

CHINA
OpenStack Days

CHINA
OpenStack Days

IT大咖说
知识分享平台

Topic: Openstack 在IBM主 机上的企业级应用和解决方案

Name:纪晨,黄锐, 黄士林



- 纪晨，IBM高级工程师，从事OpenStack主机云的设计与研发
- 黄锐，IBM高级工程师，从事OpenStack主机云的设计与研发
- 黄士林，IBM高级工程师，从事OpenStack主机云的设计与研发

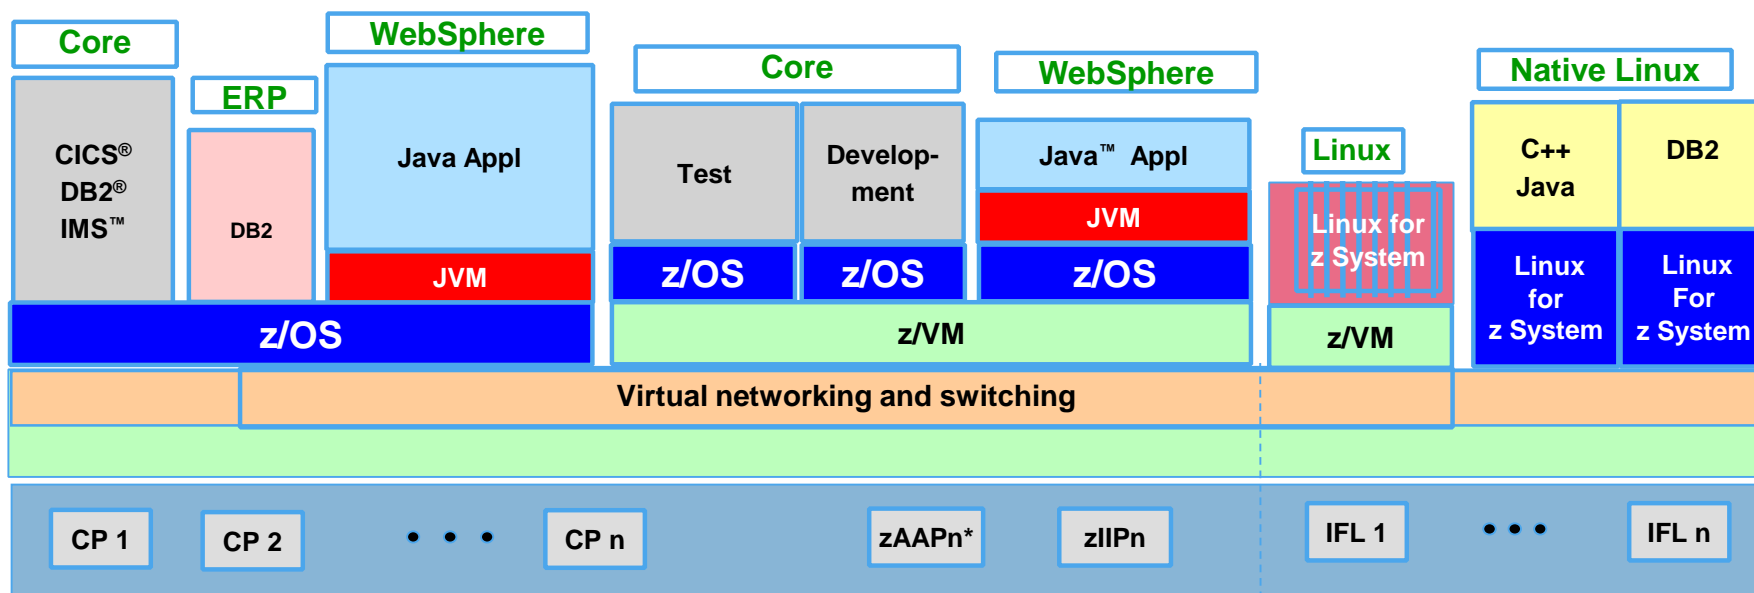
- IBM 主机以及相关介绍
- 主机上的虚拟化介绍
- 主机上的openstack解决方案介绍
- 主机上基于openstack的应用方案介绍



