


IBM Power Platform

From Core to Cloud to Edge:

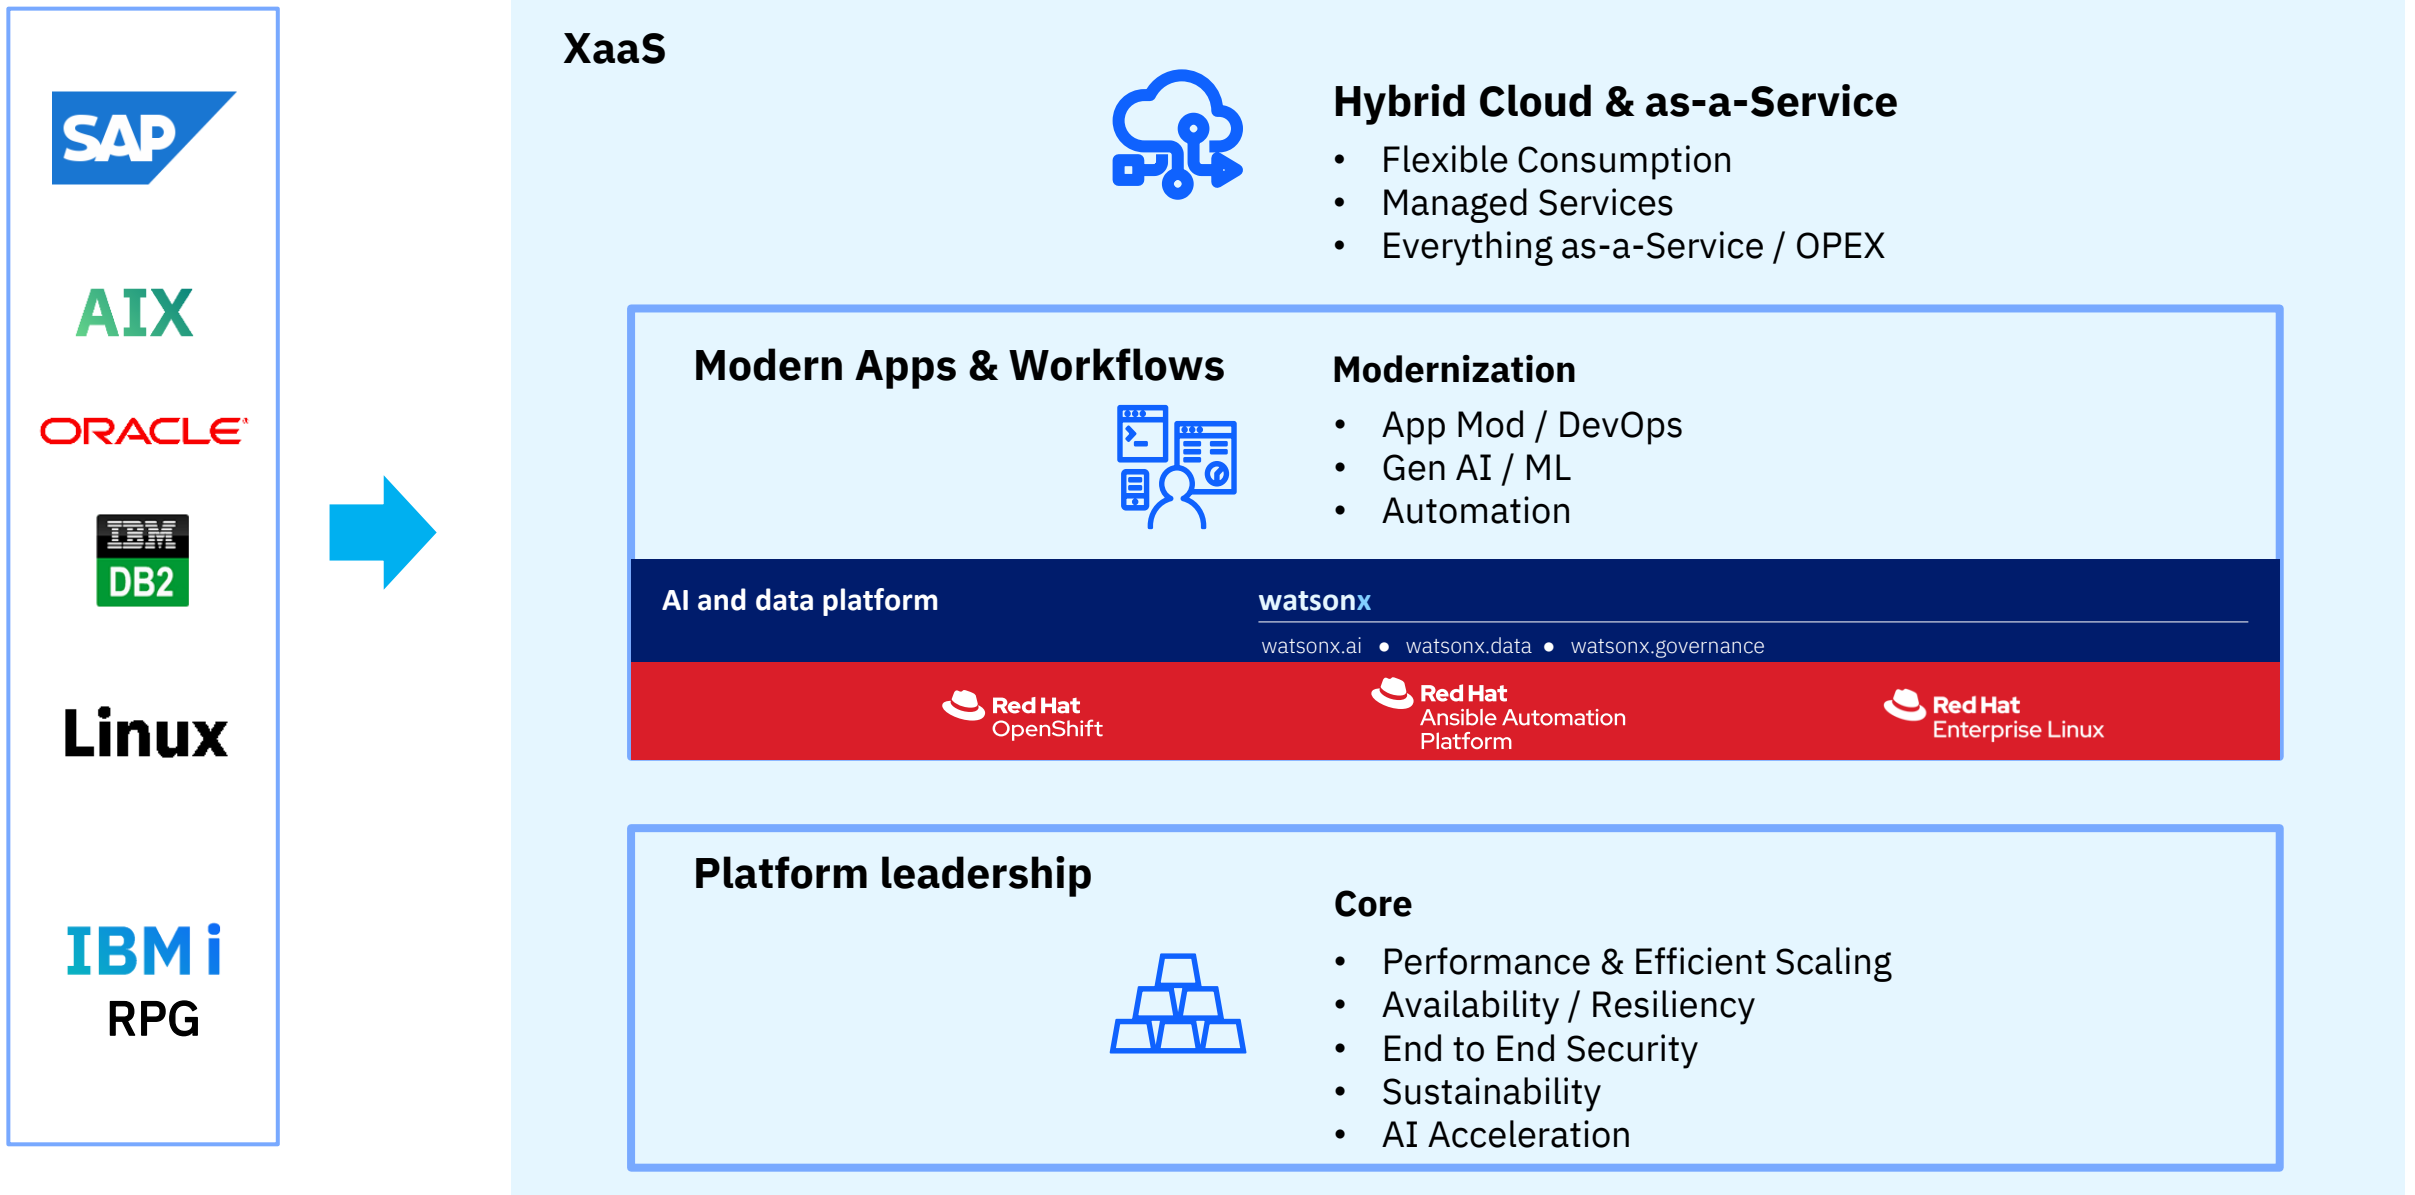
Status og nyheter

Espen Skarpsno
Technical Specialist – Storage and Power
espen.skarpsno@ibm.com
+47 98250903



September 2024

IBM Power: Accelerate digital transformation with Hybrid Cloud and AI



Meanwhile in x86 and virtualization...




