


# IBM Power Systems


## POWER10 Processor for Regular Techies



Nigel Griffiths POWER Technical Staff Member, IBM, UK  
 IBM email: [nag@uk.ibm.com](mailto:nag@uk.ibm.com) @mr\_nmon twitter  
<http://tinyurl.com/njmon> - njmon sourceforge project  
<http://tinyurl.com/AIXpert> - My Blogs  
<https://www.youtube.com/user/nigelargriffiths> - My 215 Videos

[v10]

1



## POWER10 Processor for Regular Techies

Abstract:

Nigel will cover the POWER10 processor for **big commercial server** specialists such as Systems Administrator, Performance Guru, DevOps, Solution Designer, Systems Architects, and similar. **Not for Processor Architects/Designers or kernel coders.**

We will work through the publicly available information, followed by a little speculation - how the POWER9 generation servers might look, if they were running POWER10 processors a glimpse into the future.

**No IBM Confidential information nor making Announcements in this session.**

We will not be covering **dates, prices** or  
 performance **rPerf/CPW** ratings, **GHz** or **model names**.

Relevant for AIX / IBM i / Linux on POWER environments.

2


© 2020 IBM Corporation

2

## IBM's POWER10 Processor



# 32<sup>nd</sup> Hot Chips Conference - August 2020



William Starke

- Distinguished Engineer
- POWER10 Processor Architect



Brian Thompto

- Distinguished Engineer
- POWER10 Processor Core Architect


© 2020 IBM Corporation

3

## IBM's POWER10 Processor




# In September + October → YouTube



William Starke

IBM POWER10 Processor  
OpenPOWER Summit  
25 minute version with 1100+ Views  
<https://www.youtube.com/watch?v=27VRdl2BGWg>



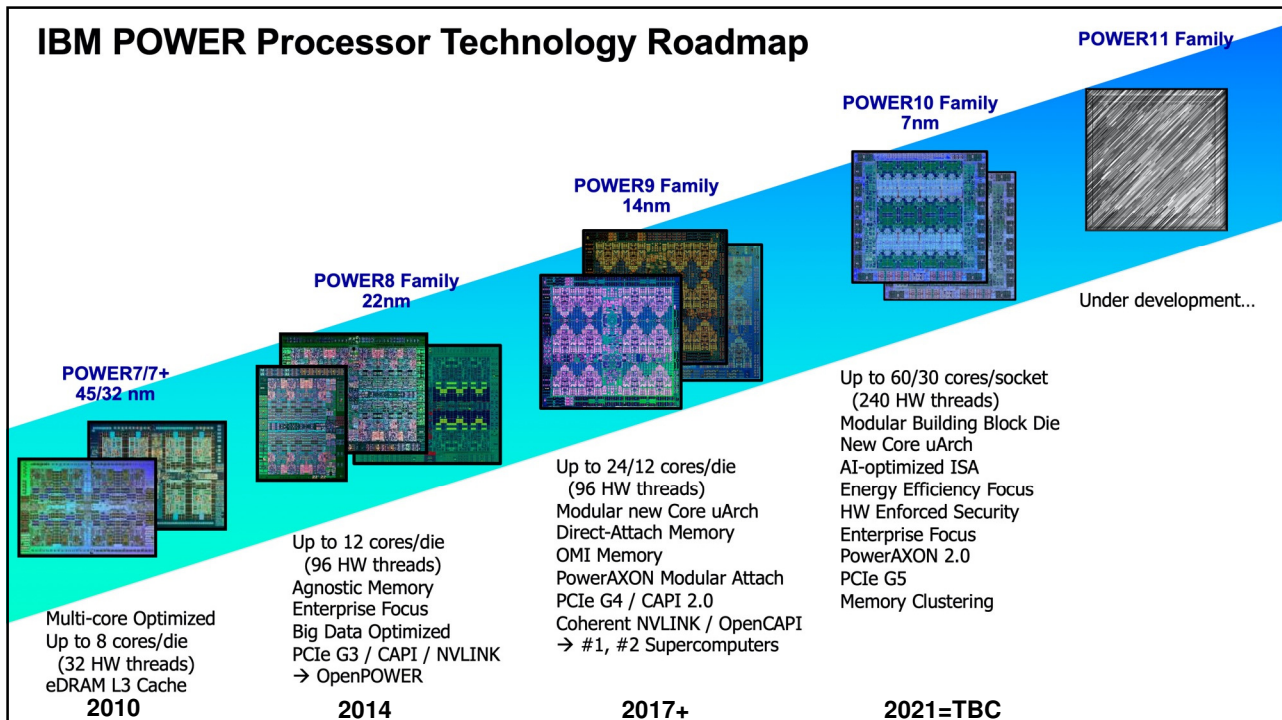
Brian Thompto

IBM POWER10 Processor: chip capabilities  
90 minute version with 180+ Views  
<https://www.youtube.com/watch?v=FMvret3p7qE>



© 2020 IBM Corporation

4



5

## POWER10 Design Focus

### Data Plane Bandwidth, Capacity, Composability, Scale

Terabyte/second sockets, Petabyte system memory capacities, 16-socket SMP → Clusters