<p>The world's premier system for enabling data as the new security perimeter</p> <ul style="list-style-type: none">• Pervasive encryption• No application changes• Protect from internal and external threats 	<p>Designed for data serving in a cognitive world</p> <ul style="list-style-type: none">• Speed, scale and reduced latency• Efficiency for managing data• Secure and flexible access to data 	<p>The best infrastructure to support an open and connected world</p> <ul style="list-style-type: none">• 'From anywhere' mobile access• Simplified sys admin of z/OS• Standardization for skills transfer 
--	---	--

https://apps.na.collabserv.com/blogs/2276ab17-3094-479d-8120-14e164e0abcf/entry/IBM_Z?lang=en_us

An integrated, highly scalable computer system that allows many different pieces of work to be handled at the same time, sharing the same information as needed with protection, handling very large amounts of information for many users with security, without users experiencing any failures in service



- Large scale, robust consolidation platform
- Built-in Virtualization
- 100's to 1000's of virtual servers on z/VM
- Intelligent and autonomic management of diverse workloads and system resources

	IaaS	Traditional IT
Consumption Model	On Demand Self Service	Request to Resource Owner
Scaling	out	up
Life Cycle	“Cattle”	“Pets”
Hardware	cost optimized	premium
Applications	expect HW outage	rely on HW HA
Reliability	via software	via hardware (+ firmware)

Note: The table shows the extremes, real life is more „mixed“...

Typically, z Systems are run in the Traditional IT model.

In an IaaS model, why use z Systems / LinuxONE ??

The Benefits of Hosting Distributed Workloads on z Systems

- Very efficient virtualization with z/VM, KVM:
 - Do more with less
 - ~100% utilization of system resources
- Business Resiliency Capabilities:
 - High Availability
 - Disaster Recovery, Serviceability
 - Reliability
 - Storage failover (HyperSwap)
 - Data replication (XRC, PPRC)
- Server hardware consolidation
- Storage consolidation
- Network consolidation
- Overall Architecture Simplification
- Security Capabilities:
 - Centralized Authentication
 - Cryptographic Acceleration
 - Image Isolation
 - Secure communications with HiperSockets and Guest LANs
- Proximity to z/OS managed data:
 - Increased transaction throughput, HiperSockets
 - Shared data access
 - Integrated storage management

Service Management Layer (provided by SmartCloud technologies)



