Break-out Session on
Red Hat OpenShift Container Platform for IBM Z® & LinuxONE

@ 13. OpenShift Anwendertreffen

Wilhelm Mild
IBM Executive IT Architect

Hendrik Brückner
IBM Manager Linux on Z Development
Red Hat Partner Engineer for RHEL & RHOCP on Z

IBM Germany Research & Development GmbH

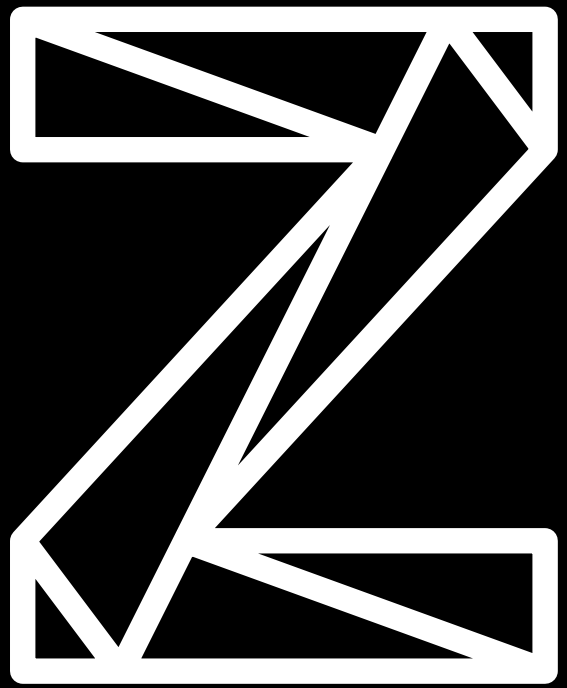

IBM

Mainframe Break-out

Speaker Introduction

Why and benefits of Hybrid Multi-Cloud environments on IBM Z & LinuxONE?

What does RH OCP look like on IBM Z & LinuxONE?

Open Discussion

Red Hat OpenShift Container Platform for IBM Z & LinuxONE

Hendrik Brückner

Manager Linux on Z Development

Red Hat Partner Engineer for RHEL & RHOCP on Z

IBM DE R&D GmbH

Wilhelm Mild

IBM Executive IT Architect

Integration Architectures for Container, Mobile, IBM Z, and Linux

IBM DE R&D GmbH

News and Updates
Red Hat OpenShift Container Platform for IBM Z® & LinuxONE

Wilhelm Mild
IBM Executive IT Architect

Hendrik Brückner
IBM Manager Linux on Z Development
Red Hat Partner Engineer for RHEL & RHOCP on Z

IBM Germany Research & Development GmbH

Trademarks

The following are trademarks of the International Business Machines Corporation in the United States and/or other countries.

API Connect*	Cognos*	IBM*	IBM Z*	Power*	z13s*
Aspera*	DataPower*	IBM (logo)*	IMS	Power Systems	z/OS*
CICS*	Db2*	IBM Cloud*	LinuxONE	WebSphere*	z/VM*
Cloud Paks	FileNet*	ibm.com*	MobileFirst*	z13*	

*** Registered trademarks of IBM Corporation**

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.
Cell Broadband Engine is a trademark of Sony Computer Entertainment, Inc. in the United States, other countries, or both and is used under license therefrom.
IT Infrastructure Library is a Registered Trade Mark of AXELOS Limited.
ITIL is a Registered Trade Mark of AXELOS Limited.
Linear Tape-Open, LTO, the LTO Logo, Ultrium, and the Ultrium logo are trademarks of HP, IBM Corp. and Quantum in the U.S. and other countries.
Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.
The registered trademark Linux® is used pursuant to a sublicense from the Linux Foundation, the exclusive licensee of Linus Torvalds, owner of the mark on a worldwide basis.
Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.
Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.
Red Hat®, JBoss®, OpenShift®, Fedora®, Hibernate®, Ansible®, CloudForms®, RHCA®, RHCE®, RHCSA®, Ceph®, and Gluster® are trademarks or registered trademarks of Red Hat, Inc. or its subsidiaries in the United States and other countries.
RStudio®, the RStudio logo and Shiny® are registered trademarks of RStudio, Inc.
UNIX is a registered trademark of The Open Group in the United States and other countries.
VMware, the VMware logo, VMware Cloud Foundation, VMware Cloud Foundation Service, VMware vCenter Server, and VMware vSphere are registered trademarks or trademarks of VMware, Inc. or its subsidiaries in the United States and/or other jurisdictions.
Zowe™, the Zowe™ logo and the Open Mainframe Project™ are trademarks of The Linux Foundation.
Other product and service names might be trademarks of IBM or other companies.

Notes:

Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated here.
IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.
All customer examples cited or described in this presentation are presented as illustrations of the manner in which some customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions.
This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area.
All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.
Information about non-IBM products is obtained from the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the 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.
Prices subject to change without notice. Contact your IBM representative or Business Partner for the most current pricing in your geography.
This information provides only general descriptions of the types and portions of workloads that are eligible for execution on Specialty Engines (e.g, zIIPs, zAAPs, and IFLs) ("SEs"). IBM authorizes customers to use IBM SE only to execute the processing of Eligible Workloads of specific Programs expressly authorized by IBM as specified in the “Authorized Use Table for IBM Machines” provided at www.ibm.com/systems/support/machine_warranties/machine_code/aut.html (“AUT”). No other workload processing is authorized for execution on an SE. IBM offers SE at a lower price than General Processors/Central Processors because customers are authorized to use SEs only to process certain types and/or amounts of workloads as specified by IBM in the AUT.

This talk is about...

Introduction

IBM's Hybrid Multicloud Strategy for IBM Z & LinuxONE

IBM Cloud Paks for IBM Z & LinuxONE

Overview of

Red Hat OpenShift Container Platform for IBM Z & LinuxONE

Creating the world's leading hybrid cloud provider

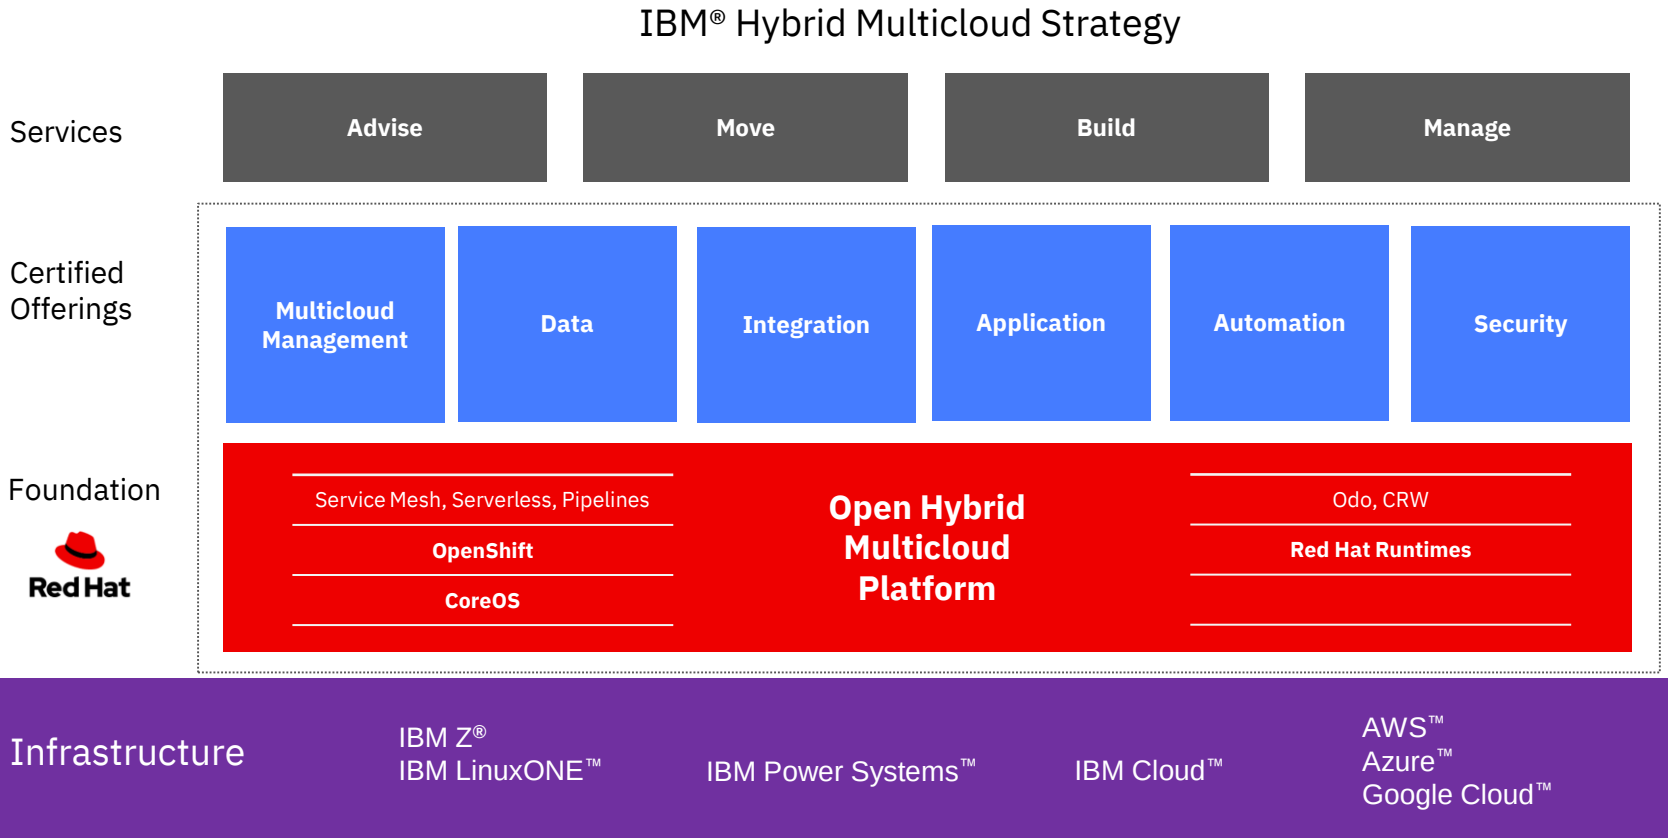

Why Hybrid Multi-Cloud with IBM Z & LinuxONE

Benefits on IBM Z

Low Latency and Large Volume **Data Serving** and **Transaction processing**

Enterprise class infrastructure **– Elastic, Scalable, Available and Resilient**

Highest levels of **Security and Compliance**

Adoption Patterns

Enterprise scale **Private Cloud-in-a-Box**

Digital Transformation and Modernization for Apps

Built-in secure enclaves for **Zero Trust Cloud Native**

Extreme consolidation and scalable **Data Serving**

*Scale-out to **2,4 million containers** on a single system*

*Reducing data-center footprint by **4:1** and power by **2:1***

*Process over **1 Trillion** encrypted transactions per day*

*Fully encrypt container data (at-rest, in-flight) and apps with **ZERO code changes***

Enterprise grade. **Open** by design. **Secured** by IBM LinuxONE

Integrate existing and new services across hybrid IT

Offerings designed for

Cloud Native Development

on IBM Z & LinuxONE

- IBM z15 & IBM LinuxONE III
- IBM z/VM 7.1
- IBM Cloud Infrastructure Center
- Red Hat OpenShift Container Platform
- IBM Cloud Paks
- IBM z/OS Cloud Broker
- IBM Hyper Protect Virtual Servers
- IBM Blockchain Platform SW

