

**IBM Connect 2015**

Seize the Moment

06/02/15

# *The OpenPOWER Foundation*

*How IBM, Google and 120 others are  
reinventing the datacenter*

**Franz Bourlet**

Power Systems Technical Sales

IBM Belgium & Luxembourg



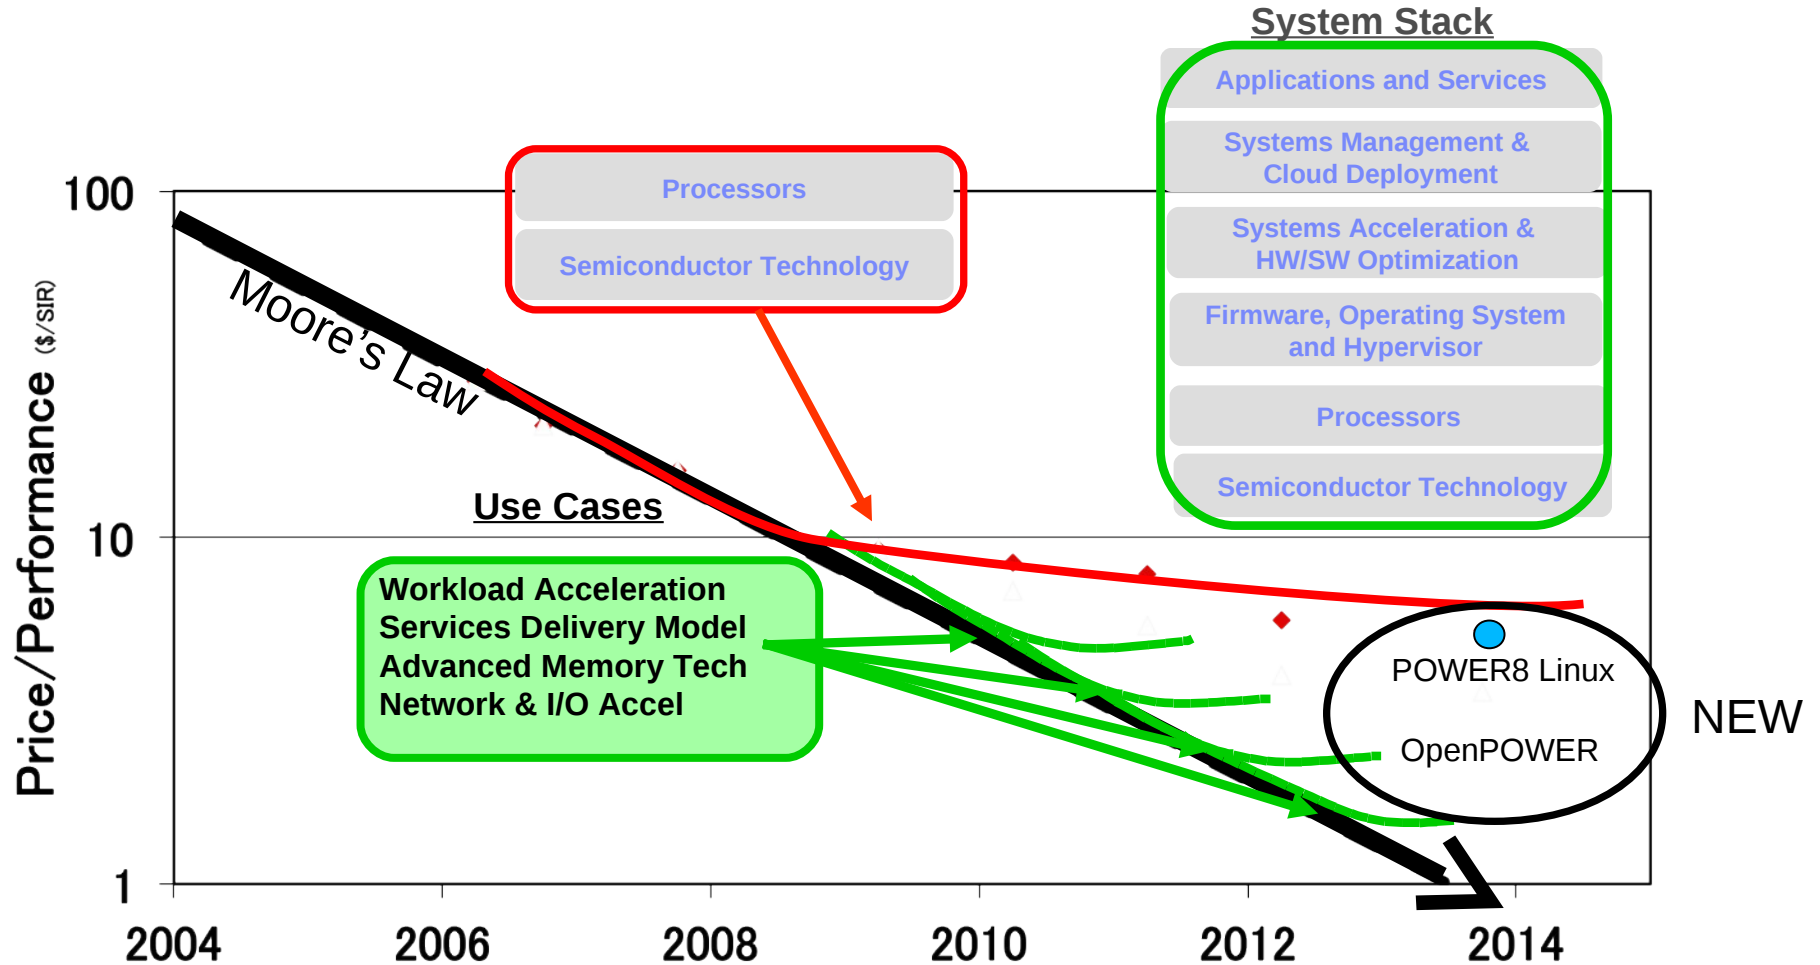
## New ways of acquiring servers

### **The number of companies designing and building servers is increasing :**

- Traditionally there have been few companies designing systems: HP, IBM, SUN, Dell, etc.
- Today there are many more: Google, Microsoft, Facebook, Rackspace, Huawei, Sugon, Inspur, etc.
- A fairly mature ecosystem including the Taiwanese ODMs is a key enabler of this trend