- z/VM
- zManager
- kvm
- Hyper-V
- Power VC
- FSM
- VMware



Datawarehousing IDAA Solution



zManager for z/OS® and zBX



Systems Director for Power System x and storage



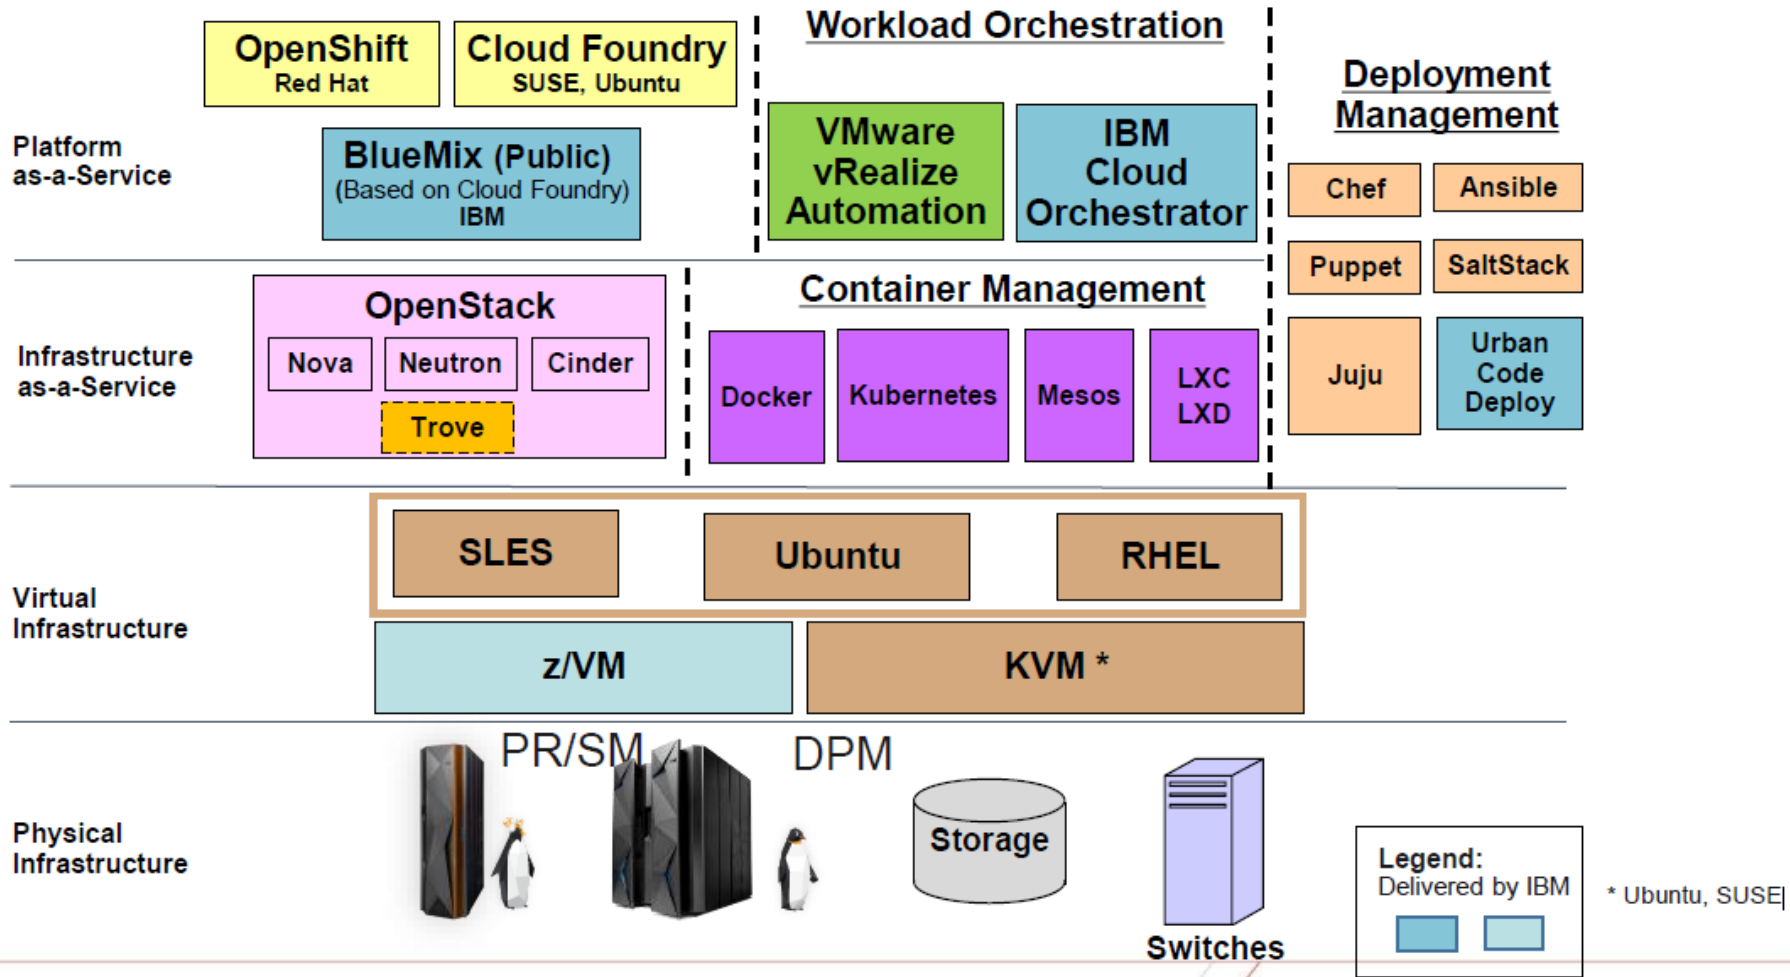
FSM for Intel® and Power ITEs



3rd party Managers and Servers



No architectural limits for integration with other platform




<https://linuxone20.cloud.marist.edu/cloud>

Test drive your apps and more in the **LinuxONE Community Cloud**

Unleash your open source apps and services with the Linux platform without limits.


