

Why virtualization is still highly used in the era of containers and cloud on IBM Z & LinuxONE

Wilhelm Mild

IBM Executive IT Architect for Containerization,
Mobile, IBM Z and Linux
IBM R&D Lab, Germany

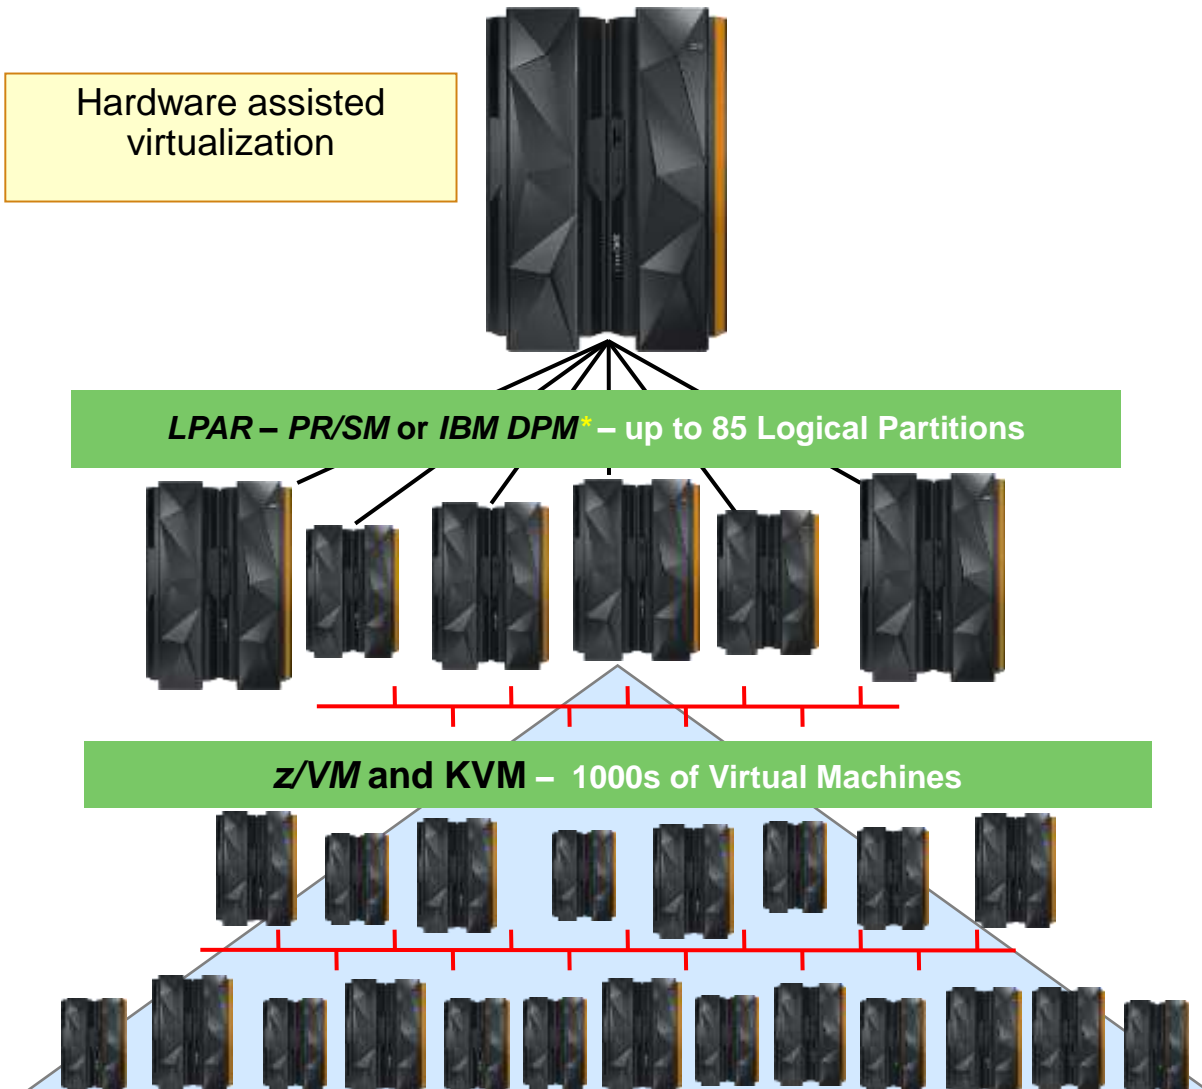


Major Waves of Technology (or, the more things change ...)



IBM Z and LinuxONE Virtualization

Built-in, Shared Everything Architecture

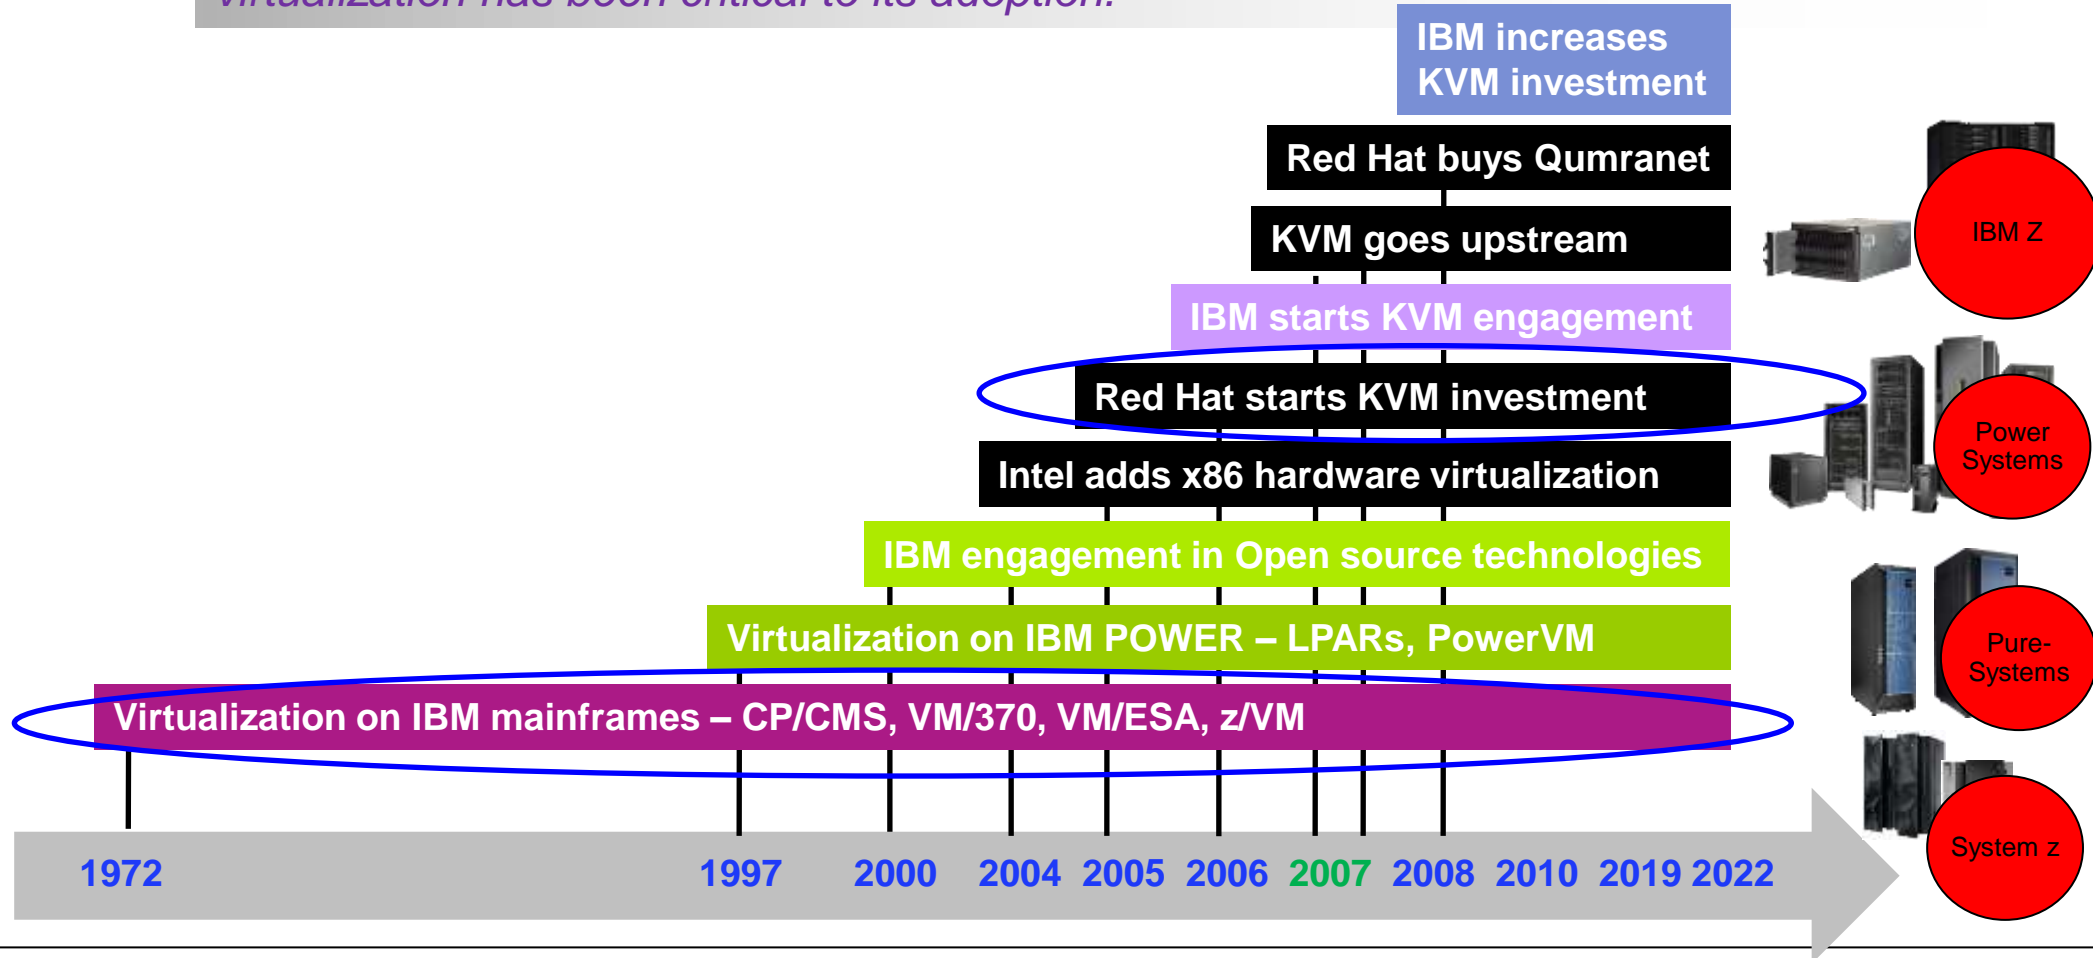


IBM® Z & LinuxONE™ Systems

- Cores are designed to run at near 100% utilization nearly 100% of the time
- Provisioning of virtual servers in seconds
- High granularity of resource sharing (<1%)
- Granular shifting resources to VMs
- Upgrade of physical resources without taking the system down
- Scalability of up to 1000's of virtual servers
- More with less: more virtual servers per core, sharing of physical resources
- Extensive life-cycle management
- HW-supported isolation, highly secure (EAL5+ on LPAR or EAL4+ certified)

A Brief History of Virtualization that IBM is engaged in

IBM has over 50 years of experience in virtualizing our servers. Virtualization was originally developed to make better use of critical hardware. Hardware support for virtualization has been critical to its adoption.

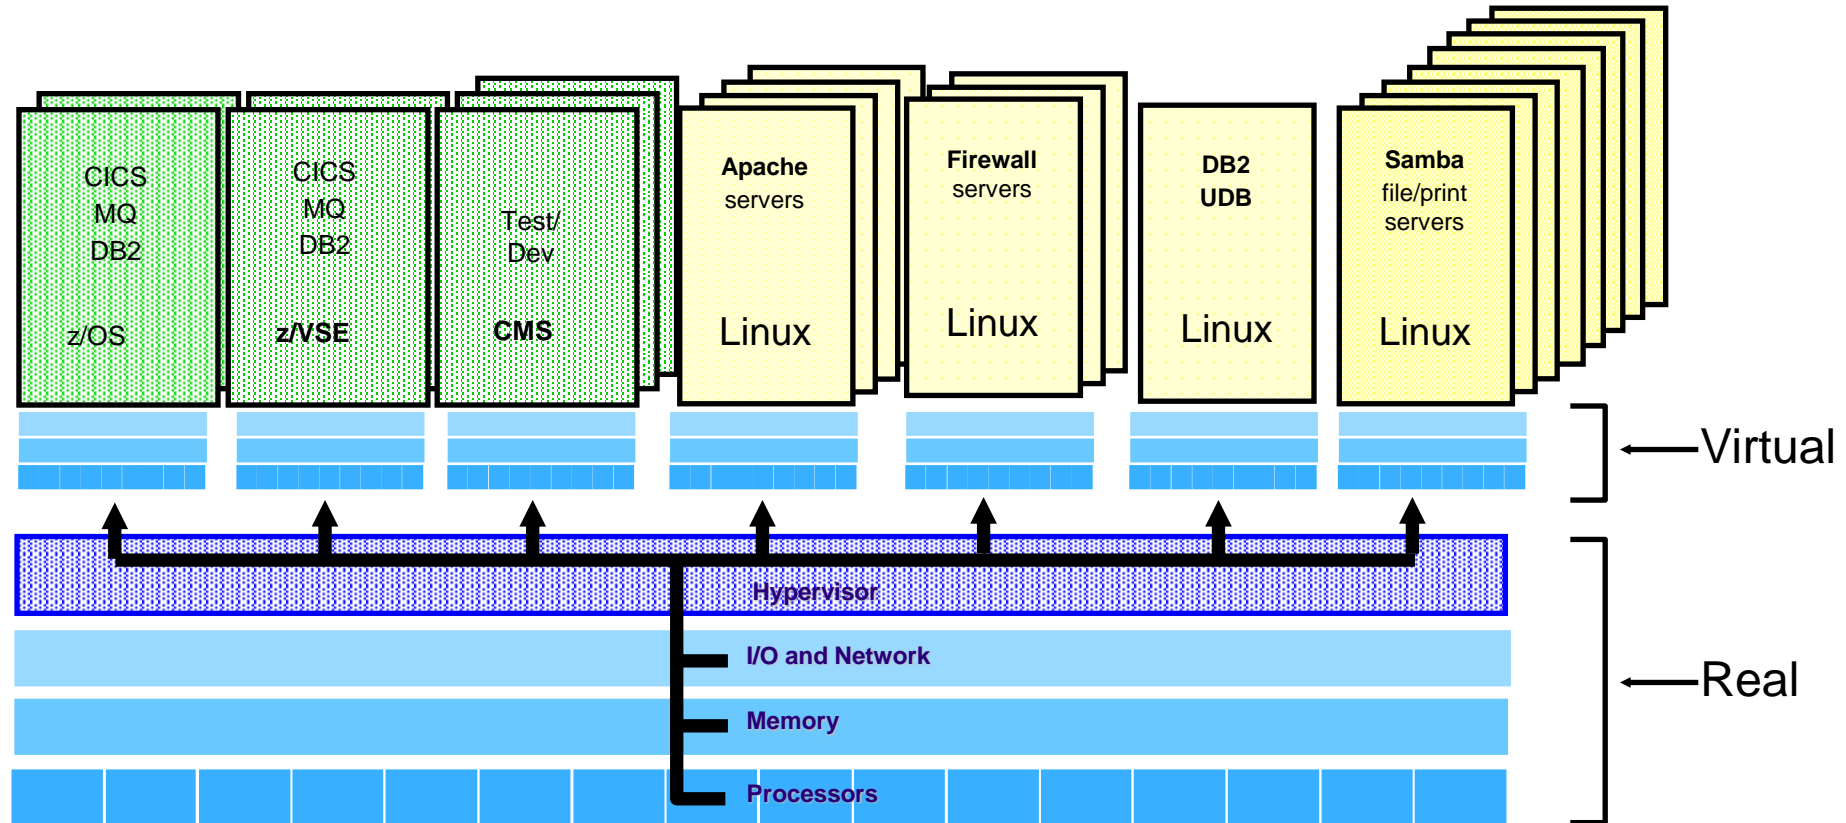


Why do we virtualize? What are the benefits of virtualization?

