

Transformation For Growth

Compete in
the Era of
SMART.

2014 Consultants &
System Integrators Interchange

Linux on Power Systems Trends and Directions

Dipankar Sarma
Distinguished Engineer, Linux Technology Center
IBM India Systems & Technology Lab



IBM

Agenda

1. Linux on Power Overview

2. Scale-out computing on Power

3. Future of Linux scale-out computing on Power



Agenda

1. Linux on Power Overview

2. Scale-out computing on Power

3. Future of Linux scale-out computing on Power



Linux momentum is growing on Power Systems



An open ecosystem, using the POWER Architecture, to **share expertise, investments, and server-class intellectual property** to serve customers' evolving needs.




Linux Centers. Offering access to hardware technical support, porting assistance, demos, toolkits, hands-on labs

Beijing | Austin | New York | Montpellier | Tokyo

500+

ISVs
enabled to
run on Linux
on Power

Global Linux ISV Development NEW. Power Development Cloud gives Linux developers who want to prototype applications access to Power
NEW. 50,000 new developers from Ubuntu development community



Regional Ecosystem Initiative

Recruiting key solutions:

- + Open Source Tools
- + Middleware
- + Industry Solutions

\$ 1B

in new Linux and open source technologies for IBM Power Systems.

Linux is Linux is Linux

Linux on Power Runs the Same Commands as Linux on x86

- **Supports Red Hat and SUSE Enterprise Linux versions consistent with x86_64**
 - POWER support available simultaneously with other platforms
 - List of packages nearly identical (except minor differences like bootloader)
 - Packages at same version/level – including kernel and device drivers
- **Leverage same open source system tools**
 - SDK: Same Free Eclipse-based development environment
<http://www14.software.ibm.com/webapp/set2/sas/f/lopdiags/sdklop.html>
 - Advance Toolchain: Same Open Source tools (GNU), IBM tested and supported on Power
- **New Little Endian platform architecture same as x86**
 - Canonical Ubuntu Linux on Power
 - SUSE SLES



Agenda

1. Linux on Power Overview

2. Scale-out computing on Power

3. Future of Linux scale-out computing on Power



Designed for
Big Data



CAPI Technology



Superior Cloud
Economics



POWER8



Open
Innovation
Platform



OpenPOWER™

Announcing a new generation of Power Systems

- **A new generation of POWER8 systems: The best Scale-out systems in the industry**
- **System Software: An intelligent IT infrastructure for Cloud**
 - Opening up the world of Linux
- **Optimized Solutions: Cloud, Big Data & Analytics, Mobile, Java and Linux applications**
 - Expanded portfolio of applications and faster time to value
- **Enterprise Pools: Transforming Enterprise IT for the Cloud**
 - POWER7+ Enterprise Pools with investment protection to POWER8



Bringing Power to Scale-out Computing

	Sandy Bridge EP	Ivy Bridge EP E5-26xx v2	Ivy Bridge EX E7-88xx v2	POWER 7+ Systems	POWER8
Clock rates	1.8-3.6GHz	1.7-3.7GHz	1.9-3.4 GHz	3.1-4.4 GHz	3.0-3.9 GHz
SMT options	1,2*	1, 2*	1, 2*	1, 2, 4	1, 2, 4, 8
Max Threads / sock	16	24	30	32	96
Max L1 Data Cache	0	32KB*	32KB*	32KB	64KB
Max L2 Cache	256 KB	256 KB	256 KB	256 KB	512 KB
Max L3 Cache	20 MB	30 MB	37.5 MB	80 MB	96 MB
Max L4 Cache	0	0	0	0	128 MB
Memory Bandwidth	31.4-51.2 GB/s	42.6-59.7 GB/s	68-85** GB/s	100 – 180 GB/sec	230 - 410 GB/sec

* Intel calls this Hyper-Threading Technology (No HT and with HT)

**32KB running in "Non-RAS mode" Only 16KB in RAS mode

**85GB running in "Non-RAS mode" = dual-device error NOT supported



POWER8 Scale-out Systems

IBM Scale-out Power Systems

- ❑ POWER8 roll-out is leading with scale-out (1-2S) systems
- ❑ Expanded Linux focus: Ubuntu, KVM, and Open Stack
- ❑ Scale-up POWER8 (>2S) systems will be rolled out over time
- ❑ PCI Gen3 right out of POWER8 processor
- ❑ OpenPOWER Innovations



