11

Deliver generative AI solutions on Power with Spyre accelerator & Ecosystem

Thibaud Besson, Technical Presales Power, IBM France

Agenda

1. Historical context
2. Economical context
3. Hardware options
4. Software stacks
5. Turn-key AI with Spyre solution

Historical context

IBM

back in 1993...
2D on CPU
or 3D in computer-aided design (CAD)

1996 – 3D Hardware acceleration on GPU

1996 – Shaders for 3D acceleration

Software code in the GPU that calculates geometries and colors, in real time.

- Vectors for positions and colors,
- Matrices to project 3D positions into the 2D space of the screen.

Leaders of GPU market :

- 3dfx Voodoo Graphics
- ATI 3D
- Nvidia NV1

3Dfx Voodoo Graphics Voodoo 1	
Clock frequency	50 MHz
Bus	64 bits
Dedicated RAM	4 Mo
Resolution	800×600
Communication	PCI
API	Direct3D Shader DirectX3

2003 – GPU for scientific computing

- Vectors and matrices are the basic mathematical objects used in linear algebra.
- Linear algebra studies the linear transformations between these objects.
- It is fundamental in many areas of physics and engineering:
  - ✓ Classical mechanics: oscillators, linear and rotational movements.
  - ✓ Electromagnetism: Maxwell's equations.
  - ✓ Signal processing: Fourier transforms, filters, antennas.
  - ✓ Materials science: stresses, deformations, structural analysis.
  - ✓ Fluid Mechanics and Heat Transfer: Navier-Stokes Equation and Thermal Diffusion Equation
  - ✓ Quantum mechanics: state vectors in Hilbert spaces, observable as Hermitian operators, solutions of the Schrödinger equation.

In the field of artificial intelligence:

- the images are represented by tensors,
- the tokens of a text are vectors in a vector database,
- Deep learning neural networks are activated according to the matrix of each layer, stimulated by an input vector
- The attention mechanism of LLMs is encoded in the form of matrices

GPU seems a natural choice for computing AI !

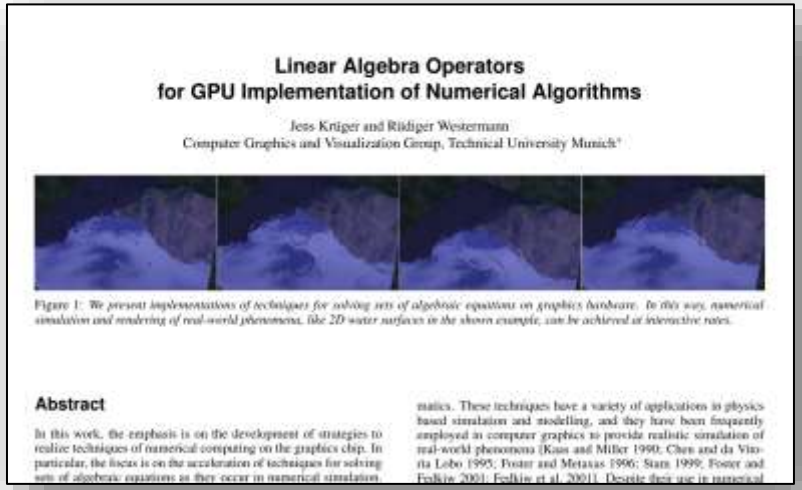

Linear Algebra Operators for GPU Implementation of Numerical Algorithms

Jens Krüger and Rüdiger Westermann
Computer Graphics and Visualization Group, Technical University Munich*

Figure 1: We present implementations of techniques for solving sets of algebraic equations on graphics hardware. In this way, numerical simulation and rendering of real-world phenomena, like 2D water surfaces in the shown example, can be achieved at interactive rates.

Abstract

In this work, the emphasis is on the development of strategies to realize techniques of numerical computing on the graphics chip. In particular, the focus is on the acceleration of techniques for solving sets of algebraic equations as they occur in numerical simulation.

matics. These techniques have a variety of applications in physics based simulation and modelling, and they have been frequently employed in computer graphics to provide realistic simulation of real-world phenomena (Kass and Miller 1990; Chen and da Vitoria Lobo 1995; Foster and Metaxas 1996; Stam 1999; Foster and Fedkiw 2001; Fedkiw et al. 2001). Despite their use in numerical

Economical context

+17 727%

NVIDIA

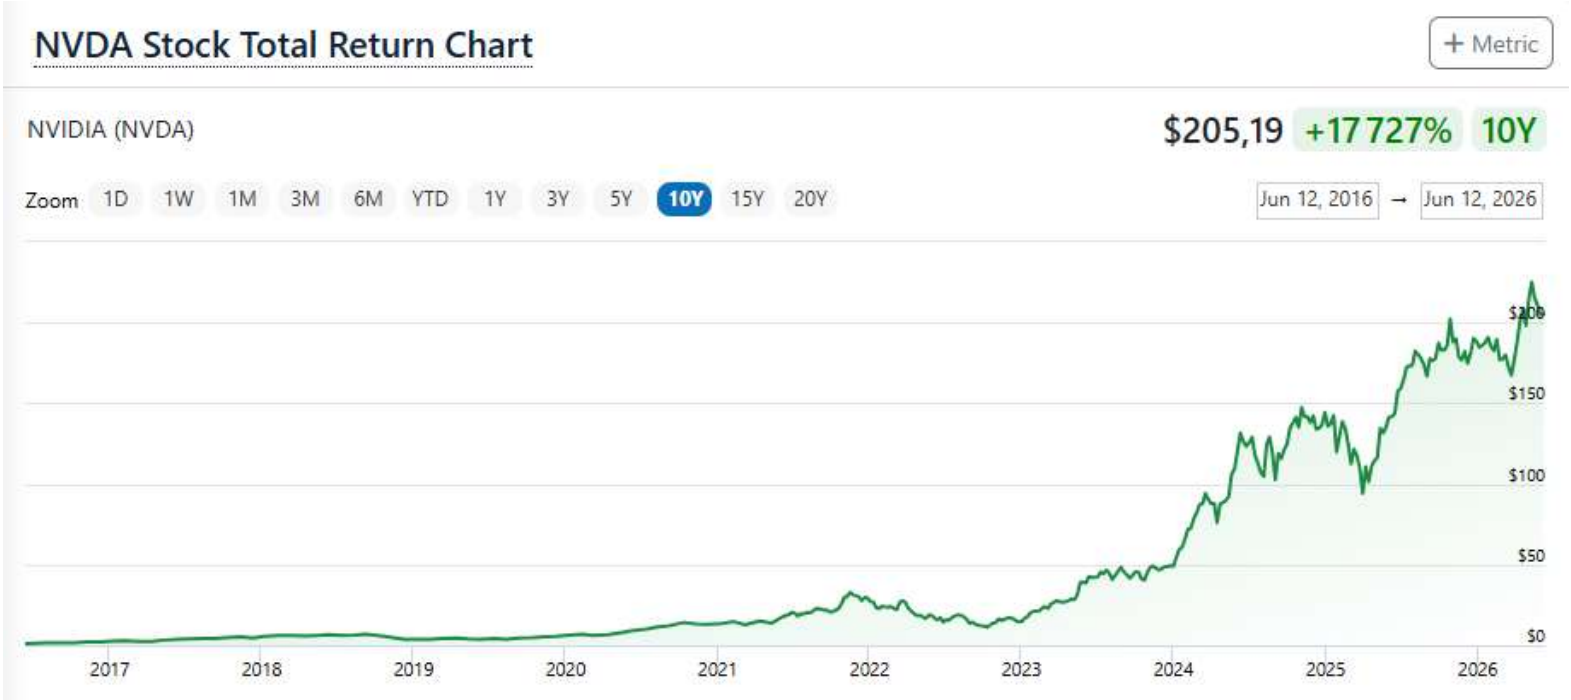

NVIDIA dominates GPU market again in 2026

- Record revenue of $57.0 billion, up 22% from Q2 and **up 62% from a year ago**
- **Record Data Center revenue** of $51.2 billion, up 25% from Q2 and up 66% from a year ago
- **Gaming and AI PC**: Third-quarter Gaming revenue was $4.3 billion, down 1% from the previous quarter and up 30% from a year ago.

Consequences for clients of the quasi-monopoly

1. Price increase on top-of-the-range models
2. Frequent stock shortages
3. Risk of slowing down innovation in the long term

