

POWER up your Linux!

May 2014

Mandie Quartly

Linux on Power technical lead UK&I

mandie_quartly@uk.ibm.com

Twitter: [@mandieq](https://twitter.com/mandieq)

020 8818 5283

Today's session: POWER up your Linux!

Will cover:

- Overview of using Linux on Power systems for techies*
- Why might this be of interest?
- What do I need to consider?
- Where do I start?

Follow on sessions

- More in depth sessions focused at those with a specific background
- Focus on certain areas of interest



Taken from tuxgallery.weebly.com

* techies = with Power and x86 background

Agenda

By the end of the hour:

- Why would I run Linux on Power?
- Linux = ?
- Power = ?
- How can I run it?
- How to install it?
- What might I use it for?
- Other things you should know
- Where to go for more information?

First, a couple of questions

- Who are you?
- Which describes you best?
- What's your favourite Linux OS for business use?



Why run Linux on Power?



Why Linux on Power?

“I know Power...”

All the benefits of running of on Power systems plus...

- Price equivalent to AIX or lower
- Access to emerging workloads available only on Linux
- Easy use of Open Source solutions e.g. Apache, Ganglia...

“I know Linux...”

All the benefits of running of Linux plus...

- Better scale out performance at lower cost than x86
- Ability to scale up
- Better reliability and resiliency than x86
- Linux is Linux

IBM Innovation & Focus

IBM Power Systems built on open innovation

Embraces open stack of software for easy development and implementation

Creates a pipeline of ecosystem innovation



Leverage industry innovation

Accelerates applications on a purpose-built infrastructure

POWER8 

IBM has a history of industry-leading collaboration on open technology



+



*15+ Years of
Collaboration
on Open Source*



1999

2000

2001

2007

2007

2011

2011

2012

2013

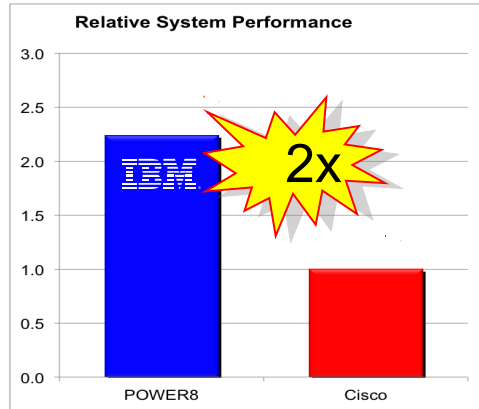
2013

IBM is heavily invested in the success of Linux

- **100%** of our systems
- **400+** software products
- **500** patents donated
- **600+** developers



Big Data / Hadoop



- POWER8 S824 will deliver **2x the performance** of the best published x86 system for Big Data / Hadoop
- POWER8 exploits additional cores, more threads, larger caches, memory bandwidth, I/O bandwidth
- Terasort is a popular benchmark to measure the performance of a Hadoop solution

IBM Analytics Stack: IBM Power System S824; 24 cores / 192 threads, POWER8; 3.5GHz, 512 GB memory, RHEL 6.5, InfoSphere BigInsights 2.1.2
http://www.cisco.com/en/US/solutions/collateral/ns340/ns517/ns224/ns944/le_tera.pdf

- Derive insights **67% faster** with Cognos Business Intelligence (Cognos v10.2) running on Linux on POWER8 versus Linux on Intel Ivy Bridge-EP
- Utilize Predictive Analytics to get the answers before your competitors

Cognos®



67% faster is based on IBM Internal Testing of sample analytic workloads; current as of March 28, 2014. Performance improvement figures are cumulative of all queries in the workload. Individual results will vary depending on individual workloads, configurations and conditions.

IBM Analytics Stack: IBM Power System S824; 24 cores / 192 threads, POWER8; 3.5GHz, 384 GB memory, DB2 10.1 and Cognos v 10.2

Competitive stack: HP DL380p; 24 cores / 48 threads; Intel E5-2697 v2; 2.7 GHz; 384 GB; DB2 10.1 and Cognos v 10.2

Lower cloud infrastructure costs

Thirty-four 2-socket Power S822L servers do the job of 100 2-socket x86 (HP DL380) servers running equal virtualised capacity

58% lower total HW TCO

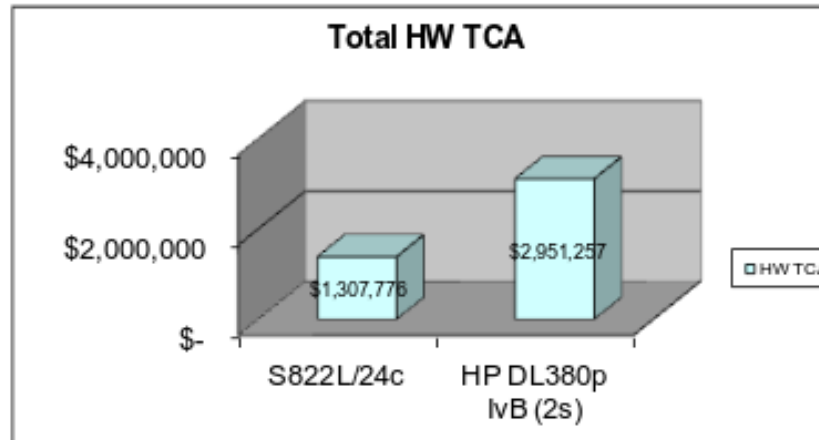
vs.

Ivy Bridge w/ VMware

66% Less Systems & Cores

Lower SW License Fees
Reduced Management Costs
Reduced Floor Space

- Power S822L**
- TCA/TCO is for 34 servers, 816 cores
 - 3x better virtualized throughput vs. an HP 2 socket Ivy Bridge
 - 2S, 24 cores each
 - POWER8, 3.0GHz
 - PowerVM



- HP DL 380 G8**
- TCA/TCO is for 100 servers, 2400 cores
 - 100 HP servers needed for ~ equal virtualized throughput of 34 Power S822L
 - 2S, 24 cores each
 - Ivy Bridge, 2.7GHz
 - VMware vSphere Ent

Source: Capacity based on IBM Sizing of typical SPECint_rate landscape and 3rd party analysis of system utilization. Pricing from www.hp.com.

This is an IBM sizing designed to replicate a typical IBM customer workload used in the marketplace. The results are calculated and not an actual customer environment. IBM's internal workload studies are not benchmark applications, nor are they based on any benchmark standard. As such, customer applications, differences in the stack deployed, and other systems variations or conditions may produce different results and may vary based on actual configuration, applications, specific queries and other variables in a production environment. Prices, where applicable, are based on published US list prices for both IBM and competitor.