Power Systems S812L

- 1-socket, 2U
- POWER8 processor
- Linux only
- CAPI support (1)
- 2H14



Power Systems S822L

- 2-socket, 2U
- POWER8 processor
- Up to 24 cores
- 1 TB memory
- 9 PCI Gen3 slot
- Linux only
- CAPI support (2)
- PowerVM & PowerKVM



Power Systems S822

- 2-socket, 2U
- Up to 20 cores
- 1 TB memory
- 9 PCIe Gen 3
- AIX & Linux
- CAPI support (2)
- PowerVM



Power Systems S814

- 1-socket, 4U
- Up to 8 cores
- 512 GB memory
- 7 PCIe Gen 3
- AIX, IBM i, Linux
- CAPI support (1)
- PowerVM



Power Systems S824L

- 2-socket, 4U
- Up to 24 cores
- Linux
- NVIDIA GPU
- 2H14



Power Systems S824

- 2-socket, 4U
- Up to 24 cores
- 1 TB memory
- 11 PCIe Gen 3
- AIX, IBM i, Linux
- CAPI support (2)
- PowerVM



PowerSC

PowerVC

PowerHA

PowerVP

PowerKVM

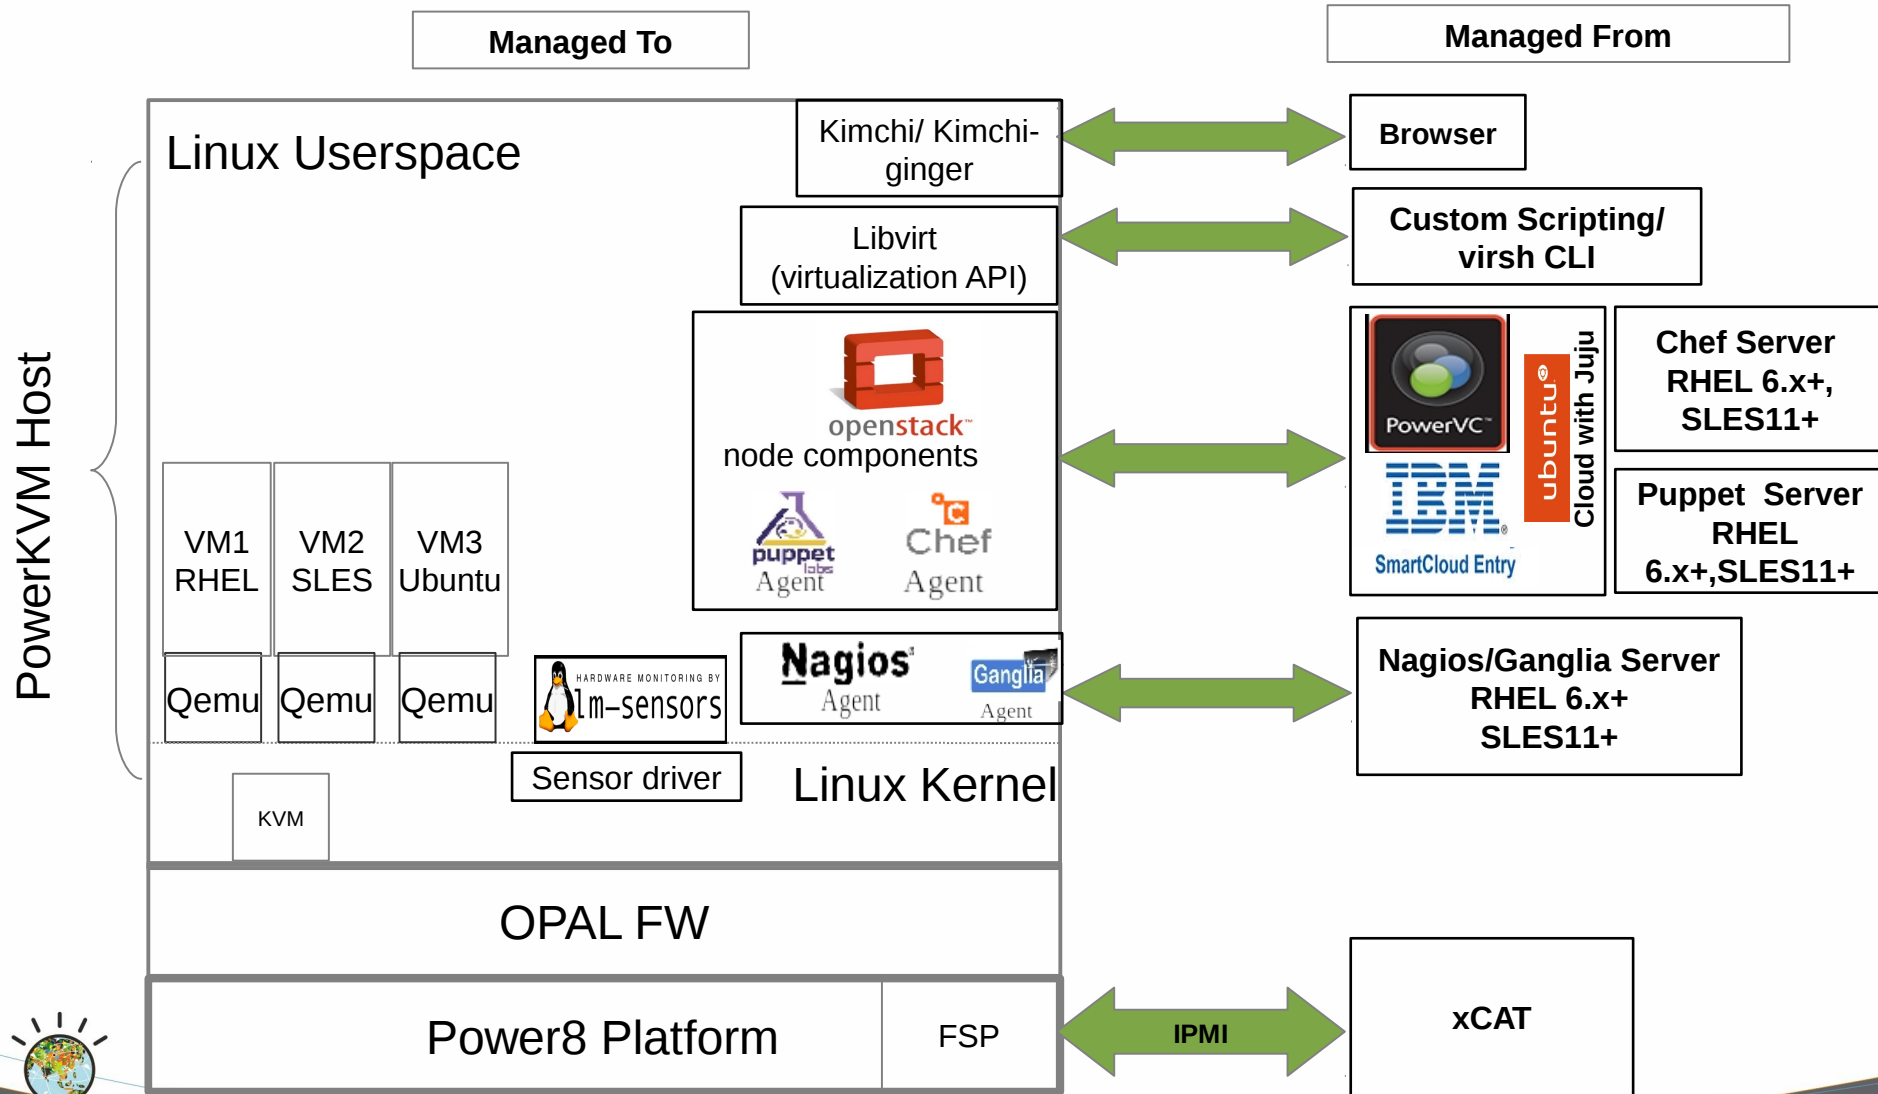
ubuntu
Supported by Canonical

SUSE

redhat

IBM

PowerKV System for Cloud



Agenda

1. Linux on Power Overview


2. Scale-out computing on Power

3. Future of Linux scale-out computing on Power



Open Platform: Ecosystem Leadership Providing Client Choice & Flexibility

OpenPOWER Foundation



- Collaborative innovation for highly advanced systems
- Produce open hardware, software, firmware and tools
- Expand industry skills and investment for Power ecosystem
- Provide alternative architectures

Open Applications and Tools



JavaScript



- Access to industry innovation from a broad community
- Optimize scripting languages & development tools
- Contribute innovation to Linux, KVM and OpenStack
- Lift and shift java apps with LE Canonical Ubuntu & SUSE Linux

Open Management



- Simplified, consistent management
- Deployment on premise or via cloud
- Seamless upgrades
- Software Defined Environments



OpenPOWER will enable data centers to rethink their approach to technology.

