Power System Technical Webinar Series

Power Systems

IBM

Today

Power8 : from hands-on experience

Starting at 10:00 am UK time by Nigel Griffiths & Gareth Coates

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

Linux

Previous Sessions:
Power up your Linux
PowerVC
PowerVP
SSP4
Best Practices
Tricks of Power Masters
IBMi and External Storage
Monitoring with ITM
Whole Machine Monitoring
Electric Server Agent
RDX Removable disks
Dynamic Platform Optimiser
And more.....

Future Sessions →

- 11th Jun - More Tricks of the Power Masters
- Suggestions Welcome

Twitter:
Gareth Coates @power_gaz
Nigel Griffiths @mr_nmon
Jyoti Dodhia @JyotiDodhia
Mandie Quartly @mandieq
Website: http://tinyurl.com/PowerSystemsTechnicalWebinars
Youtube Channel: http://tinyurl.com/IBMPowerVUGYoutubeChannel

Nigel Griffiths & Gareth Coates
EMEA Power Advanced Technology Support

Power Systems

Hands-on with POWER8 - What is it actually like?

Webinar
21st May 2014

v7

Hands-on with POWER8 - What is it actually like?

Power Systems

Open Innovation to Put Data to Work

Nigel Griffiths IBM Power Systems Gareth Coates
Advanced Technology Support, Europe

Agenda

1. ESP and beta testing life
2. It arrives → pics
3. Unpacking and Physical Install → pics
4. HMC software → demo Properties, 16 core, RAM, I/O
5. Power up → talk
6. Virtual I/O Server 2.2.3.3 & SSP4 → talk
7. Operating Systems → initial testing
8. Which OS "understand" POWER8 = SMT8 ?
   - AIX7 + sp, RHEL7, Fedora
9. LPM → as expected
10. Speed: start/stop, apps → paraworms!
11. PowerVP / PowerVC → not yet but what we expect
12. Beta FW/VIOS/OS →Indicative tests
13. Some new HMC features

2. It arrives → pics

We received the systems …

… in IBM Bedfont Lakes

… moved them to IBM Southbank

… reconfigured all of our server and networking Infrastructure …

… and started to unpack it.

3. Unpacking and Physical Install → pics

Now the main performer

Good quality Packaging

Better to use the white clips and lift off the lid, rather than using a knife.

3. Unpacking and Physical Install → pics

Rack rails for POWER8 4U

- The first job was to install the rails.
- We found a few problems in the documentation
  - which we fed back to the ESP team

ESP
Early Support Program
Pre GA field testing

- Which rail is which?
- Hard to see the markings.
- Tilt to “get the right light”