Power systems have better RAS features

RAS Feature	POWER7	x86
Application/Partition		
Move running workload partitions to another server	Yes	Yes
Partition Availability priority	Yes	No
System		
First Failure Data Capture	Yes	No
Memory Keys ensure secure hardware separation of memory	Yes	No
Processor		
Dedicated Service Element processor	Yes	No
Processor Instruction Retry	Yes	No
Alternate Processor Recovery	Yes	No
Predictive Processor Deallocation	Yes	No (1)
Dynamic Processor Sparing	Yes	No (1)
Memory		
Chipkill™ - Redundant Memory Capabilities	Yes	Yes
Chipkill™ - Survives Double Memory Failures with x8 DRAMS	Yes (2)	No
I/O		
Extended Error Handling – PCI bus, PCI bridge, GX+ bus adapter	Yes	No
VIOS - Dynamic adapter failover	Yes	No

1 = Capability built in but typically supported
 2 = Chipkill™ On Systems with DRAM sparing

Which Linux?



Red Hat support on POWER



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

- Red Hat Enterprise Linux 7
 - Public beta available for existing RHEL customers
 - Expected full support for POWER8 (native mode) and POWER 7/7+ at operating system GA
- Red Hat Enterprise Linux 6
 - POWER8 supported with U5 (P7-compatibility mode)
 - Full support of POWER6 and POWER7 (native mode)
 - Last update: U5 GA December 2013
- Fedora
 - Fedora 16 was first release to re-launch POWER
 - Fedora 20 has POWER8 support
 - Fedora remains primary community for major innovation/collaboration
- Supported add-ons
 - JBoss
 - High Performance Network Add-on

SUSE support on POWER



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

- **SUSE Linux Enterprise Server 11**
 - POWER8 supported with SP3 (P7-compatibility mode)
 - POWER7+ encryption, RNG accelerators supported with SP3
 - Full support of POWER7 (native mode)
 - Earliest supported release: SLES 11 base
 - Last update: SP3 GA July 2013
- **SUSE Linux Enterprise Server 10***
 - POWER7 supported with SP3 (P6-compatibility mode)
 - Full support of POWER6 (native mode)
 - Last update: SP4 GA April 2011
- **openSUSE**
 - openSUSE 12.2 re-launched for IBM POWER
 - openSUSE 13.2 includes POWER8 support (native mode)
- **Supported add-ons**
 - SUSE Linux Enterprise High Availability Extension (included in base Power license)

* Not supported on POWER7+ and POWER8 systems


Canonical support on POWER



- **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
- **Supported add-ons**
 - JuJu Charms
 - MaaS (Metal as a Service)
 - Landscape
- **Debian**
 - Community enablement in process

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

Power has Linux “release parity” with x86



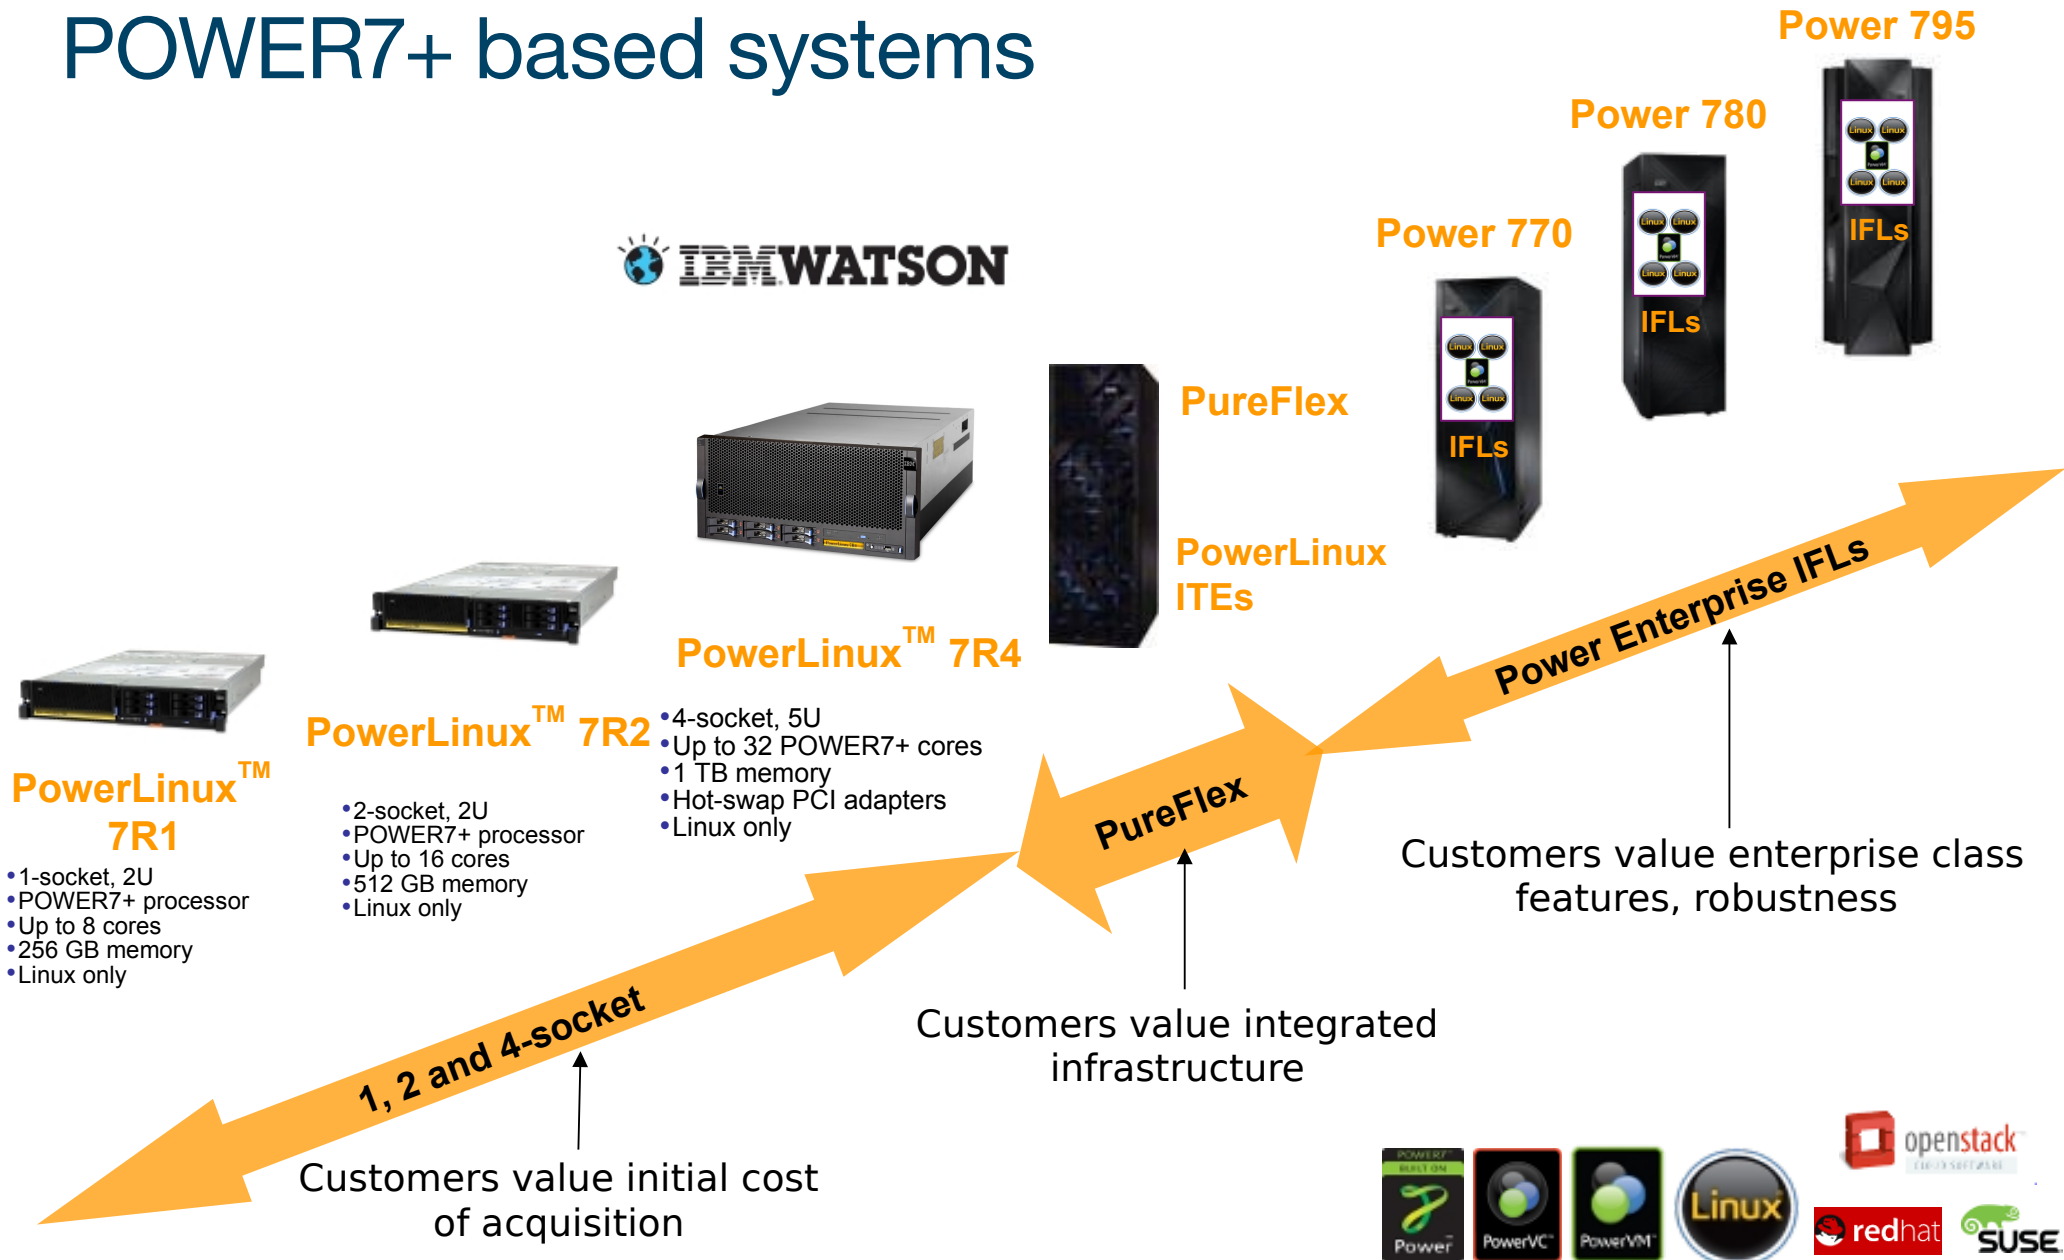
See for more details:

