

The slide features a blue background with a large 'POWER9' logo on the right, composed of two concentric arcs in green and blue. Above the logo are three circular icons: 'AIX', 'i For Business', and 'Linux'. The IBM logo is in the top right corner. The main text on the left reads: 'IBM Power Systems POWER9 Enterprise Model E980 First Look'. Below this, it says 'Version 7' and 'This is an unofficial document It is not IBM approved or IBM reviewed' in red. At the bottom, it says 'POWER Advanced Technology Support EMEA (UKI)'. Five headshots of team members are shown at the bottom with their names and contact information: Nigel Griffiths (nag@uk.ibm.com), Mike Pearson, Mike J Smith, Alain Jeantet, and Michael Poli.

IBM Power Systems
POWER9 Enterprise
Model E980
First Look Version 7
This is an unofficial document
It is not IBM approved or IBM reviewed

POWER Advanced Technology Support EMEA (UKI)

Nigel Griffiths
nag@uk.ibm.com

Mike Pearson


Mike J Smith

Alain Jeantet

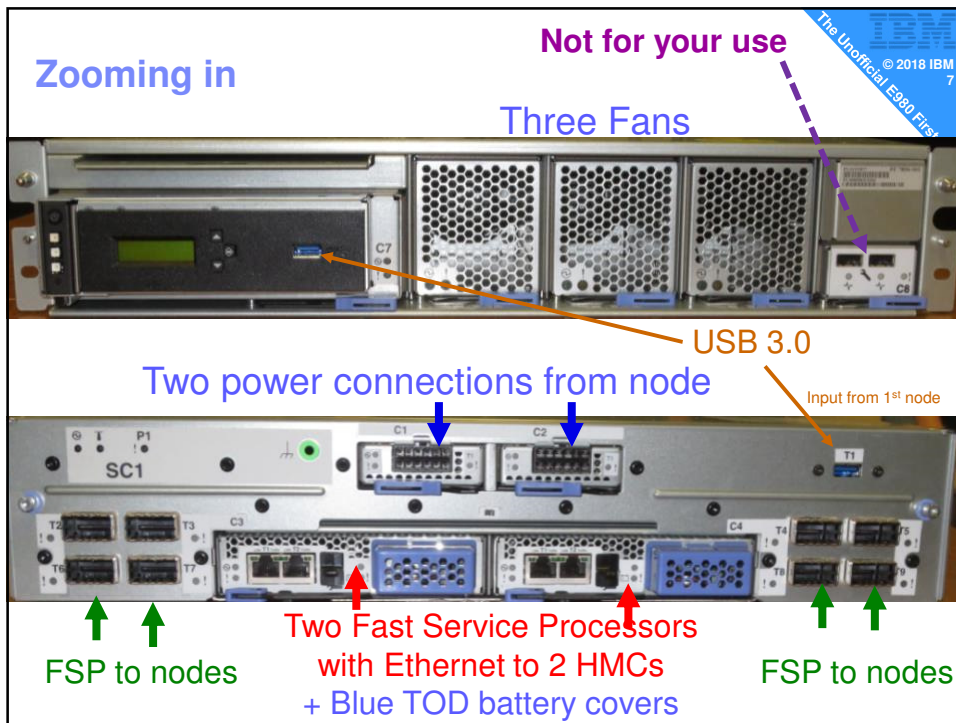
Michael Poli


Introduction

- Pictures are from an Early Ship Program (ESP) Server
 - Arrived in the ATS, London, UK → 21th June 2018
 - Yours might be slightly different
- This is a Single node and System Control Unit
 - So you will not see the inter-node cables (until the end)7
- A heads-up on what to expect with POWER9 Enterprise
- **POWER9 E980 is SSR (Client Engineer) install**
- If you have E880 experience you know 75% already



















E980 Rails & cable check




Node power to System Unit cable





USB cable



Node <-> FSP cable




Getting organised






E980 Adapter cassette





The Unofficial E980 First Look
© 2018 IBM
23

E980 Adapter cassette



The Unofficial E980 First Look
© 2018 IBM
24





Is the E980 Adapter cassette the same as the E880 Adapter cassette?