NVIDIA Announces Financial Results for Third Quarter Fiscal 2026

November 19, 2025

Hardware options

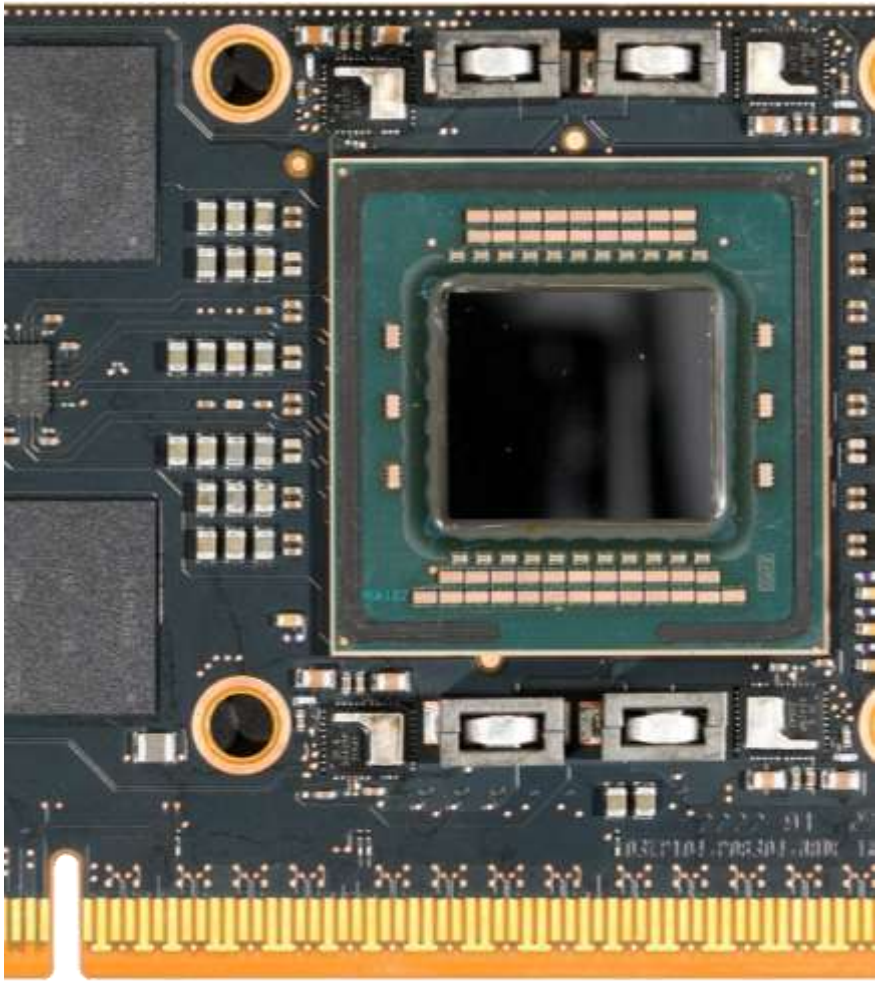

What’s inside NVIDIA hardware ?

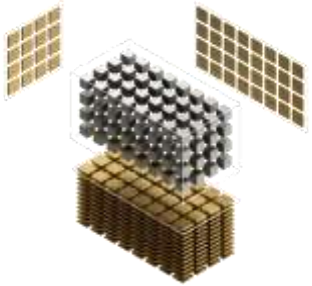

CUDA cores

General-purpose processing units that handle a wide variety of parallel computations, such as graphics rendering, video processing, and scientific simulations.

CUDA cores work primarily with single-precision (FP32) and double-precision (FP64) floating-point numbers, ensuring high accuracy across a broad range of tasks.

Tensor cores

Specialized hardware units designed specifically to accelerate matrix operations, which are central to deep learning and AI algorithms

Tensor cores are built for mixed-precision arithmetic (e.g., FP16, BF16, INT8), allowing them to process lower-precision computations much faster while still maintaining reasonable accuracy by accumulating results in FP32.

For tasks involving large neural networks or real-time AI inference, Tensor Cores are a clear advantage; for more general or non-matrix-heavy workloads, CUDA Cores are sufficient.

Quick history of the NVIDIA portfolio

GPU Model	CUDA Cores	Tensor Cores	Power Usage (W)	First Availability	Memory Size	Memory Bandwidth	Architecture	Approx Retail Price (USD)	FP32 (TFLOPS)	FP16 (TFLOPS)	Tensor FP16 (TFLOPS)	INT8 (TOPS)
NVIDIA V100	5120	640	250	Jun 2017	16–32 GB HBM2	900 GB/s	Volta	$8,000–$10,000	14.1	28.3	113	56
NVIDIA T4	2560	320	**70**	**Sep 2018**	16 GB GDDR6	320 GB/s	Turing	$600–$900	8.1	16.3	**65**	130
NVIDIA A100	6912	432	400	Jun 2020	80 GB HBM2e	2039 GB/s	Ampere	$9,500–$14,000	19.5	39	312	624
NVIDIA RTX A6000	10752	336	300	Oct 2020	48 GB GDDR6	768 GB/s	Ampere	$4,749–$5,299	38.7	77	309	619
NVIDIA A10	9216	288	150	Apr 2021	24 GB GDDR6	600 GB/s	Ampere	$2,800–$3,300	31.2	62.5	125	250
NVIDIA H100	14592	456	**700**	**Oct 2022**	80 GB HBM3	3000 GB/s	Hopper	$25,000–**$30,000**	51	102	1979	3958
NVIDIA L40S	18176	568	**350**	**Oct 2022**	**48 GB** GDDR6	864 GB/s	Ada Lovelace	$7,500–**$8,750**	91.6	183	**733**	1466
NVIDIA L4	7,424	240	**72**	**Mar 2023**	**24 GB** DDR6	300 GB/s	Ada Lovelace	$2,000–**$2,500**	30.3	30.3	**121***	242*
NVIDIA B200	18944	592	**1000**	**Jan 2024**	**192 GB** HBM3e	8000 GB/s	Blackwell	$45,000–**$50,000**	75	150	**2250**	4500

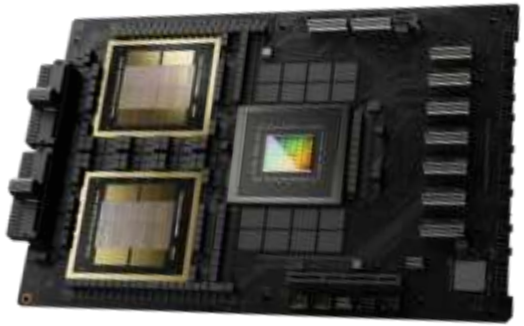

AMD INSTINCT MI300 Series

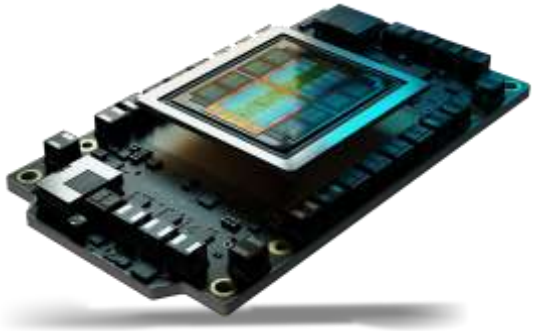

Hardware
Instinct MI Series

RDNA Architecture:

- Gaming

CDNA architecture:

- compute/AI

Instinct MI Series (MI200 / MI300 / MI300X / MI325X)

- Data center GPUs optimized for AI training and inference
- Strong performance on large models (e.g., LLMs) with high HBM memory capacity (MI300X ≈ 192 GB HBM)

Software Stack
ROCm ecosystem

ROCm (Radeon Open Compute)

- AMD's open software platform, alternative to CUDA

Framework support

- Native or optimized support for PyTorch (primary focus) less mature on TensorFlow
- Increasing support for Hugging Face, ONNX, etc.

AI libraries

- Optimized math & AI libraries (MIOpen, rocBLAS, etc.)

AMD offers a credible alternative to NVIDIA for large-scale AI, especially for PyTorch and LLM workloads, combining powerful hardware Instinct MI Series with an open but still maturing software stack (ROCm).

Google TPU

A TPU is an ASIC: application-specific integrated circuit (ASIC) specifically developed by Google for neural networks. TPUs are extremely powerful for large-scale, highly optimized ML workloads, especially in Google Cloud.

TPU is tightly linked to the Google ecosystem:

- Available only in the Google cloud: Data sovereignty concerns, not suitable for strict on-prem environments.
- TensorFlow-first ecosystem: TPUs are deeply optimized for TensorFlow, not PyTorch.
- Limited ecosystem/tooling: Designed for the Google software stack, when Many ML tools are GPU-first.

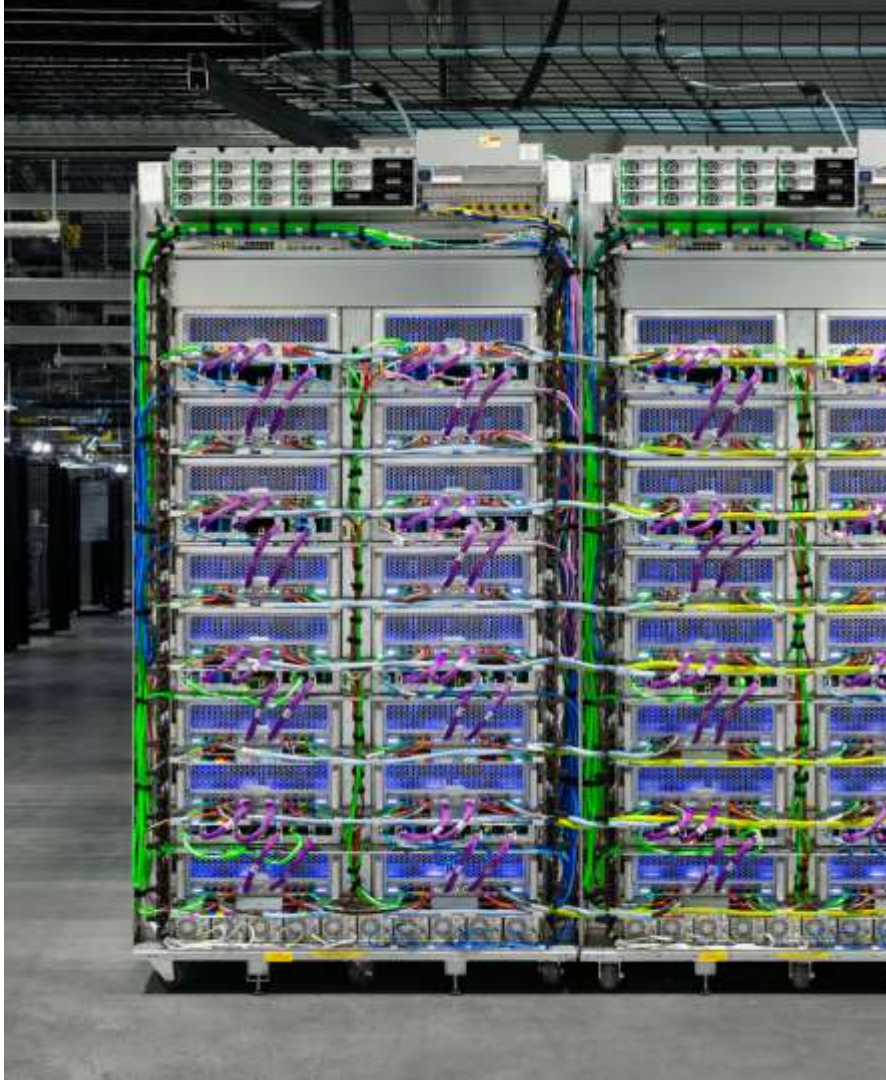

IBM Albany AI Hardware Center

launched in February 2019, with large initial investments from IBM, SUNY Polytechnic Institute, and the state of New York.

