



**IBM i**

**common**

EUROPE LUXEMBOURG

**29 septembre 2022**

**D'Braustuff – National Brewery**

**IBM Power10**

*Making secure, resilient, agile,  
hybrid multi-cloud and AI  
workloads possible*



Fabian Michel

Senior Power Systems & Storage Technical Specialist

IBM Technology, Belgium

[fabian\\_michel@be.ibm.com](mailto:fabian_michel@be.ibm.com)

@Fabian\_Michel

**IBM** | Let's Create

**IBM**



## *Agenda du jour*

*Le processeur Power10*

*La nouvelle gamme Power Systems*

*Les performances*

*Matrice de support des OS*

*Stockage NVMe*

**common**  
EUROPE LUXEMBOURG



**IBM** | Let's  
Create



common  
EUROPE LUXEMBOURG



IBM | Let's  
Create

## *Le Processeur IBM Power10*

# IBM POWER Processor Technology Roadmap

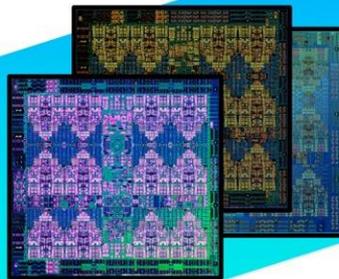
POWER11 Family

POWER10 Family  
7nm



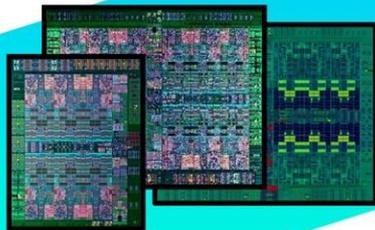
Under development...

POWER9 Family  
14nm



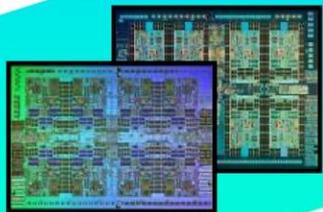
Up to 24/12 cores/die  
(96 HW threads)  
Modular new Core uArch  
Direct-Attach Memory  
OMI Memory  
PowerAXON Modular Attach  
PCIe G4 / CAPI 2.0  
Coherent NVLINK / OpenCAPI  
→ #1, #2 Supercomputers

POWER8 Family  
22nm



Up to 12 cores/die  
(96 HW threads)  
Agnostic Memory  
Enterprise Focus  
Big Data Optimized  
PCIe G3 / CAPI / NVLINK  
→ OpenPOWER

POWER7/7+  
45/32 nm



Multi-core Optimized  
Up to 8 cores/die  
(32 HW threads)  
eDRAM L3 Cache

Up to 60/30 cores/socket  
(240 HW threads)  
Modular Building Block Die  
New Core uArch  
AI-optimized ISA  
Energy Efficiency Focus  
HW Enforced Security  
Enterprise Focus  
PowerAXON 2.0  
PCIe G5  
Memory Clustering

# 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:

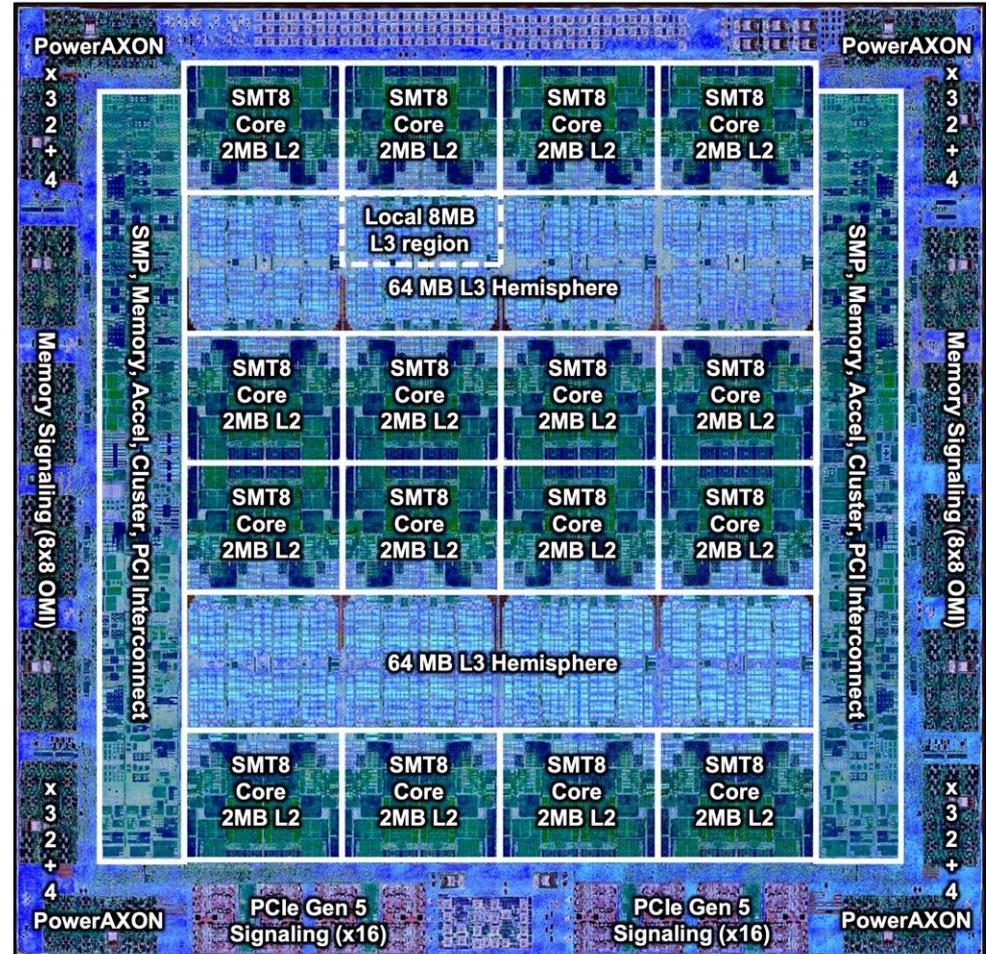
- 16 x8 at up to 32 GT/s (1 TB/s)
- Technology agnostic support: near/main/storage tiers
- Minimal (< 10ns latency) add vs DDR direct attach

## 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

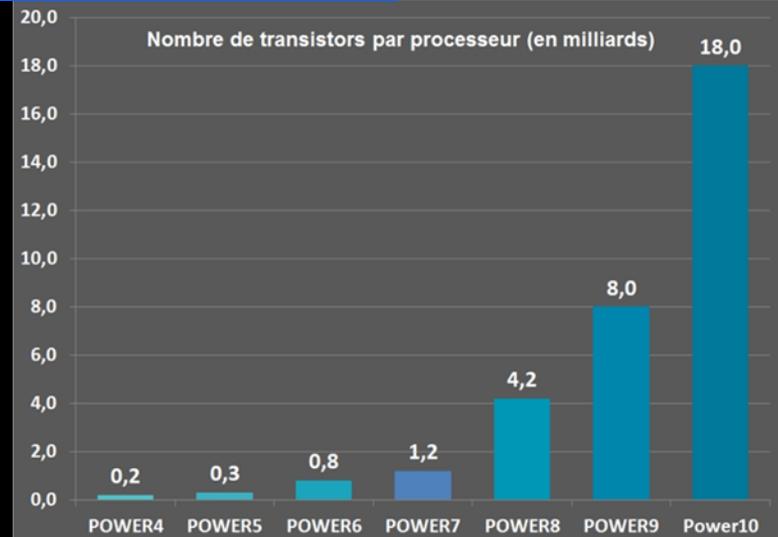
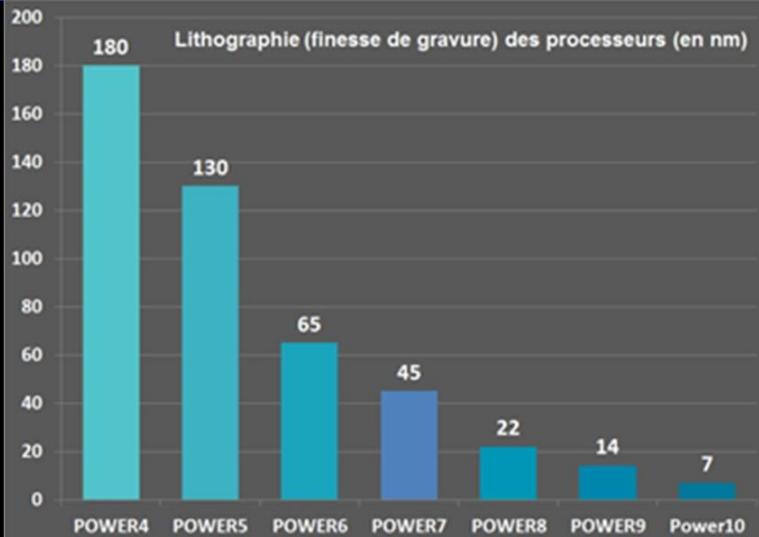


Die Photo courtesy of Samsung Foundry

7nm leads to a higher core density, enhanced per core performance, more energy efficiency and...  
More innovation!

Even more secure with transparent memory encryption!

Real time insights with built-in AI accelerators!



# Why IBM Power10?

## Hybrid Cloud Simplicity

Frictionless<sup>1</sup> deployment  
provides built-in cost  
optimization

## Security

Transparent in-memory  
encryption simplifies  
protection without  
impacting performance

## Resiliency

25% lower downtime vs.  
comparable  
high-end servers<sup>2</sup>

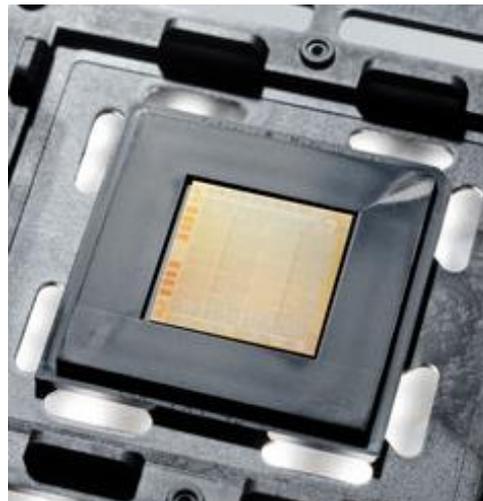
## AI

Built-in inference engine  
brings AI closer to your  
data to reduce cost and  
complexity

## Sustainability

33% lower energy  
consumption for the same  
workload with POWER9<sup>3</sup>

1. IBM Power helps deliver a frictionless experience in extending mission-critical workloads across hybrid cloud, without requiring additional middleware or application refactoring.
2. Based on "ITIC 2021 Global Server Hardware, Server OS Reliability Report", July 2021
3. Power E1080 compared to Power E980; POWER9 (12C) IS 5081 RPERF @ 16,520 WATTS (0.31 RPERF/WATT), POWER10 (15C) IS 7998 RPERF @ 17,320 WATTS (0.46 RPERF/WATT) 0.46 / 0.31 = 1.48 MORE RPERF/WATT



# IBM Power

## Engineered for Agility



# Power10 – Engineered for Agility

<b>Cloud Agility and Economics</b>	<b>Pervasive security and resilience</b>	<b>Streamline insights and automation</b>
------------------------------------	--	---

## Sept/Nov 2021 High End

- Highest Performance, most scalable system in industry
- Increased flexibility with Power Private Cloud
- Most scalable and agile SAP HANA



## Summer 2022 Scale Out & Midrange Systems

- > 2X capacity growth
- Power Private Cloud across the line
- Enhanced resilience



# IBM Power

Engineered  
for agility



Provides a  
frictionless hybrid  
cloud experience

## Respond faster to business demands

with efficient scaling and  
consistent **pay-for-use  
consumption across  
public and private clouds**

## Protect data from core to cloud

using **memory encryption  
at the processor level**  
designed to support end-  
to-end security across  
public and private clouds  
**without impacting  
performance**

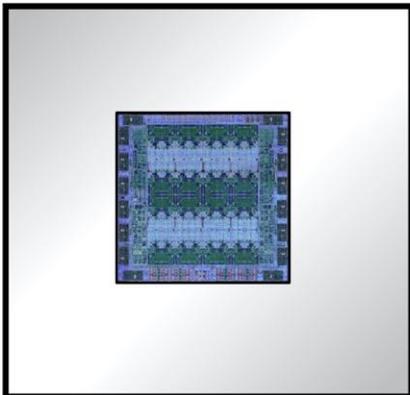
## Streamline insights and automation

by running **AI inferencing  
directly in core** and  
leveraging Watson  
services in IBM Cloud

## Maximize availability and reliability

with built-in advanced recovery and self-healing for  
infrastructure redundancy and disaster recovery in IBM Cloud

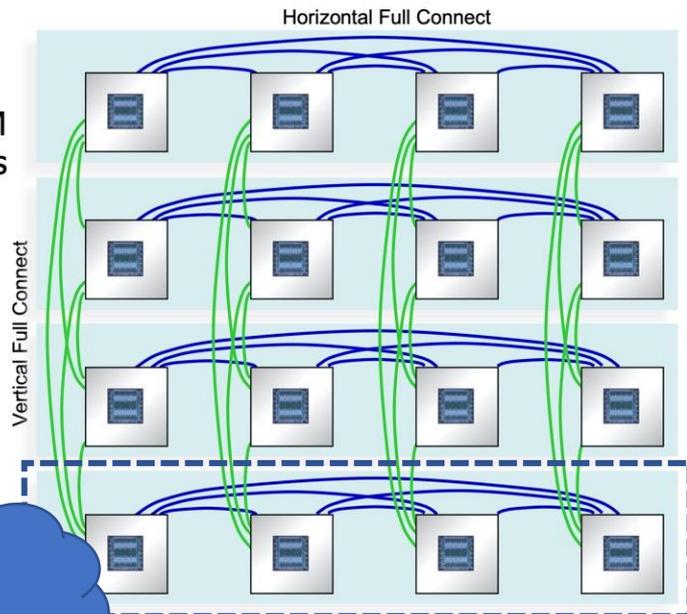
# Socket Composability: **SCM & DCM**



## 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

Up to  
16 SCM  
Sockets



16 sockets x 15c =  
Max. 240 cores  
x 8 = 1920 threads



# IBM Power E1080

## Engineered for agility



Respond faster to business demands

**World record two-tier SAP SD standard application benchmark**

for 8-sockets<sup>1</sup>

**4.1X**

containerized throughput per core than compared x86<sup>3</sup>

Protect data from core to cloud

**2.5X**

faster AES crypto performance per core over Power9<sup>2</sup>

**4X**

crypto engines in every core over Power9

Streamline insights and automation

**5X**

faster AI inferencing per socket for high precision math over Power E980<sup>4</sup>

1. IBM Power E1080, two-tier SAP SD standard application benchmark running SAP ERP 6.0 EHPS, Power10 3.55-4.0 GHz processor, 4,096 GB memory, 8p/120c/960t, 174,000 SD benchmark users (955,050 SAPs), 93 sec average response time, 19,131,000 line items/hr, 571,203,000 dialog steps/hr, 99% utilization, AIX 7.2, DB2 11.5 w/ average request time .019 sec dialog/.048 sec update. Certification # not available at time of press, all results can be found at [www.ibm.com/benchmark](https://www.ibm.com/benchmark) Valid as of 8/27/21

2. AES-256 in both GCM and XTS modes runs about 2.5 times faster per core when comparing Power10 E1080 (15-core modules) vs. Power9 E980 (12-core modules) according to preliminary measurements obtained on RHEL Linux 8.4 and the OpenSSL 1.1.1g library

3. Based on IBM internal testing of Red Hat OpenShift Container Platform 4.8.2 worker nodes running 80 pods using the Daytrader7 workload (<https://github.com/WASdev/sample.daytrader7/releases/tag/v1.4>) accessing an AIX Db2 database. Comparison E1080 running OCP accessing AIX Db2 on a 5922 versus OCP on Cascade Lake accessing AIX Db2 on a 5922. Valid as of 8/25/2021 and conducted under laboratory conditions. Individual result can vary based on workload size, use of storage subsystems & other conditions.

