

ppc64le – An alternative architecture for the most discriminating workloads on OpenShift

Manoj Kumar

IBM

Michael Daitzman & Robert Baron

Mass Open Cloud

RED HAT SUMMIT

May 08, 2019



Topics

- ppc64le the pedigree
- ppc64le AND Red Hat OpenShift/OpenStack
- ppc64le for Data intensive workloads
- ppc64le for AI/ML
- ppc64le in the MASS OPEN CLOUD
- Demo

ppc64le the pedigree

ppc64le AND Red Hat







November 2018



Rank	System	Cores	Rmax (TFlop/s)	Rpeak (TFlop/s)	Power (kW)
1	Summit - IBM Power System AC922, IBM POWER9 22C 3.07GHz, NVIDIA Volta GV100, Dual-rail Mellanox EDR Infiniband , IBM DOE/SC/Oak Ridge National Laboratory United States	2,397,824	143,500.0	200,794.9	9,783
2	Sierra - IBM Power System S922LC, IBM POWER9 22C 3.1GHz, NVIDIA Volta GV100, Dual-rail Mellanox EDR Infiniband , IBM / NVIDIA / Mellanox DOE/NNSA/LLNL United States	1,572,480	94,640.0	125,712.0	7,438



Meet the POWER9 family

Big Data Workloads		Enterprise AI Workloads		Mission Critical Workloads		
LC922		 AC922		Power Scale-Out Systems		Power Enterprise Systems
				S922/S914/S924 H922/H924/L922		E950
						



ppc64le AND Red Hat OpenShift/OpenStack

Red Hat OpenStack Platform AND ppc64le



OpenStack 14
Available

Add Power compute nodes to your existing OpenStack infrastructure.

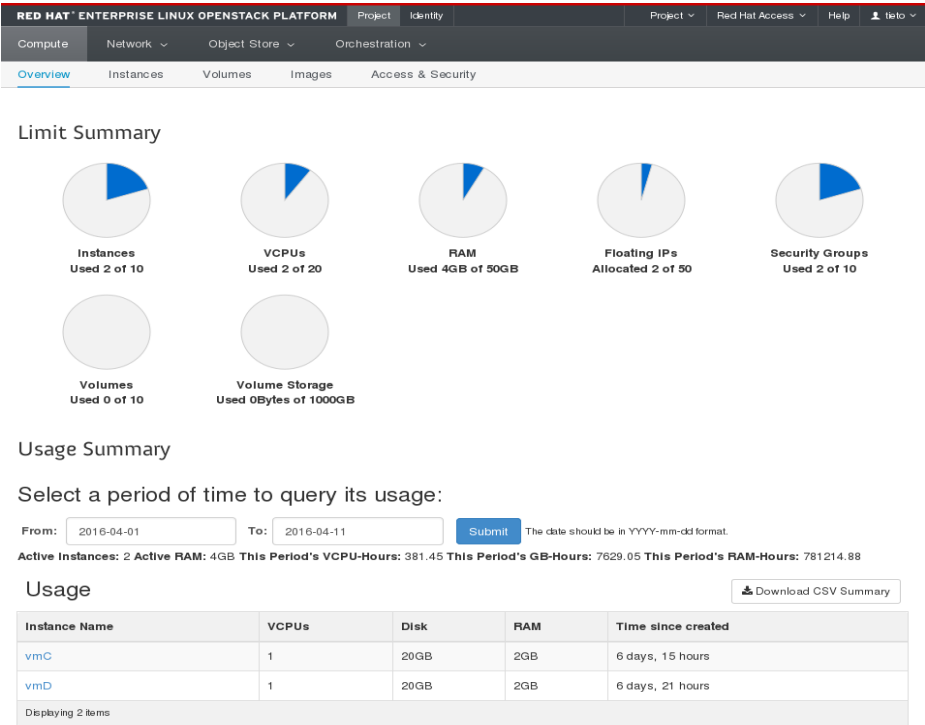
Deploy VM based applications faster on POWER for an on-premises Private Cloud.

Expand the capabilities of your Private Cloud, with matching infrastructure with the workloads (Fit for purpose).