RED HAT® OPENSHIFT Container Platform

IBM Cloud Paks

enterprise-ready, containerized software solutions

IBM Cloud Infrastructure Center
Architecture Overview

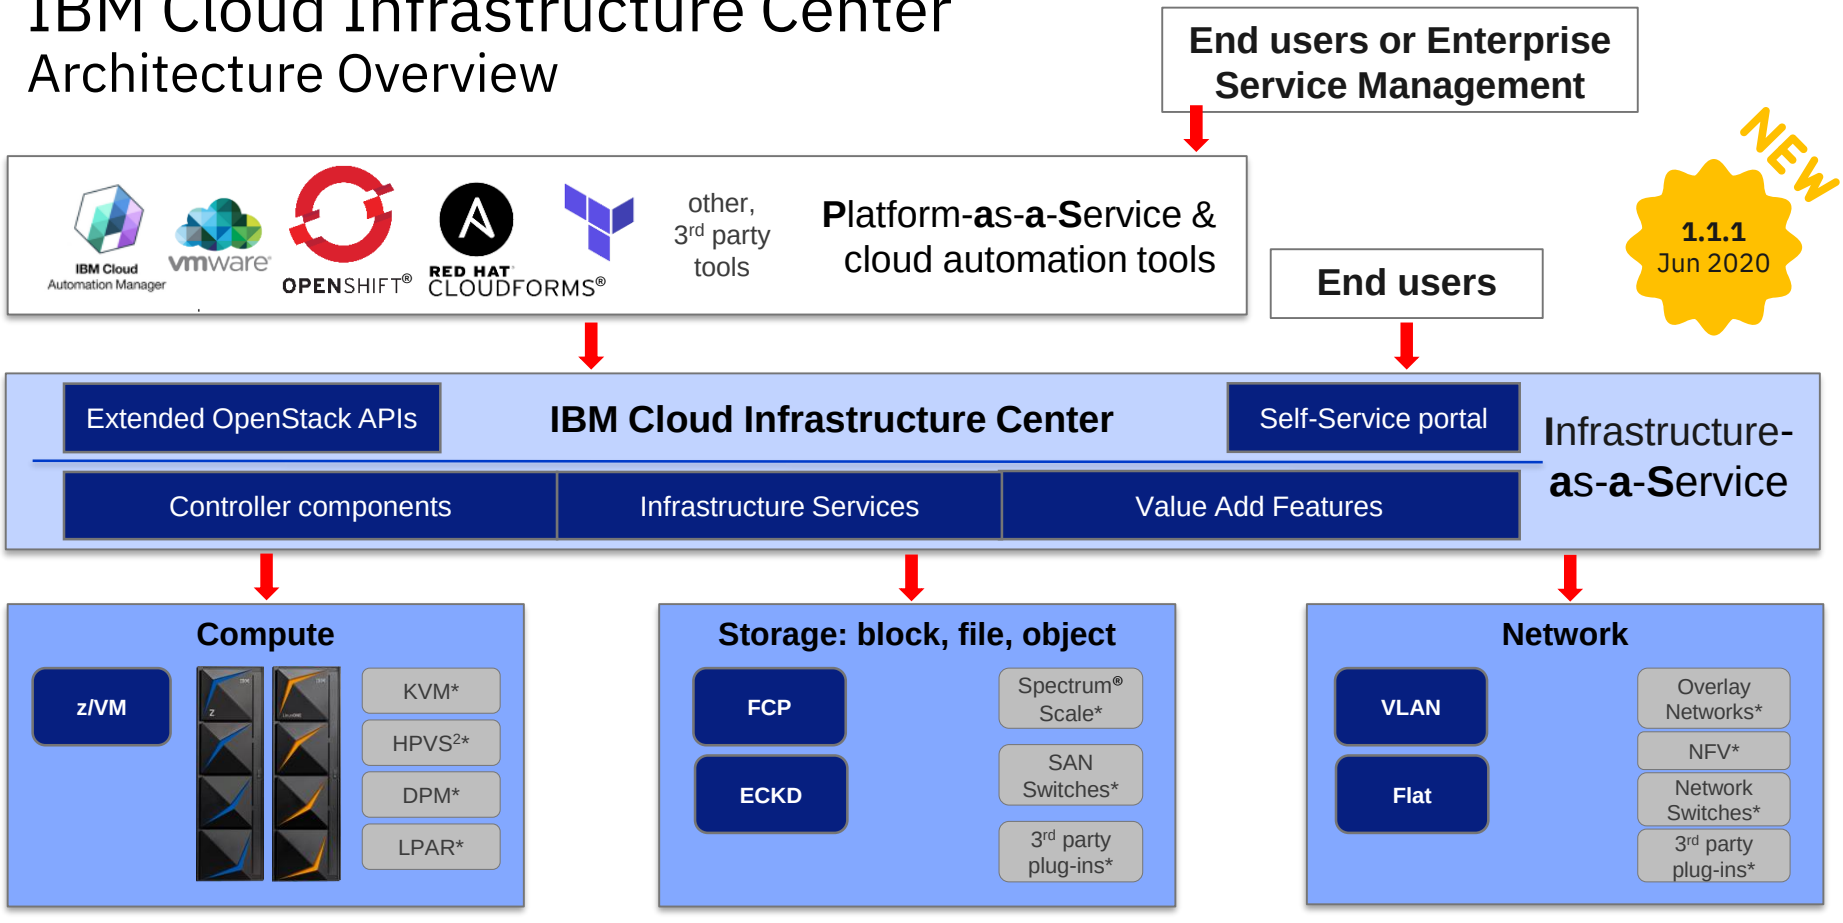

[1] *IBM Hyper Protect Virtual Server*

** All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.*

Provision RHEL, CoreOS, SLES 15 SP1 und Ubuntu 20.04 guests

IBM Software as Cloud Paks – *Middleware anywhere*

A faster, more secure way to move your core business applications to any cloud through enterprise-ready containerized software solutions

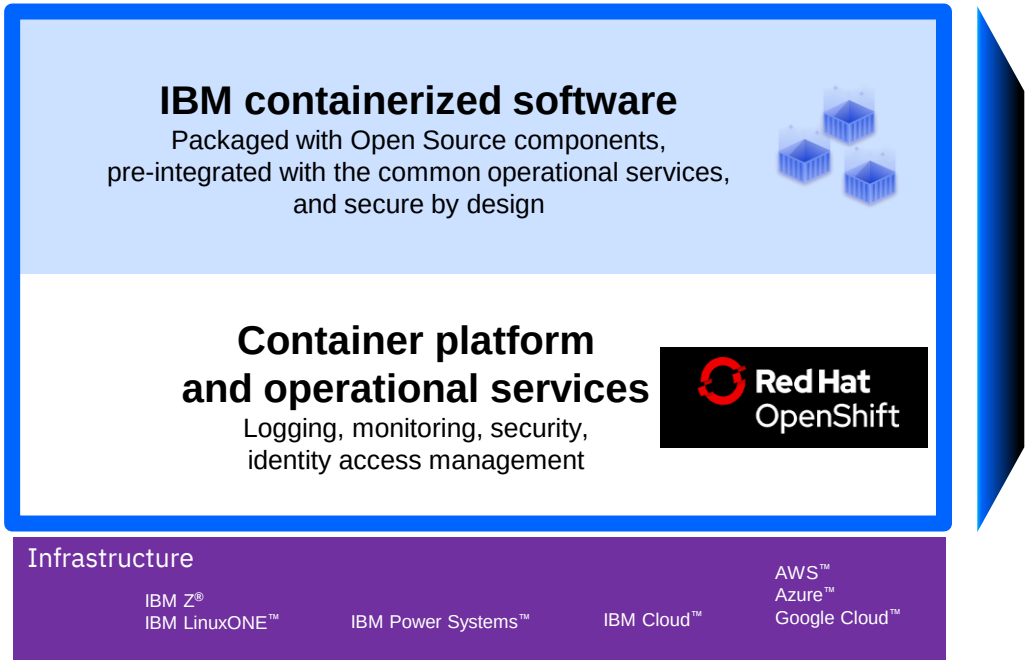

Complete yet simple
Application, data and AI services, fully modular and easy to consume

IBM certified
Full software stack support, and ongoing security, compliance and version compatibility

Run anywhere
On-premises, on private and public clouds, and in pre-integrated systems

IBM Cloud Paks on IBM Z and LinuxONE - *Roadmap*

All Cloud Paks are coming to IBM Z and LinuxONE in Various Phases!

Available except Accelerator for Teams Full release – 1Q 2021*	Phase 1 - 4Q 2020	Phase 1 – GA 9/25 Next Release - 4Q 2020*	Manage-to-Z (MCM) – GA 8/07 Manage-to-Z (RHACM) – 1Q 2021 Manage from-Z – 2H 2021	Phase 1 (target) in 1H 2021	Target – TBD
Cloud Pak for Applications Build, deploy and run applications	**Cloud Pak for Data** Collect, organize, and analyze data	**Cloud Pak for Integration** Integrate applications, data, cloud services, and APIs	**Cloud Pak for Multicloud Management** Multicloud visibility, governance, and automation	**Cloud Pak for Automation** Transform business processes, decisions, and content	**Cloud Pak for Security** Connect security data, tools, and teams
IBM containerized software	IBM containerized software	IBM containerized software	IBM containerized software	IBM containerized software	IBM containerized software
Operational services	Operational services	Operational services	Operational services	Operational services	Operational services
Container platform RH OpenShift 4.x	Container platform RH OpenShift 4.x	Container platform RH OpenShift 4.x	Container platform RH OpenShift 4.x	Container platform RH OpenShift 4.x	Container platform RH OpenShift 4.x

Runs on choice of Linux on IBM Z (z13 or later) and LinuxONE

IBM z/VM
IBM Cloud Infrastructure Center – IaaS (optional)

***Dates listed here are targets only**

Red Hat OpenShift Container Platform 4.x for IBM Z & LinuxONE

Red Hat OpenShift Container Platform (RHOCP) V4 on IBM Z and LinuxONE

- takes advantage of the underlying enterprise capabilities
  - grow to **thousands of Linux guests**
  - and **millions of containers**
- non-disruptively grow, vertical and horizontal scalability
  - including advanced security
  - **confidential Cloud Computing,** including **FIPS 140-2 Level 4** certification

These capabilities were highlighted with the recent announcement of the IBM z15 and IBM LinuxONE III. Running Red Hat OpenShift on IBM Z and LinuxONE also enables cloud native applications to easily integrate with existing data and applications on these platforms, reducing latency by avoiding network delays.

https://www.ibm.com/blogs/systems/get-ready-for-red-hat-openshift-on-ibm-z-and-linuxone/

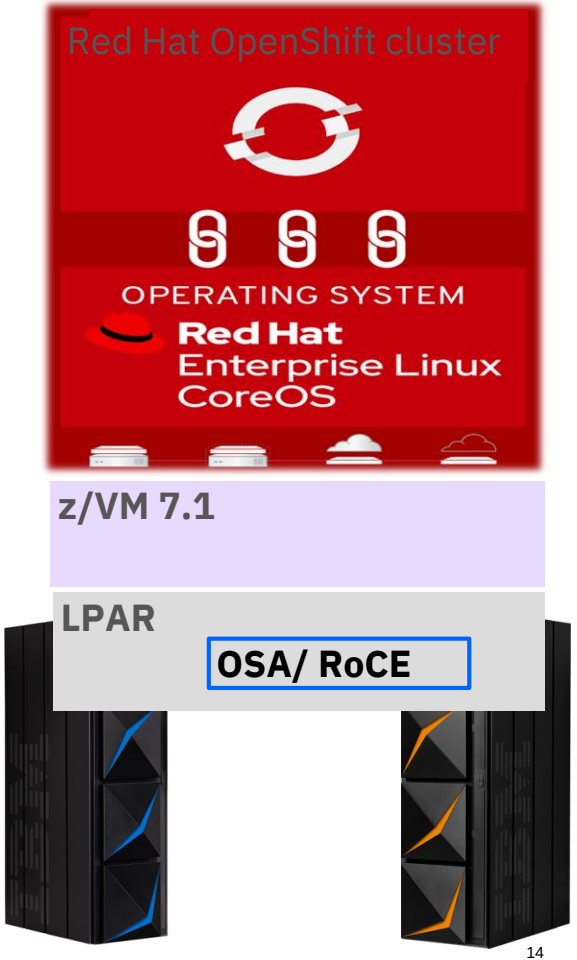

Where can you download RHOCP?

try.openshift.com
cloud.redhat.com

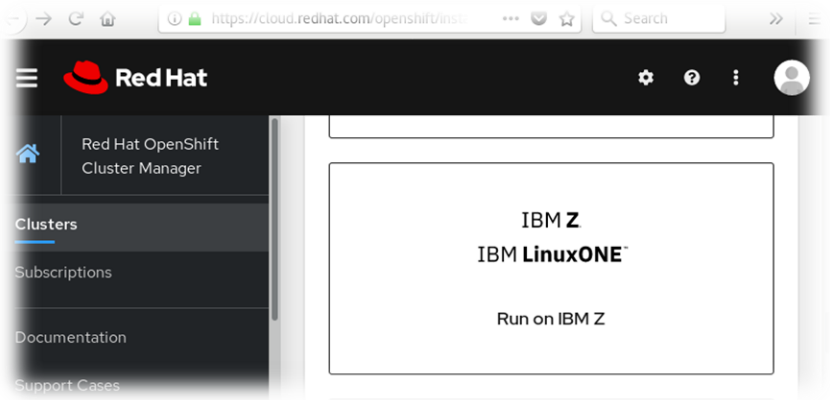

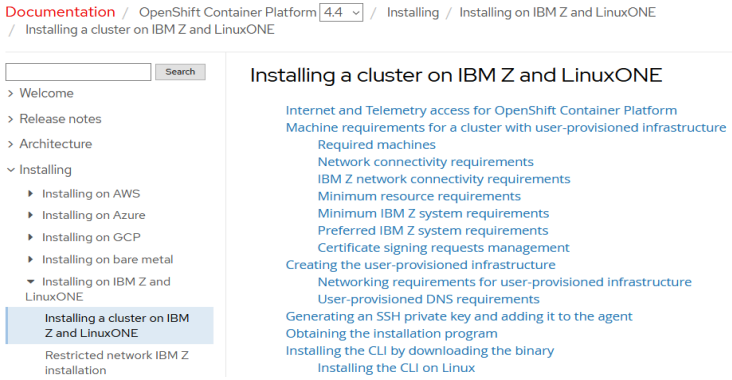

OCP 4.5 on Z was released on 7/30/20
OCP 4.4 on Z was released on 6/22/20
OCP 4.3 on Z was released on 4/30/20
OCP 4.2 on Z was released on 2/11/20

https://docs.openshift.com/container-platform/4.5/installing/installing_ibm_z/installing-ibm-z.html
https://docs.openshift.com/container-platform/4.5/release_notes/ocp-4-5-release-notes.html