Founding members included Samsung, Synopsys, Applied Materials, and Tokyo Electron Limited (TEL).

University at Albany are also crucial partners, and a sizable portion of the AI Hardware Center's work happens at the Albany NanoTech Complex.

The IBM Center’s goal is to develop the next generation of chips, systems, and software to support the future of AI.

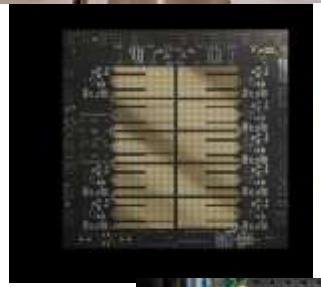

IBM Spyre™ Accelerator is not a GPU.

Message to clients & ISV: *Don't expect it to replace a GPU as a drop-in substitute.*

Message to Business Partners: *Don't position it as a hardware replacement for GPUs.*

IBM Spyre™ Accelerator card

- ASIC for AI Inference
- 360 INT8/FP8 TOPS
- 75W PCIe gen5 x16 adapter
- 128GB of LPDDR5 memory at 205 GB/s

IBM Spyre™ Accelerator is the Hardware for IBM Spyre solution.

Differing infrastructure needs across the AI model lifecycle

	Training	Tuning	Inference
Compute	Large GPU or accelerator clusters	Fewer GPUs (1-10) compared to full training	• GPUs - large models • CPUs • On-die or ASIC accelerators for small optimized models.
Memory	Large memory requirements for GPUs and CPUs	Lower memory needs than training – only updating subsets of parameters	Memory for storing model weights and managing key-value caches
Storage	Sustained high throughput with advanced data sharding and preprocessing of massive datasets	Datasets stored locally or in memory	Datasets stored locally or in memory
Networking	Ultra-fast, low-latency	Minimal networking needs unless doing multi-node tuning	Networking focused on handling incoming user requests
Infrastructure focus	Maximum throughput and scalability, prioritizing hardware utilization and resilience	Fine-tuning prioritizes cost efficiency and iteration speed, enabling rapid experiments and flexible scheduling	• Real-time inference prioritizes low latency, while handling user requests • Batch inference - focus on throughput and cost efficiency.

Software stack

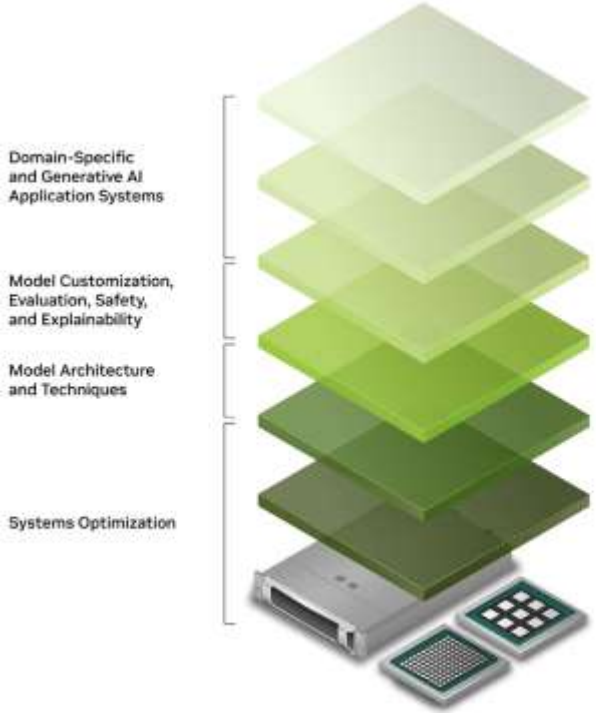

Generative AI stack and options, by NVIDIA

This is NVIDIA's vision

Other options exist, commercial or open source, cloud or on-premise.

At every level of the stack, there are many possibilities, all of which are constantly evolving:

- ✖ So many choices to make!
- ✖ Strong skills required

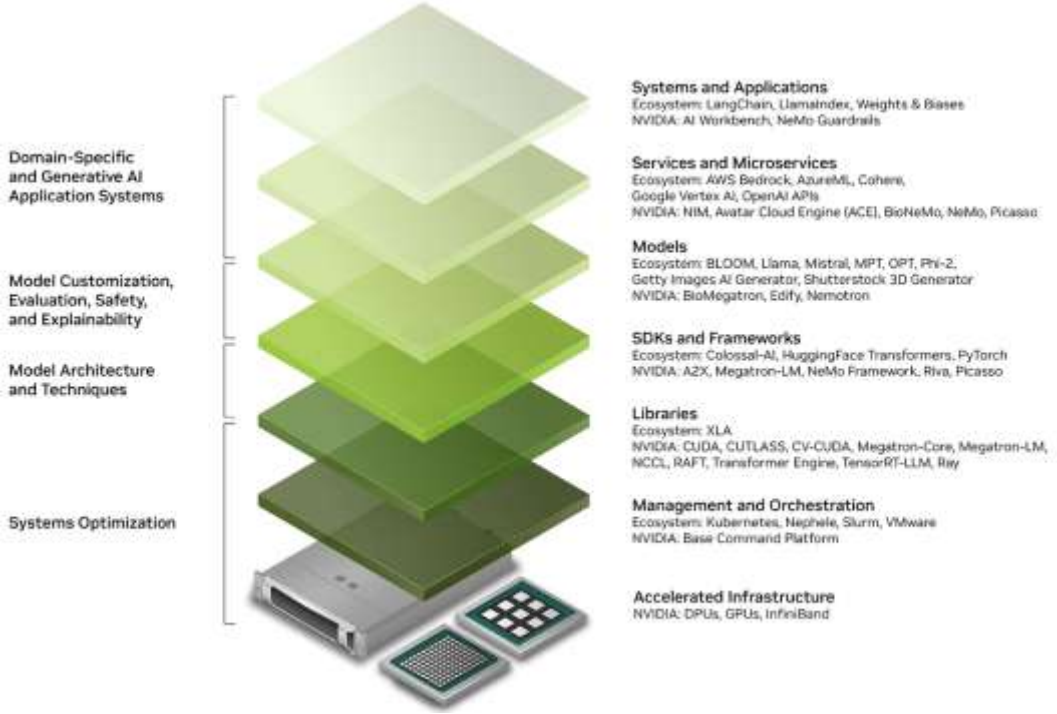

Open source Generative AI stack

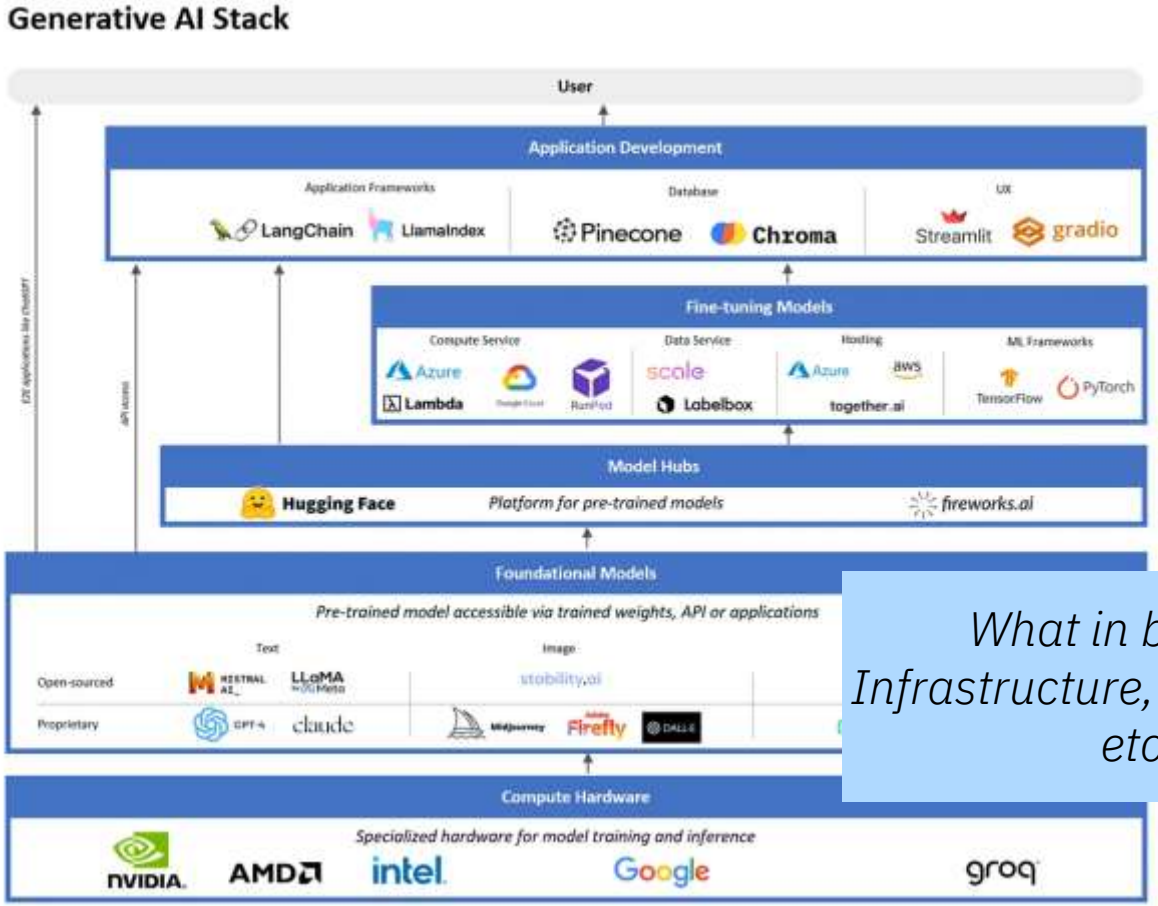

What in between? Infrastructure, orchestration, etc. ?