POWER Cloud Benefits

- Manage POWER nodes just like any other compute infrastructure.
- Superior VM density and accelerated performance for modern data & AI workloads



Red Hat OpenShift AND ppc64le



**OpenShift 3.11
Available!**

PowerVM



Handles DevOps for cloud-native & traditional applications and services on a single platform

Containerize and manage existing apps, modernize on your own timeline, and work faster with new, cloud-native applications

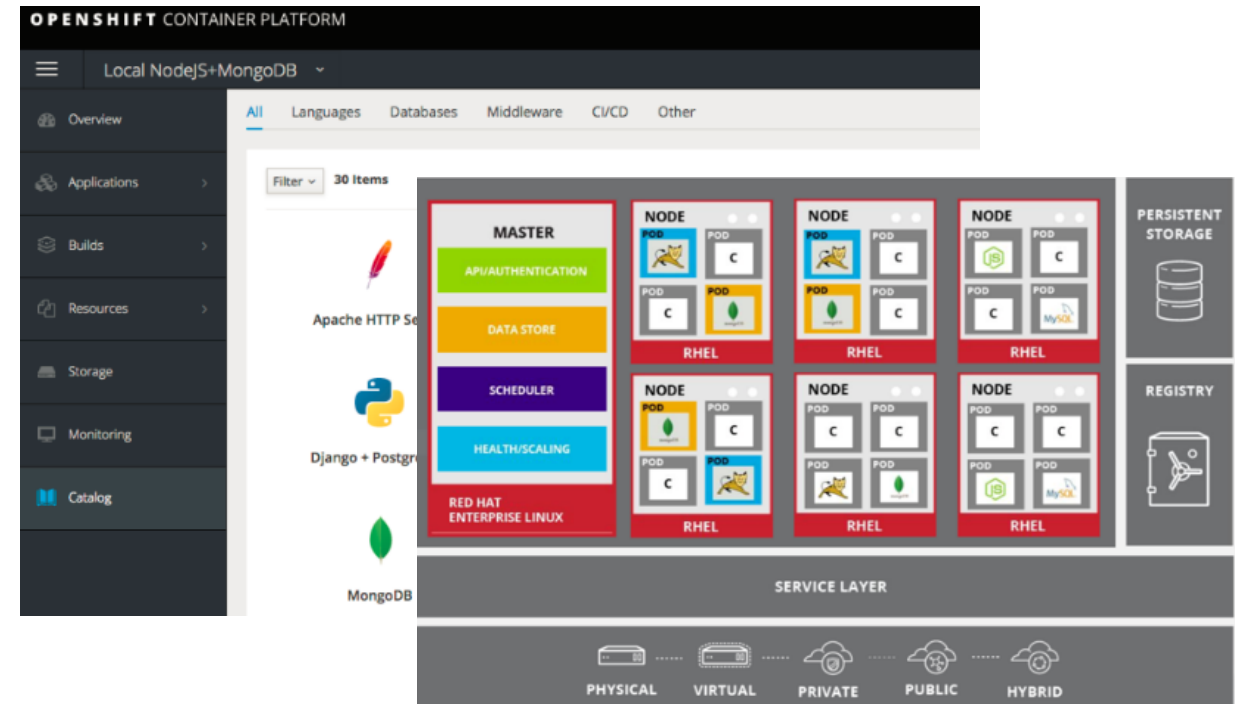
Deliver apps faster on Power for on-premises private cloud (bare-metal, PowerVM, and KVM)