### Powerful Enterprise Core

New Core Architecture, Flexibility, Larger caches, Reduced Latencies

### End-to-end Security

Hardware enabled and co-optimized with PowerVM hypervisor

### Energy Efficiency

3x improvement over POWER9

### AI-Infused Core

10-20x matrix-math performance / socket compared to POWER9

6

## POWER10 Processor Chip

### Technology and Packaging:

- 602mm<sup>2</sup> 7nm Samsung (18B devices)
- 18 layer metal stack, enhanced device
- Single-chip or Dual-chip sockets

### Computational Capabilities:

- Up to 15 SMT8 Cores (2 MB L2 Cache / core)  
(Up to 120 simultaneous hardware threads)
- Up to 120 MB L3 cache (low latency NUCA mgmt)
- 3x energy efficiency relative to POWER9
- Enterprise thread strength optimizations
- AI and security focused ISA additions
- 2x general, 4x matrix SIMD relative to POWER9
- EA-tagged L1 cache, 4x MMU relative to POWER9

### Open Memory Interface:

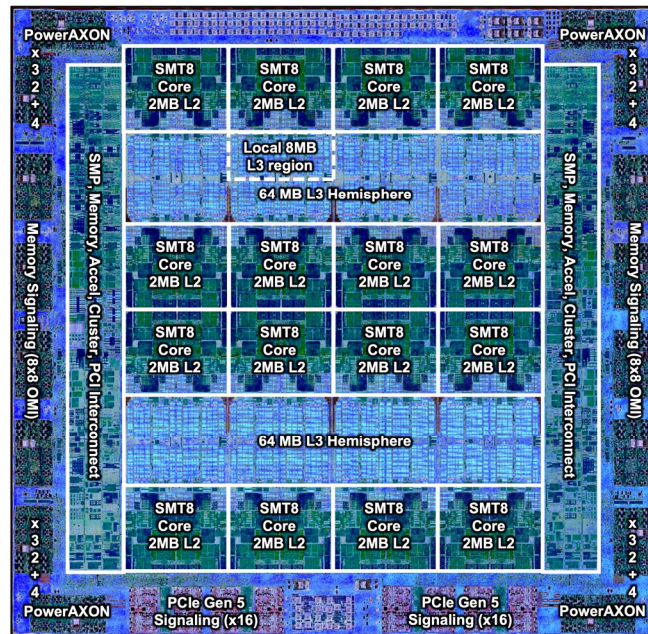
- 16 x8 at up to 32 GT/s (1 TB/s)
- Technology agnostic support: near/main/storage tiers
- Minimal (< 10ns latency) add vs DDR direct attach

### PowerAXON Interface:

- 16 x8 at up to 32 GT/s (1 TB/s)
- SMP interconnect for up to 16 sockets
- OpenCAPI attach for memory, accelerators, I/O
- Integrated clustering (memory semantics)

