



IBM LinuxONE and Virtualization Technology Advancement



Rui Yang
System z Virtualization
IBM China System Lab

Agenda

- Background
- LinuxONE introduction
- LinuxONE virtualization technology
- IBM mainframe team's contribution to Linux kernel



Background



IBM Z[®] Generations

N-4



z10 Enterprise Class

- Announced 2/2008
- 4.4 GHz
- Up to 64 cfg cores
- CP, IFL, ICF, zAAP, zIIP
- Up to 1.5 TB Memory



z10 Business Class

- Announced 10/2008
- 3.5 GHz
- Up to 10 cfg cores (5 CP)
- CP, IFL, ICF, zAAP, zIIP
- Up to 248 GB Memory

N-3



zEnterprise 196

- Announced 7/22/2010
- 5.2 GHz
- Up to 80 cfg cores
- CP, IFL, ICF, zAAP, zIIP
- Up to 3 TB Memory



zEnterprise 114

- Announced 7/12/2011
- 3.8 GHz
- Up to 10 cfg cores (5 CP)
- CP, IFL, ICF, zAAP, zIIP
- Up to 248 GB Memory

N-2



zEnterprise EC12

- Announced 8/28/2012
- 5.5 GHz
- Up to 101 cfg cores
- CP, IFL, ICF, zAAP, zIIP
- Up to 3 TB Memory



zEnterprise BC12

- Announced 7/23/2013
- 4.2 GHz
- Up to 13 cfg cores (6 CP)
- CP, IFL, ICF, zAAP, zIIP
- Up to 496 GB Memory

N-1



IBM z13

- Announced 1/14/2015
- 5.0 GHz
- Up to 141 cfg cores
- CP, IFL, ICF, zIIP
- Up to 10 TB Memory



IBM z13s

- Announced 2/16/2016
- 4.3 GHz
- Up to 20 cfg cores (6 CP)
- CP, IFL, ICF, zIIP
- Up to 4 TB Memory

N



IBM z14

- Announced 7/17/2017
- 5.2 GHz
- Up to 170 cfg. cores
- CP, IFL, ICF, zIIP
- Up to 32 TB cfg. Memory



Linux is

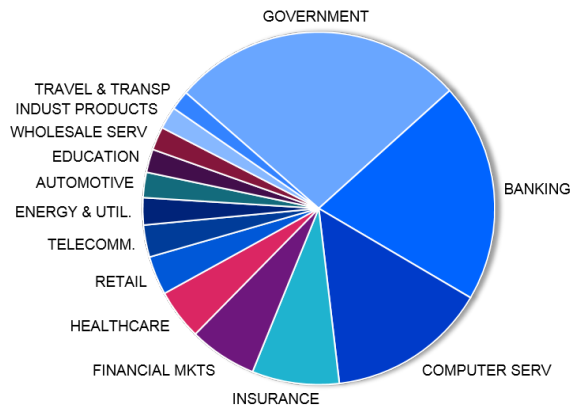
EVERYWHERE

Driving an

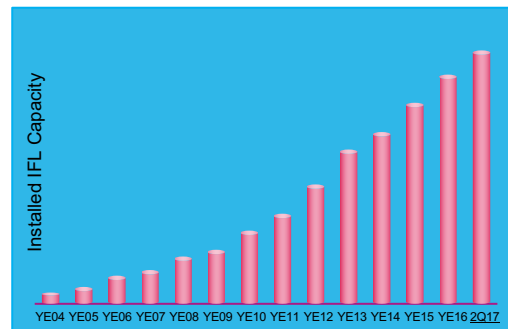
OPEN REVOLUTION

Linux on z

- The inception of Linux on z was around 2000
- Today, Linux on z is used in over 60 countries across 22 industries around the globe
- During the last decade, Linux on z grew rapidly



Installed Capacity Over Time



The new OPEN BREAKTHROUGH

The best of **IBM Z**

- Dynamic Resource Allocation
- Non-disruptive Scalability
- Continuous Business Availability
- Operational Efficiency
- Trusted Security
- Data and Transaction Serving

The best of **LINUX & OPEN**

- Freedom & Agility
- Standards based
- Speed to Innovate
- Developer Productivity
- Community Collaboration
- Open source SW & applications