Another Open source generative AI stack

And below? Infrastructure, orchestration, scaling, etc. ?

And another one...

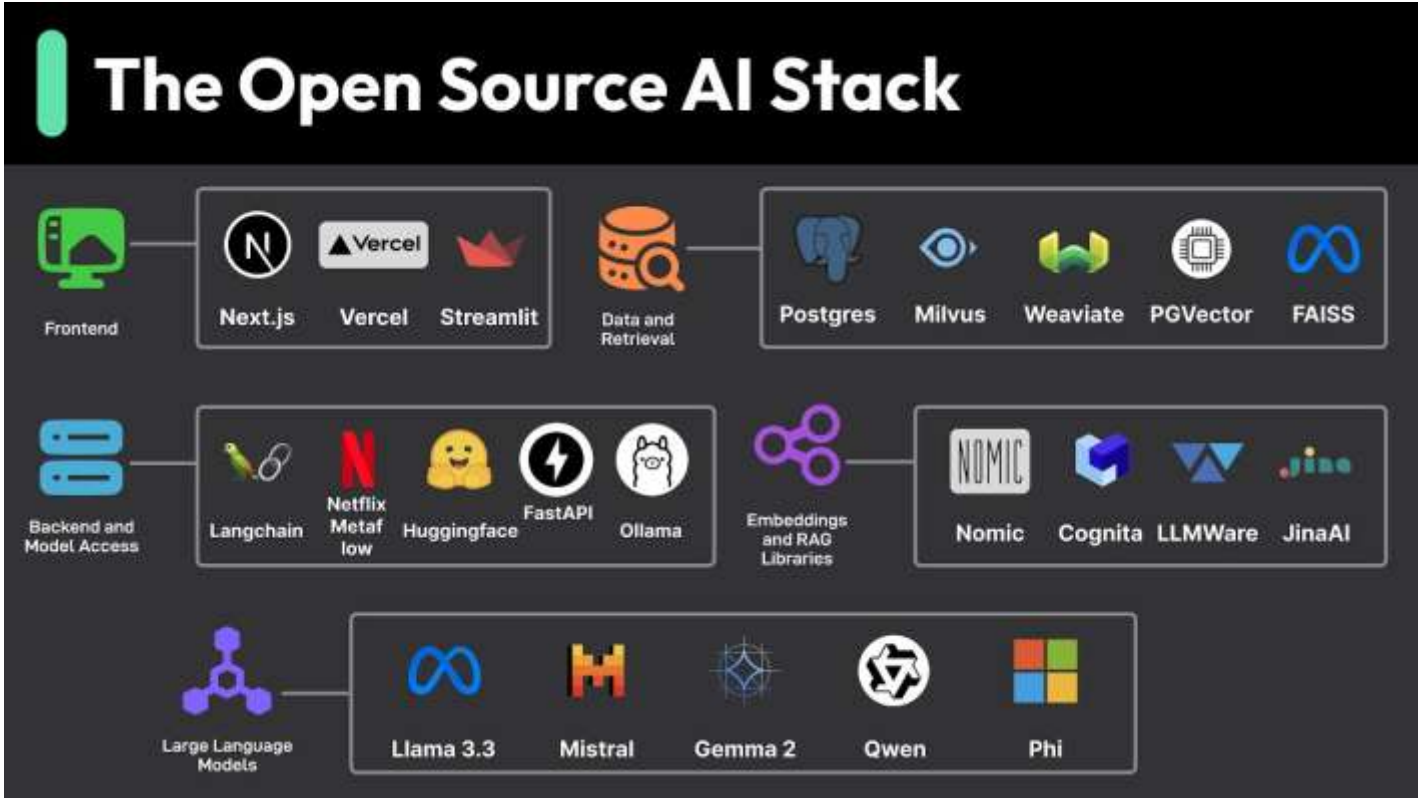

What is happening below? Infrastructure, orchestration, scaling, etc. ?

Only select enterprises have moved AI into production from PoCs, and those leaders are reaping big benefits[1]

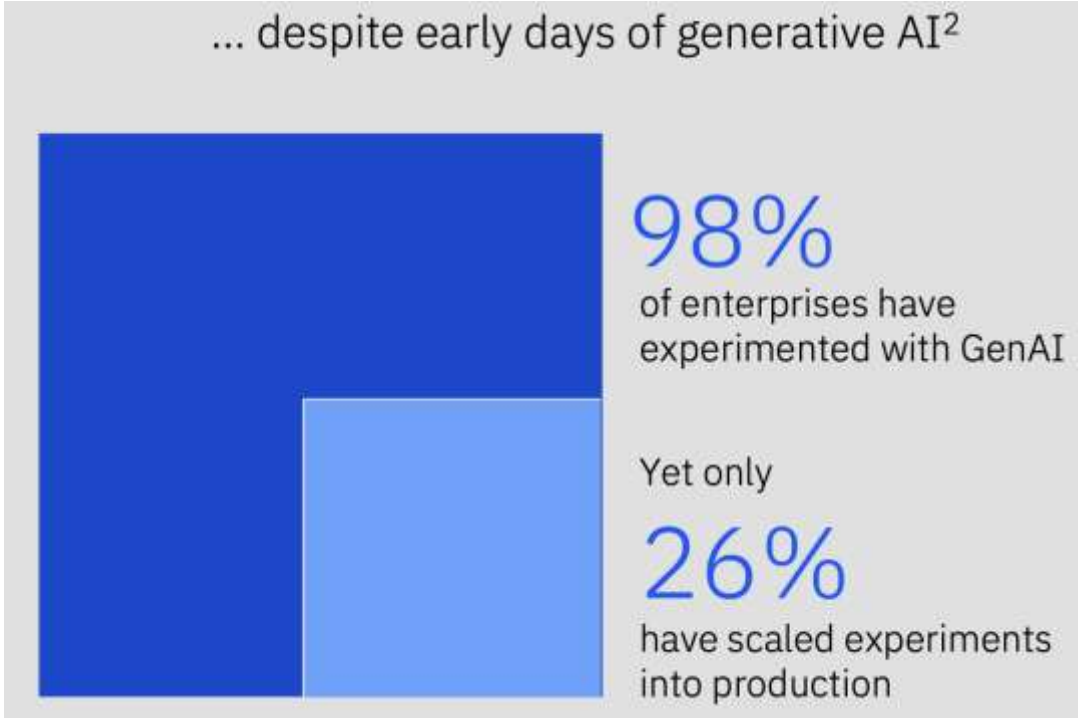

1, 2. BCG “where’s the value in AI” Oct 2024, N=1000

With a poor strategy, AI journeys can take months & fail.

We're **worried** to remain competitive; the market is moving **fast**.

Identify AI needs and objectives

- Clarify needs and pains
- Outline opportunities
- Align compliance context

We are **unsure** where to **start** and **skeptical** about **ROI**!

Align and prioritize use cases

- Assess ROI potential
- Assess feasibility
- Specify success criteria

We need **months** to cope with this **complexity**!

Plan and engineer solutions

- Scope and run POC/MVP
- Design and test solution
- Review continuously

We're iterating since **months** and **miss our ROI** targets!

Evaluate ROI and day 1 ops

- Assess success criteria and incorporate feedback
- Deploy and initialize
- Implement change management

We **cannot** go into production without meeting **compliance** goals!

Use and day 2 ops

- Employ innovation within everyday business tasks
- Validate in-production ROI
- Ensure day-2 operations

Top 5 Barriers to get AI projects started or scale into production

Skills

- Lack of data science or AI engineering skills to integrate AI into current workflows
- No AI-powered assistance to help transform exiting processes

30%

Reduction in coding time when leveraging GenAI Assistants in combination with human coders.[2]

Data Complexity

- Data resides in different locations, different formats with no standardized connectors
- Integrating and transforming enterprise data for AI use cases is complex

75%

Companies investing more in data management due to GenAI[2]

AI Integration Complexity

- Concerns of re-designing existing workflows and disrupting established processes
- Limited AI software availability

Over 90% of organizations face difficulties integrating AI with existing systems[3]

Cost

- Concerns of expensive infrastructure needed to meet performance requirements
- Concerns of cost associated with operationalizing and developing newer skills for AI services

77%

Executives say they need GenAI, only

25% strongly agree that their IT infrastructure can support AI at scale[4]

Privacy & Security

- Data used for training or fine tuning can include sensitive information and moving data off the trusted platform can lead to vulnerabilities
- AI models need to be protected to prevent adversarial attacks

75%

Executives are most concerned with data privacy and security breaches risks with GenAI[5]

1 Emerging Tech: Generative AI Code Assistants Are Becoming Essential to Developer Experience, Gartner, May 2023
2 IBM IBV "The CEO's guide to generative AI" Jul 2023
3 Aura Overcome AI Integration Challenges: Complexity to Competitive Edge
4. IBM IBV "5 Trends for 2025" Dec 2024
5. Data iku "5 GenAI Trends for 2025" 2025

Put AI to work with **IBM Power**

AI-optimized portfolio of capabilities and products to accelerate the impact of AI in core workflows to drive business transformation

IBM Power PoV
to help scale AI in production

AI-ready Infrastructure

Meet performance, security and use case flexibility requirements with AI acceleration and integrated systems

→ HW Acceleration
→ Hybrid Integrations

Optimized Ecosystem Stack

Achieve desired AI outcomes with a range of software stack options, optimized to meet diverse requirements

→ AI optimized runtime and frameworks
→ IBM, RH, and open-source AI software
→ ISV solutions

Data Fabric

Easily access, transform, and manage all enterprise data to scale analytics and AI while keeping it secure

→ Databases
→ Data Intelligence
→ Data Integration
→ Data Security

Assistants and Agents

Leverage new ways of getting work done with enterprise-ready solutions that automate processes and workflows

→ Code Assistants
→ IT Ops Assistants
→ AI-powered Agents

Spyre Solution is here

Watsonx.data is here

Bob assistant is here

Simplify with pre-built open-source AI services...