- Red Hat lifecycle information - <https://access.redhat.com/support/policy/updates/errata/>
- SUSE lifecycle information - <http://support.novell.com/lifecycle/>
- Ubuntu lifecycle information - <https://wiki.ubuntu.com/Releases>

Which Power boxes?



POWER7+ based systems



New Power Systems scale-out portfolio



Power **SC**

Power **VC**

Power **HA**

Power **VP**

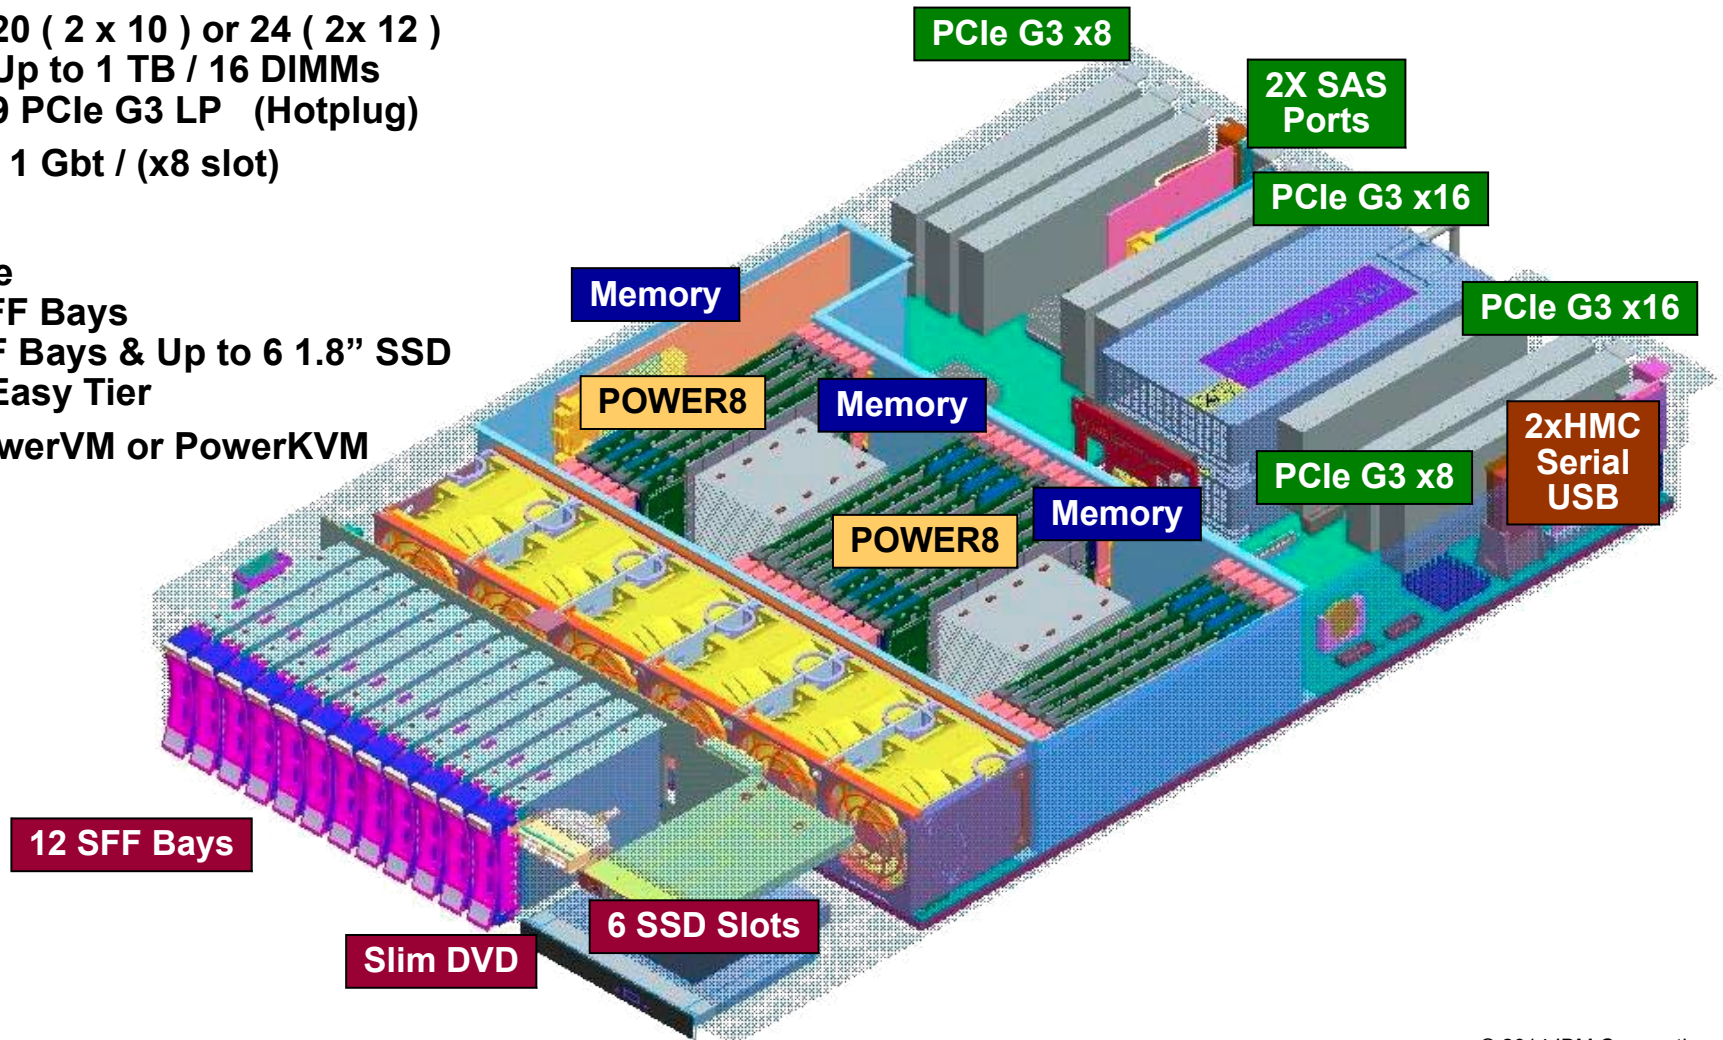
Power **KVM**

ubuntu[®]
Supported by Canonical



POWER8 Linux 2S 2U Scale Out System

- Form Factor: 2U
- Sockets: 2S
- Dual Sockets
 - Cores: 20 (2 x 10) or 24 (2x 12)
 - Memory: Up to 1 TB / 16 DIMMs
 - Slots: 9 PCIe G3 LP (Hotplug)
- Ethernet: Quad 1 Gbt / (x8 slot)
- Media: DVD
- Internal Storage
 - Up to 12 SFF Bays
 - Up to 8 SFF Bays & Up to 6 1.8" SSD Bays with Easy Tier
- Hypervisor: PowerVM or PowerKVM



**OK, I want to have a go –
where do I start?**



Things you will need...

- A POWER based system to run on
- An OS to run
 - Might already have a site favourite / strategy?
 - If not:
 - Get a 60 day trial version of [RHEL](#) or [SLES](#), or try [Ubuntu](#) (POWER8 only)
 - Download one of the community supported versions at no cost e.g. [Fedora](#) / [OpenSUSE](#) / [Debian](#) (links to ppc download sites)
 - Check the flavour / version you choose works with
 - the Power system you are using (look at the release notes)
 - the application stack you are planning to use
- To decide how to set up your system (quite a few options)
 - To virtualise or not to virtualise?
 - HMC availability?
 - Volume of virtual machines / partitions required: a one off or 100s

Installation / hosting options

Bare metal / native install

If no graphics adapter installed, then use RS232 interface (see gotcha)
Doesn't take advantage of flexibility and high utilisation of virtualisation

Using Integrated Virtualisation Management (IVM)

No need for HMC
Access mgmt partition via browser to work with partitions + virtual infrastructure

Using an HMC and PowerVM

Needs an HMC
Business as usual for larger Power servers.

Using PowerKVM

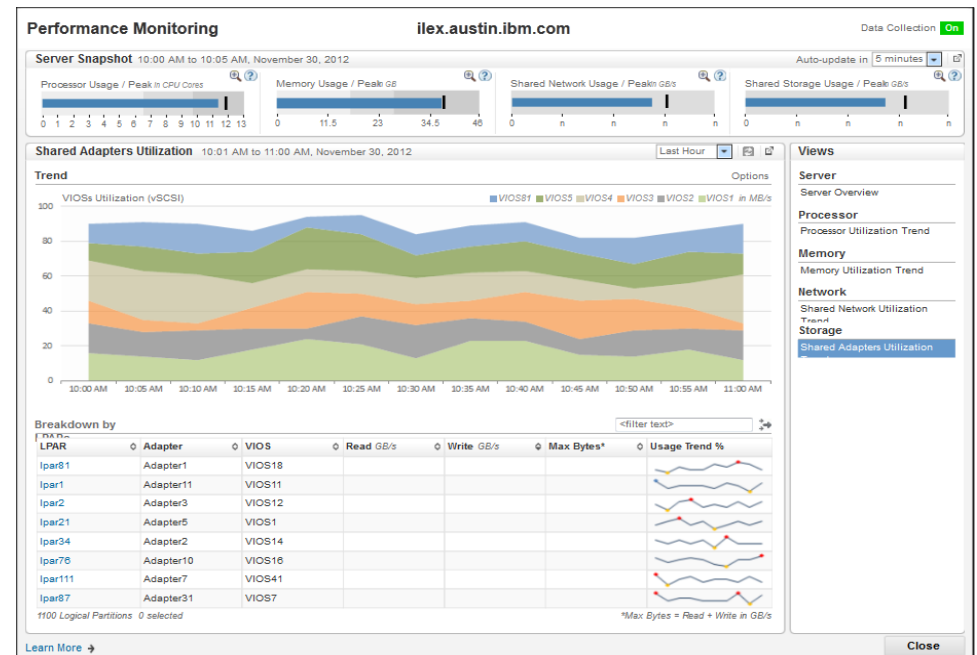
No need for HMC
Only on POWER8 Linux only servers

With virtualisation

Larger installations

PowerVM: virtualisation without limits

- New monitoring views help simplify performance and capacity management
- Support for latest POWER8 servers and capabilities
- SRIOV NIC support for enhanced network virtualisation