Red Hat Runtimes for RHOCP on Z & LinuxONE

Lightweight middleware runtimes and frameworks for developing cloud-native applications on RHOCP

RH Runtimes are available for RH OCP on Z/LinuxONE on 06/15/20

- JBoss Enterprise Application Platform (EAP) 7.3
- JBoss Web Server 5.3
- Red Hat Single Sign-on (SSO) 7.4
- Red Hat Data Grid 7.3
- AMQ (Broker) 7.5
- Quarkus 1.3.4
- Vert.x 3.9.1
- Thorntail 2.5.1
- Spring Boot 2.2.6
- Node.js 10 & 12

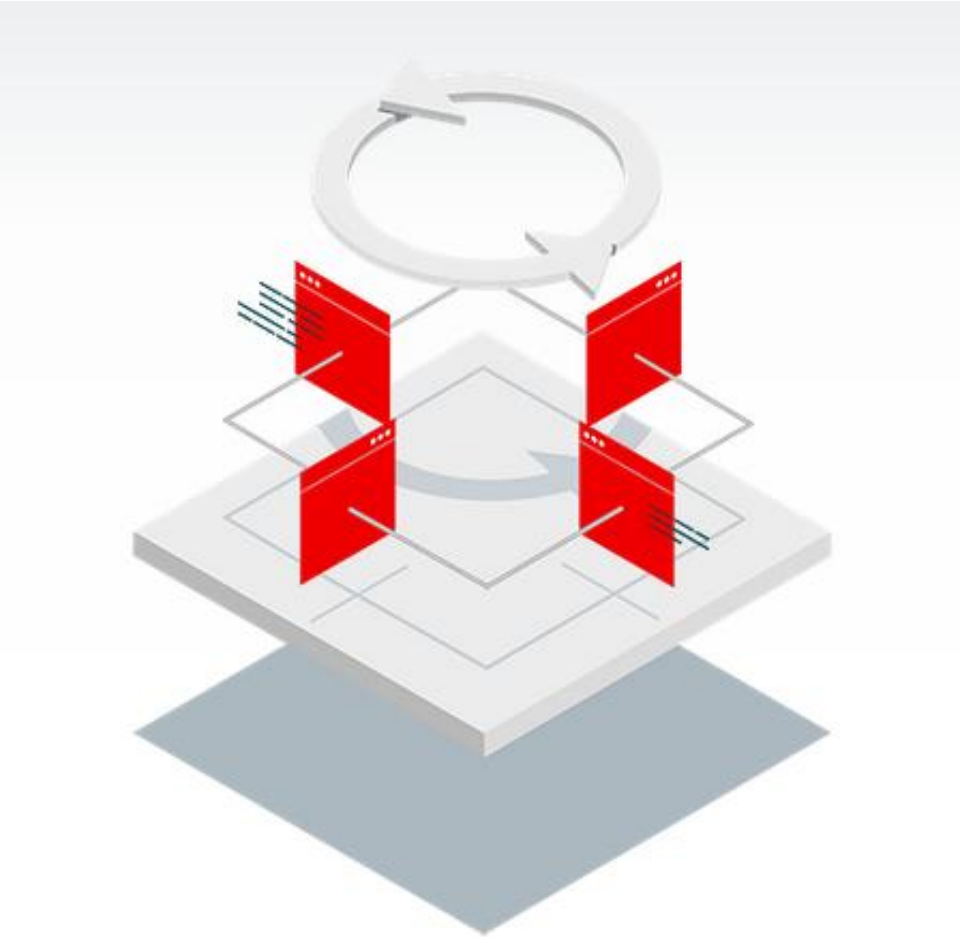

https://catalog.redhat.com/software/containers/search?p=1&architecture=s390x

Developer Experience for RHOCP on Z & LinuxONE

Developer CLI – OpenShift do (odo)

- You can use odo for creating applications on OpenShift Container Platform and Kubernetes
- **odo 1.2.6 available for IBM Z & LinuxONE on 9/23/20**
  - https://docs.openshift.com/container-platform/4.5/cli_reference/developer_cli_odo/installing-odo.html#installing-odo-on-linux-on-ibm-z

What's next? .... CodeReady Workspaces

Red Hat OpenShift Container Platform *on IBM Z/LinuxONE*

Day 1 – Installation and Setup

Planning & Installation tasks

User-Provisioned Infrastructure (UPI) – Platform administrator has to pre-provision infrastructure components

- Planning for required services
- Planning for cluster network
- Planning for storage

Day 2 – Operation and Management

Operational tasks

- (Optionally) Setting up infrastructure nodes
- Establishing etcd backup procedure
- Adding additional worker nodes
- Configuring monitoring and logging
- Integrating and authenticating with LDAP

Operator Framework

Red Hat OpenShift V4 Installation Options

On IBM Z and LinuxONE

OPENSHIFT CONTAINER PLATFORM

Full Stack Automated (IPI)

Simplified opinionated "Best Practices" for cluster provisioning

Fully automated installation and updates including host container OS.

Red Hat Enterprise Linux CoreOS

Pre-existing Infrastructure (UPI)

Customer managed resources & infrastructure provisioning

Plug into existing DNS and security boundaries

Red Hat Enterprise Linux CoreOS

HOSTED OPENSHIFT

Red Hat OpenShift on IBM Cloud *

Deploy directly from the IBM Cloud console. An IBM service, master nodes are managed by IBM Cloud engineers.

Azure Red Hat OpenShift **

Deploy directly from the Azure console. A MSFT service, jointly managed by Red Hat and Microsoft Azure engineers.

OpenShift Dedicated **

Get a powerful cluster, fully managed by Red Hat engineers and support; a Red Hat service.

* Based on RHOCP v4.3 GA slated for March; public beta available now

** Entitlements of RHOCP obtained through a Cloud Pak purchase are not transferable to these environments

Planning: RHOCP on LinuxONE or IBM Z co-located with z/OS

IBM LinuxONE

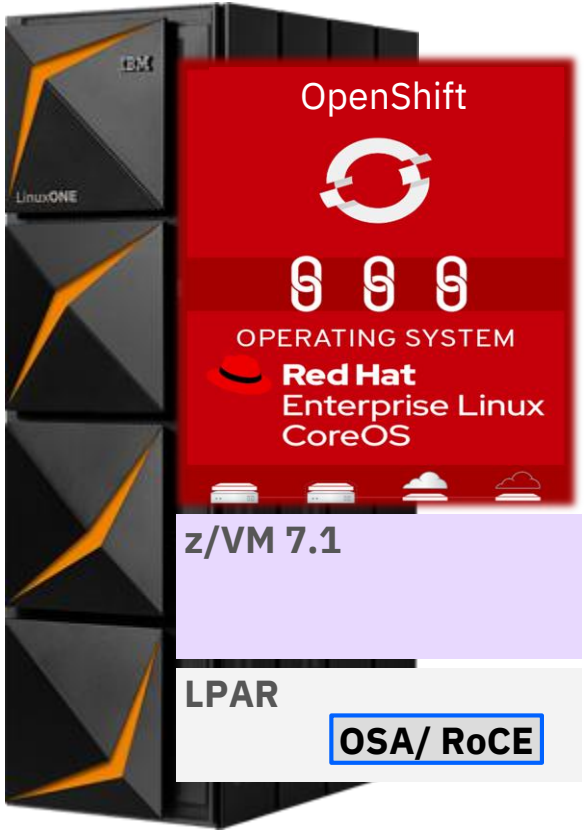

RHOCP standalone

IBM Z

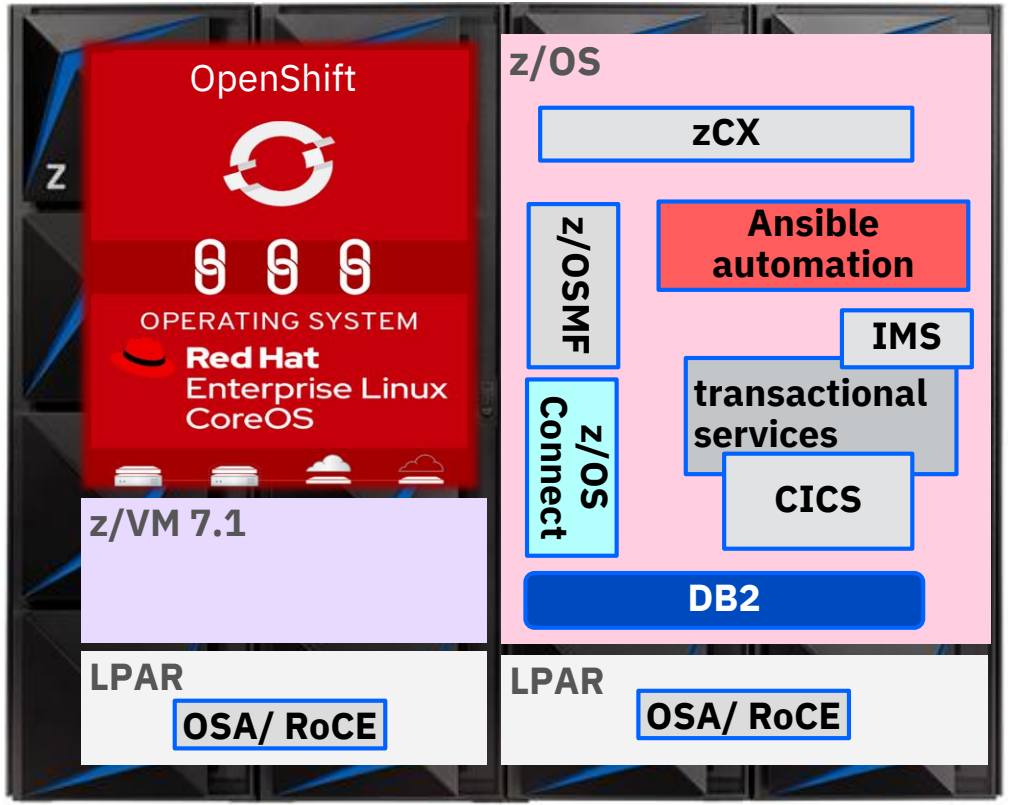

RHOCP co-located with z/OS

RHOCP co-location environment with z/OS

RHOCP co-location to z/OS major use cases:

- Dynamic workload in RHOCP accesses z/OS services
- RHOCP logic access to DB2 z/OS
- RHOCP uses z/OS Cloud Broker to provision z/OS subsystems
- RHOCP Web environment integrates with z/OS transactional services
- RHOCP running Open Source technologies extends z/OS services
- Batch workload executed in RHOCP with z/OS data

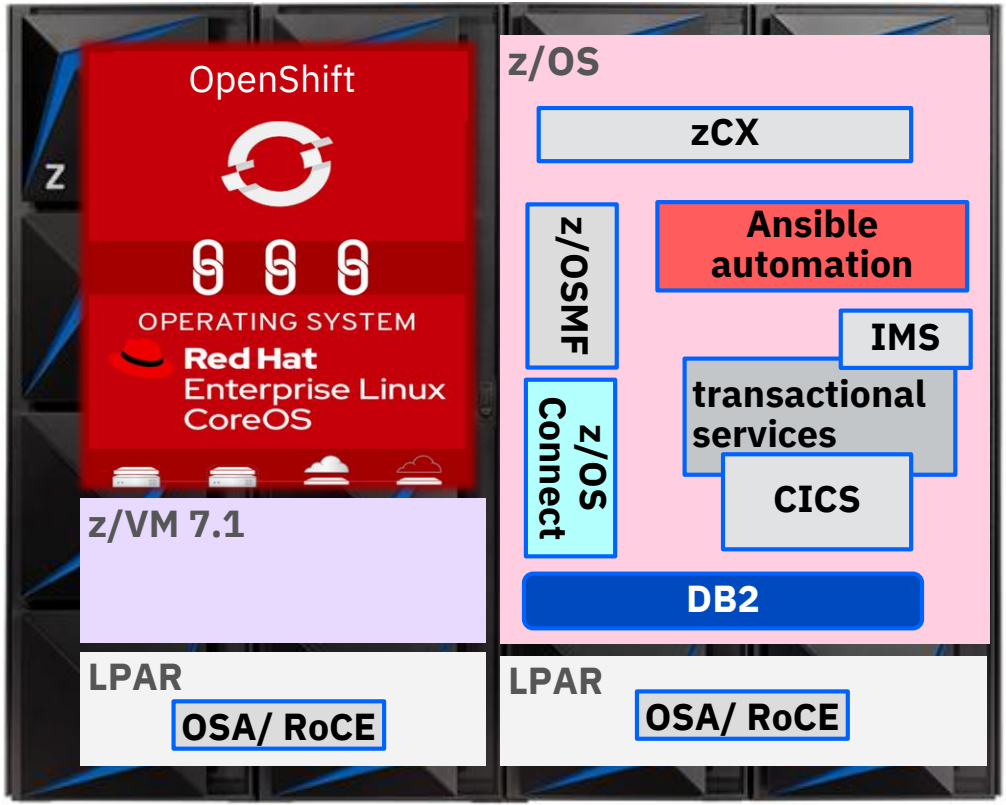

Network options:

- Shared OSA
- Hipersockets (HS) with VSWITCH Bridge (VB)

Red Hat OpenShift Container Platform on IBM Z

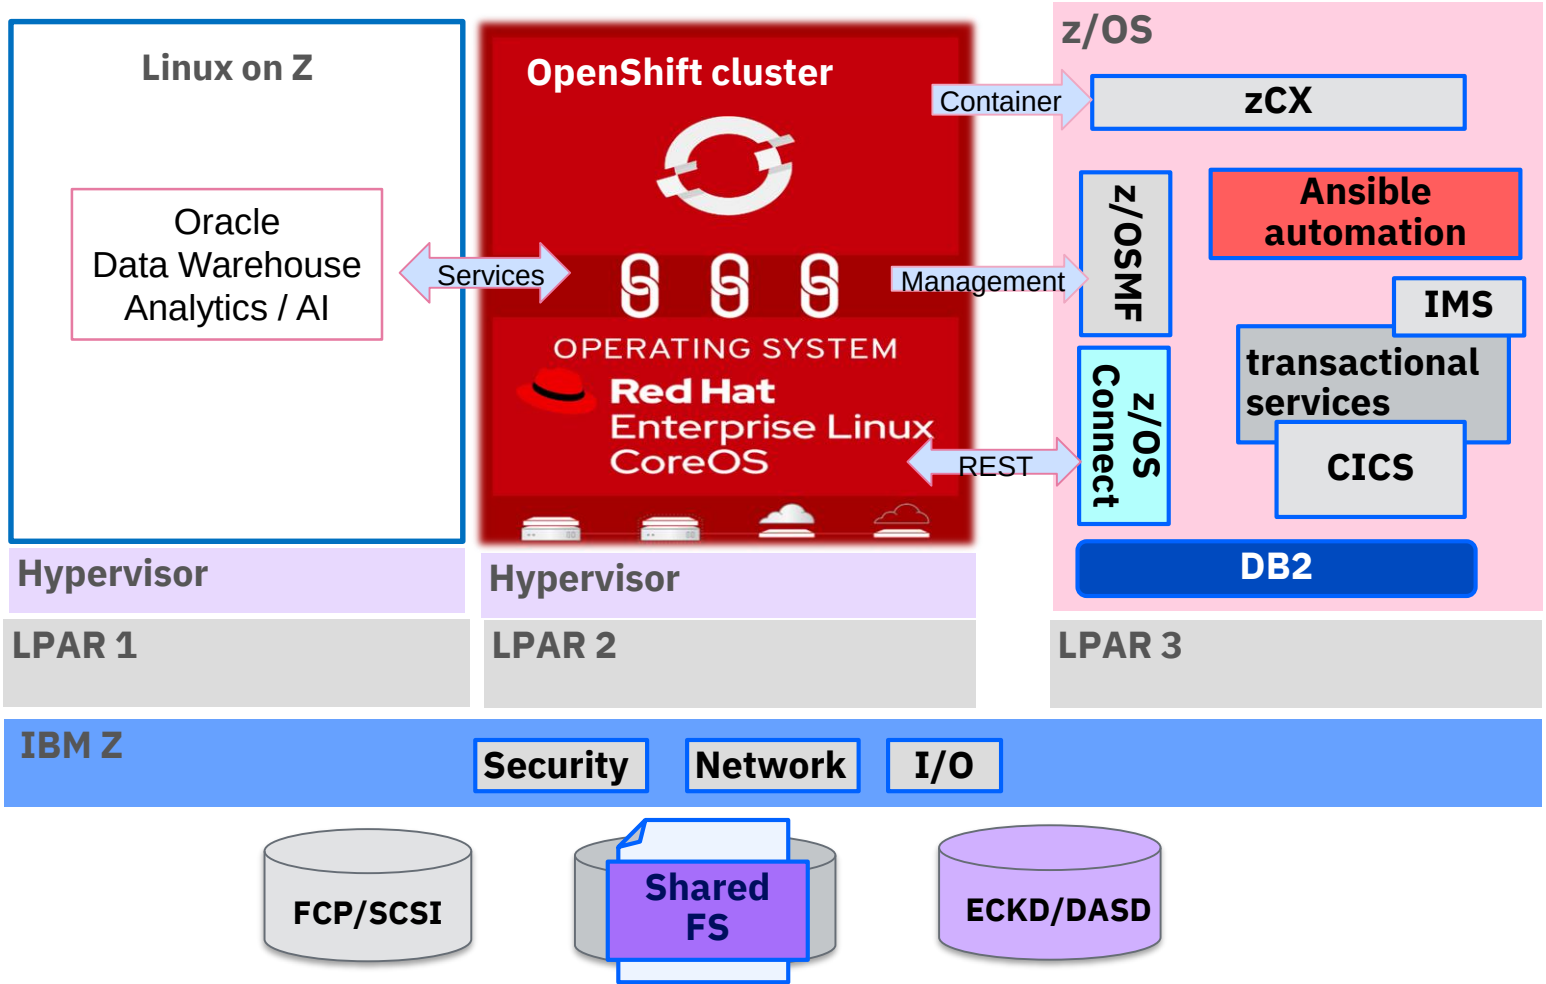

Planning: RHOCP on IBM Z & LinuxONE implementation topology

A) What is the Use Case
- PoC environment
  - *less resources*
- Productive like env.
  - *SLA based,*
  - *HA / DR*

B) What are SLAs
- DevOps integration
  - *automation*
  - *shared content*
- Transactional load
  - *performance*
- HA variants
  - *availability*
  - *resiliency*

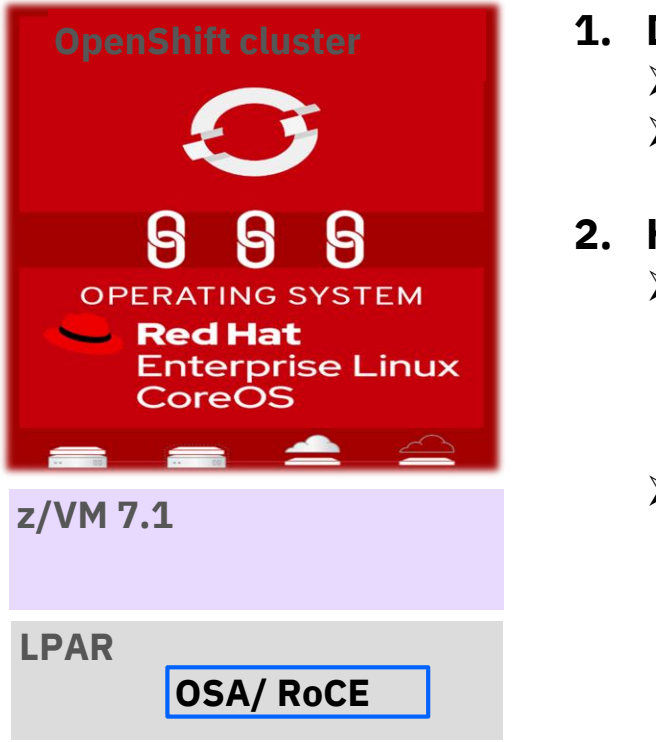

1. Deployment topology
- *RHOCP Standalone*
- *Co-located with z/OS*

2. HW topology
- On one HW machine
  - *One cluster / 1LPAR (PoC)*
  - *multiple LPARs*
- Multiple HW machines
  - *in same DC*
  - *across DC (with synchronous replication only)*

What does Red Hat OpenShift Container Platform on IBM Z & LinuxONE look like?

Sample operational RHOCP on z/VM Layout

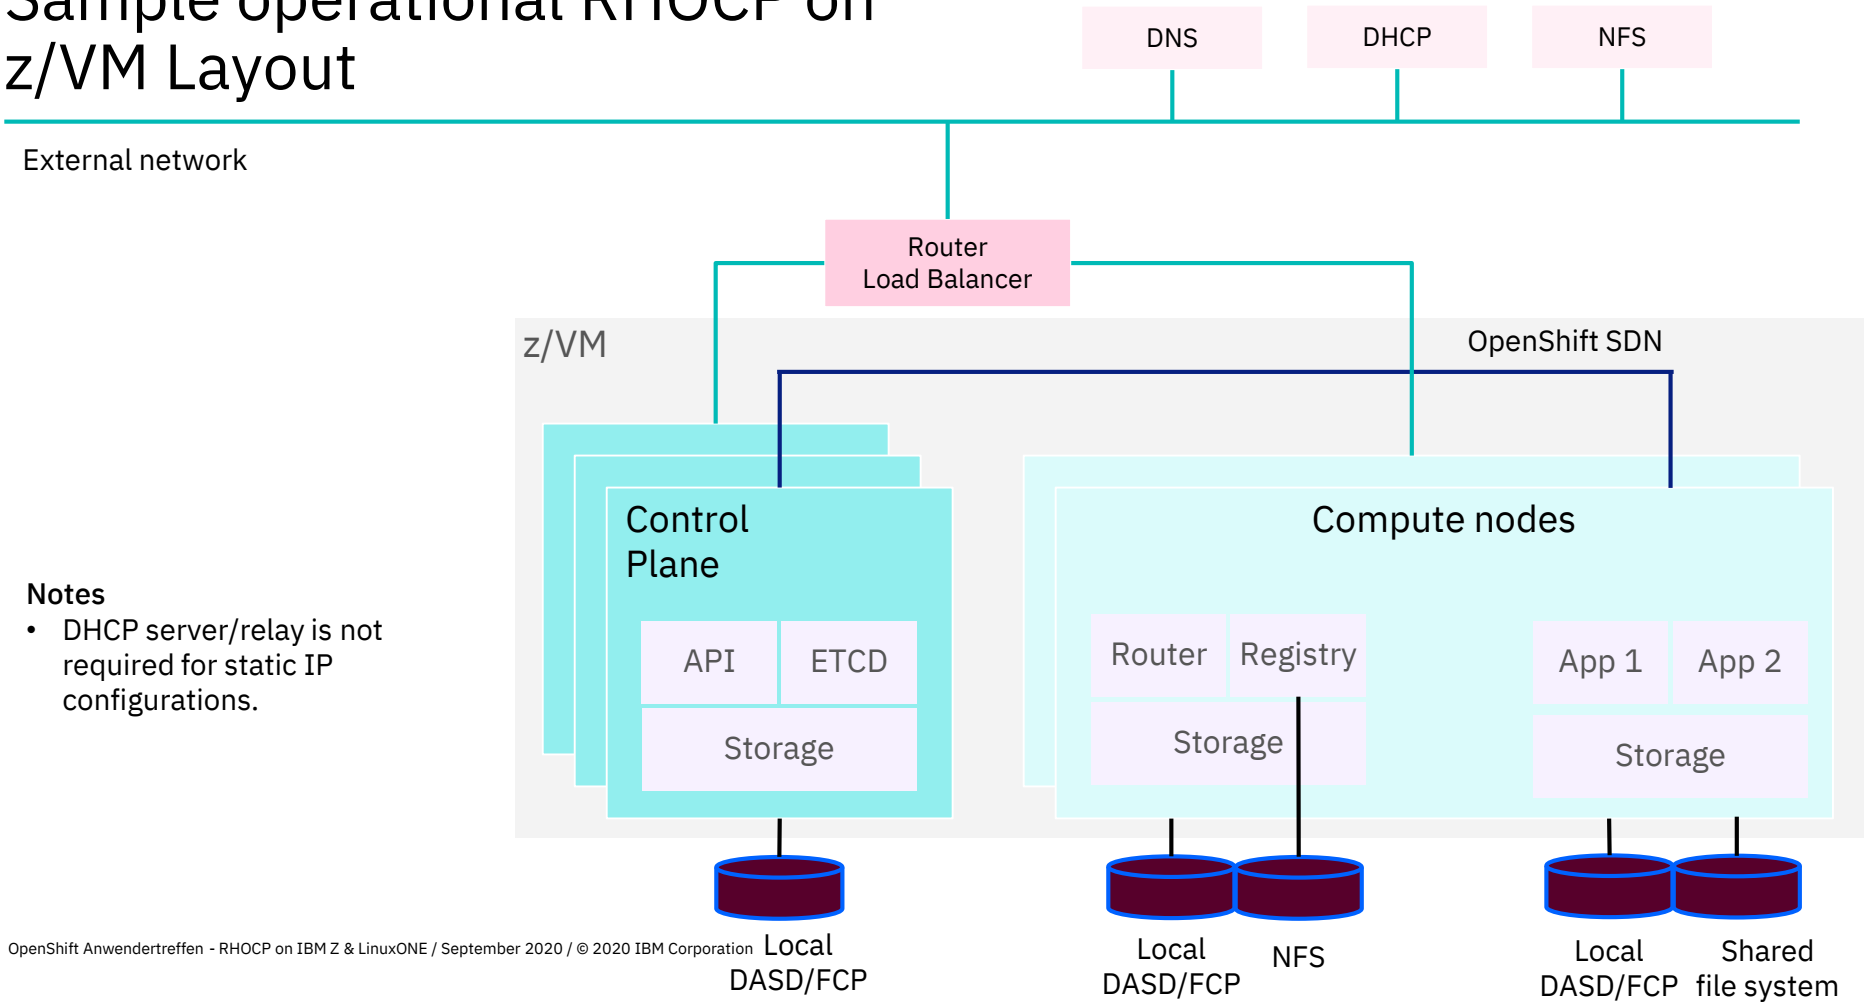

Notes

- DHCP server/relay is not required for static IP configurations.