3. Unpacking and Physical Install → pics

It looks like a Power server

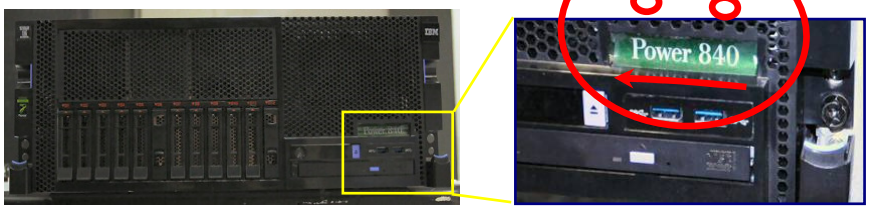

3. Unpacking and Physical Install → pics

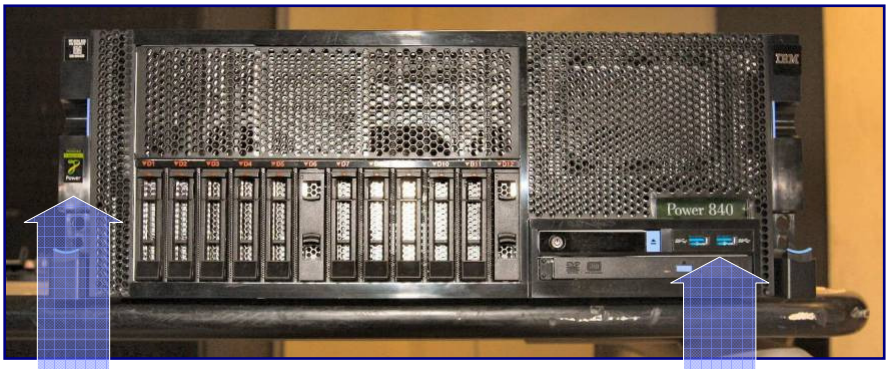

Server Bezels

POWER8

USB 3

3. Unpacking and Physical Install → pics

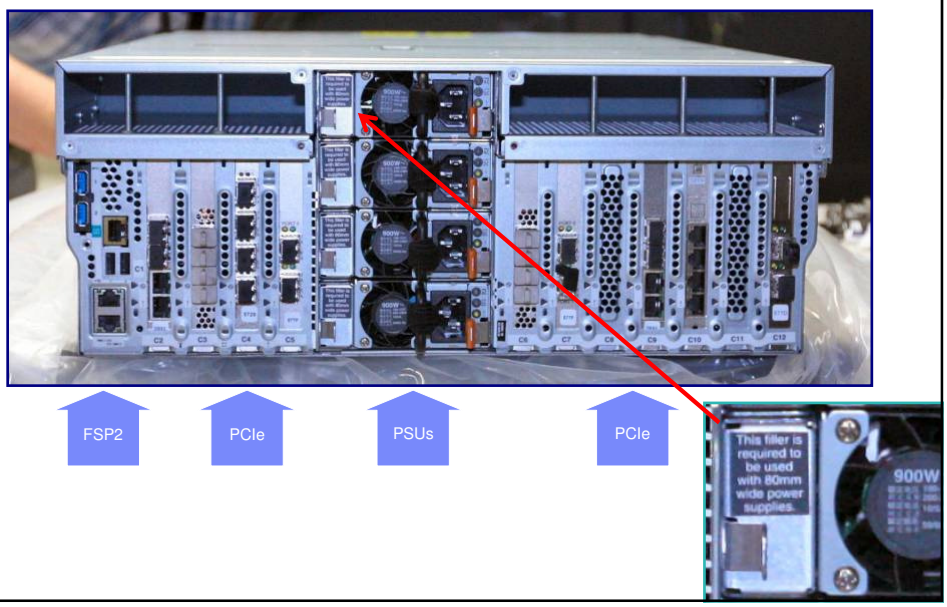

Some changes expected in PSUs for GA

3. Unpacking and Physical Install → pics

We needed to take it to pieces, well, boys and their toys ….

PSUs

Back

Front

3. Unpacking and Physical Install → pics

The bezel comes off

The operator panel is much better looking than the old one used to be.

