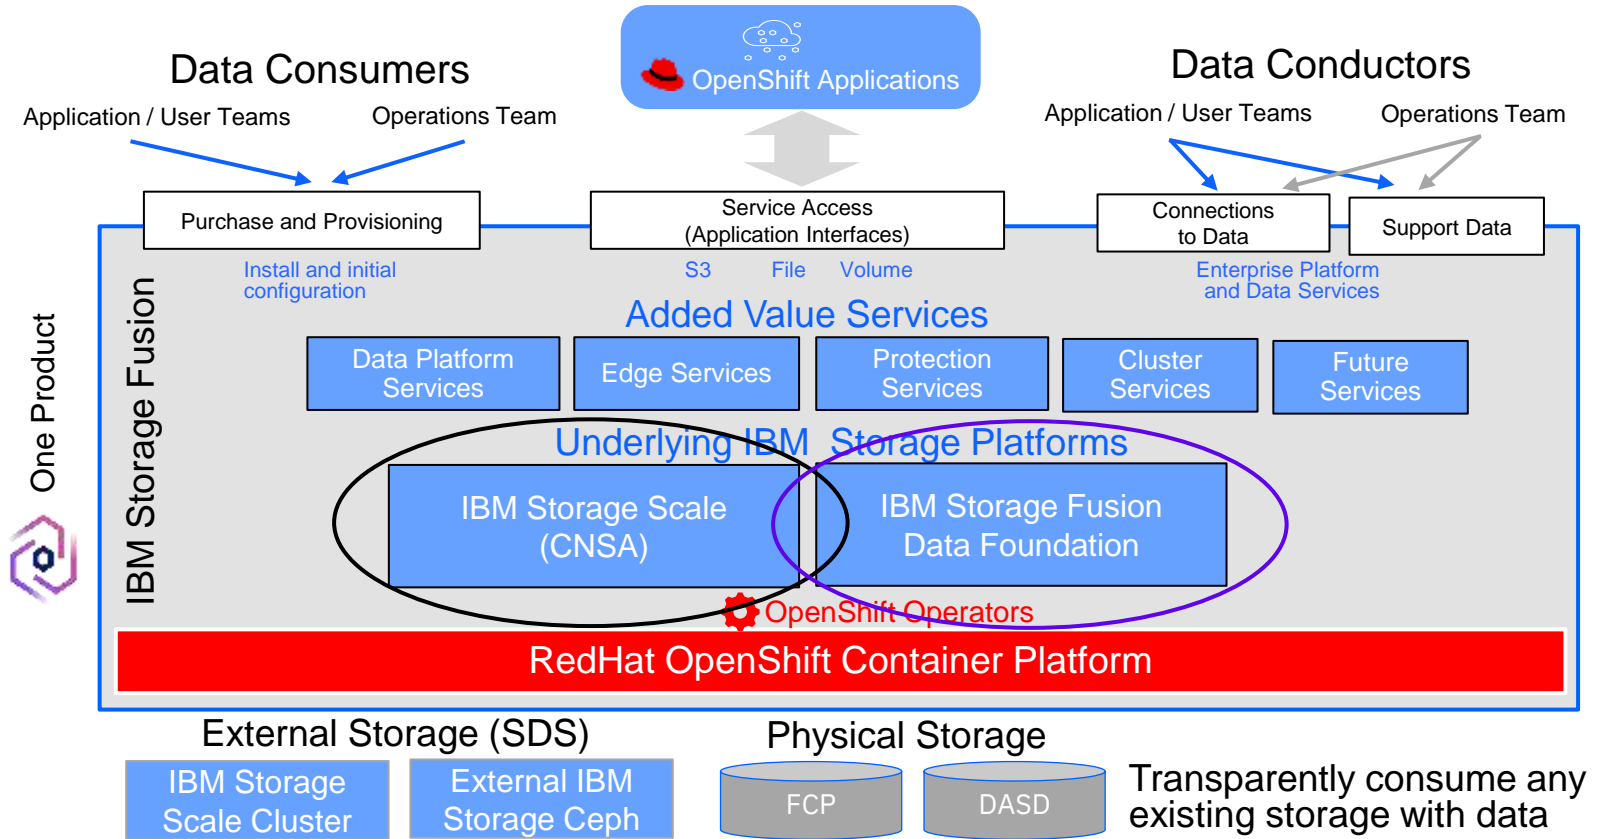
- Storage planning - which storage implementation fits best**

- IBM Storage Fusion Data Foundation
- IBM Storage Scale (is now part of Storage Fusion bundle)
- Is back-end data used in z/OS Db2, you still require a shared Software Defined Storage

Physical 'Hypervisor' & 'Container' Storage options with RH OpenShift



IBM Storage Fusion: Software Stack – as Storage Cloud



IBM Fusion: Two delivery options

Customer Apps



IBM Cloud Paks



Databases



Off the shelf



Custom apps

Offerings

On-prem Solution with IBM Fusion and IBM LinuxONE

Data services for stateful
OpenShift applications



Fusion HCI (appliance)

Integrated application
platform for Red Hat OpenShift



Built on

Red Hat OpenShift Container Platform

Deployments

IBM
LinuxONE

- Fully integrated stack on IBM LinuxONE
- Sales Play SDS

External
HW unit

- Supports x86 workload and GPU
- Pre-assembled

3. Plan for deployment option, infrastructure and operational / management model

- Decide about the deployment option**

- Single Node, Three Node, Multi-Node
- In LPARs or in a virtualized environment with z/VM or KVM