Technical Fact Sheet for RHOCP 4.x on Z & LinuxONE

Feature Overview

- RHEL CoreOS
- User-provisioned infrastructure (UPI)
- Disconnected / air-gapped installation
- Shared persistent storage with NFS or IBM Spectrum Scale (Beta)

Minimum System Requirements

- IBM z13/z13s and later, and IBM LinuxONE
- 1 LPAR with z/VM 7.1 using 3 IFLs, 80+GB
- FICON or FCP attached disk storage
- OSA, RoCE, z/VM VSwitch networking

Preferred Systems Requirements for High-Availability

- 3 LPARs with z/VM 7.1 using 6 IFLs, 112+GB

What are the hardware and software requirements?

Minimum System Requirements

Hardware Capacity

- 1 LPAR with 3 IFLs supporting SMT2
- 1 OSA and/or RoCE card

Operating System

- z/VM 7.1
  - 3 VMs for OCP Control Plane Nodes
  - 2 VMs for OCP Compute Nodes
  - 1 VM for temporary OCP Bootstrap Node

Disk storage

- FICON attached disk storage (DASDs)
  - Minidisks, fullpack minidisks, or dedicated DASDs
- FCP attached disk storage

Network

- Single z/VM virtual NIC in layer 2 mode, one of
  - Direct-attached OSA, HiperSockets, or RoCE
  - z/VM VSwitch

Memory

- 16GB for OCP Control Plane Nodes
- 8GB for OCP Compute Nodes
- 16GB for OCP Bootstrap Node (temporary)

Preferred Architecture Overview

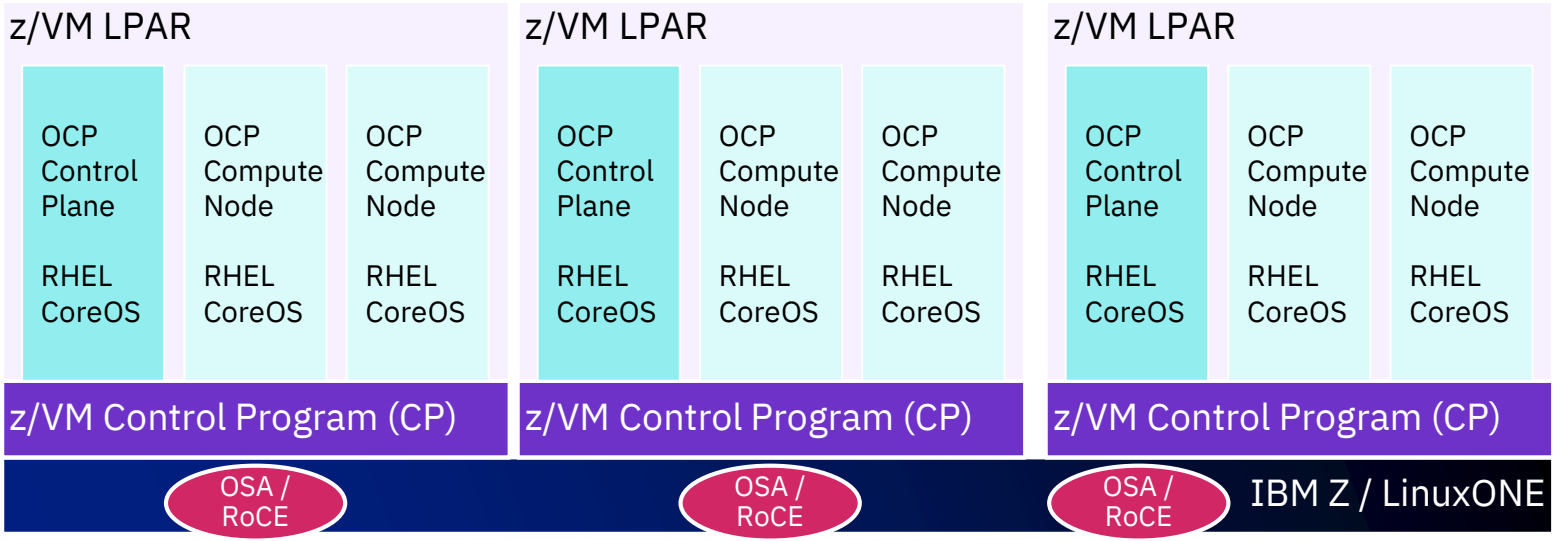

Notes

- Distribute OCP control planes to different z/VM instances on one or more IBM Z / LinuxONE servers to achieve HA and cover service outages/windows

Preferred System Requirements

Hardware Capacity

- 3 LPARs with 6 IFLs supporting SMT2
- 1-2 OSA and/or RoCE card

Operating System

- z/VM 7.1 – 3 instances for HA purposes
  - 3 VMs for OCP Control Planes (one per instance)
  - 6+ VMs for OCP Compute Nodes (across instances)
  - 1 VM for temporary OCP Bootstrap Node

Disk storage

- FICON attached disk storage (DASDs)
  - Minidisks, fullpack minidisks, or dedicated DASDs
- FCP attached disk storage

Network

- Single z/VM virtual NIC in layer 2 mode, one of
  - Direct-attached OSA or RoCE
  - z/VM VSwitch (using OSA link aggregation to increase bandwidth and high availability)

Memory

- 16+ GB for each OCP Control Plane Node
- 8+ GB for each OCP Compute Node
- 16GB for the OCP Bootstrap Node (temporary)
- For sizing details, see also https://docs.openshift.com/container-platform/4.5/scalability_and_performance/recommended-host-practices.html#master-node-sizing_

Software Configuration for OpenShift Container Platform

Infrastructure Services (Pre-requisites)

- DHCP server or static IP addresses
- DNS server
- Load balancers (optional but preferred)
- Deployment server for installation (temporary)
- Internet connectivity

Operating System

- RHEL CoreOS for Control Plane and Bootstrap Nodes
- RHEL CoreOS only for Compute Nodes

Persistent Storage

- NFSv4 server with >100GB disk storage
  - 100GB for internal registry at minimum

Bootstrap and Master Nodes (Control Planes)

- 4 vCPUs
- 16+ GB main memory
- 120GB disk storage

Worker Nodes (+ depending on workload)

- 2+ vCPUs (1+ IFLs with SMT2 enabled)
- 8+GB main memory
- 120GB disk storage

Reference about OCP cluster limits

- https://docs.openshift.com/container-platform/4.5/scalability_and_performance/planning-your-environment-according-to-object-maximums.html

Minimum
Architecture Overview – Network Option 1

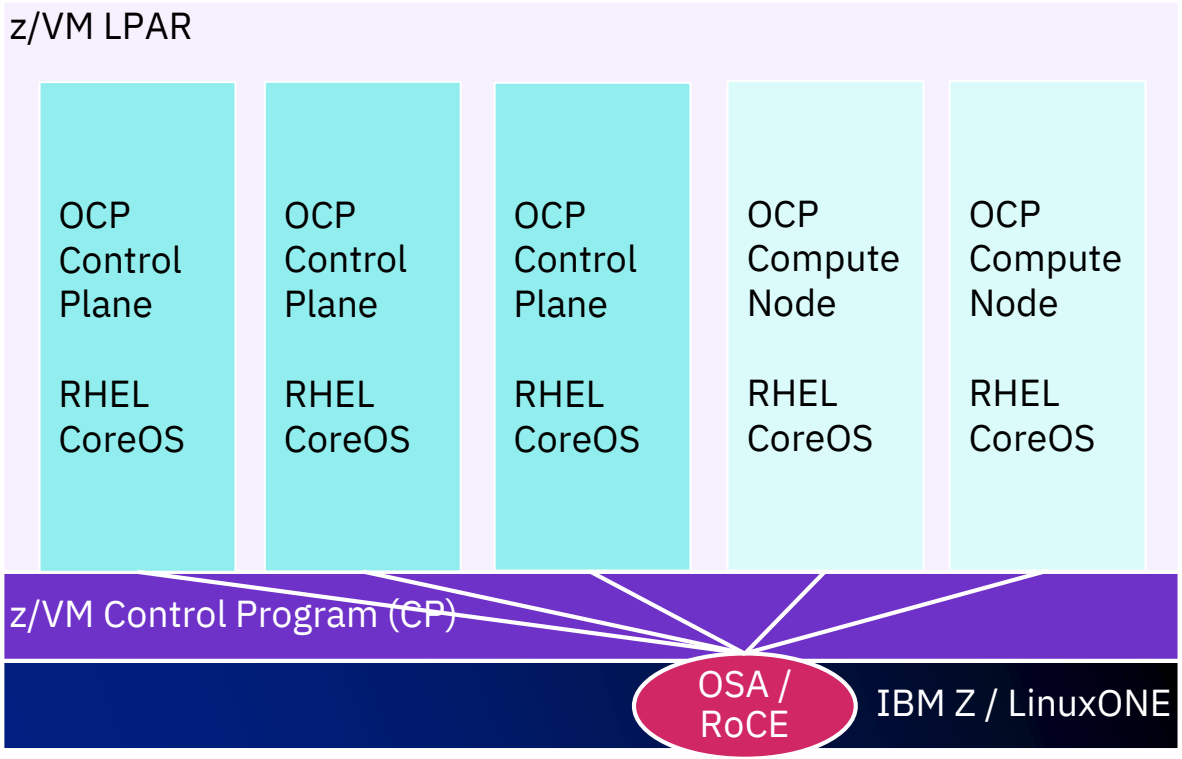

Use single vNIC for z/VM guest virtual machines

- Direct-attached OSA or RoCE to each guest virtual machine

RHOCP uses this 1 vNIC for two networks

- External communication
- Internal communication – software-defined network for Kubernetes pod communication

Minimum
Architecture Overview – Network Option 2

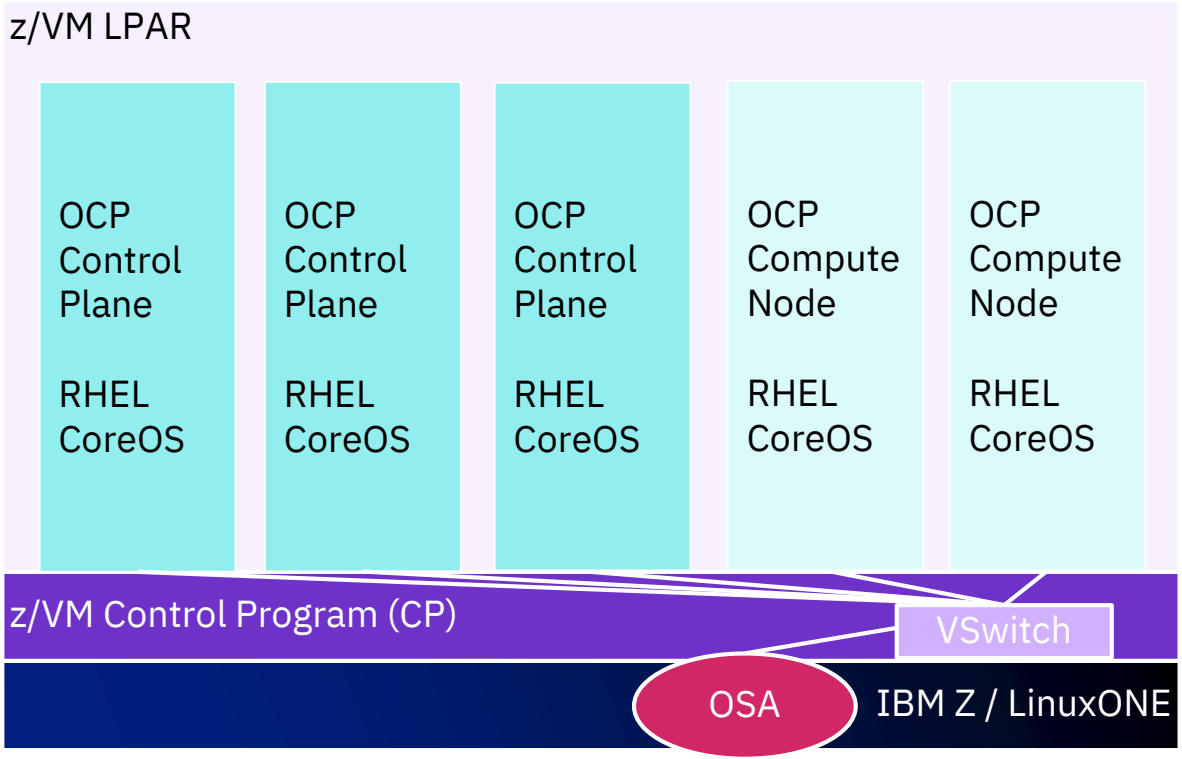

Use single vNIC for z/VM guest virtual machines

- z/VM VSwitch with OSA (optionally, using link aggregation)

RHOCP uses this 1 vNIC for two networks

- External communication
- Internal communication – software-defined network for Kubernetes pod communication