Member companies may use POWER for custom open servers and components for Linux based cloud data centers.

OpenPOWER ecosystem partners can optimize the interactions of server building blocks – microprocessors, networking, I/O & other components – to tune performance.

Buy
Build
Innovate

Platinum Members



ALTERA

Google

IBM



SAMSUNG

Suzhou
PowerCore
Technology

TYAN

IBM

How will the OpenPOWER Foundation benefit clients?

- OpenPOWER technology creates greater choice for customers
- Open and collaborative development model on the Power platform will create more opportunity for innovation
- New innovators will broaden the capability and value of the Power platform

What does this mean to the industry?

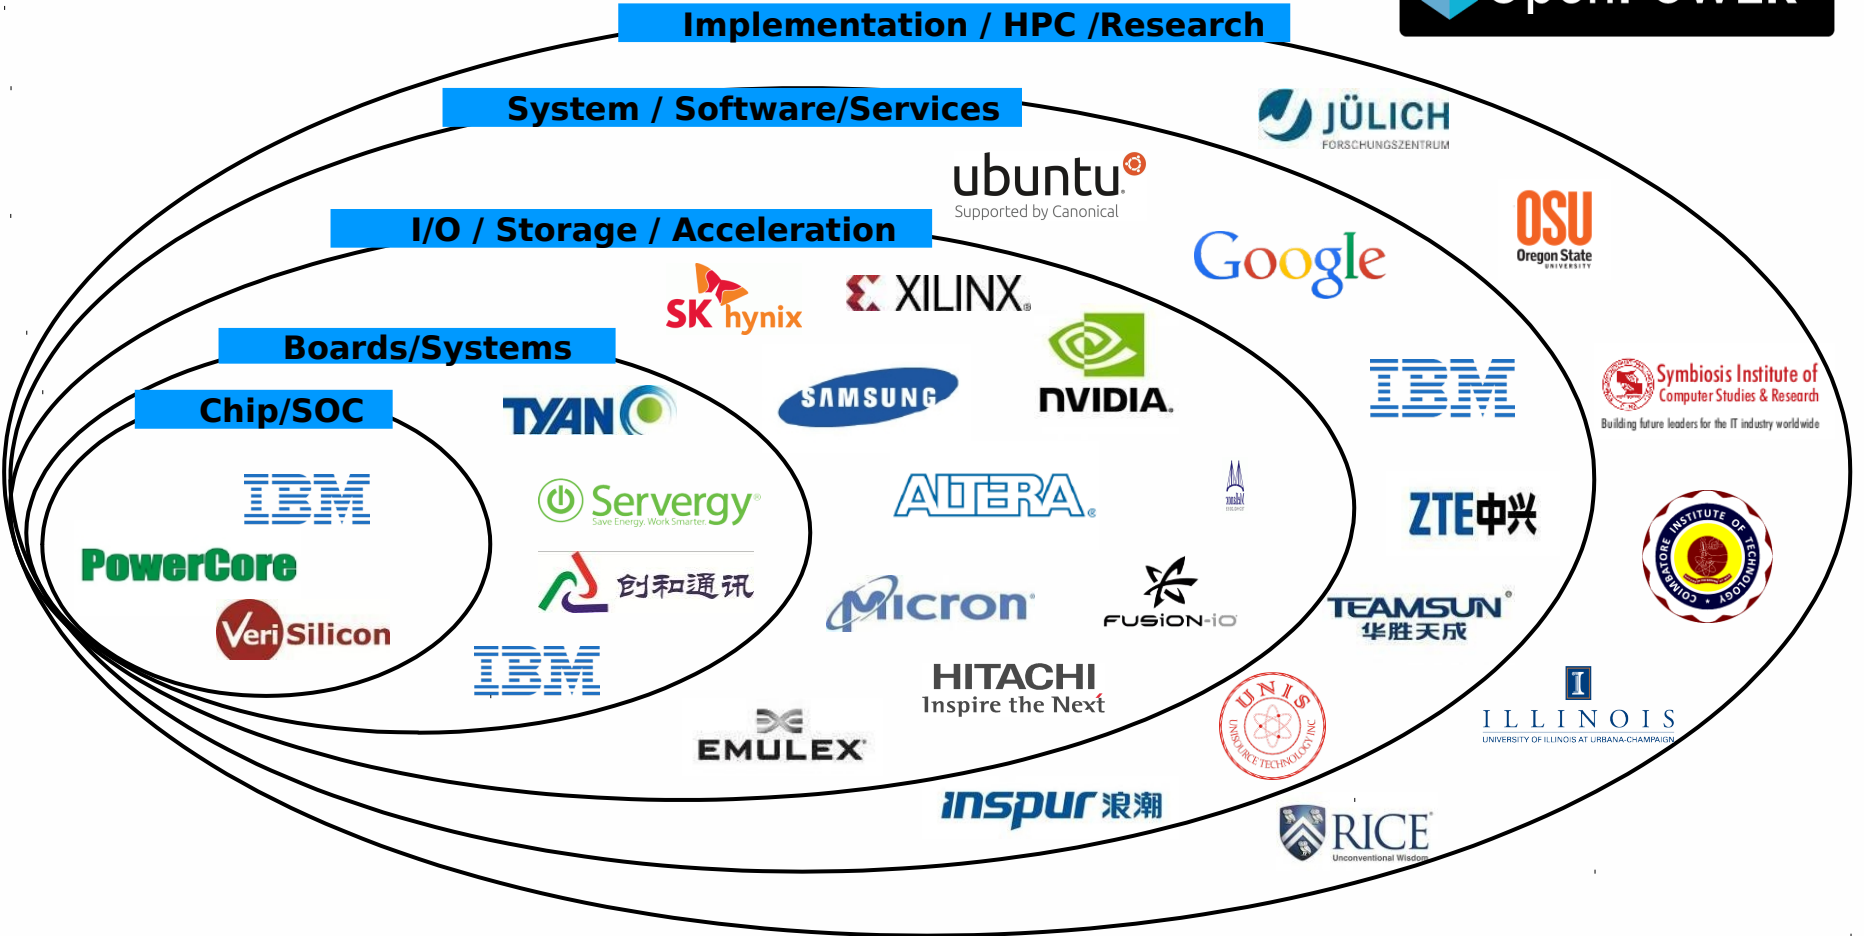
- **Game changer** on the competitive landscape of the server industry
- Will enable and drive innovation in the industry
- Provide more choice in the industry