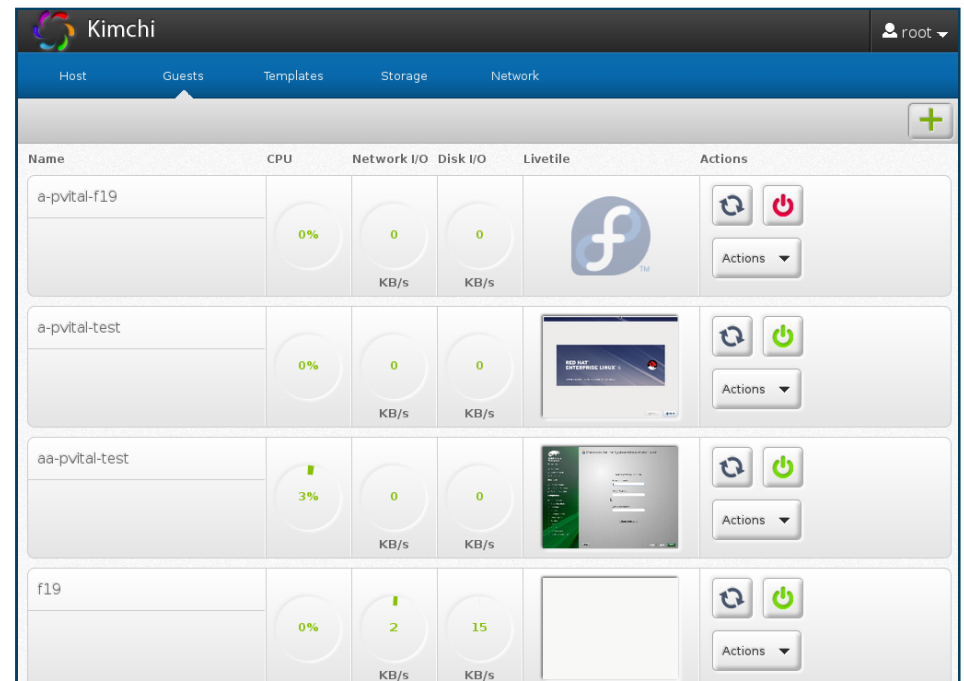


PowerVM



PowerKVM: open virtualisation for scale-out servers

- Kernel-Based Virtual Machine(KVM) virtualisation targets new Linux workloads
- Provides simplicity and familiarity for VMware and KVM Intel Linux admins
- Enables cloud providers to integrate Linux on Power into OpenStack environments
- Managed by PowerVC or open source tools such as Kimchi
- Exploits POWER8 micro-threading for performance efficiency



PowerKVM



