


# IBM POWER9 Enterprise E950

MTM: 9040-MR9 **Unofficial Deep Dive**





**This is Nigel's Personal Deck & Opinions – It is not an IBM Announcement or IBM Statement**

**Nigel Griffiths** Power Systems: Advanced Technology Support EMEA  
[nag@uk.ibm.com](mailto:nag@uk.ibm.com)  
<https://www.youtube.com/user/nigelargriffiths>


Twitter: @mr\_nmon  
 Version 7 (6<sup>th</sup> August 2018)

**Gareth Coates**  
[gaz@uk.ibm.com](mailto:gaz@uk.ibm.com)  
[@power\\_gaz](https://twitter.com/power_gaz)

11
IBM Power Systems POWER9 Enterprise E950 Unofficial Deep Dive
© IBM Corporation, 2018


# POWER9 Enterprise

**Power Systems POWER9 Enterprise E980**  
 192 POWER9 CPU cores @ 4.0 GHz  
 64 TB DDR4 CDIMM Memory  
 16 Internal NVMe for VIOS or OS booting  
 32 PCIe Gen4 adapters  
 30% jump in performance from E880  
 - due to whole system enhancements



**Power Systems POWER9 Enterprise E950**  
 48 POWER9 CPU cores @ 3.8 GHz  
 16 TB DDR4 ISDIMM Memory  
 4 Internal NVMe for VIOS or OS booting  
 10 PCIe Gen4 adapters + 1 Gen3  
 42% jump in performance from E850

This Presentation



**POWER9 Server performance:**

- Super strength SMT threads
- CPU memory caching
- SMP bus bandwidth
- Memory sizes
- PCIe Gen4 boost

12
IBM Power Systems POWER9 Enterprise E950 Unofficial Deep Dive
© IBM Corporation, 2018

**Announcement Date:** August 7<sup>th</sup> both **E950 & E980**  
**eConfig Date:** August 8<sup>th</sup>  
**General Availability Date:** August 17<sup>th</sup> **E950**

Sept 21<sup>st</sup> **E980 2 nodes= 96 core**  
 Nov 16<sup>th</sup> **E980 4 nodes=192 core**

**Upgrade from POWER8 &  
 POWER9 2 to 4 node  
 MES upgrades:**

**In Q4 for easy upgrades**

13

IBM Power Systems POWER9 Enterprise E950 Unofficial Deep Dive

© IBM Corporation, 2018

**Before we start**

**PDF of todays slides & replay from**

**<http://tinyurl.com/PowerVUG>**

- Not going to cover every point
- Not covering all the slides
- Not covering the market slides

**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

**Plus loads of link to more information**

14

IBM Power Systems POWER9 Enterprise E950 Unofficial Deep Dive

© IBM Corporation, 2018

There will be typo's, mistakes and more information becoming available

Check for updates on the Power System VUG website

If you spot mistakes please report them to me Nigel Griffiths  
→ [nag@uk.ibm.com](mailto:nag@uk.ibm.com)

15

IBM Power Systems POWER9 Enterprise E950 Unofficial Deep Dive

© IBM Corporation, 2018



POWER9 requires HMC 920 software  
(an update to HMC 910)

- Supported on CR7, CR8, CR9 or
- HMC 7063-CR1 POWER8 based
  - More CPU, RAM, Disk & slightly lower cost
  - Simple: remote start/stop with ipmitool +WebUI

#### Gotchas

- HMC Enhanced+ GUI (No Classic GUI)
- No POWER6 support
- HMC hardware too old
- IVM not supported
- Intel based HMC out of stock in many countries



HMC News  
Hardware Management Console


16

IBM Power Systems POWER9 Enterprise E950 Unofficial Deep Dive

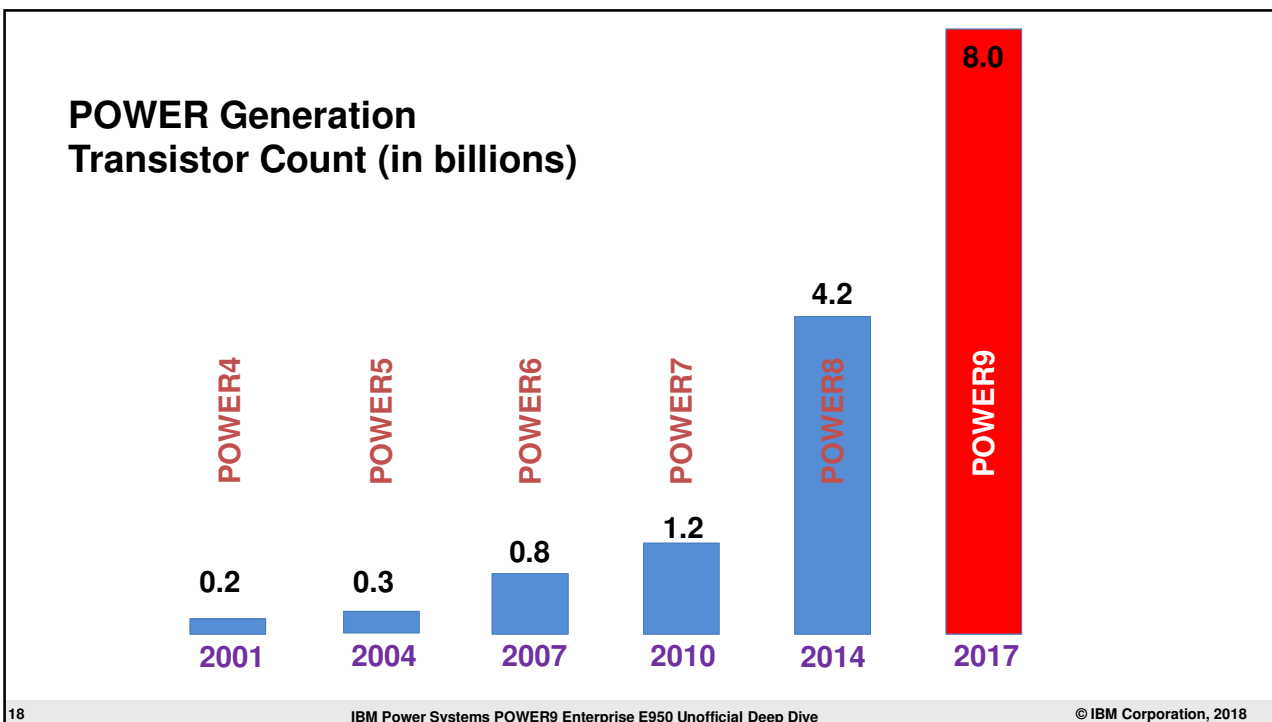
© IBM Corporation, 2018

# POWER9 Processor

features that effect the Servers



17 IBM Power Systems POWER9 Enterprise E950 Unofficial Deep Dive © IBM Corporation, 2018





# POWER9 Chip Variations





Models  
AC922  
LC922  
LC921

**December**  
OpenPOWER  
SMT4  
Up to 22 Cores  
DDR4 RAM  
Linux Only



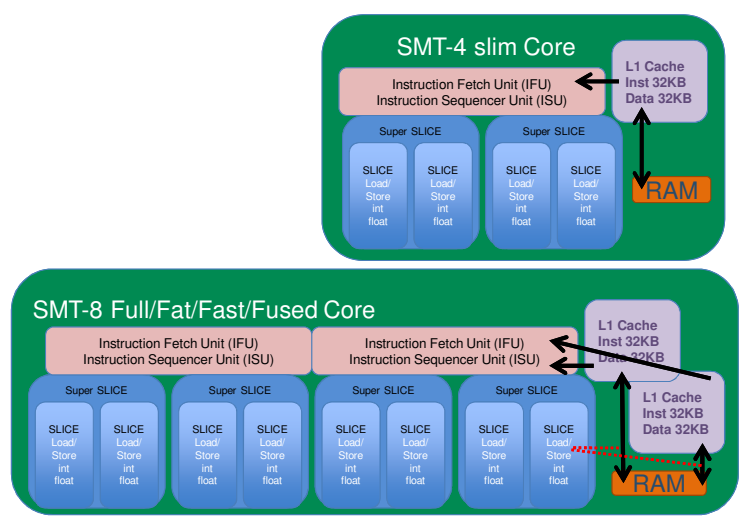
**IBM Scale-Out SMT8**  
Up to 12 cores  
Centaur with DDR4 RAM  
AIX + IBM i + Linux



Models  
E950  
  
E980  
- 2 CEC  
- 4 CEC

The E950 does not support IBM i

## Wikipedia on POWER9 fast core



Each of the eight CPU core threads gets at least a slice & up to 8 slices

## POWER9 Fused core STRENGTH

### POWER8 SMT8 Core Resources

#### Issue of VSU and AGEN

- 2x load AGEN / simple-ALU
- 2x load/store AGEN
- 2x scalar-64b / vector-128b
- 2x FXU

#### Vector Scalar Unit (VSU) Pipes

- 2x FP (64b/128b) + Complex (128b)
- 2x ALU (128b)
- 2x Permute (128b)
- 1x Decimal FP
- 1x Cryptography