Looks very similar - working the same way


But not the same:

Mechanically slightly different = slimmer
Additional RAS & usability features
Now PCIe Gen4 based

Note there are electronics in the cassette








2 inch NVMe device (works as a fast SSD)



The Unofficial E980 First Look
© 2018 IBM
32

**Cable support bracket side pieces screw on
(unlike the E950!)**

The Unofficial E980 First Look
© 2018 IBM
33




E980 front


- Five air movement devices (fans)
- Four power supplies
- - Power Cable run down channels to the rear

The Unofficial E980 First Look
© 2018 IBM
34











How different to E880?

E980 2018

- Memory DDR4
- 15 VRMs




E870/E880 2015

- Memory DDR3 initially
- 12 VRMs




© 2018 IBM
42
The Unofficial E980 First Look





The Memory card with Centaur L4 & DDR4 chips

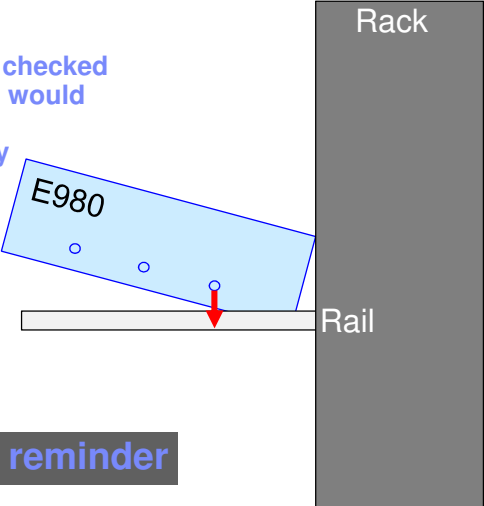


Installing the E980 node

Used the E880 approach

Nail heads nearest the rack locked in first then lower front down


Don't let go until you have triple checked
- I am not clear whose insurance would pay out if it got dropped!
- for E980 damage or client injury




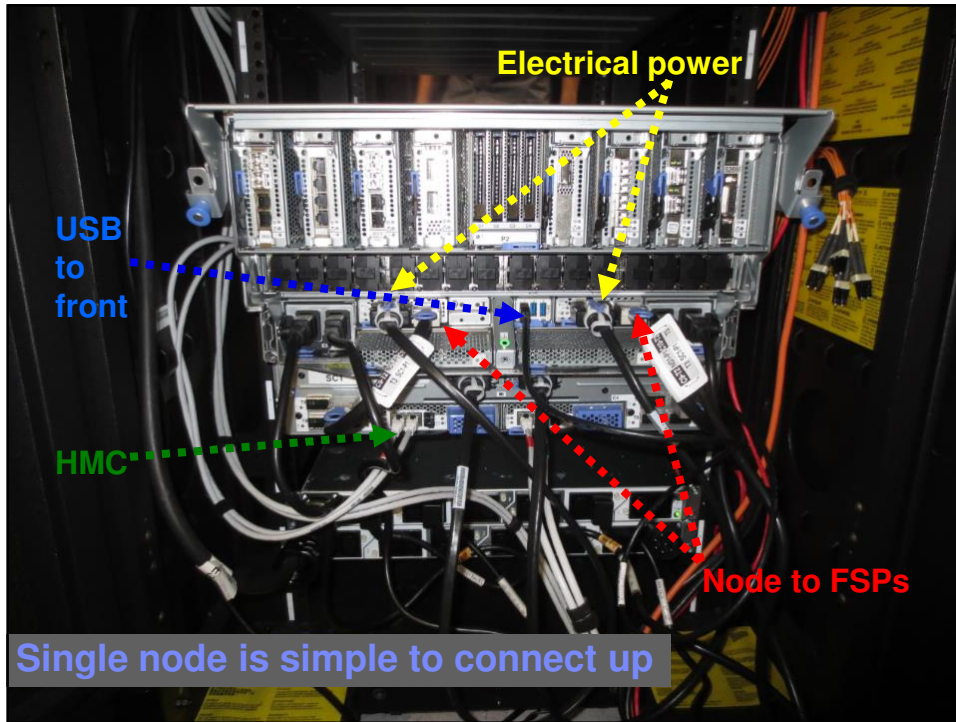
Installing the E980 node reminder

The Unofficial E980 First Look
© 2018 IBM
47









The Unofficial E980 First Look
© 2018 IBM
55

Yes, that is far too easy but we did not cheat!