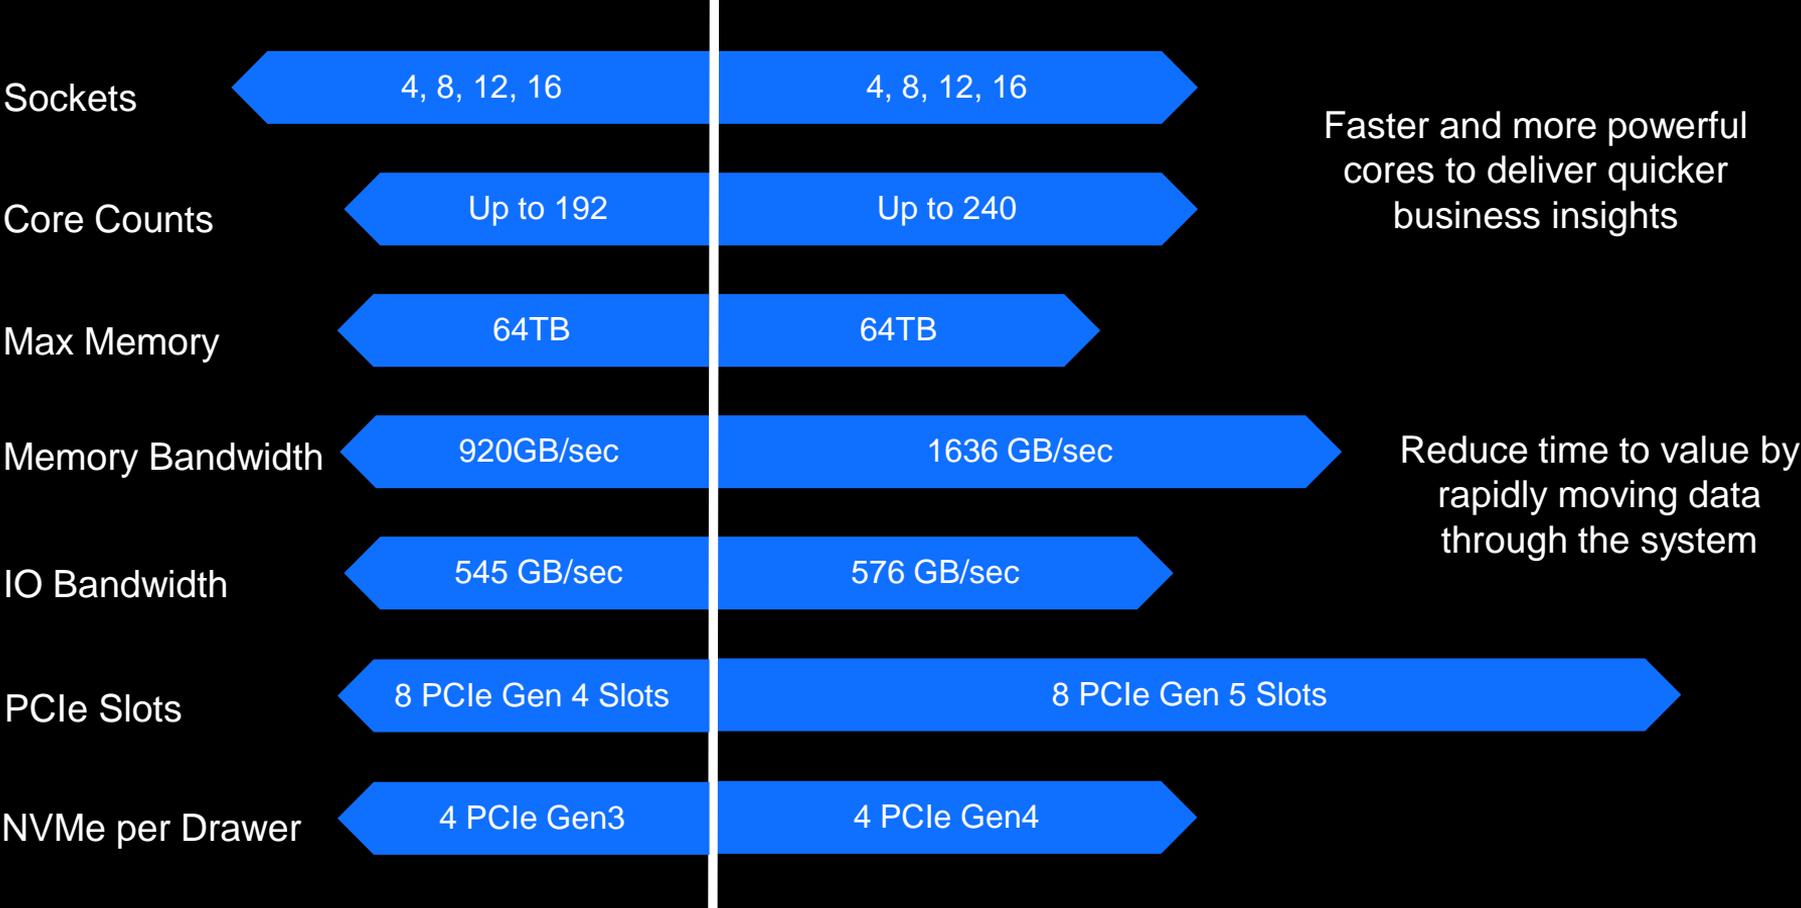
IBM Power E1080 (40 cores/3.8 GHz/2 TB memory) in maximum performance mode, 25 Gb two-port SRIOV adapter, 1 x 16gbps FCA, Websphere Liberty 21.0.0.6,Java(TM) SE Runtime Environment (build 8.0.6.36 - pxi6480sr6fp36-20210824\_02(SR6 FP36)), CoreOS Linux 4.18.0-305.10.2.el8\_4 with PowerVM. Configuration consists of 2 OCP worker lpar each with 10 cores running SMT8 with 256GB of memory, x86 s system: Intel(R) Xeon(R) Gold 6248 CPU (Cascade Lake) in performance mode, 40 cores/3.9GHz/2512GB memory), 25Gb two-port SRIOV adapter, 1 x 16gbps FCA, Websphere Liberty 21.0.0.6, Java(TM) SE Runtime Environment (build 8.0.6.36 - pxi6480sr6fp36-20210824\_02(SR6 FP36)), CoreOS Linux 4.18.0-305.10.2.el8\_4, RHEL 8.4 KVM. Configuration consists of 2 OCP worker KVM guests each with 20 cores running hyperthreading with 256GB of memory.

Database system 5922: Model 9009-22G with 22 cores and 1TB of memory. Configuration consists of 2 AIX lpar each with 8 cores running SMT8 with 131GB of memory, and a WOS lpar with 2 cores and 16GB of memory.

4. 5x improvement in per socket inferencing throughput for large size 32b floating point inferencing models from Power9 E980 (12-core modules) to Power10 E1080 (15-core modules ).Based on IBM testing using Pytorch, OpenBLAS on the same BERT Large with SqUAD v1.1 data set

5. Based on 3408 FLOPS for Power10 15c SCM (3.55 GHz x 64 x 15 SMT8 Cores) compared to 3162 FLOPS for 3rd Generation Intel Xeon Platinum 8368Q CPU (2.6 GHz x 32 x 38 Cores) based on x86 Scalable Platform Intel product brief: <https://www.mca.com/en/press/3rd-gen-intel-xeon-scalable-mpc-ess-product-brief.pdf>

# Performance – E980 vs. E1080



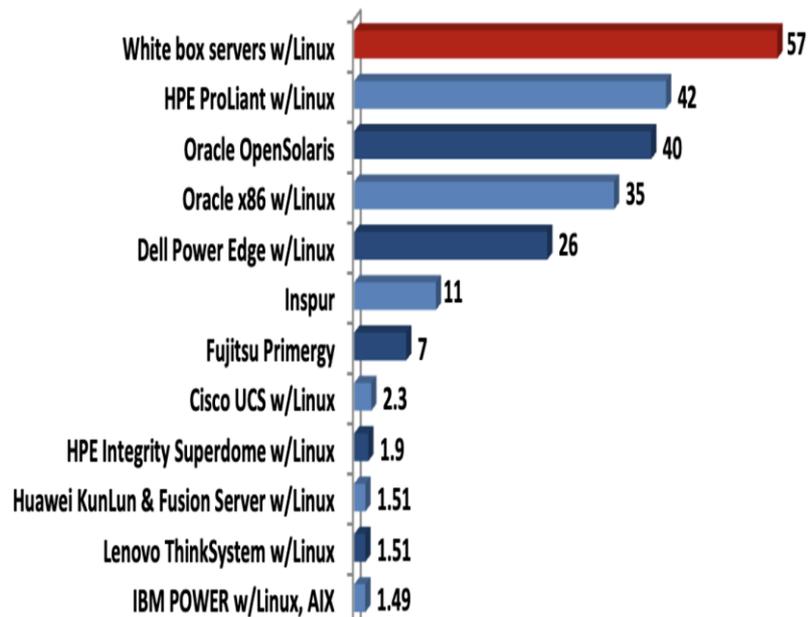
IBM Power ranked number 1  
in every major reliability category by ITIC  
for the 13<sup>th</sup> straight year

[Link to the ITIC report](#)

*“The Power9 and latest Power10 models deliver a minimum of five and six nines availability/uptime.”*



Unplanned Downtime Server Hardware Per Minute/per server



# IBM Power Security

*Design, architecture, and integration*

Security is architected into Power for all types of threats: traditional, new, and emerging

- Processor
- Network
- Applications
- Firmware
- Operating systems
- Middleware
- Hypervisors
- Containers
- AI
- Management



## Base Platform Security & Integrity

Continuously protect platform integrity across main processor, service processor and peripherals

## End to End Hybrid Cloud Security

Offer all platform capabilities with the highest level of security from enterprise through Cloud

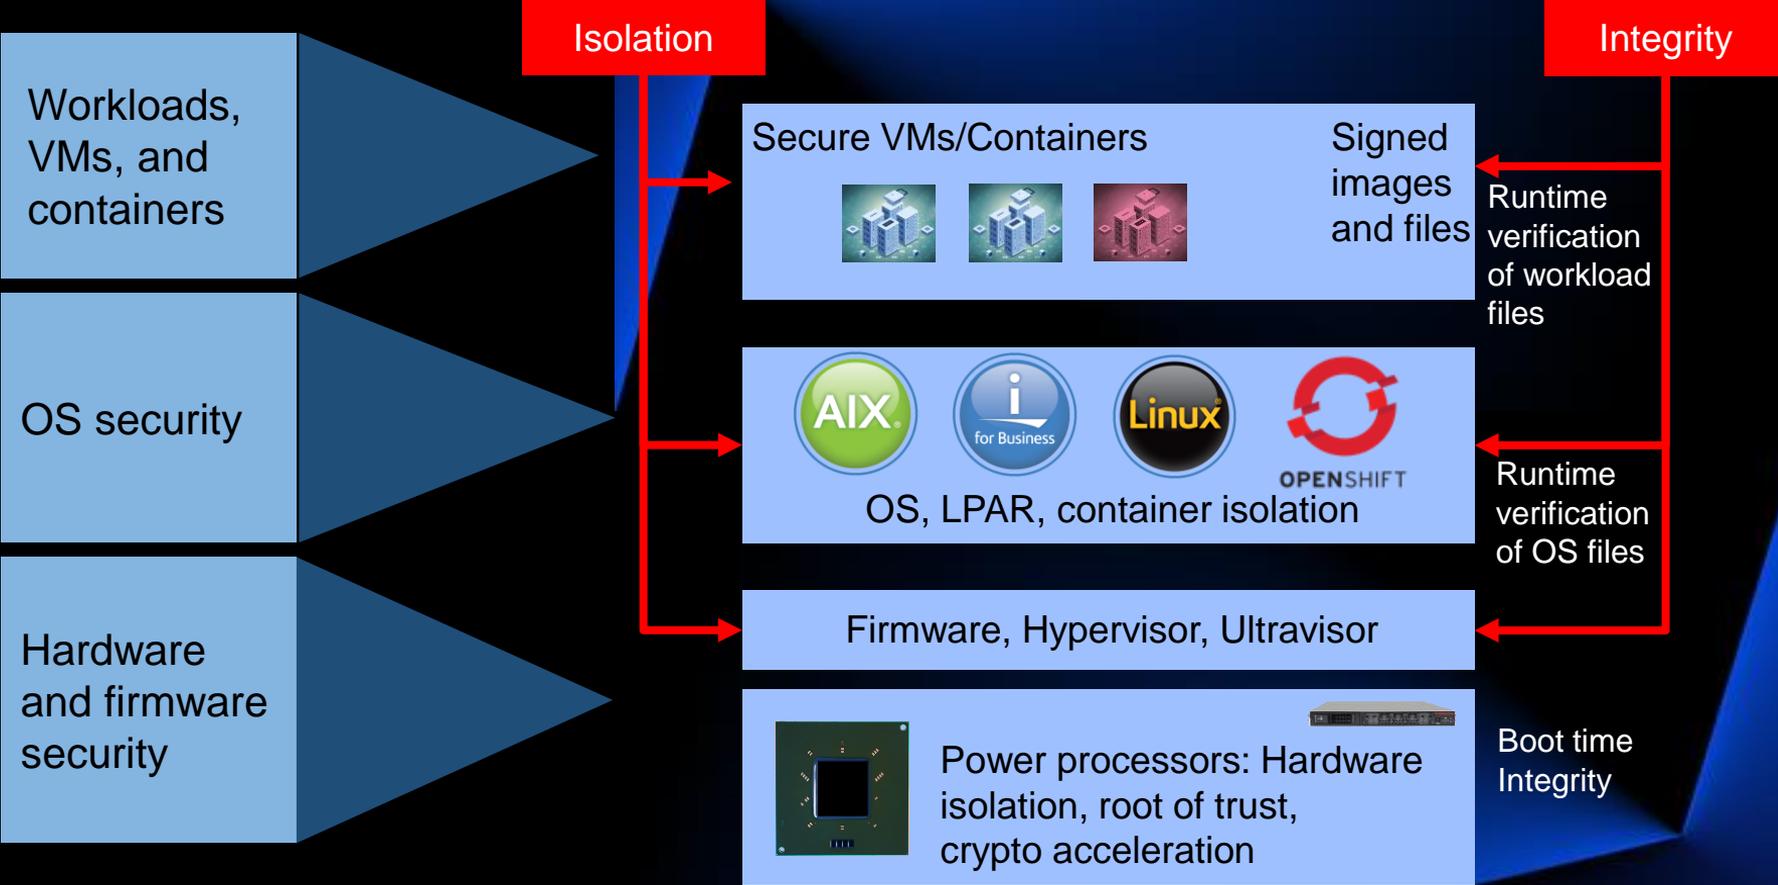
## Workload Security Enablement

Provide features to secure client workloads: HW, firmware, and OS support for isolation, integrity, encryption, event monitoring, ...

## Simplified Security Management

Automated security management to simplify security operations and compliance: patching, integrity monitoring, health checking, ...

# System level security



# Implementing Encryption can be Complex

Comprehensive data protection requires a huge investment to deploy point solutions and / or enable encryption directly in the application



Organizations struggle with questions such as:

## Where

should data  
encryption occur?

## What

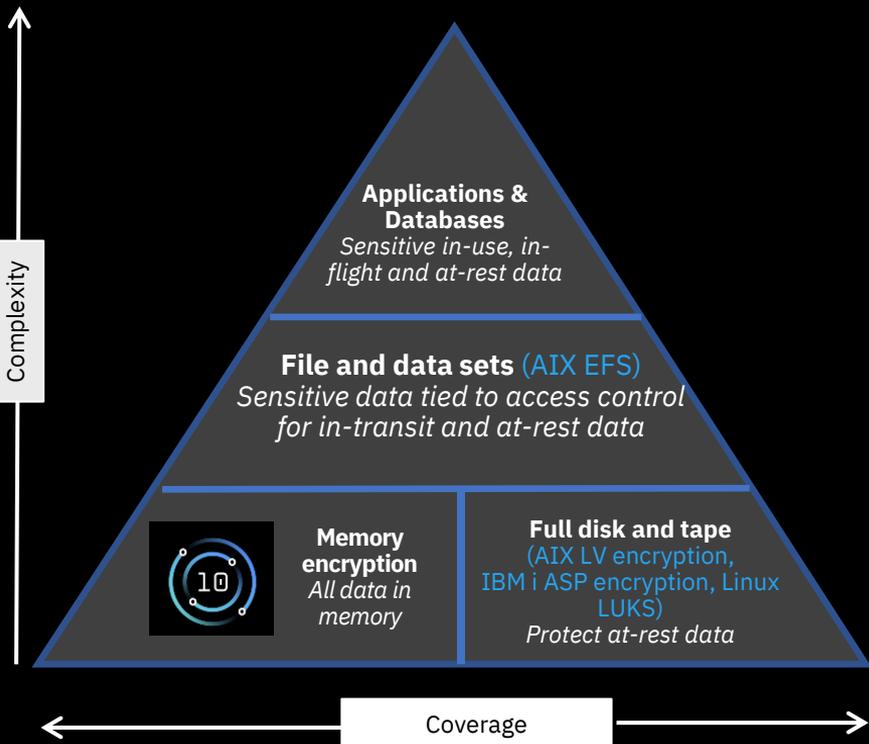
data should be  
encrypted?

## Who

 is responsible for encryption?

# Protect Data:

End to end security with full stack encryption



Transparent memory encryption with:

- No additional management setup
- No performance impact

Blazing fast hardware-accelerated encryption compared to Power9

- 4X crypto engines in every core
- 2.5X faster AES crypto performance per core\*

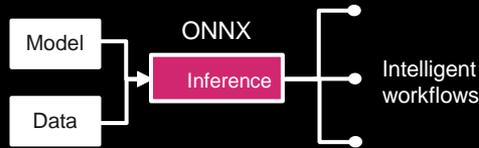
Stay ahead of current and future threats with support for:

- Quantum-safe cryptography
- Fully homomorphic encryption

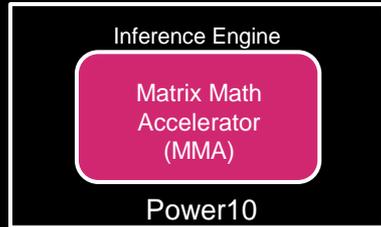
\*AES-256 in both GCM and XTS modes runs about 2.5 times faster per core than comparable Power9 systems according to preliminary measurements obtained on RHEL Linux 8.4 and the OpenSSL1.1.1g library

# Streamline insights and automation with *In-core AI inferencing and machine learning*

Bring your own models and run inference where your operational data resides



AIX IBM i Linux  Red Hat OpenShift



**4** MMA Engines per Core

# 5X

Faster AI inferencing per socket over Power E980\*

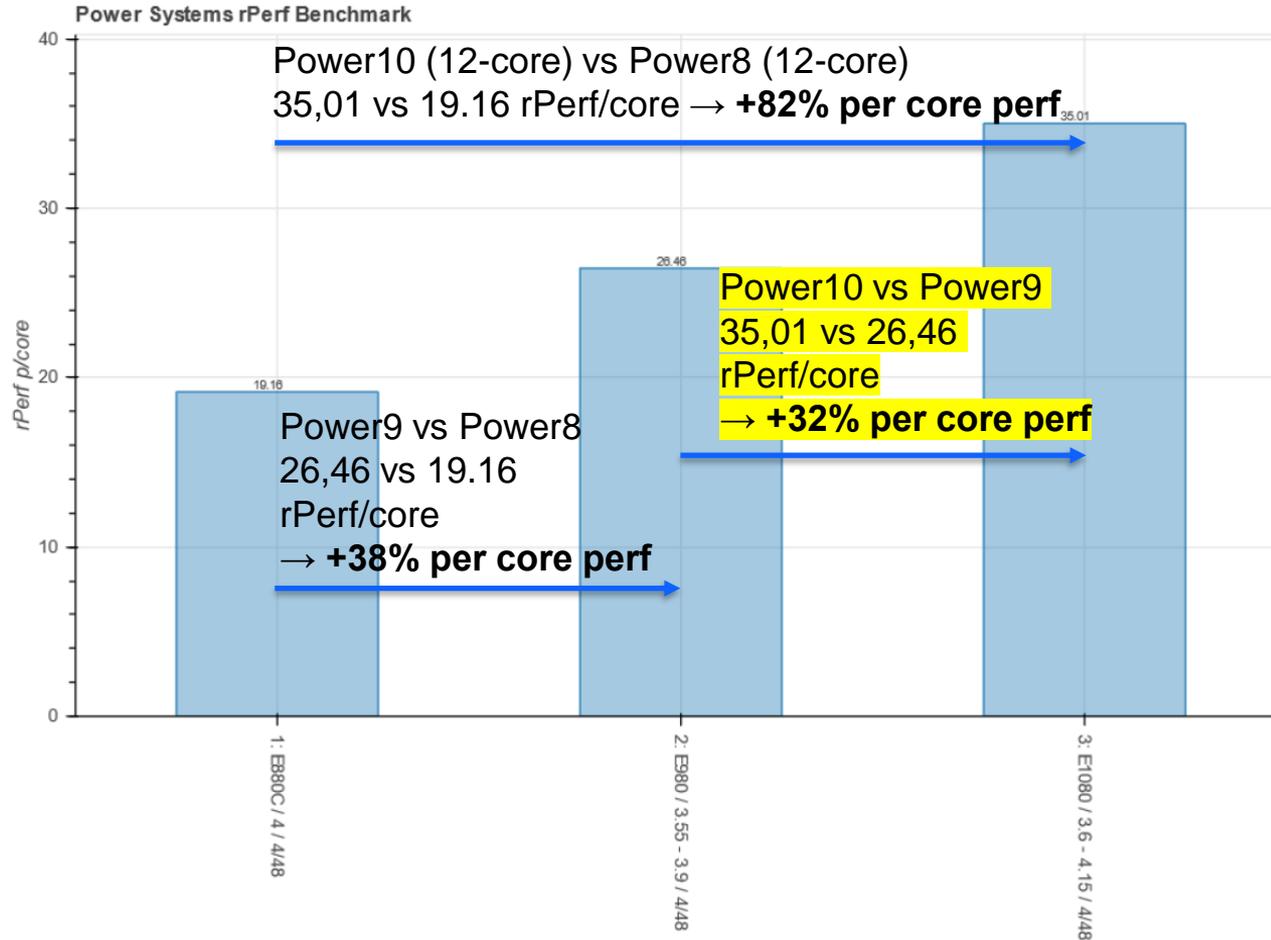
- Perform in-core AI inferencing and ML where the data resides
- Train AI models anywhere, deploy on Power without changes for AI
- Support for popular libraries, AI frameworks and ONNX runtime
- Provides alternative to using separate GPU systems