3. Unpacking and Physical Install → pics

3. Unpacking and Physical Install → pics

Good diagrams on the lid

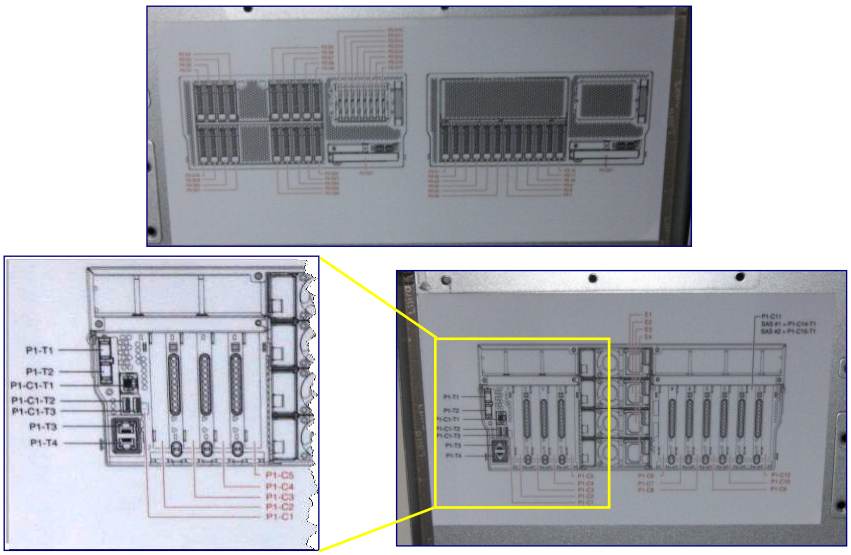

3. Unpacking and Physical Install → pics

The processors and DIMMs have a cover to optimise air flow for cooling

A DIMM

Remove the fan tray to access SAS Raid adapters

3. Unpacking and Physical Install → pics

… insides …

PCIe slots

which are hot swappable

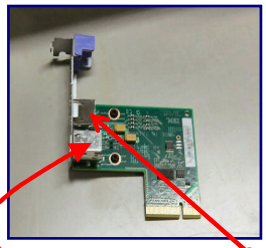

FSP2 RS232 (RJ45) and two USB

VPD "lollipop"

VPD "lollipop"

