

Today

## Linux on POWER Field Experience

Starting at 10:00 am UK time by Tom Watts



Smart Meeting → Put questions into the Chat box or AT&T Toll Free phone for better audio

- 0800-368-0638 = UK Toll Free
- 0203-059-6451 = UK but you pay for the call
- Then 6403785# Participant Code
- Other countries see chat box for the website
- Please Mute with \*6



Previous Sessions:  
POWER8 E850 from exper  
PowerSC Tools for IBM i  
HMC 8.20 Tech Preview  
The "Key" to IBM i Licensing  
POWER8 E870 from exper  
Linux on Power: Best Pract  
Linux for AIX/IBM i guys  
PowerKVM Deep Dive  
More Tricks Power Masters  
Power8 from hands-on  
Power up your Linux  
PowerVC  
And more.....

Future Sessions →

- TBC – 7<sup>th</sup> Oct - Remote Restart
- More Being planned
- Suggestions Welcome



Webinar wiki: <http://tinyurl.com/PowerSystemsTechnicalWebinars>  
Youtube Channel: <http://tinyurl.com/IBMPowerVUGYoutubeChannel>

Twitter:   
Gareth Coates @power\_gaz  
Jyoti Dodhia @JyotiDodhia

Nigel Griffiths @mr\_nmon  
Mandie Quartly @mandieq

## Linux on POWER Field Experience



Tom Watts

Senior Systems Consultant, IBM Lab Services

watts1t@uk.ibm.com

# Design Options

## Mixed Workloads

- Utilise the POWER processors as much as possible, this means multiple partitions / virtual machines and not bare metal.
- In PowerVM set up sensible processor and memory values
- In PowerKVM don't try to over commit processor resources as this will slow down the virtual machines
- Use platform management tools where possible – PowerVC for example – to capture and deploy multiple images quickly.

## High Performance

- Generally for High Performance situations (HPC, time critical trading etc.) don't use virtualisation just use a bare metal option (distribution limited at this time).
- Use PowerKVM but only define one or two virtual machines and in the configuration pin each virtual machine to a set number of processor cores and definitely don't over commit.
- If using highly multi-threaded applications ensure SMT is on and set to 8 in the virtual machines.

# Installation in the Real World™ Part 1...

## HMC Connected Systems

- Pre-Requisites:
  - Connect the service processor to the management network that also talks to the HMC – you would be surprised!
  - Once the HMC has detected the new server connect to the server and follow the prompts for passwords and eventually you will have a server you can do stuff with.
  - **ALWAYS** check and install latest firmware **BEFORE** anything else!!!
  - Create a partition and start it up to SMS and open a console session to make sure all of the basics are working.
  - Download the latest version of the desired Linux distribution – current supported versions are on the next slide (sorry for lots of text).

## Standalone / Compute Node

- Pre-Requisites:
  - Connect to the first serial interface using a terminal emulator.
  - Once into the Service Processor menus check date and time are set appropriately and set up the management port network settings (optional).
  - Connect Service Processor to management network and log in using a web browser (optional)
  - In the Service Processor menus select the option to power on the server.
  - Follow installation procedures for DVD or network installation.

# Installation in the Real World™ Part 2...



## ▪ RHEL 7 (BE/LE)

- POWER8 (native mode) and POWER 7/7+ at GA
- Available June 2014
- 7.1 adds LE distribution

## ▪ 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 (LE)

- POWER8 (native mode, LE) and POWER 7/7+
- Technology preview of KVM host support
- Available October 2014

## ▪ SLES 11 (BE)

- 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 11 High Availability Extension

## ▪ Ubuntu 15.04 (LE)

- Docker enablement
- FPGA support

## ▪ Ubuntu 14.10 (LE)

- S824L GPU enablement

## ▪ Ubuntu 14.04 (LE)

- POWER8 enabled (native mode, LE)
- No official support for POWER7+ and older systems
- 64-bit only.
- KVM hosting enabled in 14.04.02

## ▪ Supported add-ons

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

## ▪ Debian

- Supports Power as of Sid release
- Official architecture for Jessie

# Performance Considerations

## ■ Processors:

- In a PowerVM environment you have the option of specifying dedicated processor cores and shared processor fractions of a core – DO set sensible values for these, don't think that setting 0.1 (10<sup>th</sup>) of a core as the desired value will be enough to run heavy (transactional) work loads – Use the same rules for partition design as for AIX.
- In a PowerKVM environment to get the best / most efficient use of the processors try not to set up many virtual machines with low processor values and expect it to work out the best way to manage the resources – if high performance is important don't overcommit, pin each virtual machine to a set number of cores.