### PCIe Gen 5 Interface:

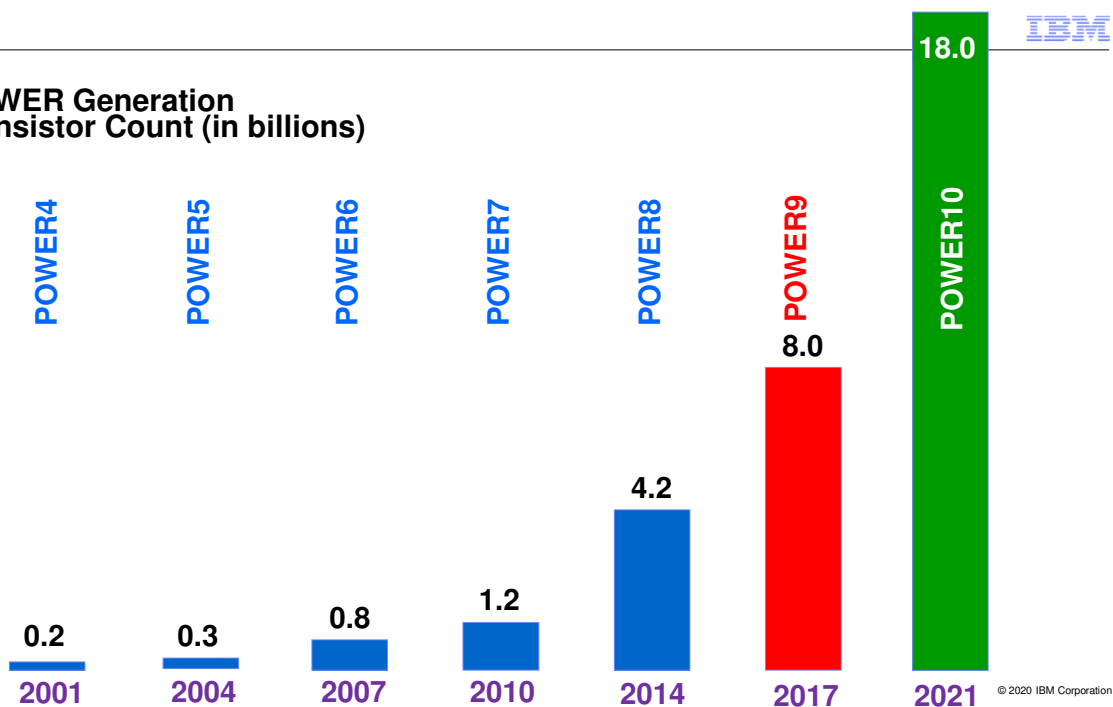
- x64 / DCM at up to 32 GT/s



Die Photo courtesy of Samsung Foundry

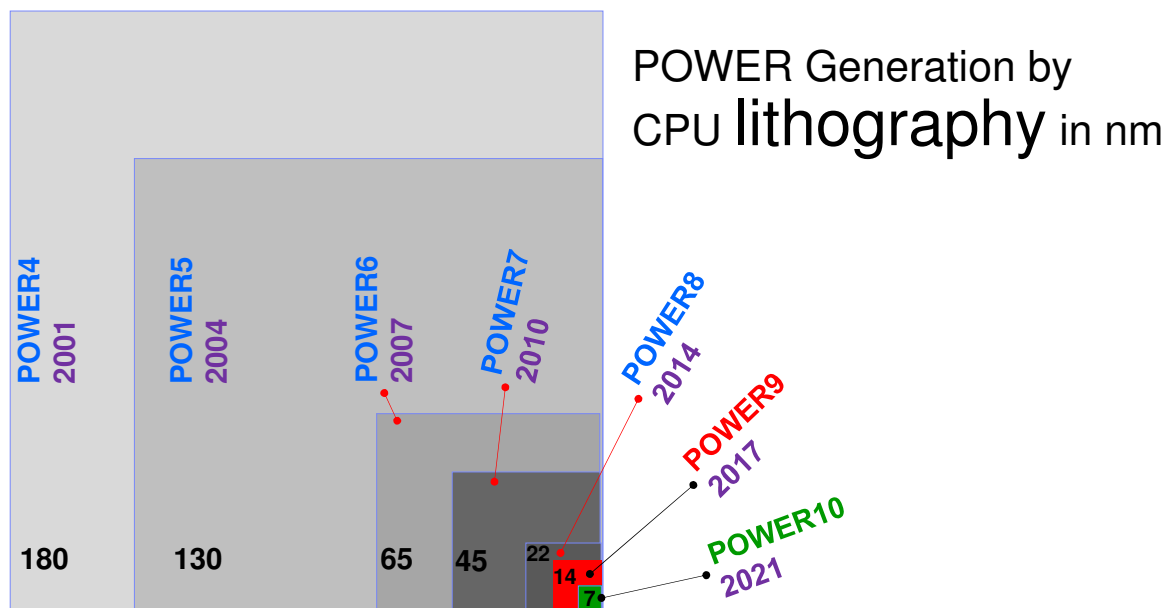
7

## POWER Generation Transistor Count (in billions)



8





© 2020 IBM Corporation

9

## POWER10 Processor Chip

### Technology and Packaging:

- 602mm<sup>2</sup> 7nm Samsung (18B devices)
- 18 layer metal stack, enhanced device
- Single-chip or Dual-chip sockets

### Computational Capabilities:

- Up to 16 SMT8 Cores (2 MB L2 Cache / core) (Up to 120 simultaneous hardware threads)
- Up to 120 MB L3 cache (low latency NUCA mgmt)
- 3x energy efficiency relative to POWER9
- Enterprise thread strength optimizations
- AI and security focused ISA additions
- 2x general, 4x matrix SIMD relative to POWER9
- EA-tagged L1 cache, 4x MMU relative to POWER9

### Open Memory Interface:


- 16 x8 at up to 32 GT/s (1 TB/s)
- Technology agnostic support: near/main/storage tiers
- Minimal (< 10ns latency) add vs DDR direct attach

### PowerAXON Interface:

- 16 x8 at up to 32 GT/s (1 TB/s)
- SMP interconnect for up to 16 sockets
- OpenCAPI attach for memory, accelerators, I/O
- Integrated clustering (memory semantics)

### PCIe Gen 5 Interface:

- x64 / DCM at up to 32 GT/s




Die Photo courtesy of Samsung Foundry

10

Core Yield = not to scale = for illustration purposes only

IBM

POWER chips with the highest numbers of cores are the most popular



By not offering the 16 cores option, there are a lot more chips available (with the 15 cores)

Over the years of manufacturing yield of the top core chips slowly improve!

11

© 2020 IBM Corporation

11

## POWER10 Processor Chip

### Technology and Packaging:

- 602mm<sup>2</sup> 7nm Samsung (18B devices)
- 18 layer metal stack, enhanced device
- Single-chip or Dual-chip sockets

### Computational Capabilities:

- Up to 15 SMT8 Cores (2 MB L2 Cache / core) (Up to 120 simultaneous hardware threads)
- Up to 120 MB L3 cache (low latency NUCA mgmt)
- 3x energy efficiency relative to POWER9
- Enterprise thread strength optimizations
- AI and security focused ISA additions
- 2x general, 4x matrix SIMD relative to POWER9
- EA-tagged L1 cache, 4x MMU relative to POWER9