\* 5x improvement in per socket inferencing throughput for large size 32b floating point inferencing models from Power9 E980 (12-core modules) to Power10 E1080 (15-core modules). Based on IBM testing using PyTorch, OpenBLAS on the same BERT Large with SQuAD v1.1 data set

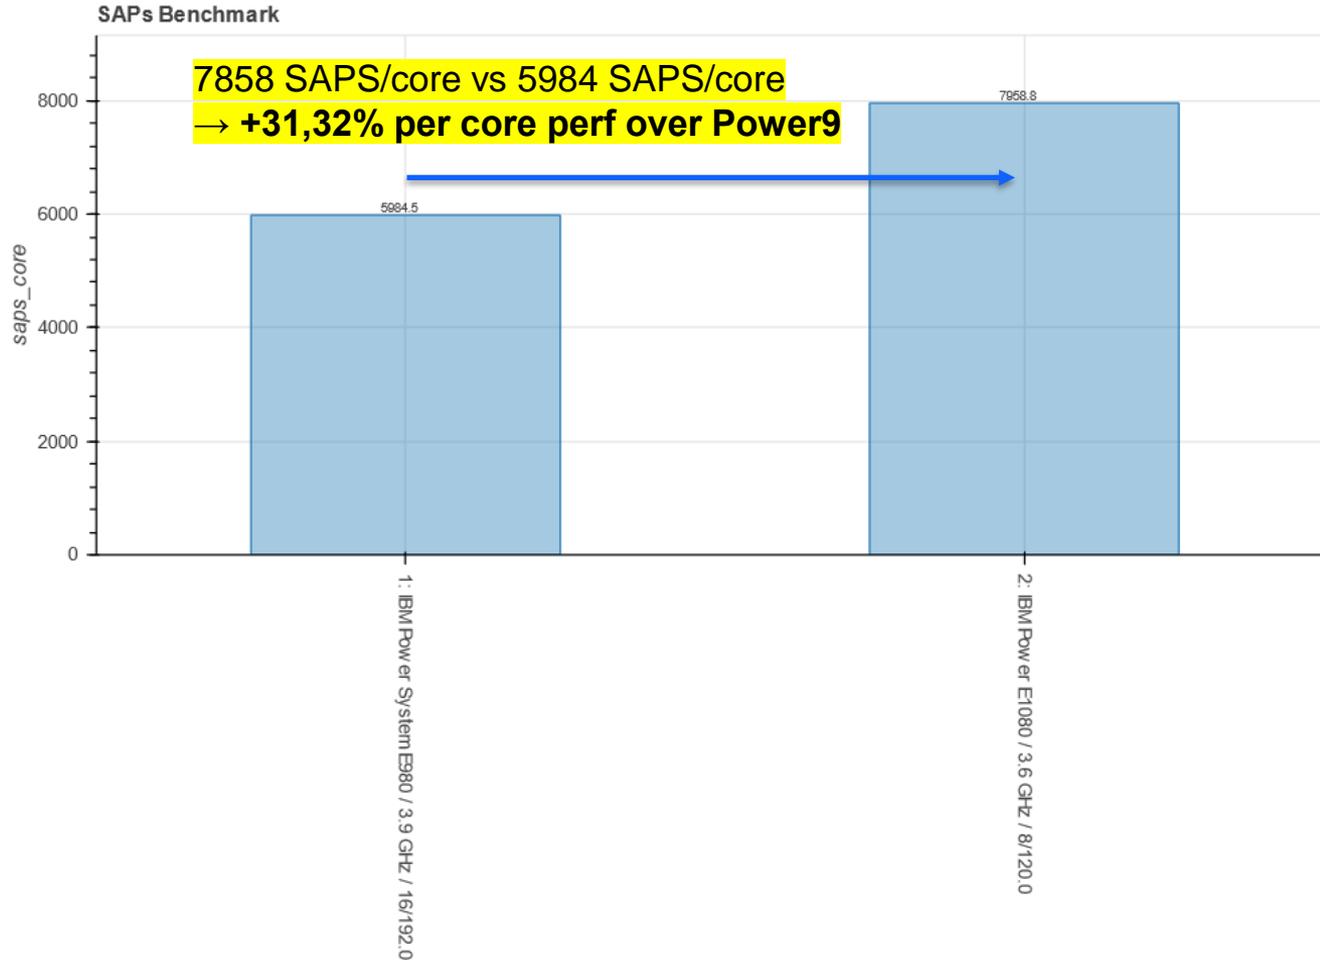
# Using Power10's Superfast AI With Event Streaming



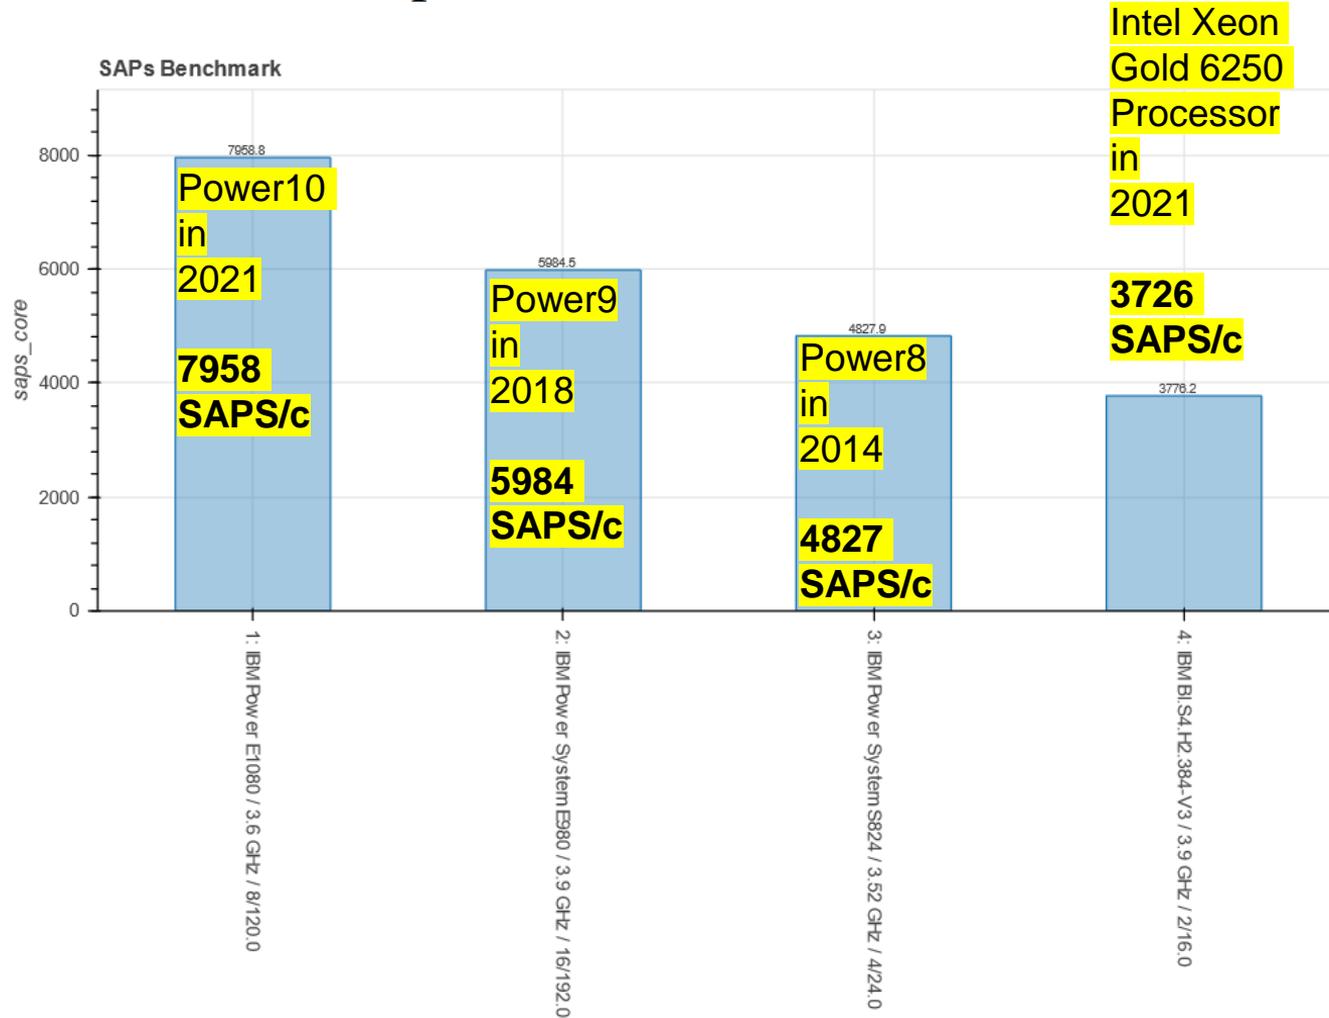
## Benchmark's Graphs



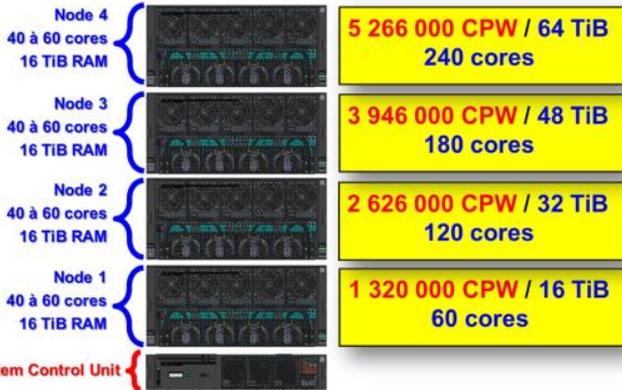
# Benchmark's Graphs



# Benchmark's Graphs



## Power E1080



4x10c E980 = 611300

→ 15283 CPW/core

4x10c E1080 = 92200

→ 23050 CPW/core

**+ 50,82% more  
performance per core**

4x12c E980 = 687500

→ 14323 CPW/core

4x12c E1080 =  
1 129000

→ 23521 CPW/core

**+ 64,22% more  
performance per core**

### Power E980

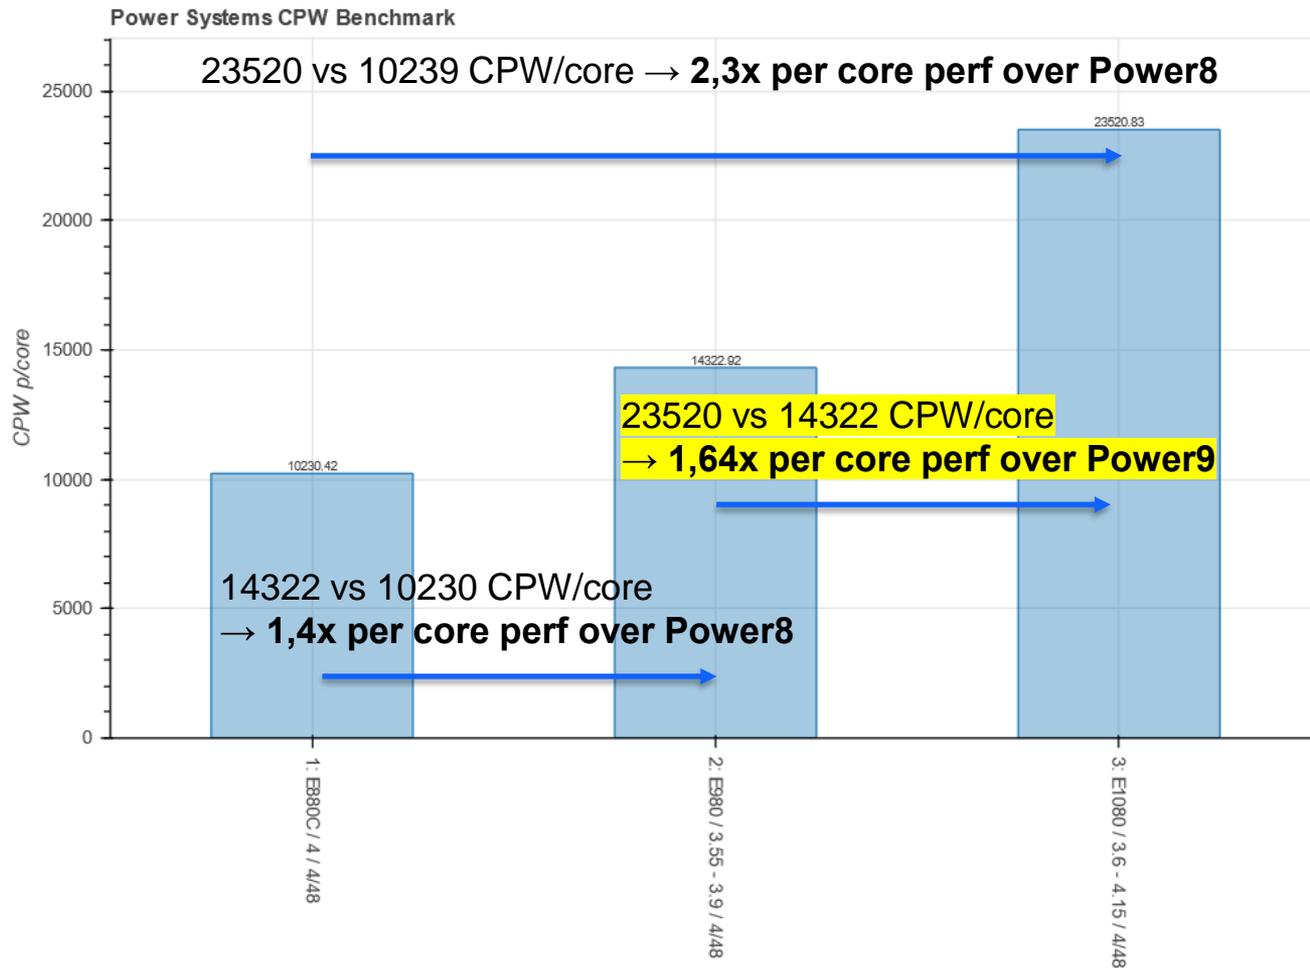
Modèle	Fréquence GHz	# cores	# sockets	CPW	CPW par core
6-core	3,58 - 3,90	24	4	381 000	15 875
		48	8	759 000	15 813
		72	12	1 140 000	15 833
		96	16	1 522 000	15 854
8-core	3,90 - 4,00	32	4	508 900	15 903
		64	8	1 012 000	15 813
		96	12	1 521 000	15 844
		128	16	2 030 000	15 859
10-core	3,70 - 3,90	40	4	611 300	15 283
		80	8	1 216 000	15 200
		120	12	1 827 000	15 225
		160	16	2 439 000	15 244
11-core	3,58 - 3,90	44	4	639 000	14 523
		88	8	1 271 000	14 443
		132	12	1 910 000	14 470
		176	16	2 549 000	14 483
12-core	3,55 - 3,90	48	4	687 500	14 323
		96	8	1 368 000	14 250
		144	12	2 055 600	14 275
		192	16	2 743 000	14 286

### Power E1080

Modèle	Fréquence GHz	# cores	# sockets	CPW	CPW par core
10-core	3,65 - 3,90	40	4	922 000	23 050
		80	8	1 834 000	22 925
		120	12	2 756 000	22 967
		160	16	3 678 000	22 988
12-core	3,60 - 4,15	48	4	1 129 000	23 521
		96	8	2 246 000	23 396
		144	12	3 375 000	23 438
		192	16	4 504 000	23 458
15-core	3,55 - 4,00	60	4	1 320 000	22 000
		120	8	2 626 000	21 883
		180	12	3 946 000	21 922
		240	16	5 266 000	21 942



# Benchmark's Graphs





common  
EUROPE LUXEMBOURG



IBM | Let's  
Create

## *Les gammes Power9 et Power10 Scale-Out (et midrange)*

# Power10 – Engineered for Agility

<b>Cloud Agility and Economics</b>	<b>Pervasive security and resilience</b>	<b>Streamline insights and automation</b>
------------------------------------	--	---

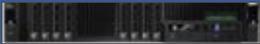
## Sept/Nov 2021 High End

- Highest Performance, most scalable system in industry
- Increased flexibility with Power Private Cloud
- Most scalable and agile SAP HANA



## Summer 2022 Scale Out & Midrange Systems

- > 2X capacity growth
- Power Private Cloud across the line
- Enhanced resilience



# Rappel de la Gamme Power 9



**S914**  
9009-41G

**S922**  
9009-22G

**S924**  
9009-42G

**H922**  
9223-22H

**H924**  
9223-42H

**L922**  
9008-22L

- 1-socket, 4U Rack & Tower
- 4,6,8 cores/ socket
- 16 IS RDIMM slots
- 1To mémoire
- 2 CAPI 2.0 Slots
- Internal RDX Media
- 8 PCIe adapters slots
- 11 NVMe devices

- 1,2-socket, 2U
- 1, 4, 8,10, 11 cores/ socket
- 32 IS RDIMM slots
- 4To mémoire
- 4 CAPI 2.0 Slots
- 11 PCIe adapters slots
- 14 NVMe devices

- 1,2-socket, 4U
- 8,10, 11, 12 cores/ socket
- 32 IS RDIMM slots
- 4To mémoire
- 4 CAPI 2.0 slots
- Internal RDX Media
- 11 PCIe adapters slots
- 14 NVMe devices

- 1,2-socket, 2U
- 4, 8,10 cores/ socket
- 32 IS RDIMM slots
- 4To mémoire
- 4 CAPI 2.0 Slots
- 9 PCIe adapters slots
- 5 NVMe devices
- AIX, IBM i up to 25%

- 1,2-socket, 4U
- 8,10,12 cores/ socket
- 32 IS RDIMM slots
- 4To mémoire
- 4 CAPI 2.0 slots
- Internal RDX Media
- 11 PCIe adapters slots
- 5 NVMe devices
- AIX, IBM i up to 25%

- 1,2-socket, 2U
- 8,10,12 cores/ socket
- 32 IS RDIMM slots
- 4To mémoire
- 4 CAPI 2.0 Slots
- 9 PCIe adapters slots
- 5 NVMe devices



**+ PowerVM**



## Technology Leadership

- L'hyperviseur PowerVM est intégré sans frais de licence, de sorte que chaque charge de travail est virtualisée, mobile et prête pour le cloud.
- Jusqu'à 4 To dans 2 sockets –
- mémoire RDIMMs DDR4 aux normes industrielles
- Ports externes à haut débit 25Gb/s - un par socket
- Des analyses et des algorithmes intégrés à la puce permettent de faire fonctionner POWER9 à une fréquence toujours optimisée.
- Modèles G : 2 fois plus de bande passante et jusqu'à 180 % de capacité NVMe en plus pour une intégration transparente entre on-prem et public cloud avec une architecture PCIe Gen4 complète.