Meet the IBM LinuxONE Systems

The most trusted, efficient and high performance
enterprise-grade Linux platform

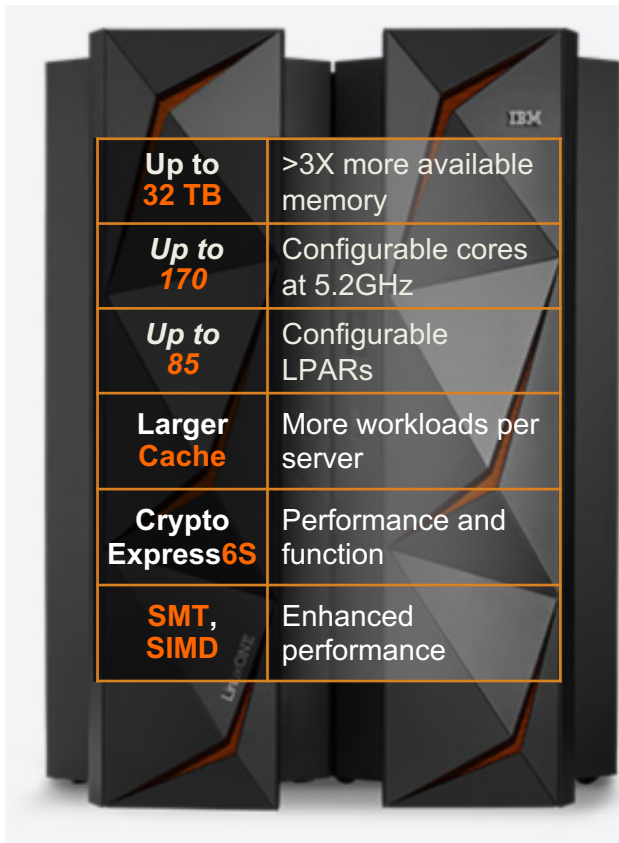


IBM **LinuxONE**
Emperor™



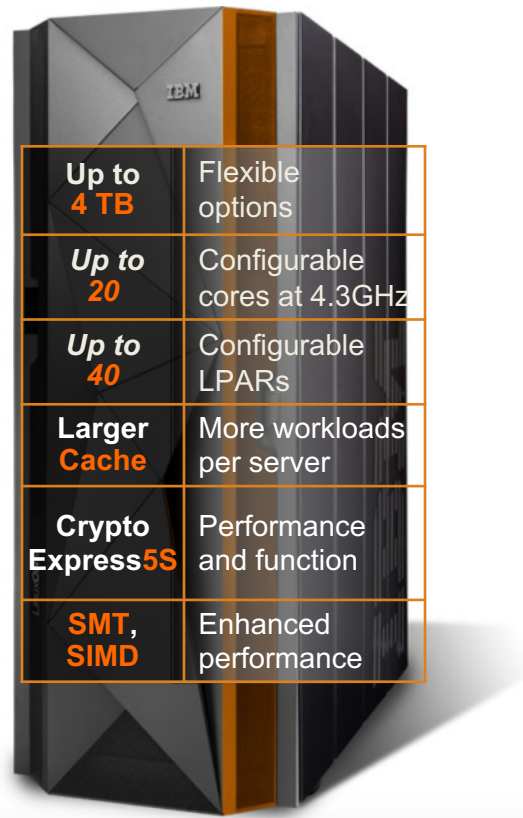
IBM **LinuxONE**
Rockhopper™

Emperor II



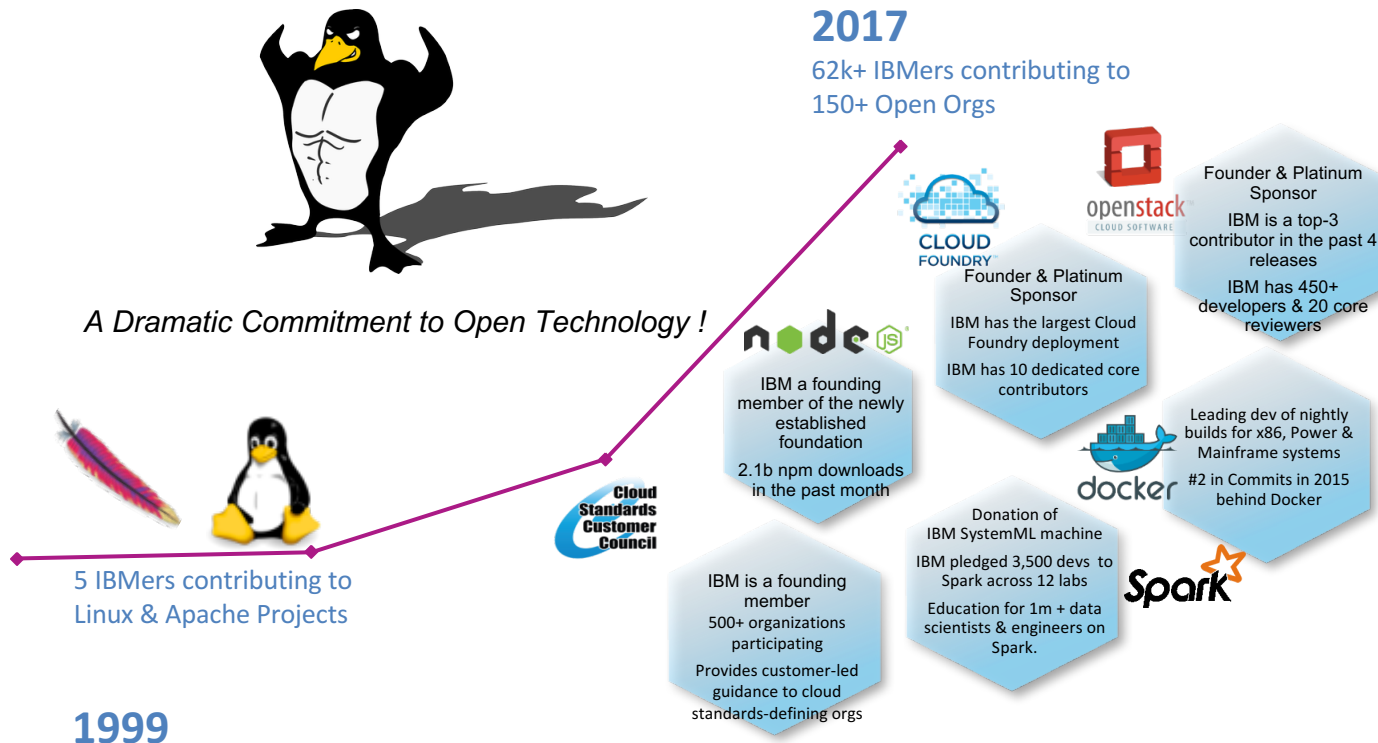
Up to 32 TB	>3X more available memory
Up to 170	Configurable cores at 5.2GHz
Up to 85	Configurable LPARs
Larger Cache	More workloads per server
Crypto Express6S	Performance and function
SMT, SIMD	Enhanced performance

IBM GDPS solutions <i>Continuous availability & Disaster recovery</i>
IBM Blockchain solution <i>The nature of transactions is being redefined</i>
IBM DBaaS solution <i>Database-as-a-service, leverage database technology</i>
IBM z/VM + IBM Wave <i>Virtualization with efficiency at scale + Intuitive virtualization management</i>
KVM on IBM Z <i>Open source virtualization</i>
Logical Partitions / Dynamic Partition Mgr <i>technology & capacity</i>
IBM LinuxONE <i>Unmatched server technology & capacity</i>



Up to 4 TB	Flexible options
Up to 20	Configurable cores at 4.3GHz
Up to 40	Configurable LPARs
Larger Cache	More workloads per server
Crypto Express5S	Performance and function
SMT, SIMD	Enhanced performance

Open Source participation is crucial to the IBM Strategy



Linux your Way - Greater flexibility and choice

Choose the distribution, runtime, hypervisor, database and analytics – it's the Linux you know and love with the openness, flexibility and agility you need for you business.

Distributions	Hypervisors	Languages	Runtimes	Management	Other	Database	Analytics
 Supported by Canonical   Community Versions    	  LPAR 	              	   OpenJDK     	    Cloud Manager      NGINX Apache ZooKeeper™ 	       WildFly   	       DB2   	        

LinuxONE virtualization technology



LinuxONE virtualization technology

▪ PR/SM-LPARs and IBM DPM

- Virtualization is built into the DNA of LinuxONE
- PR/SM manages and virtualizes all the installed and enabled system resources as a single large SMP system
- IBM Dynamic Partition Manager simplifies provisioning and management experience

▪ z/VM

- IBM proprietary server virtualization that is completely integrated into the full stack
- Complete hardware awareness
- VMSSI enables multiple z/VM systems to share and coordinate resources within a Single System Image (SSI) structure

▪ KVM

- KVM provides an open source choice for LinuxONE virtualization for Linux workloads. Best for clients that are Linux centric admins