### Open Memory Interface:

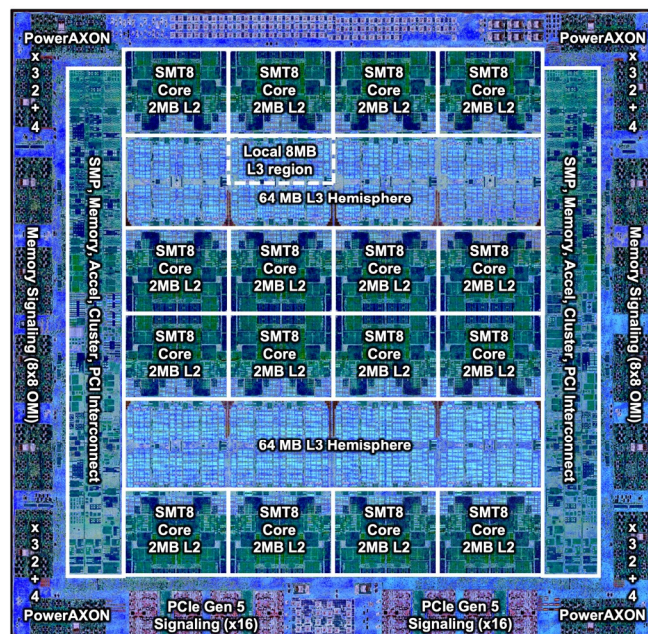
- 16 x8 at up to 32 GT/s (1 TB/s)
- Technology agnostic support: near/main/storage tiers
- Minimal (< 10ns latency) add vs DDR direct attach

### PowerAXON Interface:

- 16 x8 at up to 32 GT/s (1 TB/s)
- SMP interconnect for up to 16 sockets
- OpenCAPI attach for memory, accelerators, I/O
- Integrated clustering (memory semantics)

### PCIe Gen 5 Interface:


- x64 / DCM at up to 32 GT/s



Die Photo courtesy of Samsung Foundry

12






13




14






15



16




## Socket Composability: **SCM & DCM**



### Single-Chip Module Focus:


- 602mm<sup>2</sup> 7nm (18B devices)
- **Core/thread Strength**
  - Up to 15 SMT8 Cores (4+ GHz)
- **Capacity & Bandwidth / Compute**
  - Memory: x128 @ 32 GT/s
  - SMP/Cluster/Accel: x128 @ 32 GT/s
  - I/O: x32 PCIe G5
- **System Scale (Broad Range)**
  - 1 to 16 sockets




### Dual-Chip Module Focus:

- 1204mm<sup>2</sup> 7nm (36B devices)
- **Throughput / Socket**
  - Up to 30 SMT8 Cores (3.5+ GHz)
- **Compute & I/O Density**
  - Memory: x128 @ 32 GT/s
  - SMP/Cluster/Accel: x192 @ 32 GT/s
  - I/O: x64 PCIe G5
  - 1 to 4 sockets

Up to  
16 SCM  
Sockets



Up to  
4 DCM  
Sockets




(Multi-socket configurations show processor capability only, and do not imply system product offerings)