- Infrastructure topology**

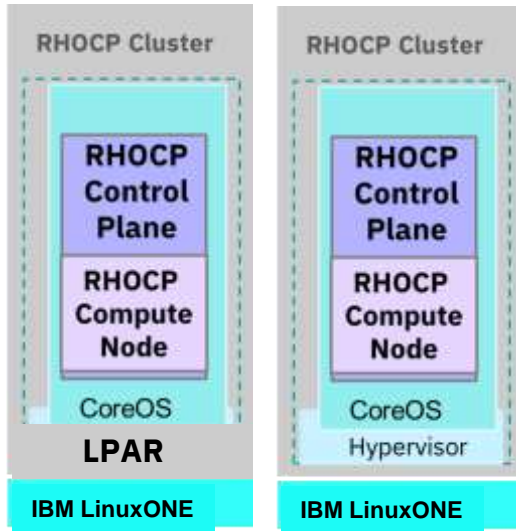
- HA requirements for the applications and nodes
- DR requirements and failover capabilities between clouds

- Operational Model - Is a centralized management for the Multi-Arch solution desired**

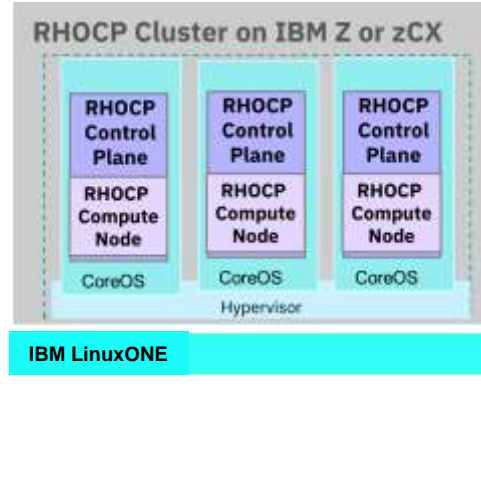
- Can be using Hosted Control 'Planes
- Can be using Advance Cluster Management for Kubernetes (ACM)

Red Hat OpenShift Single Node and HA deployment options

Single Node Cluster (SNO)