Some AMD EPYC server CPUs have a serious security flaw, so patch now

News By Sead Fadilpašić published November 15, 2023

A recently discovered AMD EPYC flaw allows for privilege escalation

<https://www.techradar.com/pro/security/some-amd-epyc-server-cpus-have-a-serious-security-flaw-so-patch-now>



Killing VMware

Ever since Broadcom completed its acquisition of VMware last November, customers have been navigating the fallout — and they're not happy.



By Steven J. Vaughan-Nichols
Computerworld | FEB 23, 2024 11:10 AM PST

<https://www.computerworld.com/article/3713242/killing-vmware.html>



Intel and AMD Chips Found Vulnerable to Data Theft

Intel and AMD have sold billions of such vulnerable chips over the past decade, particularly affecting the security of servers.

<https://www.spiceworks.com/it-security/data-security/news/intel-and-amd-chips-found-vulnerable-to-data-theft/>



VMware customers are jumping ship as Broadcom sales continue - here's where they're moving to

News By Craig Hale published March 4, 2024

Many ex-VMware customers are going open-source

<https://www.techradar.com/pro/vmware-customers-are-jumping-ship-as-broadcom-sales-continue-heres-where-theyre-moving-to>



Intel emits patch to squash chip bug that lets any guest VM crash host servers

Sapphire Rapids, Alder Lake, Raptor Lake chip families treated for 'Red

https://www.theregister.com/2023/11/14/intel_outofband_patch/



The self-created risk in Broadcom's big VMware kiss-off

Arise, the new Migration Nation

Rupert Goodwins

Mon 26 Feb 2024 10:15 UTC

https://www.theregister.com/2024/02/26/opinion_column_broadcom_vmware/



Why Broadcom Is Killing off VMware's Standalone Products

VMware's roots are everywhere, and pulling them up is going to unearth a lot of dirt. The acquisition seriously changes the relationship customers will have with the new version of the company going forward.







Jan 23rd, 2024 11:59am by Chris J. Preimesberger

<https://thenewstack.io/why-broadcom-is-killing-off-vmwares-standalone-products/>



Delivering Cloud-to-Edge Excellence on Scale-Out Systems

New

<h2>Edge Computing</h2> <ul style="list-style-type: none"> Up to 8 cores and 256 GB memory footprints Reduced IT footprint Fully active, static servers 	<h2>Mainstream Computing</h2> <ul style="list-style-type: none"> Up to 16 cores and 2TB memory footprints <ul style="list-style-type: none"> Up to 32% perf/price increase vs. P9 <ul style="list-style-type: none"> Fully active, static servers Lower technical requirements and competitive deals 			<h2>Core Computing</h2> <ul style="list-style-type: none"> Up to 48 cores and 8TB memory footprints <ul style="list-style-type: none"> Up to 1.4x system performance vs. P9 Dynamic Capacity consumption with CUoD and PEP2.0 Value-driven solutions and higher technical standards 	<h2>Cloud Computing</h2> <ul style="list-style-type: none"> Unlimited cores and memory Dynamic Capacity Computing aaS
<p>S1012 9028-21B</p>  <ul style="list-style-type: none"> 1-socket, 2U, half-wide Up to 8 cores per system 1, 4, 8 SMT8 cores/socket 4 ISDIMM slots 256 GB memory 4 PCIe HHHL slots (All Gen5) 4 NVMe U.2 bays Max of 6.4 TB of internal storage Optional Internal RDX Media Bay Rack and Tower form factors 	<p>S1014 9105-41B</p>  <ul style="list-style-type: none"> 1-socket, 4U Up to 24 cores per system 4, 8, 24 SMT8 cores/socket 8 DDIMM slots 1 TB memory 5 PCIe FHHL slots (4 Gen5) 16 NVMe U.2 bays Max of 102.4 TB of internal storage Optional Internal RDX Media Bay Rack and Tower form factors 	<p>S1022s 9105-22B</p>  <ul style="list-style-type: none"> 1,2-socket, 2U Up to 16 cores per system 4, 8 SMT8 cores/socket 16 DDIMM slots 2 TB memory 10 PCIe HHHL slots (8 Gen5) 8 NVMe U.2 bays Max of 51.2 TB of internal storage 	<p>S1022 & L1022 9105-22A / 9786-22H</p>  <ul style="list-style-type: none"> 1,2-socket, 2U Up to 40 cores per system 12, 16, 20 SMT8 cores/socket 32 DDIMM slots 4 TB memory 10 PCIe HHHL slots (8 Gen5) 8 NVMe U.2 bays Max of 51.2 TB of internal storage L1022: max 25% of cores with other OS 	<p>S1024 & L1024 9105-42A / 9786-42H</p>  <ul style="list-style-type: none"> 1,2-socket, 4U Up to 48 cores per system 12, 16, 24 SMT8 cores/socket 32 DDIMM slots 8 TB memory 10 PCIe FHHL slots (8 Gen5) 16 U.2 NVMe bays Max of 102.4 TB of internal storage Optional Internal RDX Media Bay L1024: max 25% of cores with other OS 	<p>PowerVS</p>  <ul style="list-style-type: none"> Delivered aaS off-premise Unlimited cores Unlimited memory Unlimited storage and I/O Add-on backup services including VTL