Installation / hosting options - revisit

Bare metal / native install

If no graphics adapter installed, then use RS232 interface (see gotcha)
Doesn't take advantage of flexibility and high utilisation of virtualisation

Using Integrated Virtualisation Management (IVM)

No need for HMC
Access mgmt partition via browser to work with partitions + virtual infrastructure

Using an HMC and PowerVM

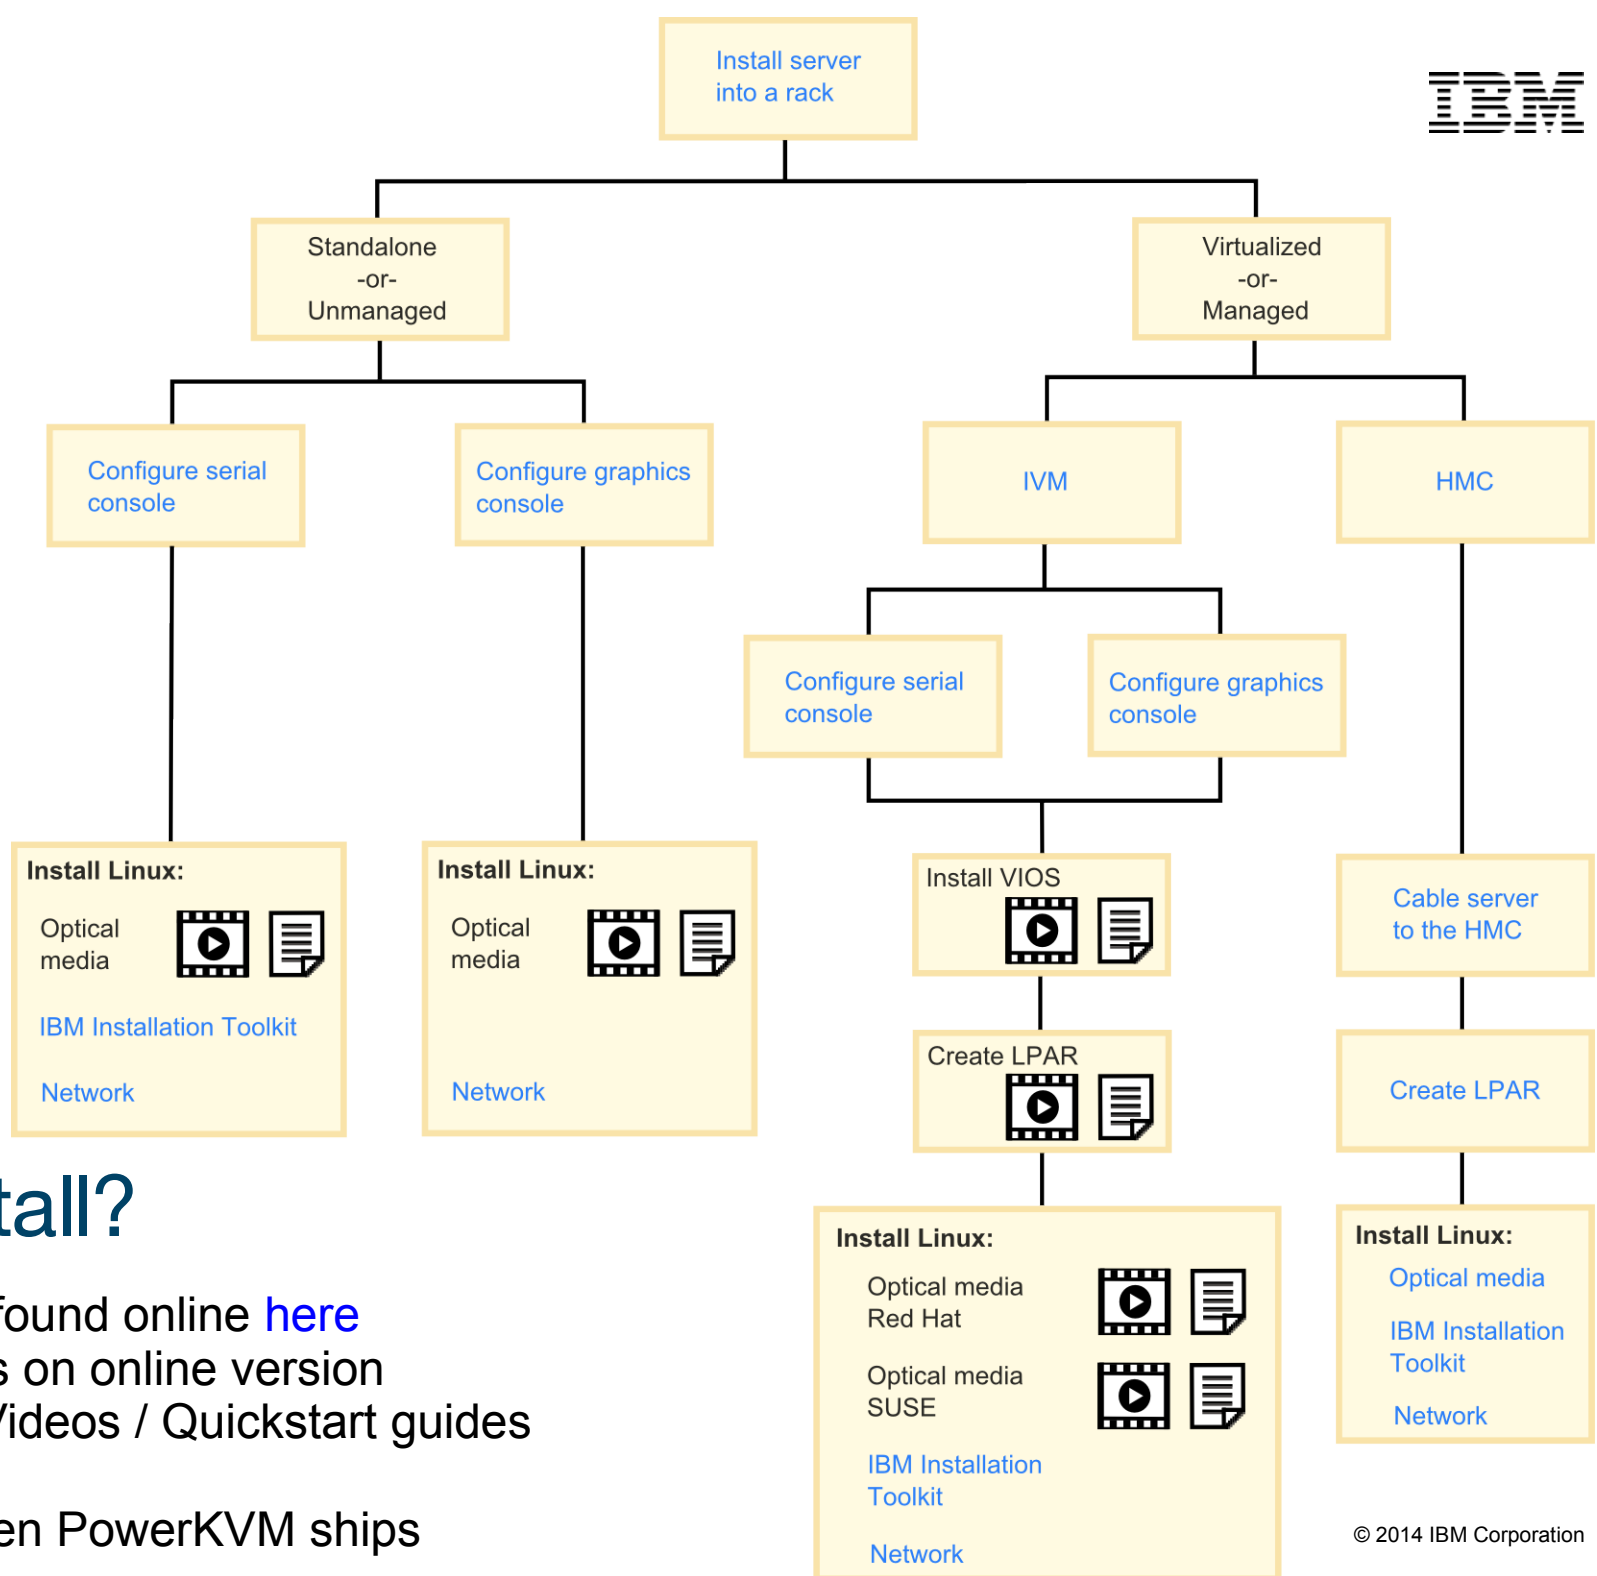
Needs an HMC
Business as usual for larger Power servers.