# Industry trends drive innovation beyond the chip

Microprocessors alone no longer drive sufficient Cost / Performance improvements



System stack innovations are required to drive Cost/Performance

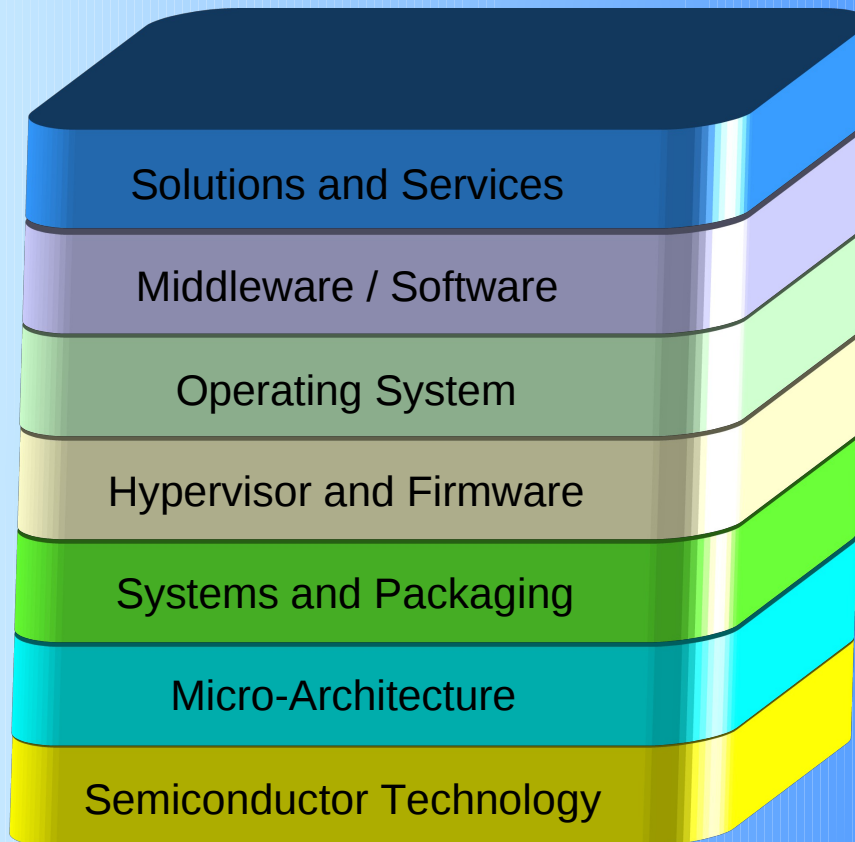
# ARM Ltd : another economic model

- Founded in November 1990 – [www.arm.com](http://www.arm.com)
  - Spun out of Acorn Computers
- Runs 95% of mobile phones and tabs
- Arm designs a range of RISC processor cores but does not fabricate silicon itself
- Licenses ARM core designs to semiconductor partners who fabricate and sell to their customers
  - Examples include Apple, Samsung, Nvidia, Qualcomm,...
- Also develops technologies to assist with the design-in of the ARM architecture
  - Software tools, boards, debug hardware, application software, bus architectures, peripherals, etc



POWER has evolved into the ARM of the datacenter : from this...

### IBM Research: Innovation Strengthens the Stack



#### IBM is Unique in the World

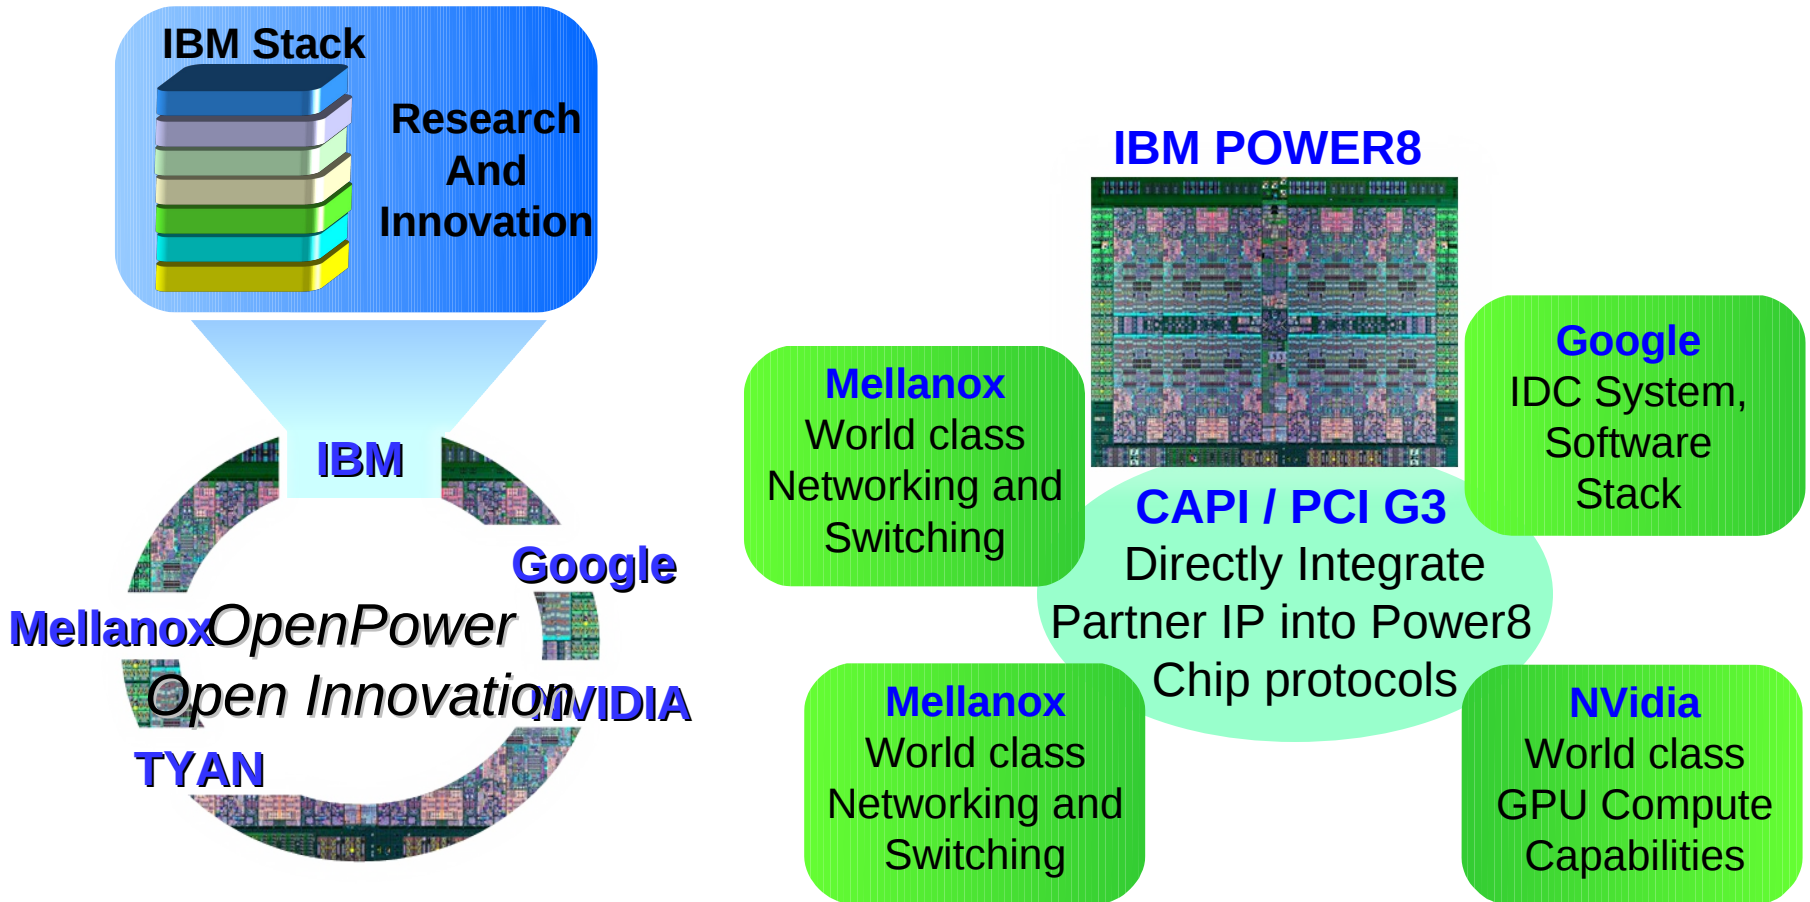
- No other company plays in all the layers. IBM is world-class in all layers
- Those that play in multiple layers are world-class only in few.