**IBM POWER10**


17

## System Composability: **PowerAXON & Open Memory Interfaces**




**PowerAXON corner**  
4x8 @ 32 GT/s

**Multi-protocol**  
"Swiss-army-knife"  
Flexible / Modular Interfaces



**Built on best-of-breed**  
Low Power, Low Latency,  
High Bandwidth  
Signaling Technology



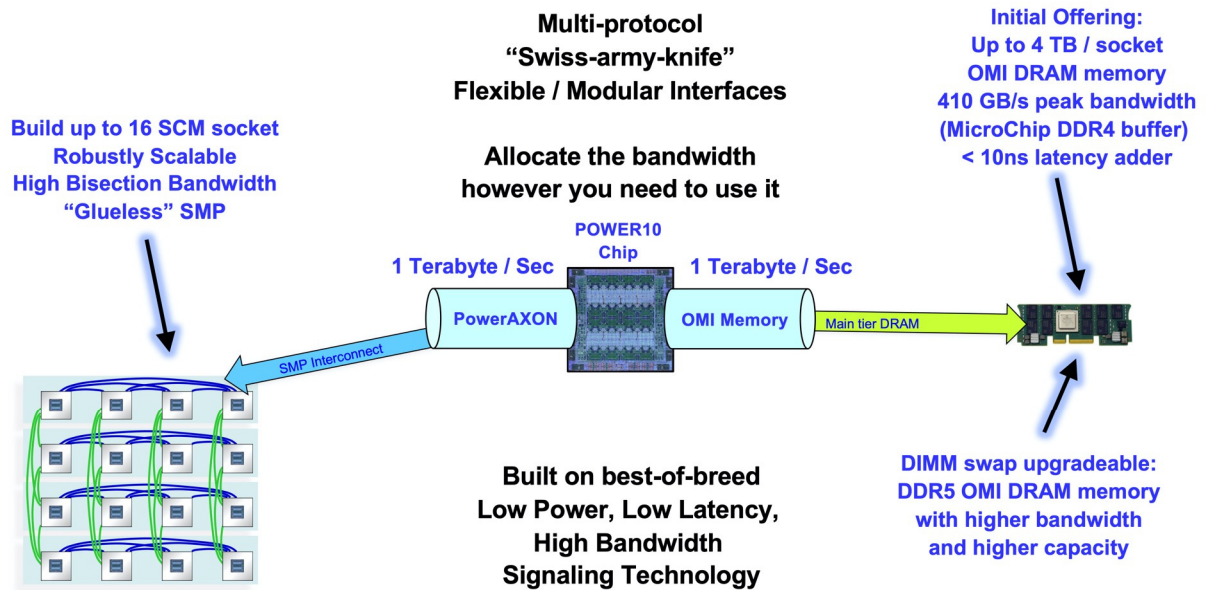
**OMI edge**  
8x8 @ 32 GT/s

**6x bandwidth / mm<sup>2</sup>**  
compared to DDR4  
signaling

**IBM POWER10**

18

## System Enterprise Scale and Bandwidth: SMP & Main Memory



(PowerAXON and OMI Memory configurations show processor capability only, and do not imply system product offerings)

**IBM POWER10**

19

**IBM**


### Future Memory 101

#### DIMM DDR4

- Parallel
- Clunky protocol
- Available for a few year

#### OMI

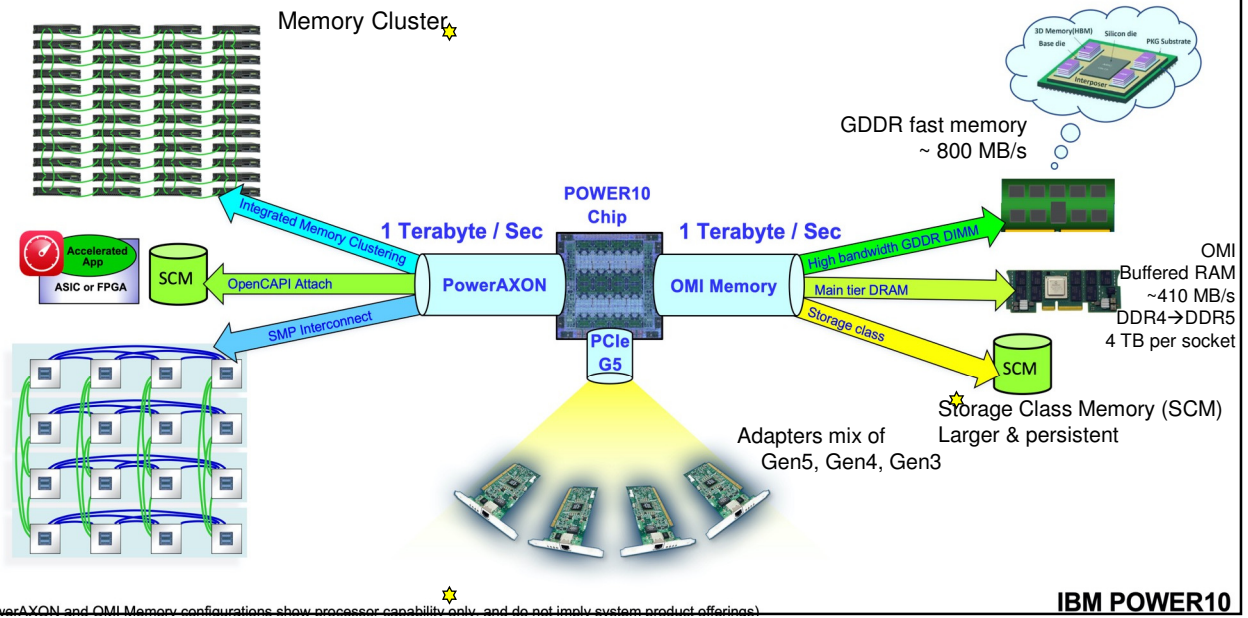
- Open Memory Interface
- MicroChip DDR4 buffer
- Serial interface (fewer pins)
- Smaller = less space
- Just became available