#### Fixed Point (FXU) Pipes

- 2x ALU (64b)
- 2x FX-MUL + Fixed Divide (64b)

#### Load Store Unit (LSU) Slices

- 64kB, 8-way Data Cache
- Up to 4 DW load or 2 store
- 1x Store complete

x4

x4

x2

x2

x2

### POWER9 SMT8 Core Resources

#### Issue of VSU and AGEN

- 8x scalar-64b / 4x vector-128b
- 8x load/store AGEN

#### Vector Scalar Unit (VSU) Pipes

- 8x FP + FX-MUL + Complex (64b slice)
- 8x ALU + Simple (64b slice)
- 4x Permute (128b)
- 4x Quad Fixed (128b)
- 4x Fixed Divide (64b)
- 2x Quad FP / Decimal FP
- 2x Cryptography

#### Load Store Unit (LSU) Slices

- 64kB, 8-way Data Cache
- Up to 8 DW load or store
- 2x Store complete



# POWER9 Processor Peripherals



## New 19" Rack 7965-S42




**GA 4Q17**  
**POWER8 & POWER9**

|  | S42               | T42              | 94Y         |
|--|-------------------|------------------|-------------|
| 42U  | Yes               | Yes              | Yes         |
| 600mm Wide (datacenter floor tile)                     | Yes ★             | No               | Yes         |
| Ship Loaded from Factory                               | Yes               | Yes              | No          |
| Flat surface for mounting H2O Manifolds and Strip PDUs | Yes ★             | No               | Yes         |
| 1200mm Depth (rack w/ covers)                          | 1070+130cvrs      | 1016+cvrs        | 1040 + cvrs |
| Rear door heat exchanger                               | Yes               | Yes              | Yes         |
| # Vertical, 1U Pockets                                 | 4                 | 4                | 6           |
| Height Reduction – fit standard doorways               | Yes               | Yes              | No          |
| Back cable depth (mm)                                  | 280 ★             | 246              | 261         |
| Earthquake certified                                   | Yes – 45lbs / EIA | Yes – 35lbs/ EIA | No          |

OK. it is only a Rack! Perhaps, it is time to retire the older T42. **Warning: E950 is tight in a T42!**

## POWER9 Adapters

- PCIe GEN 4
  - Good timing for future proofing
  - Double the I/O rates
  - GEN 4 for next generation adapters like of 40+ Gb /sec
- Initially most adapters will be GEN 3
  - Fast enough for line speeds
  - Can use GEN2 or GEN3 in GEN4 adapters slots
- Warning: Not all adapters are supported at initial GA
  - See later lists in this deck



## POWER9 U.2 NVMe support for E950

E950 has four external storage bays

Independent of SAS controllers

NVMe ~price of a SAS adapter

Concurrent 8 SAS disks and 4 NVMe drives

SAS takes 1 or 2 PCIe slots

U.2 NVMe

- A maximum of four x U.2 NVMe drives
- Higher performance than SAS SSD
- Support concurrent maintenance (unlike Scale-Out S922/S924)
- Write endurance is 2.4 drive write per day DWPD (5 years)
- Intended primarily to store and boot OS (AIX / VIOS) images
- Each NVMe device is a separate PCIe endpoint = assign to different LPARs
  - On the HMC, it looks like each has own adapter slots
- NVMe drives may be assigned to the VIOS & virtualized to client OS
- Warranty: 5 years if not worn out. A “fuel gauge” to monitor wear is provided for AIX/Linux `nvmmemgr`



Different to Scale-Out  
→ M.2  
→ these have higher DWPD



25

IBM Power Systems POWER9 Enterprise E950 Unofficial Deep Dive

© IBM Corporation, 2018

## POWER9 No internal DVD support

This should not be surprising with 22 year old Tech!

- It is old, slow, hot, unreliable = dead = get over it!

Alternative is a USB Memory Key

1. Faster: USB 3.0 reads at **90 MB/s**
2. Larger: lowest GB per buck now is 32 GB USB
3. Memory Key is €\$£ ~10

If you must go DVD:

- Use external USB DVD or USB DVD-RAM **but at your own risk**
- IBM now offered a External USB DVD (at a stiff price but supported)
  - FC#EUA5 Standalone USB DVD drive w/cable €\$£122.82
- Use the front USB sockets (provide more electricity power for mechanical DVD drive)

POWER9 OS Install:

- AIX, VIOS & Linux all supported installing from USB memory key



Same as Scale-Out



26

IBM Power Systems POWER9 Enterprise E950 Unofficial Deep Dive

© IBM Corporation, 2018





# POWER9 The Server Range

27

IBM Power Systems POWER9 Enterprise E950 Unofficial Deep Dive

© IBM Corporation, 2018

## POWER8 range

8 - 192 Cores 18U  
256 GB – 32 TB RAM  
8 - 192 PCI Adapters

8 - 80 Cores 10U  
256 GB – 16 TB RAM  
8 - 96 PCI Adapters

16 - 48 Cores 4U  
128 GB – 4 TB RAM  
7 - 51 PCI Adapters

**E880**

**E870**

**E850**

**S824(L) – 4U**  
6 - 24 cores

**S822(L) – 2U**  
6 - 20 cores

**S812LC**  
10 core

**S822LC – 2U**  
8 - 20 core


**S821LC – 1U**  
**S822LC – 2U**  
8 - 20 core

28


IBM Power Systems POWER9 Enterprise E950 Unofficial Deep Dive

© IBM Corporation, 2018

# POWER9 range in 2018




**Enterprise Q3+Q4**





7U to 22U  
4 to 16 socket  
8 to 192 cores  
16TB RAM

**Midrange Q3**



4U  
2 or 4 socket  
16 to 48 cores  
16TB RAM



**Scale-Out Q2**


**S924**  **S922** 

2U or 4U  
1 or 2 socket  
4 to 24 cores  
4TB RAM

**AC922 Q4 2017**  **LC922/LC921 Q2** 

29 IBM Power Systems POWER9 Enterprise E950 Unofficial Deep Dive © IBM Corporation, 2018

# POWER9 Model Details



30 IBM Power Systems POWER9 Enterprise E950 Unofficial Deep Dive © IBM Corporation, 2018

# POWER9 E950



3 Year 24x7 Warranty

### Software Stack

- HMC 920+
- PowerVM 2.2.6.23+
- AIX 6 TL9 sp12, 7.1 TL5 sp4 & 7.2 TL3 (full P9 support)
- Linux
  - SLES 11 sp4, 12 sp3, 15
  - RHEL 7.4, 7.5
- IBM i not supported


|                                   | System (9040-MR9) – 4U   |
|-----------------------------------|--|
| POWER9 Fast/Fused Core Processors | Two or four processors → 16,20,22,24,32,40,44,48 CPU cores<br>8, 10, 11 or 12 CPU Cores per socket<br>Maximum 3.8 GHz  |
| Sockets                           | 2 or 4 Field Upgradable  |
| Memory                            | 8 Riser Cards each with 16 x DDR4 ISDIMMs = 128 total DIMM sizes 8, 16, 32, 64,128 GB<br>16 TB max Memory (four times the E850)  |
| Media Bays                        | DVD via external USB DVD or USB flash key  |
| Integrated PCIe - Full Height     | PCIe Gen 4: eight x16 + two x8 (2 processors = 4 x16)<br>PCIe Gen 3: one x8 (default Ethernet @ 2x10Gb + 2x1Gb)<br>PCIe slots are Full High, Half Length and Blind swap  |
| Internal SAS Bays                 | - Up to 8 SAS 2.5 inch, 15mm via 1 or 2 SAS adapters (x8)<br>Split disk capable JBOD, RAID 1,5,6 or 10<br>- U.2 format 1 to 4 NVMe slots for 1 to 4 NVMe devices<br>2.5 inch 7 mm 800GB, 1.6TB 3.2TB<br>- USB 3.0 (2 front and 2 rear) |
| Max I/O Drawers                   | 4 EMX I/O Drawer (PCIe Adapter drawer)   |
| External Storage Drawers          | EXP12SX, ESP24SX & EXP24S<br>Each requires: PCIe SAS adapter or SAS port   |

31

IBM Power Systems POWER9 Enterprise E950 Unofficial Deep Dive

© IBM Corporation, 2018

# POWER9 4 Socket Server E950




32

IBM Power Systems POWER9 Enterprise E950 Unofficial Deep Dive

© IBM Corporation, 2018

## POWER9 Enterprise E950 Server - (Front view w/o Bezel)




33

IBM Power Systems POWER9 Enterprise E950 Unofficial Deep Dive

© IBM Corporation, 2018

## POWER9 Enterprise E950 Server - (Rear view)



34

IBM Power Systems POWER9 Enterprise E950 Unofficial Deep Dive

© IBM Corporation, 2018

## Processors

### Initially

- 2 or 4 Processors =  
2 or 4 POWER9 chips
- Can 2 to 4 MES upgrade


35

IBM Power Systems POWER9 Enterprise E950 Unofficial Deep Dive

© IBM Corporation, 2018

## E950 Processor Architecture Highlights

### Four directly connected POWER9 Processors = 1 hop



36

IBM Power Systems POWER9 Enterprise E950 Unofficial Deep Dive

© IBM Corporation, 2018

## E950 Processor options

| CPU cores per POWER9 | CPU cores per 2 Proc Server | CPU cores per 4 Proc Server | Comment           |
|----------------------|-----------------------------|-----------------------------|-------------------|
| 12                   | 24                          | 48                          | Max throughput    |
| 11                   | 22                          | 44                          |                   |
| 10                   | 20                          | 40                          |                   |
| 8                    | 16                          | 32                          | Max core strength |