- **Simplification** – use of standardized images, virtualized hardware, and automated configuration of virtual infrastructure
- **Migration** – one of the first uses of virtualization, enable coexistence, phased upgrades and life migrations. It can also simplify hardware upgrades by making changes transparent.
- **Efficiency** – reduced hardware footprints, better utilization of available hardware resources, and reduced time to delivery. Reuse of deprovisioned or released resources.
- **Resilience** – run new versions and old versions in parallel, avoiding service downtime
- **Cost savings** – having fewer machines translates to lower costs in server hardware, networking, floor space, electricity, administration, and more effective failover
- **To accommodate growth** – virtualization allows the IT department to be more responsive to business growth, in most cases avoiding interruption


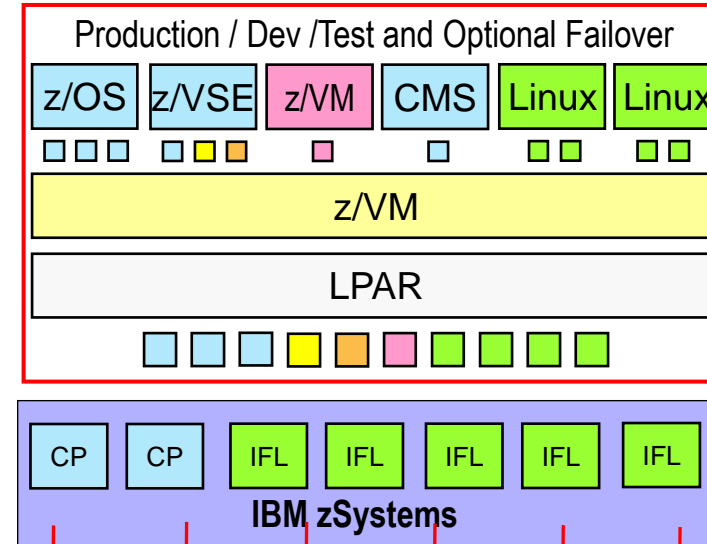

Virtualization Technology on zSystems - more than partitioning

- A *Virtual Machine* simulates the existence of a dedicated real machine, including processor functions, storage, and input/output resources.



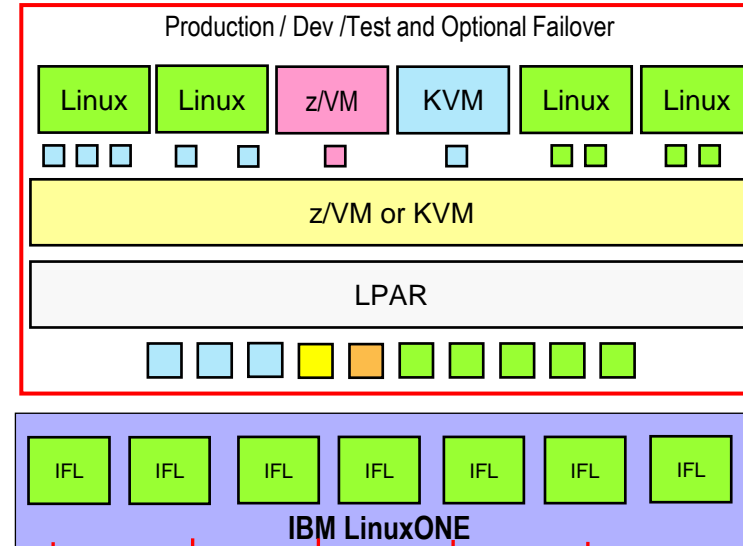

- IBM Z provides the unique capability to run hundreds of Virtual servers on one machine
- Resource sharing and virtualization are key features and provide unmatched flexibility
- Accounting of virtual resources (CPU, I/O, Network,...)

Global Virtualization – with z/VM on IBM zSystems

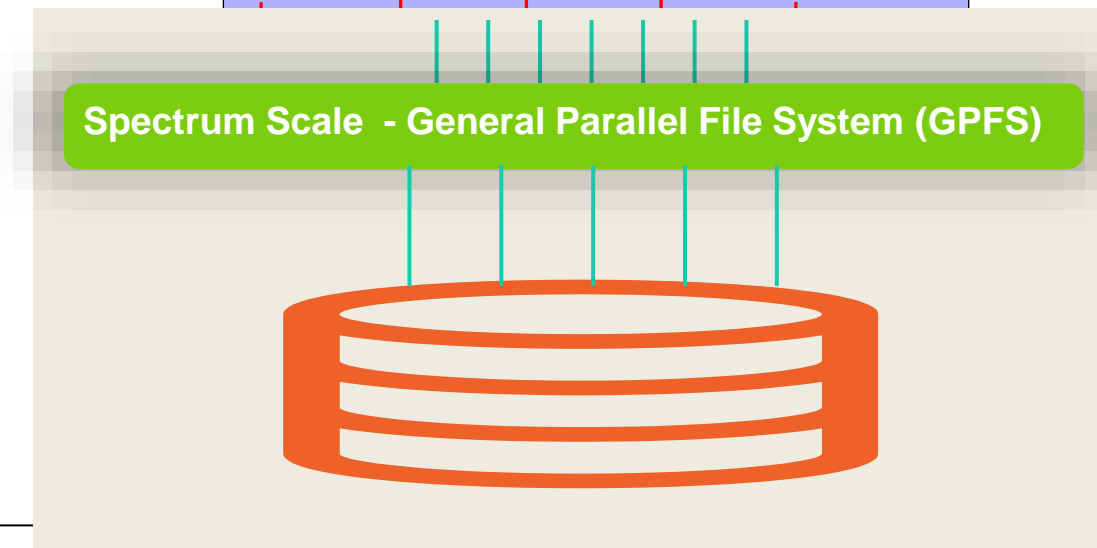


- Network Virtualization
- Memory Virtualization
- Processor Virtualization
- System Virtualization
- Mixed virtualization
- Disk Virtualization

IBM Global Virtualization with Linux on IBM zSystems & LinuxONE




- Network Virtualization
- Memory Virtualization
- Processor Virtualization
- System Virtualization
- Disk Virtualization / Cluster




IBM Virtualization on IBM zSystems and LinuxONE


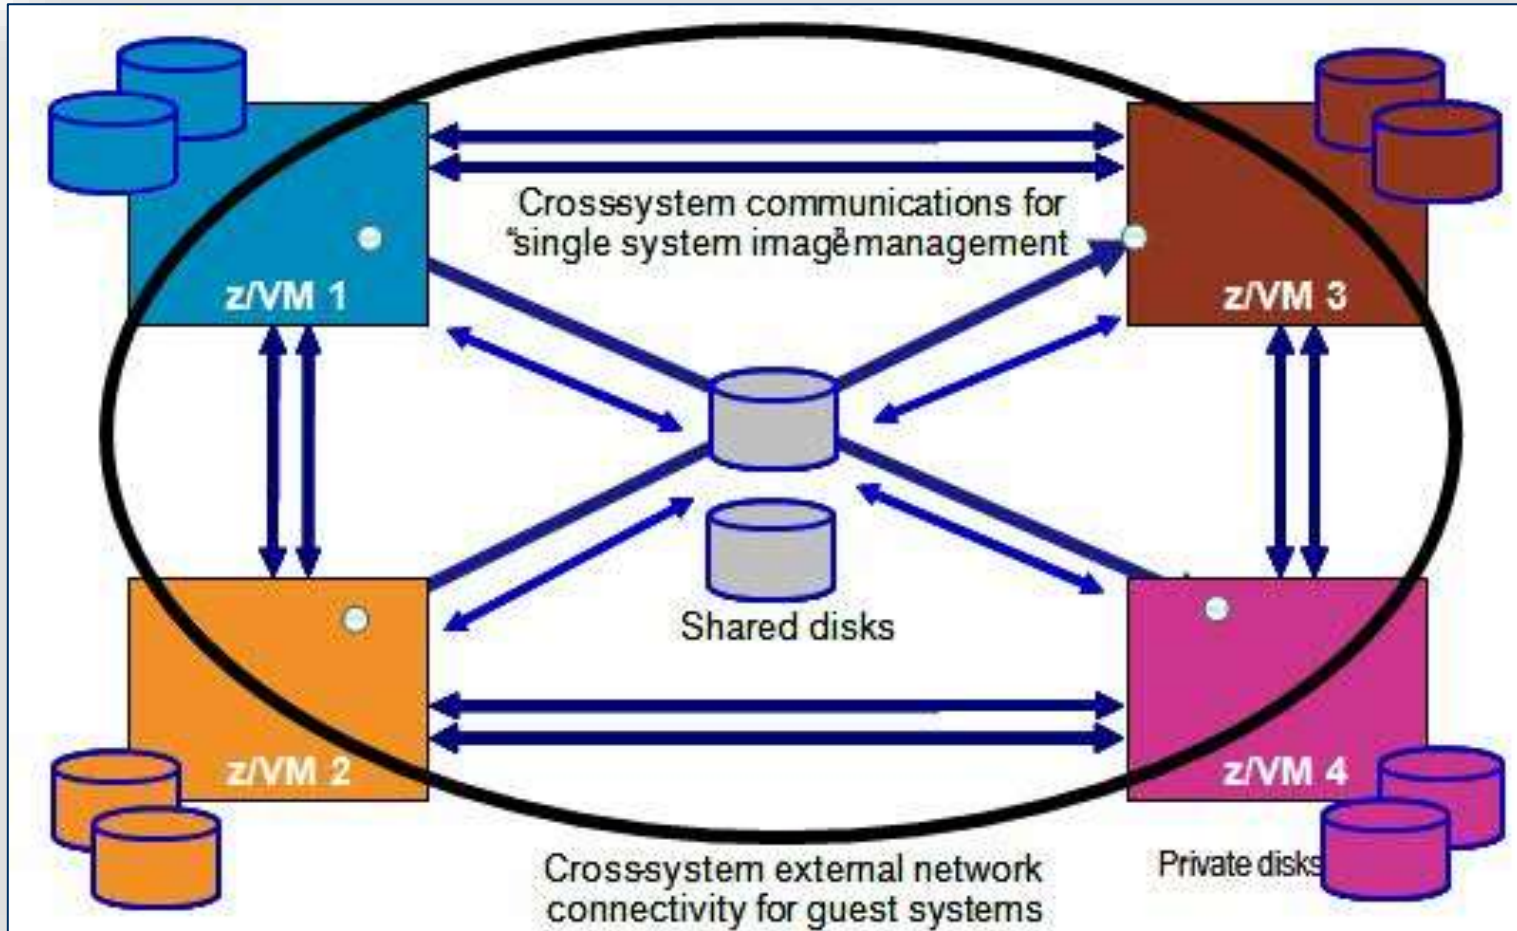
- **Vertical virtualization** - Grow workloads without growing number of virtual guest machines
 - one guest can be increased by allocating more resources (CPUs, memory)
- **Horizontal virtualization** – for isolation between servers
 - isolation of guests in a network
 - High availability for applications
- **3D Virtualization** – combine vertical and horizontal virtualization
- **Dynamically** add, remove and shift physical resources to optimize business results



Penguins got proud ! And started multiply like Rabbits



And z/VM keeps growing by adapting



Which approach is designed for the higher volume of traffic? Which road is faster?


** Illustrative numbers only*

Containers in Linux – for application isolation – competition for virtualization ?

- linuxcontainers.org is the umbrella project behind Linux Containers (LXC), Linux Container management (LXD), Linux Container FileSystem (LXCFS) and Linux cgroup manager daemon (CGManager).
- The main focus is system containers, that offer an environment as close as possible as the one you'd get from a VM but without the overhead that comes with running a separate kernel and simulating all the hardware.
This is achieved through a combination of kernel security features such as namespaces, mandatory access control and control groups (cgroups).

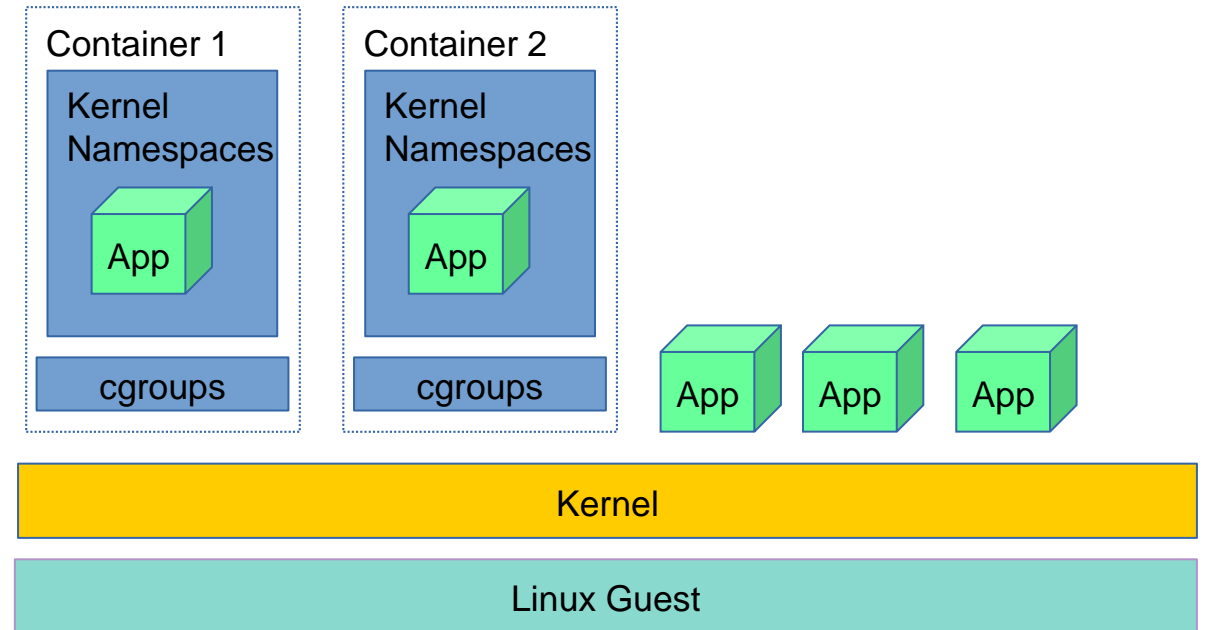
- **Container goals and characteristics:**

- Isolated application environments within a Linux OS instance
- Each container has its own, different address space but same kernel
- Serve a single task
- Self contained set of files for applications
- Startup time and efficiency compared to native execution



Linux control groups and namespaces are used for isolation

- To simplify:
 - “**cgroups**” will allocate & control resources in your container
 - CPU
 - Memory
 - Disk I/O throughput
 - “**namespace**” will isolate
 - process IDs
 - Hostnames
 - User IDs
 - network access
 - Inter-process communication
 - filesystems