PSU

3. Unpacking and Physical Install → pics

Putting it on the rails

Nail head in slot

Nail head

Clip to release rail

Front Right

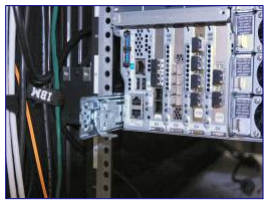

Rear Left

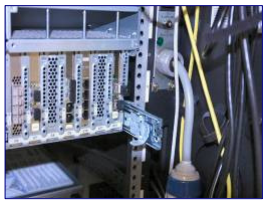

Rear Right

Cable Management Arm

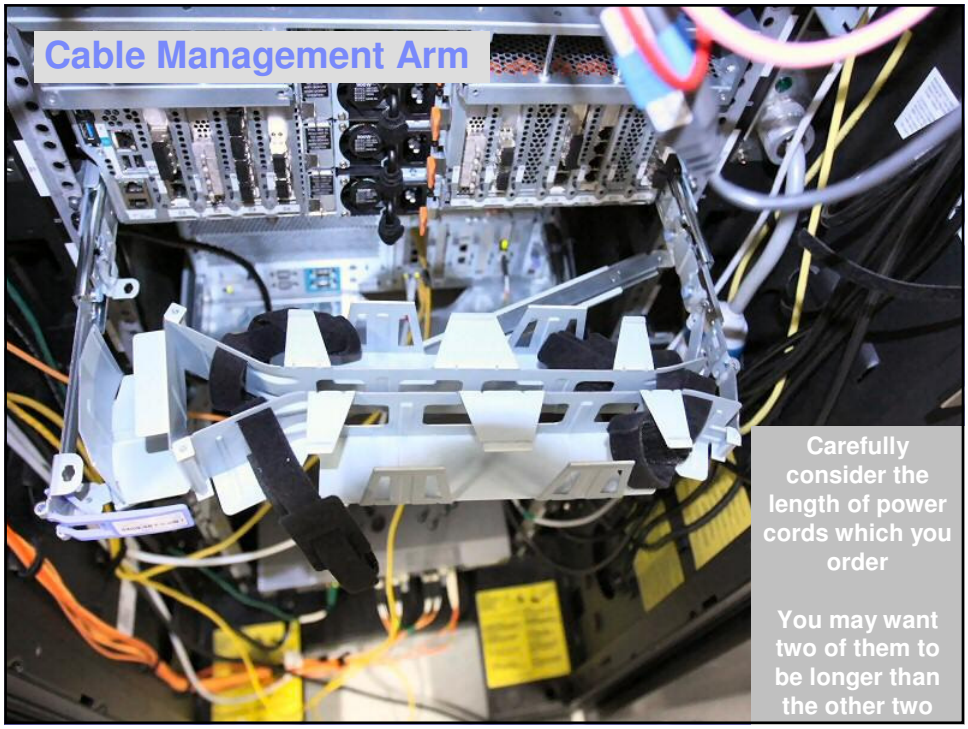

Carefully consider the length of power cords which you order

You may want two of them to be longer than the other two

3. Unpacking and Physical Install → pics

We started to cable up and installed the HMC

Four power cords
Plugs do NOT go through brackets

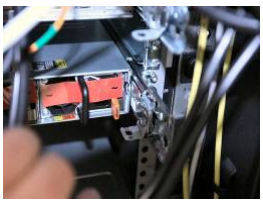

HMC PSUs

HMC Network ports

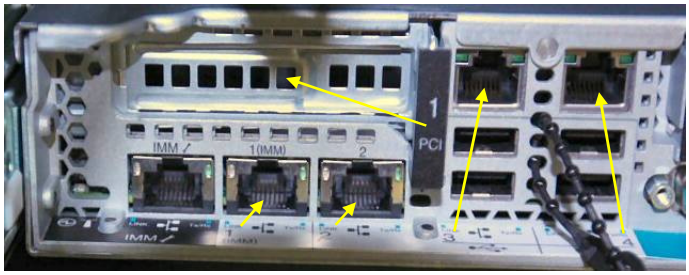

3. Unpacking and Physical Install → pics

We started to cable up and installed the HMC

4. HMC software

- The HMC (7042-CR8) was supplied with software installed ☺
- HMC V7R780 – which does not support POWER8 ☹
- So, we upgraded
  - Over the network
  - Not supported at this stage – but it worked☺
  - We installed the second HMC from media

We test as many things as possible

- Now running:

```
"version= Version: 8
 Release: 8.1.0
 Service Pack: 0
HMC Build level 20140121.2
HMC Driver FRZ
","base_versio
"
```

```
"version= Version: 8
 Release: 8.1.0
 Service Pack: 0
HMC Build level 20140401.2
","base_version=V8R8.1.0
"
```

5. Power up → talk

The interface is very familiar

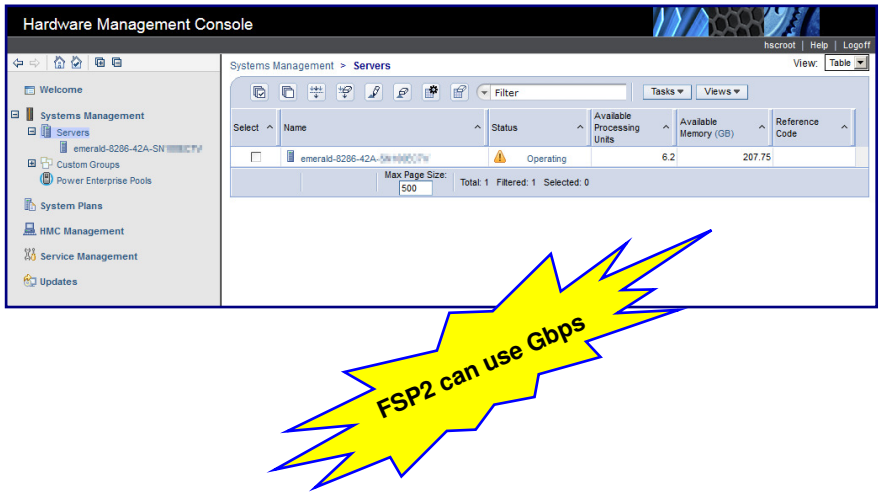

5. Power up → talk

… and it all seems to work well ☺

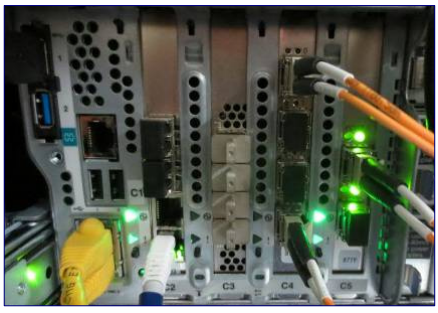

Back Left

Back Right

http://tinyurl.com/p8hotswapblog

developerWorks Technical topics Evaluation software Community Events

My home Forums Blogs Communities Profiles Podcasts Wikis Activities IBM Champion program

My Blogs Public Blogs My Updates

AIXpert Blog

POWER8 Scale-Out models have Adapter Hot-Plug (Hot-Swap)

nagger | Apr 29 | Comments (2) | Visits (914)