© 2020 IBM Corporation

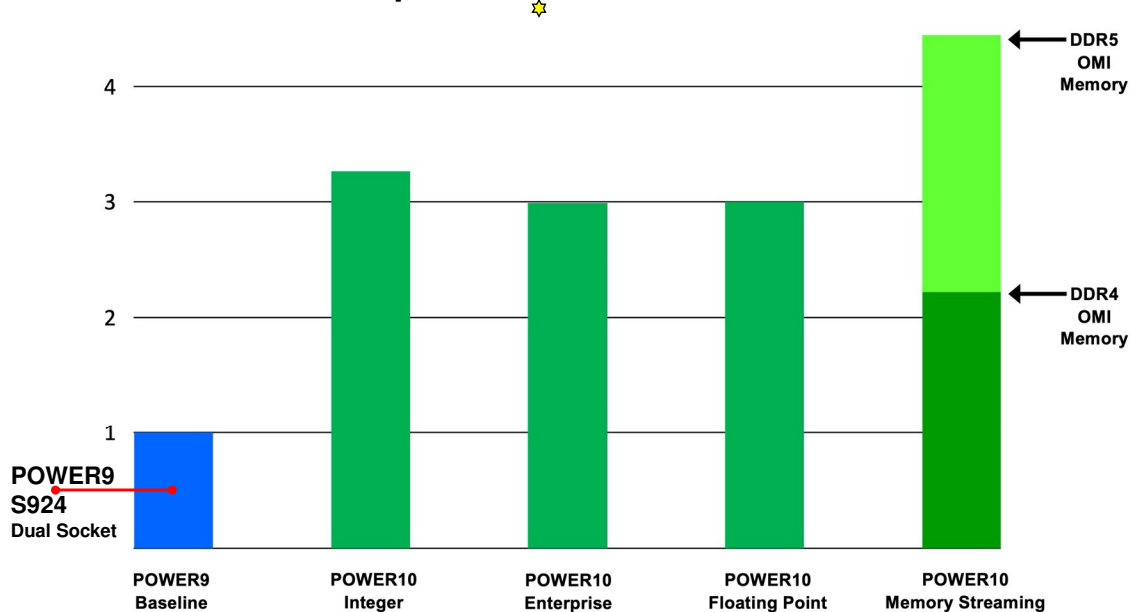
20

## System Composability: PCIe Gen 5 Industry I/O Attach



21

## POWER10 General Purpose Socket Performance Gains



(Performance assessments based upon pre-silicon engineering analysis of POWER10 dual-socket server offering vs POWER9 dual-socket server offering)

**IBM POWER10**

22



# More Details

- May be more details than you really want!

© 2020 IBM Corporation

23

## Powerful Core : Enterprise Flexibility

### Multiple World-class Software Stacks

#### Resilience and full stack integrity

- PowerVM, KVM
- AIX, IBMi, Linux on Power, OpenShift

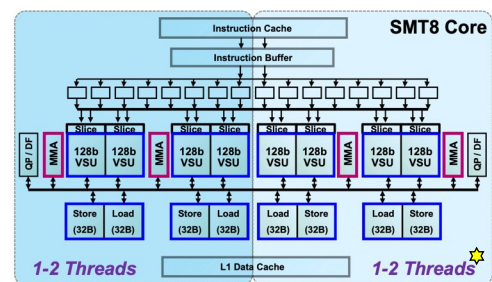


#### Partition flexibility and security

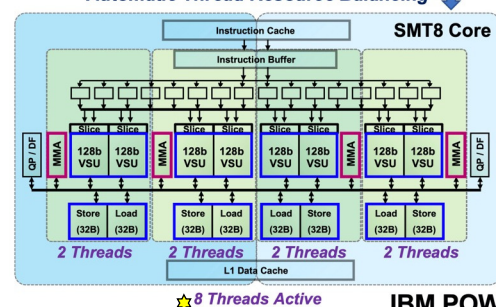
- Full-core level LPAR
- Thread-based LPAR scheduling
- **NEW:** With PowerVM Hypervisor
  - Nested KVM + PowerVM
  - Hardware assisted container/VM isolation