#### IBM R&D investment

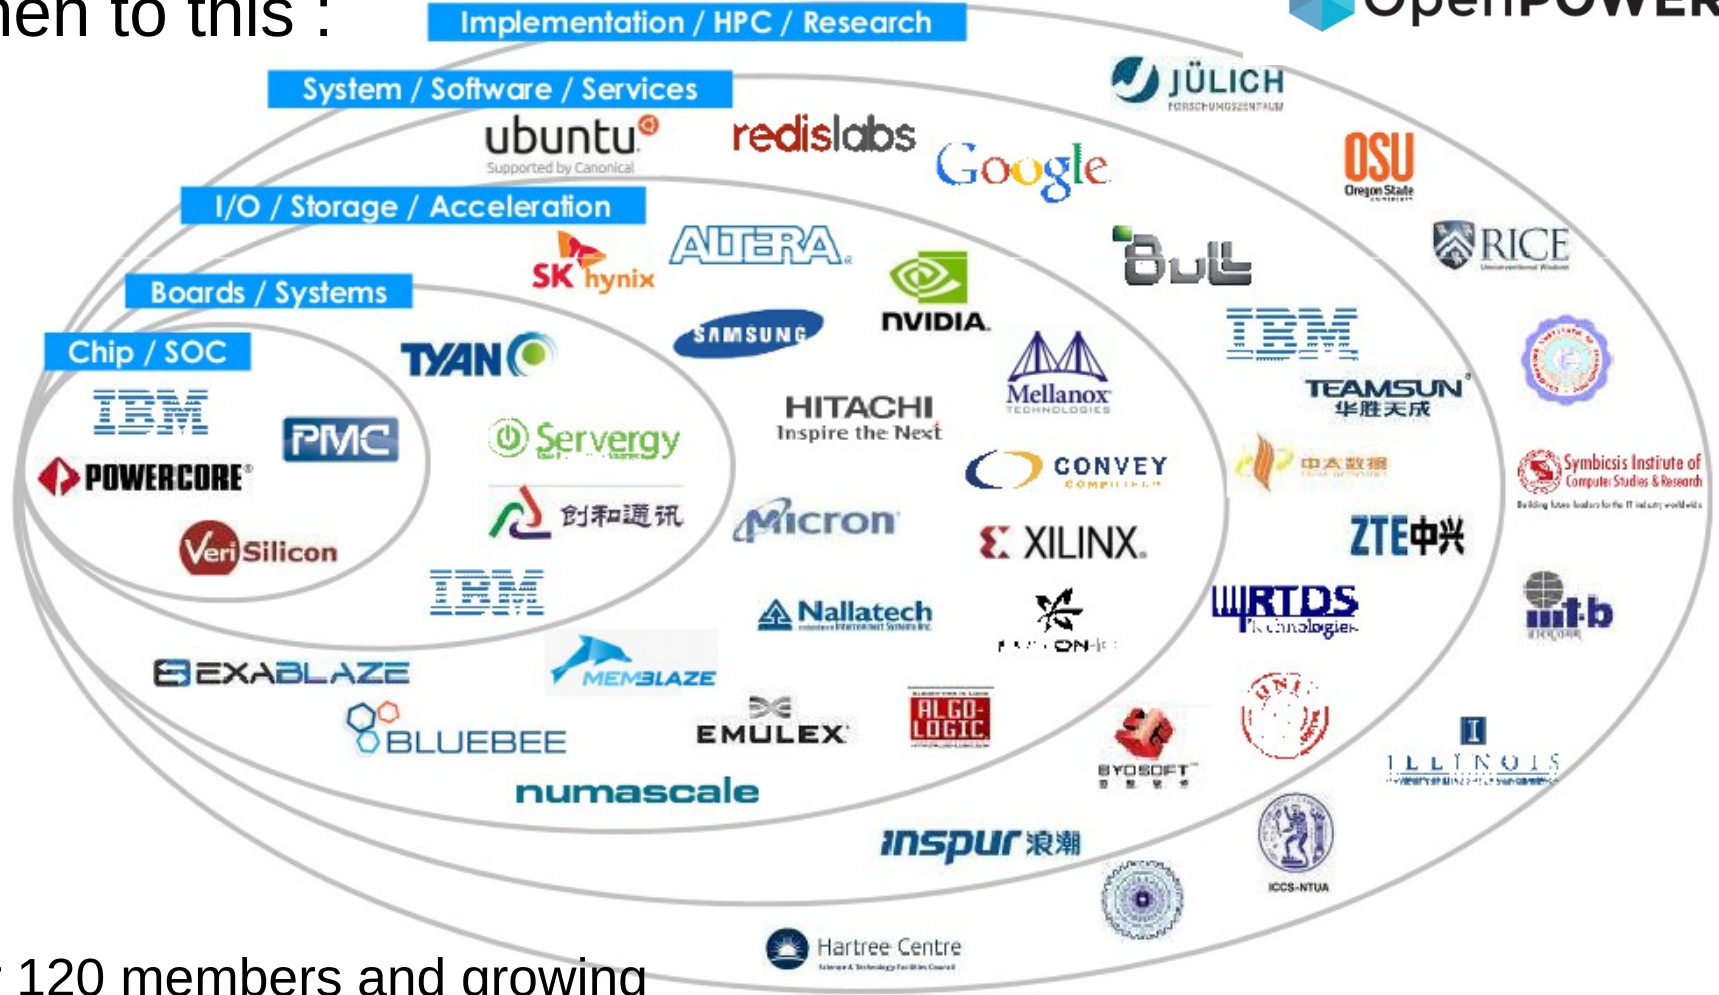
- \$6 Billion annually
- Human capital and culture
- 3000+ Researchers
- 6478 U.S. patents in 2012
- Patent leader past 20 years

As CMOS scaling slows, Value comes from Innovating across the Stack

To this...