Infrastructure



Virtual Servers

Quickly and easily deploy a Linux instance with the Linux distribution of your choice

[Manage Instances](#)




Kubernetes Cluster

Quickly and easily get a Kubernetes cluster on LinuxONE

[Coming soon](#)


Containers



MongoDB Enterprise

Get a MongoDB instance to access it remotely.


[Coming soon](#)



PostgreSQL

Get a PostgreSQL instance to access it remotely.

[Coming soon](#)




Db2

Get a Db2 instance to access it remotely.

[Coming soon](#)

[Show all the containers](#)


Explore more



Journey

Everything you need to quickly solve real problems


[Go](#)



Develop

Innovate with the latest technologies

[Go](#)



Learn

Resources to help you get started quickly

[Go](#)





LinuxONE Virtualization Options



LinuxONE has three strategic virtualization platforms

- KVM for LinuxONE
- IBM z/VM
- IBM Processor Resource/System Manager (PR/SM)



KVM for LinuxONE provides an open source choice for LinuxONE virtualization for Linux workloads. Best for clients that are not familiar with z/VM and are Linux centric admins.

z/VM

Proprietary Server Virtualization that is deeply integrated into System z. Complete hardware awareness. Supported on all IBM z Systems and LinuxONE servers. z/VM will continue to be enhanced to support Linux Workloads.

PR/SM

Divide one physical server into up to 85 logical partitions (LPAR) running a mix of multiple z/VM, Linux and KVM for LinuxONE instances isolated and secured in parallel. Share resources across LPARs or dedicated to a particular LPAR. Instances are isolated and secured.

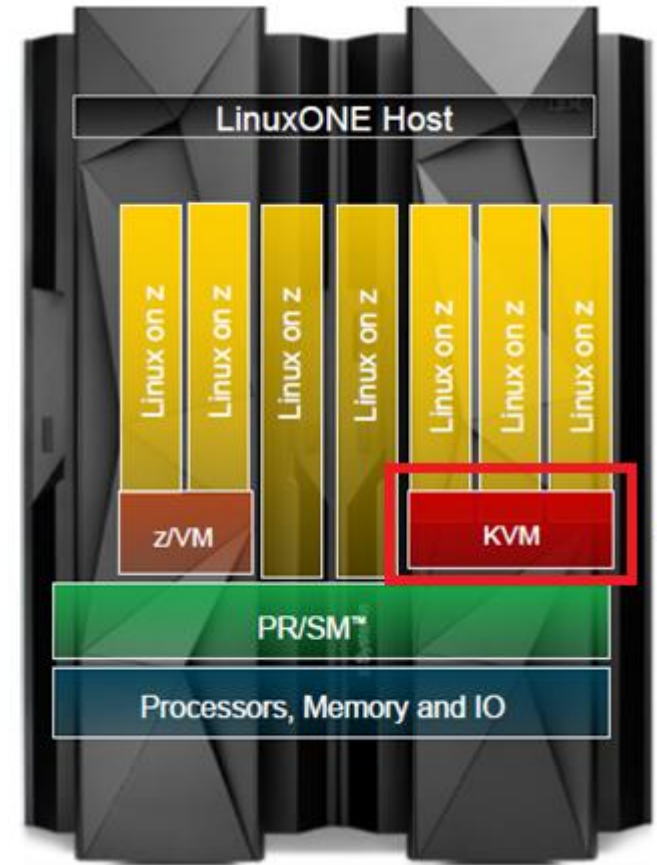
Summary

Feature	Status	Hyper-V	Ironic	Libvirt KVM (ppc64)	Libvirt KVM (s390x)	Libvirt KVM (x86)
Attach block volume to instance	optional	✓	✗	✓	✓	✓
Detach block volume from instance	optional	✓	✗	✓	✓	✓
Attach virtual network interface to instance	optional	✓	✗	✓	✓	✓