Minimum
Architecture Overview – Disk Storage Options for Installation

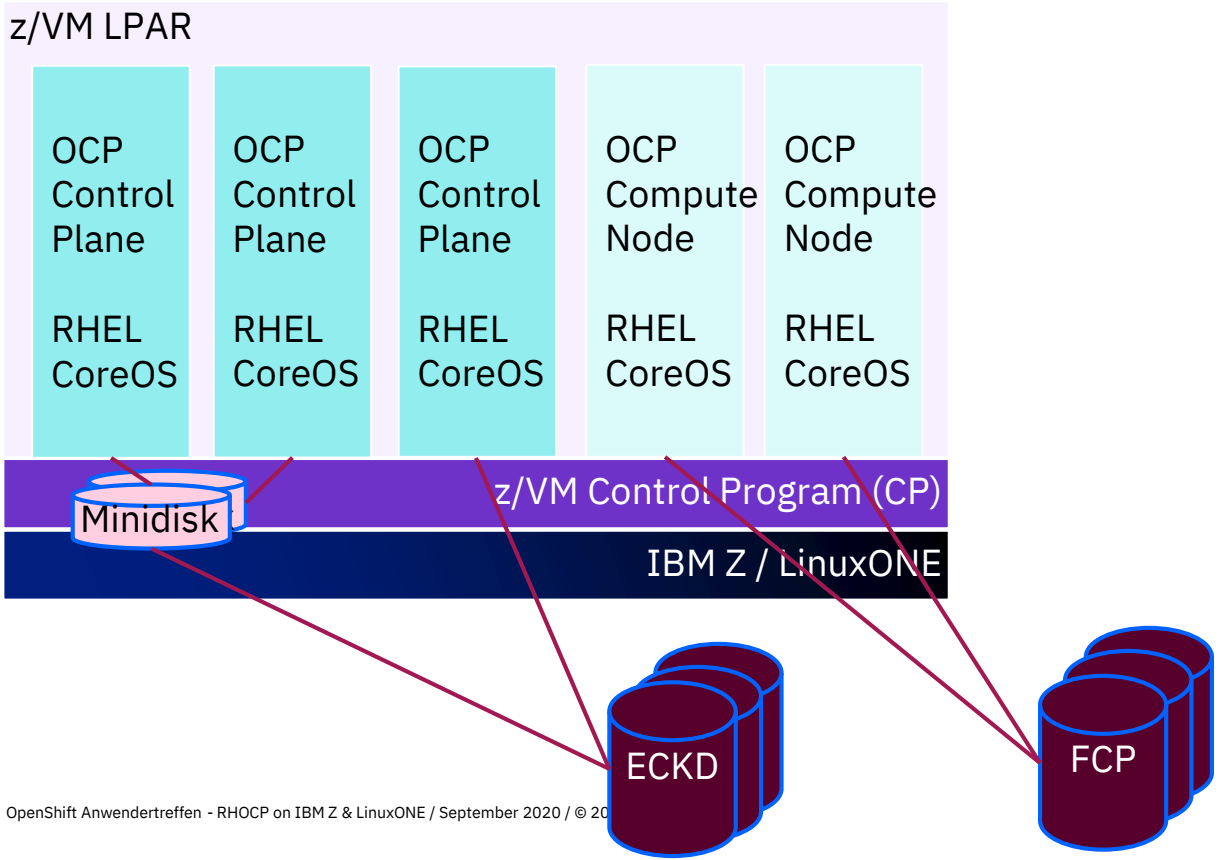

Disk storage considerations

- Minidisks are a z/VM virtual resources and represent smaller chunks on a DASD; Linux sees them as individual disks (DASDs)
- Consider HyperPAV for ECKD storage
- DASDs/FCP devices can be dedicated to a z/VM guest ("pass-through")
- Consider using FCP multipath installations (future)

Shared disk storage considerations

Required shared disk storage

- Internal registry (container images)

Use cases for shared disk storage

- Shared data pool for container instances (persistent container storage)
- Application or workload specific use cases

Shared disk storage options in the initial release

- NFS only

Shared disk storage options in future releases

- IBM Spectrum Scale (Beta available)
  - Register for the *Cloud Native Deployment of Spectrum Scale Beta Program:* https://www-355.ibm.com/technologyconnect/cna/index.xhtml
- IBM CSI Block Plugin 1.1.0
  - https://www.ibm.com/support/knowledgecenter/SSRQ8T_1.1.0/csi_block_storage_kc_welcome.html
    https://access.redhat.com/containers/#/registry.connect.redhat.com/ibm/ibm-block-csi-operator
- Red Hat OpenShift Container Storage

News on RHOCP 4.x Performance on IBM LinuxONE III LT1

OpenShift Container Platform (OCP)

Acme Air Performance on OpenShift Container Platform 4.2 on LinuxONE III LT1 vs. x86 Skylake

Achieve up to 2.7x more throughput per core and up to 2.9x lower latency on OpenShift Container Platform 4.2 on LinuxONE III using z/VM versus on compared x86 platform using KVM, when running 12 Acme Air benchmark instances on 3 worker nodes

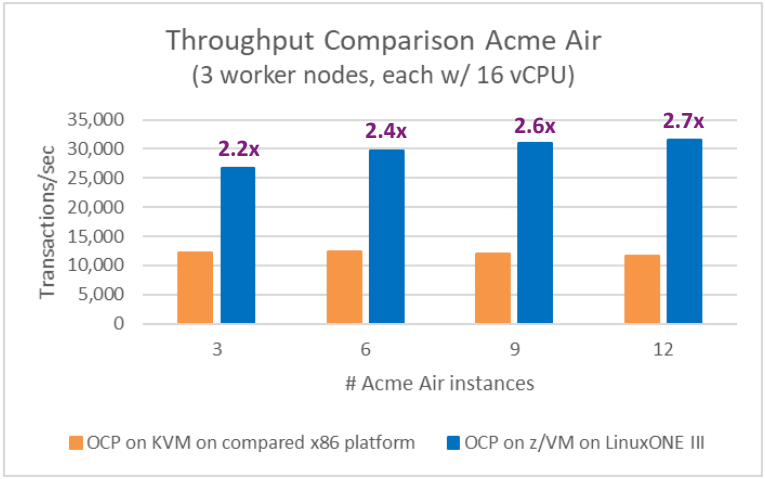

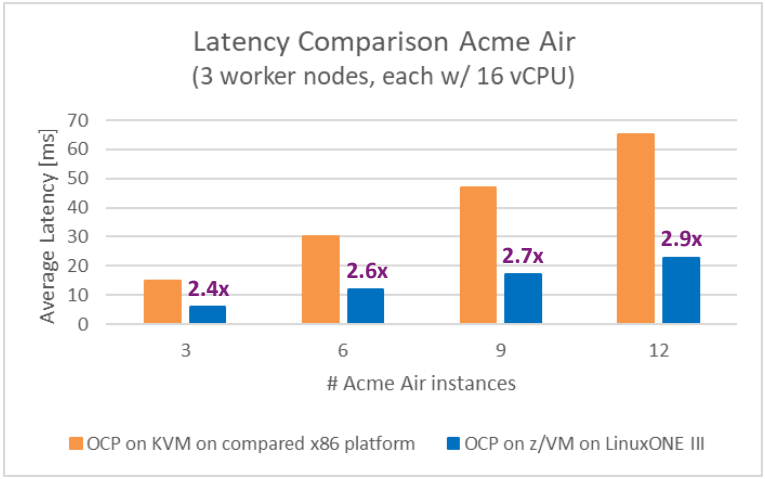

DISCLAIMER: Performance results based on IBM internal tests running the Acme Air microservice benchmark (https://github.com/blueperf/acmeair-mainservice-java) on OpenShift Container Platform (OCP) 4.2.19 on LinuxONE III using z/VM versus on compared x86 platform using KVM. On both platforms 12 Acme Air instances were running on 3 OCP Worker nodes. The z/VM and KVM guests with the OCP Master nodes were configured with 4 vCPUs and 16 GB memory each. The z/VM and KVM guests with the OCP Worker nodes were configured with 16 vCPUs and 32 GB memory each. Results may vary. LinuxONE III configuration: The OCP Proxy server ran native LPAR with 4 dedicated cores, 64 GB memory, RHEL 8.1 (SMT mode). The OCP Master and Worker nodes ran on z/VM 7.1 in a LPAR with 30 dedicated cores, 160 GB memory, and DASD storage. x86 configuration: The OCP Proxy server ran on 4 Intel® Xeon® Gold 6126 CPU @ 2.60GHz with Hyperthreading turned on, 64 GB memory, RHEL 8.1. The OCP Master and Worker nodes ran on KVM on RHEL 8.1 on 30 Intel® Xeon® Gold 6140 CPU @ 2.30GHz with Hyperthreading turned on, 160 GB memory, and RAID5 local SSD storage.

Acme Air Density on OpenShift Container Platform 4.2 on LinuxONE III LT1 vs. x86 Skylake

Run up to 2.5x more Acme Air benchmark instances per core on OpenShift Container Platform 4.2 on LinuxONE III using z/VM versus on a compared x86 platform using KVM, each processing an identical transaction load

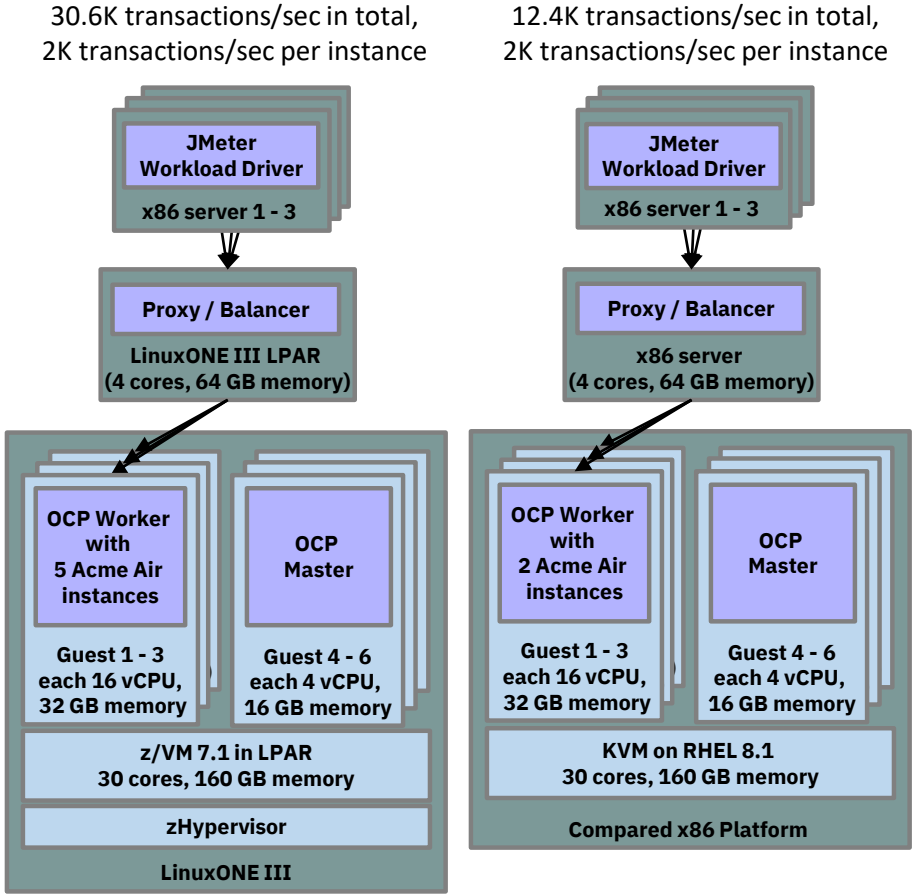

DISCLAIMER: Performance results based on IBM internal tests running the Acme Air microservice benchmark (https://github.com/blueperf/acmeair-mainservice-java) on OpenShift Container Platform (OCP) 4.2.19 on LinuxONE III using z/VM versus on compared x86 platform using KVM. The z/VM and KVM guests with the OCP Master nodes were configured with 4 vCPUs and 16 GB memory each. The z/VM and KVM guests with the OCP Worker nodes were configured with 16 vCPUs and 32 GB memory each. Results may vary. LinuxONE III configuration: The OCP Proxy server ran native LPAR with 4 dedicated cores, 64 GB memory, RHEL 8.1 (SMT mode). The OCP Master and Worker nodes ran on z/VM 7.1 in a LPAR with 30 dedicated cores, 160 GB memory, and DASD storage. x86 configuration: The OCP Proxy server ran on 4 Intel® Xeon® Gold 6126 CPU @ 2.60GHz with Hyperthreading turned on, 64 GB memory, RHEL 8.1. The OCP Master and Worker nodes ran on KVM on RHEL 8.1 on 30 Intel® Xeon® Gold 6140 CPU @ 2.30GHz with Hyperthreading turned on, 160 GB memory, and RAID5 local SSD storage.