Then to this :

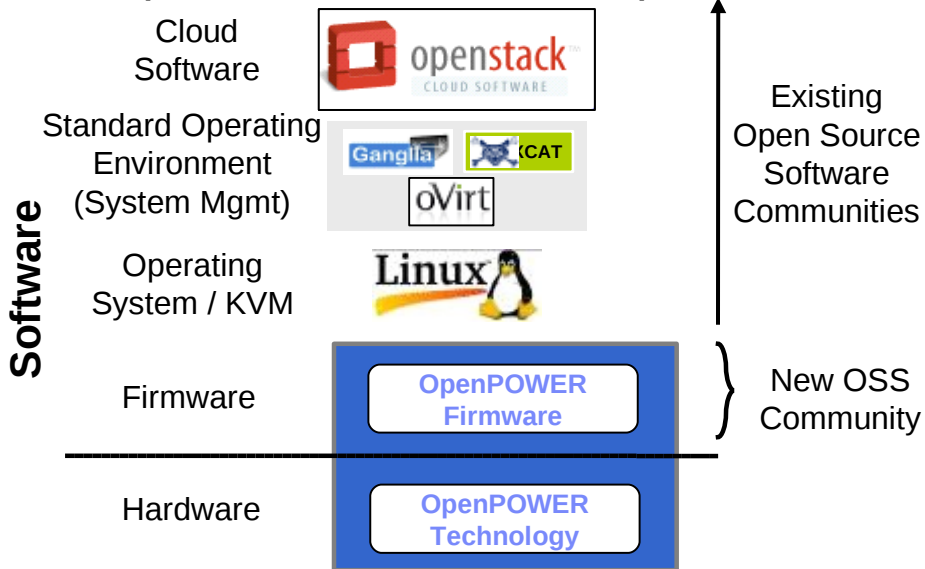


Over 120 members and growing

The goal of the OpenPOWER Foundation is to create an open ecosystem, using the POWER Architecture to share expertise, investment and server-class intellectual property to serve the evolving needs of customers.

# The most open server architecture in the industry

## Power Open Source Software Stack Components




## System Operating Environment Software Stack

A modern development environment is emerging based on this type of tools and services




## Multiple Options to Design with POWER Technology Within OpenPOWER

Hardware



“Standard POWER Products” – 2014

Framework to Integrate System IP on Chip



“Custom POWER SoC” – Future

Industry IP License Model





# POWER8 vs Intel Haswell

	Sandy Bridge EP E5-x6xx	Ivy Bridge EP E5-26xx v2	Ivy Bridge EX E7-88xx v2	Haswell EP E5-26xx v3	POWER 7+ Systems	POWER8
Clock rates (GHz)	1.8-3.6	1.7-3.7	1.9-3.4	1.6-3.5	3.1-4.4 GHz	3.0-4.1 GHz
SMT options	1,2*	1, 2*	1, 2*	1, 2*	1, 2, 4	1, 2, 4, 8
Max Threads / sock	16	24	30	36	32	96
Max L1 Cache	32KB	32KB*	32KB*	64 KB	32KB	64KB
Max L2 Cache	256 KB	256 KB	256 KB	256KB	256 KB	512 KB
Max L3 Cache	20 MB	30 MB	37.5 MB	45 MB	80 MB	96 MB
Max L4 Cache	0	0	0	0	0	128 MB
Memory Bandwidth	31.4-51.2 GB/s	42.6-59.7 GB/s	68-85** GB/s	51-68 GB/s	100 - 180 GB/sec	230 - 410 GB/sec

1 = The Micro-architecture features above contributed to premium performance for the POWER8 systems compared to Ivy Bridge systems

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

**\*32KB running in "Non-RAS mode" 16KB results in better RAS**

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

# POWER8 vs Intel Haswell

	Sandy Bridge EP E5-x6xx	Ivy Bridge EP E5-26xx v2	Ivy Bridge EX E7-88xx v2	Haswell EP E5-26xx v3	POWER 7+ Systems	POWER8
Clock rates (GHz)	1.8-3.6	1.7-3.7	1.9-3.4	1.6-3.5	3.1-4.4 GHz	3.0-4.1 GHz
SMT options	1, 2*	1, 2*	1, 2*	1, 2*	1, 2, 4	1, 2, 4, 8
Max Threads / sock	16	24	30	28	32	96
Max L1 Cache	32KB	32KB	32KB	32KB	32KB	64KB
Max L2 Cache	256 KB	307.2 KB	307.2 KB	256 KB	256 KB	512 KB
Max L3 Cache	20 MB	30 MB	30 MB	30 MB	80 MB	96 MB
Max L4 Cache	0	0	0	0	0	128 MB
Memory Bandwidth	31.4-51.2 GB/s	42.6-59.7 GB/s	42.6-59.7 GB/s	31-68 GB/s	100 - 180 GB/sec	230 - 410 GB/sec

How fast the engine

1 = The Micro-architecture features above contributed to premium performance for the POWER8 systems compared to Ivy Bridge systems

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

\*32KB running in "Non-RAS mode" 16KB results in better RAS

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

# POWER8 vs Intel Haswell

	Sandy Bridge EP E5-x6xx	Ivy Bridge EP E5-26xx v2	Ivy Bridge EX E7-88xx v2	Haswell EP E5-26xx v3	POWER 7+ Systems	POWER8
Clock rates (GHz)	1.8-3.6	1.7-3.7	1.9-3.4	1.6-3.5	3.1-4.4 GHz	3.0-4.1 GHz
SMT options	1, 2*	1, 2*	1, 2*	1, 2*	1, 2, 4	1, 2, 4, 8
Max Threads / sock	16	24	30	36	32	96
Max L1 Cache	32KB	32KB	32KB	32KB	32KB	64KB
Max L2 Cache	256 KB	307.2 KB	307.2 KB	256 KB	256 KB	512 KB
Max L3 Cache	20 MB	30 MB	30 MB	30 MB	80 MB	96 MB
Max L4 Cache	0	0	0	0	0	128 MB
Memory Bandwidth	31.4-51.2 GB/s	42.6-59.7 GB/s	42.6-59.7 GB/s	51-68 GB/s	100 - 180 GB/sec	230 - 410 GB/sec

How large the team

1 = The Micro-architecture features above contributed to premium performance for the POWER8 systems compared to Ivy Bridge systems