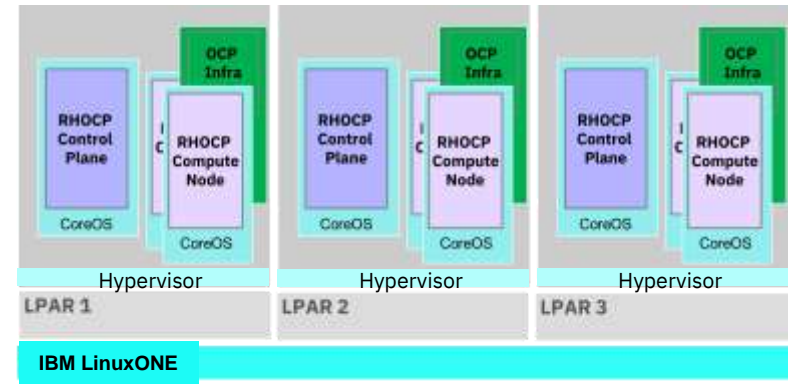


3 Node Cluster



All deployments can be in **LPARs** without hypervisor

Multi Node Cluster

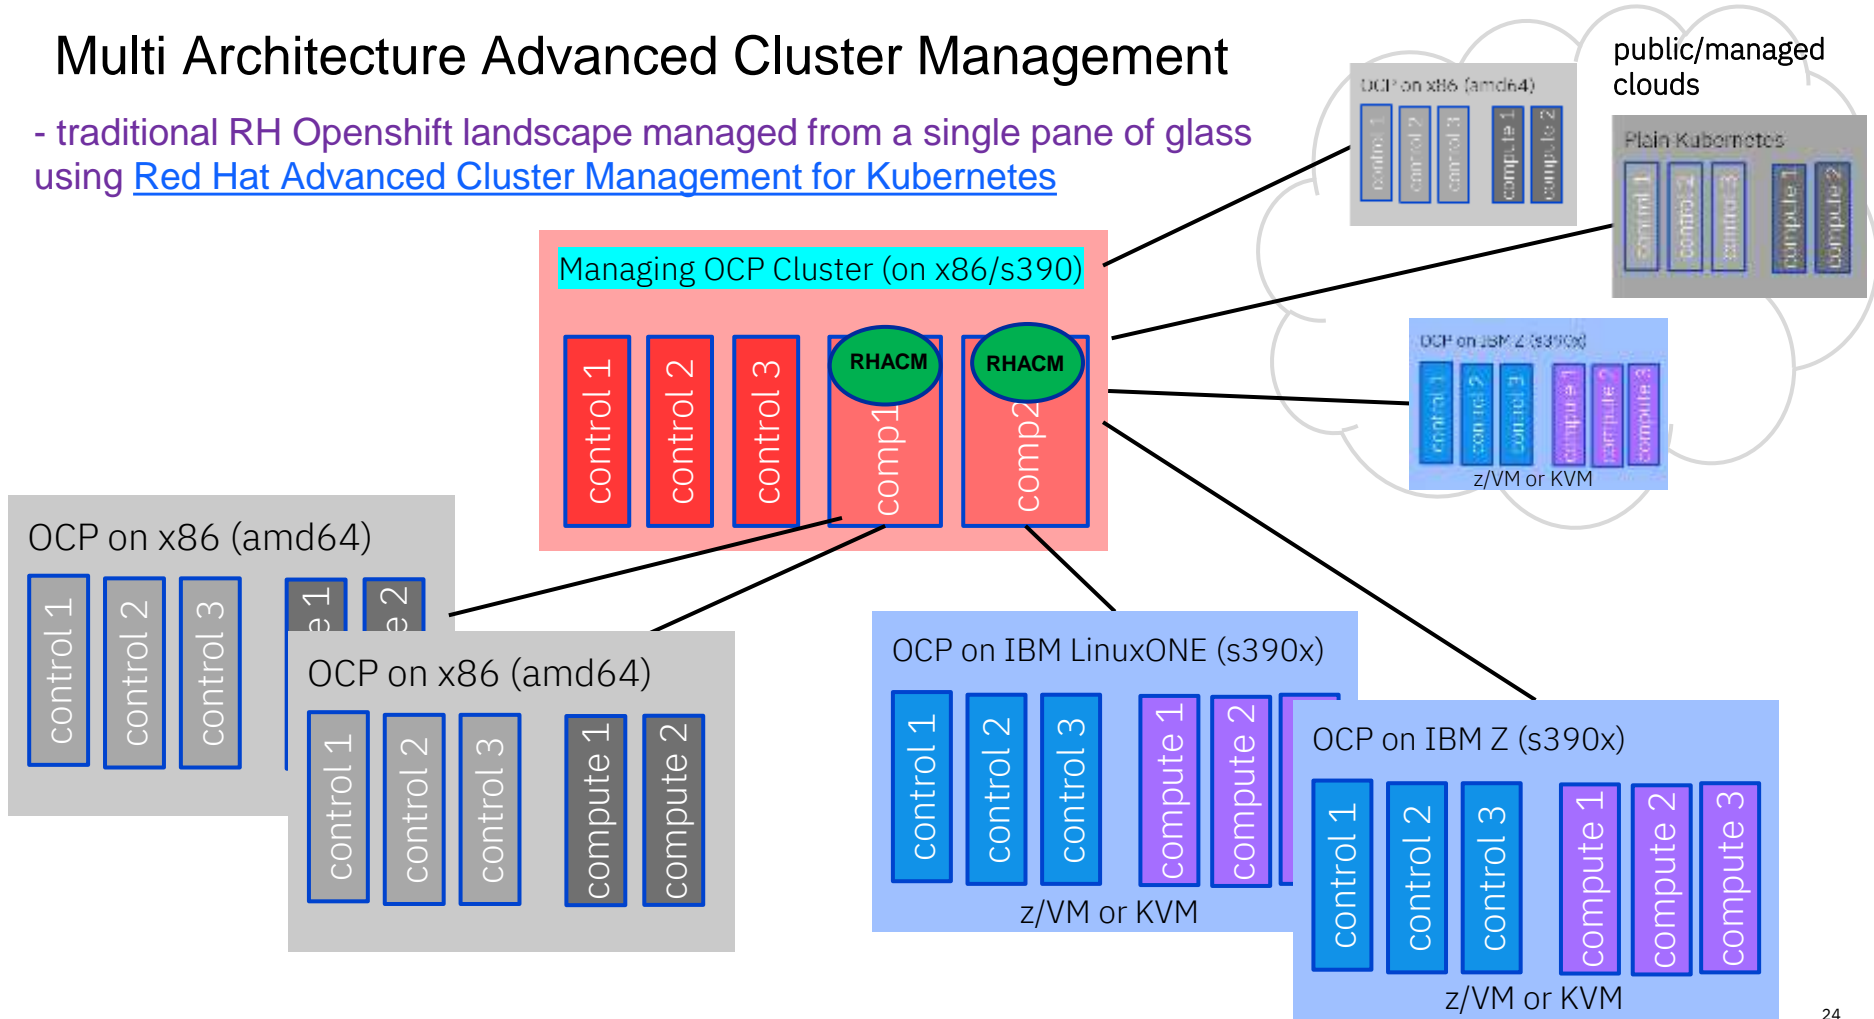


SNO Value: Lower resource foot-print due to smaller cluster size

- 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 / ZIIPs**, with SMT-2 enabled **on installation time** and requires **0.5 IFLs SteadyState on average**; with **1.5 IFLs** required for **keeping SNO OCP stable** and for running some workloads
https://docs.redhat.com/en/documentation/openshift_container_platform/4.18/html/installing_on_a_single_node/install-sno-installing-sno#install-sno-with-ibmz

Multi Architecture Advanced Cluster Management

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



Red Hat Advanced Cluster Manager Overview (RHACM)