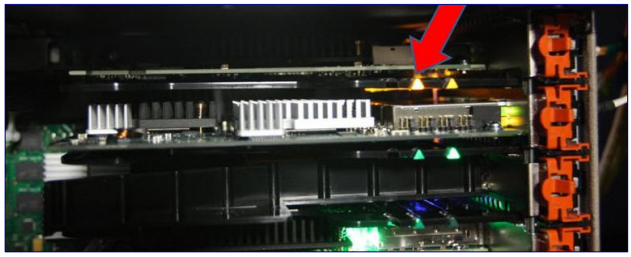

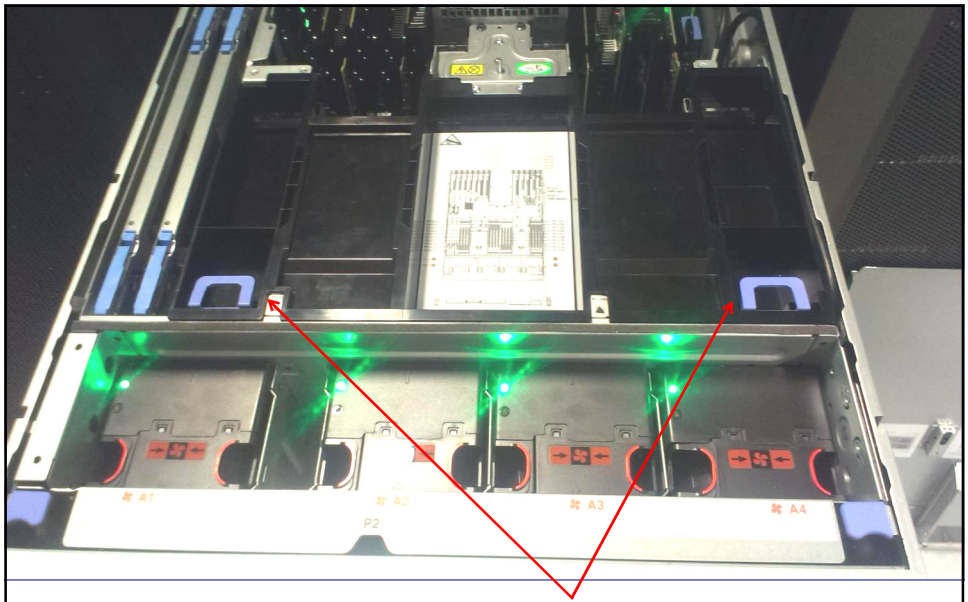

•To raise the baffle to access the Memory and Processors, push the blue tabs from the front towards the back of the machine

•Do NOT pull the blue tabs upward!

•The DIMMs and CPUs are not hot swappable so the machine would need to be powered down.

•Lifting the cover is OK as you need to do this to hot swap the PCIe adapters and Air Movement Devices (fans)

•Follow the exchange FRU process via the HMC.

6. Virtual I/O Server 2.2.3.3 & SSP4 → talk

- Small update to VIOS
- Large update with new function every Q4

- Q2 – support for new FW & adapters only
- SSP4 functions the same – fixed a few documentation errors!

- We are heavy SSP4 users
- Allows total FW, OS and SW refresh then pulling in the same LPARs in 20 seconds.
- Have a SSP between POWER6 and POWER8

7. Which OS "understand" POWER8 = SMT8 ?

1. AIX 7 TL3 with extra SP
   - For POWER8 native adapters
   - SMT=8
   - Interfaces/libraries for apps to compile in
2. RHEL7 public beta
3. SLES 12 private beta
4. Fedora 20

8. OS - initial testing

- nmon on AIX – what else!!
- smtctl – Check SMT=8
- nmon PMR about the PCPU and SCPU filed stats