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


\*32KB running in "Non-RAS mode" 16KB results in better RAS

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

# POWER8 vs Intel Haswell

	Sandy Bridge	Ivy Bridge	Haswell EP E5-26xx v3	POWER 7+ Systems	POWER8	
Clock rates (GHz)			1.6-3.5	3.1-4.4 GHz	3.0-4.1 GHz	
SMT options			1, 2*	1, 2, 4	1, 2, 4, 8	
Max Threads / sock	16	20	36	32	96	
Max L1 Cache	32KB	32KB*	32KB*	64 KB	32KB	64KB
Max L2 Cache	256 KB	256 KB	256 KB	256KB	256 KB	512 KB
Max L3 Cache	20 MB	30 MB	37.5 MB	45 MB	80 MB	96 MB
Max L4 Cache	0	0	0	0	0	128 MB
Memory Bandwidth	31.4-51.2 GB/s	42.6-59.7 GB/s	68-85** GB/s	51-68 GB/s	100 - 180 GB/sec	230 - 410 GB/sec

How abundant the food



1 = The Micro-architecture features above contributed to premium performance for the POWER8 systems compared to Ivy Bridge systems  
 \* Intel calls this Hyper-Threading Technology (No HT and with HT)  
 \*32KB running in "Non-RAS mode" 16KB results in better RAS  
 \*\*85GB running in "Non-RAS mode" and dual-device error NOT supported

# POWER8 vs Intel Haswell

	Sandy Bridge EP E5-x6xx	Ivy Bridge EP E5-26xx v2	Ivy Bridge EX E7-88xx v2	Haswell EP E5-26xx v3	POWER 7+ Systems	POWER8
Clock rates (GHz)	<b>1.8-3.6</b>	<b>1.7-3.7</b>	<b>1.9-3.4</b>	<b>1.6-3.5</b>	<b>3.1-4.4 GHz</b>	<b>3.0-4.1 GHz</b>
SMT options	<b>1,2*</b>		<b>1, 2*</b>	<b>1, 2*</b>	<b>1, 2, 4</b>	<b>1, 2, 4, 8</b>
Max Threads / sock				<b>36</b>	<b>32</b>	<b>96</b>
Max L1 Cache				<b>64 KB</b>	<b>32KB</b>	<b>64KB</b>
Max L2 Cache				<b>256KB</b>	<b>256 KB</b>	<b>512 KB</b>
Max L3 Cache	<b>20</b>			<b>45 MB</b>	<b>80 MB</b>	<b>96 MB</b>
Max L4 Cache	<b>0</b>		<b>0</b>		<b>0</b>	<b>128 MB</b>
Memory Bandwidth	<b>31.4-51.2 GB/s</b>	<b>42.6-59.7 GB/s</b>	<b>68-85** GB/s</b>	<b>51-68 GB/s</b>	<b>100 - 180 GB/sec</b>	<b>230 - 410 GB/sec</b>

How fast the data

51-68 GB/s, 100 - 180 GB/sec, 230 - 410 GB/sec

1 = The Micro-architecture features above contributed to premium performance for the POWER8 systems compared to Ivy Bridge systems  
 \* Intel calls this Hyper-Threading Technology (No HT and with HT)  
 \*32KB running in "Non-RAS mode" 16KB results in better RAS  
 \*\*85GB running in "Non-RAS mode" and dual-device error NOT supported

# In a nutshell



# IBM SWG : 70 PVU on all POWER systems running Linux

- Applies to SWG PVU licensed software products for **ALL** Power cores **running Linux**
- 70 PVU pricing previously only available for 1-2 socket servers
  - PowerLinux 7R1/7R2, p260/p270 and Power 710/720/730/740
  - Intel SandyBridge/IvyBridge servers
- 70 PVU pricing includes:
  - Power IFLs
  - Power S8xx
  - Power E8xx
  - Flex System
- Applies to POWER8 and POWER7


PVU Table per Core (section 1 of 2 - RISC and System z)

Processor Technologies													
		Processor Brand		Processor Type									
Processor Vendor	Processor Name	Server model numbers	Maximum number of sockets per server	Cores per socket						IFL Engine	Proc. Model Number	PVUs per Core	
				(1)	(2)	(4)	(6)	(8)	(12)				
IBM	POWER Systems cores running Linux OS	7R1, 7R2, 7R4 POWER IFL, p24L	All			■	■	■	■		All	70	
		Any POWER System core running Linux											
	POWER7 <sup>4</sup>	770, 780, 795	> 4				■	■	■			All	120
		750, 755, 760, 775, PS704, p460, Power ESE	4					■	■	■		All	100
		PS700-703, 710-740, p260, p270	2						■	■	■	All	70
POWER6	550, 560, 570, 575, 595	All							■		All	120	
	520, JS12, JS22, JS23, JS43	All							■		All	80	

PVU Table link: [http://www-01.ibm.com/software/lotus/passportadvantage/pvu\\_licensing\\_for\\_customers.html](http://www-01.ibm.com/software/lotus/passportadvantage/pvu_licensing_for_customers.html)

# What's in it for me?

Consider a mix of workloads requiring 250 cores on x86 :



- Scale out approach (2-socket) :

70 PVU X 250 = **17.500**




- Scale out approach (2-socket) :

70 PVU X 100 = **7.000**



# What's in it for me?

Consider a mix of workloads requiring 250 cores on x86 :



- Scale out approach (2-socket) :

70 PVU X 250 = **17.500**

- Optimized scale out (4-socket) :

100 PVU X 250 = **25.000**



- Scale out approach (2-socket) :


70 PVU X 100 = **7.000**

- Optimized scale out (4-socket) :

70 PVU X 100 = **7.000**

# What's in it for me?

Consider a mix of workloads requiring 250 cores on x86 :



- Scale out approach (2-socket) :

70 PVU X 250 = **17.500**

- Optimized scale out (4-socket) :

100 PVU X 250 = **25.000**

- Scale up approach (> 4-socket) :

120 PVU X 250 = **30.000**



- Scale out approach (2-socket) :

70 PVU X 100 = **7.000**

- Optimized scale out (4-socket) :

70 PVU X 100 = **7.000**

- Scale up approach (> 4-socket) :

70 PVU X 100 = **7.000**

# POWER8 combined with business accelerators

Altera FPGA acceleration and IBM CAPI

Monte Carlo 250x faster than POWER8 core alone, reduced C code 40x over non-CAPI FPGA



Data Engine for NoSQL 24:1 server consolidation, 3x lower cost per user, 40TB CAPI-attached flash



CAPI dev kit with FPGA card from Nallatech



US Dept of Energy \$325M super computing contract awarded to IBM, Mellanox, and NVIDIA

DoE systems for science and stockpile stewardship

Sierra and Summit systems to be >100 PF, 2 GB/core main memory, local NVRAM, and science performance 4x-8x Titan or Sequoia



NVIDIA acceleration built into IBM Power S824L

8x faster than x86 Ivy Bridge on pattern extraction


82x faster for Cognos BI and DB2 BLU



Tyan OpenPOWER Customer Reference System




# Coherent attachment of accelerator cards



## Typical I/O Model Flow



## Flow with a Coherent Model



Faster application development  
Faster application execution (40 to 1)

Comparable TCA

Linux on Intel  
Ivy Bridge +  
VMware  
Vs.