OpenShift Container Platform (OCP)

Acme Air Scaling on OpenShift Container Platform 4.2 on LinuxONE III LT1

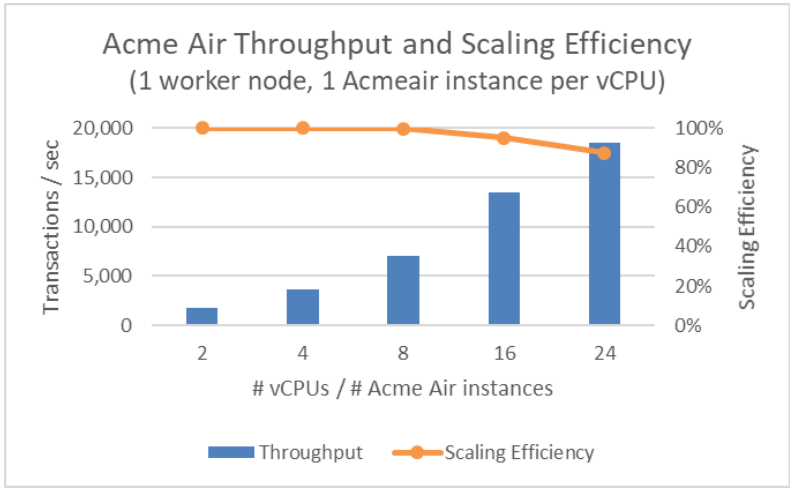

Scale-out the Acme Air benchmark to 24 virtual CPUs with up to 88% scaling efficiency on an OpenShift Container Platform 4.2 worker node on LinuxONE III LT1 using z/VM

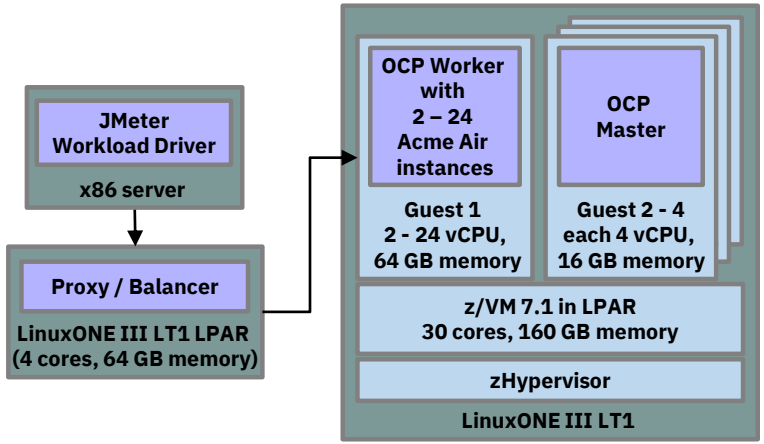

DISCLAIMER: Performance results based on IBM internal tests running the Acme Air microservice benchmark (https://github.com/blueperf/acmeair-mainservice-java) on OpenShift Container Platform (OCP) 4.2.19 on LinuxONE III LT1 using z/VM. The z/VM guests with the OCP Master nodes were configured with 4 vCPUs and 16 GB memory each. The z/VM guest with the OCP Worker node was configured with 2 - 24 vCPUs and 64 GB memory. Per vCPU one Acme Air instance was running on the OCP Worker node. The Acme Air instances were driven remotely from JMeter 5.2.1. Results may vary. LinuxONE III LT1 configuration: The OCP Proxy server ran native LPAR with 4 dedicated cores, 64 GB memory, RHEL 8.1 (SMT mode). The OCP Master and Worker nodes ran on z/VM 7.1 in a LPAR with 30 dedicated cores, 160 GB memory, and DASD storage.

What’s next?

What could be next*...?

Red Hat OpenShift Container Platform

- RHOCP release cadence with x86
- Red Hat OpenShift Service Mesh (istio) to connect, secure, control, and observe services
- Red Hat OpenShift Pipelines and Serverless
- CodeReady Workspaces

Storage Support

- IBM Spectrum Scale
- Red Hat OpenShift Container Storage (OCS) based on Ceph, Rook, and Noobaa

** All statements regarding IBM and Red Hat's future direction and intent are subject to change or withdrawal without notice and represent goals and objectives only.*

Questions?

Thank you

Hendrik Brückner
brueckner@de.ibm.com
Manager Linux on Z Development & Service
Red Hat Partner Engineer for RHEL and RHOCP on Z & LinuxONE

Wilhelm Mild
wilhelm.mild@de.ibm.com
IBM Executive IT Architect
Integration Architectures for Container, Mobile, IBM Z, and Linux

More Information

IBM LinuxONE Community Cloud – Free access to virtual servers on LinuxONE
https://developer.ibm.com/components/ibm-linuxone/gettingstarted/

Red Hat OCP portal
cloud.redhat.com

Install OCP on IBM Z
https://docs.openshift.com/container-platform/4.5/installing/installing_ibm_z/installing-ibm-z.html

Step by step sample installations and environment setup
https://www.openshift.com/blog/installing-ocp-in-a-mainframe-z-series
https://www.openshift.com/blog/red-hat-openshift-installation-process-experiences-on-ibm-z-linuxone

Ross Mauri's Blog
http://www.ibm.com/blogs/systems/red-hat-openshift-now-available-ibm-z-linuxone

IBM Systems Magazine Article
https://ibmsystemsmag.com/01/2020/cutting-edge-ibm-z-innovations

IDC Whitepaper
https://www.ibm.com/it-infrastructure/linuxone/capabilities/linux-containers

Acme Air Performance on OpenShift Container Platform 4.2 on LinuxONE III LT2 vs. x86 Skylake

Achieve up to 2.4x more throughput per core and up to 2.5x lower latency on OpenShift Container Platform 4.2 on LinuxONE III LT2 using z/VM versus on compared x86 platform using KVM, when running 12 Acme Air benchmark instances on 3 worker nodes

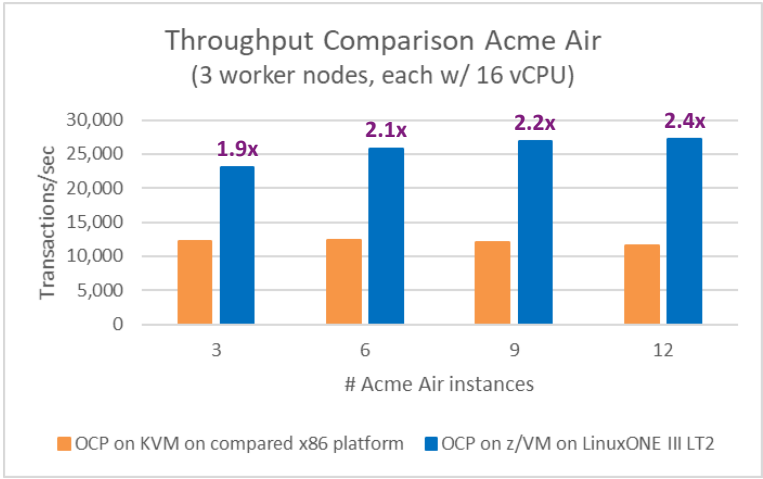

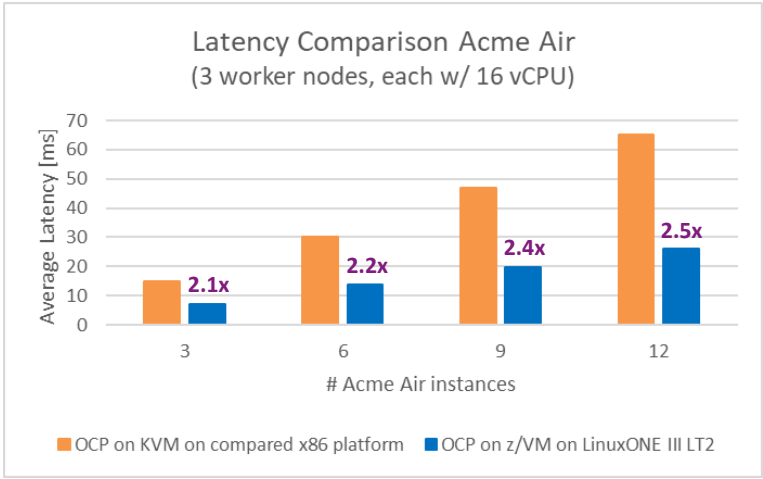

DISCLAIMER: Performance result is extrapolated based on a clock ratio of 0.8654 from IBM internal tests running the Acme Air microservice benchmark (https://github.com/blueperf/acmeair-mainservice-java) on OpenShift Container Platform (OCP) 4.2.19 on LinuxONE III LT1 using z/VM versus on compared x86 platform using KVM. On both platforms 12 Acme Air instances were running on 3 OCP Worker nodes. The z/VM and KVM guests with the OCP Master nodes were configured with 4 vCPUs and 16 GB memory each. The z/VM and KVM guests with the OCP Worker nodes were configured with 16 vCPUs and 32 GB memory each. Results may vary. LinuxONE III LT1 configuration: The OCP Proxy server ran native LPAR with 4 dedicated cores, 64 GB memory, RHEL 8.1 (SMT mode). The OCP Master and Worker nodes ran on z/VM 7.1 in a LPAR with 30 dedicated cores, 160 GB memory, and DASD storage. x86 configuration: The OCP Proxy server ran on 4 Skylake Intel® Xeon® Gold CPU @ 2.60GHz with Hyperthreading turned on, 64 GB memory, RHEL 8.1. The OCP Master and Worker nodes ran on KVM on RHEL 8.1 on 30 Skylake Intel® Xeon® Gold CPU @ 2.30GHz with Hyperthreading turned on, 160 GB memory, and RAID5 local SSD storage.

Acme Air Container Density on OpenShift Container Platform 4.2 on LinuxONE III LT2 versus x86 Skylake

Run up to 2x more Acme Air benchmark instances per core on OpenShift Container Platform 4.2 on LinuxONE III LT2 using z/VM versus on a compared x86 platform using KVM, each processing an identical transaction load

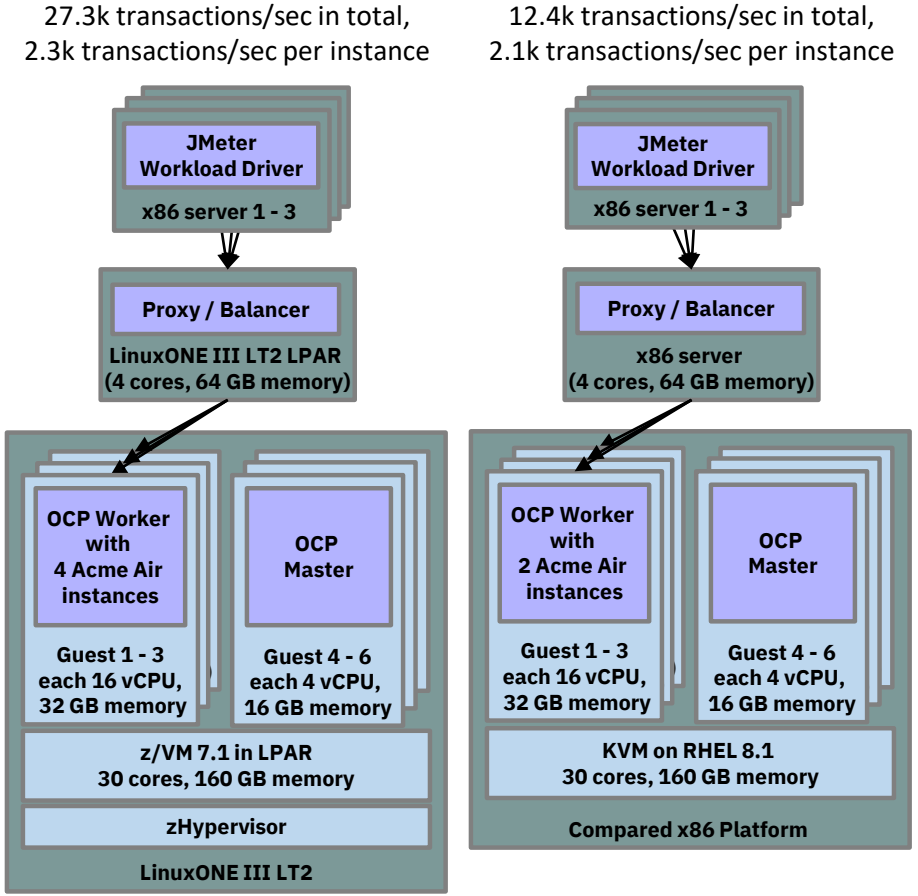

DISCLAIMER: Performance result is extrapolated based on a clock ratio of 0.8654 from IBM internal tests running the Acme Air microservice benchmark (https://github.com/blueperf/acmeair-mainservice-java) on OpenShift Container Platform (OCP) 4.2.19 on LinuxONE III LT1 using z/VM versus on compared x86 platform using KVM. The z/VM and KVM guests with the OCP Master nodes were configured with 4 vCPUs and 16 GB memory each. The z/VM and KVM guests with the OCP Worker nodes were configured with 16 vCPUs and 32 GB memory each. Results may vary. LinuxONE III LT1 configuration: The OCP Proxy server ran native LPAR with 4 dedicated cores, 64 GB memory, RHEL 8.1 (SMT mode). The OCP Master and Worker nodes ran on z/VM 7.1 in a LPAR with 30 dedicated cores, 160 GB memory, and DASD storage. x86 configuration: The OCP Proxy server ran on 4 Skylake Intel® Xeon® Gold CPU @ 2.60GHz with Hyperthreading turned on, 64 GB memory, RHEL 8.1. The OCP Master and Worker nodes ran on KVM on RHEL 8.1 on 30 Skylake Intel® Xeon® Gold CPU @ 2.30GHz with Hyperthreading turned on, 160 GB memory, and RAID5 local SSD storage.