### Hardware Based Workload Balance



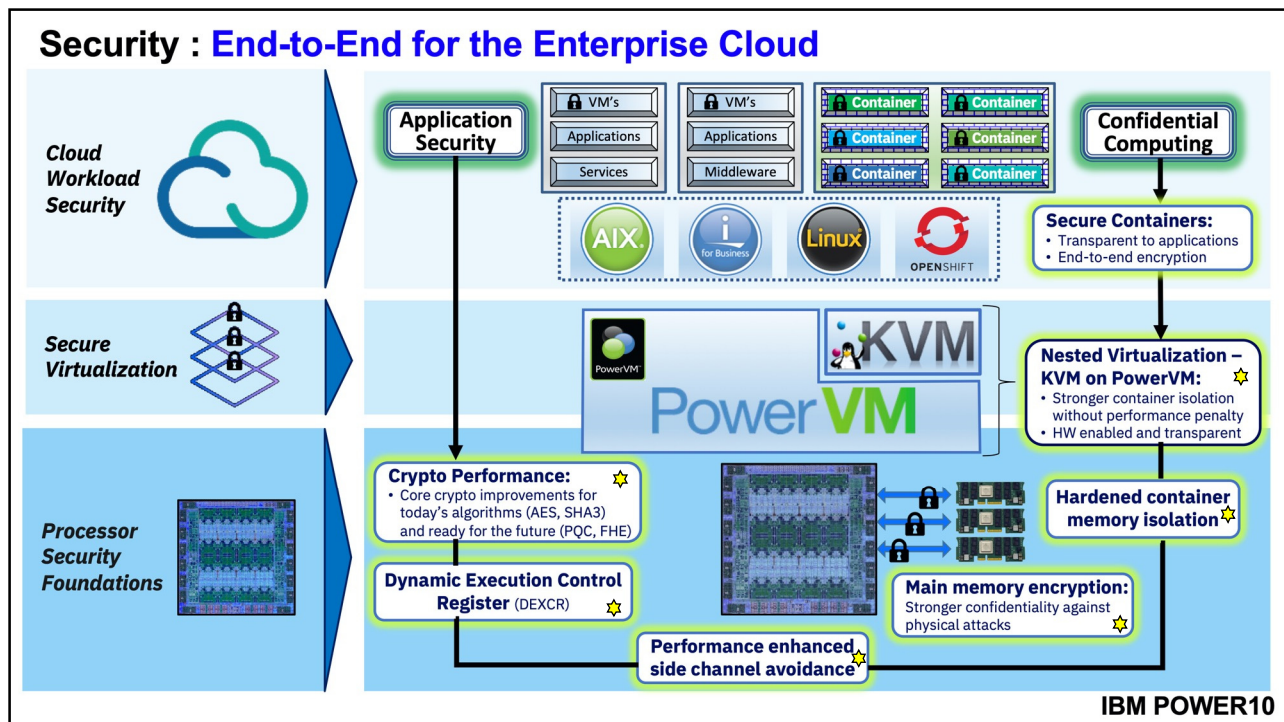
### Automatic Thread Resource Balancing



IBM POWER10

24






25

# POWER10 Servers


- Assumes some familiarity with the POWER9 servers
- Public POWER10 presentations make references to POWER10 servers with two, four and sixteen sockets

© 2020 IBM Corporation

26




# POWER9 Servers



**Scale-Out**

2U or 4U  
1 or 2 socket  
SCM 12 cores SMT=8  
4 to 24 cores  
4TB RAM




**S924**

**S922**

**Midrange**


4U  
2 or 4 socket  
SCM 12 cores SMT=8  
16 to 48 cores  
16TB RAM



**E950**

**Enterprise**


7U to 22U  
4 to 16 socket  
SCM 12 core  
SMT=8  
8 to 192 cores  
64TB RAM




**E980**

© 2020 IBM Corporation

27




# POWER9 Servers



**Scale-Out**

2U or 4U  
1 or 2 socket  
SCM 12 cores SMT=8  
4 to 24 cores  
4TB RAM




**S924**

**S922**

**Midrange**


4U  
2 or 4 socket  
SCM 12 cores SMT=8  
16 to 48 cores  
16TB RAM



**E950**

**Enterprise**

7U to 22U  
4 to 16 socket  
SCM 12 core  
SMT=8  
8 to 192 cores  
64TB RAM




**E980**

**POWER8** and **POWER9** had a similar set of servers  
 Well understood and popular with clients  
 Let us **assume** that this continues . . . for **POWER10**  
**This is NOT an IBM statement but Nigel guesswork**

© 2020 IBM Corporation


28

IBM



**POWER9 Servers**


**Enterprise**



**E980**

**Scale-Out**


2U or 4U  
1 or 2 socket  
SCM 12 cores SMT=8  
4 to 24 cores  
4TB RAM



**S924**  
**S922**

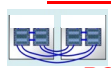
**Midrange**


4U  
2 or 4 socket  
SCM 12 cores SMT=8  
16 to 48 cores  
16TB RAM



**E950**

2U or 4U with 1 or 2 sockets  
**DCM 15 cores** SMT=8  
Up to **60 cores**  
**OMI** RAM  
PCIe Gen5 + Gen4





Two Dual Chip Module Sockets holding four POWER10  
processors with up to 15 cores each = **60 POWER10 cores**


**POWER10 Servers**

This can be deduced from POWER10 chip  
details via the Hot Chips conference

© 2020 IBM Corporation


29

IBM



**POWER9 Servers**


**Enterprise**



**E980**

**Scale-Out**


2U or 4U  
1 or 2 socket  
SCM 12 cores SMT=8  
4 to 24 cores  
4TB RAM



**S924**  
**S922**

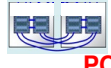
**Midrange**


4U  
2 or 4 socket  
SCM 12 cores SMT=8  
16 to 48 cores  
16TB RAM



**E950**

2U or 4U with 1 or 2 sockets  
**DCM 15 cores** SMT=8  
Up to **60 cores**  
**OMI** RAM  
PCIe Gen5 + Gen4





4 DCM Sockets 8 POWER10's  
15 cores = **120 POWER10 cores**


**POWER10 Servers**

This can be deduced from POWER10 chip  
details via the Hot Chips conference

© 2020 IBM Corporation

30

IBM



# POWER9 Servers

**Scale-Out**

S924  
S922

2U or 4U  
1 or 2 socket  
SCM 12 cores SMT=8  
4 to 24 cores  
4TB RAM

**Midrange**

E950

4U  
2 or 4 socket  
SCM 12 cores SMT=8  
16 to 48 cores  
16TB RAM


**Enterprise**

E980


7U to 22U  
4 to 16 socket  
SCM 12 core  
SMT=8  
8 to 192 cores  
64TB RAM

---


2U or 4U with 1 or 2 sockets  
**DCM 15 cores** SMT=8  
Up to **60 cores**  
**OMI** RAM  
PCIe Gen5 + Gen4




4U with 4 sockets  
**DCM 15 cores** SMT=8  
Up to **120 cores**  
**OMI** RAM  
PCIe Gen5



22U with 4 to 16 sockets  
**SCM 15 cores** SMT=8  
Up to **240 cores**  
**OMI** RAM  
PCIe Gen5



**240**  
**POWER10**  
cores



## POWER10 Servers

This can be deduced from POWER10 chip details via the Hot Chips conference

© 2020 IBM Corporation

31

IBM

Reminder – this is not an announcement

From public facts on the POWER10 processor chips

- + some assumptions
- + some guesswork

**POWER8** and **POWER9** had a similar set of servers

Well understood and popular with clients

Let us **assume** that this continues . . . for **POWER10**

**This is NOT an IBM statement but Nigel guesswork**


32
© 2020 IBM Corporation

32





## POWER10 SIMD / AI Socket Performance Gains



(Performance assessments based upon pre-silicon engineering analysis of POWER10 dual-socket server offering vs POWER9 dual-socket server offering)