# LA GAMME Power 10

## Systeme d'entrée de gamme

- Jusqu'à 16 cœurs et 2 To de mémoire.
- Jusqu'à 32 % d'augmentation des performances par rapport au prix du Power9
- Serveurs statiques et entièrement actifs
- Exigences techniques réduites et offres compétitives

## Performances et évolutivité accrues

- Jusqu'à 48 cœurs et 8 To de mémoire
- Jusqu'à 50 % d'augmentation des performances par rapport au prix et 1,4 fois plus de performances système par rapport au Power9.
- Consommation dynamique de la capacité avec CUoD et PEP2.0
- Solutions axées sur la valeur et normes techniques plus élevées

### S1014

9105-41B



- 1-socket, 4U
- Jusqu'à 8 cœurs par système
- 4, 8 cœurs / socket SMT8
- 8 emplacements DDIMM
- 1 To de mémoire \* (GA: 512 Go)
- 5 emplacements PCIe FHHL (capacité 4 Gen5)
- 16 baies NVMe U.2
- Max de 102,4 To de stockage interne
- Optionnel RDX Media Bay
- Facteurs de forme de l'armoire et de la tour

### S1022S

9105-22B



- 1,2-socket, 2U
- Jusqu'à 16 cœurs par système
- 4, 8 cœurs / socket SMT8
- 16 emplacements DDIMM
- 2 To de mémoire \* (GA: 1 To)
- 10 emplacements PCIe HHHL (capacité 8 Gen5)
- 8 baies NVMe U.2
- Max de 51,2 To de stockage interne

### S1022 & L1022

9105-22A / 9786-22H



- 1,2-socket, 2U
- Jusqu'à 40 cœurs par système
- 12, 16, 20 cœurs SMT8 / socket
- 32 emplacements DDIMM
- 4 To de mémoire \* (GA: 2 To)
- 10 emplacements PCIe HHHL (capacité 8 Gen5)
- 8 baies NVMe U.2
- Max de 51,2 To de stockage interne
- L1022: max 25% des cœurs avec d'autres systèmes d'exploitation

### S1024 & L1024

9105-42A / 9786-42H



- 1,2-socket, 4U
- Jusqu'à 48 cœurs par système
- 12, 16, 24 cœurs SMT8 / socket
- 32 emplacements DDIMM
- 8 To de mémoire \* (GA: 2 To)
- 10 emplacements PCIe FHHL (8 Gen5)
- 16 baies NVMe U.2
- Max de 102,4 To de stockage interne
- Optionnel RDX Media Bay
- L1024: max 25% des cœurs avec d'autres systèmes d'exploitation

AIX

IBM i

Linux



Tous les serveurs basés **PowerVM**

\* Support prévu après la GA

# POWER10 E1050 4-Socket Server Highlights\*

- 4U Server - 19" Rack Enclosure
- POWER10 Enterprise DCM processor (12-core, 18-core, 24-core processor options)
  - Min # of cores that must be activated is one socket's worth
- 2, 3 or 4 Processor Socket offerings
  - Minimum configuration is 2 Processor Sockets populated
- Processor DCMs at 32Gbs X-bus and efficient 1-Hop processor fabric interconnect
- Up to 16TB Total DDR4 DDIMMs – up to 4TB per processor
  - 4U DDIMMs: 32GB, 64GB, 128GB, 256GB
  - Minimum configuration – 4 DDIMMs per socket, 50% of installed memory must be active
- 11 PCIe Gen4/5 slots (3 Gen4 only)
- 6 High Speed OpenCAPI acceleration ports
- Internal Storage 10 NVMe 7mm/15mm GEN4 Drives
- Enhanced DC-DC Voltage Regulator Redundancy
- Full Fan Concurrent Maintenance
- I/O Expansion and Storage drawers support
- Seamless Growth with Capacity on Demand and PEP 2.0 for Processor and Memory
- Active Memory Mirroring for Hypervisor (optional)
- Built-in PowerVM virtualization
- Power to Cloud Rewards
- IBM Cloud Management Console
- 3 Years Warranty Service Bundled<sup>1</sup>



PowerVC  
PowerVM

	1socket	2socket	3socket	4socket
E1050	-	48	72	96
	-	36	54	72
	-	24	36	48

\* IBM's statements regarding its plans, directions and intent are subject to change or withdrawal without notice at IBM's sole discretion

1. Select Geographies Only, may include Power Expert Care and Proactive Support options

# PROCESSEUR P10 et GROUPE PROCESSEUR IBM i

S1014 Code dispositif	Processeur Max SMT8	Standard Plage de fréquences	IBM i Groupe P
EPG2	8 coeurs	3,00 à 3,90 GHz	P10
EPG0	4 coeurs	3,00 à 3,90 GHz	P05

S1022S Code dispositif	Processeur Max SMT8	Nombre maximal de coeurs de système	Standard Plage de fréquences	IBM i Groupe P
EPGQ	8 coeurs	16 coeurs	3,00 à 3,90 GHz	P10
EPGR	4 coeurs	4 coeurs	3,00 à 3,90 GHz	N/A

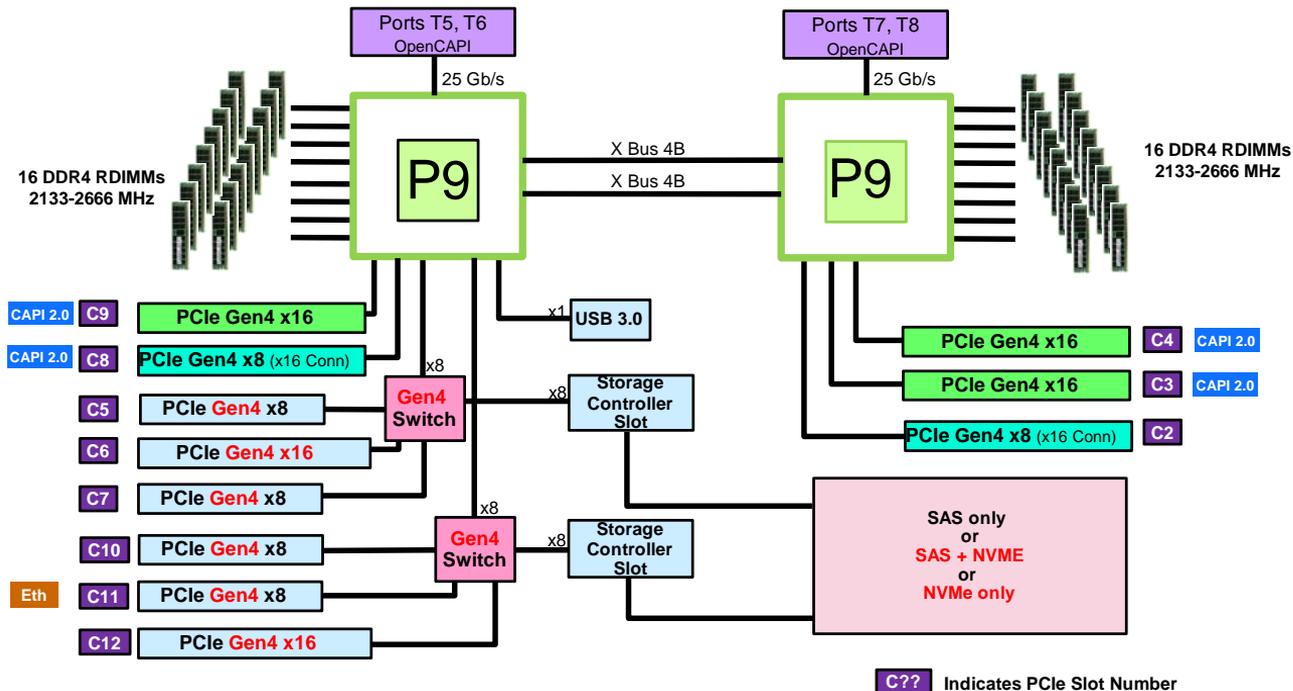
Sur les S1022/S1022s, plusieurs partitions IBM i sont prises en charge, et chaque partition peut avoir jusqu'à quatre coeurs (VIOS requis)  
4-cores non pris en charge pour IBM i dans les S1022s

S1022 Code dispositif	L1022 Code dispositif	Processeur Max SMT8	Nombre maximal de coeurs de système	Standard Plage de fréquences	IBM i Groupe P	
EPGA	EPGH	20 coeurs	40 coeurs	2,45 à 3,90 GHz	P10	Offre de débit maximale
EPG8	EPGF	16 coeurs	32 coeurs	2,75 à 4,00 GHz	P10	
EPG9	EPGG	12 coeurs	24 coeurs	2,90 à 4,00 GHz	P10	Offre de performances maximales par cœur/thread

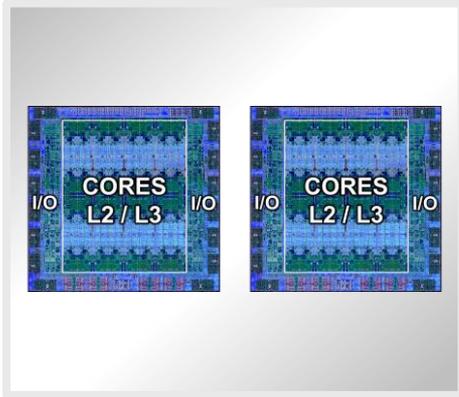
S1024 Code dispositif	L1024 Code dispositif	Processeur Max SMT8	Nombre maximal de coeurs de système	Standard Plage de fréquences	IBM i Groupe P	
EPGD	EPGK	24 coeurs	48 coeurs	2,75 à 3,90 GHz	P30	Offre de débit maximale
EPGC	EPGJ	16 coeurs	32 coeurs	3,10 à 4,00 GHz	P20	
EPGM	EPGN	12 coeurs	24 coeurs	3,40 à 4,00 GHz	P20	Offre de performances maximales par cœur/thread

# S924 GA4 (9009-42G) Refresh System Topology

Les systèmes S924 2S ont 80 voies PCIe G4 fonctionnant à un débit max de 16 Gbps. La bande passante théorique totale maximale en duplex intégral des E/S du S924 2S est de 320 Go/s.

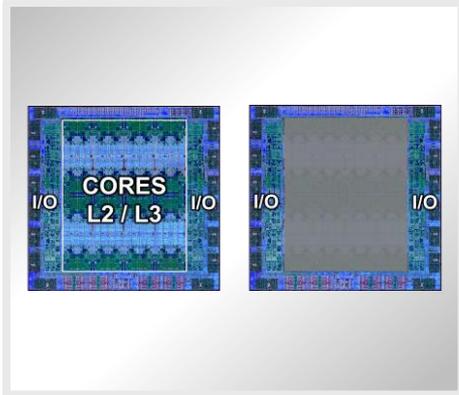


# Processeurs Power10 Scale Out: DCM et eSCM



## DCM - Module à double socket (Dual-Chip Module)

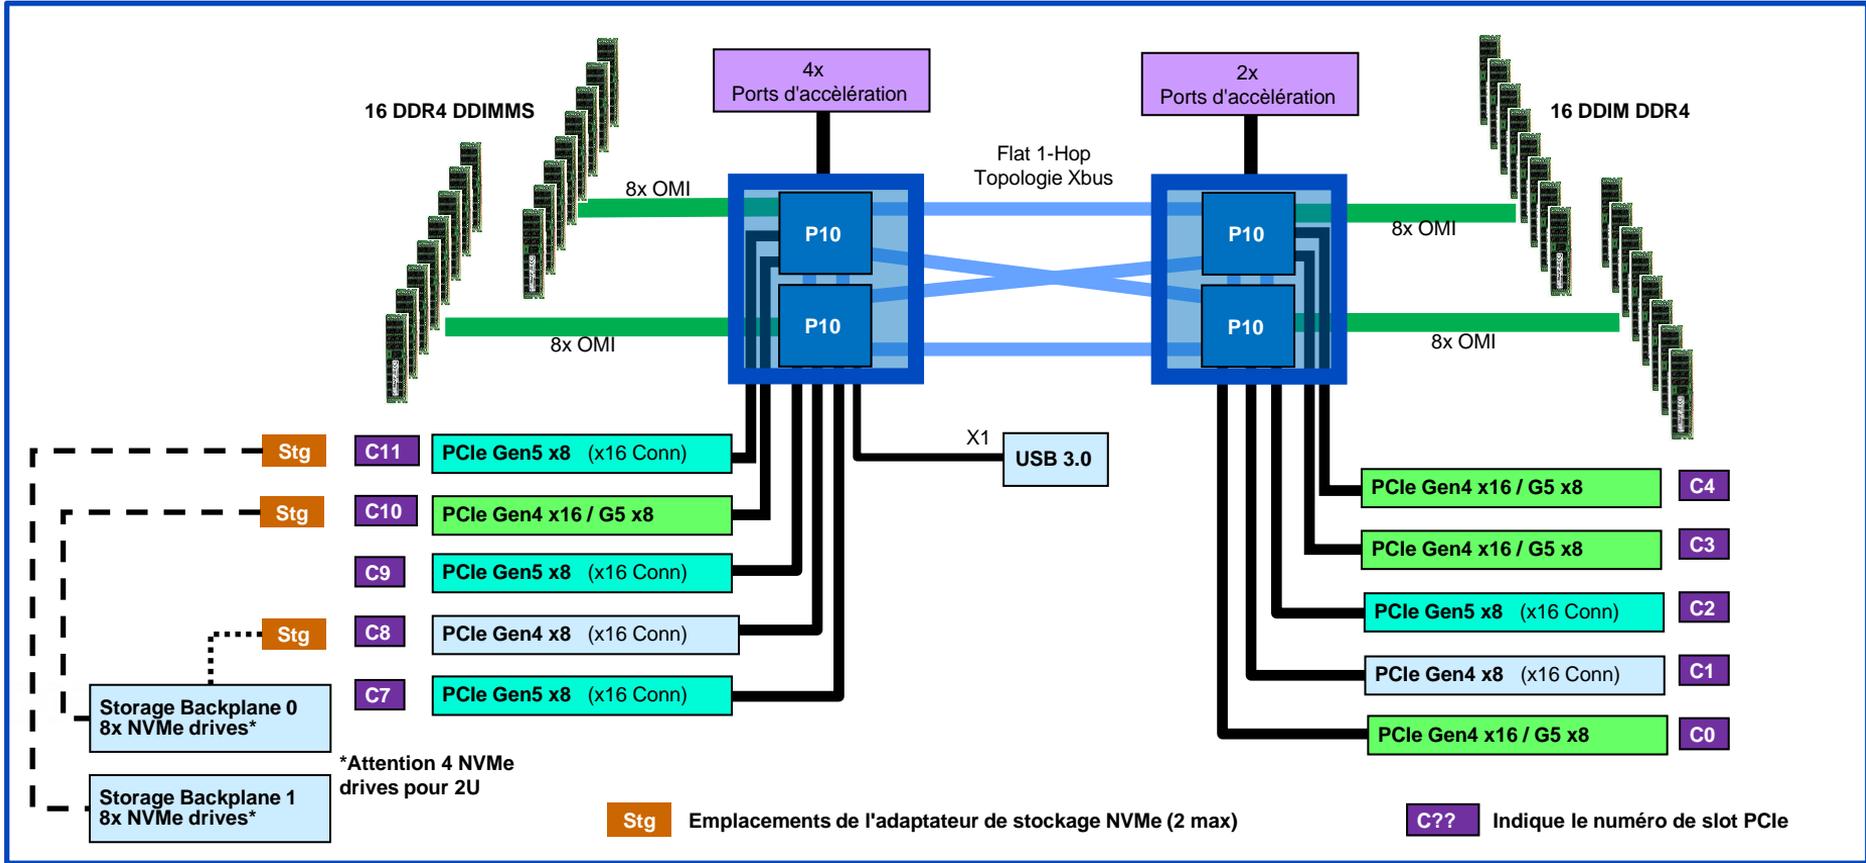
- Permet une gamme complète d'offres
- Fournit jusqu'à 24 cœurs par socket
- 16 canaux de mémoire OMI (8 par puce)
- Utilisé sur S1024 et S1022
- Coud disponible sur l'ensemble des modèles S1022 et S1024



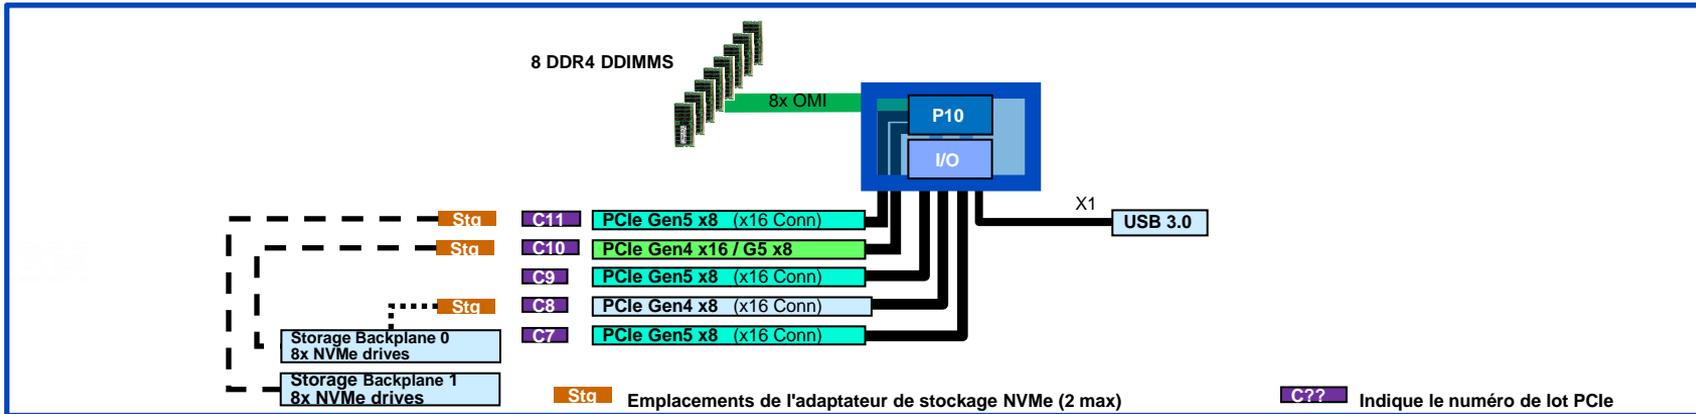
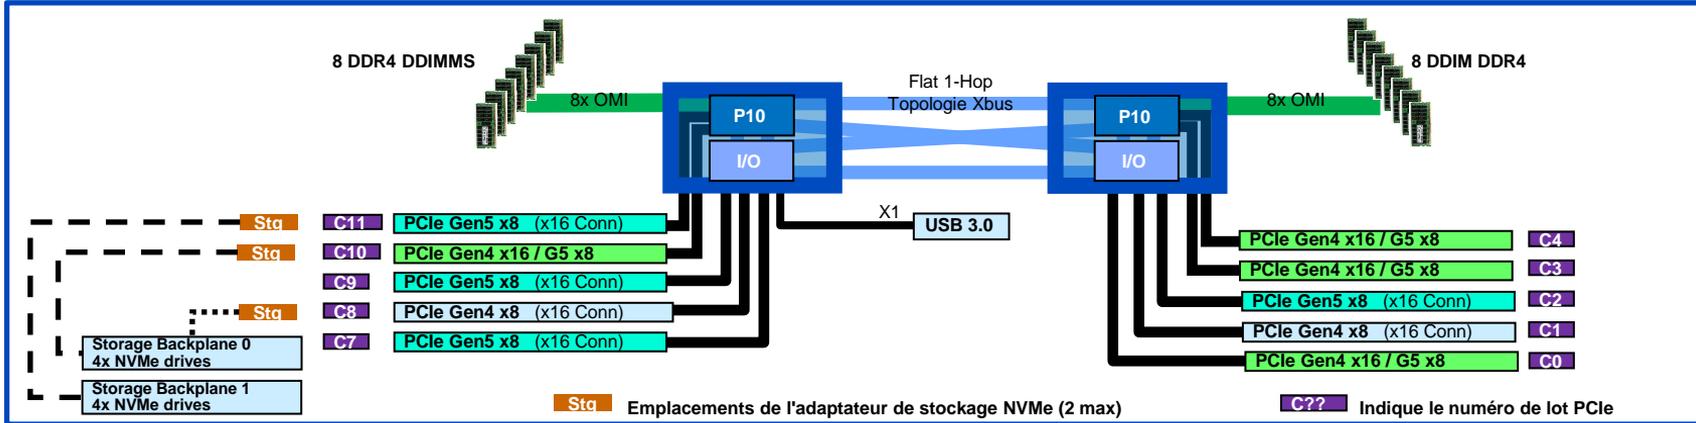
## eSCM – Module monosocket d'entrée (entry Single-Chip Module)