Much higher performance  
Full 16 TB memory available

These CPU cores counts are available on the S924

E950 makes sense if you:

- a) plan to upgrade later or
- b) need 8 TB memory, S924 max is 4TB

37

IBM Power Systems POWER9 Enterprise E950 Unofficial Deep Dive

© IBM Corporation, 2018

## E950 Processor options

| CPU cores per POWER9 | CPU cores per 2 Proc Server | CPU cores per 4 Proc Server | rPerf | rPerf / core |
|----------------------|-----------------------------|-----------------------------|-------|--------------|
| 12                   | 24                          | 48                          | 1146  | 23.8         |
| 11                   | 22                          | 44                          | 1072  | 24.3         |
| 10                   | 20                          | 40                          | 1034  | 25.9         |
| 8                    | 16                          | 32                          | 870   | 27.2         |

Much higher performance  
Full 16 TB memory available

These CPU cores counts are available on the S924

E950 makes sense if you:



- a) plan to upgrade later or
- b) need 8 TB memory, S924 max is 4TB

POWER9 Performance Report <https://www-01.ibm.com/common/ssi/cgi-bin/ssialias?htmlfid=POO03017USEN&>

38

IBM Power Systems POWER9 Enterprise E950 Unofficial Deep Dive

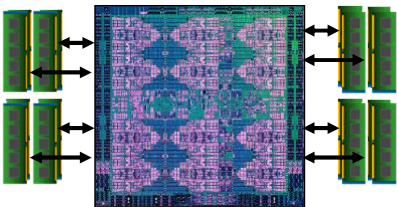
© IBM Corporation, 2018



## POWER9 Processor Family

~~Two~~ **Three** Memory Architectures

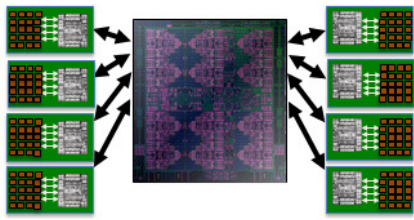
**Scale Out**  
Direct Attach Memory



8 Direct DDR4 Ports

**Max 2-Socket Systems**

**Scale Up**  
Buffered Memory



8 Buffered Channels

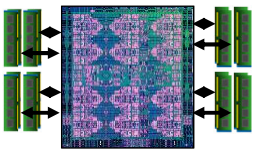
**4 to 16 Socket Systems**

41 IBM Power Systems POWER9 Enterprise E950 Unofficial Deep Dive © IBM Corporation, 2018

## POWER9 Processor Family

### Three Memory Architectures

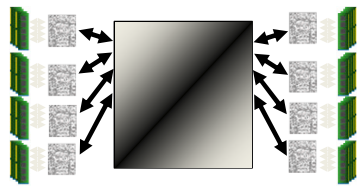
**Scale Out**  
Direct Attach Memory



8 Direct DDR4 Ports  
POWER9 → DDR4  
Lower cost

**Max 2-Socket Systems**

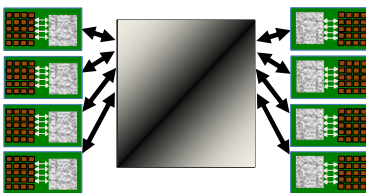
**Scale Up**  
Buffered Memory



8 Buffered Channels  
POWER9 → Centaur chip → DDR4  
Best of both!

**E950 4 Socket Systems**

**Scale Up**  
Buffered Memory




8 Buffered Channels  
POWER9 → Centaur chip → RAM  
Higher performance & higher RAS

**E980 16 Socket Systems**

42 IBM Power Systems POWER9 Enterprise E950 Unofficial Deep Dive © IBM Corporation, 2018





### E950 Memory Subsystem Highlights

DDR4 DIMMS sizes

- ✓ 8, 16, 32, 64, 128GB

Per E950 Server


- ✓ Max 920 GB/s Memory Bandwidth
- ✓ Min 2 Socket: 64 GB Memory
- ✓ Min 4 Socket: 128 GB Memory
- ✓ Max of 16TB Memory
- ✓ 8 Riser Cards with 128 DIMMs slots
- ✓ Minimum of 50% Memory Activated

Per POWER9 processor

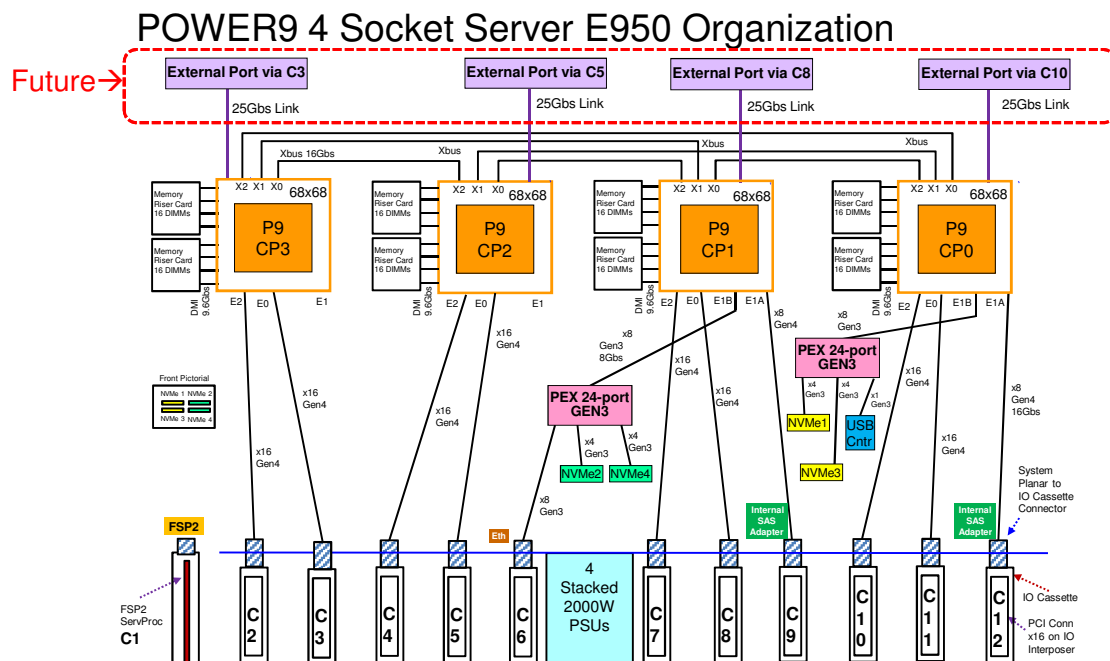
- ✓ Max 230 GB/s Memory Bandwidth
- ✓ Up to 4TB Memory
- ✓ 2 Memory Riser cards (like below) each with
  - ✓ 4 Centaur chips
  - ✓ 16 DIMM slots