IBM POWER10

35

IBM

## Summary for Regular Techies

36

© 2020 IBM Corporation

36



## POWER10 Speeds and Feeds Summary

### Performance Gains

- POWER10 with 15 cores = 25% Jump
- + Dual Chip Modules with 30 CPU cores = 100% Jump
- + CPU core thread strength (SMT=8) improvement = 20% Jump
- = POWER10 servers are going to be

**~3 times faster**

With PowerAXON, OMI, PCIe Gen5 scaling and NVMe disks  
- No internal bandwidth limitations to slow down the processor

37

© 2020 IBM Corporation

37



## POWER10 Speeds and Feeds Summary

### Given:

- same size servers &
- expected similar electricity use &
- higher performance
- = **Superb Green Credentials** (transactions per Watt)

38

© 2020 IBM Corporation

38



## POWER10 Client Practical Benefits

### Stronger (faster) Cores and Threads

- Reduce core counts = reduced Software Licence costs

### Larger Virtual Machines per server

- Scale-out: two 10 core VM plus VIOSs → two 28 core VMs plus VIOSs
- Mid-range: two 20 core VM plus VIOSs → two 58 core VMs plus VIOSs
  - Larger VM = faster Apps = less performance issues = reduced SysAdmin

### More Virtual Machines per server

- Scale-out: two 10 core VM plus VIOSs → seven 8 core VMs plus VIOSs
- Mid-range: two 20 core VM plus VIOSs → seven 16 core VMs plus VIOSs plus spare 4 cores
  - More VM per server = consolidation of servers = less SysAdmin time
  - Reduced rack space/floor space, electricity, network connections

Other arguments require prices and rPerf/CPW for comparisons

39

© 2020 IBM Corporation

39

## Questions?

No questions on

- Dates
- Prices
- Performance
- rPerf or CPW ratings
- GHz
- Model names
- Future Lotto  
winning numbers!



40

40





## Special notices

This document was developed for IBM offerings in the United States as of the date of publication. IBM may not make these offerings available in other countries, and the information is subject to change without notice. Consult your local IBM business contact for information on the IBM offerings available in your area.

Information in this document concerning non-IBM products was obtained from the suppliers of these products or other public sources. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. Send license inquiries, in writing, to IBM Director of Licensing, IBM Corporation, New Castle Drive, Armonk, NY 10504-1785 USA.

All statements regarding IBM future direction and intent are subject to change or withdrawal without notice and represent goals and objectives only.

The information contained in this document has not been submitted to any formal IBM test and is provided "AS IS" with no warranties or guarantees either expressed or implied.

All examples cited or described in this document are presented as illustrations of the manner in which some IBM products can be used and the results that may be achieved. Actual environmental costs and performance characteristics will vary depending on individual client configurations and conditions.

IBM Global Financing offerings are provided through IBM Credit Corporation in the United States and other IBM subsidiaries and divisions worldwide to qualified commercial and government clients. Rates are based on a client's credit rating, financing terms, offering type, equipment type and options, and may vary by country. Other restrictions may apply. Rates and offerings are subject to change, extension or withdrawal without notice.

IBM is not responsible for printing errors in this document that result in pricing or information inaccuracies.

All prices shown are IBM's United States suggested list prices and are subject to change without notice; reseller prices may vary.

IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.

Any performance data contained in this document was determined in a controlled environment. Actual results may vary significantly and are dependent on many factors including system hardware configuration and software design and configuration. Some measurements quoted in this document may have been made on development-level systems. There is no guarantee these measurements will be the same on generally-available systems. Some measurements quoted in this document may have been estimated through extrapolation. Users of this document should verify the applicable data for their specific environment.

Revised September 26, 2006

© 2020 IBM Corporation



## Special notices (continued)

IBM, IBM (logo), AIX, AIX (logo), EnergyScale, IBM i, i for business (logo), Power, POWER, PowerVM, PowerVM (logo), PowerLinux, PowerLinux (logo), Power Architecture, Power ISA, POWER9, and POWER10 are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries.

A full list of U.S. trademarks owned by IBM may be found at: <http://www.ibm.com/legal/copytrade.shtml>.

Red Hat, OpenShift, and the OpenShift logo are registered trademarks of Red Hat, Inc. in the United States and other countries.

The OpenPOWER word mark and the OpenPOWER logo mark, and related marks, are trademarks and service marks licensed by OpenPOWER Foundation.

OpenCAPI and the OpenCAPI logo are trademarks of the OpenCAPI Consortium.

Linux is a registered trademark of Linus Torvalds in the United States, other countries or both.

PowerLinux™ uses the registered trademark Linux® pursuant to a sublicense from LMI, the exclusive licensee of Linus Torvalds, owner of the Linux® mark on a world-wide basis.

Other company, product and service names may be trademarks or service marks of others.

© 2020 IBM Corporation