Multicluster lifecycle management



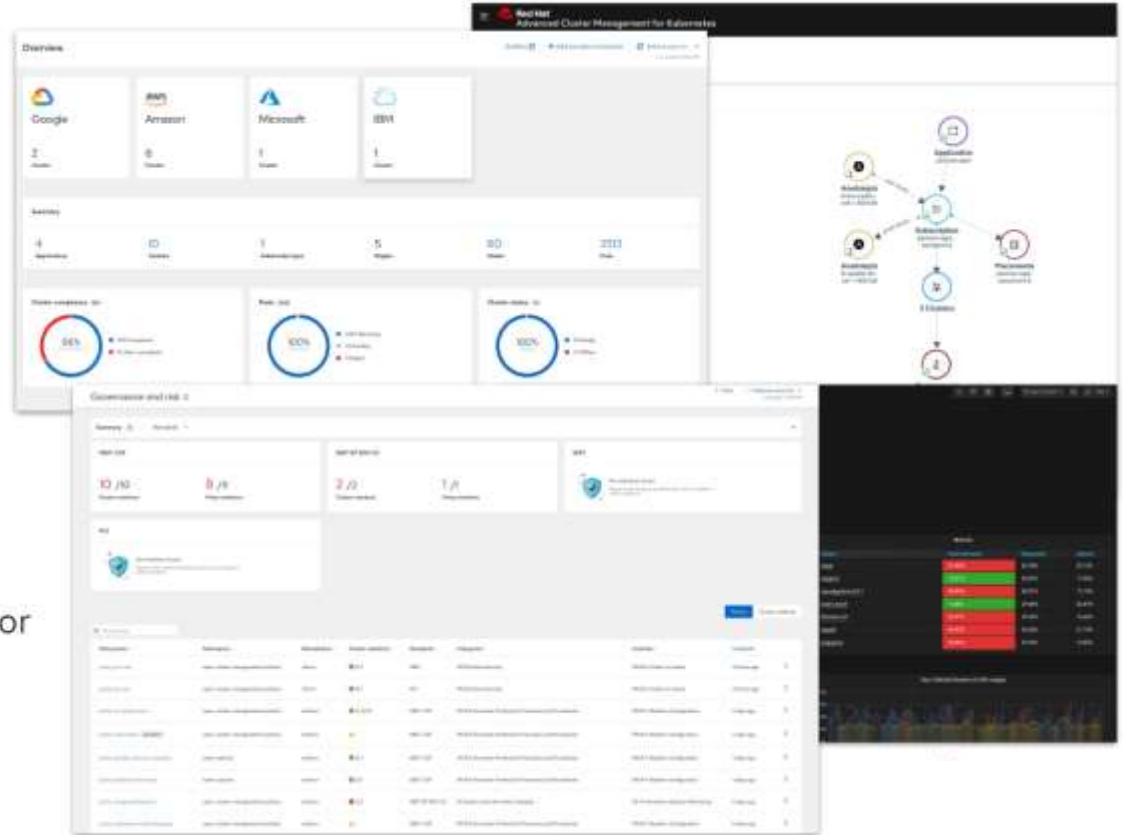
Policy driven governance, risk, and compliance



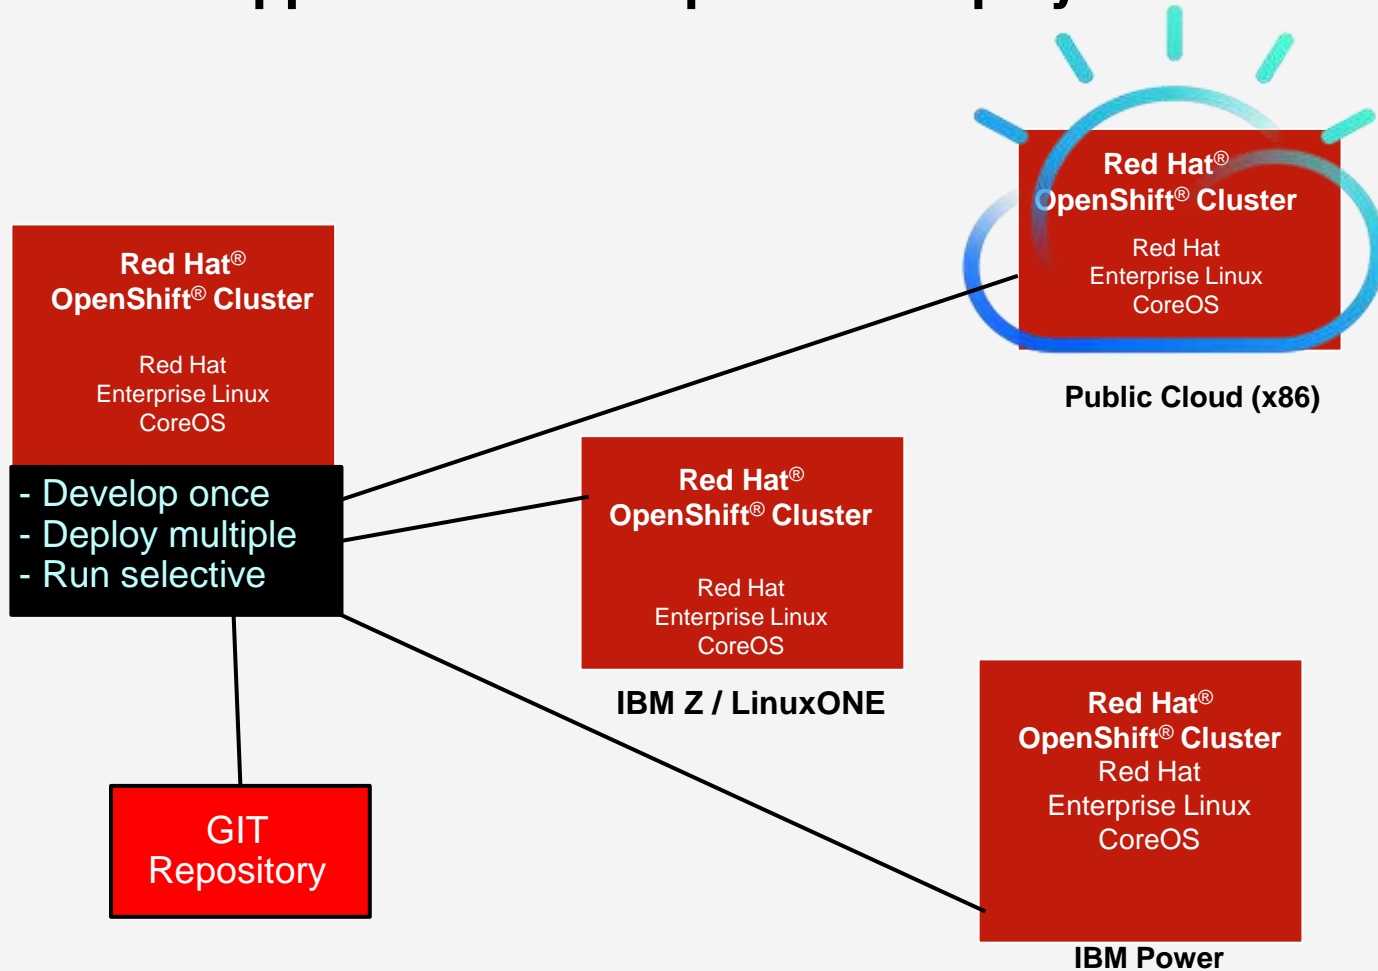
Advanced application lifecycle management