- Un Processeur avec calcul, mémoire et E/S
- Second Processeur ajoute E/S seulement (emplacements PCIe)
- Permet des offres de processeurs à faible coût
- Juqu'à 8 cœurs par socket
- 8 emplacements mémoires OMI
- Utilisé sur S1022s et S1014
- Deconfiguration de core possible

# Architecture des serveurs 2 socket 4U et 2U DCM



# Architecture des serveur S1022s et S1014 eSCM



# Options de mémoire des serveurs 2U ET 4U

- jusqu'à 409 Go/s de bande passante mémoire de pointe par socket
- DDIMM à base de DDR4
- 8/16 emplacements DDIMM par socket, 32 emplacements DDIMM au total
- Capacité mémoire maximale de 8 To (2 To à la GA initiale) **64 GO Max pour le 4 core IBM i S1014**
- **Règles du plug-in DDIMM par socket**
  - La configuration minimale est de 2 DDIMMs par socket
  - Toutes les barrettes DIMM derrière un socket doivent être du même type
  - Les barrettes DIMM 2U et 4U ne peuvent pas être mélangées dans le même système
- **NOUVEAU ! Prise en charge de la fonction Active Memory Mirroring : mise en miroir de la mémoire de l'hyperviseur pour assurer la résilience face aux erreurs de mémoire non corrigibles.**

S1014 / S1024 / S1022 / S1022S Code dispositif	L1024 / L1022 Code dispositif	Barrette DIMM	DRAM Vitesse	Bande passante de la mémoire	Disponibilités
EM6N	EM7N	16GB 2U DDIMM	3200 MHz	409 Go / s	
EM6W	EM7W	32GB 2U DDIMM	3200 MHz	409 Go / s	
EM6X	EM7X	64GB 2U DDIMM	3200 MHz	409 Go / s	
<b>EM6Y</b>	<b>EM7Y</b>	<b>128 GB 2U DDIMM</b>	<b>2666 MHz</b>	<b>341 Go / s</b>	4Q2022
<b>EM78</b>	<b>EM7F</b>	<b>256GB 4U DDIMM</b>	<b>2933 MHz</b>	<b>375 Go / s</b>	4Q2022
EM8G	EM8G	Active Memory Mirroring (AMM) <b>sauf S1014</b>			Fonctionnalité complémentaire

# Options de stockage des serveurs 2U ET 4U

## Options de stockage interne

- Toutes les unités de stockage internes sont PCIe NVMe (pas de SAS)
- Les emplacements PCIe généraux (C10/C8, C11) prennent en charge la carte JBOF NVMe et sont câblés sur le fond de panier NVMe
- Nombre maximal de 16 unités NVMe U.2 prises en charge (8 pour les 2U)
- **carte EJ2A pour connectique EMX0**

**Pour le S1014 4 core IBM i Stockage interne limité à 6,4 TB (3,2 TB miroiré) supportant 2-8 NVMe PCIe devices incluant les modules U2 et AIC.**

Options de stockage interne S1024 / L1024

FC	Modele	Description
EJ1Y	S1014 /S1024	Carte NVMe JBOF avec fond de panier U.2 8-Pak
EJ1X	S1022/S1022S	Carte NVMe JBOF avec fond de panier U.2 4-Pak
EUA0	S1014 /S1024	Station d'accueil RDX

Options de stockage externe S1024 / L1024

FC / MTM	Modele	Description
ESLS	S1014 /S1024	19 Tiroir d'extension de disque 24 SFF Gen2-Baies
EUA5	S1022/S1022S	DVD USB w / Câble
7226-1U3	S1014 /S1024	Tiroir de support 19 avec 2 baies
TBD		Tiroir NVMe 24 U.2 Bays (Future GA)

## Présentation des supports pris en charge

- **Unités NVMe U.2**
  - 0.8TB 4K U.2 7mm 18W PCIe Gen4 Mainstream
  - 0.8TB 4K U.2 15mm 18W PCIe Gen4 Enterprise Class
  - 1.6TB 4K U.2 15mm 18W PCIe Gen4 Enterprise Class
  - 3.2TB 4K U.2 15mm 18W PCIe Gen4 Enterprise Class
  - 6.4TB 4K U.2 15mm 18W PCIe Gen4 Enterprise Class
- **Cartouche de disque RDX → S1014 et S1024**
  - Cartouche de disque 320Go (EU08)
  - Cartouche de disque 500 Go (1107)
  - Cartouche de disque 1 To (EU01)
  - Cartouche de disque 1,5 To (UE15)
  - Cartouche de disque 2To (EU2T)



16 NVMe U.2 15 mm Bays/



# Tiroir d'extension PCIe sur serveur 2U et 4U

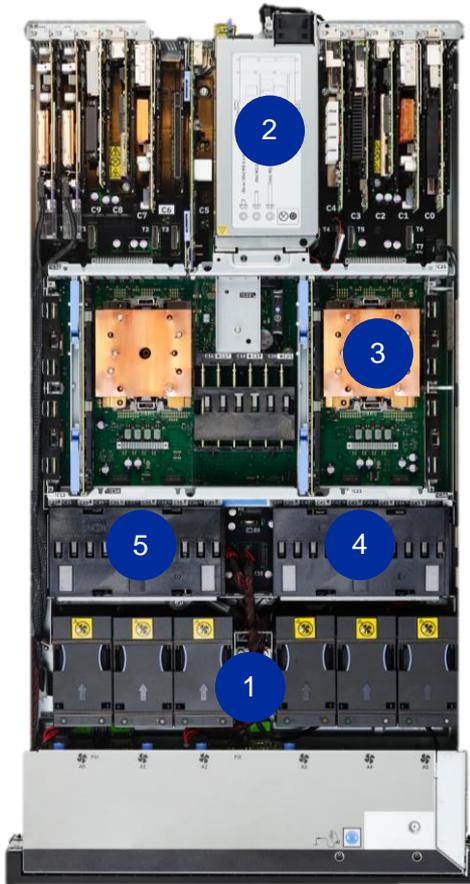
Nombre CPU	Modèle	Nombre maximal de tiroirs d'extension d'E-S (EMX0)	Nombre maximal de modules d'échange d'E-S (EMXH)	Nombre total de slots PCIe
1	S1024 /S1022 / S1022S	1	1	10
2	S1024 /S1022 / S1022S	2	4	30

Nombre d'UC	Modèle	Nombre maximal de tiroirs d'extension d'E-S (EMX0)	Nombre maximal de modules d'échange d'E-S (EMXH)	Nombre total de slots PCIe
1	S1014	0.5	1	10

**ATTENTION EMX0 NON SUPPORTE SUR S1014 4 CORE IBM I**

\*

# IBM Power redundancy and reliability built-in



- 1 Redundant system cooling with multiple fans & concurrent maintenance
- 2 Redundant/bulk power - even if 2 power supplies fail system keeps operating on 4U systems
- 3 Processors include FFDC, instruction retry, core contained checkstops, power and cooling monitoring
- 4 Chipkill technology with advanced ECC protects from memory chip failure - plus spare
- 5 Active Memory Mirroring for hypervisor feature to allow system to read from mirror if a memory read error is detected

# New Power10 Memory Technologies

## ***New Differential DIMM Technology (DDIMM)***

*Moving from Industry Standard to Differential DIMMs in Power10*

## ***New Open Memory Interface (OMI)***

*Providing higher bandwidth and flexibility for future memory technologies*



2X better memory RAS  
than industry standard  
DIMMs<sup>1</sup>

2.4X higher memory  
bandwidth than  
scalable x86  
processors

DDR4 running at up  
to 3200 Mbps data  
rate provides 409  
GB/s peak memory  
bandwidth per socket

Transparent memory  
encryption with no  
additional  
management setup  
and no performance  
impact

Chipkill technology  
with advanced ECC  
protects from memory  
chip failure - plus spare

Active Memory  
Mirroring (AMM)  
feature supported -  
Mirrors hypervisor  
memory to provide  
resiliency from  
uncorrectable memory  
errors

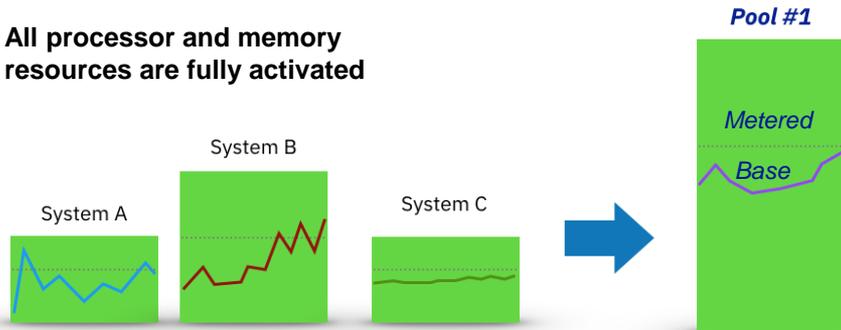
<sup>1</sup> Based on IBM's internal analysis of the IBM product failure rate of DDIMMS vs Industry Standard-DIMMS

# Power private cloud with shared utility capacity

Deploy shared utility capacity across a pool of Power E1080/E980 systems, Power E1050/E950 systems or S922/S924, S1022/S1024 systems

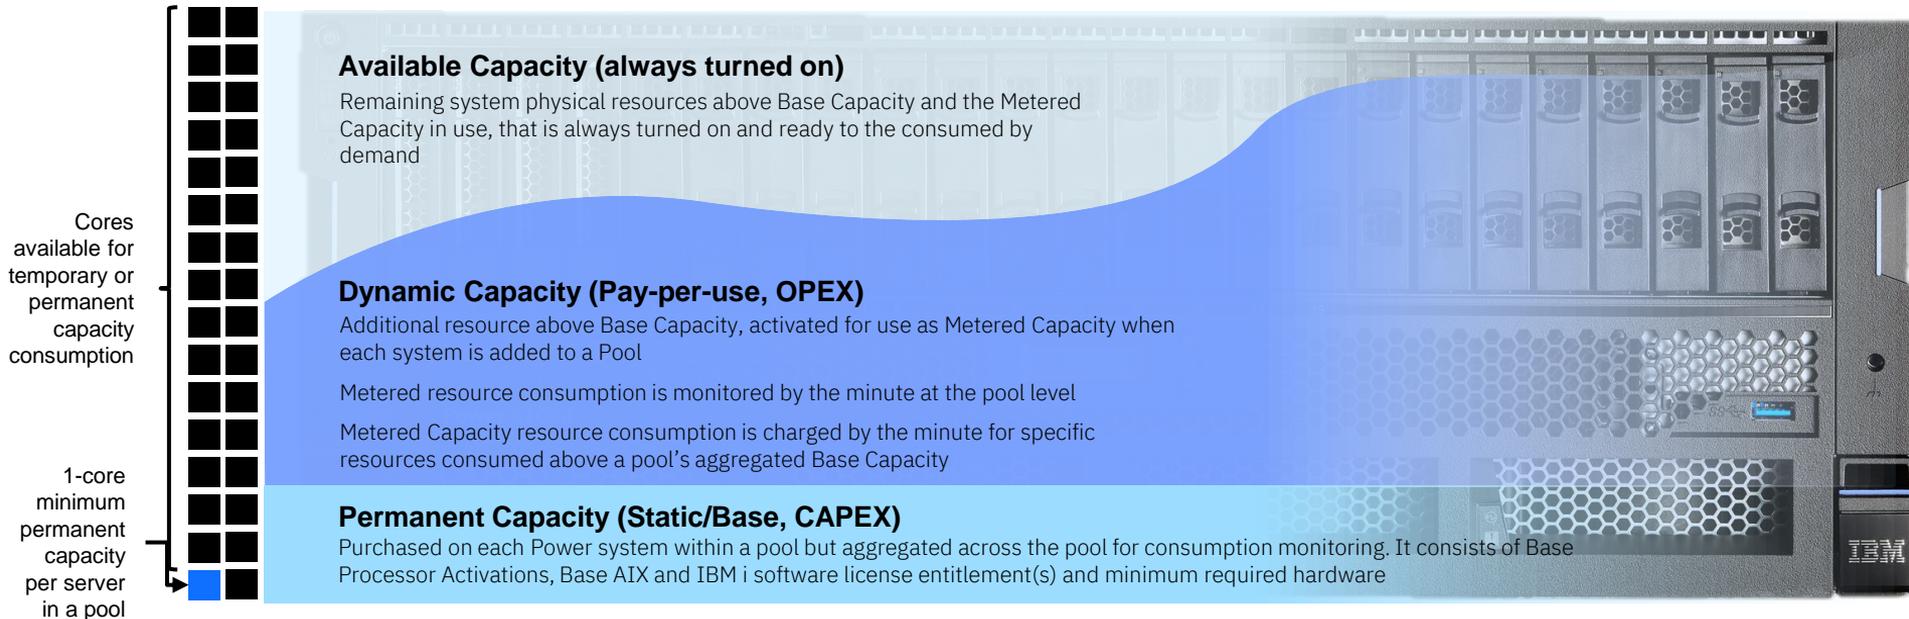
IBM Cloud Management Console with HMC automatically meters any resource use that exceeds the pool's base capacity, and debits minutes real-time against capacity credits on account

All processor and memory resources are fully activated



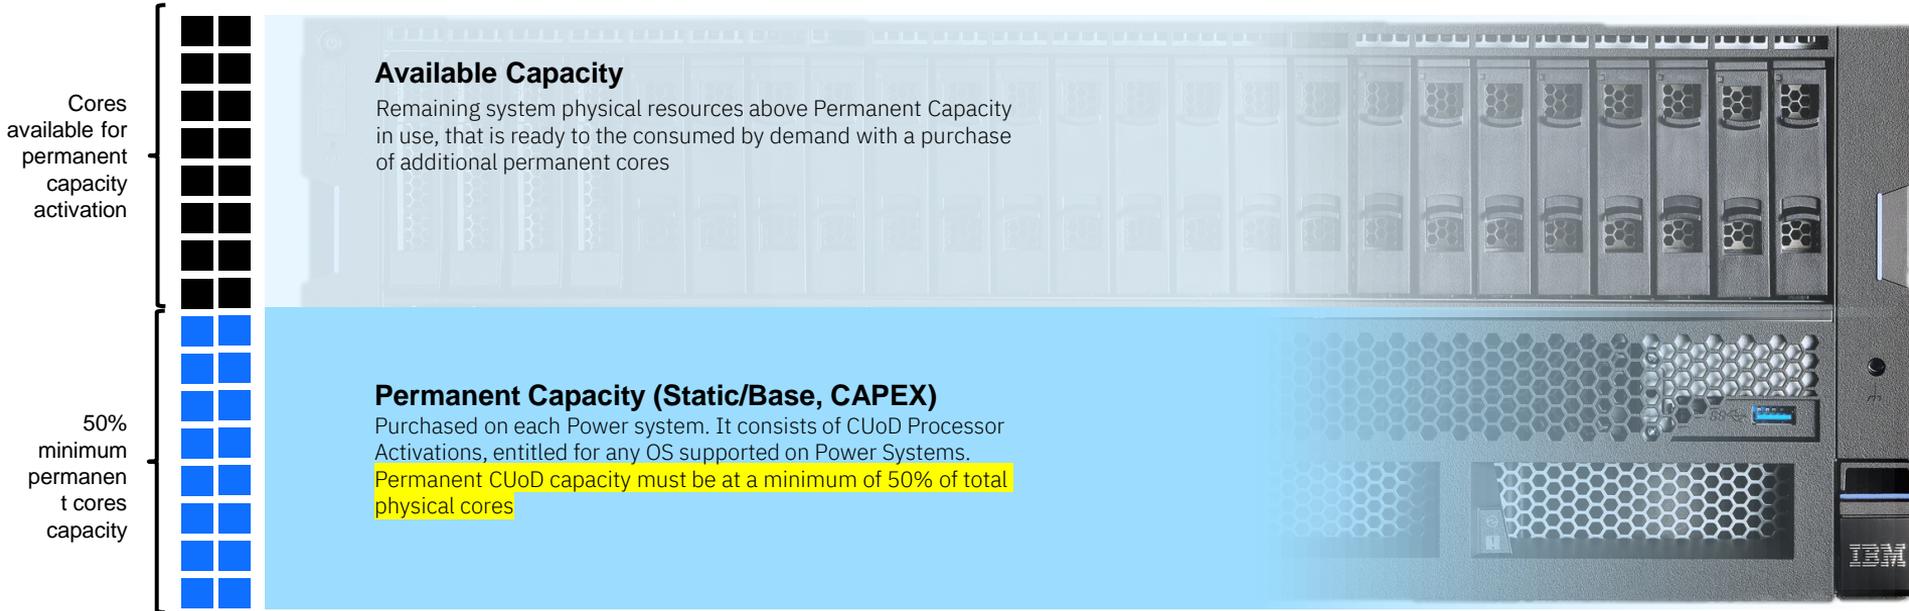
# Capacity Consumption within an Enterprise Pool

- No base monthly fees: **pay for only what use with metering by the minute, pre-paid or post-paid**
- Share resources across Power9 and Power10 systems in the same pool
- Mix S1022 and S1024 servers in the same pool
  - Except for IBM i P30 tier
- Buy 1, Move 1 base core activations migration from Power9 to Power10 servers