...integrated with trusted, consistent, and comprehensive inferencing platforms...

...empowered by reliable & safe infrastructure options.

IBM Spyre™ for Power

Catalog with pre-built AI services
(through open-source *or* IBM **Open-Source AI Foundation for Power**)

Red Hat AI

Red Hat AI Inference Server
(on RHEL)

or

Red Hat OpenShift AI
(on Red Hat OpenShift)

IBM watsonx

IBM
watsonx.ai
(on IBM Cloud)

On-premises

IBM **Power11**
on-chip Acceleration

+

IBM **Spyre**
Accelerator for Power

Cloud

IBM **Power**
Virtual Server

+

IBM **watsonx**
as a Service

Spyre Stack

external

mandatory

optional

*removal allowed if purchased via Red Hat. Support for those components will then come via Red Hat.

Services — Support

Applications — Core applications, ISV solutions, ...

Application integration

5639-SAI
IBM Open-Source AI Foundation for Power
(catalog with pre-build AI services, adoption patterns, and integration packages & tools such as Python wheels)

	2025	2026	1Q 2026	2026+
Inference server	5639-RIS Red Hat AI Inference Server*	5639-SAI Open-source AI foundation (vLLM + 5639-SPY; to be committed)	5639-OAI Red Hat OpenShift AI (committed)	PIDs TBD IBM watsonx products (to be committed)
Operating system	5639-1RE/3RE/5RE RHEL 9.6+*	5639-1RE/3RE/5RE RHEL 9.6+*	CoreOS	CoreOS

Middleware

5639-SPY
IBM Spyre™ Enablement Stack for Power
(driver, backend, runtime, firmware, models, ...)

5639-SPA
IBM Spyre™ Stack Updates for Power
(for driver, backend, runtime, firmware, models, ...)

HMC

PowerVM

Hardware

9080-HEU, 9043-MRU, 9856-42H, 9824-42A, 9856-22H, 9824-22A
IBM Power11 server

ENZ0 + 2xENZA + ECLR / ECLS / ECLX / ECLY / ECLZ + EJ24 / EJ2A
I/O drawer + cabling + adapter

8xECSE
8x IBM Spyre™ Accelerator for Power

Spyre Security

Control: 100% control to unlock trusted AI services where enterprise data resides (on-premises or IBM PowerVS).

Air-gapped: Install & operate AI services without internet access.

Support: One-stop-shop supported & integrated enterprise AI stack.

Network security: 100% native IBM Power ecosystem without data exposure to any network between Spyre drawer and IBM Power server.

Enterprise Support

ibm.com/support

File tickets for **any** component of the Spyre stack to get help!

Example 1: Software issue

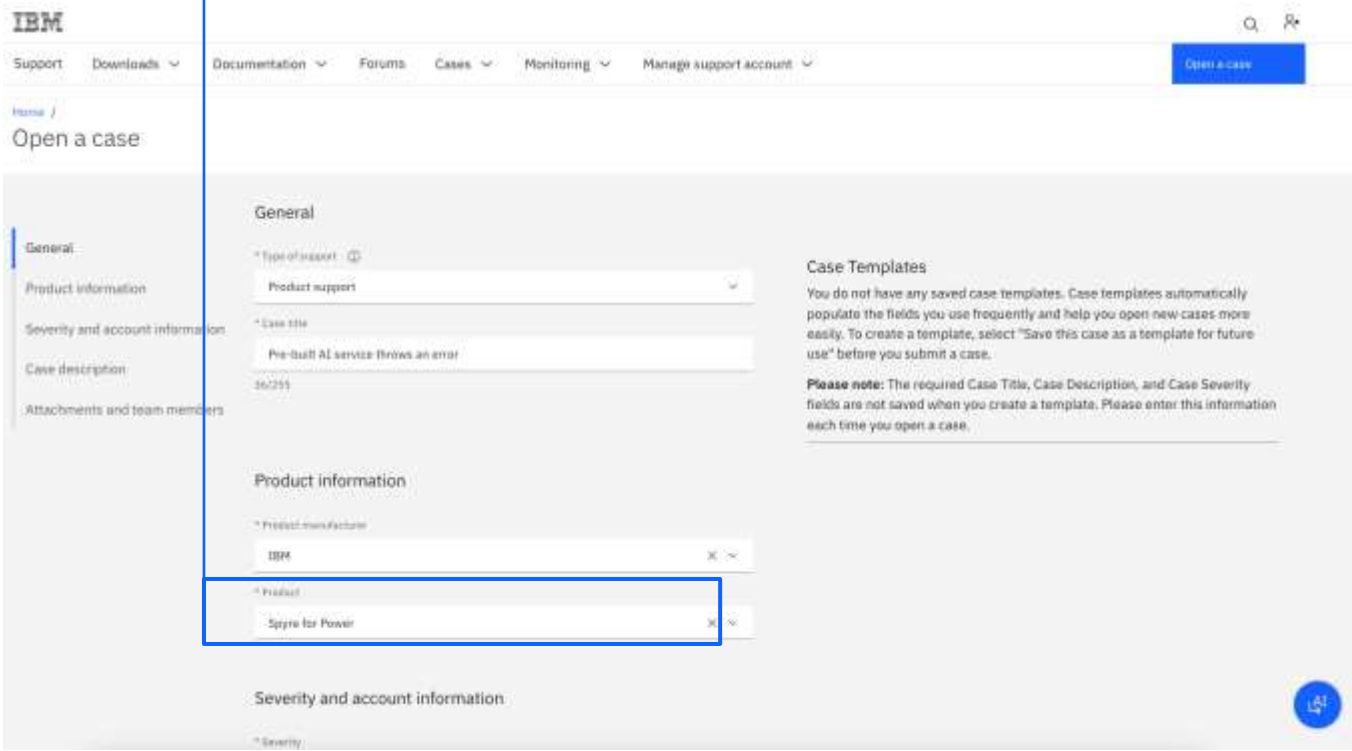

Enterprise Support

ibm.com/support

File tickets for **any** component of the Spyre stack to get help!

Example 2: Hardware issue

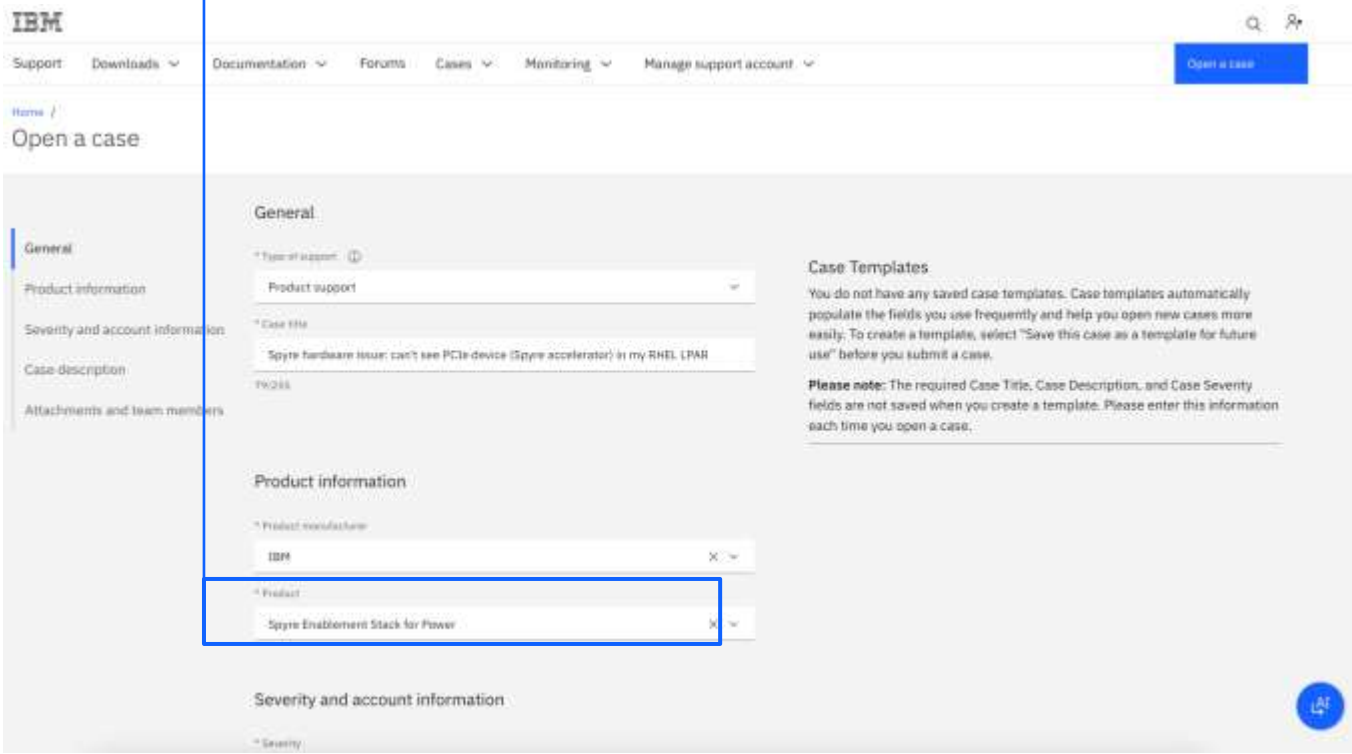

IBM Spyre™ Solutions catalog

IBM has optimized accuracy & performance for enterprise use cases.