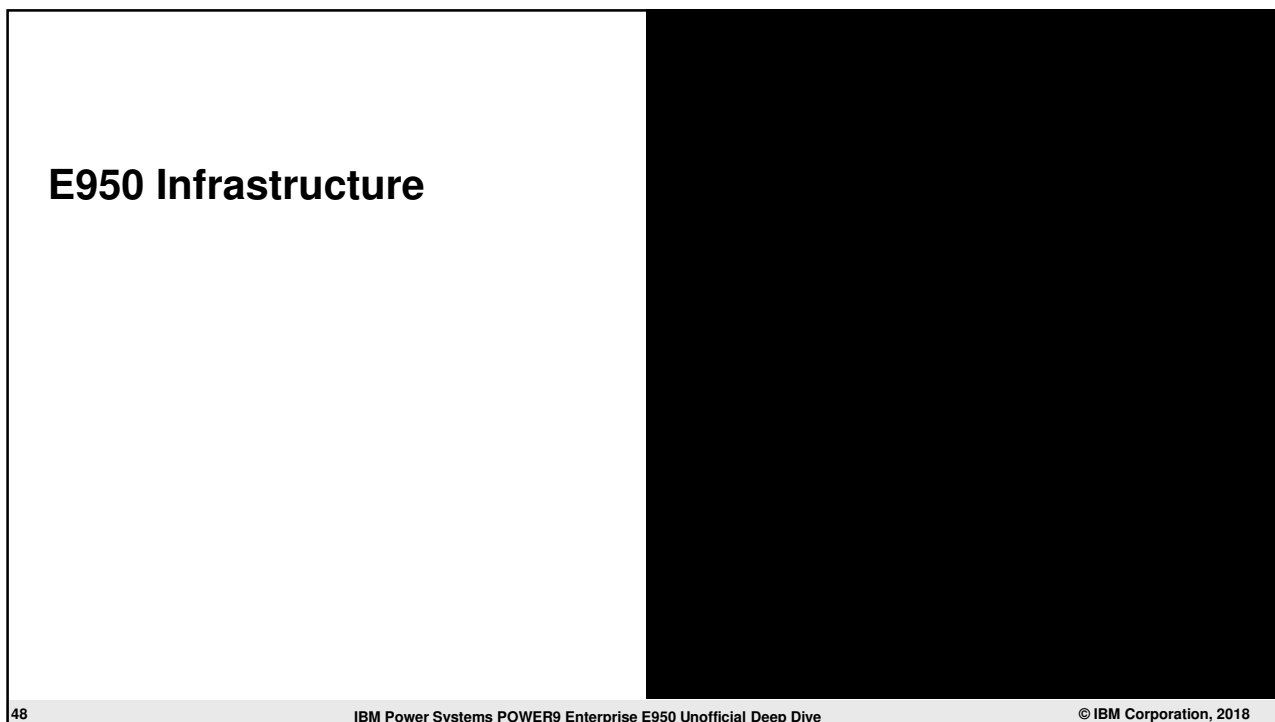
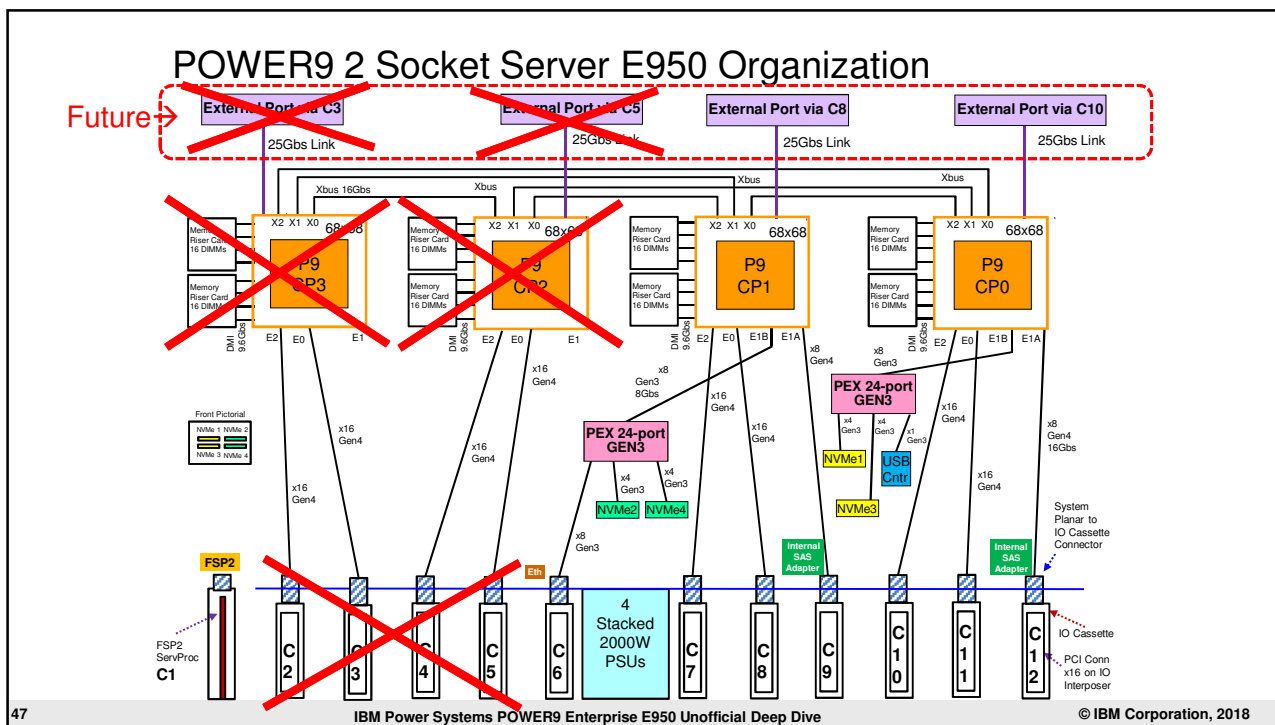
Memory increment

- ✓ 8 DIMMs



# E950 I/O Subsystem





## Key Power E950 technical enhancements over Power E850C

| Features                     | E850C                                      | E950   | Enhancements   |
|------------------------------|--|--|--|
| Processor                    | P8 (DCM)                                   | P9 (SCM)   | ✓ <b>More efficient processor communication</b>                |
| Sockets                      | 2 to 4                                     | 2 to 4   |  |
| Cores                        | 32 or 40 or 48                             | 32 or 40 or 44 or 48                                     | ✓ <b>More processor options</b>                                |
| Maximum Memory               | 4 TB                                       | 16 TB  | ✓ <b>4X Maximum Memory</b>                                     |
| DIMM Type/DIMM slots count   | Up to 32 CDIMMs                            | Up to 128 ISDIMMs  | ✓ <b>More Memory Capacity and Flexibility</b>                  |
| Memory Bandwidth             | 768 GB/sec                                 | 920 GB/sec   | ✓ <b>20% more Memory Bandwidth</b>                             |
| IO Expansion Slots           | Yes  | Yes  |  |
| PCIe slots                   | 11   | 11   | ✓ <b>8 Gen4 x16 + 2 Gen4 x8 + 1 Gen3</b>                       |
| Acceleration Ports           | Yes (CAPI 1.0)                             | Yes (CAPI 2.0 + OpenCAPI)                                | ✓ <b>Enabled for enhanced acceleration for AI environments</b> |
| PCIe Hot Plug Support        | Yes  | Yes + Blindswap  |  |
| IO bandwidth                 | 315 GB/sec                                 | 630 GB/sec   | ✓ <b>~2X IO bandwidth</b>                                      |
| Ethernet ports               | Quad 1 Gbt (x8 Slot)                       | Quad 1 Gbt (x8 Slot)                                     |  |
| Internal Storage Bays        | 12 (8 x 2.5" SAS bays + 4 x 1.8" SSD bays) | 12 (8 x 2.5" SAS bays + 4 x NVMe bays)                   | ✓ <b>NVMe for high-bandwidth, low latency local boot</b>       |
| Internal Storage Controllers | Integrated                                 | Optional & Conc Maintainable                             | ✓ <b>Improved cost and seviceability</b>                       |
| RAS                          |  | Enhanced DC-DC Reg Redundancy<br>Enhanced Fan Conc Maint | ✓ <b>More reliable power/cooling</b>                           |

49

IBM Power Systems POWER9 Enterprise E950 Unofficial Deep Dive

© IBM Corporation, 2018

## E950 follows the Enterprise Server package traditions

Included with the Server

- **PowerVM** at no cost\*
  - Includes Hypervisor, VIOS, Firmware and HMC\*\* support for server management, virtualization and RAS
- **PowerVC** at no cost\*
- **Power to Cloud Rewards** (Education or Lab Services days)
- **Cloud Management Console** (Server status reports on you mobile/cell phone or tablet)
- **Warranty**: 3 years of 24x7 service included

Regular & popular Power Server options with a cost:

- Elastic Capacity on Demand(CoD) for both processor cores & memory
- Enterprise Pool Capacity (n-1 so the pool can be POWER8 and POWER9)
- IBM Active Memory Expansion for AIX
- IBM Active Memory Mirroring for Hypervisor

\* There is a SWMA cost

\*\* HMC is orderable separately at a cost


50

IBM Power Systems POWER9 Enterprise E950 Unofficial Deep Dive

© IBM Corporation, 2018

# POWER9

## eConfig view of the Features



51
IBM Power Systems POWER9 Enterprise E950 Unofficial Deep Dive
© IBM Corporation, 2018

### E950 Processor GHz

| Feature code | CPU cores per POWER9 | CPU cores per Server (max) | Nominal GHz | No-name GHz | Max GHz |
|--------------|----------------------|----------------------------|-------------|-------------|---------|
| EPWT         | 12                   | 48                         | 2.8         | 3.15        | 3.8     |
| EPWY         | 11                   | 44                         | 2.85        | 3.28        | 3.8     |
| EPWS         | 10                   | 40                         | 3.0         | 3.4         | 3.8     |
| EPWR         | 8                    | 32                         | 3.3         | 3.6         | 3.8     |

Number of Processor sockets filled

- 1 – not allowed
- 2 – minimum
- 3 – possible future offering
- 4 – maximum

↑

GHz reported by some commands

↑

Default and GHz in Practice  
AIX: lparstat -E 1 3

52
IBM Power Systems POWER9 Enterprise E950 Unofficial Deep Dive
© IBM Corporation, 2018

## POWER9 Energy & CPU GHz Balance

Enables higher dynamic operational frequencies

**Important Frequencies**

- **High** – Overclocking (~15%)
- **Medium** – Overclocking (~10%)
- **Nominal** – Fixed normal GHz
- **Power Saver** – Fixed reduce GHz
  - Reduces electrical power use = saves money
- **Zero GHz** – The server is powered off!