Enabled on IBM enterprise systems

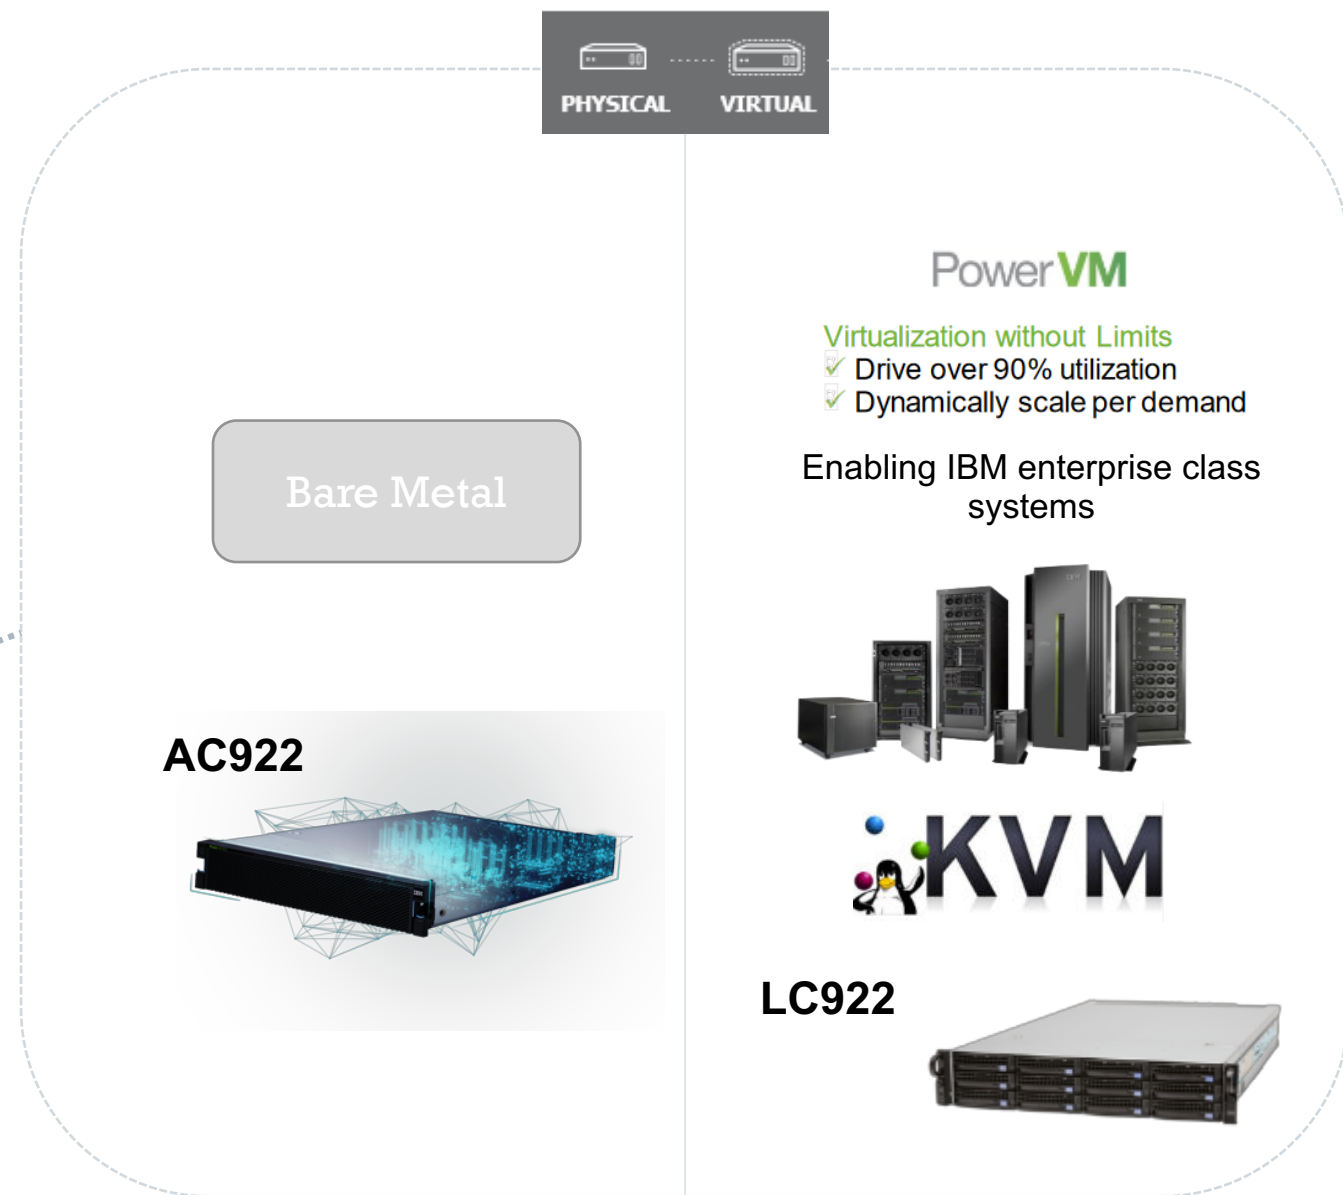
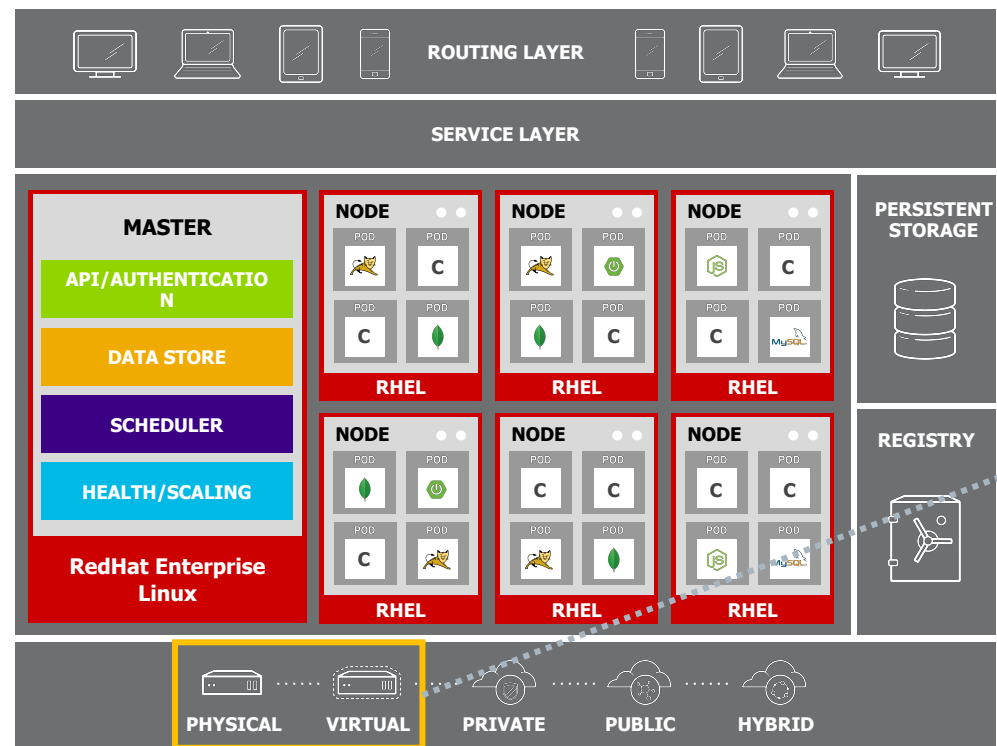