Patterns serving Enterprise use cases

Services based on AI capabilities

Predefined Models

Optimized

Available

Functionally working and exploring for enterprise use cases.

Tested

DIY exploration and experimentation of AI capabilities.

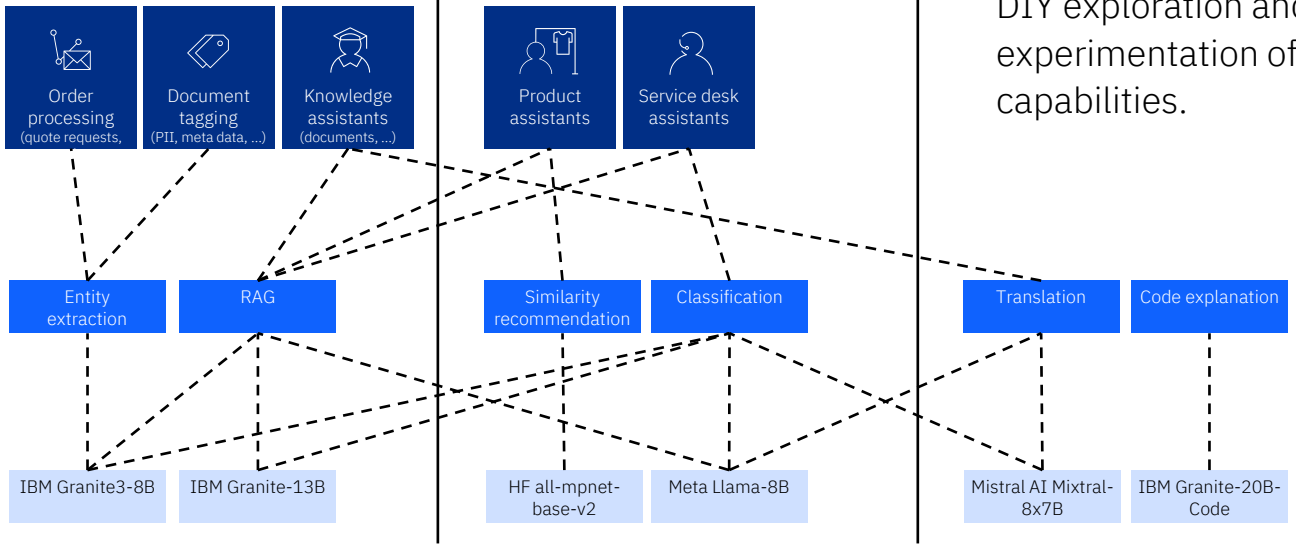

Digital Assistant pattern

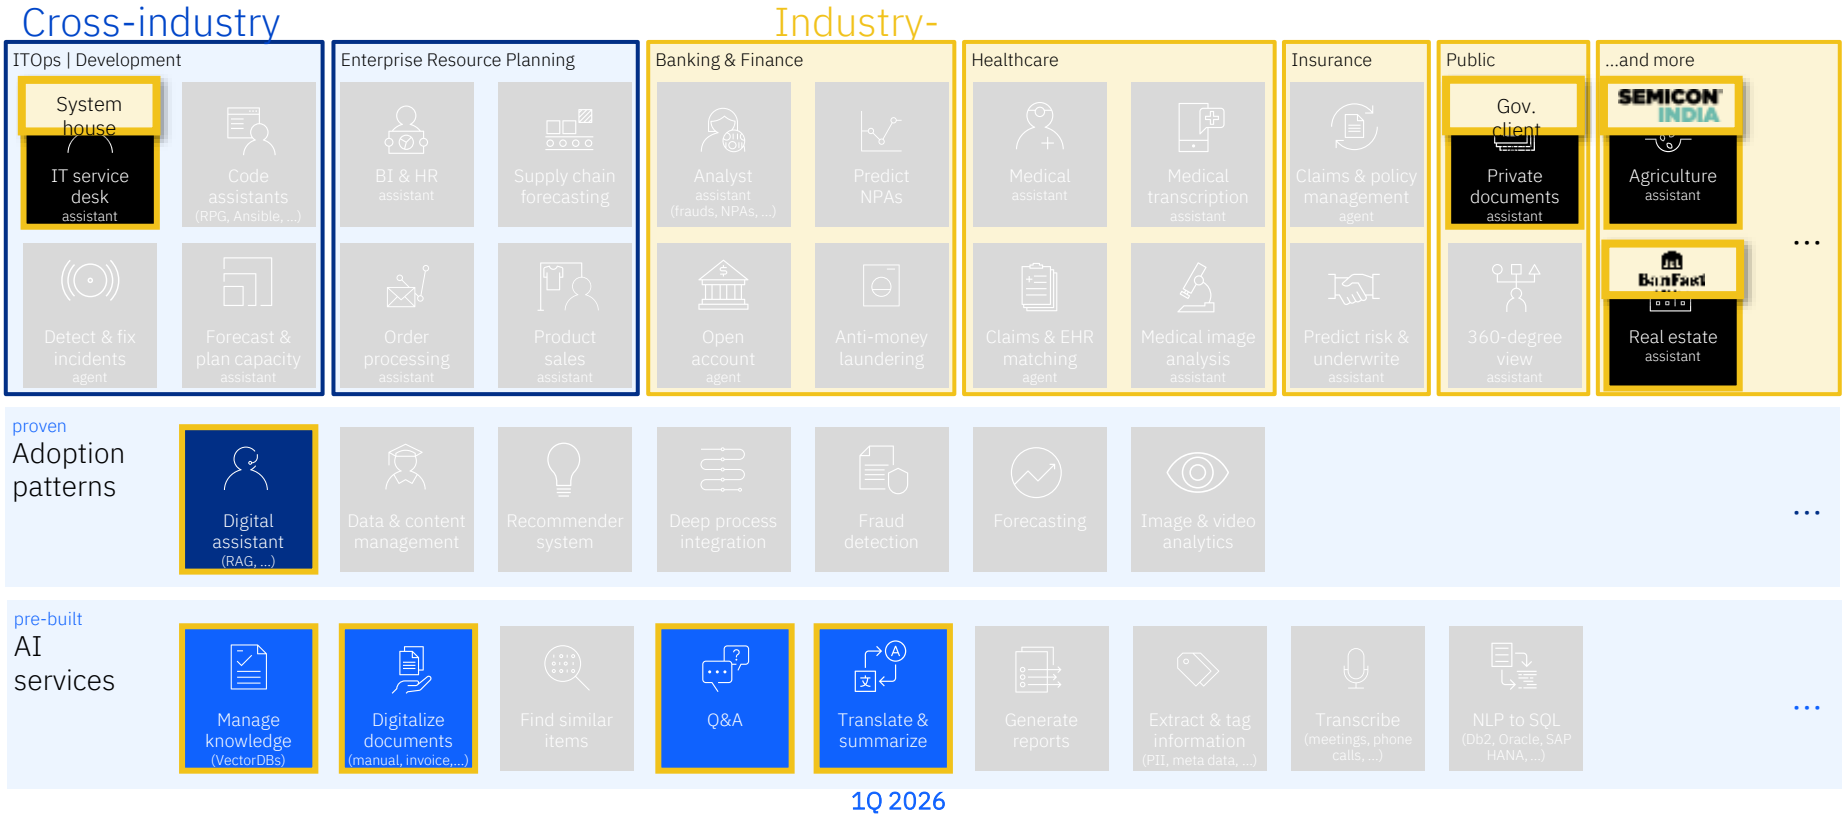

Client examples & demos

aligned with core workloads; not restricted to Spyre™.

Cross-industry

ITOps | Development

- System house (DACH) — IT service desk assistant
- Code assistants (RPG, Ansible, ...)
- Detect & fix incidents agent
- Forecast & plan capacity assistant

Enterprise Resource Planning

- BI & HR assistant
- Large retailer US — Supply chain forecasting
- Order processing assistant
- Large retailer US — Product sales assistant

Industry-specific

Banking & Finance

- Demo TechXchange 2024 — Analyst assistant (frauds, NPAs, ...)
- Finacle — Predict NPAs
- Open account agent
- Anti-money laundering

Healthcare

- SHIBUYA by PULSEN — Medical assistant
- Medical transcription assistant
- Claims & EHR matching agent
- Medical image analysis assistant

Insurance

- Demo TechXchange 2025 — Claims & policy management agent
- Predict risk & underwrite assistant

Public

- Gov. client (DACH) — Private documents assistant
- 360-degree view assistant

...and more

- SEMICON INDIA — Agriculture assistant
- BanFast — Real estate assistant
- ...

proven

Adoption patterns

- Digital assistant (RAG, ...)
- Data & content management
- Recommender system
- Deep process integration
- Fraud detection
- Forecasting
- Image & video analytics
- ...

pre-built

AI Services

Open Source containers available from in IBM cloud registry

- Manage knowledge (VectorDBs)
- Serve models
- Digitalize documents (manual, invoice,...)
- Find similar items
- Q&A
- Translate & summarize
- Generate reports
- Extract & tag information (PII, meta data, ...)
- Transcribe (meetings, phone calls, ...)
- NLP to SQL (Db2, Oracle, SAP HANA, ...)
- ...