Using PowerKVM

No need for HMC
Only on POWER8 Linux only servers

With virtualisation

Larger installations



How to install?

Flowchart can be found online [here](#)
 Blue writing = links on online version
 Other symbols = Videos / Quickstart guides

To be updated when PowerKVM ships

Installation of a Linux operating system

Either natively or into a partition / virtual machine

- DVD
- Network
- IBM Installation toolkit
 - “hides” distribution front end, automated install, includes IBM service and productivity tools
 - installation natively or on a pre-configured LPAR
 - minimal options, optionally select server type e.g. LAMP
- To make life easier:
 - Use VNC installation mode for headless nodes
 - **boot: linux vnc**
 - Needs network connection and assumes DHCP available, otherwise specify network settings also
 - Automated installation using config file available on network
 - Kickstart for Redhat
 - AutoYAST for SUSE

Linux is Linux

(but there are some differences related to the architecture)

- Partitioning:
 - additional PreP partition which holds bootloader
 - MBR is present, but only holds partition table
- System firmware
 - System Management Services (SMS)
 - Press 1 at boot time to enter, x to exit
 - Config boot order, bootp network boot, ...
- Bootloader: yaboot (compared to grub / lilo)
- No PXE boot, network boot uses BOOTP protocol (grub2 future direction)
- Package names: ppc / ppc64
- Additional value-add packages: e.g. lprutils, Powerpc-utils, ppc64-utils, servicelog

```

-----
Navigation keys:
M = return to Main Menu
ESC key = return to previous screen          X = eXit System Management Services
-----
Type menu item number and press Enter or select Navigation key:M

PowerPC Firmware
Version AL740_095
SMS 1.7 (c) Copyright IBM Corp. 2000,2008 All rights reserved.
-----
Main Menu
1.  Select Language
2.  Setup Remote IPL (Initial Program Load)
3.  Change SCSI Settings
4.  Select Console
5.  Select Boot Options
    
```

```

-----
Navigation Keys:
-----
X = eXit System Management Services
-----
Type menu item number and press Enter or select Navigation key:
    
```

Potential gotcha

- Beware of the RS232 serial interface – dig out those old cables!
 - For native / IVM installations with no graphics adapter, first part of install requires input of information via serial port
 - Need adapter from Power machine (port marked S1, looks like ethernet / RJ45 socket) to RS232 serial port
 - Need laptop with a serial port (!) or a RS232 to USB adapter
 - A female to female adapter to attach your adapters together
 - A terminal emulator on your laptop e.g. Teraterm (Windows) or minicom (Linux)
 - See Nigel's AIXpert [related blog](#) for more details

Machine State	Port function
Powered off	Service processor text interface
Powered on	System console for 1 st partition
HMC Managed	Disabled

Serial port attribute	Default setting
Line speed	19200
Word Length (bits)	8
Parity	No
Stop Bits	1