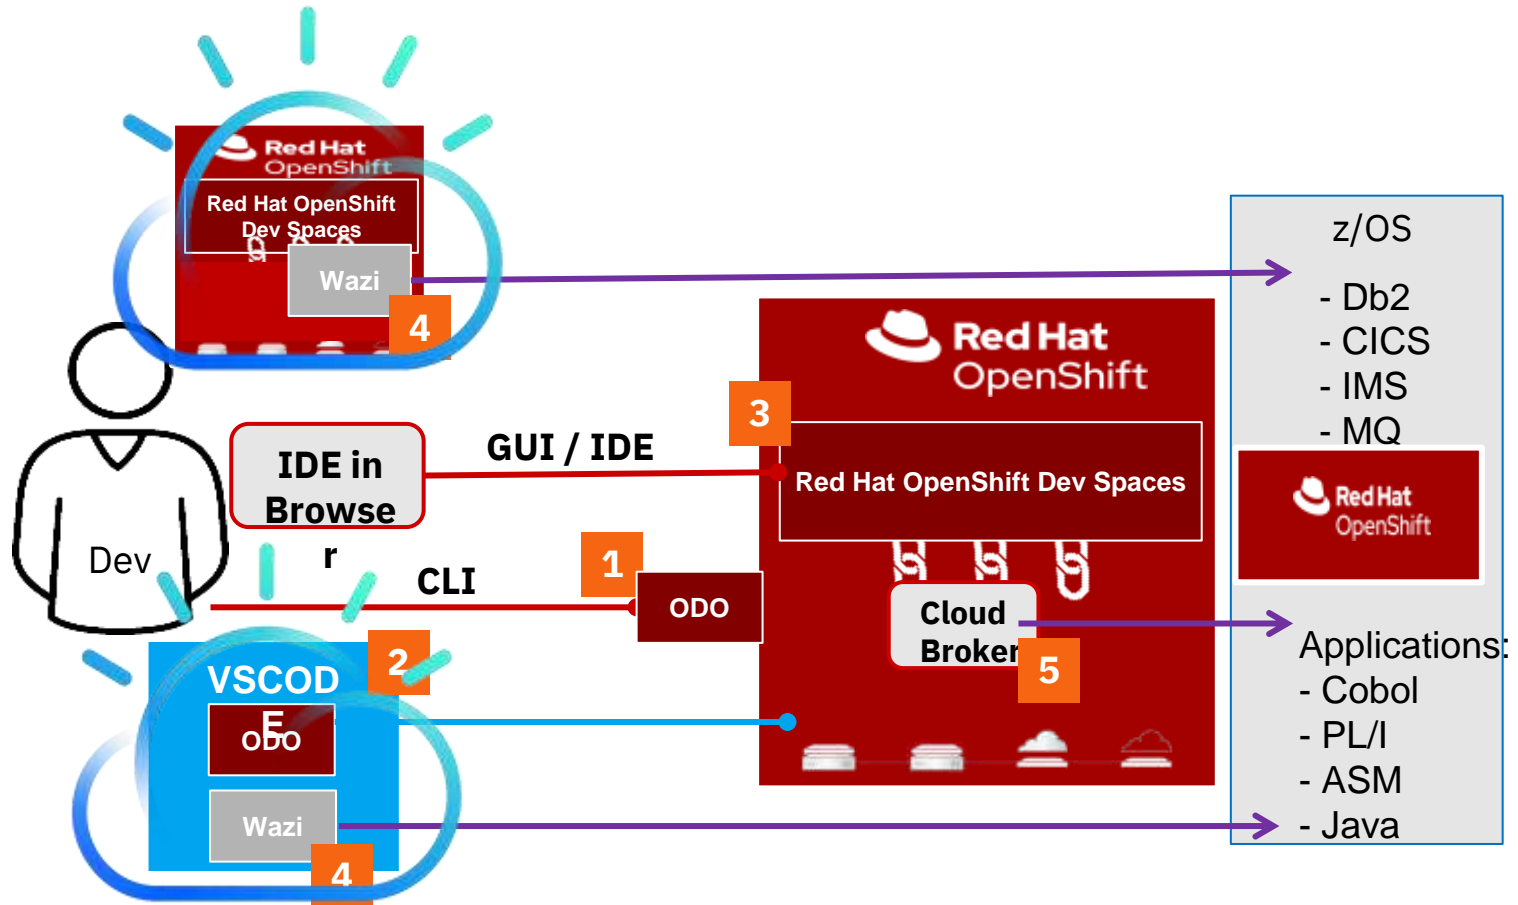
Multicluster observability for health and optimization