★ **Static Power Saver Mode**  
– Reduced electrical use to lower costs – lower GHz

★ **Disabled All Mode** → “overclocking” disabled  
– Fixed Frequency

🏠 **Dynamic Performance Mode**  
– Variable Frequency (Nominal to High GHz) based in workload  
– Higher the workload, the lower the GHz

🔴 **Maximum Performance Mode** – E950 default  
– Variable Frequency (Medium to High GHz) based in workload  
– If necessary, speeds up fans  
– In a hot 27+C computer room, can lower GHz to Nominal GHz

53 IBM Power Systems POWER9 Enterprise E950 Unofficial Deep Dive © IBM Corporation, 2018

### On Twitter?

@mr\_nmon

```

# lparstat -E 1 3
System configuration: type=Shared mode=Uncapped smt=8 lcpu=16 mem=16384MB ent=1.00
Physical Processor Utilisation:
-----Actual-----
user  sys  wait  idle   freq  %usr  %sys  %wait  %idle
1.126 0.016 0.000 0.000 3.7GHz [113%] 1.272 0.018 0.000 0.000
1.127 0.016 0.000 0.000 3.7GHz [113%] 1.274 0.018 0.000 0.000
1.124 0.016 0.000 0.000 3.7GHz [113%] 1.271 0.018 0.000 0.000
#
# lsattr -El proc0
frequency 3300000000 Processor Speed False
smt_enabled true Processor SMT enabled False
smt_threads 8 Processor SMT threads False
state enable Processor state False
type PowerPC_POWER9 Processor type False
#
                    
```

3.7 GHz = 13% Over-clocking (oops!)

3.3 billion Hz = 3.3 GHz Nominal

P9 GHz **part 1**: #POWER9 servers in practice run at (max) ~3.7-4 GHz, other server chips eat our dust! I see: normal GHz+overclocking, I am told to not use the "o" word, oops! #EnergyScale guys say run full speed but will lower GHz, if getting hot like your air-conditioning fails!

P9 GHz **part 2**: #POWER9 servers **How to get too hot!** If you don't have: max CPU count+ max memory size+ max disks+ max high-speed adapter AND max server workload+ computer room is warm then your server may never get too hot and still be at that (max) GHz. I know as I tried!!!

P9 GHz **part 3**: One quirk on AIX: commands like lparstat -E 1 9 report the varying current GHz but others report the non-overclocking (oops!) GHz value called Nominal So don't worry is you buy 3.9 GHz but nmon or lsattr -El proc0 reports a lower Nominal GHz between 2.3 to 3.3 GHz

54 IBM Power Systems POWER9 Enterprise E950 Unofficial Deep Dive © IBM Corporation, 2018



## Memory Options

| Feature Code | DIMM Size  |
|--------------|------------|
| EM6A         | 8GB DIMM   |
| EM6B         | 16GB DIMM  |
| EM6C         | 32GB DIMM  |
| EM6D         | 64GB DIMM  |
| EM6E         | 128GB DIMM |

- ✓ Up to 230 GB/s peak memory bandwidth per socket
- ✓ Up to 920 GB/s peak memory bandwidth per Server
- ✓ Memory Riser card: 16 DDR4 DIMMs & imbedded 4 Centaur L4 Cache chips
- ✓ Max Memory Riser card: 8
- ✓ Maximum 2 TB per Riser so a total of 16 TB
- ✓ Industry Standard DDR4 memory Registered DIMMs (server class)
- ✓ DDR4 ISDIMMs rated 2400MHz but running at 1600MHz behind Centaurs
- ✓ Only IBM DIMMs are supported
- ✓ Minimum config is 1 Riser Card with 8 DIMMs x 8 GB per Processor (64 GB / processor)
- ✓ DIMM plug rules are complex
  - ✓ A single Riser card recommended all memory the same size
  - ✓ A single Riser card can have 2 DIMM sizes (8 one size + 8 another size)

55

IBM Power Systems POWER9 Enterprise E950 Unofficial Deep Dive

© IBM Corporation, 2018

## Memory rules min and max for 2 or 4 procs

- All memory is DDR4 - No POWER8 E850 RAM is supported
- All memory is the same speed
- Only difference is the DIMM size
- Each Processor socket has 2 Riser cards, each with 16 DIMM slots
- 2 Processor Socket: min 64 GB\* & max 8 TB Memory
- 4 Processor Socket: min 128 GB\* & max 16 TB Memory
- Add DIMMs in packs of eight = half a Riser Card
- More DIMMs = more performance

- \* Poor performance as  $\frac{3}{4}$  of the ports are not used.
- Recommend not going below 256 GB and 512GB

56

IBM Power Systems POWER9 Enterprise E950 Unofficial Deep Dive

© IBM Corporation, 2018

## NVMe storage Options



| Feature Code | DIMM Size |
|--------------|-----------|
| EC5J         | 800 GB    |
| EC5K         | 1.6TB     |
| EC5L         | 3.2TB     |



- ✓ Three sizes of internal NVMe boot devices
- ✓ Ideal for VIOS boot or OS boot devices (two for redundancy)
- ✓ Faster than SSD - skipping SAS protocol level
- ✓ 2.4 drive write per day DWPD (5 years warranty)
  - ✓ Meaning complete Drive Writes Per Day
- ✓ “nvmmgr” command to determine drive wear
- ✓ They will issue warning messages at 100% used
  - ✓ Customer to backup and replace

57

IBM Power Systems POWER9 Enterprise E950 Unofficial Deep Dive

© IBM Corporation, 2018

## Three Backplane options



Depends on the SAS internal disk support that you need

- All have four NMVe bays

One of these SAS options must be selected:

1. Zero DASD Backplane no disks (only NMVe), so no SAS adapter  
–FC#EJ0B
2. Base DASD backplane with 1 SAS PCIe adapter + selected SAS drives  
–FC#EJBB - all disks connect to the 1 adapter
3. Split DASD backplane with 2 SAS PCIe adapters + selected SAS drive  
–FC#EJ0J - disks are split between the two SAS adapters 4 + 4

58

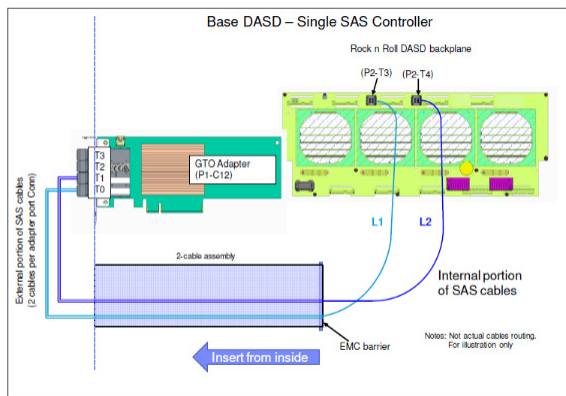
IBM Power Systems POWER9 Enterprise E950 Unofficial Deep Dive

© IBM Corporation, 2018

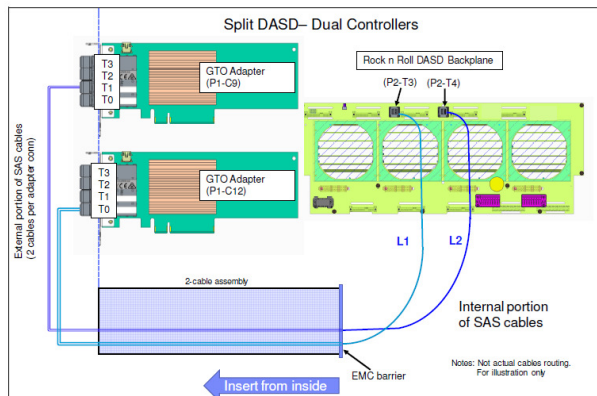


## Internal SAS cabling

All 8 disks on one SAS EJ0K adapter

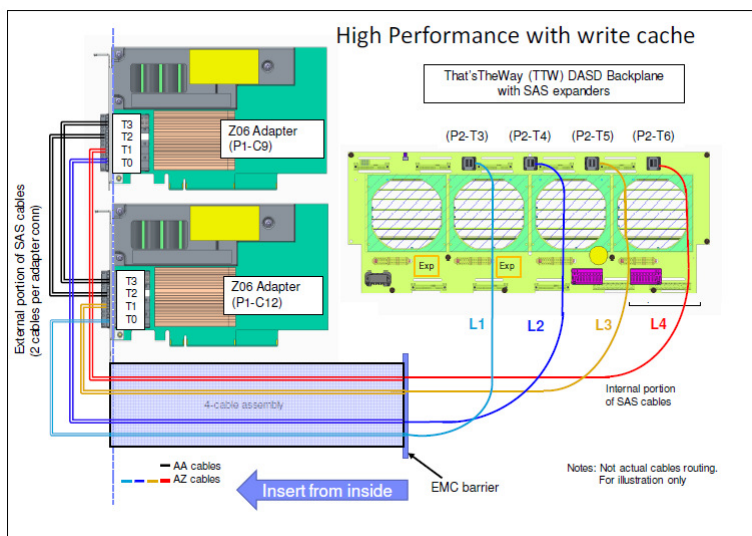


Split Disk 4 + 4 with EJ0K SAS adapters



## Internal SAS cabling

High Performance with large shared adapter disk cache – 2 x EJ14 adapters





## Dual VIOS suggestions

### Network

- 2 or 4 (for redundancy) Ethernet adapters

### Storage

#### A. 1 to 4 NVMe drives

- Fast (possible AIX flash Cache use too)
- No system downtime to replace NVMe drives