# Introducing Capacity Upgrade-on-Demand (CUoD)

- CUoD allows customers the flexibility to get what they need and add permanent per-core capacity in the future as needed
- Benefit from having all the memory, PCIe slots, and system resources w/o concerns about processor availability
- No CMC or HMC required, and no PowerVM swma charges on non-activated cores
- No Elastic CoD available (ECoD)





common  
EUROPE LUXEMBOURG



IBM | Let's  
Create

*Performance*

# Test de performances CPW

projection CPW Pour les systèmes basés sur SMT8

Power S1022 CPW Projections (SMT8)				
Offres	Nombre maximal de coeurs	1-socket	2 sockets	CPW/CORE
S1022 / 20c	40 coeurs	N/A	1 047 000	26 175
S1022 / 16c	32 coeurs	N/A	850 400	26 575
S1022 / 12c	24 coeurs	252 183	504 367	21 015

Power S1024 CPW Projections (SMT8)				
Offres	Nombre maximal de coeurs	1-socket	2 sockets	CPW/CORE
S1024 / 24c	48 coeurs	N/A	947 500	19 740
S1024 / 16c	32 coeurs	N/A	725 000	22 656
S1024 / 12c	24 coeurs	312 500	585 100	24 379

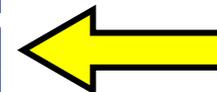
Projections CPW Power S1022s (SMT8)				
Offres	Nombre maximal de coeurs	1-socket	2 sockets	CPW/CORE
S1022s / 8c	16 coeurs	212 600	425 200	26 575
S1022s / 4c	4 coeurs	N/A	N/A	N/A

Power S1014 CPW Projections (SMT8)			
Offres	Nombre maximal de coeurs	1-socket	CPW/CORE
S1014 / 8c	8 coeurs	205 300	25 663
S1014 / 4c	4 coeurs	106 300	26 575

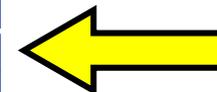
D'après le document IBM Power Performance Report du 12/07/2022

# Performances IBM i en CPW Evolution de P8 à P10

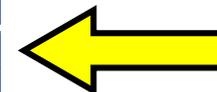
POWER10 with CVE					POWER9 with CVE					% evolution
Modèles	Cores	GHz	CPW Total	CPW /Core	Modèles	Cores	GHz	CPW Total	CPW /Core	GAIN max (%) P9 =>P10
<b>S1022 8/16</b>	4	3.0-3.9	106 300	26 575	<b>S922 8/16</b>	4	3.4 - 3.9	68 000	17 000	56%
<b>S1014</b>	4	2.3 - 3.8	106 300	26 575	<b>S914</b>	4	2.3 - 3.8	52 500	13 125	102%
<b>S1014</b>	8	2.8 - 3.8	205 300	25 663	<b>S914</b>	8	2.8 - 3.8	122 500	15 313	68%
<b>S1024</b>	24	3.4-4.0	585 100	24 379	<b>S924</b>	24	3.4 - 3.9	370 700	15 446	58%



POWER9 with CVE					POWER8 with CVE					% evolution
Modèles	Cores	GHz	CPW Total	CPW /Core	Modèles	Cores	GHz	CPW Total	CPW /Core	GAIN max (%) P8 =>P9
<b>S922 8/16</b>	4	3.4 - 3.9	68 000	17 000	<b>S822 8/16</b>	4	4,15	49 960	12 490	36%
<b>S914</b>	4	2.3 - 3.8	52 500	13 125	<b>S814</b>	4	3,02	37 440	9 360	40%
<b>S914</b>	8	2.8 - 3.8	122 500	15 313	<b>S814</b>	8	3,72	81 050	10 131	51%
<b>S924</b>	24	3.4 - 3.9	370 700	15 446	<b>S824</b>	24	3,52	218 510	9 105	70%



POWER10 with CVE					POWER8 with CVE					% evolution
Modèles	Cores	GHz	CPW Total	CPW /Core	Modèles	Cores	GHz	CPW Total	CPW /Core	GAIN max (%) P8 =>P10
<b>S1022 8/16</b>	4	3.0-3.9	106 300	26 575	<b>S822 8/16</b>	4	4,15	49 960	12 490	113%
<b>S1014</b>	4	2.3 - 3.8	106 300	26 575	<b>S814</b>	4	3,02	37 440	9 360	184%
<b>S1014</b>	8	2.8 - 3.8	205 300	25 663	<b>S814</b>	8	3,72	81 050	10 131	153%
<b>S1024</b>	24	3.4-4.0	585 100	24 379	<b>S824</b>	24	3,52	218 510	9 105	168%



# Power10 IBM i QMODEL, QPRCFEAT, Software Tiers

Model	Machine Type Model	QMODEL	Cores	Processor Feature	QPRCFEAT	IBM i Processor Group / SW tier
S1014	9105-41B	41B	8	EPG2	EPG2	P10
S1014	9105-41B	41B	4	EPG0	EPG0	P05
Model	Machine Type Model	QMODEL	Cores	Processor Feature	QPRCFEAT	IBM i Processor Group / SW tier
S1022s	9105-22B	22B	16	EPGQ	EPGQ	P10
S1022s	9105-22B	22B	8	EPGQ	EPGQ	P10
S1022s	9105-22B	N/A	4	EPGR	N/A	NOT supported by IBM i
Model	Machine Type Model	QMODEL	Cores	Processor Feature	QPRCFEAT	IBM i Processor Group / SW tier
S1022	9105-22A	22A	40	EPGA	EPGA	P10
S1022	9105-22A	22A	32	EPG8	EPG8	P10
S1022	9105-22A	22A	24	EPG9	EPG9	P10
S1022	9105-22A	22A	12	EPG9	EPG9	P10
Model	Machine Type Model	QMODEL	Cores	Processor Feature	QPRCFEAT	IBM i Processor Group / SW tier
S1024	9105-42A	42A	48	EPGD	EPGD	P30
S1024	9105-42A	42A	32	EPGC	EPGC	P30
S1024	9105-42A	42A	24	EPGM	EPGM	P20
S1024	9105-42A	42A	12	EPGM	EPGM	P20

# Power10 IBM i Software Tiers

S1014 Feature Code	Processor SMT8 Cores	Typical Frequency Range	IBM i P Group
EPG2	8 cores	3.00 to 3.90 GHz	P10
EPG0	4 cores	3.00 to 3.90 GHz	P05

S1022s Feature Code	Processor SMT8 Cores	Maximum system cores	Typical Frequency Range	IBM i P Group
EPGQ	8 cores	16 cores	3.00 to 3.90 GHz	P10
EPGR	4 cores	4 cores	3.00 to 3.90 GHz	N/A

S1022 Feature Code	L1022 Feature Code	Processor SMT8 Cores	Maximum system cores	Typical Frequency Range	IBM i P Group
EPGA	EPGH	20 cores	40 cores	2.45 to 3.90 GHz	P10
EPG8	EPGF	16 cores	32 cores	2.75 to 4.00 GHz	P10
EPG9	EPGG	12 cores	24 cores	2.90 to 4.00 GHz	P10

S1024 Feature Code	L1024 Feature Code	Processor SMT8 Cores	Maximum system cores	Typical Frequency Range	IBM i P Group
EPGD	EPGK	24 cores	48 cores	2.75 to 3.90 GHz	P30
EPGC	EPGJ	16 cores	32 cores	3.10 to 4.00 GHz	P30
EPGM	EPGN	12 cores	24 cores	3.40 to 4.00 GHz	P20

Note: On 2U systems, multiple IBM i partitions are supported, and each partition can have up to four cores. VIOS is required.

# Power10 Scale-Out CPW Performance

Offering	Maximum Cores	1-socket	2-socket
S1024 / 24c	48-core	500,500	947,500
S1024 / 16c	32-core	383,400	725,000
S1024 / 12c	24-core	312,500	585,100

Offering	Maximum IBM i Cores (1)	1-socket	2-socket
S1022 / 20c	4-core	-	104,700
S1022 / 16c	4-core	-	106,300
S1022 / 12c	4-core	106,300	106,300

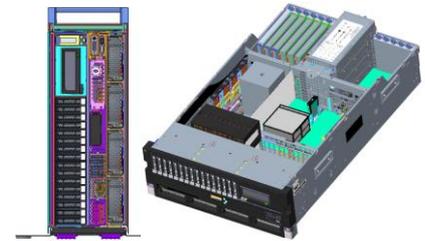
Offering	Maximum IBM i Cores (2)	1-socket	2-socket
S1022s / 8c	4-core	106,300	106,300

- (1) S1022 is maximum of 4 IBM i cores with VIOS  
 (2) S1022s is maximum of 4 IBM i cores with VIOS

Offering	Maximum IBM i Cores	1-socket
S1014 / 8c	8-core	205,300
S1014 / 4c	4-core	106,300

# IBM i Migration Options / Technology Refresh S1014

- Migrate IBM i clients from Power8 and Power9 systems to Power10
- System growth to support more workload on the same number of cores OR core counts can be reduced by up to ½ (based on Power8)



CPW (CPW/core)

P10 SW Tier	S914 / 8c	122,500 (15,313)
P10 SW Tier	S914 / 6c	78,500 (13,083)
P05 SW Tier	S914 / 4c	52,500 (13,125)



CPW (CPW/core)

P10 SW Tier	S814 / 8c	81,050 (10,131)
P10 SW Tier	S814 / 6c	56,400 (9,400)
P05 SW Tier	S814 / 4c	37,440 (9,360)



CPW (CPW/core)

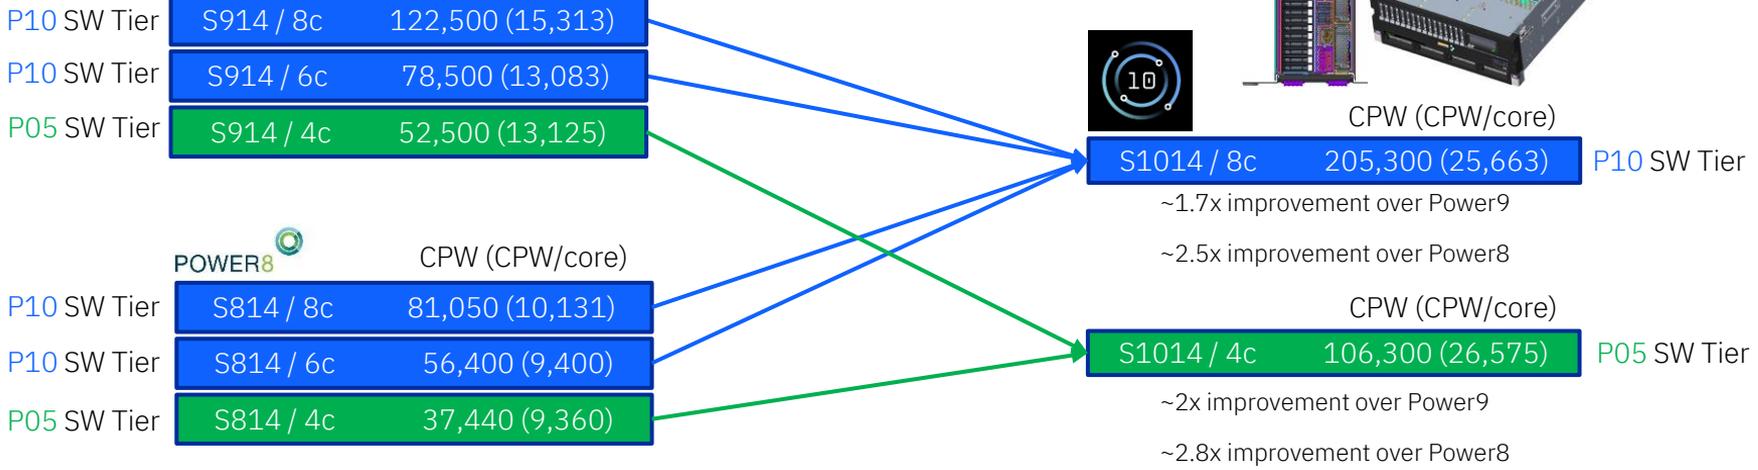
S1014 / 8c	205,300 (25,663)	P10 SW Tier
------------	------------------	-------------

~1.7x improvement over Power9  
~2.5x improvement over Power8

CPW (CPW/core)

S1014 / 4c	106,300 (26,575)	P05 SW Tier
------------	------------------	-------------

~2x improvement over Power9  
~2.8x improvement over Power8



# IBM i Migration Options / Technology Refresh S1024

- Migrate IBM i clients from Power8 and Power9 systems to Power10
- System growth to support more workload on the same number of cores OR core counts can be reduced by up to ½ (based on Power8)



CPW (CPW/core)

P20 SW Tier	S924 / 24c 3.4-3.9 370,700 (15,446)
P20 SW Tier	S924 / 22c 3.4-3.9 350,000 (15,909)
P20 SW Tier	S924 / 11c 3.4-3.9 195,299 (17,755)
P20 SW Tier	S924 / 20c 3.5-3.9 318,000 (15,900)
P20 SW Tier	S924 / 10c 3.5-3.9 174,500 (17,450)
P20 SW Tier	S924 / 16c 3.8-4.0 268,500 (16,781)
P20 SW Tier	S924 / 8c 3.8-4.0 145,500 (18,188)

CPW (CPW/core)

P30 SW Tier	S1024 / 48c 2.75-3.9 947,500 (19,740)
P30 SW Tier	S1024 / 24c 2.75-3.9 500,500 (20,854)

CPW (CPW/core)

P30 SW Tier	S1024 / 32c 3.1-4.0 725,000 (22,656)
P30 SW Tier	S1024 / 16c 3.1-4.0 383,400 (23,963)

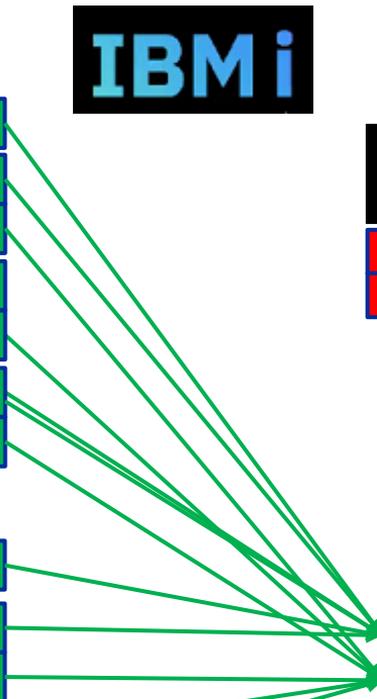


CPW (CPW/core)

P20 SW Tier	S824 / 24c 3.52 218,510 (9,105)
P20 SW Tier	S824 / 16c 4.15 164,470 (10,279)
P20 SW Tier	S824 / 8c 4.15 89,580 (11,198)
P20 SW Tier	S824 / 12c 3.89 123,240 (10,270)
P20 SW Tier	S824 / 6c 3.89 68,250 (11,375)

CPW (CPW/core)

P20 SW Tier	S1024 / 24c 3.4-4.0 585,100 (24,379)
P20 SW Tier	S1024 / 12c 3.4-4.0 312,500 (26,042)



# IBM i Migration Options / Technology Refresh S1022/S1022s

- Migrate IBM i clients from Power8 and Power9 systems to Power10
- System growth to support more workload on the same number of cores OR core counts can be reduced by up to ½ (based on Power8)



CPW (CPW/core)

P10 SW Tier S922 / 4c 3.4-3.9 GHz 68,000 (17,000)

P10 SW Tier S922 / 4c 2.9-3.8 GHz 60,000 (15,000)

P05 SW Tier S922 / 1c 2.3-3.8 GHz 19,000 (19,000)



CPW (CPW/core)

P10 SW Tier S822 / 4c 4.15 GHz 49,960 (12,490)

P10 SW Tier S822 / 4c 3.89 GHz 45,220 (11,305)

P10 SW Tier S822 / 4c 3.42 GHz 42,470 (10,618)

P05 SW Tier S812 / 1c 3.02 GHz 9,360 (9,360)



CPW (CPW/core)

S1022 / 4c 2.45-3.9 104,700 (26,175) P10 SW Tier

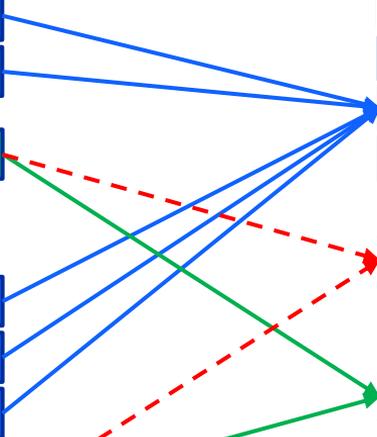
S1022 / 4c 2.75-4.0 106,300 (26,575) P10 SW Tier

S1022 / 4c 2.9-4.0 106,300 (26,575) P10 SW Tier

S1022s / 4c 3.0-3.9 106,300 (26,575) P10 SW Tier

CPW (CPW/core)

S1014 / 4c 106,300 (26,575) P05 SW Tier



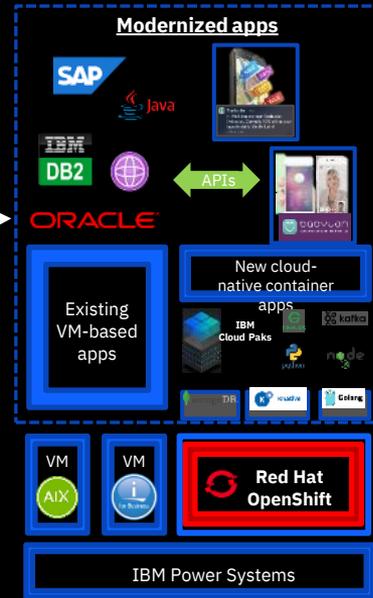
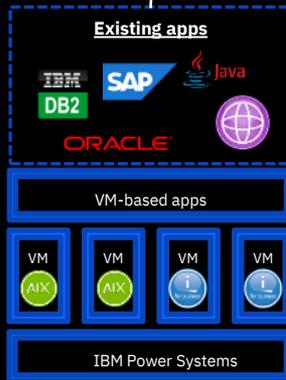
# What Does App Modernization Look Like on IBM Power Systems?