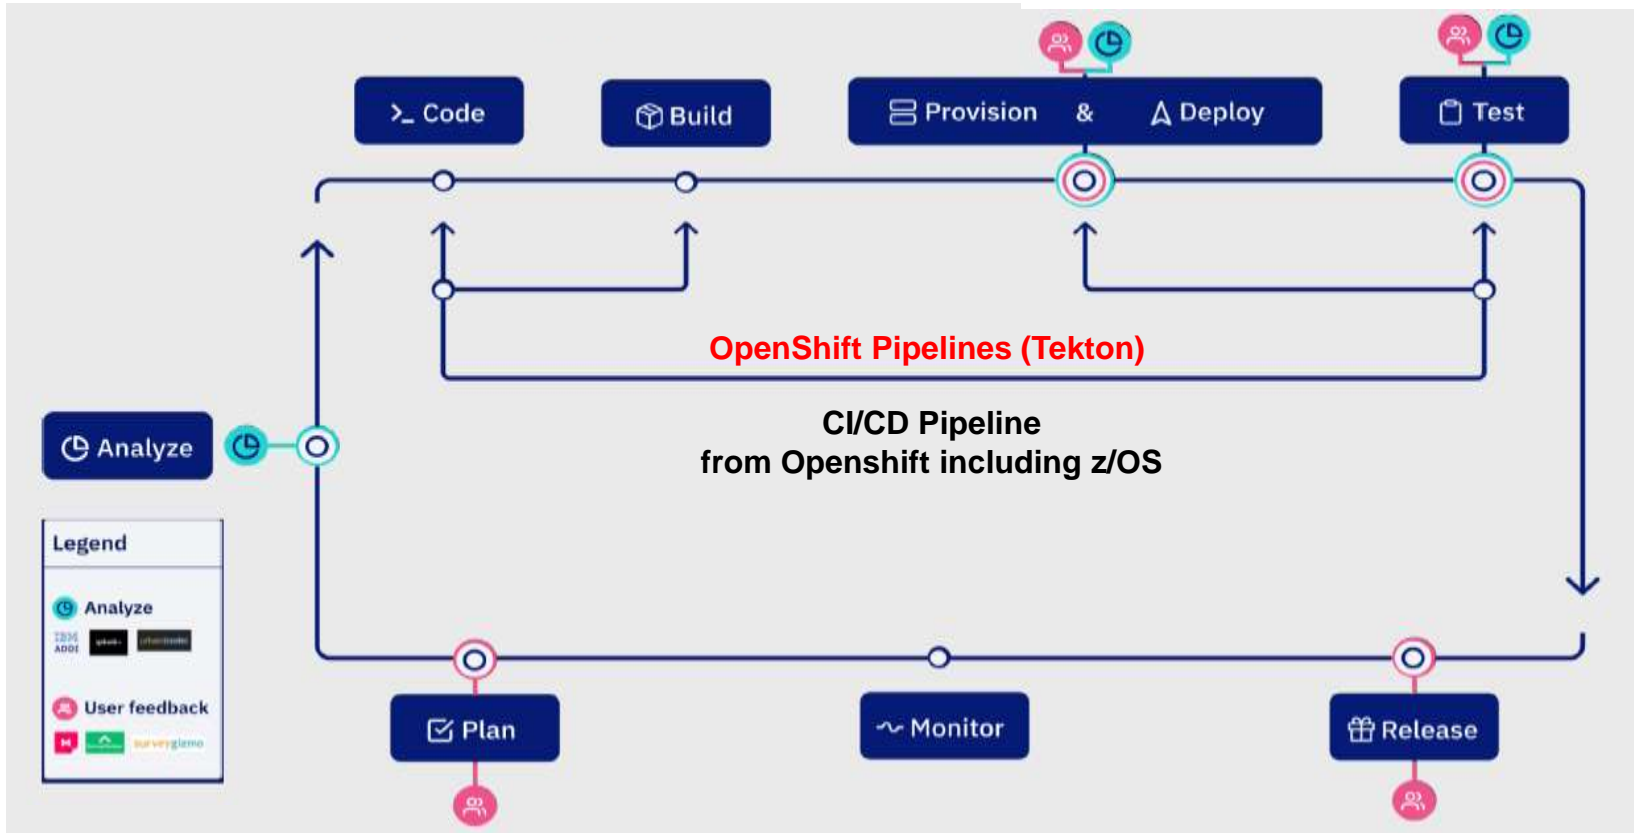
Cross platform Application development & deployment consistency