Transformation For Growth

2014 Consultants &
System Integrators Interchange

Building collaboration and innovation at all levels

Compete in
the Era of
SMART.



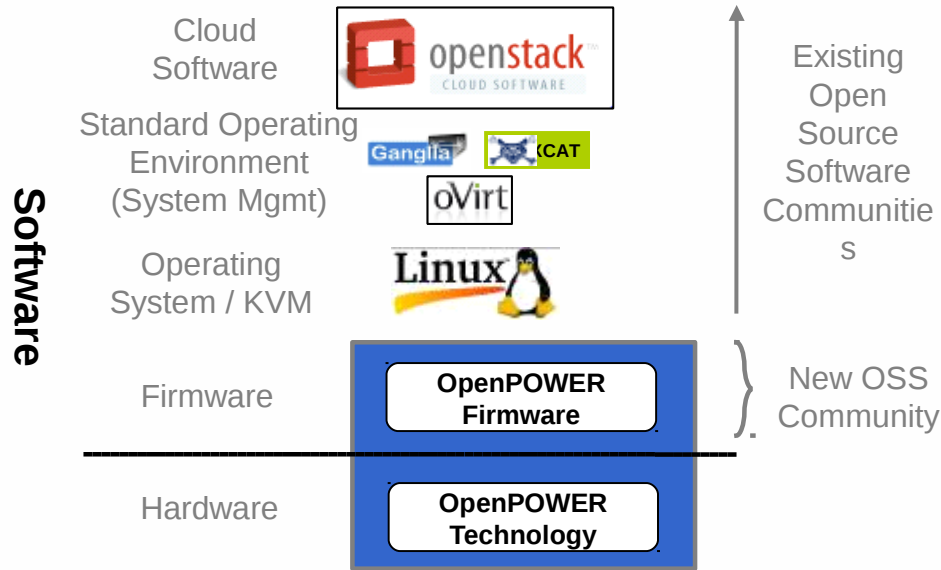
Welcoming new members in all areas of the ecosystem