Check out the **Application Modernization Field Guide**

<https://www.ibm.com/downloads/cas/D9POQ3YR>



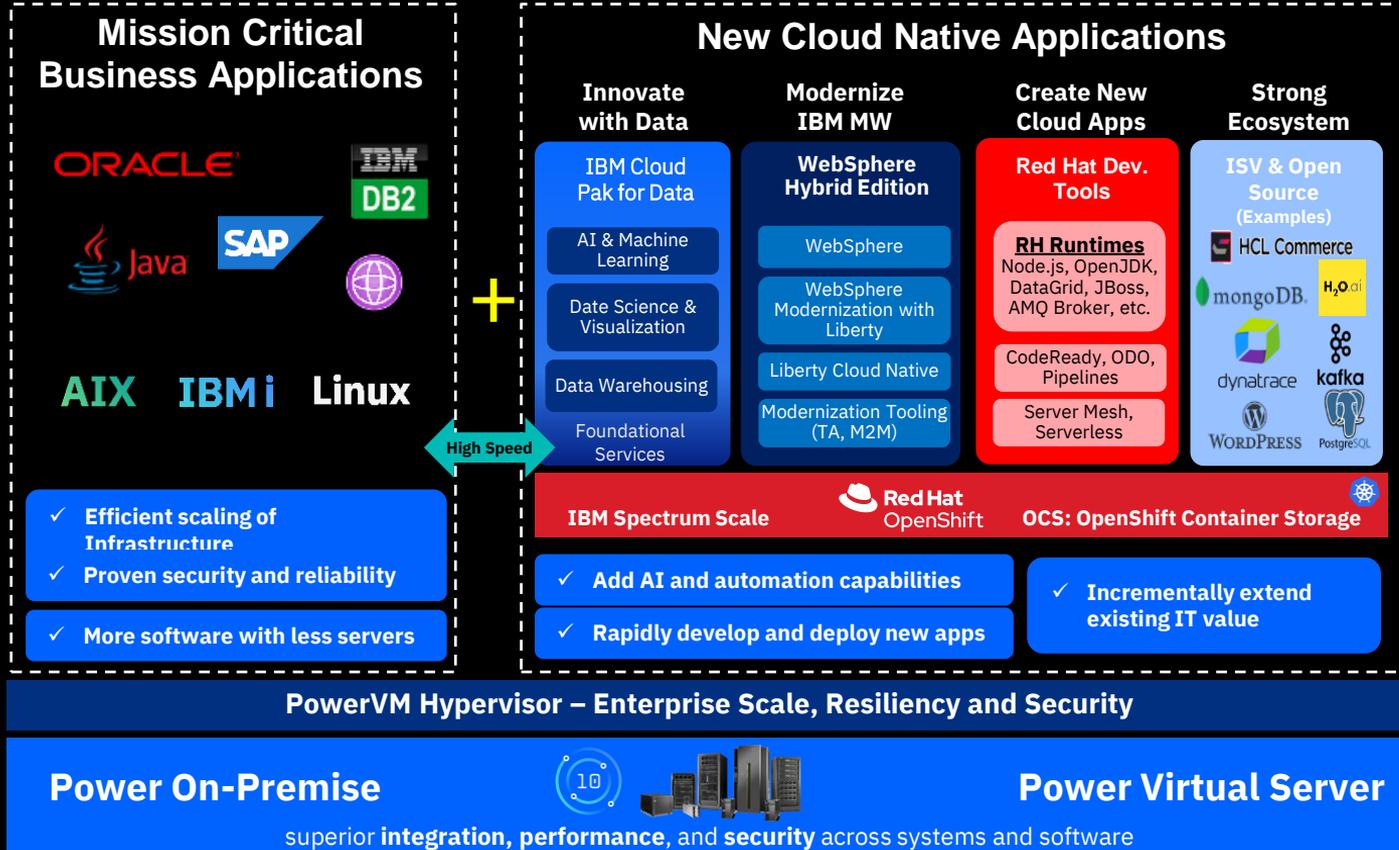
Incrementally modernize



**Deploy anywhere:**  
on-premises or in  
the public cloud



# IBM Power – One Platform for Digital Transformation



Exploit databases on Power

Co-locate apps & data

Infuse AI into apps

Containerize ISV apps

# Provide an efficient, agile data fabric with Cloud Pak for Data + Power S1022

*3.5x more throughput per core and 57% lower 3-yr TCO running Db2 Warehouse in Cloud Pak for Data on Power S1022 vs compared x86-based server*

		<b>IBM Power S1022</b> (32-core, 1TB)	<b>Intel Xeon Platinum 8380</b> (80-core, 1TB)
<b>Solution Cost</b> -Server + (CP4D base @ \$25,164/core and DB2 @ \$17,920/core) with 3 years support		<b>\$1,537,007</b>	<b>\$3,542,552</b>
<b>Server price</b> - Power S1022 - Xeon Platinum 8380-based server - Includes OS/Virt - Both with 3-yr warranty		<b>\$166,127</b>	<b>\$115,352</b>
<b>QpH</b> Total Queries/Hour		<b>4,700 QpH</b>	<b>3,337 QpH</b>
<b>QpH per \$1000</b>		<b>3.1 QpH/\$K</b>	<b>0.94 QpH/\$K</b>

**3.5X**  
per core performance

**3.3X**  
Better Price-Performance

**57%**  
Lower solution costs over 3 years

Based on IBM internal testing of Db2 Warehouse executing a sample analytics workload of 30 distinct queries of varying complexity (intermediate and complex). Results valid as of May 3, 2022 and conducted under laboratory conditions, individual results can vary based on workload size, use of storage subsystems and other conditions. Comparison is based on queries per hour on IBM Power S1022 (2x16core/1TB) versus Intel Xeon Platinum 8380-based (2x40-core/1TB) systems. Test was run with IBM Db2 Warehouse on Cloud Pak for Data 4.0.6 and Red Hat OpenShift 4.9 utilizing Db2 performance best practices (pages 34 and 35) at <https://wdug.com/resources/Documents/LUW%20Presentations/March%202019/Db2%20BLU%20Internals%20Concepts%20and%20Best%20Practices%20Updated%20for%20V11%201.pdf>

Total Cost of Ownership is defined as hardware, software, and maintenance costs over a period of three years and Price-Performance is performance/3-yr TCO where hardware pricing is based on IBM Power S1022 <https://www.ibm.com/products/> and extrapolated industry standard x86 pricing based on IBM internal industry knowledge and Software pricing available at <https://marketplace.redhat.com/en-us/products/ibm-cloud-pak-for-data/pricing> and <https://www.allbluesolutions.com/store>



common  
EUROPE LUXEMBOURG



IBM | Let's  
Create

*Matrices de support des OS*

# Software solutions that grow with you



## IBM AIX

Scalable and robust enterprise open standards-based UNIX operating system for the Power Systems architecture. AIX has a history of consistently delivering a high-performance secure environment.

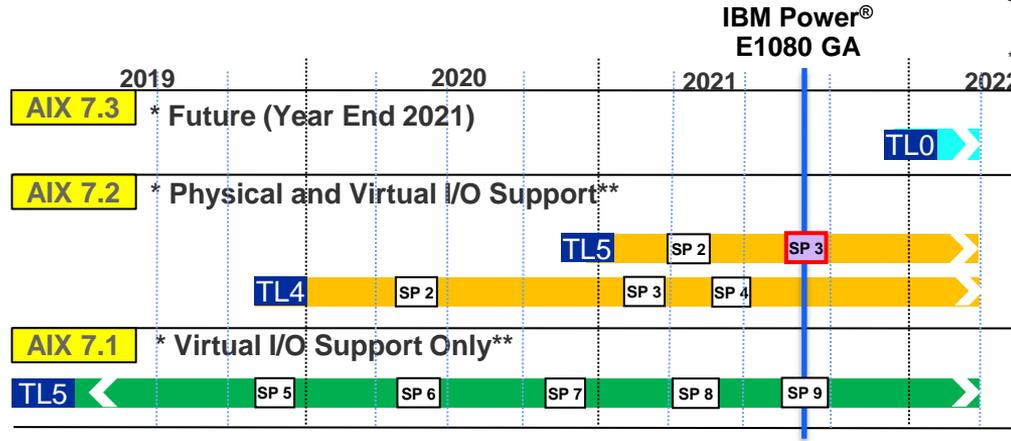
## IBM i

An integrated platform enabling flexibility and dependability with robust architecture, exceptional security and business resilience.

## Linux

An open operating system built by the open source community, resulting in faster processing speed, bandwidth and inherent security.

# AIX Release Support for Power10



## Run-time Considerations

- AIX 7.1 runs fully virtualized\*\*
- AIX 7.2 supports full virtualization and physical IO

\*\* Excludes Native VF (SR-IOV) directly in the client LPAR

SP timelines are for illustration only

## Supported levels at Power10 GA Sept 17, 2021

### If installing the AIX operating system LPAR with any I/O configuration:

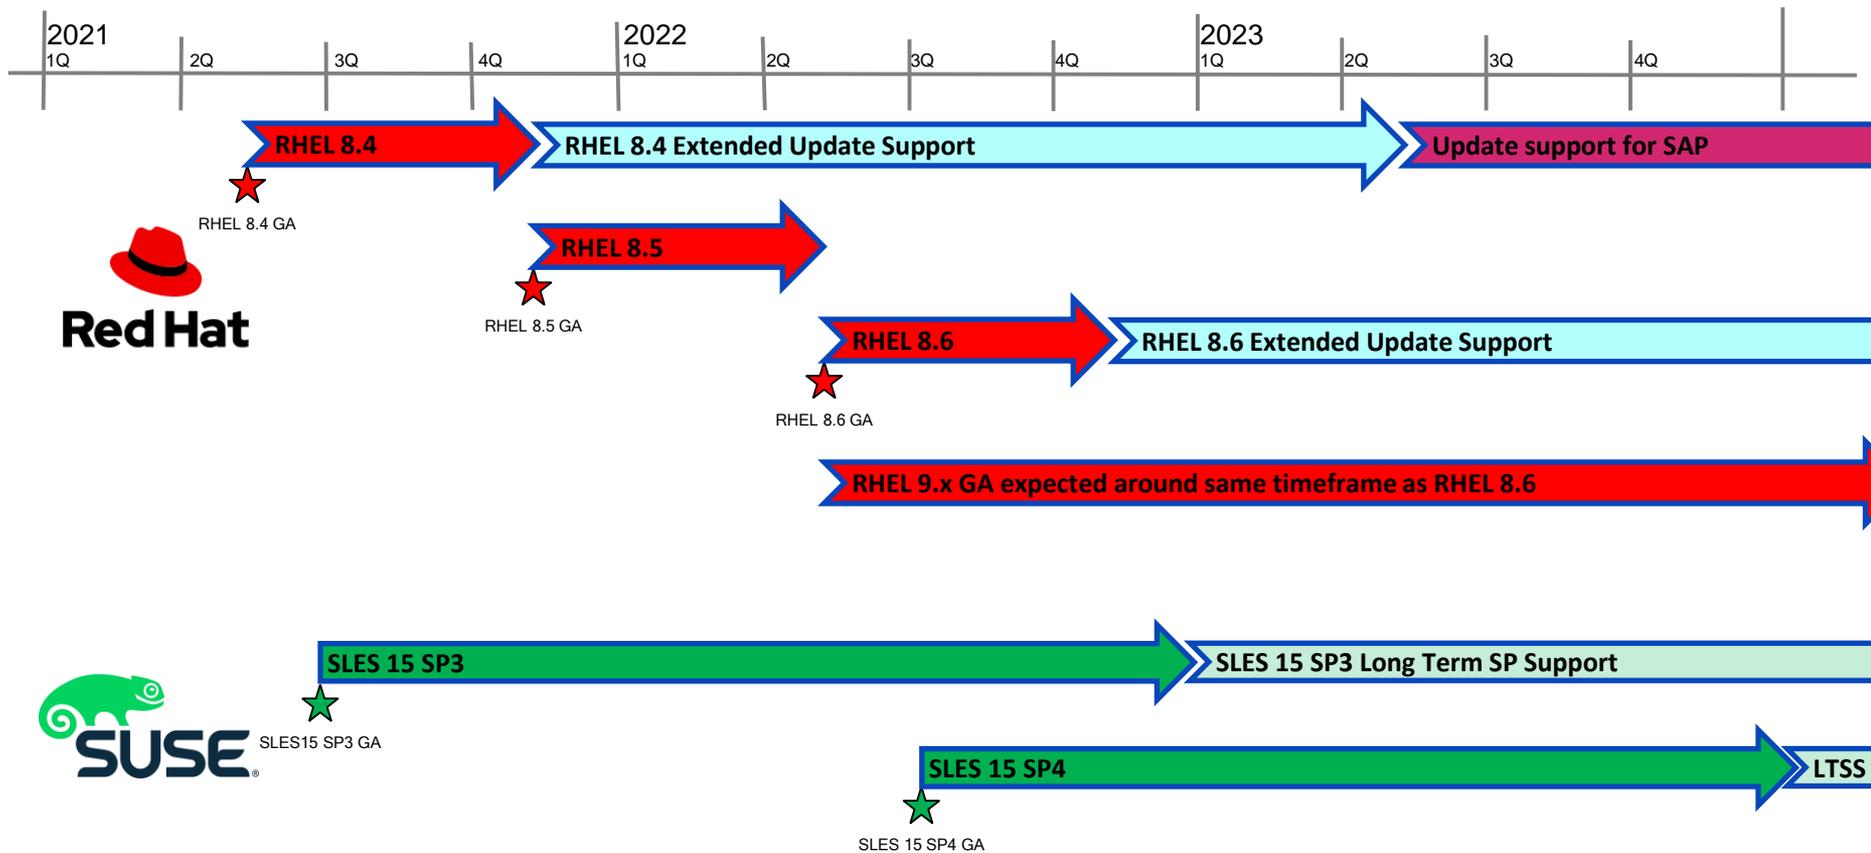
- AIX Version 7.2 with the 7200-05 Technology Level and Service Pack 7200-05-03-2136 or later
- AIX Version 7.2 with the 7200-04 Technology Level and Service Pack 5 (planned availability February 11, 2022)

### If installing the AIX operating system Virtual I/O only not including Native VF (SR-IOV) directly in the client LPAR :

- AIX Version 7.2 with the 7200-05 Technology Level and Service Pack 7200-05-01-2038 or later
- AIX Version 7.2 with the 7200-04 Technology Level and Service Pack 7200-04-01-1939 or later (minimum 7.2 level)
- AIX Version 7.1 with the 7100-05 Technology Level and Service Pack 7100-05-05-1937 or later (minimum 7.1 level)

Detailed Minimum and Recommended AIX Levels per Machine-Type-Model are available on the [IBM System to AIX Maps website](#)

# Native Linux Support for Power10 processor-based systems



# IBM i et gamme Power



IBM announced back in July that **The IBM Power E850 (not the newer E850c) will reach End of Service on the 31st of October 2023.**



Servers	IBM i 6.1 <sub>1</sub>	IBM i 7.1 <sub>1</sub>	IBM i 7.2 <sub>1</sub>	IBM i 7.3	IBM i 7.4	IBM i 7.5
<b>Power10</b> S1014, S1022/L1022/S1022s (VIOS only), S1024/L1024, E1080				✓	✓	✓
<b>Power9</b> S914, S922 (VIOS only), H922 (VIOS only), S924, H924, E980 S922 Entry (single native I/O partition only)		✓ <sup>4</sup> ⚠	✓ ⚠	✓	✓	✓
<b>Power8</b> S814, S822 (VIOS only), S824, E870, E870C, E880, E880C S812 (no IBM i 7.1, native I/O only)		✓	✓	✓	✓ ⚠	
<b>Power7/7+</b> 710, 720, 730, 740, 750 (VIOS only), 760 (VIOS only), 770, 780, 795	✓	✓	✓	✓		
<b>Power7 BladeCenter PS700/701/702/703/704 (all VIOS only)</b> <b>Power7/7+ Flex p260/270/460 (all VIOS only)</b>	✓	✓	✓			
<b>Power6+ 520, 550, 560, 570, JS23/43 (VIOS only)</b> <b>Power6 520, 550, 570, 595, JS12/22 (VIOS only)</b>	✓	✓	✓ <sup>3</sup>			
<b>Power5/5+</b> 515, 520, 525, 550, 570, 595	✓	✓				
800, 810, 825, 870, 890	✓					

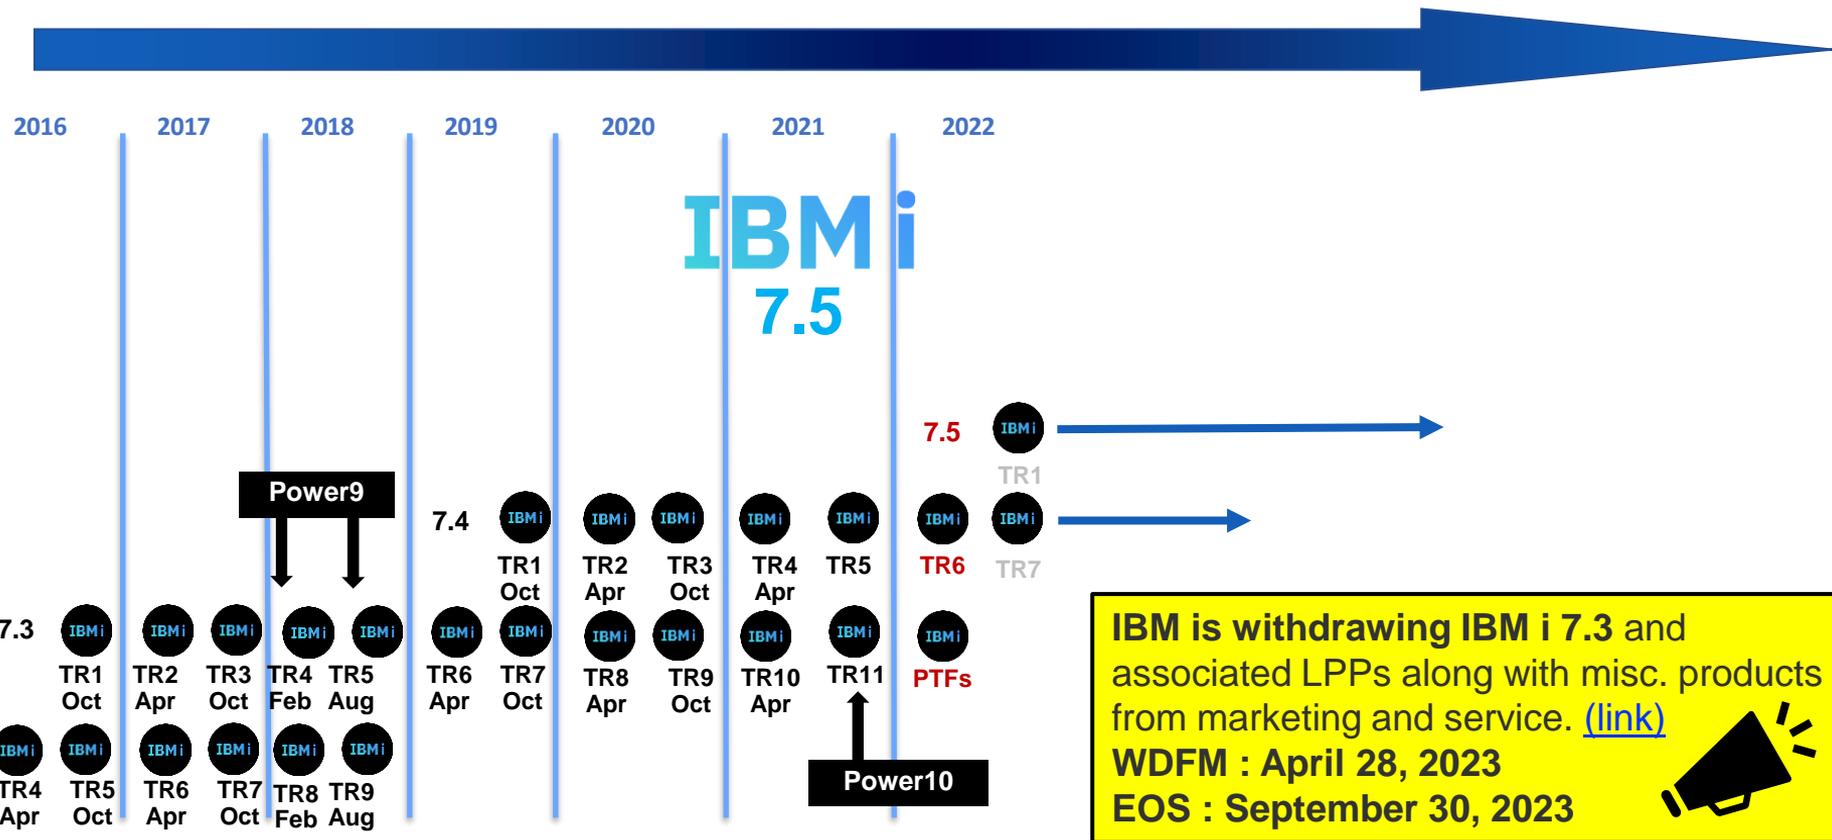
1 - IBM i 6.1, 7.1 et 7.2 : reportez-vous au cycle de vie de la version IBM i pour obtenir des informations sur le support étendu.