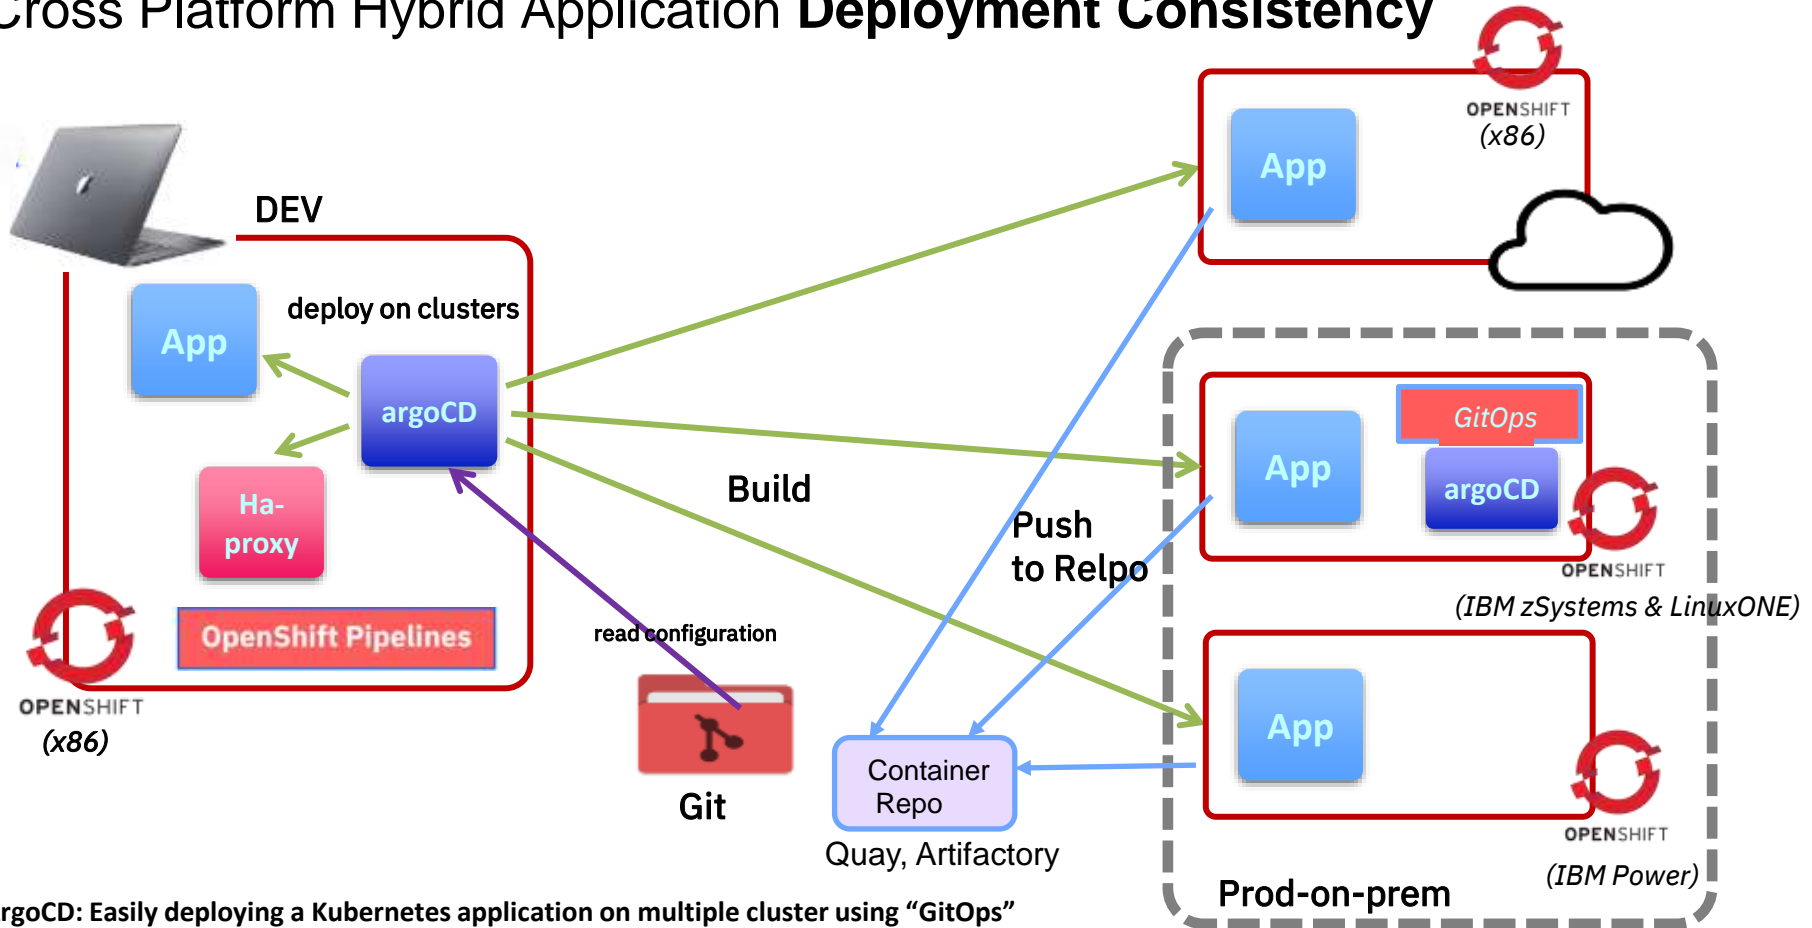
Cloud-native development tools in Red Hat OpenShift for IBM Z & LinuxONE



What makes up DevOps and CI/CD



Cross Platform Hybrid Application 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



Use a Multi-Arch build pipeline with OpenShift Pipelines

Project: multiarch-build-dev ▾

PipelineRuns > PipelineRun details

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

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

PipelineRun details

```
graph LR; setup[setup] --> code_lint[code-lint]; code_lint --> dockerfile[dockerfile...]; dockerfile --> simver[simver]; simver --> build_x86[build-x86]; simver --> build_power[build-power]; simver --> build_z[build-z]; build_x86 --> manifest[manifest]; build_power --> manifest; build_z --> manifest; manifest --> deploy[deploy]; deploy --> health[health]; health --> img_scan[img-scan]; img_scan --> tag_release[tag-release];
```

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>

Use Operators for ease of deployment in Multi-Arch:

<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-arch images. The 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 each architecture.