Linux on POWER8  
+ with PowerVM

**Dell PowerEdge  
R720**

**\$28,366**



**HP ProLiant  
DL380 G8**

**\$29,829**



**IBM Power  
822L**

**\$29,264**



**Server list price\***  
-3-year warranty, on-site

**\$12,605**

**Virtualization**  
- OTC + 3yr. 9x5 SWMA

**\$10,064**

VMware vSphere Enterprise 5.1

**Linux OS list price**  
- RHEL, 2 sockets, unlimited  
guests, 9x5, 3 yr. sub./ supp.

**\$5,697**

Red Hat subscription and Red  
Hat support

**Total list price:  
(Total cost of acquisition)**

**\$28,366**

**\$14,068**

**\$ 10,064**

VMware vSphere Enterprise 5.1

**\$5,697**  
Red Hat subscription and Red  
Hat support

**\$29,829**

**\$14,895**

**\$9,880**

PowerVM for IBM PowerLinux

**\$4,489**  
Red Hat subscription and IBM  
support

**\$29,264**

**Server model**

Dell R720

HP ProLiant DL380p G8

IBM Power 822L

Processor / cores

**Two 2.7 GHz , E5-2697, Ivy Bridge, 12-core processors**

**Two 3.4 GHz POWER8, 10-core**

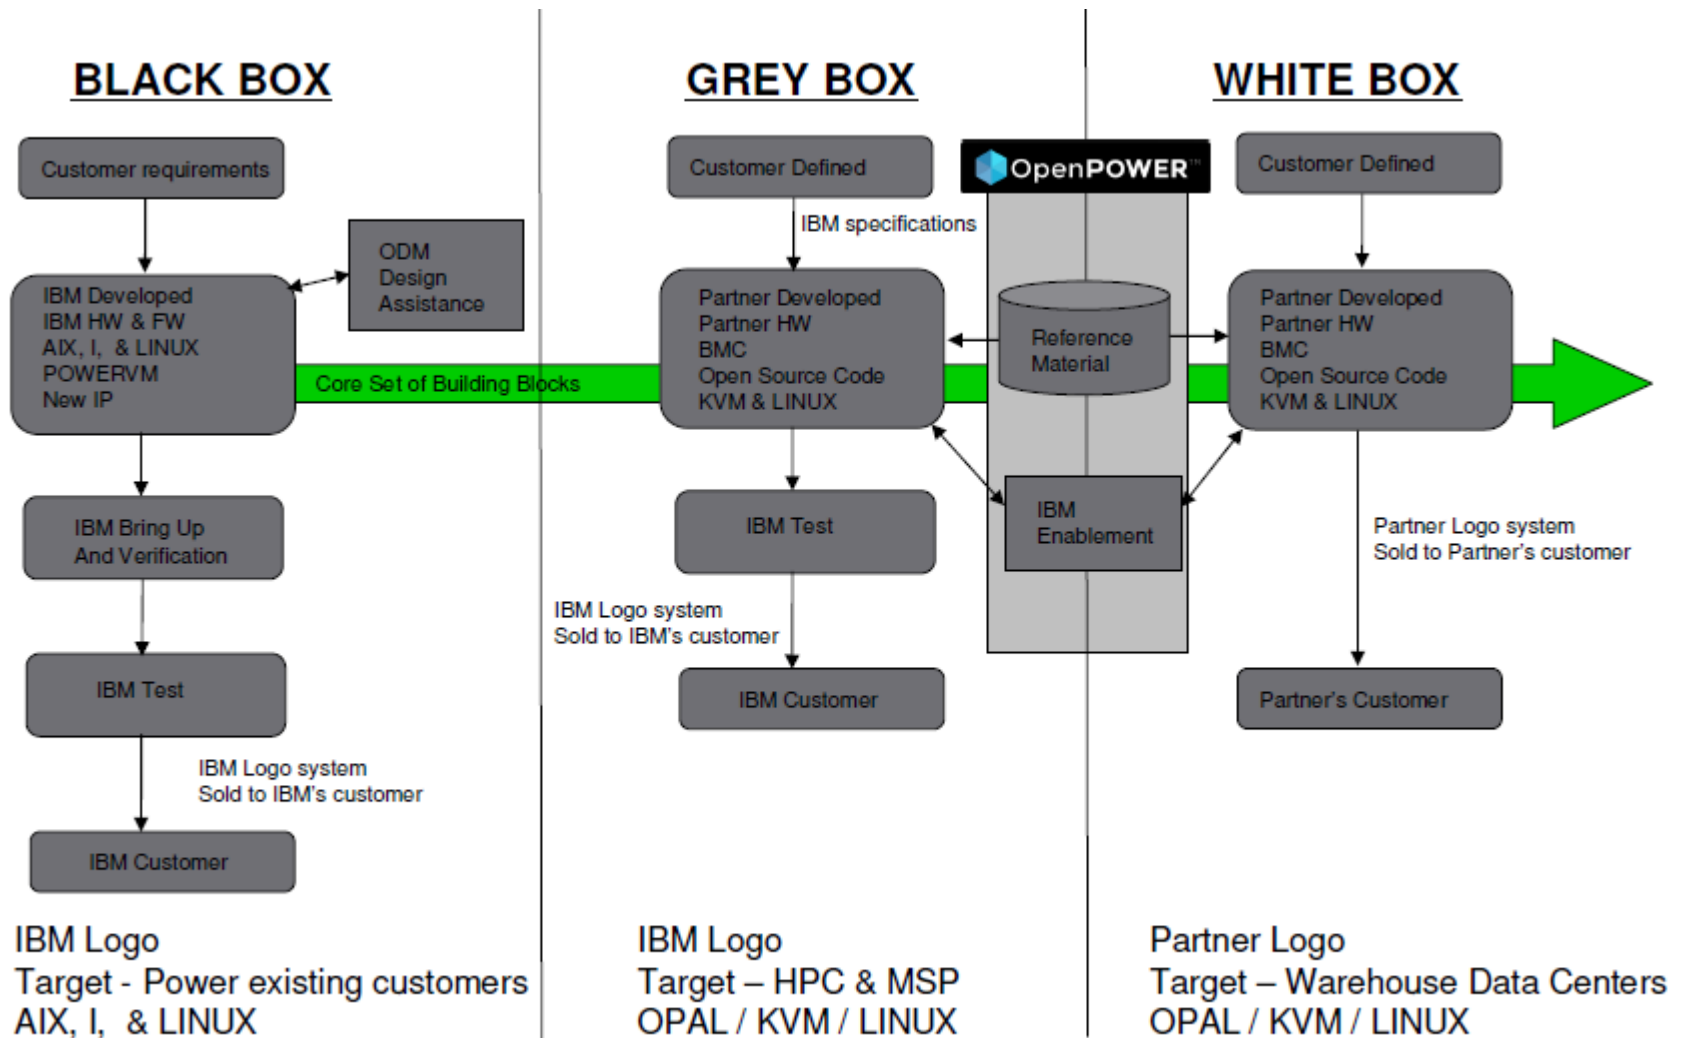
Configuration

64 GB memory, 2 x 300GB 15k HDD, 10 Gb two port

Same memory, HDD, NIC

\* Based on US pricing for Power S822L announcing on April 28, 2014 matching configuration table above. Source: hp.com, dell.com, vmware.com