Only one node = zero node to node cables

Next chart is borrowed from Gareth's E980 Deep Dive for high numbers of nodes

The Unofficial E980 First Look
© 2018 IBM
56


SMP cabling

SMP Cable Lengths

- 655mm
- 780mm
- 1000mm
- 1200mm
- 1400mm

Each connection shown in the diagram represents a pair of SMP cables

SMP cabling



Summary: Hardware Install 😊

- **Easy**
 - OK we only had 1 low down node & no inter node cables
- If previously installed a E880 = no brainer
- If new → follow manual
- If not new → follow the manual

- No clock cable makes the E980 simpler
- Inter node cables are clever:
 - on plugging-in LEDs flash to highlight if you got it right/wrong

- Regular E880/980 rules apply
 - HMC's running 920 or above
 - Prepare for arrival: site survey, weight, height, power, network, SAN
 - Not client install
 - Its heavy so buy / order a Lift Tool
 - These are high end servers – so get it right first time → RTFM
 - Ear protection on power up fan noise

Software Install 😐

- Connect up the HMC running 920+ software

- Quickly setup the basics
 - Quad VIOS
 - VIOS have FC SAN adapters to Shared Storage Pool (SSP)
 - A major asset for our crash & burn environment
 - Our record beating AIX LPAR with 11,000 LPMs

- Bad news
 - It's a 3 month loaner
 - E980 (like the E880) take quite a long time to cold boot-self test
 - These servers only start once and then run for five years

Initial power - up

Hardware Management Console

hmc15 Resources > All Systems

All Systems

View and monitor the state, health, and capacity information of all the systems that are connected to the management console.

Select All | Actions | Total: 4 Selected: 0

System Name	State	CPU Available	Memory Available
P9-E950-silver	Operating	0.6%	63%
P9-S922-amber	Operating	0.3%	13%
P9-S924-red	Operating	0.4%	22%
Server-9080.M9S-SN130A148	Initializing	32 CPU Available	5120 GB Available

Bug: The HMC guys promised to fix the picture

Server renamed to "brick" – all servers are colours!

Hardware Management Console

hmc15 Resources > All Systems

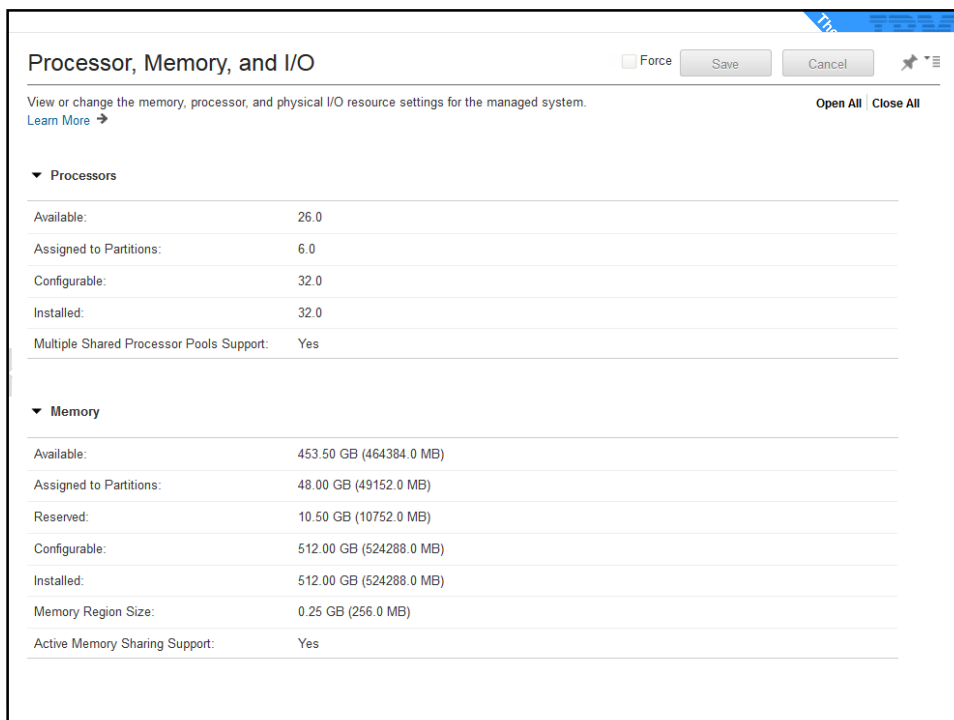
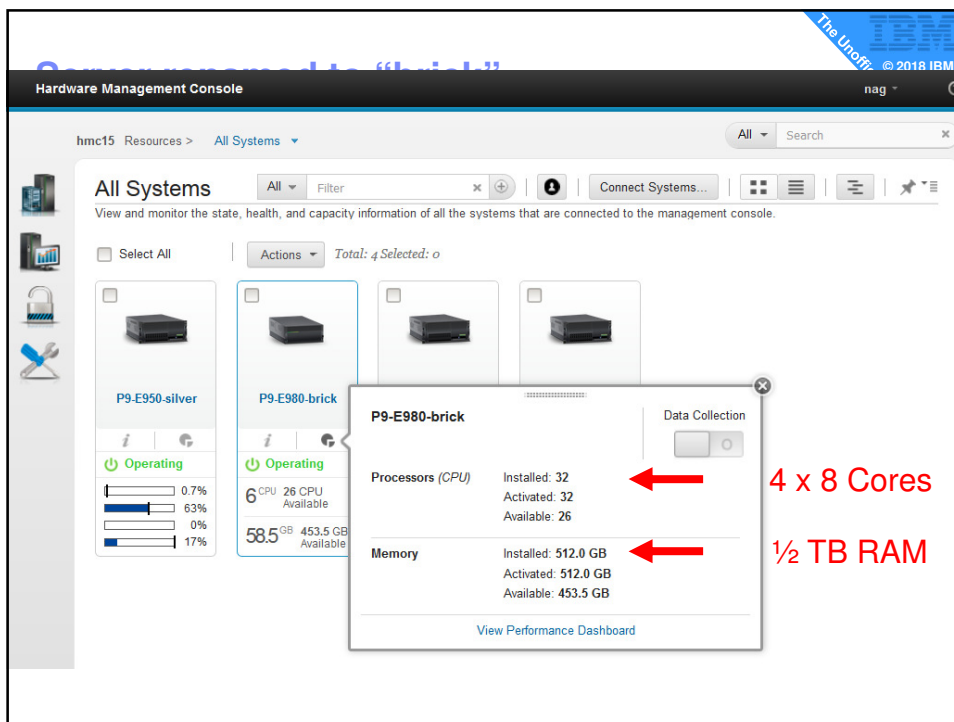
All Systems