## ■ Storage:

- DO NOT ALLOCATE A SINGLE DISK / LUN TO A PARTITION / VIRTUAL MACHINE AND EXPECT THE DISK PERFORMANCE TO BE AMAZING!!!!
- Allocate a number of disks / lun's to a virtual machine and use LVM to create logical volumes to work with.
- Multipathing is VERY useful for more efficient performance and help against path failure but set up can be "exciting" – make sure SAN zoning is correct for multipathing (you can use a single test lun to prove the pathing) before allocating lots of disks / lun's.
- Monitor the disk queues using **iostat -xm** during normal application operation to make sure they are not constantly higher than the disk queue depth values.

# Performance Considerations

## ■ Memory:

- Allocate as much memory as possible to each virtual machine up front and minimise changing the memory values of all virtual machines to give them the best possible resource placement and performance.
- Swap space should be set up to allow for a kernel crash dump to be held and to handle applications over committing memory, but you do not need to set up 100's of Gigabytes worth of space, normally a distribution installation routine will calculate an appropriate swap size if you let it choose default partitioning so let it do that first and get a value before changing to custom partitioning (if you so choose to).



## ■ Applications / Compilers:

- When using commercial software (ISV supplied including IBM software) make sure you follow the recommended settings from those organisations – this may include memory tuning values, disk layout recommendations, minimum processor allocation needed etc. You will be surprised how many customers don't initially follow those and get caught out.
- Custom / Bespoke application software, possibly in house created, may require re-compilation so make sure the latest development packages are installed, the IBM Advanced Toolchain is in place and try to use the IBM xLC/xLF compilers to get better performance over the GCC compilers.





# 2015

## IBM Power Systems & System Storage Technical University

26-30 October | Cannes, France



Registration &  
info at  
[ibm.com/training](http://ibm.com/training/events)  
/  
events



The **2015 IBM Power Systems & System Storage Technical University** will explore the Power Systems & System Storage capabilities in depth and allow you to take a deep dive into smarter computing implementation, product announcements, as well as key topics across the Power Systems & System Storage platform.

### WHY ATTEND

**Expand your knowledge and expertise** on key topics by attending business and technical sessions, trends & directions sessions, technology updates in a variety of formats such as lectures, hands on labs, live demo's and Q&A sessions.

**Exchange ideas and interact** with colleagues and competitors, leading-edge thinkers, technology gurus, executives and business experts from IBM.

**Join us at the IBM Power Systems & System Storage Technical University in Cannes.  
Watch. Listen. Learn more about IBM technologies.**

**Next Time**

**TBC – 7<sup>th</sup> Oct – Remote Restart**



Previous Sessions:  
POWER8 E850 from exper  
PowerSC Tools for IBM i  
HMC 8.20 Tech Preview  
The "Key" to IBM i Licensing  
POWER8 E870 from exper  
Linux on Power: Best Pract  
Linux for AIX/IBM i guys  
PowerKVM Deep Dive  
More Tricks Power Masters  
Power8 from hands-on  
Power up your Linux  
PowerVC  
And more.....

Future Sessions →

- More being planned
- Suggestions Welcome



Webinar wiki: <http://tinyurl.com/PowerSystemsTechnicalWebinars>  
Youtube Channel: <http://tinyurl.com/IBMPowerVUGYoutubeChannel>

Twitter:   
Gareth Coates @power\_gaz  
Jyoti Dodhia @JyotiDodhia

Nigel Griffiths @mr\_nmon  
Mandie Quartly @mandieq