# New POWER economical model

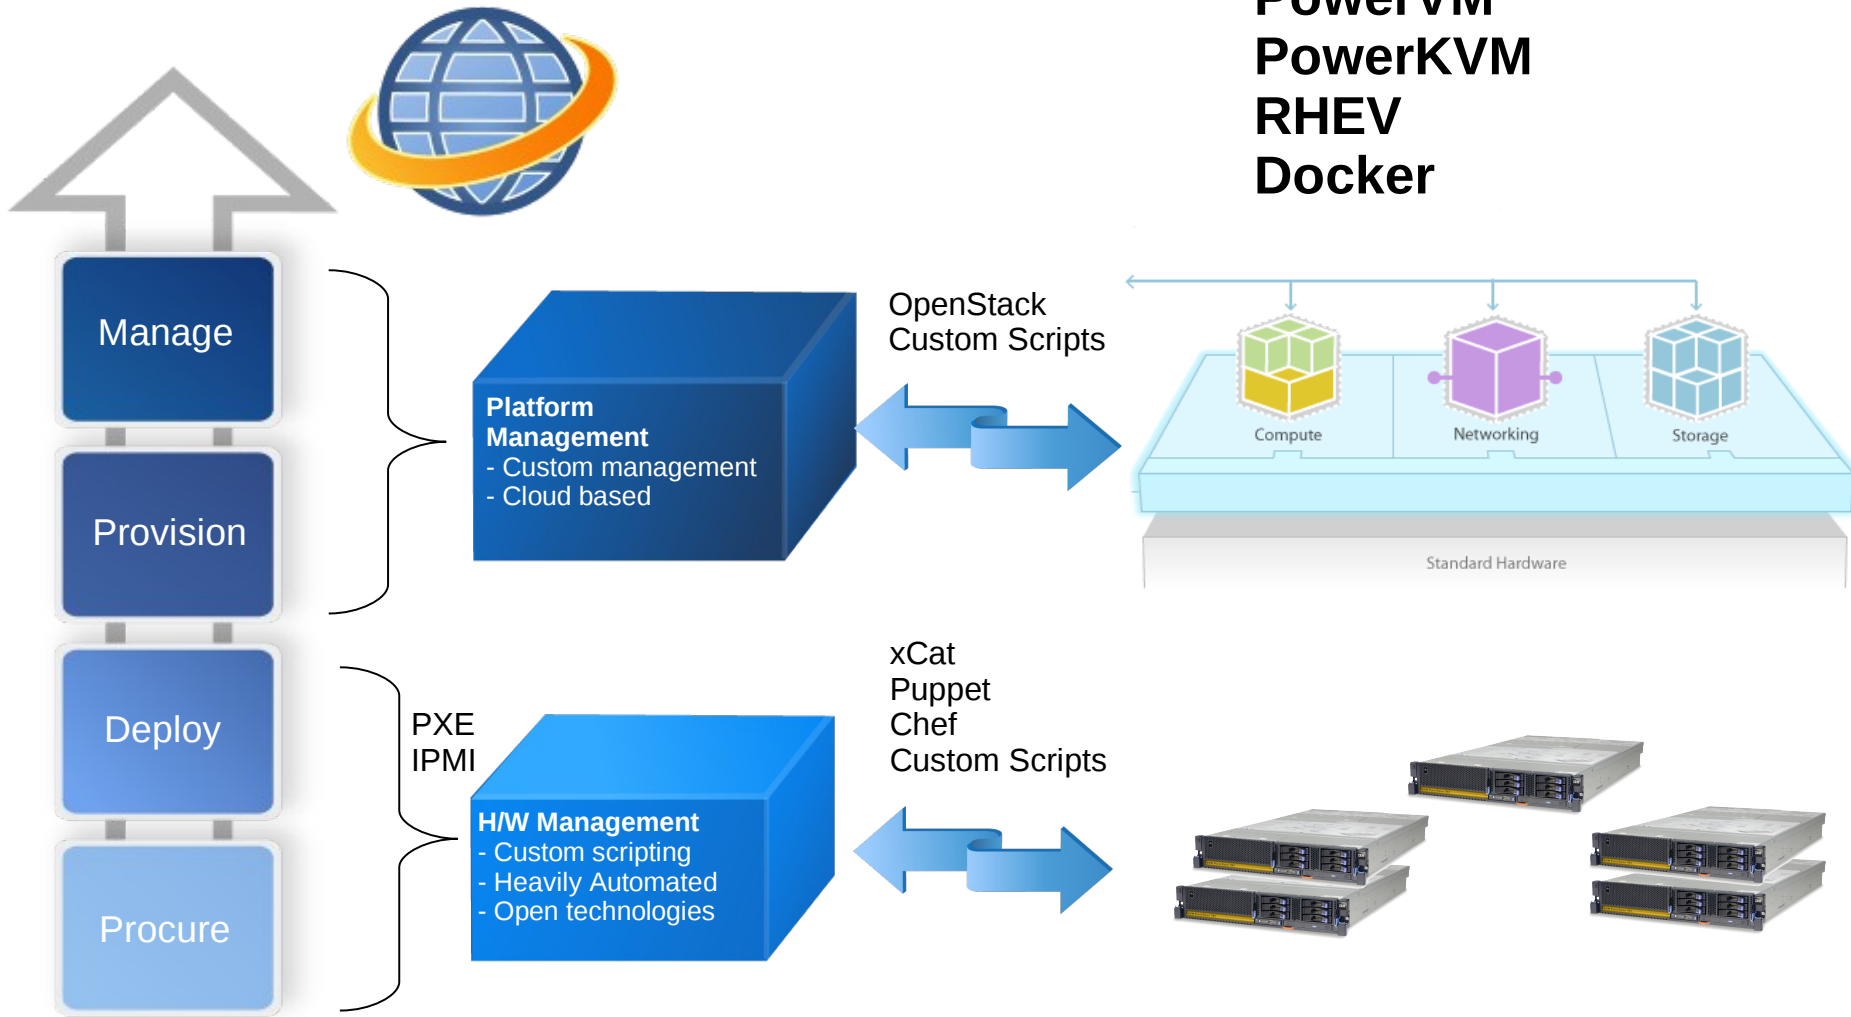




- RHEL 7
  - POWER8 (native mode) and POWER 7/7+
  - Available June 2014
- RHEL 6
  - POWER8 supported with U5 (P7-compatibility mode)
  - Full support of POWER6 and POWER7 (native mode)
- Fedora
  - Fedora 16 was first release to re-launch POWER
  - Fedora 20 has POWER8 support
- Supported add-ons
  - JBoss
  - High Performance Network Add-on
- SLES 12
  - Anticipated to support POWER8 (native mode) and POWER 7/7+
  - Available October 2014
- SLES 11
  - POWER8 with SP3 (P7-compatibility mode)
  - POWER7+ encryption, RNG accelerators with SP3
  - Full support of POWER7 (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
- Ubuntu 14.10
  - Continued support for POWER8
  - Anticipate 4Q14 availability
- 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
  - LTS
- Supported add-ons
  - JuJu Charms
  - MaaS (Metal as a Service)
  - Landscape
- Debian
  - Debian community now supports Power as of