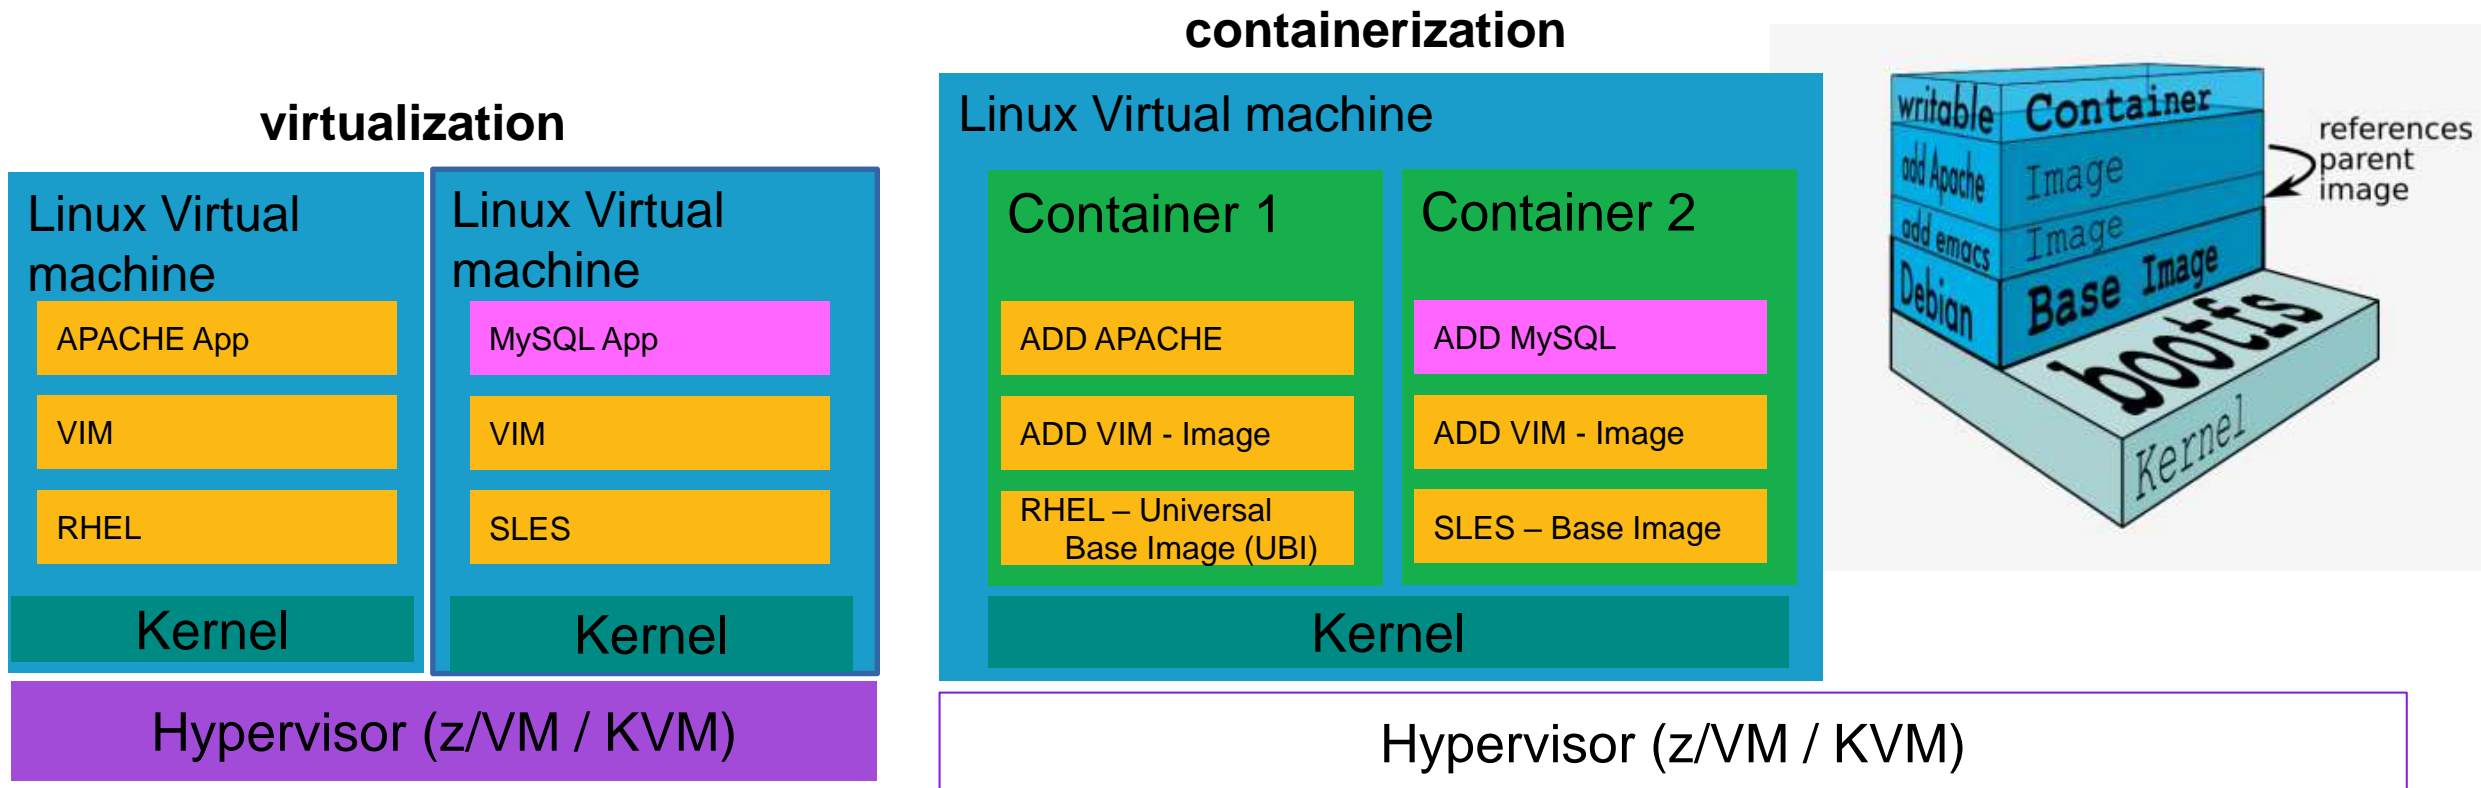
Linux Containers vs. virtual server

Virtualization, usually provides a **high level of isolation** and security as all communication between the guest and host is **through the hypervisor**.

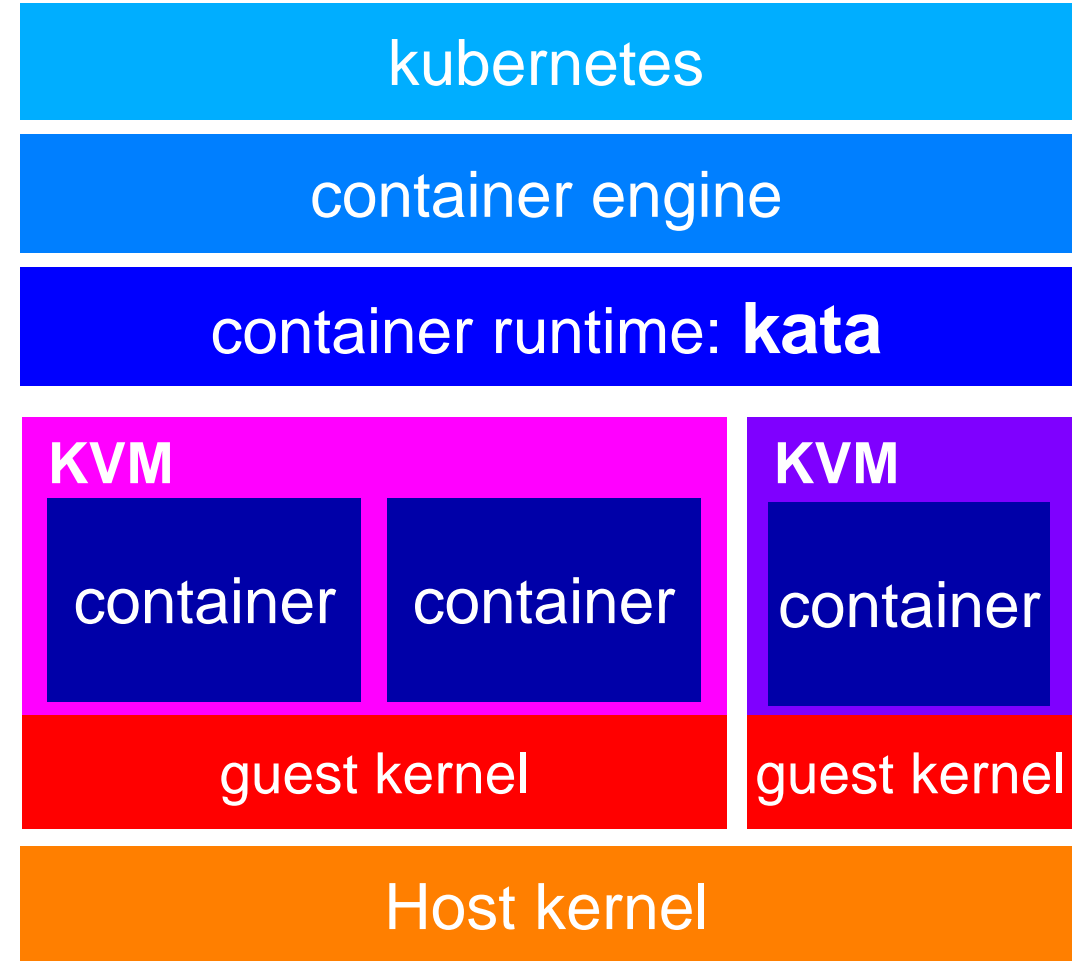

- It has usually some **overhead due to the infrastructure emulation**.

Containers, reduce the virtualization overhead, the level of virtualization called "**container virtualization**" was introduced which allows to run **multiple isolated user space instances on the same kernel**.

- **Containers is a layered approach and uses copy-on-write filesystems**



Newest Trend: From Classic containers to Kata Container technology – secured via virtualization for Confidential computing

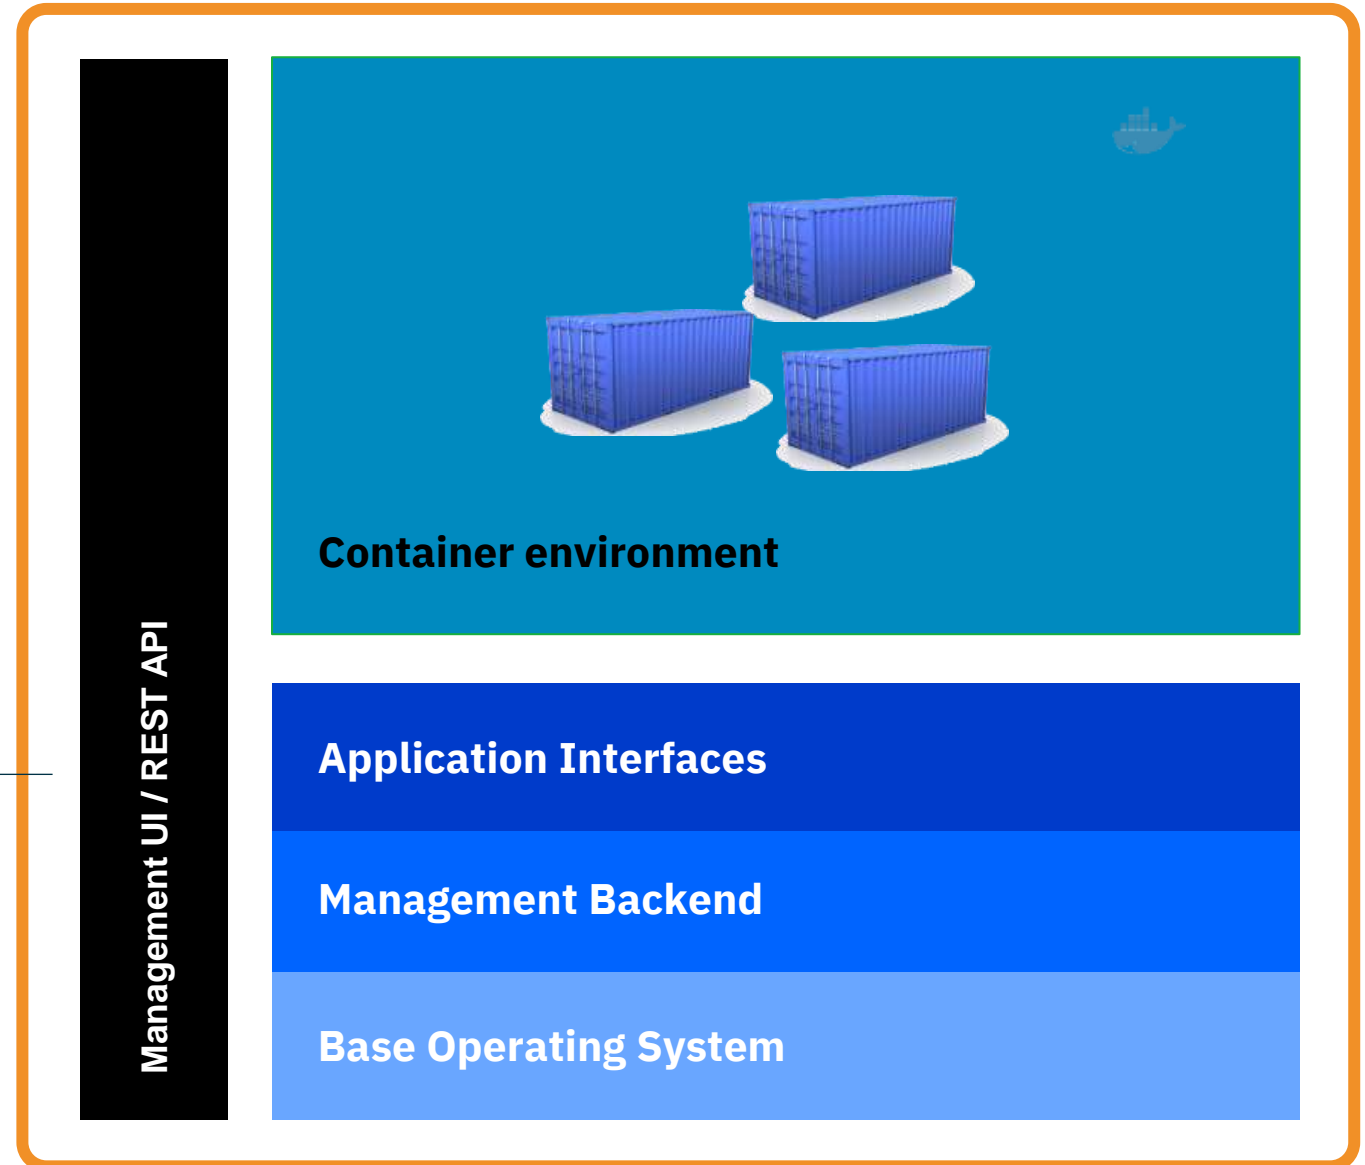


Most secure environment for container services is on IBM Z: IBM Secure Service Container (SSC)

- **SSC is a special LPAR** and provides simplified mechanism for fast deployment and management of secure solutions
- Provides tamper protection during installation and runtime
- Ensures confidentiality of data and code -at flight and at rest
- Management provided via Remote APIs (RESTful) and web interfaces only
- Enables containers to be delivered via distribution channels

IBM Secure Service Container Appliance

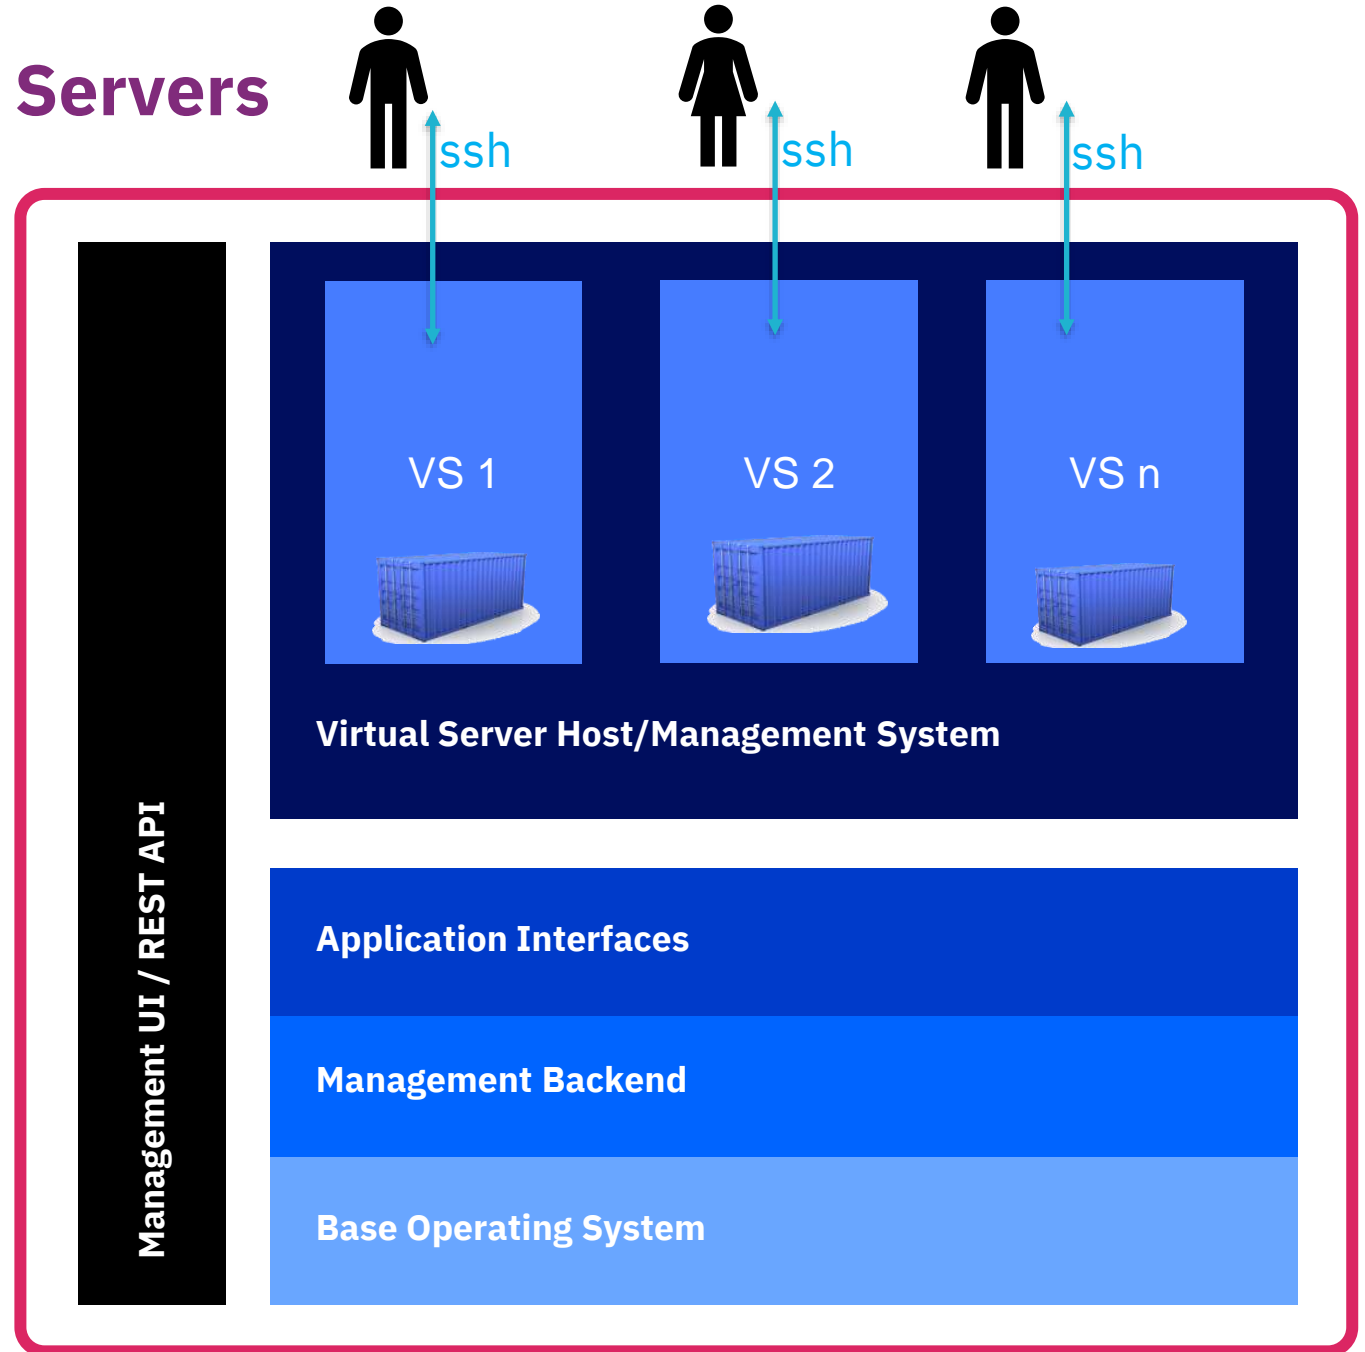
Deploy your container workload in a highly secure environment