#### B. 2 or 4 (for redundancy) Fibre Channel SAN

- SAN booting VIOS is standard these days & fast. Assumes you have a SAN

#### C. Internal disks or SSD via two SAS adapter (Split backplane)

- 8 disk at the front (redundancy using AIX mirroring)
- If not SSD its pretty slow I/O
- Requires 2 SAS adapters & uses 2 PCIe slots

#### D. Remote Disk Drawer

61

## Internal storage numbering

### SAS disks

- 1 2 3 4 5
- 6 7 8


–SAS adapter 1 + 2

### NVMe

- 1 2
- 3 4

–VIOA

–VIOB



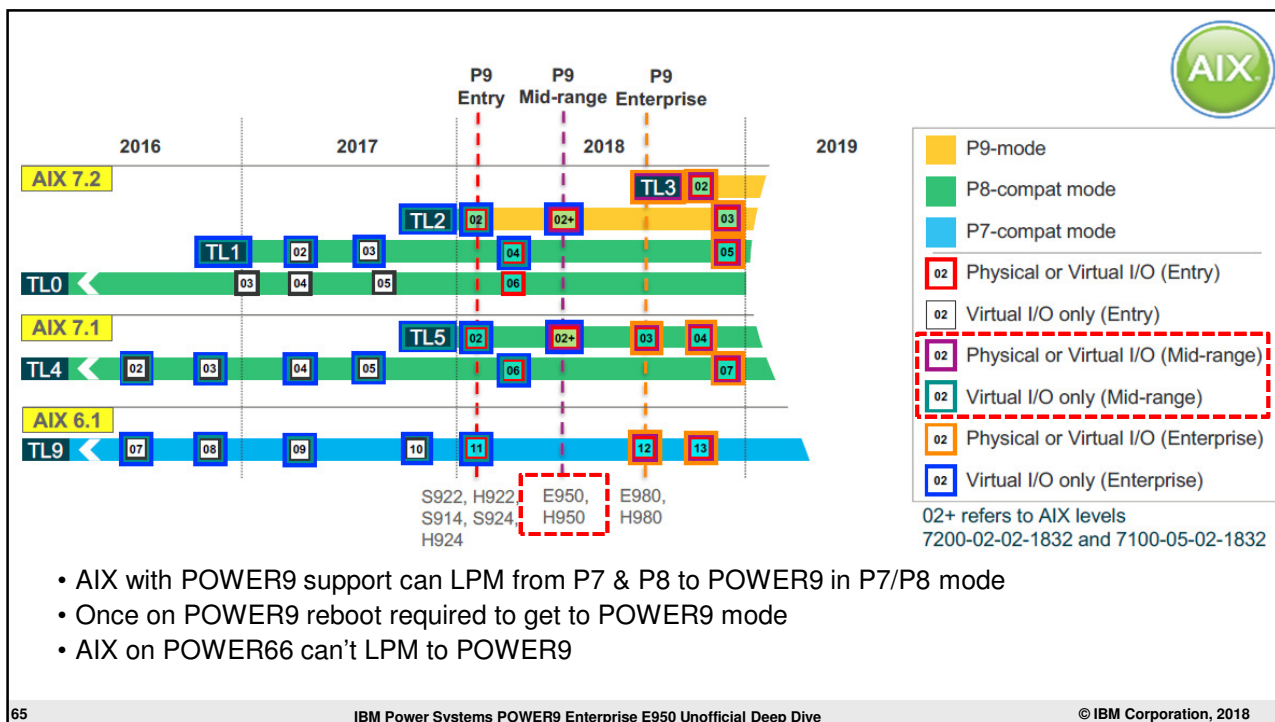
62



# Operating Systems for Enterprise

# AIX





## AIX for new POWER9 Generally

*Reference Chart not E950 specific*

### AIX 7.2 TL2 + SP2 (Mar 2018)

- Random number generator
  - User-mode accessible
  - For application exploitation
- Power 3.0B ISA extensions
  - VSX3, string, video encode, quad floating point, pc relative addressing, 32-bit overflow, Memory Atomics
  - For application exploitation
- NVMe support

### AIX 7.2 TL3 (2H 2018)

- Dynamic System Optimizer
  - Enabled for P9 systems
- 1536-thread single LPAR
  - =192-core/SMT8 single system image support
- 32 TB max RAM in an LPAR

## AIX Level Support at **E950** **RELEASE-TL-SP-YYWW**

### Virtualised and LPM

- AIX 7200-02-01-1732 or later
- AIX 7200-01-01-1642 or later
- AIX 7100-05-01-1731 or later
- AIX 7100-04-02-1614 or later
- AIX 6100-09-07-1614 or later\*

### AIX with Adapters at 17<sup>th</sup> August

- AIX 7200-02-02-1832 or later
- AIX 7100-05-02-1832 or later

### AIX with Adapters later in 2018

- AIX 7200-01-05-1837 Dec 2018
- AIX 7100-04-07-1837 Dec 2018
- AIX 6100-09-12-1838 Sept 2018\*

\* AIX 6 service extension required

67

IBM Power Systems POWER9 Enterprise E950 Unofficial Deep Dive

© IBM Corporation, 2018

# Linux



68

IBM Power Systems POWER9 Enterprise E950 Unofficial Deep Dive

© IBM Corporation, 2018

## Linux on POWER9 E950



- SUSE SLES
- Red Hat RHEL

- Linux on E950 only available under PowerVM
- Older Linux version run fine in POWER8 mode
  - SLES 11 sp4 and 12 sp3
  - RHEL 7.4
- Refreshed releases in 2018 have some POWER9 support – ask the vendor
- For POWER9 optimisation best chance is:
  - SLES 15
  - RHEL 7.5 for POWER9 also know as the ALT version  
like RHEL-ALT-7.5-20180315.0-Server-ppc64le-dvd1.iso

69

IBM Power Systems POWER9 Enterprise E950 Unofficial Deep Dive

© IBM Corporation, 2018

## Linux on POWER9 E950



- Canonical Ubuntu

- Ubuntu on PowerVM no longer being supported by IBM
  - Includes 16.04 and 18.04
  - Include POWER8 and POWER9
  - Does actually work – just no support
- If you previously purchased support **Ubuntu 16.04 in POWER8 mode on PowerVM**, then support continues and is available for 2 more years
- **Native Ubuntu** on base metal & for KVM hosting is **fully supported by IBM**
  - Native meaning “not PowerVM” environment (also known as OPAL mode)
  - For example: POWER9 Servers **AC922 & LC922** and **LC921**

70

IBM Power Systems POWER9 Enterprise E950 Unofficial Deep Dive

© IBM Corporation, 2018

# IBM i



IBM i is not supported on the E950

# VIOS

2.2.6.23 or the latest on 17<sup>th</sup> August 2018

Do not use older versions due to new E950 devices like NVMe

VIOS 3.1 is coming soon (based on AIX 7.2)

# Storage Remote Drawers Adapter Remote Drawers

Same drawers & connections as POWER9 Scale-Out Servers


- this should all be business as usual
- also used in POWER8

73

IBM Power Systems POWER9 Enterprise E950 Unofficial Deep Dive

© IBM Corporation, 2018

## Remote Disk drawers for Disks



- SAS Disk Drawers EXP12SX | EXP24SX | EXP 24S
  - FC#ESLL | FC#ESLS | FC#5887 (slider, slider & homerun)
  - Supported via PCIe SAS adapter
    - FC#EJ0J (full height) & FC#EJ0M (low profile) (GTO)
  - As used on POWER8
- Care needed if the old “migrating” drawer is back level

EXP12SX = 12x 3.5 inch disks

EXP24SX = 24 x 2.5 inch disks



74

IBM Power Systems POWER9 Enterprise E950 Unofficial Deep Dive


© IBM Corporation, 2018



## Remote Adapter drawers - FC#EMX (know as MEX drawers)

- Four Processor E950 supports
  - 8 FC#EJ08 cards (Bear Mountain) in x16 PCIe Slots
  - 8 Fan-out drawer halves = Four #EMX Drawers
  - Each Fan-out has 6 adapters slots (12 per drawer)

- Eight Adapter slots → 48 Adapter slots plus 3 slower E950 slots = total 51 !