Built from the same source as x86  
Delivered on the same schedule as x86  
Supported at the same time as x86

# Integrating POWER into the datacenter





## Hot off the press – OVH MSP to offer POWER8 platform

“OVH has become the first commercial partner outside of IBM to formally announce the deployment of POWER8 servers inside their data centre. These will be deployed as bare metal servers for customers who have big data and analytics problems and want the most powerful solution that they can get access to. As part of this announcement, OVH will become an OpenPOWER Foundation member.”



- OVH [announcement](#)
- Google [announcement](#)
- Softlayer [announcement](#)
- Rackspace [announcement](#)

# 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



**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.*

## OpenPOWER drives industry innovation

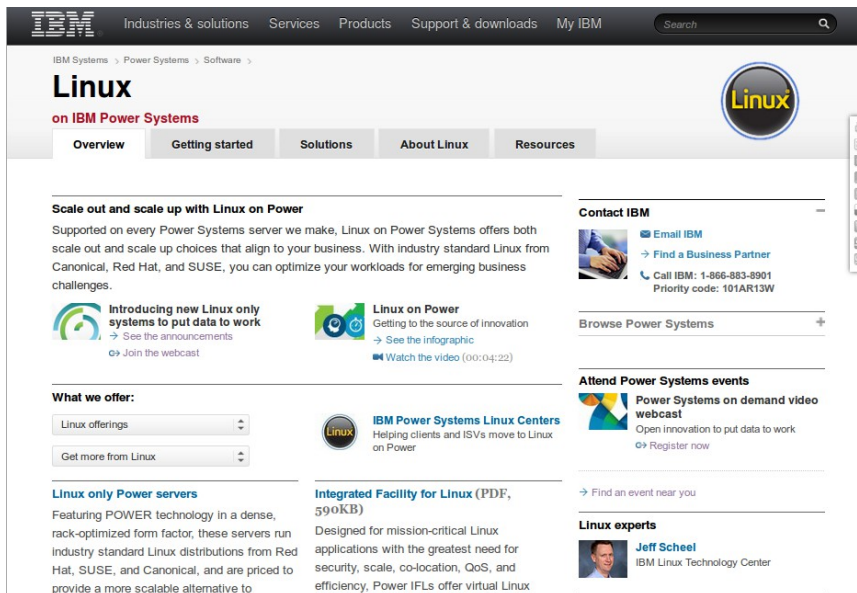
The goal of the OpenPOWER Foundation is to create an open ecosystem, using the POWER Architecture to share expertise, investment, and server-class intellectual property to serve the evolving needs of customers.

<http://openpowerfoundation.org/>

Where to find more information?

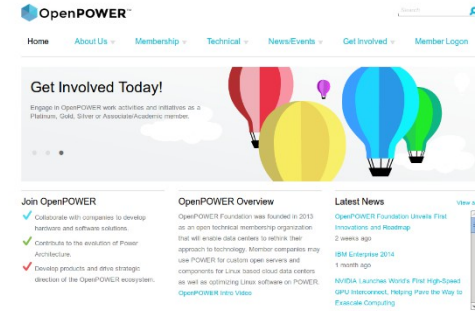
# Power Systems Linux Portal

(Product Information)



The screenshot shows the IBM Power Systems Linux Portal. The navigation bar includes 'Industries & solutions', 'Services', 'Products', 'Support & downloads', and 'My IBM'. The main content area is titled 'Linux on IBM Power Systems' and features tabs for 'Overview', 'Getting started', 'Solutions', 'About Linux', and 'Resources'. Key sections include:
 

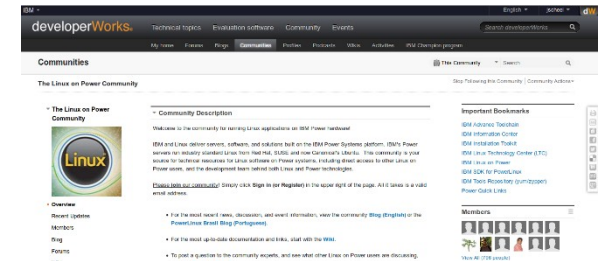
- Scale out and scale up with Linux on Power:** Supported on every Power Systems server, offering standard Linux distributions from Canonical, Red Hat, and SUSE.
- Introducing new Linux only systems to put data to work:** Includes announcements and a webcast link.
- Linux on Power:** Getting to the source of innovation with an infographic and a 4:22 video.
- What we offer:** Linux offerings and 'Get more from Linux' options.
- Linux only Power servers:** Featuring POWER technology in a dense, rack-optimized form factor.
- Integrated Facility for Linux (PDF, 590KB):** Designed for mission-critical Linux applications.
- IBM Power Systems Linux Centers:** Helping clients and ISVs move to Linux on Power.
- Contact IBM:** Email, business partner finder, and call center (1-866-883-8901).
- Attend Power Systems events:** Power Systems on demand video webcast.
- Linux experts:** Jeff Scheel, IBM Linux Technology Center.



The screenshot shows the OpenPOWER website. The navigation bar includes 'Home', 'About Us', 'Membership', 'Technical', 'News/Events', 'Get Involved', and 'Member Login'. The main content area is titled 'Get Involved Today!' and features:
 

- Join OpenPOWER:** Collaborate with companies to develop hardware and software solutions.
- OpenPOWER Overview:** OpenPOWER Foundation was founded in 2013 as an open technical membership organization.
- Latest News:** OpenPOWER Foundation Unveils First Innovations and Meeting (2 weeks ago), IBM Celebrates 2014 (1 month ago), and NVIDIA Launches World's First High-Speed GPU Interconnect.

## The OpenPOWER Foundation



The screenshot shows the developerWorks Linux on Power Community page. The navigation bar includes 'Technical topics', 'Evaluation software', 'Community', and 'Events'. The main content area is titled 'The Linux on Power Community' and features:
 

- Community Description:** Welcome to the community for sharing Linux applications on IBM Power hardware.
- Important Bookmarks:** IBM Address Book, IBM Information Center, IBM Installation Toolkit, IBM Linux Technology Center (ITC), IBM Linux on Power, IBM SDK for PowerLinux, and IBM Tech Resource Center/IBM Power Quick Links.
- Members:** A list of community members.

## The Linux on Power Community



franz\_bourlet@be.ibm.com