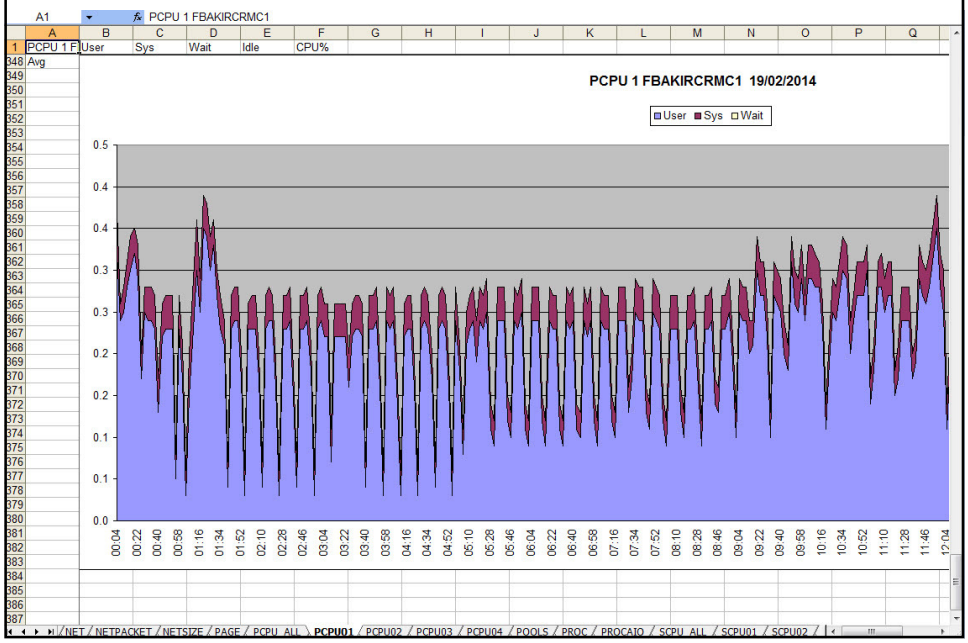

9. LPM

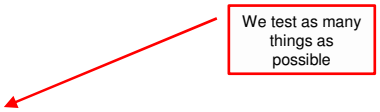

Basically, it works! ☺

1. CLI and GUI
2. POWER6 ⟷ POWER8
3. POWER7 ⟷ POWER8
4. SSP or Traditional VIOS
5. Support for as many as 16 LPMs in parallel

10. Speed: start/stop, apps → paraworms

1. LPAR start and stop seems quicker
   - But hard to test
2. As we have full POWER7/7+ support
   - Everything works
   - Rather boring ☺
3. POWER8 mode primary difference is SMT=8
   - Most application are not SMT aware
   - Available in minor SP for AIX
   - Available some Linux only – can't force synchronous GA

11. PowerVP / PowerVC

1. ESP Internal ONLY Beta's not yet available
2. Always planned for mid April or later
3. What we expect
4. PowerVP
5. PowerVC
   1. Scale
   2. PowerLinux / PowerKVM support
   3. SSP → lots