75

IBM Power Systems POWER9 Enterprise E950 Unofficial Deep Dive

© IBM Corporation, 2018

# I/O support for 17<sup>th</sup> August release of E950 Adapters, disks & related bits

Important for early migration & quick adoption

76

IBM Power Systems POWER9 Enterprise E950 Unofficial Deep Dive

© IBM Corporation, 2018

| POWER9 I/O currently planned for GA1 on Power E950 |      |   |     |            |  | E950<br>FC SAN<br>Drawer SAS |
|--|------|---|-----|------------|--|------------------------------|
| FC   | CCIN | Description   | Max | OS support |  |                              |
| 5729   | 5729 | PCIe2 8Gb 4-port Fibre Channel Adapter                  | 50  | AIX Linux  |  |                              |
| 5735   | 577D | 8 Gigabit PCI Express Dual Port Fibre Channel Adapter   | 50  | AIX Linux  |  |                              |
| EN0A   | 577F | PCIe3 16Gb 2-port Fibre Channel Adapter                 | 50  | Linux      |  |                              |
| EN12   |      | PCIe2 8Gb 4-port Fibre Channel Adapter                  | 50  | AIX Linux  |  |                              |
| EN1A   | 578F | PCIe3 32Gb 2-port Fibre Channel Adapter                 | 50  | Linux      |  |                              |
| EN1C   | 578E | PCIe3 16Gb 4-port Fibre Channel Adapter                 | 50  | Linux      |  |                              |
| FC   | CCIN | Description   | Max | OS support |  |                              |
| EJ0J   | 57B4 | PCIe3 RAID SAS Adapter Quad-port 6Gb x8                 | 34  | AIX Linux  |  |                              |
| EJ0K   | 57BF | PCIe3 RAID SAS Adapter Quad-port 6Gb x8 for MR9         | 34  |            |  |                              |
| EJ10   | 57B4 | PCIe3 SAS Tape/DVD Adapter Quad-port 6Gb x8             | 34  | AIX Linux  |  |                              |
| EJ14   | 57B1 | PCIe3 12GB Cache RAID PLUS SAS Adapter Quad-port 6Gb x8 | 34  | AIX Linux  |  |                              |
| FC   | CCIN | Description   | Max | OS support |  |                              |
| ESLL   | 78D1 | EXP12SX SAS Storage Enclosure                           | 136 | AIX Linux  |  |                              |
| ESLS   | 78D1 | EXP24SX SAS Storage Enclosure                           | 136 | AIX Linux  |  |                              |

77

| POWER9 I/O currently planned for August 17 <sup>th</sup> on Power E950 |      |   |     |            |  | E950<br>Disk |
|--|------|---|-----|------------|--|--------------|
| FC   | CCIN | Description   | Max | OS support |  |              |
| ES78   | 5B16 | 387GB SFF-2 SSD 5xx eMLC4 for AIX/Linux             | 768 | AIX Linux  |  |              |
| ES7E   | 5B17 | 775GB SFF-2 SSD 5xx eMLC4 for AIX/Linux             | 768 | AIX Linux  |  |              |
| ES7K   | 5B19 | 387GB SFF-3 SSD 5xx eMLC4 for AIX/Linux             | 8   | AIX Linux  |  |              |
| ES7P   | 5B1A | 775GB SFF-3 SSD 5xx eMLC4 for AIX/Linux             | 8   | AIX Linux  |  |              |
| ES80   | 5B21 | 1.9TB Read Intensive SAS 4k SFF-2 SSD for AIX/Linux | 768 | AIX Linux  |  |              |
| ES83   | 5B2B | 931GB Mainstream SAS 4k SFF-3 SSD for AIX/Linux     | 768 |            |  |              |
| ES85   | 5B10 | 387GB SFF-2 SSD 4k eMLC4 for AIX/Linux              | 768 | AIX Linux  |  |              |
| ES8C   | 5B11 | 775GB SFF-2 SSD 4k eMLC4 for AIX/Linux              | 768 | AIX Linux  |  |              |
| ES8F   | 5B12 | 1.55TB SFF-2 SSD 4k eMLC4 for AIX/Linux             | 768 | AIX Linux  |  |              |
| ES8J   | 5B20 | 1.9TB Read Intensive SAS 4k SFF-3 SSD for AIX/Linux | 8   | AIX Linux  |  |              |
| ES8N   | 5B13 | 387GB SFF-3 SSD 4k eMLC4 for AIX/Linux              | 8   | AIX Linux  |  |              |
| ES8Q   | 5B14 | 775GB SFF-3 SSD 4k eMLC4 for AIX/Linux              | 8   | AIX Linux  |  |              |
| ES8V   | 5B15 | 1.55TB SFF-3 SSD 4k eMLC4 for AIX/Linux             | 8   | AIX Linux  |  |              |
| ES8Y   | 5B29 | 931GB Mainstream SAS 4k SFF-2 SSD for AIX/Linux     | 768 |            |  |              |
| ES92   | 5B20 | 1.86TB Mainstream SAS 4k SFF-3 SSD for AIX/Linux    | 8   |            |  |              |
| ES96   | 5B21 | 1.86TB Mainstream SAS 4k SFF-2 SSD for AIX/Linux    | 768 |            |  |              |
| ESE1   | 5B2C | 3.72TB Mainstream SAS 4k SFF-3 SSD for AIX/Linux    | 8   |            |  |              |
| ESE7   | 5B2D | 3.72TB Mainstream SAS 4k SFF-2 SSD for AIX/Linux    | 768 |            |  |              |

8 SFF-3 disk drive slots at the front of the E950

Lime colour found in eConfig has "Withdrawn" could be migrated from POWER8 & supported

78

POWER9 I/O currently planned for August 17<sup>th</sup> on Power E950**E950**

| FC   | CCIN | Description                                       | Max | OS support |
|------|------|---|-----|------------|
| ESG5 | 5B16 | 387GB Enterprise SAS 5xx SFF-2 SSD for AIX/Linux  | 768 |            |
| ESG9 | 5B19 | 387GB Enterprise SAS 5xx SFF-3 SSD for AIX/Linux  | 8   |            |
| ESGB | 5B10 | 387GB Enterprise SAS 4k SFF-2 SSD for AIX/Linux   | 768 |            |
| ESGD | 5B13 | 387GB Enterprise SAS 4k SFF-3 SSD for AIX/Linux   | 8   |            |
| ESGF | 5B17 | 775GB Enterprise SAS 5xx SFF-2 SSD for AIX/Linux  | 768 |            |
| ESGH | 5B1A | 775GB Enterprise SAS 5xx SFF-3 SSD for AIX/Linux  | 8   |            |
| ESGK | 5B11 | 775GB Enterprise SAS 4k SFF-2 SSD for AIX/Linux   | 768 |            |
| ESGM | 5B14 | 775GB Enterprise SAS 4k SFF-3 SSD for AIX/Linux   | 8   |            |
| ESGP | 5B12 | 1.55TB Enterprise SAS 4k SFF-2 SSD for AIX/Linux  | 768 |            |
| ESGR | 5B15 | 1.55TB Enterprise SAS 4k SFF-3 SSD for AIX/Linux  | 8   |            |
| ESHJ | 5B29 | 931 GB Mainstream SAS 4k SFF-2 SSD for AIX/Linux  | 768 |            |
| ESHL | 5B21 | 1.86 TB Mainstream SAS 4k SFF-2 SSD for AIX/Linux | 768 |            |
| ESHN | 5B2F | 7.44 TB Mainstream SAS 4k SFF-2 SSD for AIX/Linux | 768 |            |
| ESHS | 5B2B | 931 GB Mainstream SAS 4k SFF-3 SSD for AIX/Linux  | 8   |            |
| ESHU | 5B20 | 1.86 TB Mainstream SAS 4k SFF-3 SSD for AIX/Linux | 8   |            |
| ESHW | 5B2E | 7.44 TB Mainstream SAS 4k SFF-3 SSD for AIX/Linux | 8   |            |
| ESM8 | 5B2D | 3.72 TB Mainstream SAS 4k SFF-2 SSD for AIX/Linux | 768 |            |
| ESMQ | 5B2C | 3.72 TB Mainstream SAS 4k SFF-3 SSD for AIX/Linux | 8   |            |