View and monitor the state, health, and capacity information of all the systems that are connected to the management console.

Select All | Actions | Total: 4 Selected: 0

System Name	State	CPU Available	Memory Available
P9-E950-silver	Operating	0.8%	63%
P9-E980-brick	Operating	6 CPU Available	58.5 GB Available
P9-S922-amber	Operating	0.3%	13%
P9-S924-red	Operating	0.4%	22%

Server renamed to "brick" – all servers are colours!



Processor, Memory, and I/O Force

View or change the memory, processor, and physical I/O resource settings for the managed system.
[Learn More](#) →

Adapter Description	info	Physical Location Code	Owner	Bus Number	SR-IOV Capable (Logical Port Limit)
PCIe2 4-port (10Gb FCoE & 1GbE) SR&RJ45 Adapter	/i	U78D5.001.CSS151D-P1-C1-C1	brickvios1	16	Yes (Dedicated mode)
1 Gigabit Ethernet (UTP) 4 Port Adapter PCIe-4x/Short	/i	U78D5.001.CSS151D-P1-C2-C1	Unassigned	17	No
PCIe2 4-port(10Gb FCoE & 1GbE) Copper&RJ45	/i	U78D5.001.CSS151D-P1-C3-C1	brickvios2	18	Yes (Dedicated mode)
Empty slot	/i	U78D5.001.CSS151D-P1-C4-C1	Unassigned	19	Yes (510)
Empty slot	/i	U78D5.001.CSS151D-P1-C5-C1	Unassigned	20	Yes (510)
PCIe3 4-Port 16Gb FC Adapter	/i	U78D5.001.CSS151D-P1-C6-C1	brickvios1	21	No
8 Gigabit PCI Express Dual Port Fibre Channel Adapter	/i	U78D5.001.CSS151D-P1-C7-C1	brickvios2	22	No
Empty slot	/i	U78D5.001.CSS151D-P1-C8-C1	Unassigned	23	Yes (510)
800GB NVMe Gen3 U.2 Slim SSD	/i	U78D5.001.CSS151D-P1-P2-C1	brickvios1	24	No
Empty slot	/i	U78D5.001.CSS151D-P1-P2-C2	Unassigned	25	Yes (240)
800GB NVMe Gen3 U.2 Slim SSD	/i	U78D5.001.CSS151D-P1-P2-C3	brickvios2	26	No
Empty slot	/i	U78D5.001.CSS151D-P1-P2-C4	Unassigned	27	Yes (240)
Universal Serial Bus UHC Spec	/i	U78D5.001.CSS151D-P1-C13	Unassigned	28	No