<https://docs.openstack.org/nova/latest/support-matrix.html>

z/vm and DPM is working on CI part

- kvm on z
 - Standard linux management and operation controls
 - Standard KVM interfaces allow for quick startup for clients who are familiar with x86
 - KVM-based virtualization LinuxONE allows businesses to reduce costs by deploying fewer systems to run more workloads, sharing resources, and improving service levels to meet demand
 - KVM open source solution for running virtual servers on LinuxONE enables cloud deployments and big data solutions while reducing complexity and cost



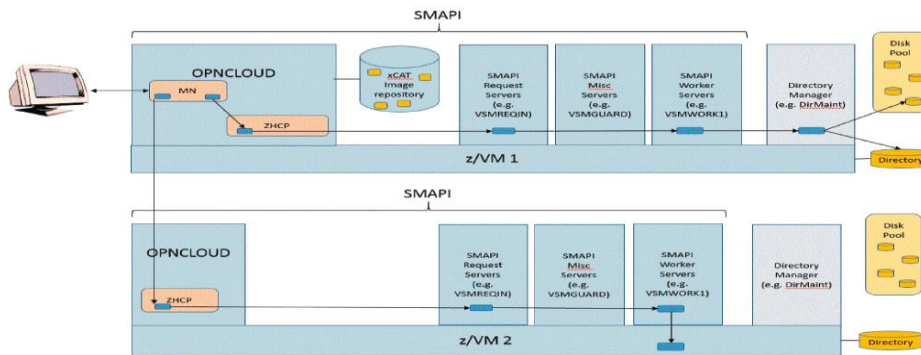
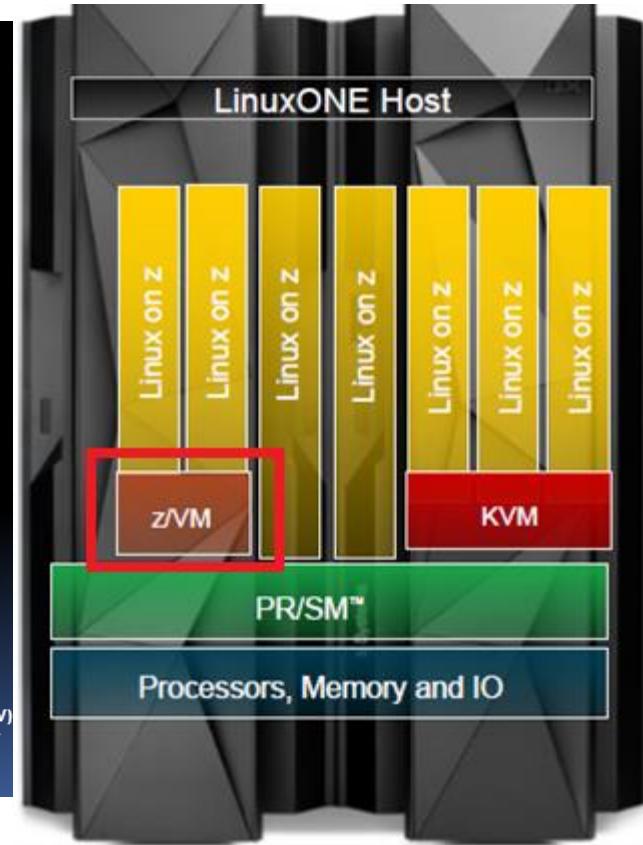
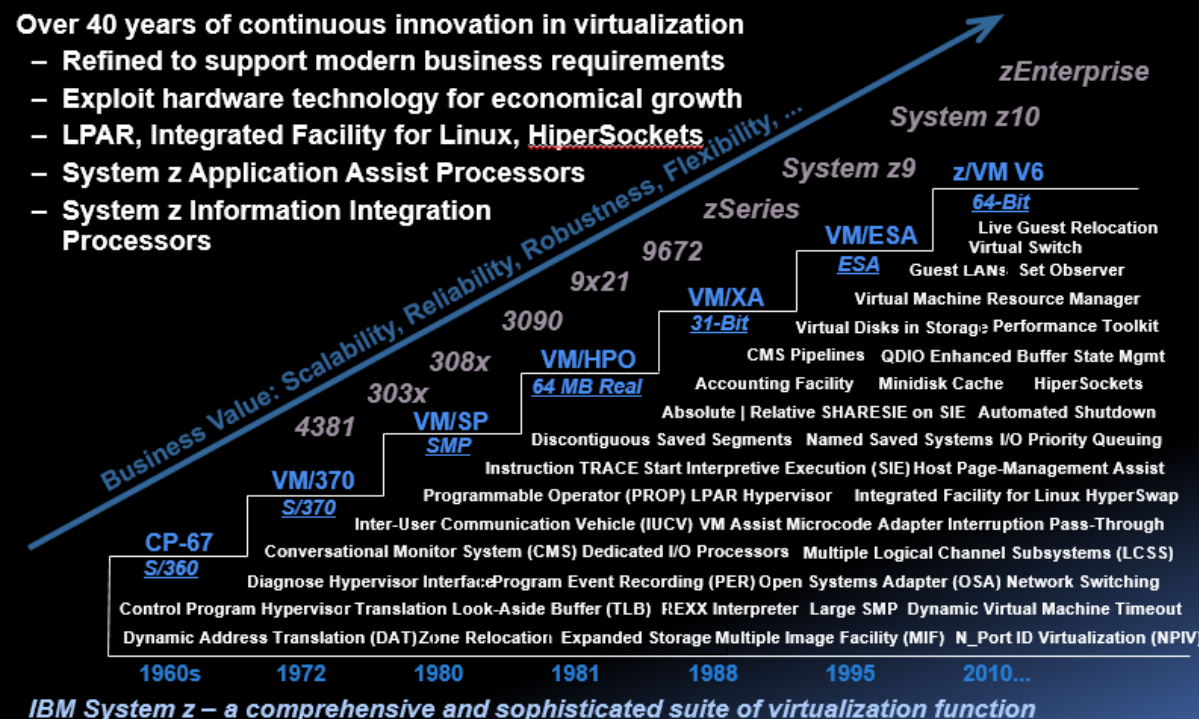
https://www.linux-kvm.org/page/Processor_support#S390