IBM Cloud Hyper Protect Virtual Servers

- Rapidly provision a Virtual Server running on LinuxONE or in the IBM Cloud
- Authentication is done via ssh keys
→ No password is exposed to IBM
- Our system administrators do not have access to the data within the Virtual Servers and the hosting OS
- Built on IBM Secure Service Container to enforce confidentiality
- Available now

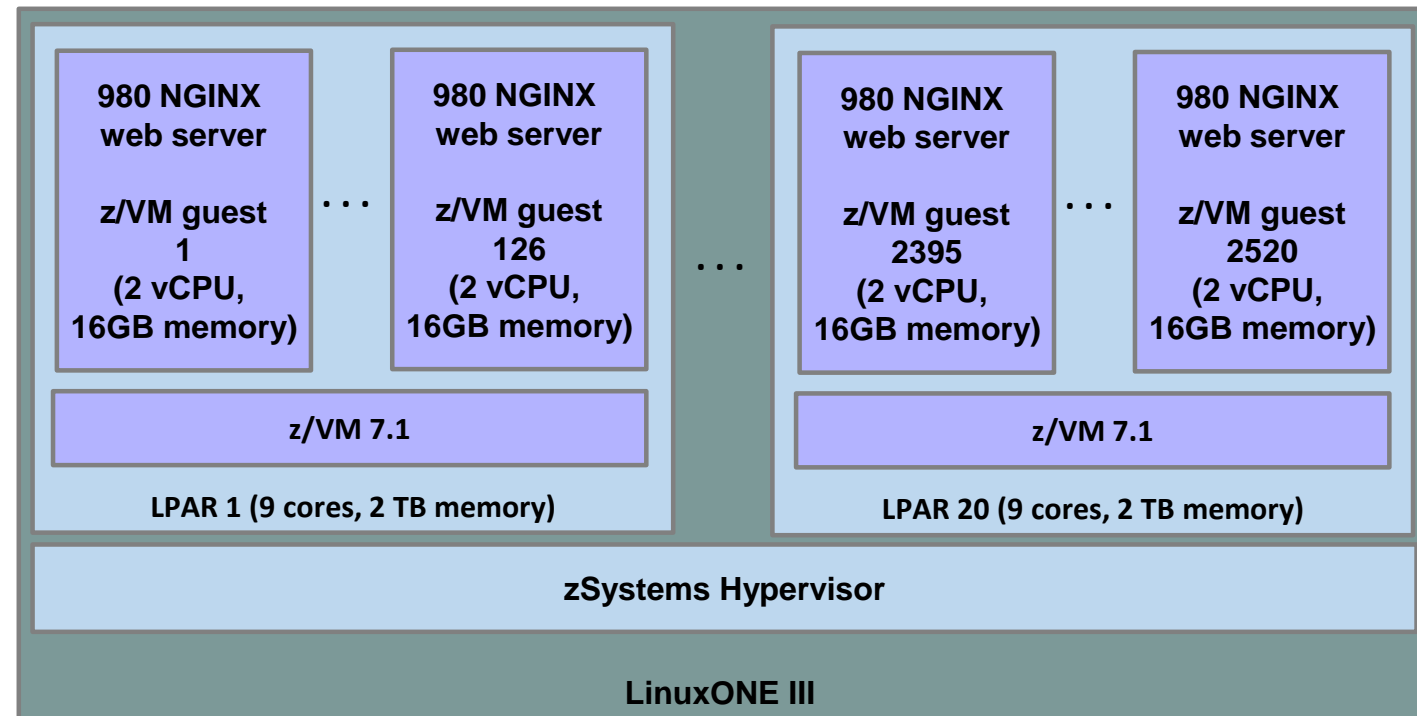
<https://www.ibm.com/cloud/hyper-protect-virtual-servers>



Container Scale-out Performance

Scale-out with Container under z/VM on LinuxONE III

- Scale-out to **2.4 million containers** in a single LinuxONE III system virtualized with z/VM

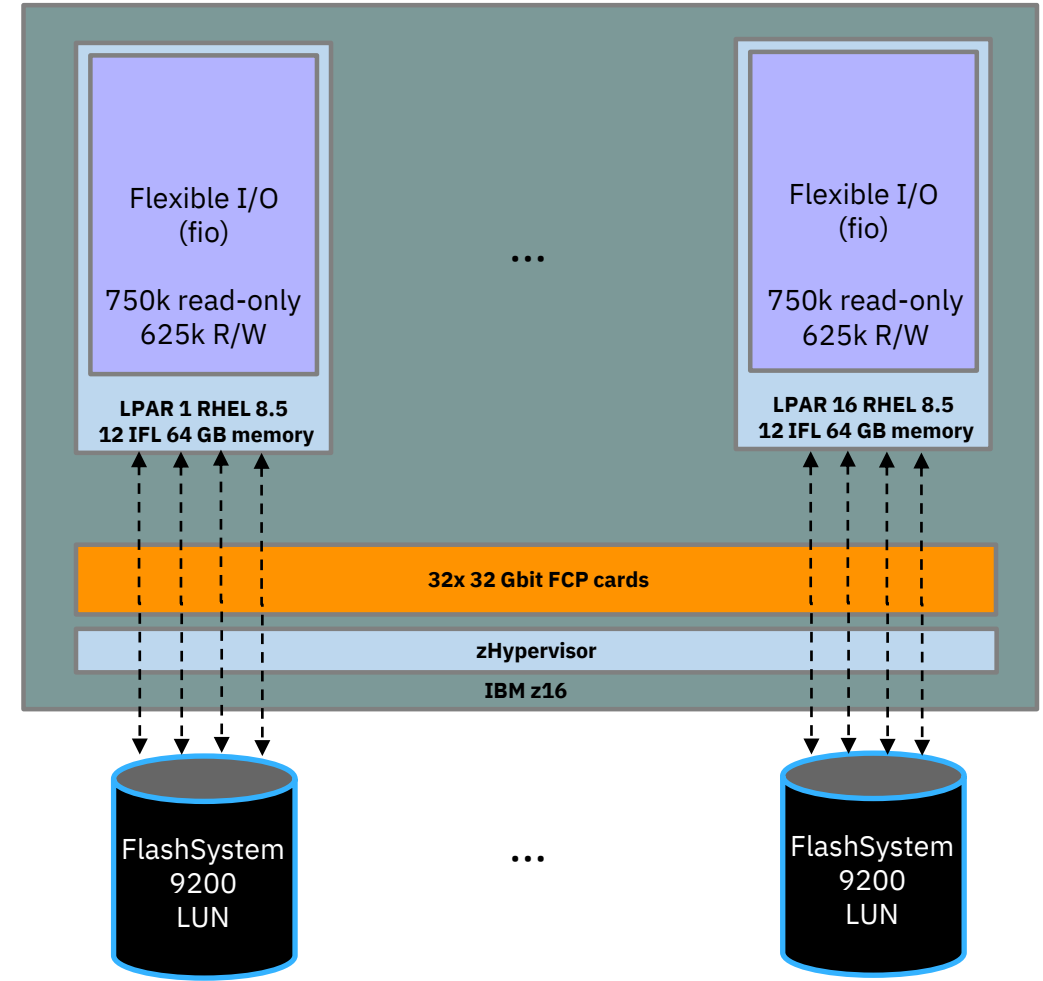


- DISCLAIMER:** Performance result is extrapolated from IBM internal tests running in a LinuxONE III LPAR with 1 dedicated core and 16 GB memory 980 NGINX Docker containers. Results may vary. Operating system was SLES12 SP4 (SMT mode). Docker 18.09.6 and NGINX 1.15.9 was used.

Scale-up encrypted I/O operations with FCP

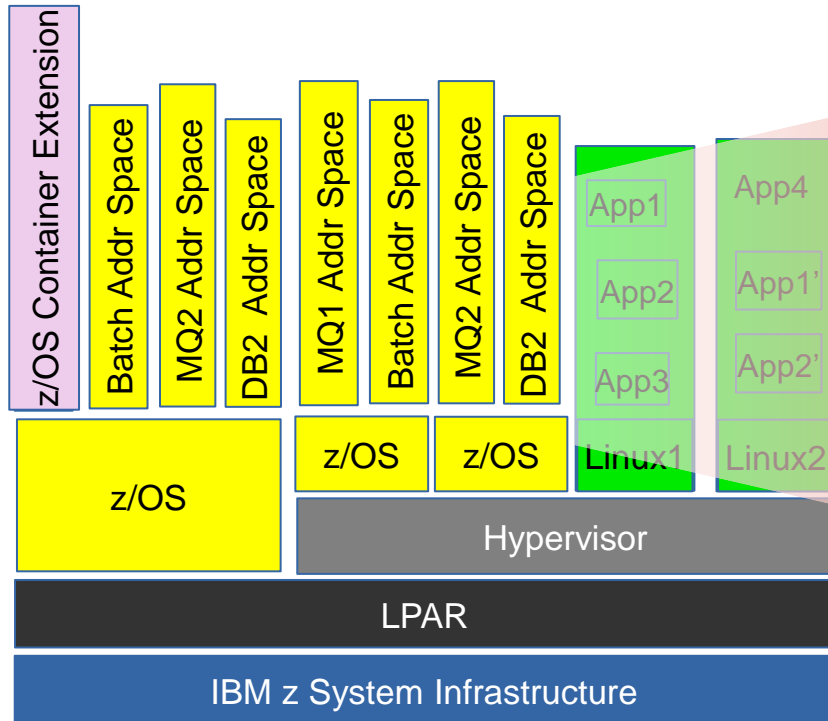
On IBM z16, scale up your I/O intensive Linux application and protect your data at rest with up to **12 million read-only I/O operations per second** and **10 million R/W operations per second** to an **encrypted filesystem** with FCP attached storage

DISCLAIMER: Performance results is extrapolated based on IBM internal tests running the fio 3.19 benchmark tool. The fio benchmarking tool was run with 128 parallel threads using 8 volumes on FS9200 equally distributed over the two nodes and file size of 150GB on each volume. z16 configuration: LPAR with 12 dedicated IFLs, 64 GB memory, RHEL 8.5 (SMT mode) running fio 3.19. Two FICON Express 32S cards and Linux XFS file system encrypted with luks2 per LPAR. Results may vary.



Application isolation is long tradition in IBM zSystems

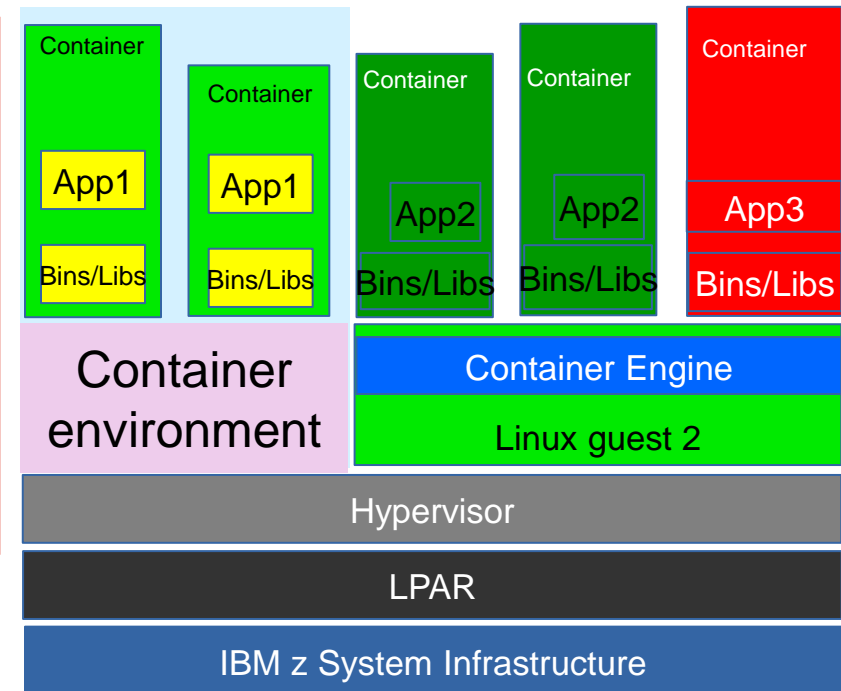
z/OS and Linux virtualization



Virtualization:

- › Infrastructure oriented
- › Virtual server resource management
- › Several applications per server
- › Isolation per virtual server

Docker Container deployment in Linux



Containers:

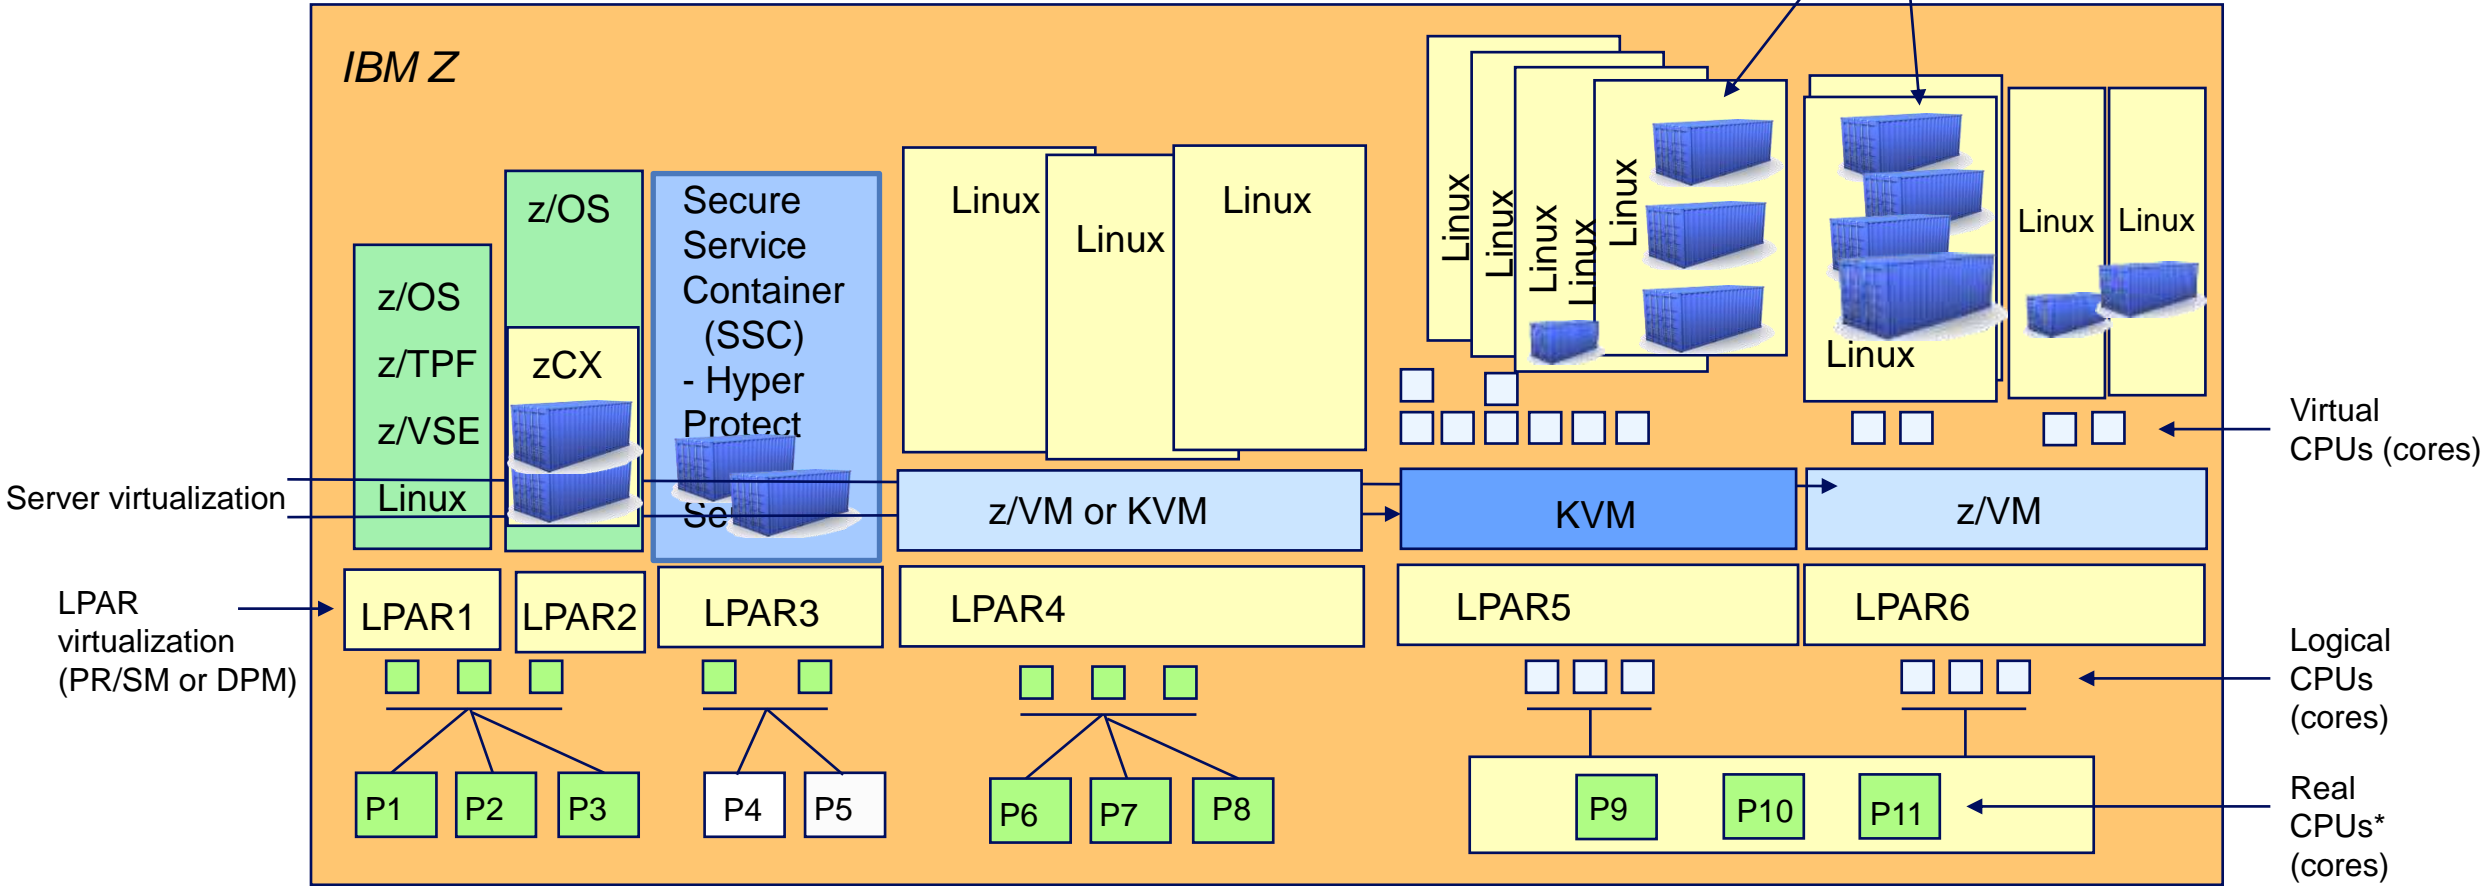
- › Service oriented
- › Application management via container
- › Solution decomposed into several units
- › Dynamic, isolation in container

IBM zSystems Virtualization and Container options



Server virtualization. There are typically dozens or hundreds of Linux servers in a LPAR virtualized using z/VM or KVM or SSC.

Application isolation. There are typically thousands of Containers in Linux on IBM Z.



P1 – P11 are Central Processor Units (CPU -> core) or Integrated Facility for Linux (IFL) Processors (IFL -> core)
 * - One shared Pool of cores per System only
 Note: - LPARs can be managed by traditional PR/SM in IBM Z and additional with Dynamic Partition Manager (DPM) in LinuxONE

Kubernetes (K8S) – container orchestration defines itself in a cluster format



cluster

Kubernetes Master Nodes (Control)

K8S Worker Node1
(Compute)

pod1

pod2

K8S Worker Node2
(Compute)

pod1

pod2

K8S Worker Node3
(Compute)