PCIe Gen 4

NVMe

USB 3.0 built-in to first node also cabled to the front of the System Control Unit

The Unofficial E980 First Look
 © 2018 IBM
 66

Call to Client Action

Good News

1. E980 is not a massive external change from E880
2. HW Install is a no brainer
3. No clock cables = easier
4. Expecting a large POWER9 performance boost

Bad News

1. Nothing like the shocks we had with the E950 😊
2. Shows confidence to loan us a E980 this early
3. SECRET:
4. Some teething issues – for sure
5. Nothing unexpected 14 weeks before GA
6. Many already addressed by developers
7. Reporting a few PMR / PMH to get fixed

NVMe Usability

- NVMe is like SSD – it wears out
- Built with plenty of “spare” capacity to replace wear
- Life depends on its use → I/O’s per day
- We don’t count that on disks – they just fail
- NVMe has better engineering & monitors/reports wear
- nmvemgr command . . .

The Unofficial E980 First Look
© 2018 IBM
67

8 NVMe Usability

```
brickvios1:/home/padmin $ lsdev | grep -i nvme
hdisk0          Available    NVMe 4K Flash Disk
nvme0          Available    PCIe3 x4 NVMe Flash Adapter
silvrvios1:/home/padmin $

# nmvemgr -M -l nvme0
Critical Warning ..... 0x0
Composite Temperature (Kelvin) ..... 306
Available Spare (%) ..... 100
Percentage of NVM subsystem life used ..... 0
Data Units Read (1000 units of 512 bytes) ..... 1928423
Data Units written (1000 units of 512 bytes) ..... 2141752
Host Read Commands ..... 14109948
Host Write Commands ..... 29866123
Number of Power Cycles ..... 30
Power On Hours ..... 1620
Unsafe Shutdowns ..... 10
Media and Data Integrity Errors ..... 0
Number of Error Information Log Entries ..... 54
#

# lsvg rootvg
VOLUME GROUP:    rootvg      VG IDENTIFIER:  00fb601f0000163b6bdc0b7
VG STATE:        active     PP SIZE:        1024 megabyte(s)
VG PERMISSION:   read/write  TOTAL PPs:     745 (762880 megabytes)
```

The Unofficial E980 First Look
© 2018 IBM
68

The screenshot shows the IBM Advanced System Management (ASMG) web interface. The top navigation bar includes the IBM logo, the title 'Advanced System Management', and copyright information. Below the navigation bar, there is a user status bar showing 'Log out', 'User ID: admin', and 'P9-E950-silver'. The main content area is divided into two sections: a navigation menu on the left and a main content area on the right. The navigation menu is expanded, showing various system management options. Red arrows point to 'System Configuration' and 'Power and Performance Mode Setup' in the menu. The main content area displays a 'Welcome' message, system information (Machine type, serial number, date, time, service processor), current hardware and firmware uptime, and a table of current users and user status.

Navigation menu items:

- Expand all menus
- Collapse all menus
- Power/Restart Control
- System Service Aids
- System Information
- System Configuration
 - System Name
 - Configure I/O Enclosures
 - Time Of Day
 - Firmware Update Policy
 - PCI Error Injection Policy
 - HSL, Opticonnect Connections
 - IO Adapter Enlarged Capacity
 - Hardware Management Consoles
 - Floating Point Unit Computation Test
 - Virtual Trusted Platform Module
 - PCIe Hardware Topology
 - Hardware Page Table Size
 - Console Type
 - Predictive Dynamic Memory Deallocation
 - Speculative Execution Control
 - Hardware Deconfiguration
 - Program Vital Product Data
 - Service Indicators
 - Power Management
 - Power and Performance Mode Setup
 - Idle Power Saver
 - Tuning Parameters
 - Security
 - Network Services

Main Content Area:

Welcome

Machine type-model: 9040-MR9
 Serial number: 13601FX
 Date: 2018-6-6
 Time: 10:52:31 UTC
 Service Processor: Primary (Location: U78D4.001.CSS149W-P1-C1)

Current hardware uptime: 0 Days 2 Hours 21 Minutes 37 Seconds
 Current firmware uptime: 0 Days 2 Hours 9 Minutes 24 Seconds

Current users

User ID	Location
admin	10.255.128.1

User Status

User ID	Status
dev	Disabled
celogin	Enabled
celogin1	Enabled
celogin2	Disabled

The screenshot shows the 'Power and Performance Mode Setup' configuration page. The page title is 'Select + Continue for instant mode change'. The current power saver mode is set to 'Enable Maximum Performance mode'. Below this, there is a list of radio button options for different power modes. A 'Continue' button is located at the bottom of the page.

Power and Performance Mode Setup

Current Power Saver Mode : Enable Maximum Performance mode

- Disable all modes ?
- Enable Static Power Saver mode ?
- Enable Dynamic Performance mode ?
- Enable Maximum Performance mode ?

Note: Enabling any of the Power Saver modes will cause changes in the processor frequencies, changes in processor utilization, changes in power consumption, and performance to vary. Other effects are possible as well. Please see the EnergyScale™ white paper for more information on power saving modes.

?

nmon on AIX is OK

- POWER9 mode at 3.4 GHz + MTM+serial OK
- AIX 7200-02-02-1810

```

N N M M OOOO N N For online help type: h
NN N MM MM O O NN N For command line option help:
N N N M MM M O O N N N quick-help nmon -?
N N N M M O O N N N full-details nmon -h
N NN M M O O N NN To start nmon the same way every time:
N N M M OOOO N N set NMON job variable, for example:
export NMON=cat

-----
TOPAS_NMON
16 - CPU currently
16 - job configured
3400 - MHz CPU clock rate
PowerPC_POWER9 - Processor
64 bit - Hardware
64 bit - Kernel
S,AIX72-TL2_on_-762b86ea-00000 - Logical Partition
7.2.2.18 TL02 - AIX Kernel Version #
vm21_AIX72 - Hostname
vm21_AIX72 - Node/WPAR Name
130A148 - Serial Number
IBM,5900-N9S - Machine Type
    
```

nmon APPROVED

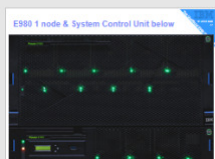
Max Perf mode at 3.9 GHz Beta chips

```

# lparstat -E 1 3
System configuration: type=Shared mode=Uncapped smt=4 cpu=16 mem=16384MB ent=2.00 Power=Dynamic-Performance
Physical Processor Utilisation:
-----Actual-----Normalised-----
user sys wait idle freq user sys wait idle
-----
0.001 0.001 0.000 1.998 3.8GHz[111%] 0.001 0.001 0.000 1.998
0.000 0.001 0.000 1.999 3.8GHz[111%] 0.000 0.001 0.000 1.999
0.000 0.001 0.000 1.999 3.8GHz[111%] 0.000 0.001 0.000 1.999
    
```


Load of detailed picture in the deck for you

E980 1 node & System Control Unit below




77

E980 2 node & System Control Unit below




78

E980 4 node System Control Unit in a rack with rack door open




79

E980 4 node System Control Unit




80

E980 4 node System Control Unit at an angle



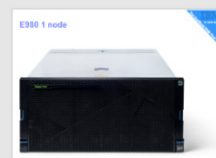
81

E980 1 node



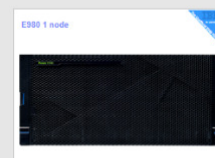
82

E980 1 node



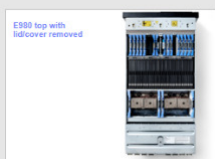
83

E980 1 node




84


E980 top with lid/cover removed




E980 1 node



E980 1 node



E980 during an explosion!





Before we end

- Slides
- Replays

- Much larger slide deck

PDF of todays slides & replay from

<http://tinyurl.com/PowerVUG>

Going to share the PowerPoint on the Power VUG website

<http://tinyurl.com/AIXpert>

I have ~250 slides including

~ 100 picture slides of the server

- Pictures are of a beta machine
- GA servers might differ slightly

+ More links to information