AIX

IBMi

Linux



All PowerVM based servers

New IBM Power S1012 System

- New base system design to explore an alternate approach toward both existing and new use cases.
- It is a one-socket, 2U height, half-rack width, equipped with a 1, 4, or 8-core eSCM module, up to 256GB of memory, up to 6.4TB internal storage, 4 PCIe slots, and optional RDX.
- It is offered in two form factors:
 - A rack-mounted option utilizing a bare-metal enclosure.
 - A desktide-mounted option that is 54% smaller than S1014.
- 3-year warranty systems with available value-added support services through Power Expert Care offerings.
- Dual system density unlocks possibilities for architecture best practices approach without compromising on space



Flexibility on Initial Order Configurations

The S1012 introduces an innovative, compact, but flexible design with powerful, secure, and reliable capabilities. It is offered in two form factors and shipped in different set-ups that best suit your customers:

Rack-mounted Options



- For requirements fully covered by one S1012 system in the first months or year, with a ready-to-go expansion provisioned.
- The 2nd bay of the enclosure comes with a filler to avoid dust and inaccurate airflow through the rack.
- 2nd bay can be further occupied by another S1012 if needed with no impact to the existing system.

- For adopting infrastructure best practices such as high availability, advanced remote management.
- For opportunities to expand the footprint beyond its current workload, i.e., with a focus on AI and container-based apps, open-source databases.

- For a backup architecture relying on Tapes, the 2nd bay of the enclosure can be filled with an RDX drive.
- Although it sits externally to the system, it provides an internal RDX-like performance through a direct-attach connection at the back of the enclosure.

S1012 Tower



- Reliable and secure computing capacity in places without racking or datacenter-like infrastructure.
- The tower desk-side enclosure fits one S1012 system and an optional RDX Tape Drive. It has a door lock function, an outside op-panel, a dust filter, and low, office-level acoustics certification.

Rack-to-Tower and Tower-to-Rack conversions are supported when the 4-core processor module is configured. The 1-core processor module is supported only on the Tower form factor. The 8-core processor module is supported only on the rack-mounted form factor

Main Use Cases for S1012

IBM i P05 & P10

- Security for mission-critical data running on IBM i
- Reduce IT footprint in ROBO/back-office
- Attractive price point comparable to Power9
- Improved SLAs to business area/local operation

Small Apps and DBs

- Increase availability and performance: up to 2.5x more performance per core vs. P8
- Lower concerns with data privacy and data loss
- Support app modernization in small spaces

Cloud-to-Near Edge & AI

- Built-in AI acceleration w/o purchasing, maintaining, or noise disturbance of GPUs
- Enhanced data security and remote management capabilities
- Data Lifecycle Management
- Manufacturing equipment management with IoT
- Suitable for Small Language Models and Vision Models



IBM i Subscription Term Licensing

IBM i Licensing Options as of May 7, 2024

Effective May 7, 2024, clients who acquire new specific IBM Power10 processor-based servers at IBM i P05 and P10 software tiers will have the following IBM i Licensing options:

- Acquire or renew IBM i Subscription Term licenses, or
- Convert IBM i non-expiring licenses to IBM i Subscription Term license at a lower-priced subscription term option.
 - Active SWMA on the donor is required.
 - Transfer of IBM i non-expiring licenses is not supported.
 - Conversion pricing applies for 3-year, 4-year, 5-year Subscription Terms.