PowerVP

PowerVC

PowerKVM

12. Beta FW/VIOS/OS →Indicative tests

Running today

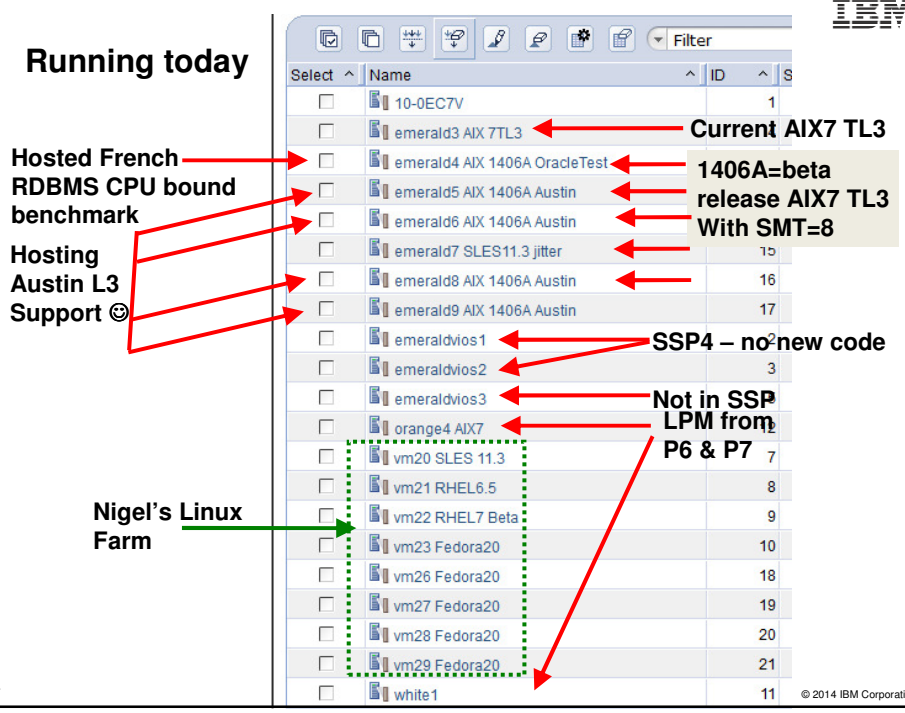

Fedora install without dhcp bug

Eventually worked around it
See Nigel's public **AIXpert blog** for all the details
"Fun with Fedora"

```
setparams 'Install Fedora 20 (64-bit kernel)' '64'

     linux /ppc/ppc${2}/vmlinuz  ro vnc vncpassword=abc12345 ip=9.137.62.23 n\
etmask=255.255.255.0 gateway=9.137.62.1
     initrd /ppc/ppc${2}/initrd.img
```

Note:
Linux on Power blogs get a lot of hits

Single Thread Performance tested by Worms

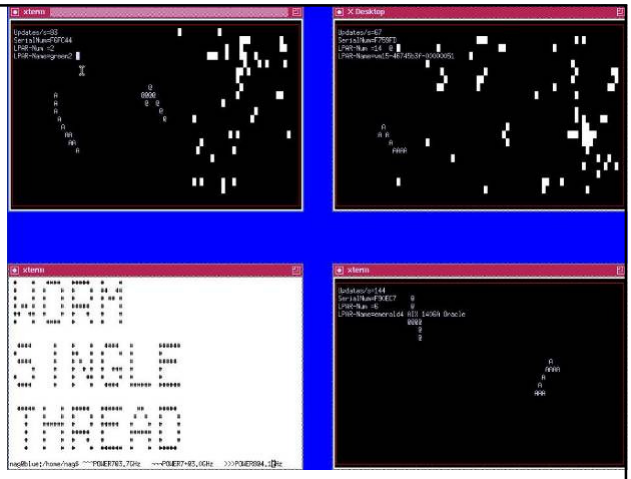

Adjusted for GHz

POWER7+ 3.0 GHz 92 updates per second
POWER8 4.1 GHz **144** updates per second

Conclusions:

- **'Wow! POWER 8 single threaded CPU bound workloads really "kick ass" at 56% faster (144/92*100) even allowing for GHz ratings.'**

39

SMT=8 tested by paraWorms

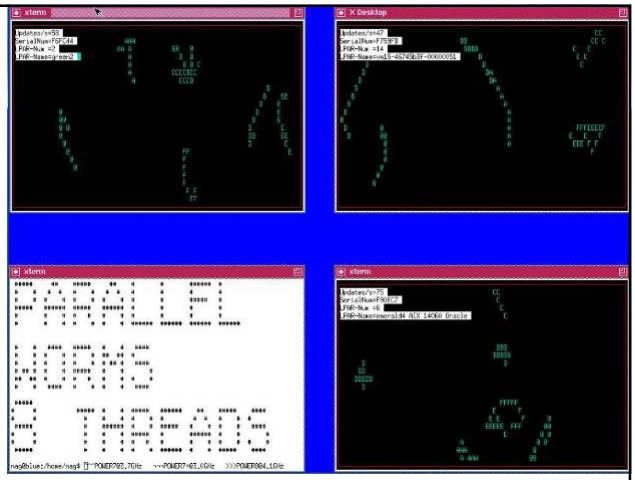

Adjusted for GHz

POWER7+ 3.0 GHz 64 updates per second
POWER8 4.1 GHz **75** updates per second
Conclusions:

- **That SMT=8 POWER8 trick has kicked in goes 17% faster**

40

Memory tests

nmem64 -m 1200 -s 10
1.2 GB random memory access
Forces real DIMM accesses
POWER8 7.9 million/sec **= 60% faster**
POWER7/7+ 4.5 to 5 million/sec

nmem64 -m 8 -s 10
8 MB random memory access
Can be cached at L3
POWER8 = 18.1 million/sec **= 63% faster**
POWER7+ = 11.1 million/sec

So Jeff Stuecheli gets to keep his job ☺
- Mr Power Memory Architect

POWER8 Wafer