**Disk**

79

IBM Power Systems POWER9 Enterprise E950 Unofficial Deep Dive

© IBM Corporation, 2018

POWER9 I/O currently planned for August 17<sup>th</sup> on Power E950**E950**

| FC   | CCIN | Description   | Max  | OS support |
|------|------|---|------|------------|
| 1953 | 19B1 | 300GB 15k RPM SAS SFF-2 Disk Drive (AIX/Linux)                    | 1536 | AIX Linux  |
| 1964 | 19B3 | 600GB 10k RPM SAS SFF-2 Disk Drive (AIX/Linux)                    | 1536 | AIX Linux  |
| ES62 | 5B1D | 3.86-4.0 TB 7200 RPM 4K SAS LFF-1 Nearline Disk Drive (AIX/Linux) | 768  |            |
| ES64 | 5B1F | 7.72-8.0 TB 7200 RPM 4K SAS LFF-1 Nearline Disk Drive (AIX/Linux) | 768  |            |
| ESD5 | 59D0 | 600GB 10K RPM SAS SFF-3 Disk Drive (AIX/Linux)                    | 8    | AIX Linux  |
| ESDB | 59E0 | 300GB 15K RPM SAS SFF-3 Disk Drive (AIX/Linux)                    | 8    | AIX Linux  |
| ESEV | 59D2 | 600GB 10K RPM SAS SFF-2 Disk Drive 4K Block - 4096                | 1536 | AIX Linux  |
| ESEZ | 59C9 | 300GB 15K RPM SAS SFF-2 4K Block - 4096 Disk Drive                | 1536 | AIX Linux  |
| ESF3 | 59DA | 1.2TB 10K RPM SAS SFF-2 Disk Drive 4K Block - 4096                | 1536 | AIX Linux  |
| ESF5 | 59D3 | 600GB 10K RPM SAS SFF-3 Disk Drive 4K Block - 4096                | 8    | AIX Linux  |
| ESF9 | 59DB | 1.2TB 10K RPM SAS SFF-3 Disk Drive 4K Block - 4096                | 8    | AIX Linux  |
| ESFB | 59E1 | 300GB 15K RPM SAS SFF-3 4K Block - 4096 Disk Drive                | 8    | AIX Linux  |
| ESFF | 59E5 | 600GB 15K RPM SAS SFF-3 4K Block - 4096 Disk Drive                | 8    | AIX Linux  |
| ESFP | 59CC | 600GB 15K RPM SAS SFF-2 4K Block - 4096 Disk Drive                | 1536 | AIX Linux  |
| ESFT | 59DD | 1.8TB 10K RPM SAS SFF-2 Disk Drive 4K Block - 4096                | 1536 | AIX Linux  |
| ESFV | 59DE | 1.8TB 10K RPM SAS SFF-3 Disk Drive 4K Block - 4096                | 8    | AIX Linux  |
| ESNK | 5B41 | 300GB 15K RPM SAS SFF-3 4k Block Cached Disk Drive (AIX/Linux)    | 8    |            |
| ESNP | 5B45 | 600GB 15K RPM SAS SFF-3 4k Block Cached Disk Drive (AIX/Linux)    | 8    |            |

**Disk**

80

IBM Power Systems POWER9 Enterprise E950 Unofficial Deep Dive

© IBM Corporation, 2018

POWER9 I/O currently planned for August 17<sup>th</sup> on Power E950

| FC   | CCIN | Description                                      | Max | OS support |
|------|------|--|-----|------------|
| 5899 | 576F | PCIe2 4-port 1GbE Adapter                        | 50  | AIX Linux  |
| EC66 | 2CF3 | PCIe4 2-port 100Gb ROCE EN CAPI adapter          | 10  |            |
| EN0H | 2B93 | PCIe3 4-port (10Gb FCoE & 1GbE) SR&RJ45          | 51  | AIX Linux  |
| EN0K | 2CC1 | PCIe3 4-port (10Gb FCoE & 1GbE) SFP+Copper&RJ45  | 51  | AIX Linux  |
| EN0S | 2CC3 | PCIe2 4-Port (10Gb+1GbE) SR+RJ45 Adapter         | 51  | AIX Linux  |
| EN0U | 2CC3 | PCIe2 4-port (10Gb+1GbE) Copper SFP+RJ45 Adapter | 51  | AIX Linux  |
| EN0W | 2CC4 | PCIe2 2-port 10/1GbE BaseT RJ45 Adapter          | 51  | AIX Linux  |
| EN15 | 2CE3 | PCIe3 4-port 10GbE SR Adapter                    | 51  | AIX Linux  |

E950

Network

**NOTE**

The adapters with FCoE in the description are supported for Ethernet usage but you require an RPQ if you intend to use FCoE in POWER9

81

IBM Power Systems POWER9 Enterprise E950 Unofficial Deep Dive

© IBM Corporation, 2018

POWER9 I/O currently planned for August 17<sup>th</sup> on Power E950

| FC   | CCIN | Description  | Max | OS support |
|------|------|--|-----|------------|
| EM6A |      | 8 GB DDR4 Memory (2666 MHz)                            | 128 |            |
| EM6B |      | 16 GB DDR4 Memory (2666 MHz)                           | 128 |            |
| EM6C |      | 32 GB DDR4 Memory (2666 MHz)                           | 128 |            |
| EM6D |      | 64 GB DDR4 Memory (2666 MHz)                           | 128 |            |
| EM6E |      | 128 GB DDR4 Memory (2666 MHz)                          | 128 |            |
| EPWR |      | 8-core 3.6 GHZ processor                               | 32  |            |
| EPWS |      | 10-core 3.4 GHZ processor                              | 40  |            |
| EPWT |      | 12-core 3.15 GHZ processor                             | 48  |            |
| EPWX |      | 11-core X.YZ GHZ processor                             | 44  |            |
| FC   | CCIN | Description  | Max | OS support |
| EJ08 | 2CE2 | PCIe3 Optical Cable Adapter for PCIe3 Expansion Drawer | 8   | AIX Linux  |
| EC5J | 59B4 | 800 GB Mainstream U.2 SSD NVMe                         | 16  | AIX/Linux  |
| EC5K | 59B5 | 1.6 TB Mainstream U.2 SSD NVMe                         | 16  | AIX/Linux  |
| EC5L | 59B6 | 3.2 TB Mainstream U.2 SSD NVMe                         | 16  | AIX/Linux  |

E950

CPU + RAM  
Other

82

IBM Power Systems POWER9 Enterprise E950 Unofficial Deep Dive

© IBM Corporation, 2018



## POWER9 E950 Power Supply

Concurrent maintenance & redundant power – four units

### Rating

- 2000 W 200- 240 VAC
- Redundancy 2+2 → if cabled correctly

### Energy Efficiency

- 80+ Platinum Power Supply Compliant
- EPA Energy Star Compliant
- Built-in Advanced Thermal & Power Management

# Physicals

## POWER9 E950




### E950 size

- Width: 448 mm (17.5 in.)
- Height: 175 mm (6.9 in.)
- Depth: 902 mm (35.6 in.)

→ 4 EIA units

→ 5 inches greater than E850

→ T42 needs a 8 inch extension for rear cables



### E950 Weight:

- 69 kg (152 lb)
- With weight reduction (removing parts) ~121 lb
- Four person lift recommended or better yet a lifting tool

→ Heavier than the E850



85

IBM Power Systems POWER9 Enterprise E950 Unofficial Deep Dive


© IBM Corporation, 2018

## POWER9 E950



### E950 Electrical Power

- Operating voltage: 200 - 240 V AC
- Operating frequency: 50 - 60 Hz +/- 3 Hz
- Power consumption: 3,850 watts maximum
- Power source loading: 3.9 kVA maximum



### Note:

- Model 9040-MR9 use four power supply units
- C19/C20 connectors for higher amps → different to the E850 C13's

86

IBM Power Systems POWER9 Enterprise E950 Unofficial Deep Dive

© IBM Corporation, 2018

## POWER9 E950



### E950 Thermal output

- 14,403 Btu/hr maximum (per system node)

### E950 Maximum altitude:

- 3,050 m (10,000 ft)

### E950 Noise level

- Acoustics A-weighted Upper-Limit Sound Power Levels

One typically configured Four 8-core or 12-core, 2 TB memory:

- 7.4 bels (operating/idle: 25 C, 500 m)

One maximum configured Four 12-core, 2 TB memory:

- 8.1 bels (heavy workload, 25 C, 500 m)



87

IBM Power Systems POWER9 Enterprise E950 Unofficial Deep Dive

© IBM Corporation, 2018

## Call to Action

→ We need to make sure every one understands

Good News

1. E950 many improvement to the already good E850
2. Up to 48 CPU cores: POWER9 performance boost ~42%
3. Memory jump → 4 TB to 16 TB
4. Adapters → Easy to use rear PCIe cassette access
5. HMC – get them ready ASAP

Challenges

1. Electricity → C19/C20 connectors for higher amps, may need PDU changes
2. Size → It's a lot longer – need to check racks & doors
3. Heavy → Loading in to a rack needs extra care, use the Lifting Device
4. Noisy → May need ear protection, check the computer room policy



88

IBM Power Systems POWER9 Enterprise E950 Unofficial Deep Dive

© IBM Corporation, 2018

## Power E950 Specification Highlights



- ✓ 4U Server - 19" Rack Enclosure
- ✓ 2 or 4 POWER9 Enterprise SMT8 processors
- ✓ Processor SCMs enables efficient 1-Hop processor fabric interconnect
- ✓ Up to 16TB Total DDR4 DIMMs – up to 4TB per processor
  - 920 GB/s memory bandwidth per System
  - 230 GB/s memory bandwidth per Processor
  - 128 DIMM slots on 8 memory riser cards
  - 16 DIMM slots on each riser
- ✓ Capacity on Demand for Processor and Memory
- ✓ 10 PCIe Gen4 slots, 1 PCIe Gen3 - Blind swap, Full Height, Half Length
- ✓ Four High Speed 25Gb/s acceleration ports to attach to future accelerators



- ✓ 4 x NVMe Flash U.2 Bays (bootable)
- ✓ 8 x internal 8 SFF (2.5") SAS bays
- ✓ Storage controller adapters plug into PCIe slots
  - Single backplane for 0 or 1 SAS storage PCIe adapters
  - Split backplane for two storage SAS PCIe adapters
  - SAS PCIe adapters are concurrently maintainable
- ✓ Enhanced DC-DC Regulator Redundancy
- ✓ Full Fan Concurrent Maintenance
- ✓ I/O Expansion and Storage drawers support
- ✓ 4U Server - 19" Rack Enclosure

No Charge: PowerVM & PowerVC