POWER9 Cloud Benefits

- Build, run cloud native apps with full control of Kubernetes on Power cloud (bare-metal, PowerVM or KVM)
- Superior container density and accelerated performance for modern data and AI workloads



Red Hat OpenShift AND ppc64le



ppc64le for Data intensive workloads

The Workload

(idea from MongoDB Manual → <https://docs.mongodb.com/manual/tutorial/geospatial-tutorial/>)

Collection #1 - Neighborhoods

- 195 rows
- 4 Megabytes raw
- Each row has an array of geospatial points that outline the neighborhood


Collection #2 - Restaurants

- 23,359 rows (restaurant listings)
- Each row has the location- single geospatial coordinate (Lat/Long) and the establishment's name
- 3.2 Megabytes raw



Each MicroService (Node.Js Application, Express, Mongoose) call runs 1 Query (Jmeter Script v4)

- Get 10 neighborhoods (MongoDB picks 10 and returns them)
- Get My Neighborhood (fixed location)
- GetNeighborhoods (pass in location coordinates)
- GetRestaurants located in a Donut shape around me (max & min range)
- GetRestaurants in range




Each “Jmeter user” calls all 5 micro services per loop. 4 users, 10 loops = 200 transactions

Short runs are about 8 mins long – 675,000 transactions

Long runs are about 30 mins long – 2,700,000 transactions

MongoDB on ppc64le

	IBM Power L922 (20-core, 256GB, 2 LPARs) 146 containers	Intel Xeon SP based 2-socket server (36-core, 256GB, 2 VMs) 72 containers