The specific IBM Power10 processor-based servers at IBM i P05 and P10 software are:

- IBM Power S1012 (9028-21B) 1-core or 4-core processor (P05 software tier)
- IBM Power S1014 (9105-41B) 4-core processor (P05 software tier)
- IBM Power S1014 (9105-41B) 8-core processor (P10 software tier)
- IBM Power S1022 (9105-22A) (P10 software tier)
- IBM Power S1022s (9105-22B) (P10 software tier)

[April 2, 2024 announcement: Conversion pricing is announced](#)

DDR5 and other Hardware Enhancements – July 9th Announcement

Effective 7/19/2024

- New DDR5 memory DDIMM available on the IBM Power10 servers
- New NVMe Enterprise and Mainstream SSD PCIe4 U.2 modules
- New PCIe 4-port 1 GbE Adapter available on IBM Power10 scale-out servers
- PCIe4 4-Port 25/10/1 GbE RoCE SFP28 Adapter with SR-IOV support now available on all Power10 servers
- NED24 NVMe Expansion Drawer – Multipath enablement
- PCIe3 SAS Tape/DVD Adapter Quad-port 6Gb x8 for support with AIX only
- Support for LTO-5, -6 tapes in addition to LTO-7, -8, and -9 tapes with PCIe3 12Gb x8 SAS Tape HBA Adapter



Effective 8/13/2024

- 2-NVMe backplane U.2 (15mm) Flash drive bays

Other Announcements:

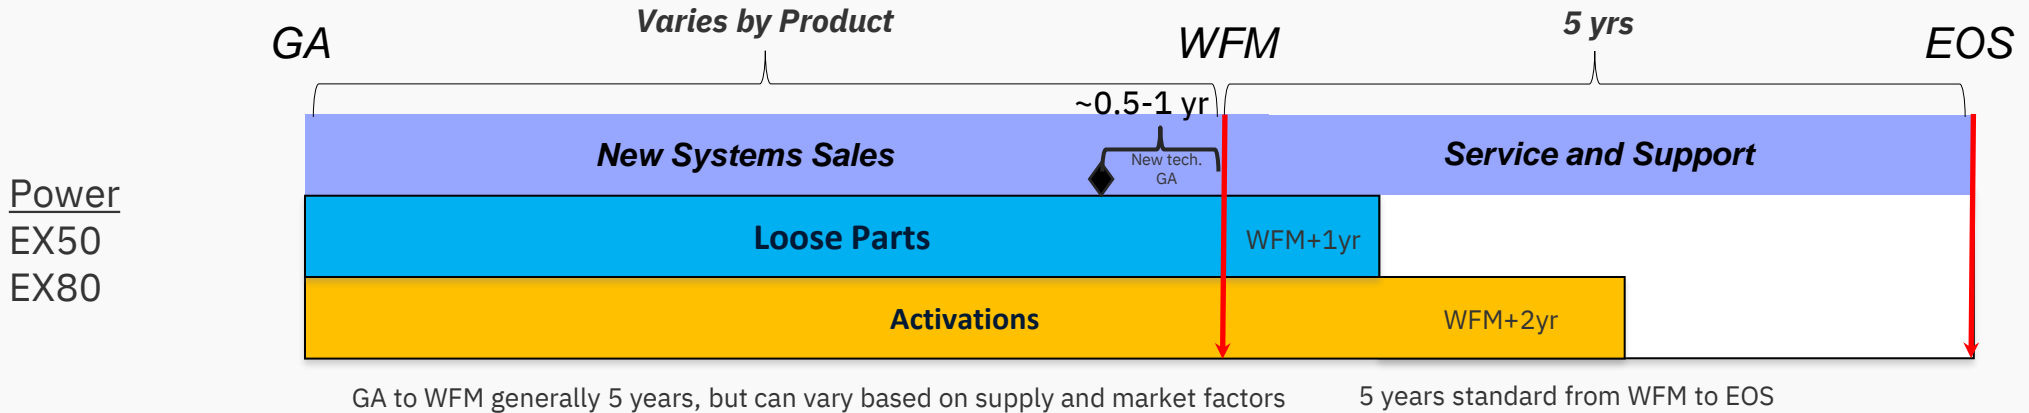
- Hardware withdrawal - IBM Power 9 Activation features (Effective – 12/31/2024)
- Expert Labs Standard Offering default attached to E1050 and E1080 (Effective - 7/9/2024)