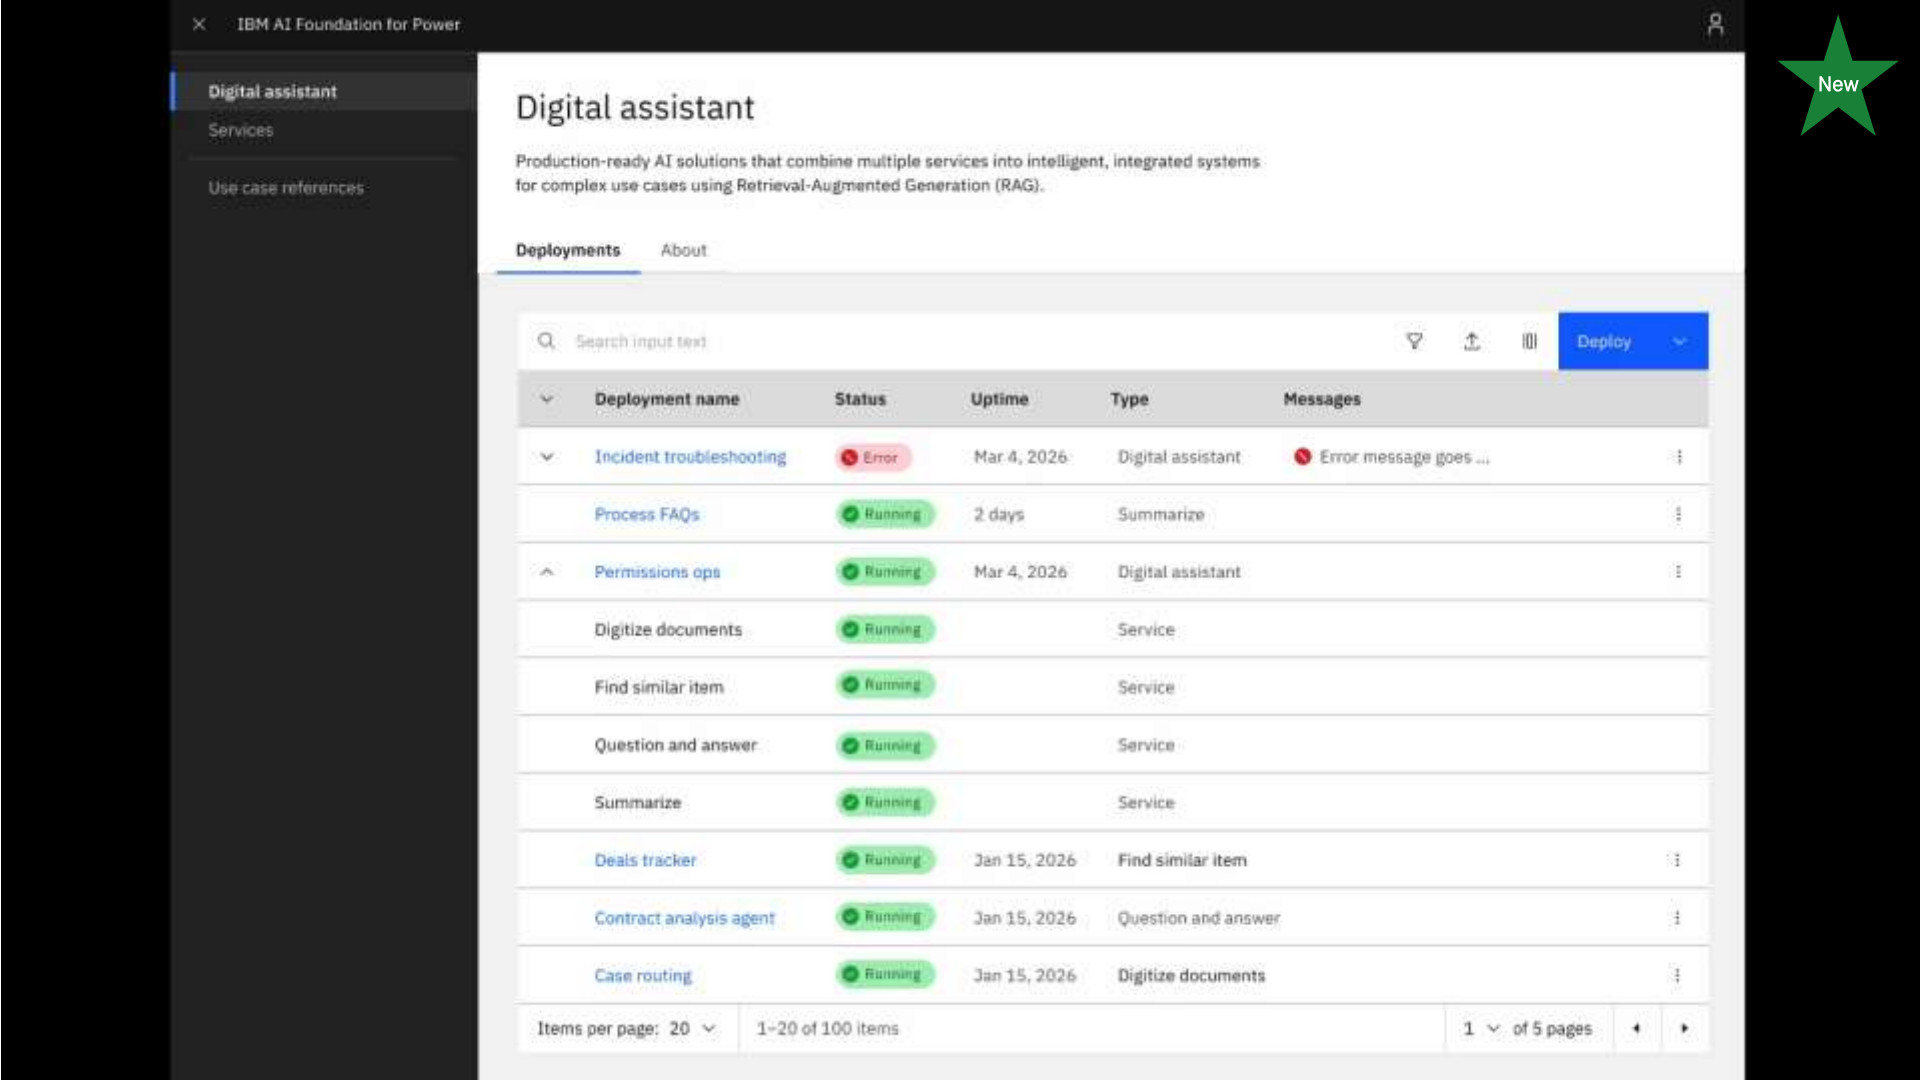
IBM AI Foundation for Power
Digital assistant
Services
Use case references
Digital assistant
Production-ready AI solutions that combine multiple services into intelligent, integrated systems for complex use cases using Retrieval-Augmented Generation (RAG).
Deployments
About
Search input text
Deploy
Deployment name
Status
Uptime
Type
Messages
Incident troubleshooting
Error
Mar 4, 2026
Digital assistant
Error message goes ...
Process FAQs
Running
2 days
Summarize
Permissions ops
Running
Mar 4, 2026
Digital assistant
Digitize documents
Running
Service
Find similar item
Running
Service
Question and answer
Running
Service
Summarize
Running
Service
Deals tracker
Running
Jan 15, 2026
Find similar item
Contract analysis agent
Running
Jan 15, 2026
Question and answer
Case routing
Running
Jan 15, 2026
Digitize documents
Items per page: 20
1–20 of 100 items
1 of 5 pages
New

New

Digital assistant
Services
Use case references

Digital assistant

Production-ready AI solutions that combine multiple services into intelligent, integrated systems for complex use cases using Retrieval-Augmented Generation (RAG).

Deployments | About

Deploy

Card title

Services
- Digitize documents
- Find similar items
- Question and answer
- Summarize

Use case references
- Agriculture assistant
- Analyst assistant
- BI assistant
- Claims & policy management agent
- Conference slide search
- Financial documents assistant
- Invoice matching assistant
- IT service desk assistant
- Open account agent
- Order processing assistant
- Private documents assistant
- Product sales assistant
- Real estate assistant

Minimum resource allocation

Required cores	Required memory	Required Spyre cards
0.5 - 2.0	15GB -25GB	4 cards

Code and architecture

View code

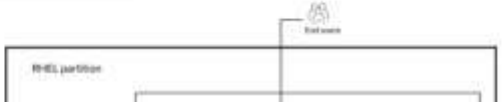

Demos and prototypes

Simplify with pre-built open-source AI services...

...integrated with trusted, consistent, and comprehensive inferencing platforms...

...empowered by reliable & safe infrastructure options.