Server price ^{2,3,4} -3-year warranty	\$28,821	\$30,286
Solution Cost ⁵ -Server + RHEL OS + Virtualization + ICP Cloud Native VPC Annual Subscription @ \$250 per core per month x 36 months	\$ 218,938 (\$28,821 + \$10,117 + \$180,000)	\$ 358,205 (\$30,286 + \$3,919 + \$324,000)
Geospatial workload ¹ Total Transactions per Second - With 2 VM's	2,559 tps	2,217 tps
\$/container	\$1,500	\$4,975
Containers/core	7.3	2.0b

3.6X
Greater containers/core



3.3X
Better price-performance
(based on number of containers)

1. Based on IBM internal testing running MongoDB's Geospatial queries at 700 users, each running 1000 transactions using jmeter v4. Each container uses MongoDB 4.0.2 & Node.js v8.14.1 (REST APIs) with socket bound containers. Testing added containers to each server until servers reached response time limit of 99% of transactions completing in under 1 second. Results valid as of 2/7/19. Conducted under laboratory condition with speculative execution controls to mitigate user-to-kernel and user-to-user side-channel attacks on both systems, Individual result can vary based on workload size, use of storage subsystems & other conditions. Details about MongoDB workload: <https://docs.mongodb.com/manual/tutorial/geospatialtutorial/>

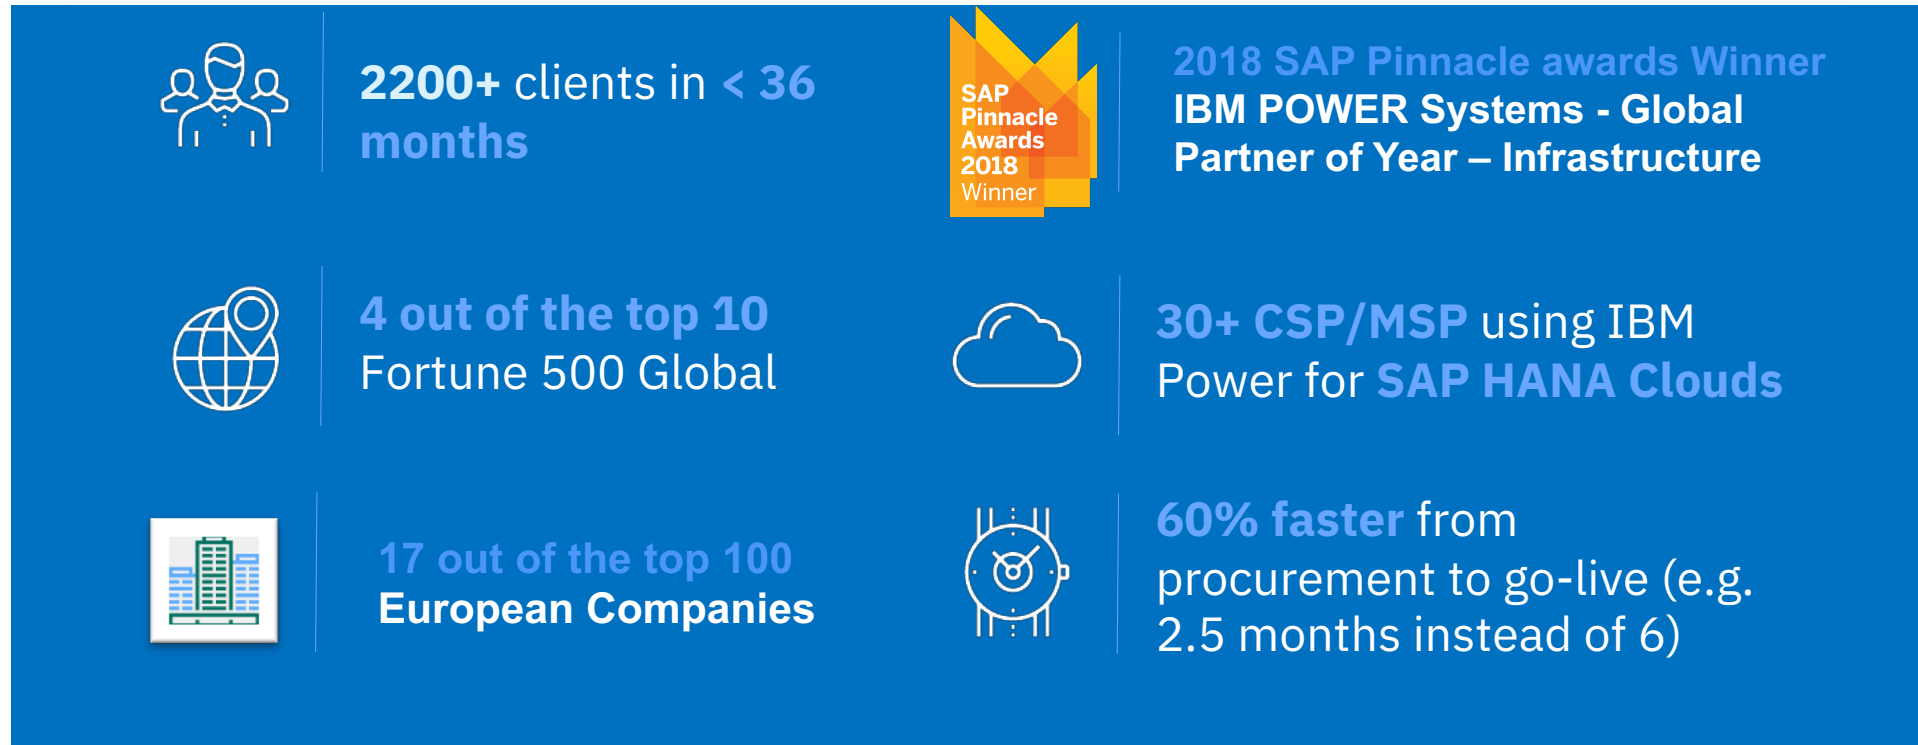
2. IBM Power L922 (2x10-core/typical 2.9 GHz/256 GB memory) 2x 388 GB SSD, 2x 10 Gb two-port network, RHEL 7.6 with PowerVM (2 partitions@10-cores each),