Power Systems Life Cycle

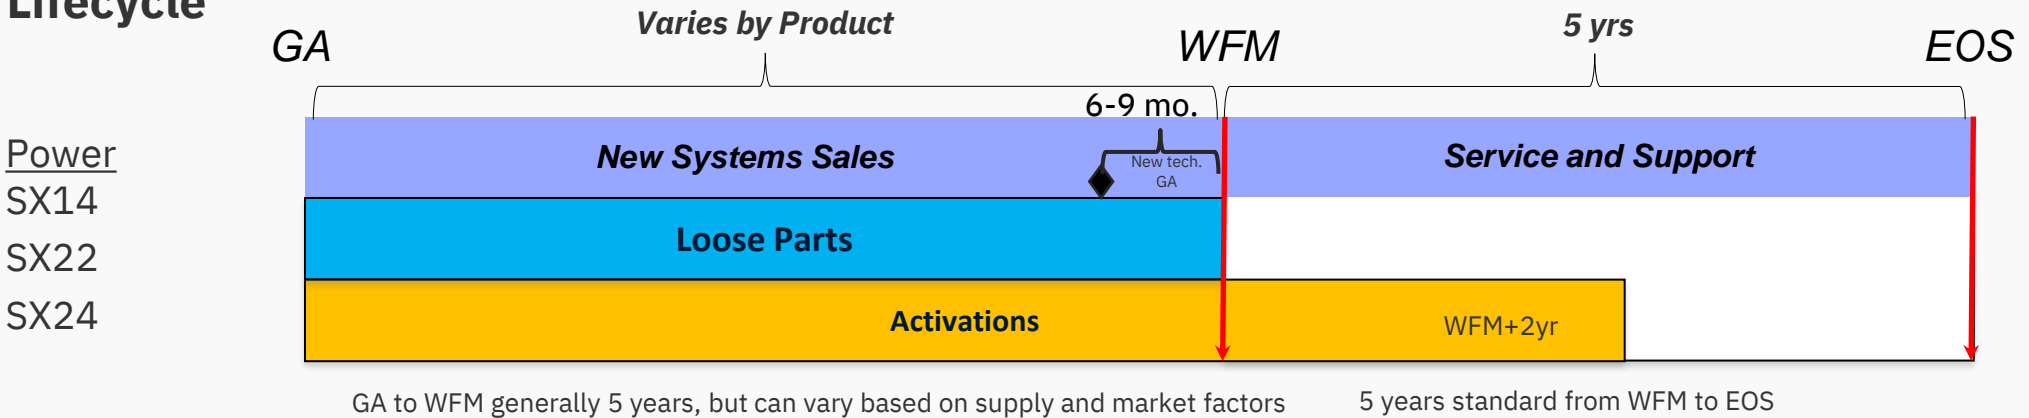
IBM will withdraw all processor, memory and related activations on Power9 systems

- Announce: 7/9/2024
- Effective: 12/31/2024
- Systems: Power E980, E950, and all S924, S922, S914, H924 and H922 models

HE and MR Lifecycle



SCO Lifecycle




Post EOS Support - Service Extensions will be offered in lieu of Hardware Maintenance Agreements or Expert Care

- Separate pricing and more limited coverage
- Coverage is based on several factors, including availability of parts, skills, etc.
- Contracts available in one-year increments
- Per Call pricing will be increased
- Exceptions for price and contract length will be limited only granted in rare circumstances

GA General Availability
WFM Withdrawal From Marketing
EOS End of Support

Accelerate AI Efficiently with IBM Power10 Technical Innovation

On-Chip AI acceleration



Per Core: 4x MMA* & 8x SIMD**

- Accelerate Matrix & Vector Math operations in the processor without GPUs.
- Minimize data movement from processor to GPU and vice-versa.
- Lower quantization for improved performance.


Large Memory & Cache



TBs Main memory per core

- Store large AI models (like LLMs), large datasets (higher batch sizes), multiple AI models (parallel inferencing) in single memory (1TB-16TB).
- Large cache (4x) to speed up execution.

Highly Parallel



4x more Threads per core

- Provides up to 8 threads per core (4x vs. Intel)
- Run parallel jobs for AI inferencing
- Improve throughput

Enterprise Server Design



Lower CO2 footprint


- Same work with lower energy usage
- Consolidate more workloads due to higher guaranteed utilization

*MMA: Matrix Math Accelerator
** SIMD: Single-Instruction Multiple-Data

Improved On-processor AI Acceleration

- 24 TOPS per chip
- 192 TOPS per drawer
- 768 TOPS per system


- Integrated as a CISC instruction
- Remote AI accelerator
- Support for LLM compute primitives
- Int8, FP16 datatypes



AI accelerator

IBM Spyre™ Accelerator PCIe attached card

- 300+ TOPS
- 75W PCIe gen5 x16 adapter
- 128GB of LPDDR5 memory



All projections are based on preview hardware.

IBM Power10 Servers



All Cores can run
Linux, IBM i & AIX

2x S1012 **128GB**

4 core
2x 2port 1/10/25GbE
2x 2port 32Gb FC

RAM

2,7-4x Performance pr
core vs x86

~500k NOK

IBM Power10 Servers

New DDR5
819GB/s per socket
Total 1,6TB/s per server



L1022 2TB
40 core (2x 20core)
2x 2port 1/10/25GbE
2x 2port 32Gb FC
DDR5 RAM

750k NOK

Inkl 5år 24/7

PowerVM Hardware Hypervisor

25%

Cores can run
IBM i & AIX


2,7-4x

Performance pr
core vs x86

3-4x

More bang for
your bucks


IBM Storage



Storage for
Data and AI



Storage for
Hybrid Cloud



Storage for
Data Resilience

Software

IBM Storage Scale
IBM Storage Ceph

IBM Storage Fusion

IBM Storage Defender
IBM Storage Insights

Hardware


IBM Storage Scale System
IBM Cloud Object Storage

IBM Storage Fusion HCI
System

IBM Storage FlashSystem |
DS8000 | Tape | Networking

IBM Storage Today...