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

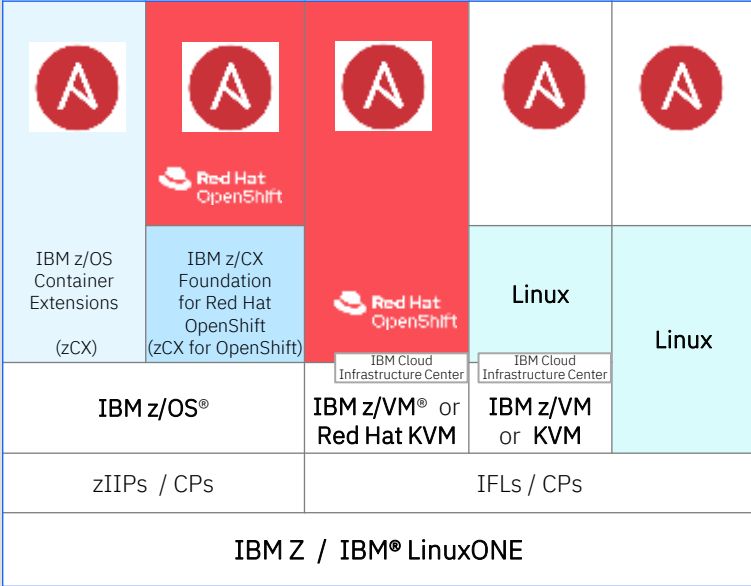
For a given container image, you can check which architectures are supported by listing them by inspecting the manifest. Piping the output to the `python json.tool` module enables pretty-printed JSON output.

Red Hat Ansible Automation Platform Deployment options on IBM Z / IBM LinuxONE

Red Hat Ansible Automation Platform runs on

- Linux
- Red Hat OpenShift in virtual machines
- IBM zCX Foundation for Red Hat OpenShift
- IBM z/OS Container Extensions

on IBM Z / IBM LinuxONE.






Ansible Galaxy – Collections

https://galaxy.ansible.com/ui/collections/?page_size=10&view_type=list&sort=name&page=1&keywords=HMC

☰ GALAXY English ?

- Ansible Galaxy
- Search
- Collections ▾
 - Collections
 - Namespaces
- Roles >
- Documentation ↗
- Terms of Use ↗

	<p>ibm_zhmc Provided by ibm</p> <p>34 Modules 0 Roles 1 Plugin 0 Dependencies</p> <p>z hmc ibm dpm infrastructure Show Less</p>	<p>Updated 14 days ago v1.9.4</p>
	<p>power_hmc Provided by ibm</p> <p>13 Modules 1 Role 22 Plugins 0 Dependencies</p> <p>hmc ibm power 1 more</p>	<p>Updated 3 days ago v1.12.0 🔗 5</p>
	<p>zhmc Provided by sabyadi</p> <p>0 Modules 0 Roles 0 Plugins 0 Dependencies</p> <p>z hmc ibm dpm Show Less</p>	<p>Updated 1 year ago v1.0.0</p>

Red Hat Ansible Certified Content for Hardware Management Console

available at

catalog.redhat.com/software/collection/ibm/ibm_zhmc

The available **IBM Z HMC collection** provides Ansible modules and sample playbooks for automating management tasks on the Hardware Management Console (HMC) of IBM Z and IBM LinuxONE machines, such as creating, updating or deleting partitions and other resources.

LPAR automation example – in the RH OpenShift installation using Ansible:

<https://ibm.github.io/Ansible-OpenShift-Provisioning/set-variables-group-vars/>

Modules (26)

zhmc_adapter	zhmc_partition
zhmc_adapter_list	zhmc_partition_list
zhmc_console	zhmc_password_rule
zhmc_cpc	zhmc_password_rule_list
zhmc_cpc_list	zhmc_session
zhmc_crypto_attachment	zhmc_storage_group
zhmc_hba	zhmc_storage_group_attachment
zhmc_ldap_server_definition	zhmc_storage_volume
zhmc_ldap_server_definition_list	zhmc_user
zhmc_lpar	zhmc_user_list
zhmc_lpar_list	zhmc_user_role
zhmc_nic	zhmc_user_role_list
zhmc_nic_list	zhmc_virtual_function

IBM Z 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



IBM trusted Registry (with registration to IBM Cloud):

<https://ibm.biz/BdfAdW>

Red Hat catalog with container images for IBM zSystems & LinuxONE

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

The screenshot displays the Red Hat Ecosystem Catalog interface. At the top, the navigation bar includes the Red Hat logo and links for Hardware, Software, and Cloud & service providers. The main heading is "Container images", with a subtext explaining that they offer lightweight and self-contained software for scalable deployment. Below this, a breadcrumb trail shows "Home > Software > Container images". A search bar and a "Search" button are present, along with a "1 - 15 of 589" result indicator. A filter sidebar on the left shows "Architecture" with "s390x" selected, and "Category" with "API Management", "Application Delivery", "Application Development", "Application Server", and "Automation" listed. The main content area features six container image cards, each with the Red Hat logo and a description:

- Red Hat**
rhel7
Red Hat Enterprise Linux 7
by Red Hat, Inc.
Provides the latest release of Red Hat Enterprise Linux 7 in a fully featured and supported base image.
Updated 3 days ago
- Red Hat**
redhat-openjdk-18/openjdk18-spendfit
Java Applications
by Red Hat, Inc.
OpenJDK 8 Image for Java Applications
Updated 5 days ago
- Red Hat**
rhe0/etcd
Etcd
by Red Hat, Inc.
A highly-available key-value store for shared configuration.
Updated 3 days ago
- Red Hat**
rhel7/atomic
RHEL Atomic Base Image
by Red Hat, Inc.
Provides the latest release of Red Hat Enterprise Linux 7 in a fully supported
- Red Hat**
rhc/httpd-24-rhel7
Apache httpd 2.4
by Red Hat, Inc.
Platform for running Apache httpd 2.4 or building httpd-based application
- Red Hat**
rhc/s2i-base-rhel7
s2i base
by Red Hat, Inc.
Base image with essential libraries and tools used as a base for builder images

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 [Red Hat 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



IT Architecture
GISEI, Lead IT Architect

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