3. Competitive stack: 2-socket Intel Xeon Skylake Gold 6150 (2x18-core/ 2.7 GHz/256 GB memory), 2 x 480 GB SSD, 3 x 10 Gb two-port network, RHEL 7.6, KVM (2 VMs@18-cores each)

4. Pricing is based on Power L922 <https://www.ibm.com/it-infrastructure/power/scale-out>, and publically available x86 pricing.

5. IBM software pricing for ICP Cloud Native VPC Monthly Subscription .

SAP HANA clients are rapidly adopting Power Systems



Source: IBM data for companies using Power for HANA in Fortune Global 500, beta.fortune.com/global500/list/ and https://en.wikipedia.org/wiki/List_of_largest_European_companies_by_revenue

ppc64le for AI/ML

PowerAI – Open Source frameworks

Caffe
3.7x


Chainer
3.8x

 TensorFlow
2.3x

**TRAIN MORE. BUILD MORE.
KNOW MORE.**

Industry leading innovations

- Large Model Support
- Distributed Deep Learning
- Coherence
- NVLINK 2.0

POWER9

ppc64le in the MASS OPEN CLOUD

Demo – PowerAI on OpenShift

Thank you!

Red Hat OpenShift on PowerVM: http://www-01.ibm.com/common/ssi/ShowDoc.wss?docURL=/common/ssi/rep_sm/p/877/ENUS5639-OCP/index.html&lang=en&request_locale=en

IBM Model asset exchange: <https://developer.ibm.com/exchanges/models/>

IBM PowerAI Vision:

<https://www.ibm.com/us-en/marketplace/ibm-powerai-vision>

Manoj Kumar

Chief Engineer

IBM Power Systems

kumarmn@us.ibm.com

+1-512-286-7089