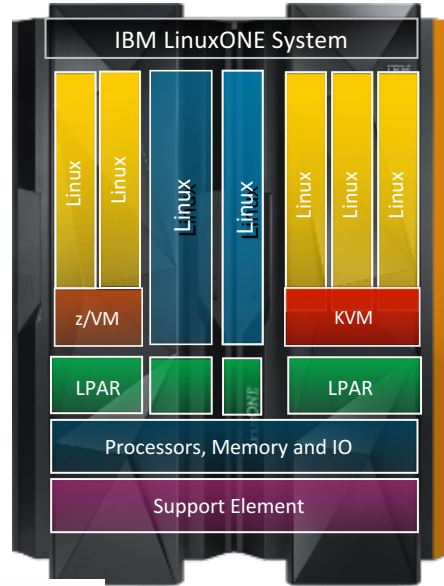
IBM LinuxONE Systems Virtualization Overview

LPAR virtualization

– Virtualization at the firmware level via **DPM** (graphical tool) or **PR/SM**

IBM z/VM

- World class quality, security, reliability - powerful and versatile
- Extreme scalability creates cost savings opportunities
- Exploitation of advanced technologies, such as:
 - Shared memory (Linux kernel, executables, communications)
- Highly granular control over-resource pool
- IBM Wave provides graphical “drag & drop” management interface



KVM virtualization

- Standardizes configuration and operation of server virtualization
- Leverage common Linux administration skills to administer virtualization
- Flexibility and agility leveraging the Open Source community
- Turns a Linux into an hypervisor
- available in your Distro:
 - SLES 12.2
 - UBUNTU 16.04