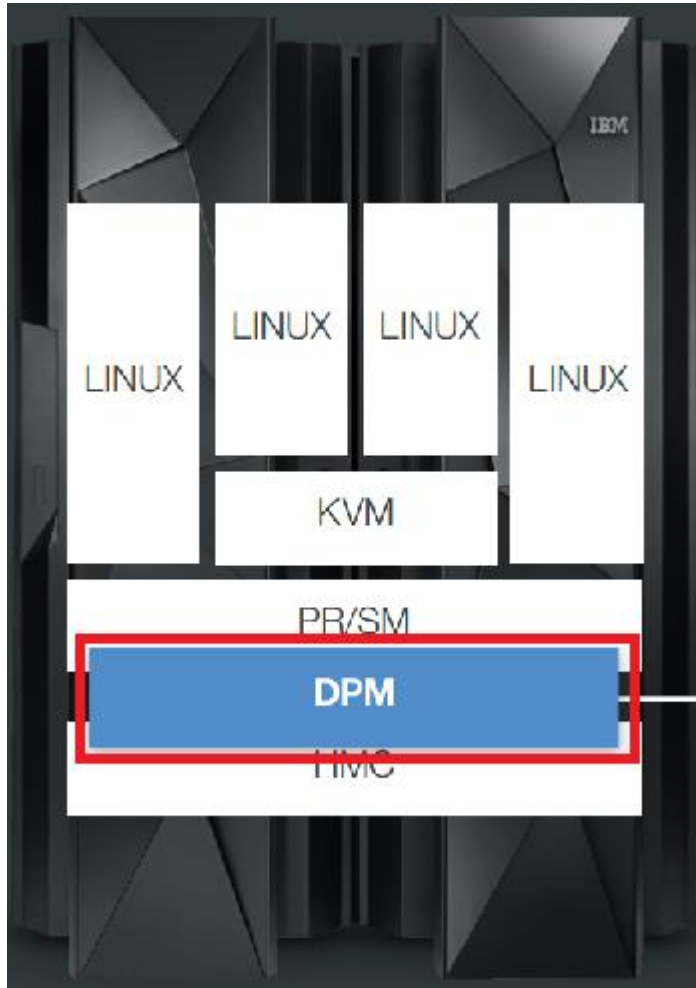
<https://specs.openstack.org/openstack/nova-specs/specs/kilo/implemented/libvirt-kvm-systemz.html>