OpenShift Container Platform (OCP)

Acme Air Scaling on OpenShift Container Platform 4.2 on LinuxONE III LT2

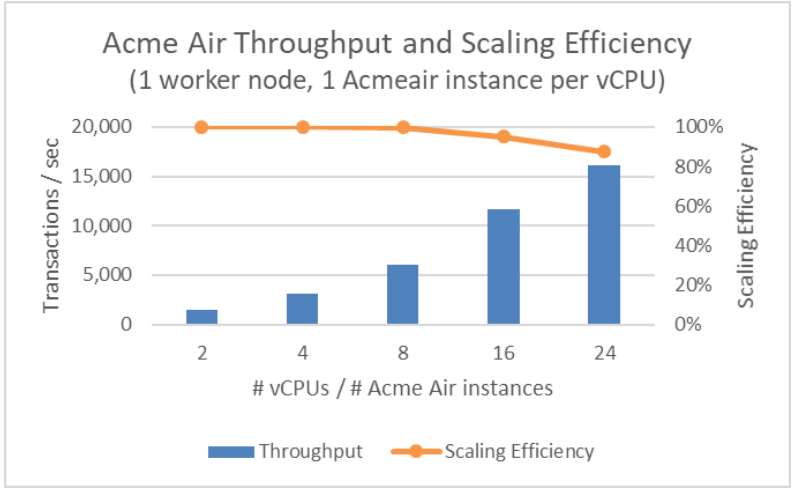

Scale-out the Acme Air benchmark to 24 virtual CPUs with up to 88% scaling efficiency on an OpenShift Container Platform 4.2 worker node on LinuxONE III LT2 using z/VM

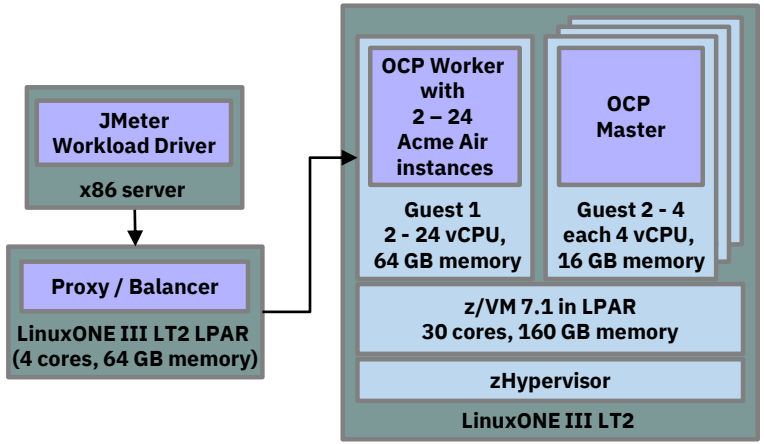

DISCLAIMER: Performance result is extrapolated based on a clock ratio of 0.8654 from IBM internal tests running the Acme Air microservice benchmark (https://github.com/blueperf/acmeair-mainservice-java) on OpenShift Container Platform (OCP) 4.2.19 on LinuxONE III LT1 using z/VM. The z/VM guests with the OCP Master nodes were configured with 4 vCPUs and 16 GB memory each. The z/VM guest with the OCP Worker node was configured with 2 - 24 vCPUs and 64 GB memory. Per vCPU one Acme Air instance was running on the OCP Worker node. The Acme Air instances were driven remotely from JMeter 5.2.1. Results may vary. LinuxONE III LT1 configuration: The OCP Proxy server ran native LPAR with 4 dedicated cores, 64 GB memory, RHEL 8.1 (SMT mode). The OCP Master and Worker nodes ran on z/VM 7.1 in a LPAR with 30 dedicated cores, 160 GB memory, and DASD storage.

Acme Air Performance on OpenShift Container Platform 4.2 on **LinuxONE III LT1** vs. x86 Skylake

Benchmark Setup

- 3 OpenShift Container Platform (OCP) Master and 3 Worker nodes on LinuxONE III under z/VM versus on x86 under KVM
- Acme Air microservice benchmark (https://github.com/blueperf/acmeair-mainservice-java) instances placed manually on the OCP Worker nodes such that each OCP Worker node ran the same number of instances
- Acme Air instances were driven remotely from 3 x86 servers running JMeter 5.2.1

System Stack

- LinuxONE III
  - LPAR with 4 dedicated cores, 64 GB memory, RHEL 8.1 (SMT mode), running the OCP Proxy server
  - LPAR with 30 dedicated cores, 160 GB memory, DASD storage, running z/VM 7.1
    - 3 guests with 4 vCPU, 16 GB memory, each running an OCP Master
    - 3 guests with 16 vCPUs, 32 GB memory, each running an OCP Worker
  - OpenShift Container Platform (OCP) 4.2.19
- x86
  - 4 Intel® Xeon® Gold 6126 CPU @ 2.60GHz w/ Hyperthreading turned on, 64 GB memory, RHEL 8.1, running the OCP Proxy server
  - 30 Intel® Xeon® Gold 6140 CPU @ 2.30GHz w/ Hyperthreading turned on, 160 GB memory, running KVM on RHEL 8.1
    - 3 guests with 4 vCPU, 16 GB memory, each running an OCP Master
    - 3 guests with 16 vCPUs, 32 GB memory, each running a OCP Worker
  - OpenShift Container Platform (OCP) 4.2.19

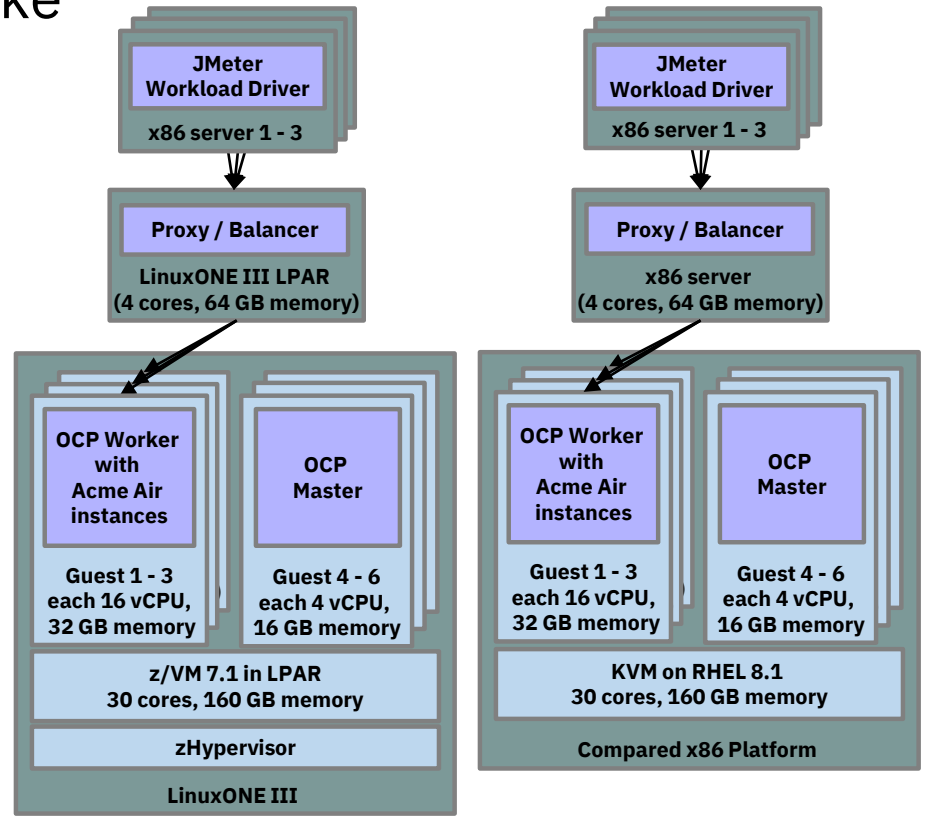

Acme Air Performance on OpenShift Container Platform 4.2 on **LinuxONE III LT2** vs. x86 Skylake

Benchmark Setup

- 3 OpenShift Container Platform (OCP) Master and 3 Worker nodes on LinuxONE III LT2 under z/VM versus on x86 under KVM
- Acme Air microservice benchmark (https://github.com/blueperf/acmeair-mainservice-java) instances placed manually on the OCP Worker nodes such that each OCP Worker node ran the same number of instances
- Acme Air instances were driven remotely from 3 x86 servers running JMeter 5.2.1

System Stack

- LinuxONE III LT2
  - LPAR with 4 dedicated cores, 64 GB memory, RHEL 8.1 (SMT mode), running the OCP Proxy server
  - LPAR with 30 dedicated cores, 160 GB memory, DASD storage, running z/VM 7.1
    - 3 guests with 4 vCPU, 16 GB memory, each running an OCP Master
    - 3 guests with 16 vCPUs, 32 GB memory, each running an OCP Worker
  - OpenShift Container Platform (OCP) 4.2.19
- x86
  - 4 Skylake Intel® Xeon® Gold CPU @ 2.60GHz w/ Hyperthreading turned on, 64 GB memory, RHEL 8.1, running the OCP Proxy server
  - 30 Skylake Intel® Xeon® Gold CPU @ 2.30GHz w/ Hyperthreading turned on, 160 GB memory, running KVM on RHEL 8.1
    - 3 guests with 4 vCPU, 16 GB memory, each running an OCP Master
    - 3 guests with 16 vCPUs, 32 GB memory, each running a OCP Worker
  - OpenShift Container Platform (OCP) 4.2.19

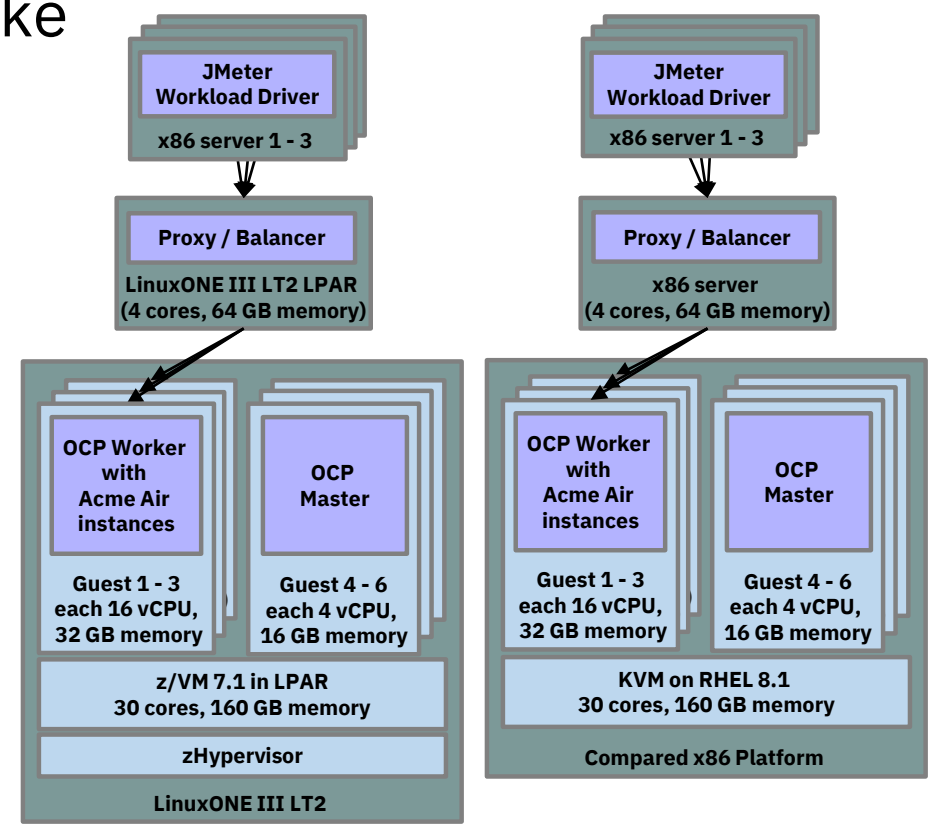

Minimum System Requirements

Hardware Capacity

- 1 LPAR with 3 IFLs supporting SMT2
- 1 OSA and/or RoCE card

Operating System

- z/VM 7.1
  - 3 VMs for OCP Control Plane Nodes
  - 2 VMs for OCP Compute Nodes
  - 1 VM for temporary OCP Bootstrap Node

Disk storage

- FICON attached disk storage (DASDs)
  - Minidisks, fullpack minidisks, or dedicated DASDs
- FCP attached disk storage

Network

- Single z/VM virtual NIC in layer 2 mode, one of
  - Direct-attached OSA, HiperSockets, or RoCE
  - z/VM VSwitch

Memory

- 16GB for OCP Control Plane Nodes
- 8GB for OCP Compute Nodes
- 16GB for OCP Bootstrap Node (temporary)