pod1

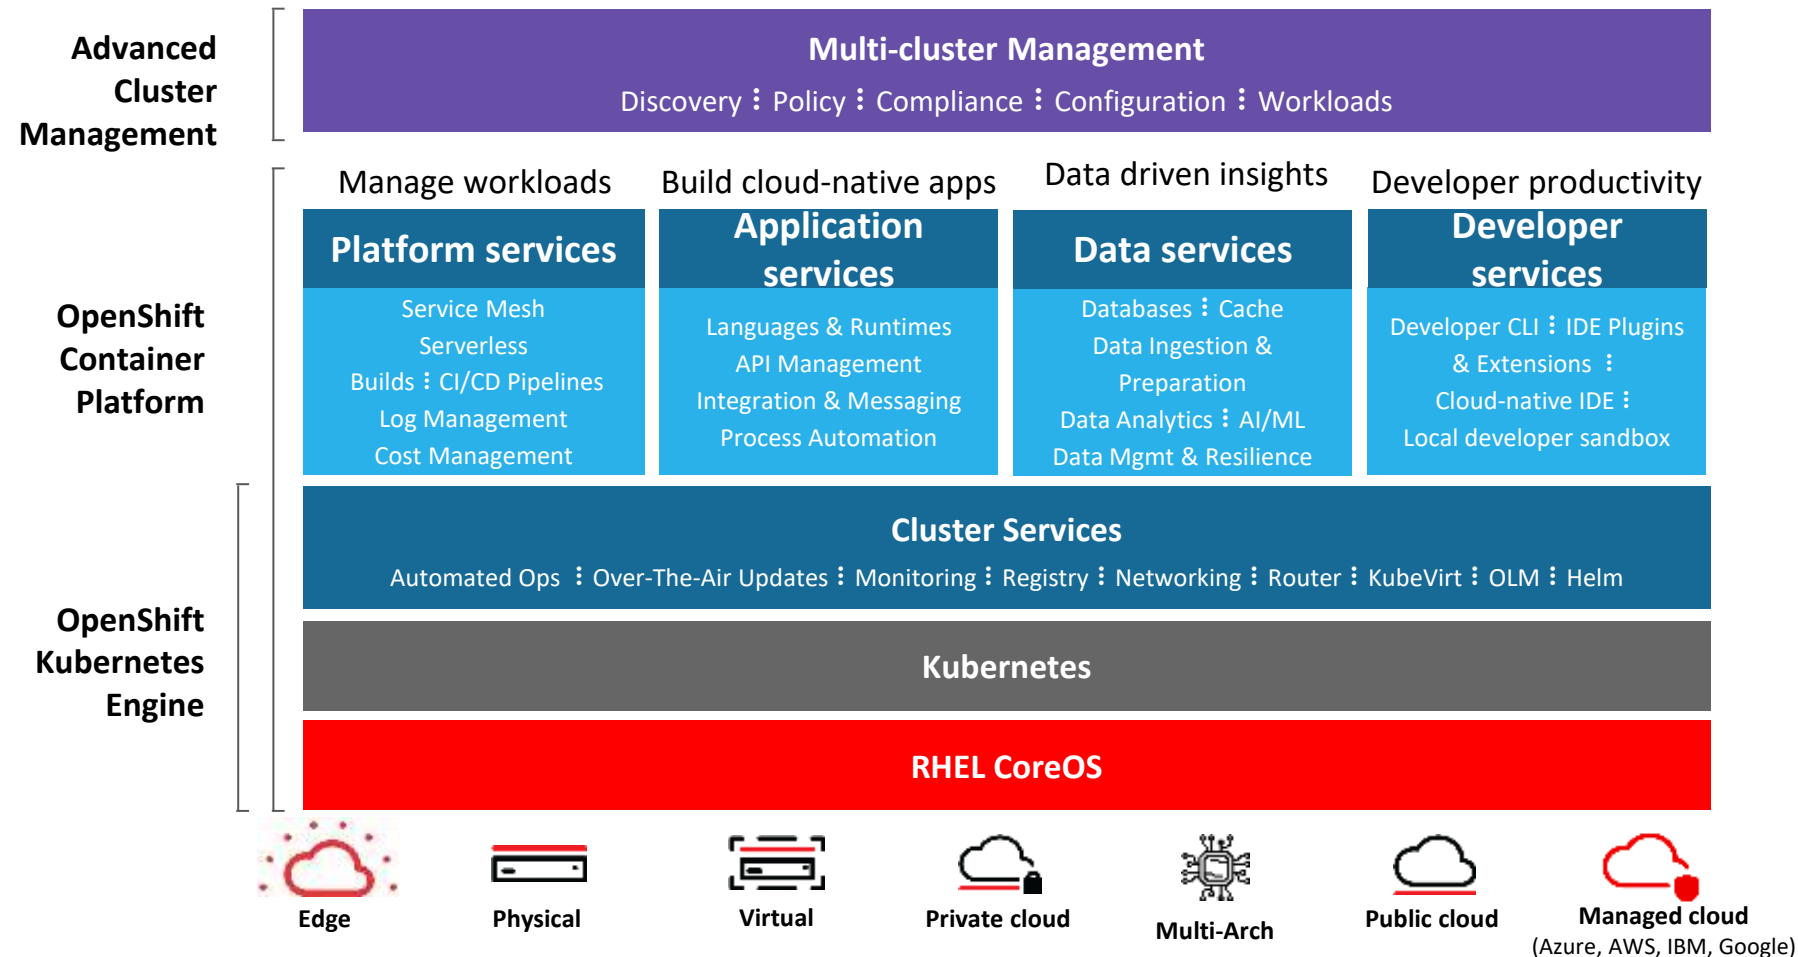
pod3

Perfect fit
for Virtual
machines
on IBM
zSystems

Kubernetes is not running container – it orchestrates them

Red Hat OpenShift is Kubernetes and additional services

- the only solution running on multiple architectures including IBM zSystems

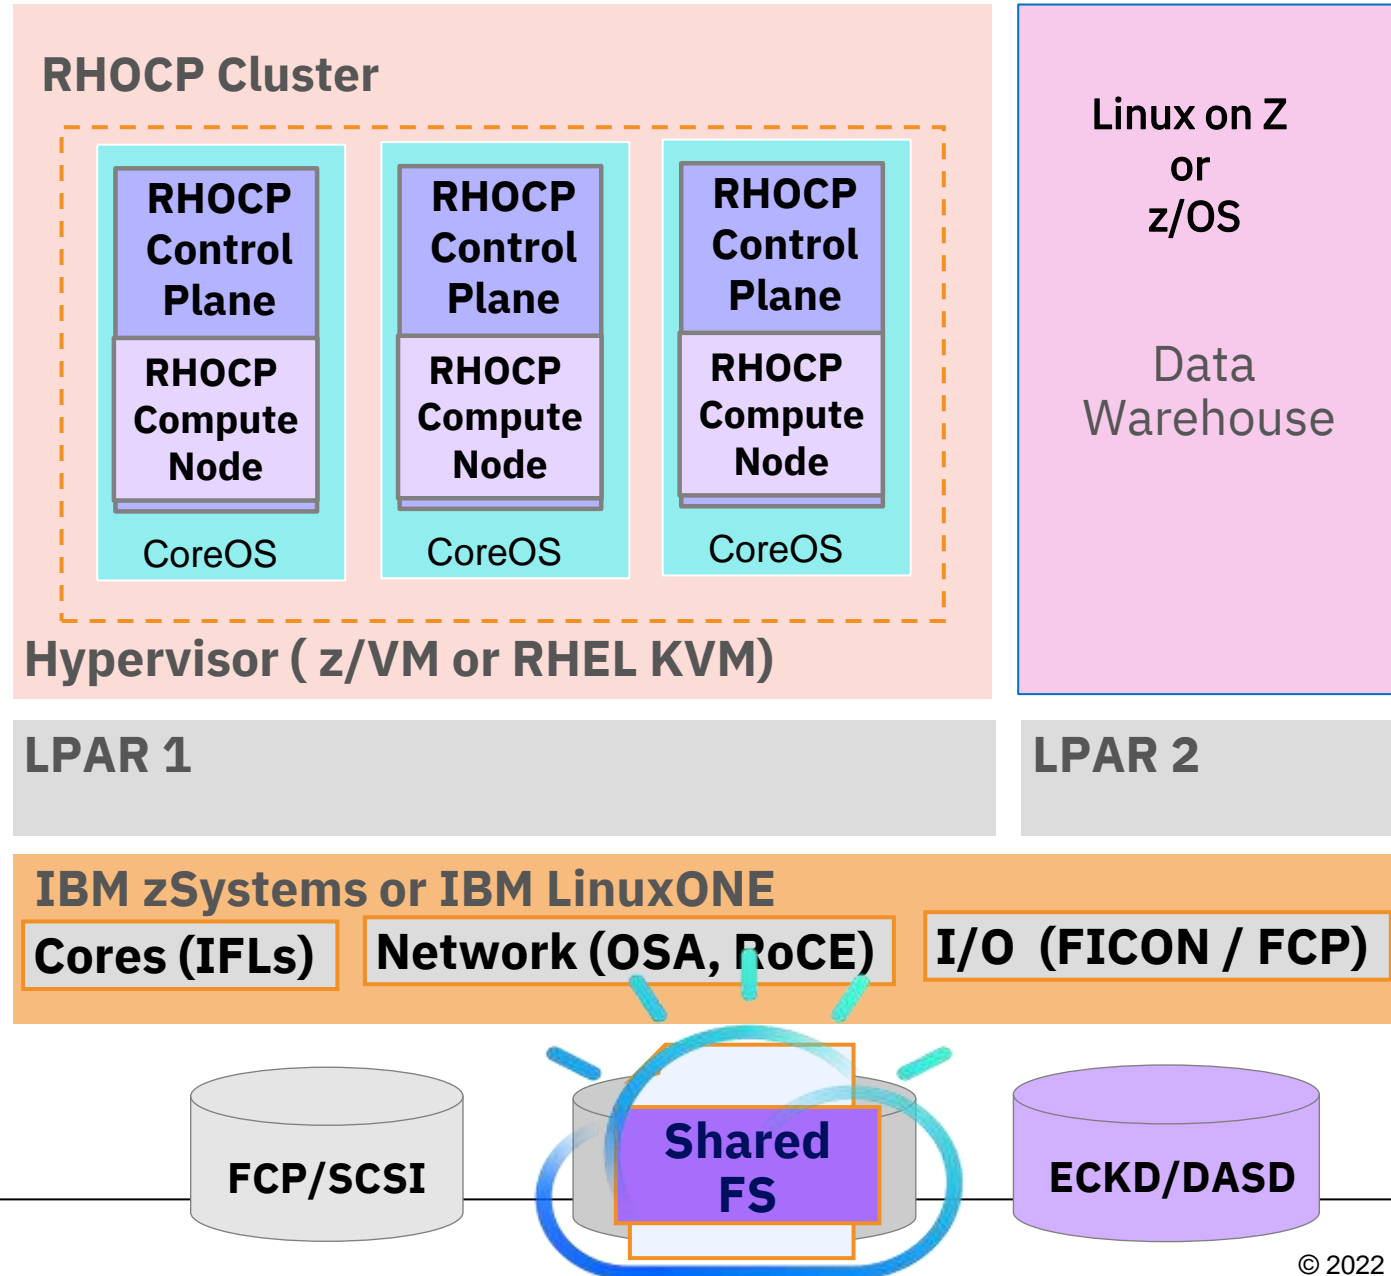


Minimum Installation Scenario of RHOCP on IBM zSystems

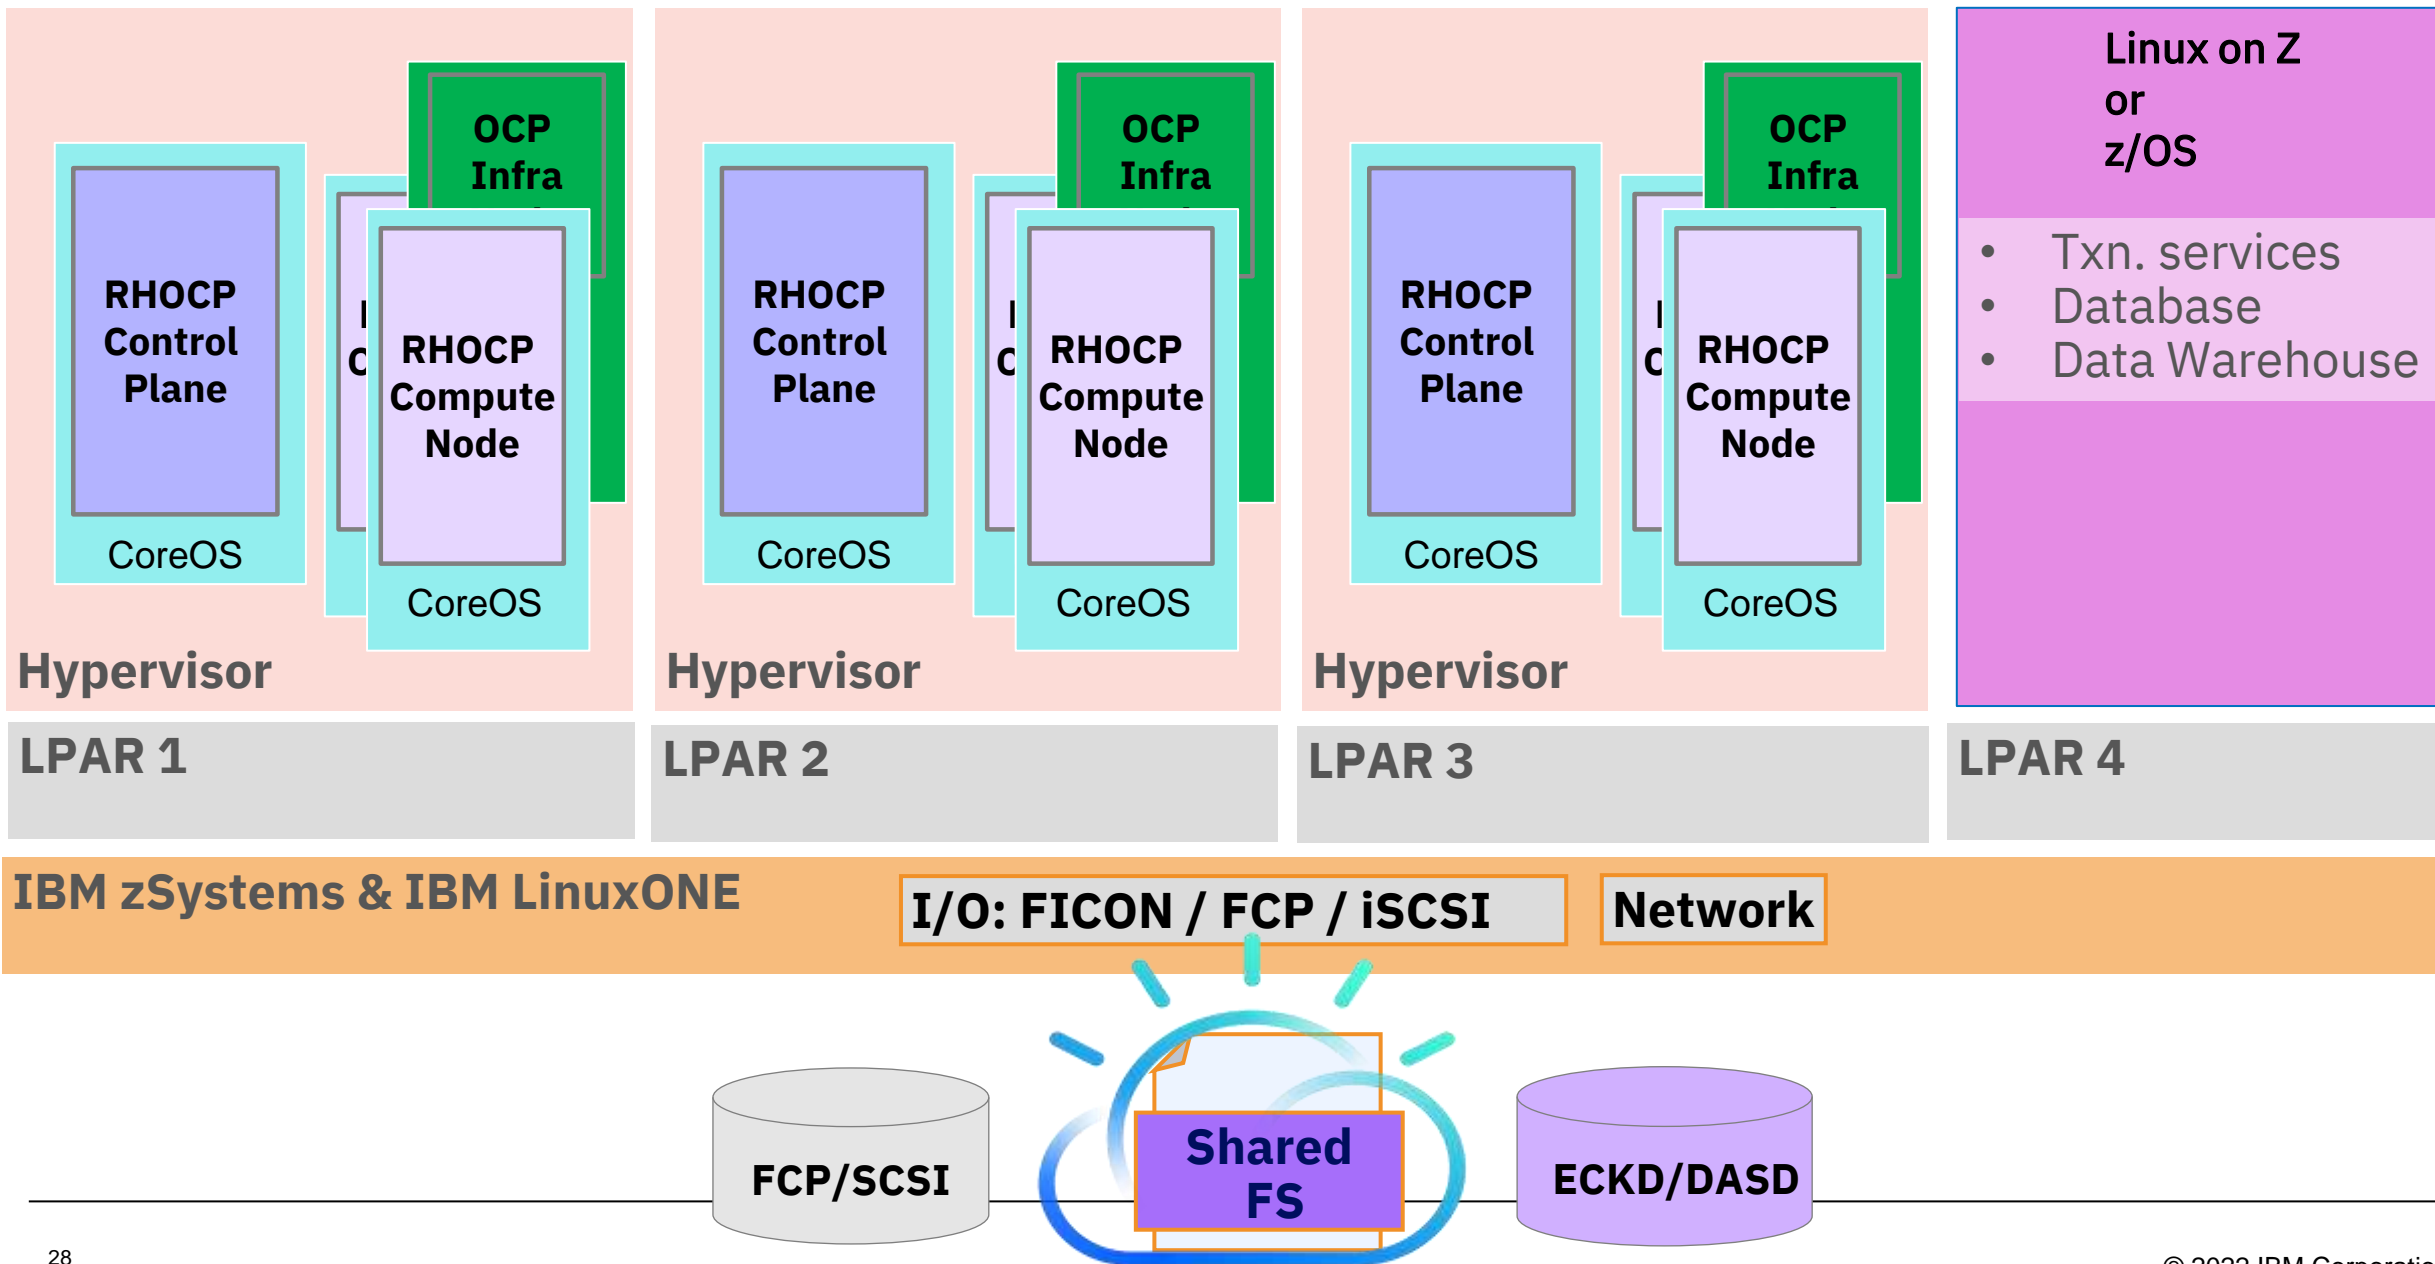


'Three Node' cluster from RHOCP 4.8

- converged Control Plane nodes and Compute Nodes



RHOCP cluster production like Overview diagram

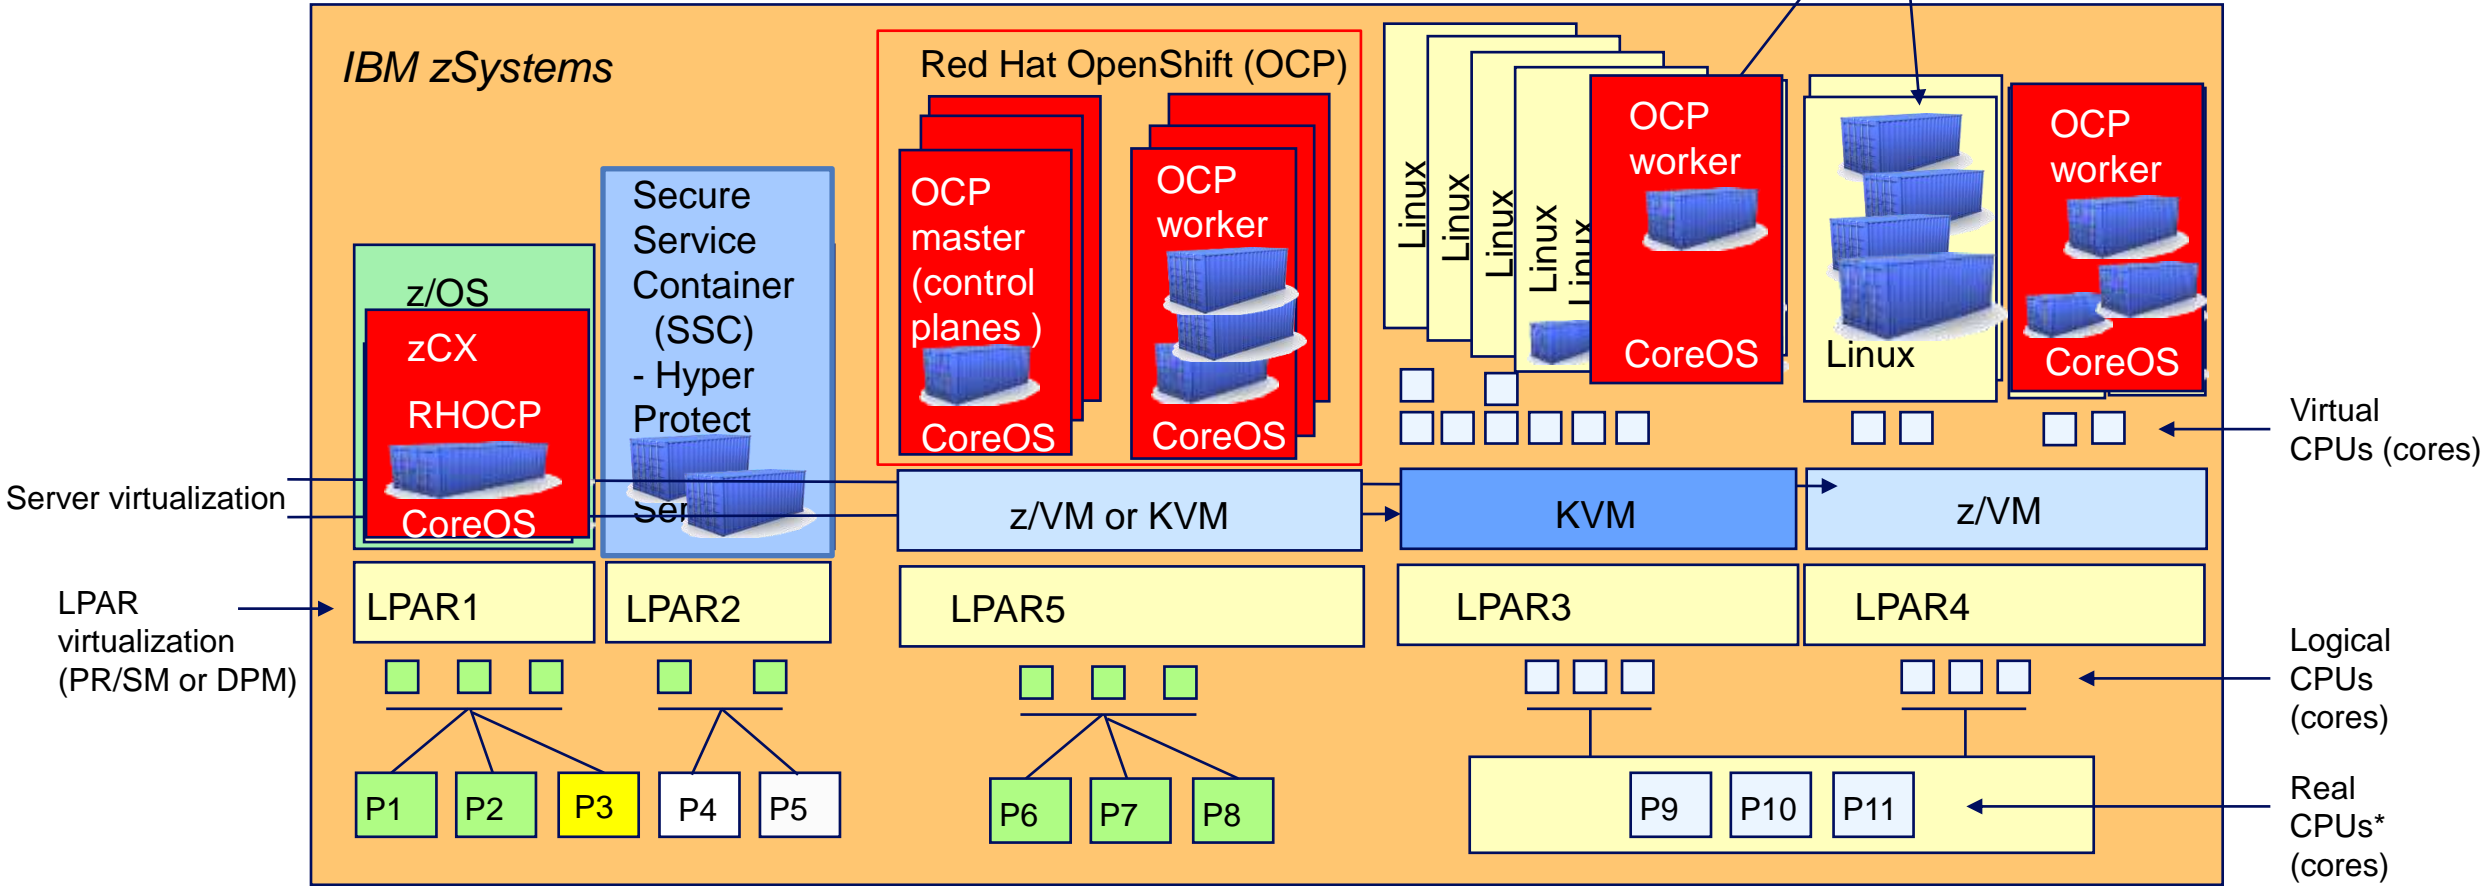


IBM Z Virtualization and Container options

Server virtualization. There are typically dozens or hundreds of Linux servers in a LPAR virtualized using z/VM or KVM or SSC.