IBM System z Virtualization Genetics

Over 40 years of continuous innovation in virtualization


- Refined to support modern business requirements
- Exploit hardware technology for economical growth
- LPAR, Integrated Facility for Linux, HiperSockets
- System z Application Assist Processors
- System z Information Integration Processors





- DPM
 - On IBM z Systems and IBM LinuxOne machines, most work loads runs faster in PR/SM hypervisor, than in a VM of a software hypervisor such as KVM or z/VM.
 - More like 'Ironic', it's baremetal solution on z
 - Reference
 - <http://nova-dpm.readthedocs.io/en/latest/>
 - <http://git.openstack.org/cgit/openstack/nova-dpm>

PR/SM = Processor Resource/System Manager



Linux bridge, macvtap



IBM DS8K, SVC
EMC VMAX
and more

https://www.ibm.com/developerworks/community/wikis/home?lang=en#!/wiki/W21ed5ba0f4a9_46f4_9626_24cbbb86fbb9/page/z%20Systems



- <https://www.suse.com/newsroom/post/2015/suse-offers-beta-preview-of-suse-openstack-cloud-6/>

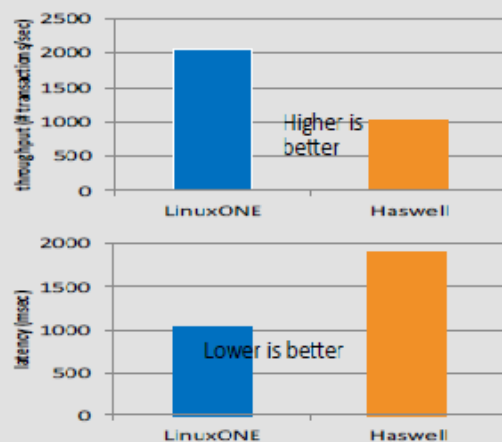


- <https://www.ubuntu.com/info/release-end-of-life>
- IBM own solution (openstack appliance), refstack compatible
 - <http://www.vm.ibm.com/sysman/osmntlvl.html>

z Systems and Containers - Facts

Extreme Virtualization with Containers


- A single LinuxONE Emperor ran more than **1 Million** containers
 - Workload: busybox httpd server (no NAT)
- LinuxONE Emperor runs **4K** containers on avg **2.0x** better than a compared Haswell-based system
 - Workload: Apache Solr
- LinuxONE Emperor can host over **10k** containers
 - Workload: 4k Apache Solr + 6k busybox httpd server (no NAT)



The throughput and response-time for a single Linux host running 4096 containers

Multi-Layer Auto Scaling

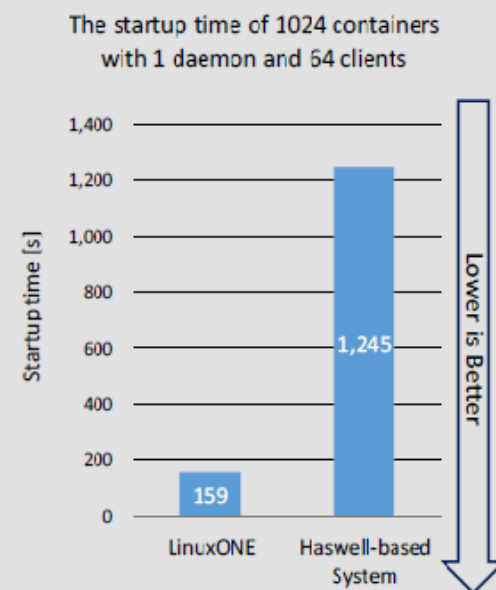
- Manage multiple virtualization layers to minimize the amount of resources to meet a SLA for a wide range of workload demand.
 - Start a set of containers when an application-level bottleneck is detected
 - Start a Docker Engine daemon in the same host when a daemon-level bottleneck is detected
 - Start an OS when an OS-level bottleneck is detected
 - Adjust the hardware resources such as CPU, memory, and I/O dynamically when a HW-level bottleneck is detected according to the workload demand