100+ inquiries and numerous active dialogues underway



OpenPower Hardware & Software

Power Open Source Software Stack Components

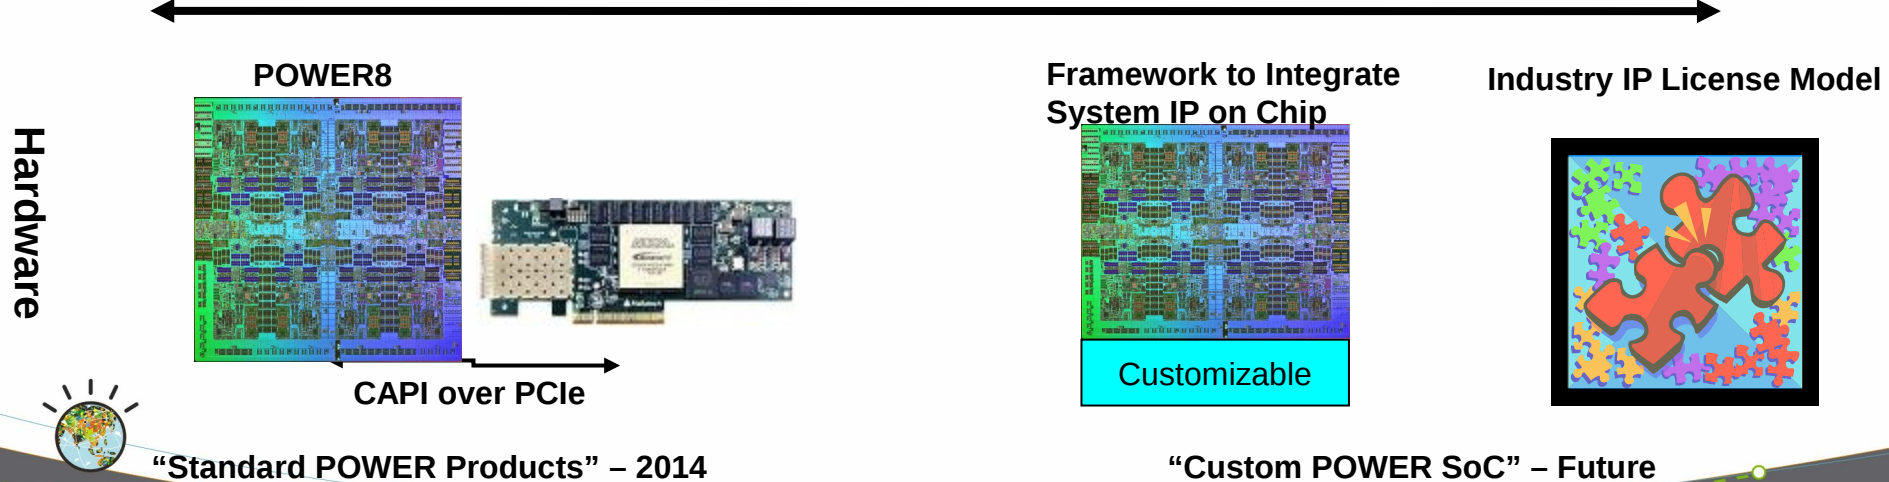


System Operating Environment Software Stack

A modern development environment is emerging based on tools and services



Multiple Options to Design with POWER Technology Within OpenPOWER



Non-IBM POWER8 products




The Tyan reference (ATX) board, SP010, measures 12" by 9.6"

- › one single-chip module (SCM)
- › four DDR3 memory slots
- › four 6 Gb/sec SATA peripheral connectors
- › two USB 3.0 ports
- › two Gigabit Ethernet network interfaces
- › keyboard and video
- › intended for developers

The Google reference board

- › two single-chip module (SCM)
- › four modified SATA ports
- › Google use only



IBM Contribution to Linux

Top 20 contributor affiliations (since 2006)	#changes	% of total
Red Hat	31,261	11.9%
SUSE	16,738	6.4%
Intel	16,219	6.2%
IBM	16,073	6.1%
Consultant	7,986	3.0%
Oracle	5,542	2.1%
Academia	3,421	1.3%
Nokia	3,272	1.2%
Fujitsu	3,156	1.2%
Texas Instruments	2,982	1.1%
Broadcom	2,916	1.1%
Linux Foundation	2,890	1.1%
Google	2,620	1.0%
Analog Devices	2,595	1.0%
SGI	2,578	1.0%
AMD	2,510	1.0%
Parallels	2,419	0.9%
Freescale	2,265	0.9%
Cisco	2,259	0.9%
HP	2,158	0.8%

Source: “Linux Kernel Development” report from Linux Foundation, March 2012

<http://www.linuxfoundation.org/publications/linux-foundation>

“None” or “Unknown” affiliations are not included.