Including
Data Lake
functionality



File/Object

IBM Storage Scale




ESS 6000
Up to **310GB/s**



ESS 3500
All Flash Models
Up to **91GB/s**

NL-SAS Expansions



Block

IBM Storage FlashSystem



FlashSystem 9500
Up to **8M IOPS**
0.05ms Latency

809%

More IOPS/Watt
compared to other
vendors



FlashSystem 7300
Up to **3.5M IOPS**
0.05ms Latency



FlashSystem 5300
Up to **2.5M IOPS**
0.05ms Latency



IBM FlashCore Modules Gen4

4.8TB, 9.6TB, 19.2TB, and 38.4TB
SLC to QLC with AI
PCI Gen 4

Compression, **QSC-encryption**
& **ransomware detection**
with no performance impact



IBM Storage Ready Nodes



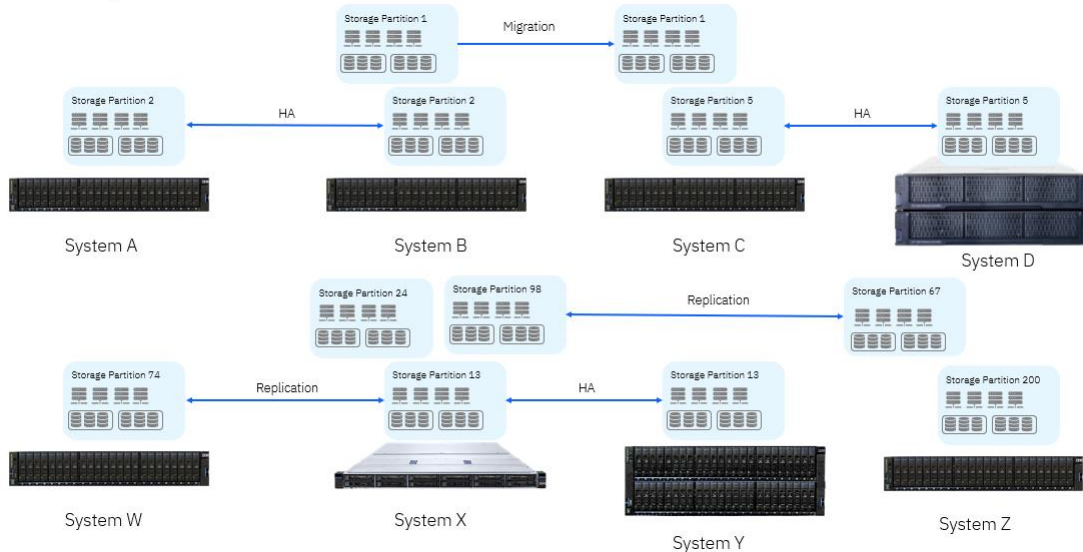
IBM Storage Ceph

IBM Cloud Object Storage

IBM Storage Virtualize 8.7.0 Technical Update

- RFA Announce – 8.7.0 – June 4, 2024
- eGA – 8.7.0 – June 14, 2024
- pGA – There is no hardware associated with this release

8.7.0 is an LTS Release. This means that it will get planned patches and updates until 8.7.0.x code goes EOS. Non-LTS releases will follow until the next LTS release.



Async DR

Introduced previously, Policy-Based Replication provides a simple to use, high performance DR solution