**OK, I've got it installed, now
what do I do with it?**



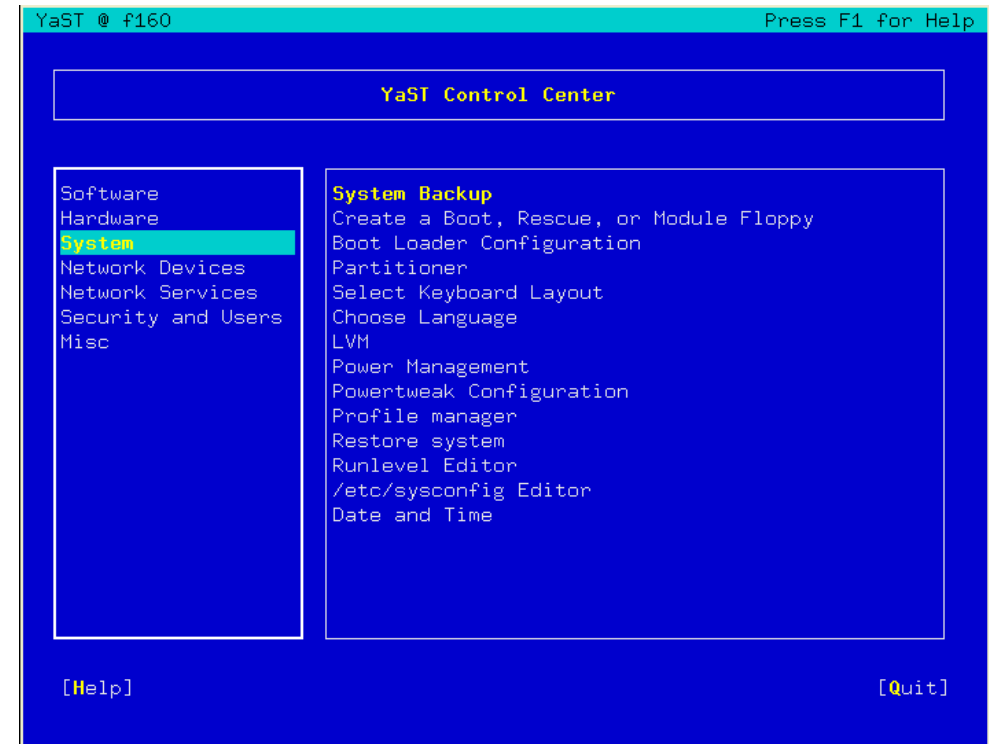
Starting off

System configuration:

- `yast2` = smitty on SLES / OpenSUSE
 - GUI / TUI
- equivalent for RedHat / Fedora found via `system-config-*` commands
 - e.g. `system-config-network`, `system-config-services` ...

Service and productivity tools ([link](#)):

- Hardware diagnostic aids and productivity tools (yum repository)
- IBM Software Development Kit
 - a free, Eclipse-based integrated development environment
 - Advance Toolchain
 - Classic Linux performance analysis tools, including OProfile and Valgrind
- `nmon` (of course!)



This is starting to look pretty good! What else should I know?

What might you use it for?

IBM Software Group apps

156 applications supported and increasing!

Big Data, Data services, Business application middleware, Infrastructure services, Development and test, Mobile, Social, Enterprise Content Management, High Availability, Security

For more details see: IBM Software Product Compatibility Reports

[link](#)

Independent Software Vendor (ISV) apps

1000s of applications available for Linux on Power systems



[IBM Global Solutions Directory](#)

Open source “killer apps”

(aka things that are much easier to do / possible on Linux compared to AIX)



Useful snippets

70 PVU rating for all Power cores running Linux



PowerKVM just announced

PowerKVM

You don't need to buy a server – switch on some dark cores!

- Power IFL


- 4 x CUoD core activations
- 32 GB CUoD memory activations
- 4 x PowerVM for PowerLinux license entitlement
- Scales in increments of 4 cores




Available on Power 770, 780 & 795



Recent Redbooks



[Link](#)

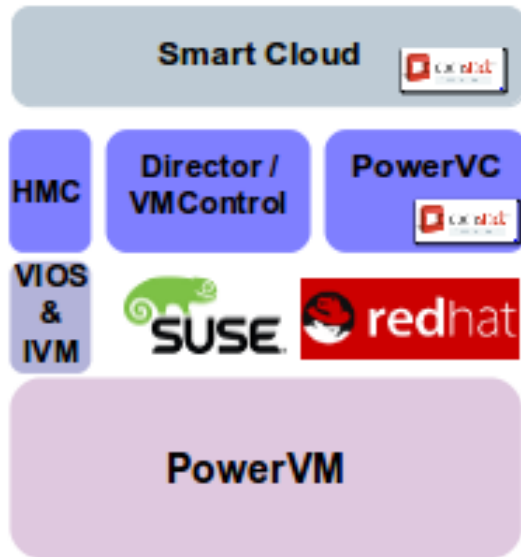


[Link](#)

[Link to search for "Linux" on the Power Systems Redbooks site](#)