Linux support for POWER

- Built from the same source as x86
 - Supported at the same time as x86
 - Delivered on the same schedule as x86
- Close development relationship with IBM



▪ RHEL 7

- Available for existing RHEL customers
- POWER8 (native mode) and POWER 7/7+ at GA
- LE / baremetal host support in future

▪ RHEL 6

- POWER8 supported with U5 (P7-compatibility mode)
- Full support of POWER6 and POWER7
- (native mode)

▪ Fedora

- Fedora supports POWER, actively developed
- Fedora 20 has POWER8 support

▪ Supported add-ons

- JBoss
- High Performance Network Add-on
- More SW in future

▪ SLES 11

- POWER8 with SP3 (P7-compatibility mode)
- POWER7+ encryption, RNG accelerators with SP3
- Full support of POWER7 (native mode)
- LE and baremetal host support in future

▪ SLES 10

- POWER7 supported with SP3 (P6-compatibility mode)
- Full support of POWER6 (native mode)

▪ openSUSE

- openSUSE 12.2 re-launched for IBM POWER
- openSUSE 13.2 includes POWER8 support

▪ Supported add-ons

- SUSE Linux Enterprise High Availability Extension
- More SW in future

▪ Ubuntu 14.04

- POWER8 enabled (native mode)
- No official support for POWER7+ and older systems
- No support for 32-bit applications. 64-bit only.
- Supported in KVM only at this time
- Baremetal / host supported as tech preview, official support in near future

▪ Supported add-ons


- Ubuntu openstack
- JuJu Charms
- MaaS (Metal as a Service)
- Landscape

▪ Debian



- Community enablement, officially supported architecture



Linux on Power – Future Direction



Open Platform for Choice



Special Notices

This document was developed for IBM offerings in the United States as of the date of publication. IBM may not make these offerings available in other countries, and the information is subject to change without notice. Consult your local IBM business contact for information on the IBM offerings available in your area.

Information in this document concerning non-IBM products was obtained from the suppliers of these products or other public sources. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. Send license inquires, in writing, to IBM Director of Licensing, IBM Corporation, New Castle Drive, Armonk, NY 10504-1785 USA.

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

The information contained in this document has not been submitted to any formal IBM test and is provided "AS IS" with no warranties or guarantees either expressed or implied.

All examples cited or described in this document are presented as illustrations of the manner in which some IBM products can be used and the results that may be achieved. Actual environmental costs and performance characteristics will vary depending on individual client configurations and conditions.

IBM Global Financing offerings are provided through IBM Credit Corporation in the United States and other IBM subsidiaries and divisions worldwide to qualified commercial and government clients. Rates are based on a client's credit rating, financing terms, offering type, equipment type and options, and may vary by country. Other restrictions may apply. Rates and offerings are subject to change, extension or withdrawal without notice.

IBM is not responsible for printing errors in this document that result in pricing or information inaccuracies.

All prices shown are IBM's United States suggested list prices and are subject to change without notice; reseller prices may vary.

IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.

Any performance data contained in this document was determined in a controlled environment. Actual results may vary significantly and are dependent on many factors including system hardware configuration and software design and configuration. Some measurements quoted in this document may have been made on development-level systems. There is no guarantee these measurements will be the same on generally-available systems. Some measurements quoted in this document may have been estimated through extrapolation. Users of this document should verify the applicable data for their specific environment.



Legal Statement

- Copyright International Business Machines Corporation 2014.
- Permission to redistribute in accordance with IBM CSI Interchange 2014 submission guidelines is granted; all other rights reserved.
- This work represents the view of the authors and does not necessarily represent the view of IBM
- IBM, IBM logo, ibm.com, POWER, POWERVM are trademarks of International Business Machines Corporation in the United States, other countries, or both.
- Openstack is a registered trademark of Openstack Foundation.
- Intel is a trademark or registered trademark of Intel Corporation or its subsidiaries in the United States and other countries.
- Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.
- Other company, product, and service names may be trademarks or service marks of others.
- References in this publication to IBM products or services do not imply that IBM intends to make them available in all countries in which IBM operates.
- INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you. This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