Red Hat OpenShift is an Enterprise grade Kubernetes environment. It can be installed in a z/VM or KVM env.

Application isolation. There are typically thousands of Containers in Linux on IBM zSystems.



P1 – P11 are Central Processor Units (CPU -> core) or Integrated Facility for Linux (IFL) Processors (IFL -> core) or ZIIP cores in z/OS
 * - One shared Pool of cores per System only
 Note: - LPARs can be managed by traditional PR/SM in IBM Z and additional with Dynamic Partition Manager (DPM) in LinuxONE

z/OS Container Extensions – a virtual container environment



Pre-packaged OpenShift environment provided by IBM

- 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

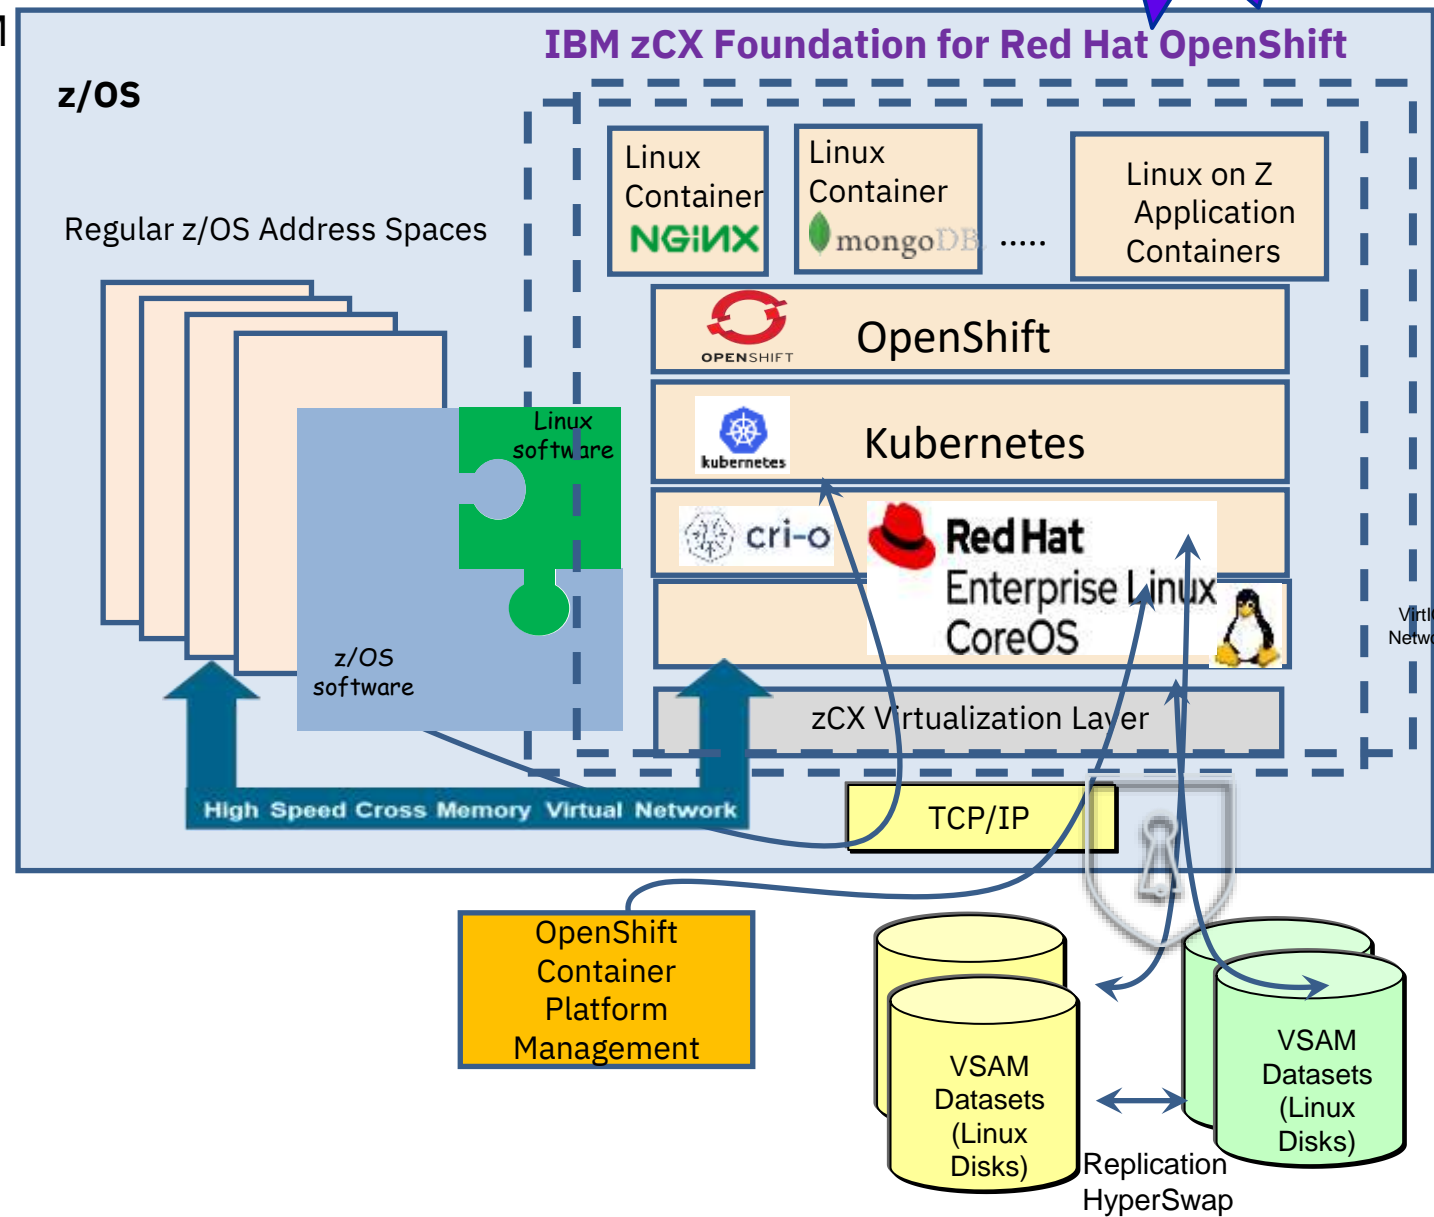
- Any software available as a Container image (s390x) - growing ecosystem
- Any home-grown Linux on Z container images
- Using standard interfaces

Access into underlying environment via RHOCAP APIs

- No root access
- Administrative tasks via RHOCAP and z/OS
- Secure virtual network - SAMEHOST

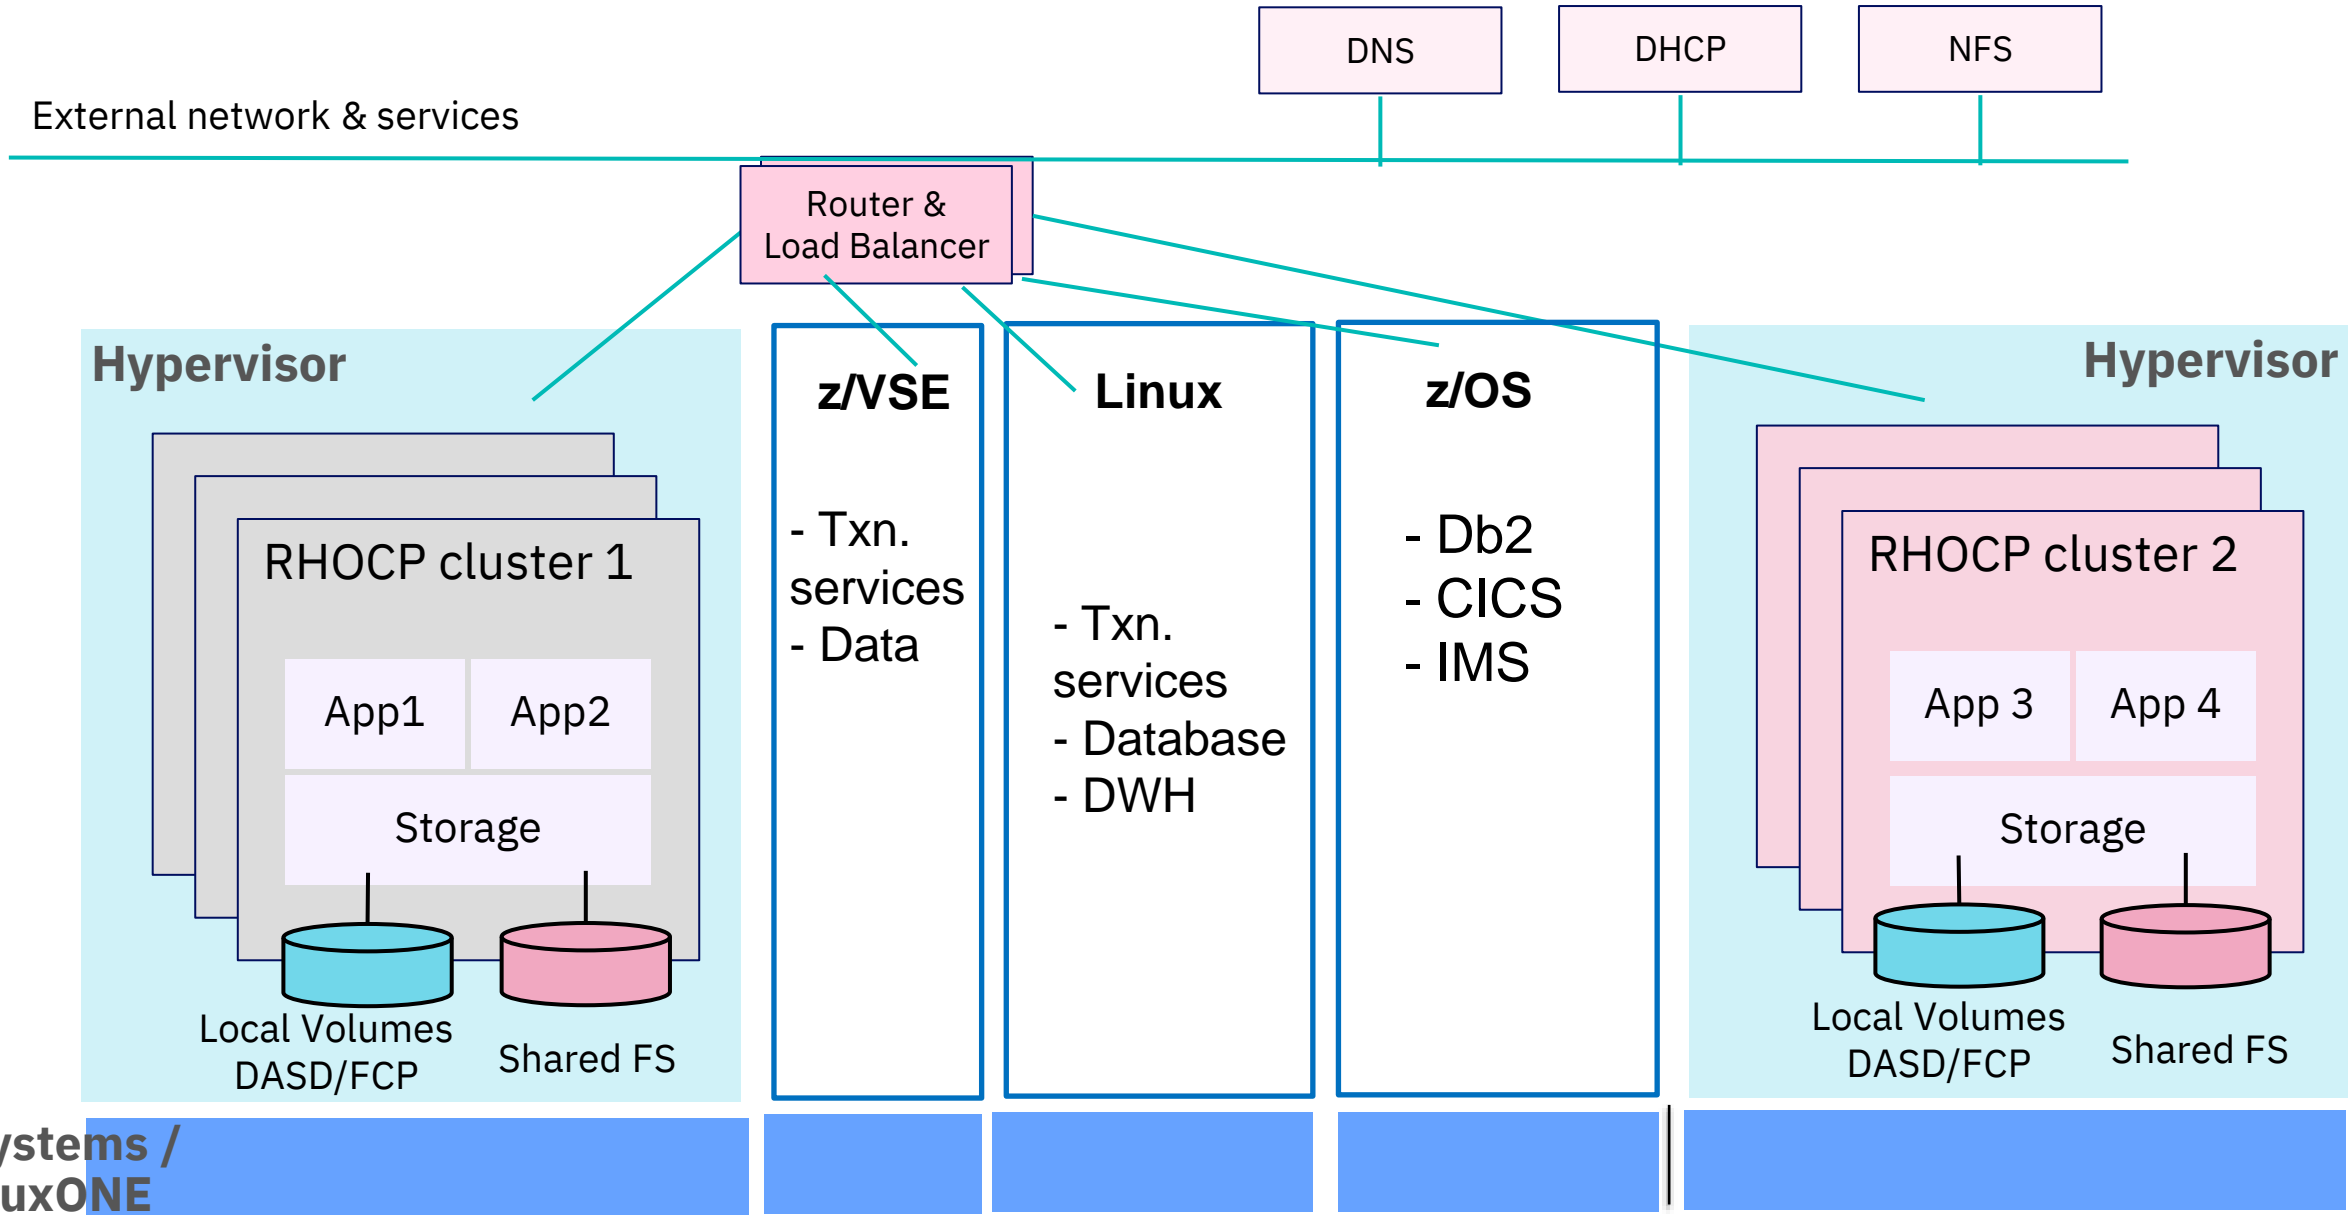
Also provides IBM and ISVs a means of delivering solutions into this environment