3 - Pas de support IOP ou HSL

2 - Client virtuel uniquement

4 - Disponible uniquement sur certains modèles S922 et E980.

# IBM i Release Roadmap – Announcement Dates



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

\*\* Arrows indicate "ongoing status and do not imply any specific dates.

# May 3: Announcing Innovations for IBM i

- IBM i 7.5
  - Enhanced options for achieving greater levels of security
  - Greater interoperability through REST APIs
  - Real-time upgrades with Db2 Mirror support for multiple releases of IBM i
- IBM i Modernization Engine for Lifecycle Integration (Merlin)
  - Automated conversion of legacy fixed-format RPG code into modern Free Form to onboard new developers faster
  - Industry-standard tools like Git, Jenkins, and Visual Studio Code integrated for cloud-native development
  - Secure deployment of new IBM i applications
- IBM i Subscription
  - Run hardware, software license, and support in 1-to-5-year terms
  - Keep client technology current, lower TCA, and provide flexibility with OpEx consumption

# IBM i



# IBM i System Subscription Offering for SMBs

## Customer Pain Points

- Traditional IT budget pressures
- It just works and simplifies operation and use
- Maintain technology currency

## Offering description

**IBM i System Subscription Details:** IBM Power S1014 with 4 cores, 64 GB of memory, 800 GB NVMe storage, 1 IBM i entitlement/25 users (P05 tier), 3 to 5-year terms and enhanced support (24x7) and services

- Cloud-like Financial Experience
  - Sign-up for IBM i System Subscription to get on-prem capacity
  - Complete PaaS for IBM i applications on-prem
- Flexible and Predictive financials
  - Lower TCA, consistent payment and potential for OpEx consumption (IBM owned)
  - No IBM i transition fees to move to new Power and IBM i technologies
- Worry-free Maintenance and operation
  - Default 24x7 coverage with value-added IBM services covering lifecycle activities
    - Security/health recommendations, remote firmware updates, assistance with OS updates/upgrade planning, critical outage support, IBM Technical Account Manager



## IBM i + Power S1014 Subscription Terms

- 3 to 5-year terms initial subscription; 1 year renewals
- Annual payment for all content delivered

**Announce: September 6, 2022**

**GA: September 23<sup>rd</sup>, 2022**

IBM i System Subscription    Smartphone Plan



**IBM i**

<  
**Cost**

~\$50 per user per month for IBM i System Subscription on-prem



~\$75 per user per month smartphone plan cost



common  
EUROPE LUXEMBOURG



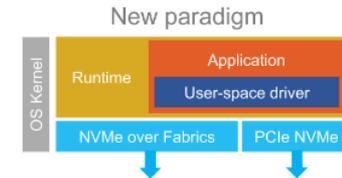
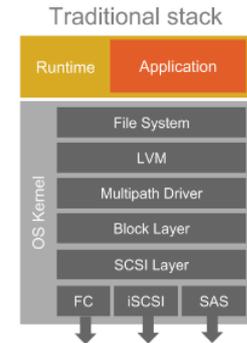
IBM | Let's  
Create

## *Hardware et IBM i le stockage NVME*

# Qu'est ce que le NVMe? (Non-Volatile Memory Express)



- **NVMe est une interface et un protocole de communication**
  - NVMe est conçu dès le départ pour offrir **une bande passante élevée et un accès au stockage à faible latence**.
  - NVM Express définit une interface efficace permettant au logiciel hôte de communiquer avec un sous-système de mémoire non volatile via PCI Express (NVMe over PCIe).
  - Fonctionnement analogue à celui de SAS et SATA, mais conçu pour réduire la charge des pilotes, du système d'exploitation et des applications.
  - Utilise des compléments d'E/S basés sur des sondages, par opposition aux compléments basés sur des interruptions.
- **NVMe utilise la matrice PCIe**
  - Plusieurs dispositifs aujourd'hui sur le marché
  - Plusieurs facteurs de forme, y compris des disques de 2,5 pouces
- **NVMe a été conçu pour les hautes performances**
  - Augmentation des IOP, de la bande passante et réduction de la latence.
  - Exploitation de la mémoire Flash et des mémoires non volatiles de nouvelle génération.
  - Exploitation des environnements multicœurs, parallélisme élevé des E/S.



# Qu'est ce que le NVMe?



## NVMe Controller

- PCIe Attached
- Parallel Architecture (Multi Q, Q pairs ...)
- Low Latency Design
- Fabrics Attach Friendly
- Self Encryption & Sanitize
- Virtualization (Multiple Namespace, SR-IOV)
- IO Determinism
- Zoned Namespace
- Management Interface Support (in & out of band)
- Computational Storage
- ... More Innovations



## Media

- NAND TLC (most widely used)
- Optane 3DXP
- LL NAND
- NAND QLC
- DRAM – Flash backed
- MRAM
- .. More to come

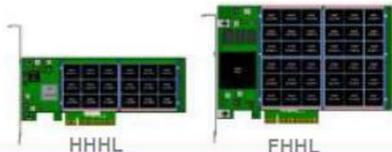
**Form Factors:** M.2, U.2 (2.5" thin & thick), EDSFF (E1.S, E1.L, E3.S, E3.L thin & thick), Add in card

**Software Infrastructure:** Investissement protégé par la réutilisation sur les appareils de plusieurs fournisseurs

# NVMe format type

## Add-in Card (AIC)

- x4 / x8 PCIe Gen3 or Gen4
- High performance and power (25-50 Watts)
- Focus on high capacity - up to 6.4TB currently
- Best performance and latency



Feature Code	PCIe3 Card NVMe Flash Adapter		
	1.6TB	3.2TB	6.4TB
AIX/Linux (LP/FH)	#EC5G / #EC5B	#EC5C / #EC5D	#EC5E / #EC5F
IBM i (LP/FH)	#EC6U / #EC6V	#EC6W / #EC6X	#EC6Y / #EC6Z

October 2019



Feature Code	PCIe4 Card NVMe Flash Adapter		
	1.6TB	3.2TB	6.4TB
AIX/Linux (LP/FH)	#EC7A / #EC5B	#EC7C / #EC7D	#EC7E / #EC7F
IBM i (LP/FH)	#EC7J / #EC7K	#EC7L / #EC7M	#EC7N / #EC7P

October 2020

## U.2

- x2 / x4 PCIe Gen3
- Higher power required to achieve max SSD performance
- Expect to be direct attached to CPU vs through a HBA



2.5" 7mm

**E950/E980**  
800GB - #EC5J  
1.6TB - #EC5K  
3.2TB - #EC5L



August 2018



2.5" 15mm

**S9xx G Models**  
800GB - #EC5X  
1.6TB - #ES1F/ES1E  
3.2TB - #ES1H/ES1G  
6.4TB - #EC5W/EC5V



July 2020

## M.2

- x2 / x4 PCIe Gen3
- Different physical sizes
- Power less than 9 watts
- Limitations: Minimal Connector Plug Capability (<10), No hot swap
- Targeted as a replacement for SATA SSDs

■ Withdrawn from marketing 3/24/2020, announcement 920-065



Card #EC59 with 1-2 #ES14



**S9xx**  
400GB #ES14



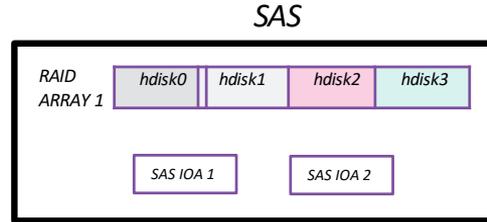
# PCIe Card NVMe/IBM i/IBM Power Systems Hardware

- NVMe est capable de fournir des performances supérieures à celles des disques SSD. La technologie NVMe peut fournir un nombre d'IOPS en lecture ou en écriture et un débit (Go/sec) nettement supérieurs à ceux des SSD SAS/SATA. Les différences de performances réelles des systèmes ou des applications varient selon le client et la charge de travail.
- NVMe offre des capacités de virtualisation supplémentaires puisque chaque dispositif est un point d'extrémité PCIe qui peut être dédié à une partition/LPAR.
- Au moins une paire d'adaptateurs NVMe identique est requise ; les paires d'adaptateurs NVMe suivantes peuvent être différentes de la première paire. Après la commande d'une paire identique, un adaptateur NVMe de capacité différente est autorisé. La meilleure pratique consiste à ne commander que des paires pour que la mise en miroir soit la plus simple possible.
- Les dispositifs NVMe nécessitent la mise en miroir du système d'exploitation IBM i car il n'y a pas de support RAID matériel. Les paires en miroir doivent être sur des dispositifs physiques différents. Les disques NVMe ne peuvent être mis en miroir que sur des disques NVMe et les disques SAS ne peuvent être mis en miroir que sur des disques SAS.
- Le remplacement à chaud n'est pas pris en charge(AIC card), mais un NVMe supplémentaire pourrait être sur le système comme un remplacement à froid pour accélérer le processus de réparation, et il est seulement un remplacement dans le fait qu'un client n'a pas à le commander / brancher dans le serveur. Le développement d'IBM i est conscient du désir d'avoir quelque chose de plus performant qu'un cold spare.

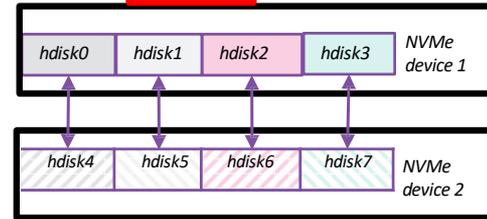


# SAS Versus NVMe Storage

Vue du groupe de volumes de l'OS



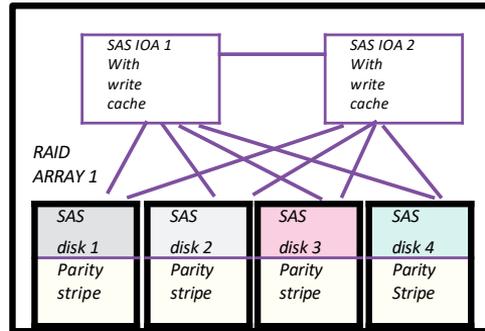
NVMe



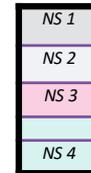
Mise en miroir du système d'exploitation  
VIOS utiliserait lvmX

Le système d'exploitation est conscient du RAID matériel

Vue physique du dispositif



NVMe device 1

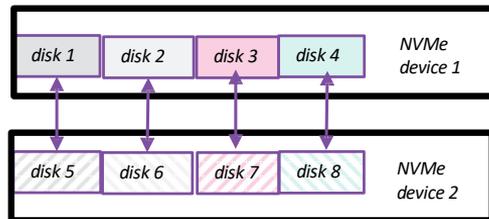


NVMe device 2



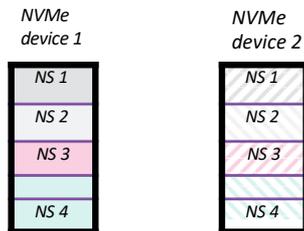
# NVMe Mirroring (2 or More Devices)

Vue du groupe de volumes de l'OS



Tous les namespaces mis en miroir ont la même taille sur tous les périphériques.

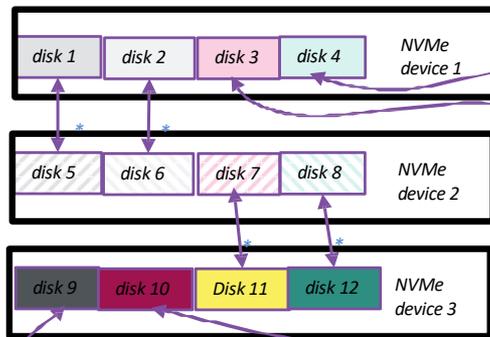
Vue physique du dispositif



La capacité physique du dispositif n'a pas besoin d'être de la même taille.  
**On miroire des namespace et pas des devices**

# NVMe Mirroring (3 or More Devices Option 1)

Vue du groupe de volumes de l'OS

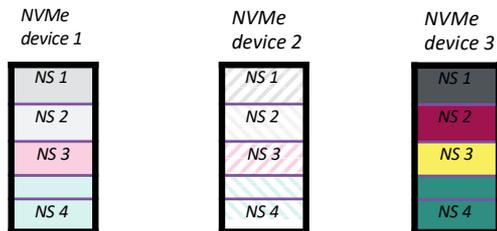


Tous les namespaces mis en miroir ont la même taille sur tous les périphériques.

Tous les périphériques NVMe ont la même taille dans cet exemple donc même nombre de space sur chaque NVMe

\*miroir OS

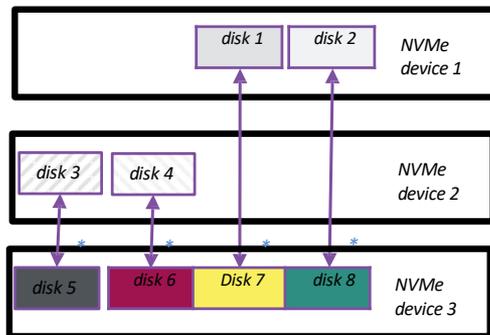
Vue physique du dispositif



La capacité physique du dispositif n'a pas besoin d'être de la même taille.  
**On miroire des namespace et pas des devices**

# NVMe Mirroring (3 or More Devices Option 2)

Vue du groupe de volumes de l'OS

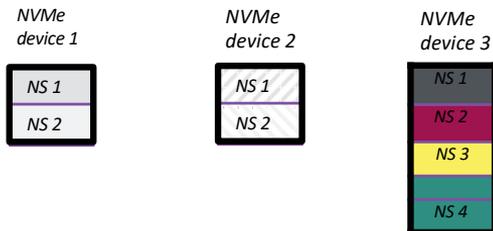


Tous les namespaces mis en miroir ont la même taille sur tous les périphériques.

Le troisième périphérique NVMe est deux fois plus grand que les autres.

\*miroir OS

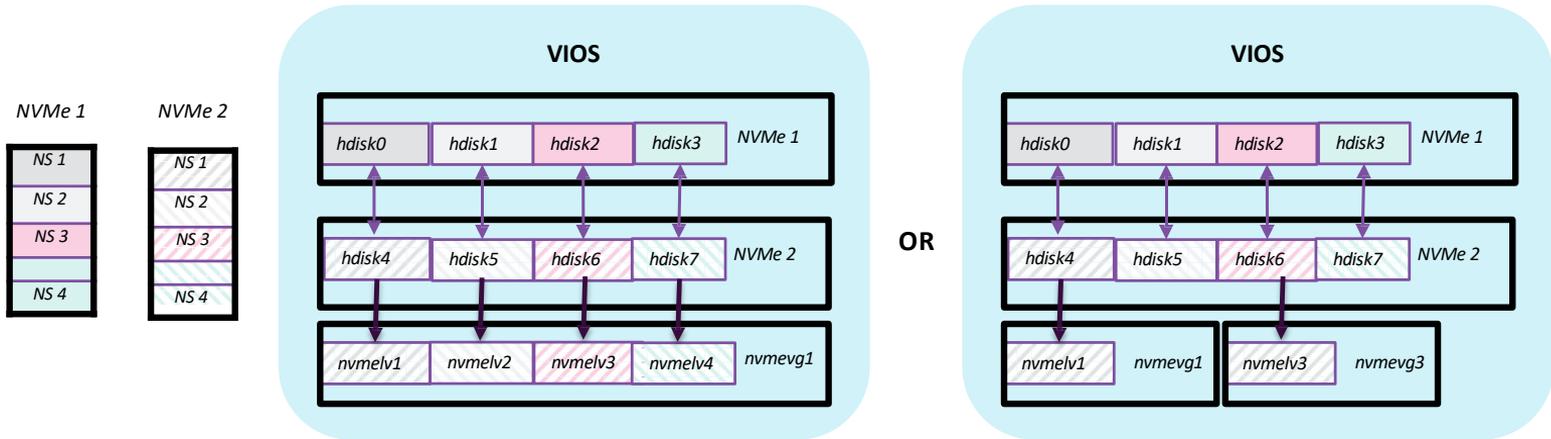
Vue physique du dispositif



La capacité physique du dispositif n'a pas besoin d'être de la même taille. On miroire des namespace et pas des devices

# VIOS and NVMe - Single VIOS and 2 NVMe

- Créez le Namespace NS1 sur NVMe1 que VIOS voit comme hdisk0.
  - Créez le Namespace NS1 sur NVMe2 que VIOS voit comme hdisk4.
  - Ajoutez hdisk0 et hdisk4 au groupe de volume nvmevg1 (configurez éventuellement un miroir dans VIOS)
  - Créez le volume logique nvmelv1
  - Créez le périphérique vscsi vnvme1 avec nvmelv1 comme cible de sauvegarde.
  - Attribuer vnvme1 à vhost de la partition IBM i guest
- 
- Vous décidez de l'emplacement du miroir (VIOS ou OS) et des configurations lv & vg.





# RPQ 8A2707 to order 2 add. 3.2 TB - PCIe4 NVMe U.2

RPQ 8A2707 provides the ability to order two additional 3.2 TB - PCIe4 NVMe U.2 Enterprise modules for IBM i (**equivalent to FC #ES1H**) for the **9105-41B server**. This is RPQ is **exclusive to IBM i customers that have purchased LPP 5733-ICC IBM Cloud Storage Solutions for i**. A total of **two RPQs are required to ensure two NVMe devices are configured to provide IBM i data protection via software mirroring by the OS**.

**Each instance of RPQ 8A2707 will deliver one 3.2 TB - PCIe4 NVMe U.2 Enterprise module** for IBM i (equivalent to FC #ES1H). A quantity of two RPQs 8A2707 are required for each 9105-41B server.

## Prereqs:

- 1) Software PID 5733-ICC.
- 2) 4-Core (9105-41B) Configuration.

[https://www.ibm.com/common/ssi/ShowDoc.wss?docURL=/common/ssi/rep\\_rp/7/ENUS8A2707/index.html](https://www.ibm.com/common/ssi/ShowDoc.wss?docURL=/common/ssi/rep_rp/7/ENUS8A2707/index.html)

MERCS