More on PowerKVM

PowerVM Stack



Cloud Software
 Sys Mgmt Software
 Operating System
 Hypervisor / Firmware

New PowerKVM Stack

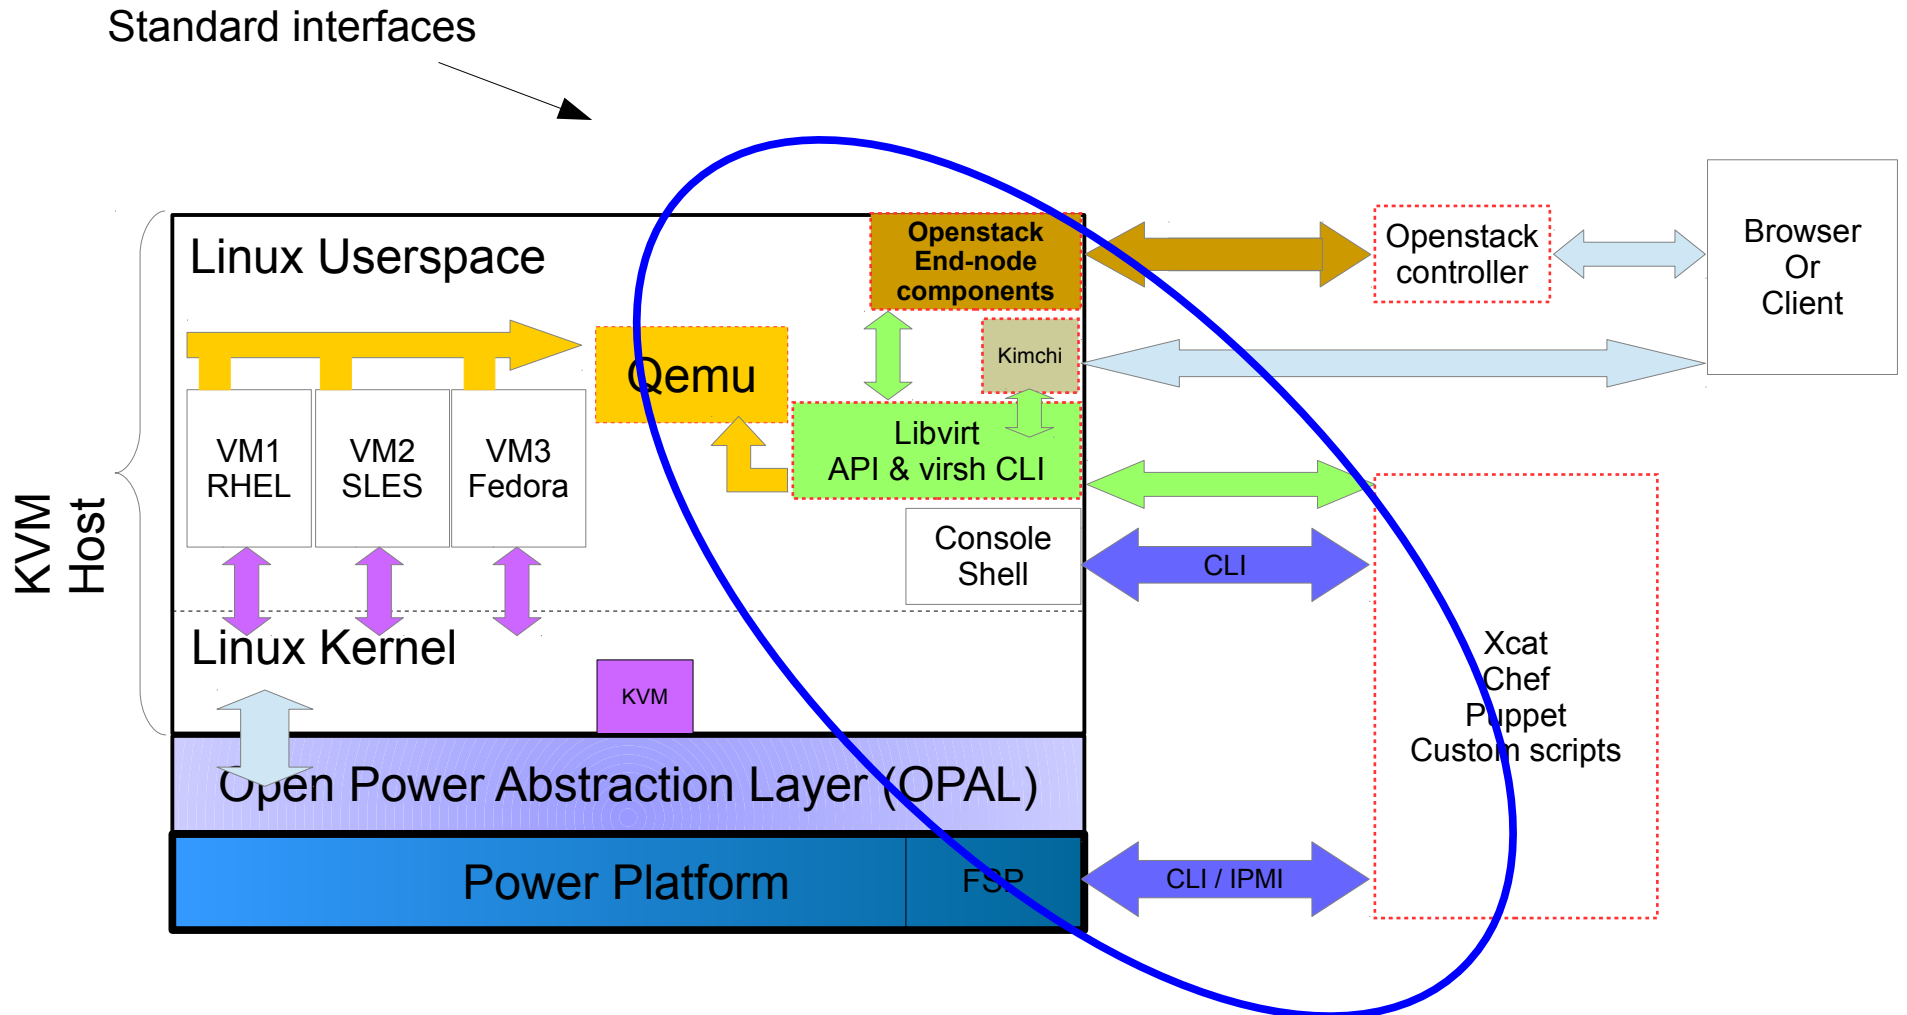


Key points:

- Only on scale out 1-2 socket Linux only servers
- Not nested in LPAR, sits on hardware
- No HMC / IVM / VIOS
- Embraces opensource clouds and other virtualisation SW through standard interfaces like oVirt (VDSM) and OpenStack

KVM on Power Management Architecture

Enabling integration into an x86 world



Note: oVirt has been enabled in version 3.4 for Linux on Power.

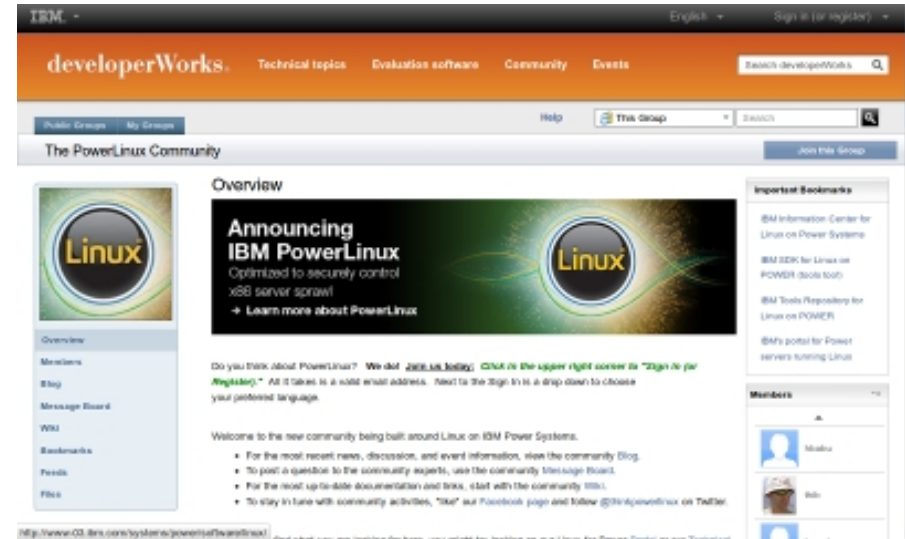
Where to go for more information?



Learn more about Linux on Power Systems

Power Systems Linux web site

www.ibm.com/systems/power/software/linux/



The PowerLinux Community
www.ibm.com/developerworks/group/tpl/

AIXpert blog tinyurl.com/AIXpert



[@ibmpowerlinux](https://twitter.com/ibmpowerlinux)



plus.google.com/communities/100156952249293416679

In summary....

- There are many benefits to running Linux on Power whether you are coming from a Power background or an x86 background.
 - For scale out workloads the default answer should not be Linux on x86. Linux on Power has comparable TCA, better TCO, better performance and reliability.
 - In an existing Power environment Linux can be better suited for emerging / open source workloads.
- Have a go!
 - Download a version of Linux and try it out on your existing hardware.
 - IBMers / partners – try it out on the [IBM Power Development Cloud](#)
 - Customers - don't have hardware but have a focus workload? Get us involved in a proof of concept
 - Unsure about your Linux strategy? Get us involved to look more closely at your existing workloads and their suitability for moving.

Thanks for your time!

Mandie Quartly

Linux on Power technical lead UK&I

mandie_quartly@uk.ibm.com

Twitter: [@mandieq](https://twitter.com/mandieq)

020 8818 5283