–vVol 2.0 replication enabled

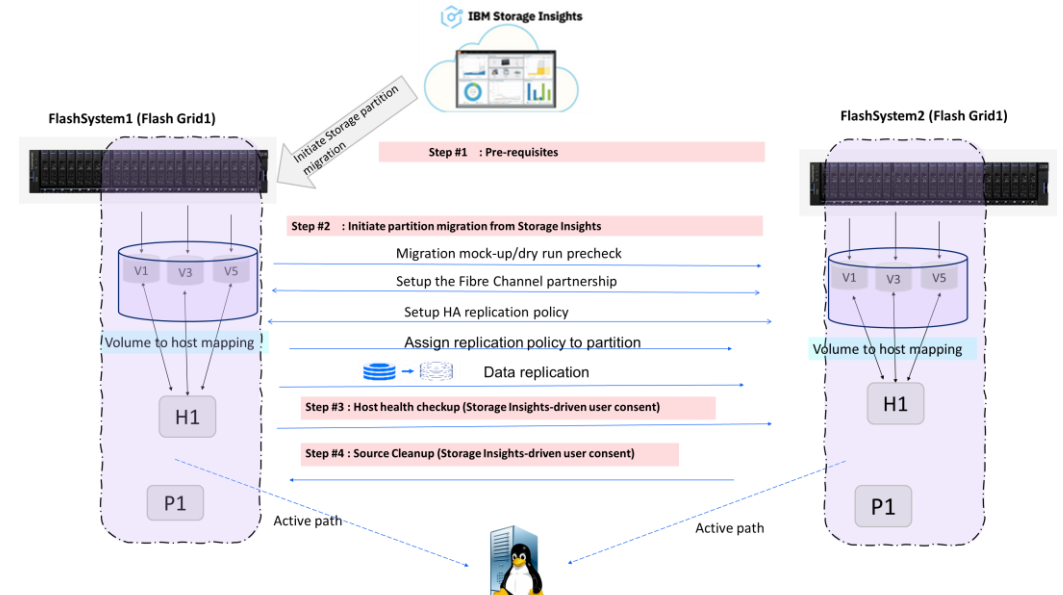
HA

New Policy-Based High Availability provides high throughput, low latency HA with no single point of failure

Coming 2H24

Policy-Based HA + DR for enterprise three-site replication

8.7.0 is the final release that supports legacy Remote Copy features like Global Mirror, Metro Mirror and HyperSwap.



Ransomware Monitoring Architectural Overview

Brought to you by:

IBM Research


IBM Storage Virtualize

IBM FlashCore

IBM Storage Insights Pro

Granular data analytics

Volume Statistics




33

Ex: Create SGC Snap to limit scope

Show Real-Time Data And Trends

Learn From Data

IBM FlashCore Modules



External Tools



Qradar SOAR / Defender

Responses / Actions

Storage Insights



IBM Fusion, two offerings

Customer Apps



IBM Cloud Paks

Data: Business automation
 Integration: Watson AIOps
 Security: Network automation
 Cloud Satellite



Databases

Cassandra
 MongoDB
 RabbitMQ
 Elastic search
 PostgreSQL
 Spark



Off the shelf

Pega
 Mulesoft
 TIBCO



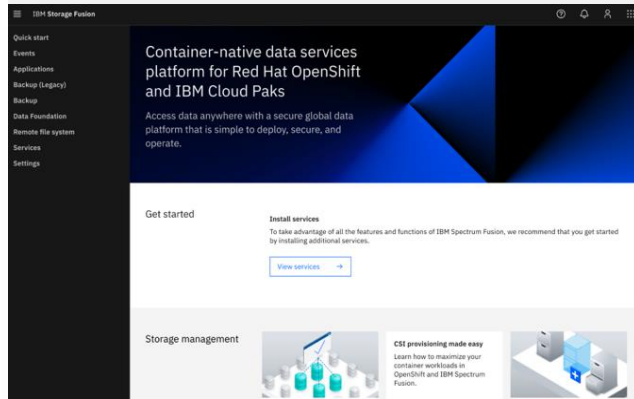
Custom apps

Home grown

Offering

Fusion software

Data services for stateful OpenShift applications

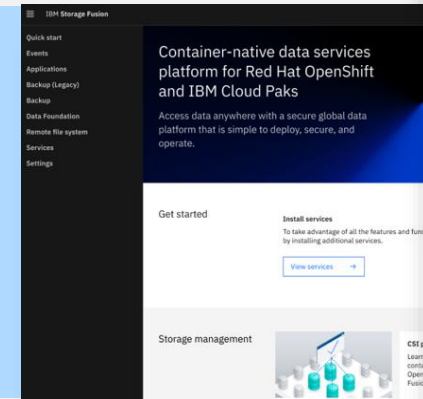


Consistent experiences

- APIs
- Fusion console
- Data protection
- Disaster Recovery
- Fusion Data Foundation (FDF)

Fusion HCI

Integrated Application Platform for OpenShift



Built on



Deployments



AWS EBS



Azure block



Persistent disk



IBM block



SAN, vSAN



DS8K



FlashSystems



Bare metal

Server-attached drives

Hyper-Converged Infrastructure for OpenShift

Switches, servers, storage

Fuel Your AI at the ultimate IBM learning event

IBM TechXchange Conference

October 21-24, 2024

Mandalay Bay – Las Vegas

Register now →


View session catalog →

Join our community and save 30% →

<https://ow.ly/mNHe50T7pes>

Join COMMON at
IBM TechXchange
Conference 2024

Save 30% on Registration
(link in Description)



IBM

COMMON
LEARN. ENGAGE. GROW.

IBM