IBM mainframe team's contribution to Linux kernel

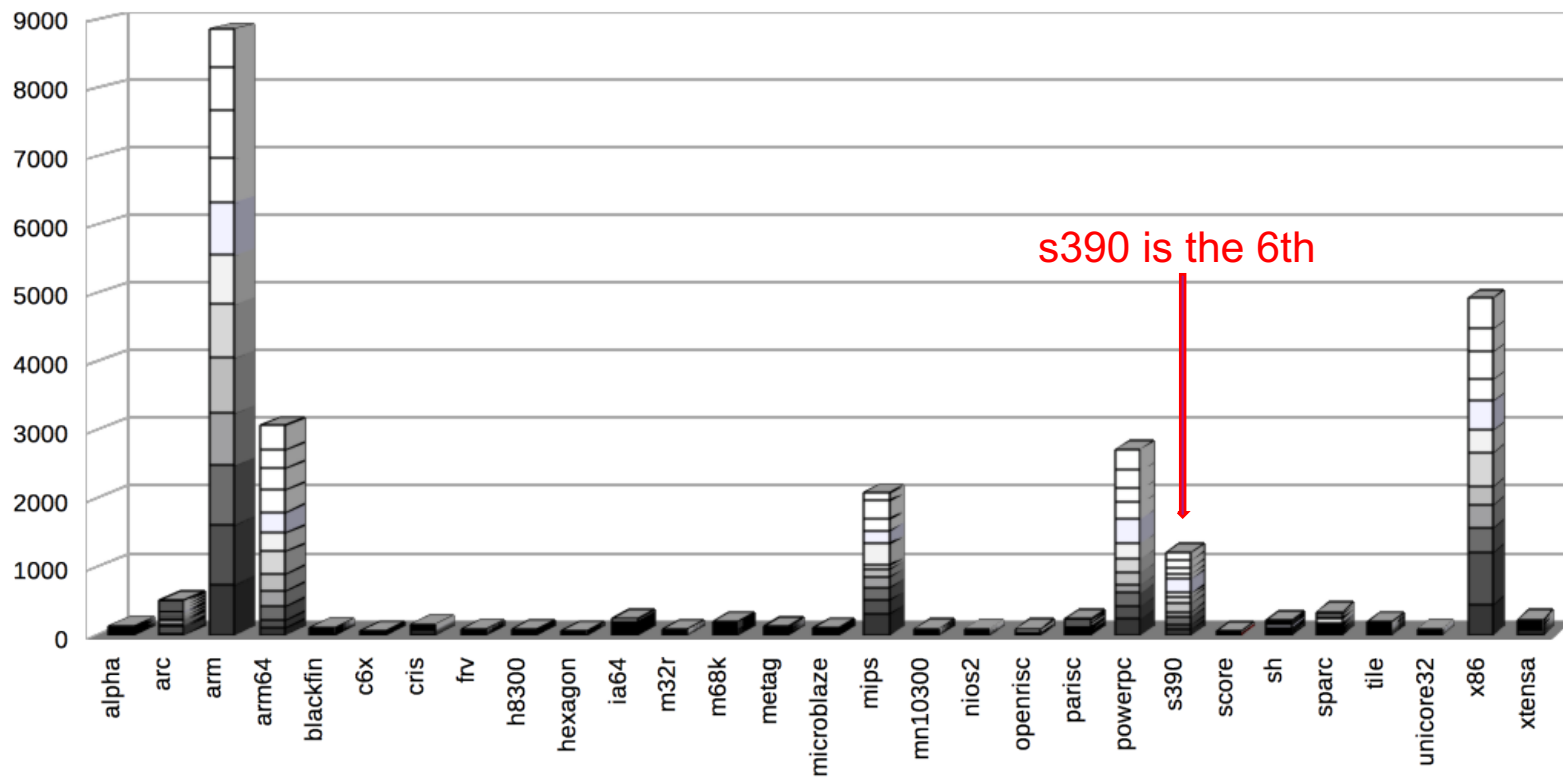


Some facts around IBM contribution to Linux Kernel

- How many git commits are there in the main Linux repository up to v4.12?
 - **679419** commits (630031 without merge commits, %7 merge commits)
- How many of these git commits are s390 related?
 - **~7354** commits (~1.1%)
- What is the maximum of lines added by a single git commit up to v4.12?
 - git commit d7e09d0397e84eef “staging: add **Lustre file system client support**”, **258994** insertions(+)
- What is the maximum of lines added by a single git commit for s390?
 - git commit 4a71df50047f0db6 “**new qeth device driver**”, **13498** insertions(+)
- What is the average size of a git commit in v4.x (patch lines)?
 - **148.33** over all git commits in v4.0 – v4.12
- What is the average size of a s390 commit in v4.x (patch lines)?
 - **215.03** lines for s390 related git commits in v4.0 – v4.12



Git commits per architecture in 4.x




Some selected Linux kernel features on LinuxONE

- **Toleration for Crypto Express 6 cards (kernel 4.10)**
 - Allow to use the new crypto hardware in CEX5 compat mode
- **Instruction execution protection (kernel 4.11)**
 - Also know as non-executable mappings or short “noexec”
 - New bits in the segment and page tables can be used to forbid code execution for a 1M segment or a 4K page
- **Support for the Guarded Storage Facility (kernel 4.12)**
 - Designed to improve the performance of Java while garbage collection is active
 - Up to 64 regions of memory can be marked as guarded
 - Reading a pointer with the new LGG or LLGFSG instruction will do a range check on the loaded value and automatically invoke a user space handler if one of the guarded regions is affected
- **True random number generator (kernel 4.12)**
 - The MSA-7 CPACF extension provides a new function for true random numbers



Some selected Linux kernel features on LinuxONE

- **Shared Memory Communications over RDMA (SMC-R)** is a protocol that allows applications to exploit RDMA (RoCE) with the socket interface
 - RDMA technology provides the capability to allow hosts to logically share memory. The SMC-R protocol defines a means to exploit the shared memory for communications - transparent to the applications
 - A first version of the Linux code is now upstream with kernel 4.11-rc1



The world is changing,
SO ARE WE



Questions?



Rui Yang

Senior Software Engineer

IBM China System Lab

ruiyangr@cn.ibm.com



Notices and Disclaimers

Copyright © 2016 by International Business Machines Corporation (IBM). 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 according to 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 is in compliance with any law

Notices and Disclaimers Con't.

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 in connection with 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 PARTICULAR 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, Aspera®, Bluemix, Blueworks Live, CICS, Clearcase, Cognos®, DOORS®, Emptoris®, Enterprise Document Management System™, FASP®, FileNet®, Global Business Services®, Global Technology Services®, IBM ExperienceOne™, IBM SmartCloud®, IBM Social Business®, Information on Demand, ILOG, Maximo®, MQIntegrator®, MQSeries®, Netcool®, OMEGAMON, OpenPower, PureAnalytics™, PureApplication®, pureCluster™, PureCoverage®, PureData®, PureExperience®, PureFlex®, pureQuery®, pureScale®, PureSystems®, QRadar®, Rational®, Rhapsody®, Smarter Commerce®, SoDA, SPSS, Sterling Commerce®, StoredIQ, Tealeaf®, Tivoli®, Trusteer®, Unica®, urban{code}®, Watson, WebSphere®, Worklight®, X-Force® and System z® Z/OS, 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.