IBM Spyre™ for Power

Catalog with pre-built AI services
(through open-source *or* IBM **Open-Source AI Foundation for Power**)

Red Hat AI

Red Hat AI Inference Server
(on RHEL)

or

Red Hat OpenShift AI
(on Red Hat OpenShift)

IBM watsonx

IBM
watsonx.ai
(on IBM Cloud)

On-premises

IBM **Power11**
on-chip Acceleration

+

IBM **Spyre**
Accelerator for Power

Cloud

IBM **Power**
Virtual Server

+

IBM **watsonx**
as a Service

The future of AI is open

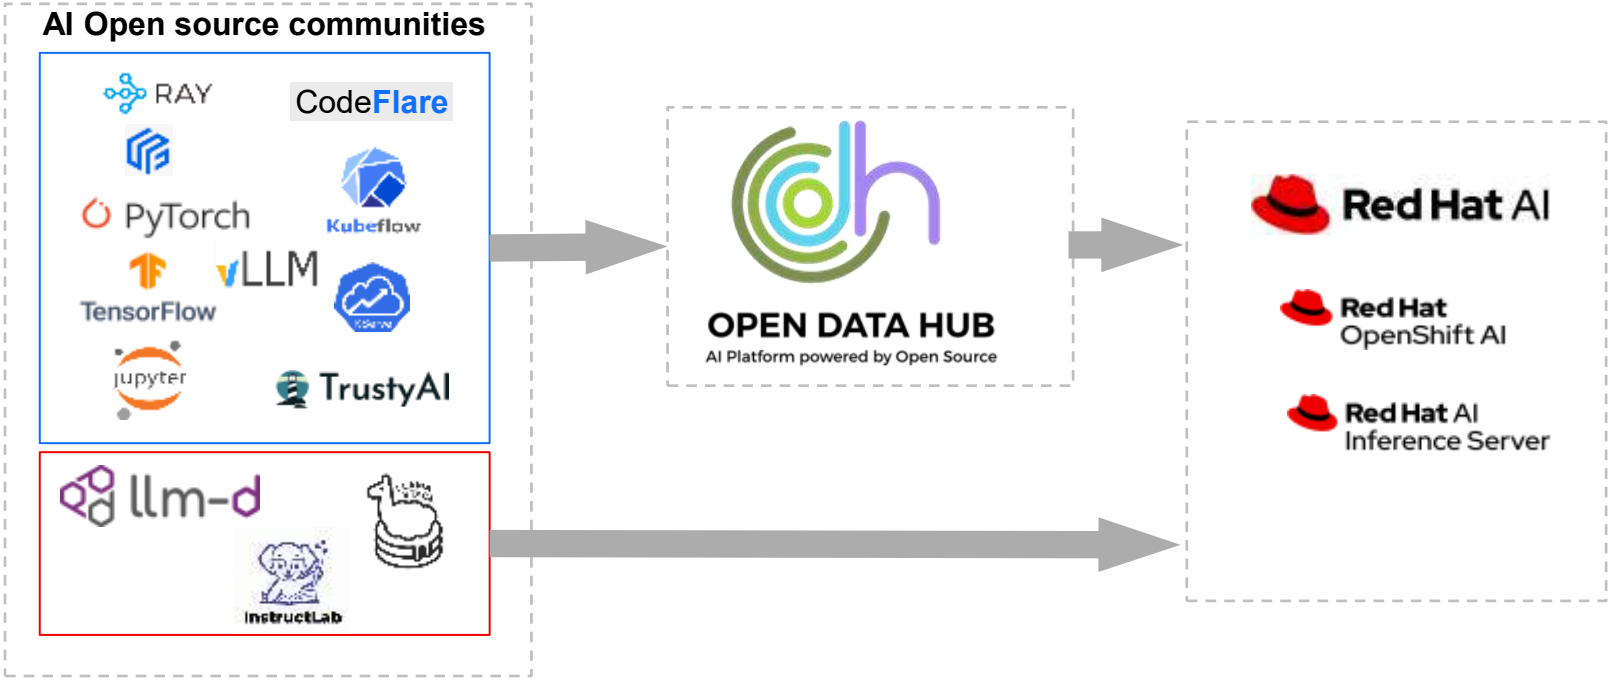

Leverage speed of open-source → Enhance UX for open-source → Make enterprise-ready

Red Hat

Red Hat AI (RHAI) - Overview

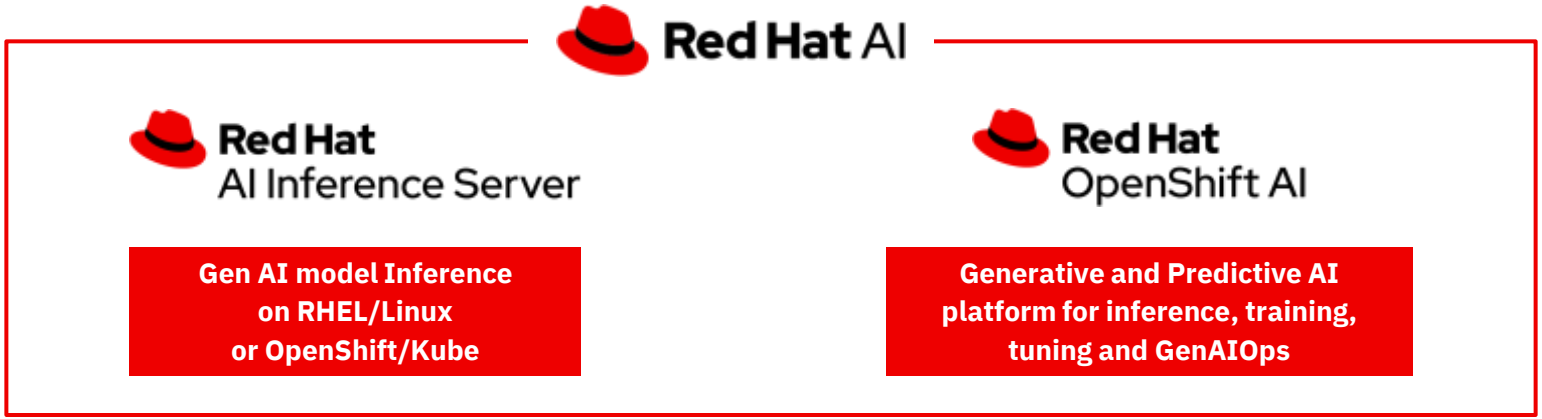

Trusted, Consistent and Comprehensive foundation

NVIDIA AMD intel Hardware Acceleration Google aws IBM

Physical Virtual Private Cloud Public Cloud Edge

Red Hat AI Inference Server (RHAIIS) – What is it?

- → Supports IBM Spyre™ for Power only *(not for MMA)*
- → Runs on RHEL
- → "Get-started" option for Power clients

RHAIIS – Why vLLM?

High-performance

- → Natively optimized for Spyre
- → Advanced algorithms to improve inference throughput

Drive developer and IT Productivity

- → Simple APIs for online (real-time) and offline (batch) inference
- → Standard API format (Open-AI API compatible)

vLLM is emerging as *the Linux* for GenAI Inference.

Red Hat OpenShift AI (RHOAI) – What is it?

One platform for generative and predictive AI across the *hybrid cloud*.

Features

Value

All AI/ML services in one product →

Red Hat OpenShift AI

→ *Improve data scientist & developer productivity*

Platform management services →

Red Hat OpenShift

→ *Improve DevOps productivity*

In-built support for multiple HW accelerators →

NVIDIA intel AMD Google IBM aws

→ *Native fit-for-purpose accelerator support*

Flexible deployment →

Virtual machine (on-prem)

Private Cloud (PowerVS)

Others (public cloud, sovereign cloud, etc)

→ *Easily move AI to where the data lives*

Red Hat

Red Hat SW for MMA & Spyre – when to use *Red Hat Openshift AI (RHOAI) & Red Hat AI Inference Server (RHAIIS)*

RHOAI

Client Value	*Build, run* and *manage* AI flexibly	*Easily integrate AI* with modernized applications	*Accelerate time-to-value* with fit-for-purpose hardware acceleration.
Key Capability	Single software for Model training*, inference and monitoring on-premise or on IBM PowerVS.	Run AI on the same Power platform as modernized apps	Unified software stack to seamlessly move workloads between MMA (on-chip) and Spyre (off-chip) accelerators for optimal performance.

RHAIIS

Client Value	*Simple to get started*	*Flexible*	*Optimized*
Key Capability	Single container (vLLM) to run genAI use cases	Can run on RHEL or Openshift	Tuned to run on IBM Spyre for targeted genAI use cases

Simplify with pre-built open-source AI services...

...integrated with trusted, consistent, and comprehensive inferencing platforms...

...empowered by reliable & safe infrastructure options.

IBM Spyre™ for Power

Catalog with pre-built AI services
(through open-source *or* IBM **Open-Source AI Foundation for Power**)

Red Hat AI

Red Hat AI Inference Server
(on RHEL)

or

Red Hat OpenShift AI
(on Red Hat OpenShift)

IBM watsonx

IBM **watsonx.ai**
(on IBM Cloud)

On-premises

IBM **Power11** on-chip Acceleration

+

IBM **Spyre** Accelerator for Power

Cloud

IBM **Power** Virtual Server

+

IBM **watsonx** as a Service

Simplify AI on-boarding

by moving between cloud & on-premises seamlessly.

Explore.

minutes
change one configuration

Explore cloud-based AI services
IBM Cloud, AWS, Azure, GCP, OpenAI, ...

Test.

minutes
change one configuration
(when Spyre * infrastructure is available)

Co-locate AI services with system of records
IBM Power on-chip acceleration

Launch.

Promote to production
IBM Power11
+ IBM Spyre™ for Power

Discover turnkey AI with IBM Spyre™ for Power.

INFORM YOURSELF

Explore AI on IBM Power yourself:

ibm.com/products/power/ai

30 minutes or less

GET BRIEFED

Meet with an IBM AI expert for custom demonstrations of IBM Spyre™. Understand where AI can help boost productivity & drive growth without long pilots.

2 – 4 hours

SEE VALUE

Run a use case alignment workshop, try AI services for your use cases at the same day, and plan the enterprise integration.

1 day

---

common EUROPE

Sebastian Lehrig

WW AI on IBM Power Leader Germany

Monday, June 15

17:05 CEST

2026 Client Outcomes for AI with IBM Power
06 - RHÔNE 3B
Sebastian Lehrig

Tuesday, June 16

11:40 CEST

AI Acceleration for Power with Spyre
08 - ST CLAIR 3A
Sebastian Lehrig

Thank you