- Requires packaging of software as Container images



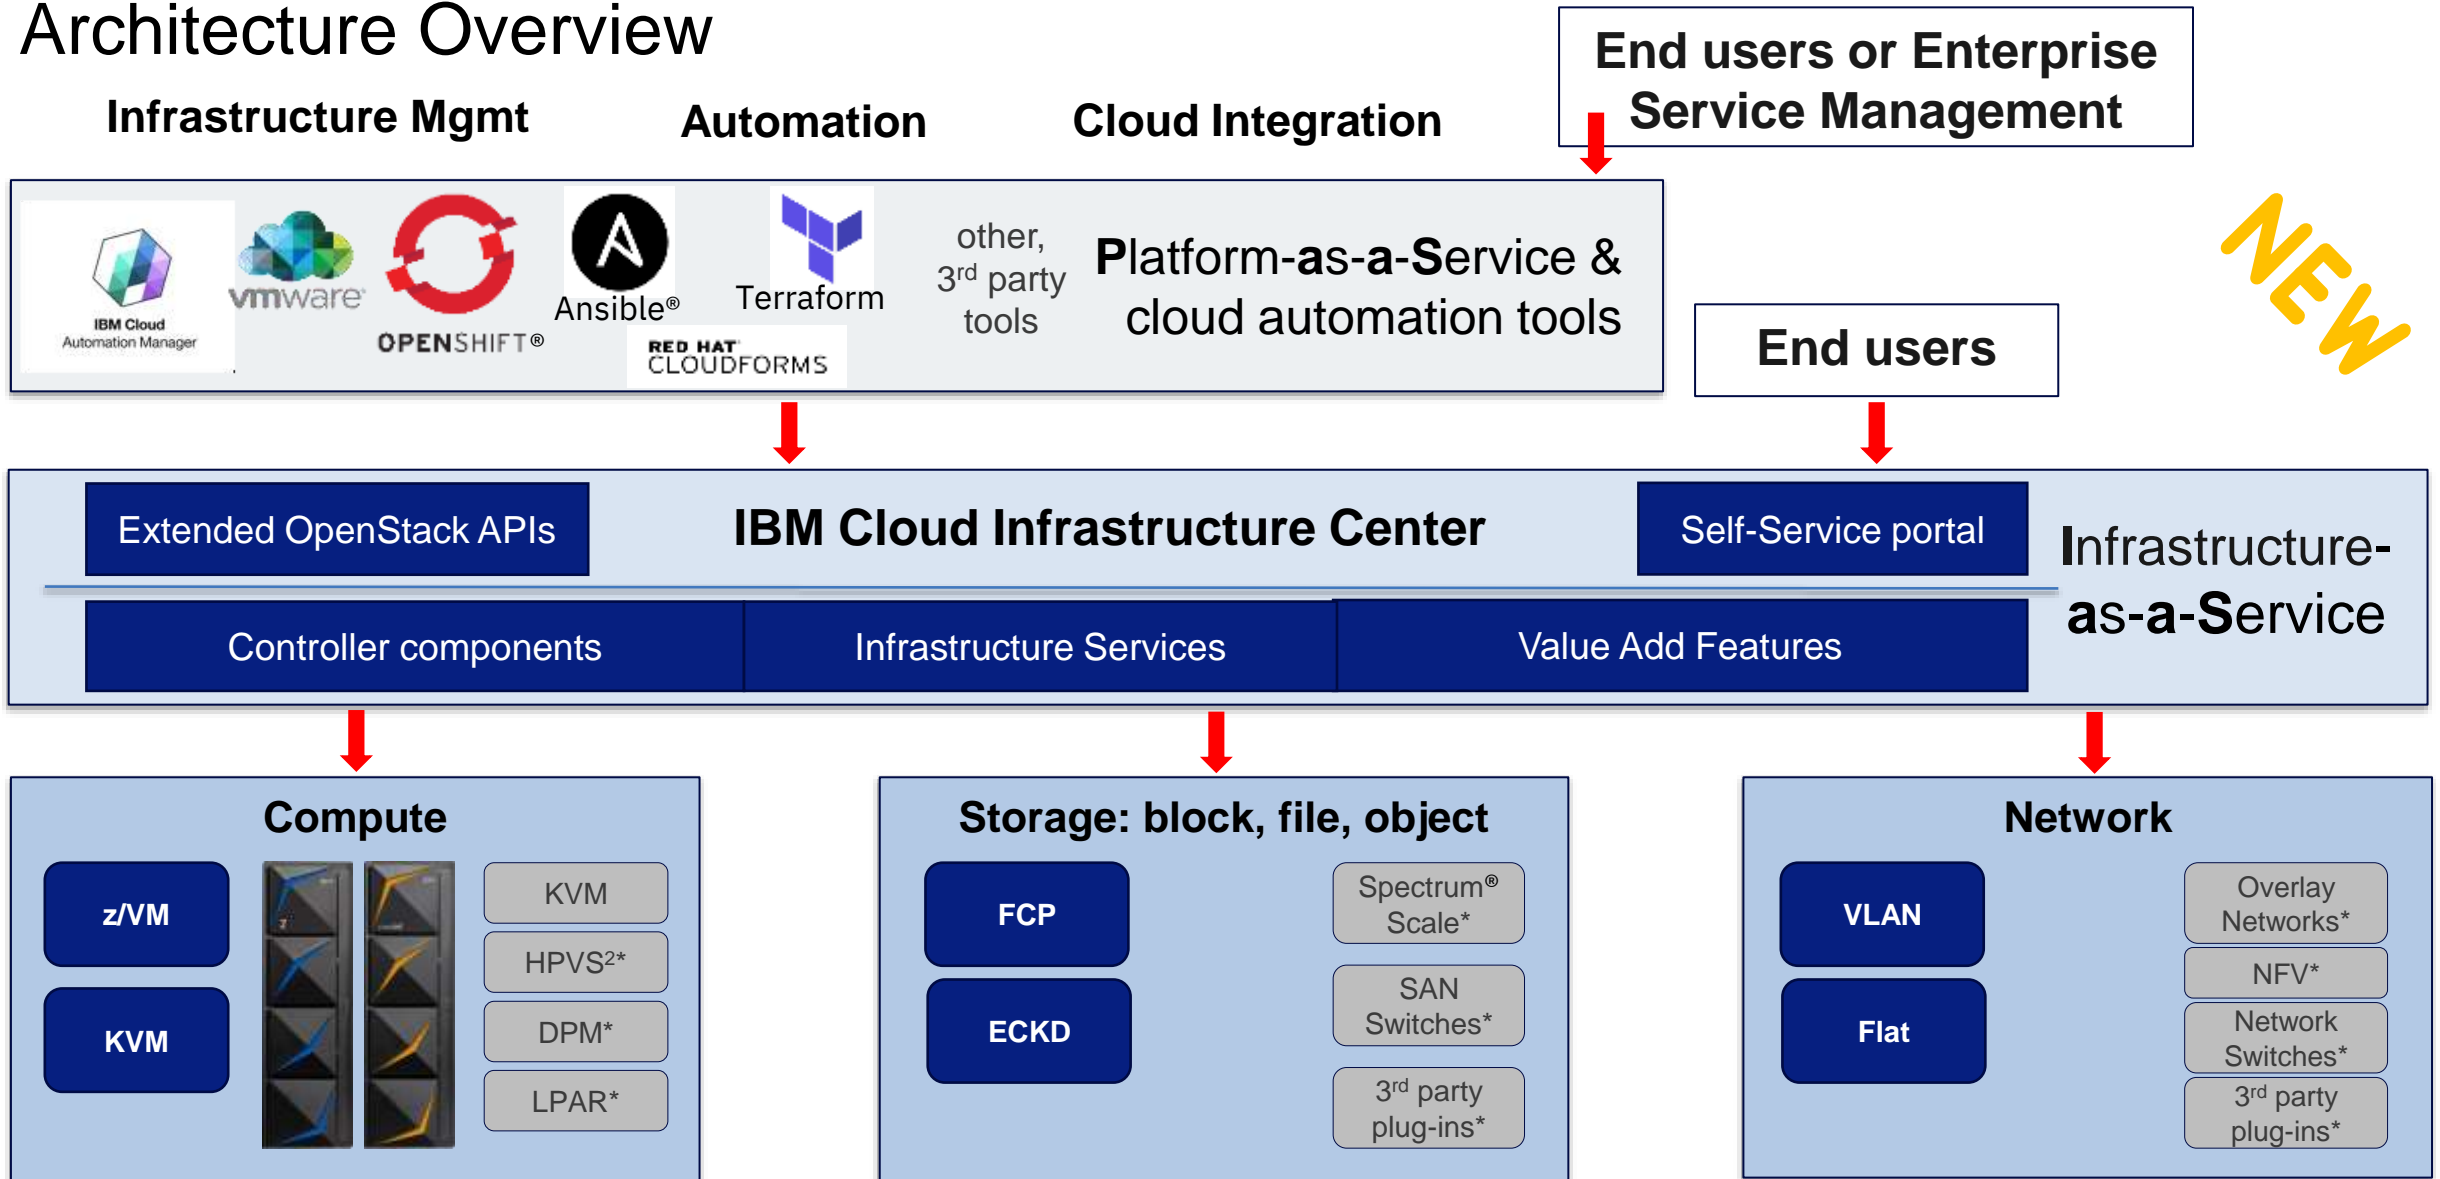
Why RHOCP on IBM Z – cause of operational capabilities for hybrid

THE platform for Hybrid workload and multiple RHOCP environments on the same HW machine



Infrastructure management with IBM Cloud Infrastructure Center

Architecture Overview



¹ IBM Hyper Protect Virtual Server

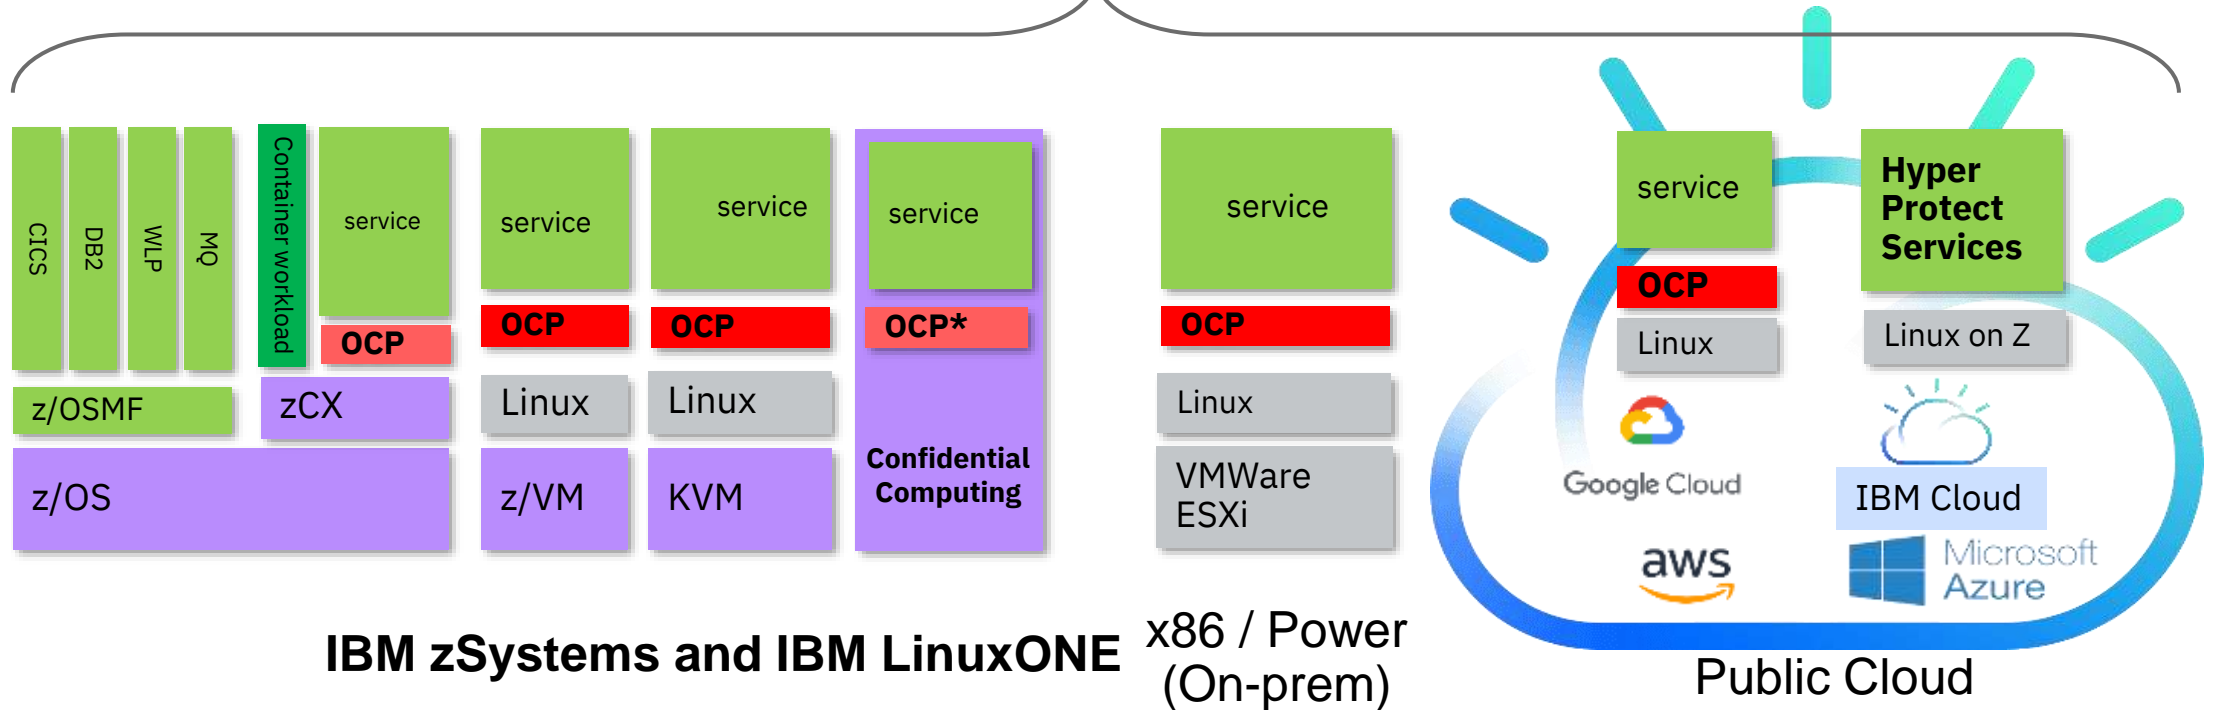
* All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

Can provision RHEL, CoreOS, SLES 15 SP1 und Ubuntu 20.04 guests

Outlook: The Hybrid Multicloud Vision with OpenShift



Advanced Cluster Management (RHACM)
for multi cloud solutions and App. Life cycle
- [Can manage Kubernetes and OpenShift cluster](#)



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

Designed for competitive advantage

The Strengths and Benefits of Virtualization for containers and cloud on IBM zSystems

<https://www-03.ibm.com/systems/z/os/linux/success/>

1-Operational

Efficiency
advantages

Simplicity, high performance and same arrangements for HA / DR and administration

2-Co-location

Integration
options

Designed to run multiple different environments in a single LPAR and isolate at same time

3-Security

Multi envir.
control

Pervasive encryption enablement, RACF for z/VM and EAL 5+ certification, Secure boot

4-Eco. benefits


Business
continuity

Avsailability, Reliability & Management cost can be considerably less vs x86 or public clouds

Virtualization on IBM zSystems

– continues the way of effectively run and manage different workloads since many years, incl. container and clouds

- From the beginning sympathy on all environment levels
- Solid Virtualization
- High Reliability
- Scaling on demand with highest flexibility



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*



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