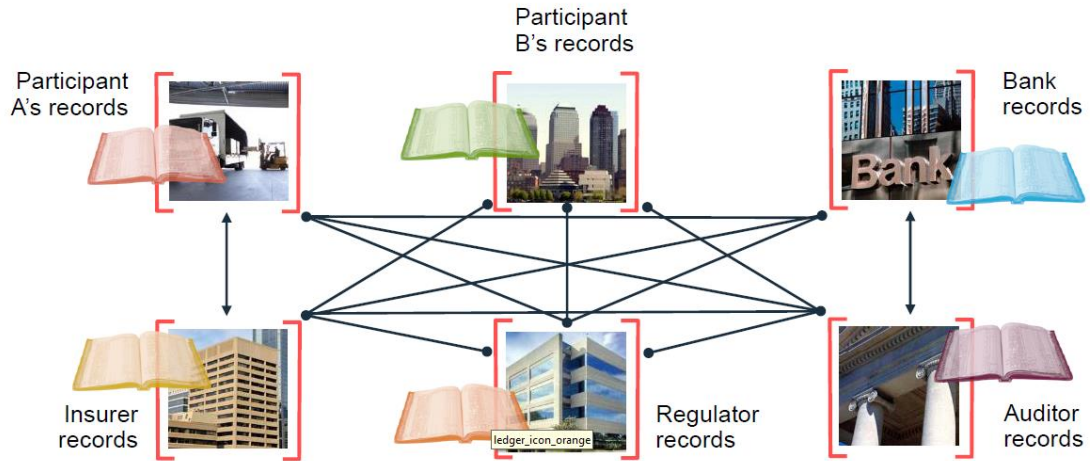
Financial Demo:
<https://www.youtube.com/watch?v=VWBNolwGEjo&t=649s>

Extreme Agility with Containers

- LinuxONE Emperor can start containers **7.8x faster** than a compared Haswell-based system.
 - Workload: nginx
- Significant agility to adapt to dynamic workload behavior




Problem




... inefficient, expensive, vulnerable

A shared replicated, permissioned ledger...



... with consensus, provenance, immutability and finality



Hyperledger Project Members



Premier

General



IBM Blockchain Offerings



All based upon Hyperledger Fabric

IBM Blockchain -aaS on Bluemix

Starter

- Start writing chaincode in seconds
- Integrated dashboard, logs and tools
- Community samples, tutorials, and quickstarts

Developers
v0.6 fabric

BETA

High Security Business Network

- Differentiated compute
- Best in Industry security, isolation
- Proven Audit environment for compliance

Single Organization **GA**
v0.6 fabric

Generally Available

High Security Business Network v1.0



- Create a business network
- Invite other Orgs to join
- Instantiate chaincode and begin transacting

Multiple Organizations **BETA**
v1.0 fabric

BETA

Self managed peers



* * any Docker environment
IBM offers technical support for x86, Power and System z

Support for Hyperledger Project fabric v0.6 **GA**

Generally Available

<https://hub.docker.com/r/ibmblockchain/fabric/>

Why Blockchain solutions on IBM z Systems?

- High Availability Requirements
- Open Standards and Linux
- Disaster Recovery Requirements
- Scalability for growth
- Economics of Linux (IFL) Specialty Engines
- TCO versus Total Cost of Acquisition
- z Systems have the highest security rating or classification for any commercial server (EAL4+)
- HW encryption
- Fast container deployment and management
- Colocation with existing applications (e.g. CICS based solution)
- Keeping data in place



THANK YOU