

Sovereign AI Factory Strategic and technical benefits for business

Wojciech Stramski, Beyond.pl CEO



Certifications and standards



Partners



Organizations

Available options to launch AI on the market

On-premise



- Full control over data
- Private/secure



- Significant investment costs (hardware + potential data center, people)
- Less scaling flexibility
- Technical complexity
- Lack of specialists on the market

Hyperscalers



- Lower cost of entry
- Global presence
- Scalability



- Data privacy
- Agreement based on foreign jurisdiction (non- EU)
- Geopolitical risks
- Inflexible SLA
- High operating costs, including exit costs
- Vendor dependency (technological lock-in)
- Limited GPU availability




beyond.pl


AlaaS | GPUaaS | GPU Colocation



- Agreement based on EU/Polish law
- Access to the most efficient infrastructure (B200) + scalability
- Limited technological risk
- Predictable costs, up to 50% lower than with hyperscalers
- No exit costs
- Hybrid solutions (integration with hyperscalers possible)
- English technical support (24/7)
- Flexible SLAs




AI Service Ecosystem: Hyperscalers vs. Beyond.pl's AI Factory



Strategic AI challenges for every institution and business



Added value of an AI Factory



Cost savings




**Accelerated
implementation
(time-to-market)**



**Data security and
control**



**Compliance
with legal
regulations**



Flexibility

Cost savings

- No upfront investment in infrastructure or specialists to maintain it
- Ability to **scale the infrastructure** relative to current needs
- Access to **software** that ensures optimal **utilization** of the infrastructure
- No outgoing data transfer costs (no egress fees)
- No **exit costs**
- **Transparency and predictability**



Acceleration of deployment (Time to Market)

- Fully commissioned, integrated, optimized based on NVIDIA reference design infrastructure
- Deployed and proven software
- Pre-trained, validated, and optimized AI models
- Reduced technological barriers of entry
- Automated implementation processes



Data security (data sovereignty)

- Provider origin and jurisdiction
- Infrastructure **location (EU)**
- Full control of data **migration and deletion**
- **Legal jurisdiction** applicable to the signed agreement



Compliance with legal regulations

- Compliance with applicable European and national regulations (including [AI Act](#), [NIS2](#), [DORA](#), [GDPR](#), etc.)
- Ready for [future regulatory changes](#)
- Transparency and auditability of AI systems



Security and compliance of services with legal requirements (AI Act, NIS2, DORA)

	GPU Colocation GPU Servers	AI Services
Business Case	<i>depending on the division of responsibility</i>	Sandbox/POC Training Inference
Data Security	at rest in transit in use	at rest in transit in use
	<i>depending on the division of responsibility</i>	
Compliance DORA, AI Act, NIS2, GDPR		

The compliance of the solution with regulations depends on the components used

Flexibility

- SLA guarantee
- Hybrid solutions
- Flexible access model: short, medium, or long-term
- Separate access to storage array – data preparation without unnecessary GPU costs
- Additional services
- Agreement based on Polish law (negotiable)

NO
VENDOR
LOCK-IN



Why do implemented AI projects not generate ROI?

- Projects left in the Proof-of-Concept phase, without production implementation
- Garbage in, garbage out – poor quality of data leading to poor-quality model outputs
- Poor integration with organizational processes and systems
- Lack of AI project management strategy
- Excessively high implementation costs and or high operating costs
- Lack of competences – implementation experience
- Lack of innovative culture – organizational resistance

5%

**AI projects
generate millions in
savings or value gains**

95%

**AI projects
do not provide
measurable returns**

Source: MIT NANDA: Report The GenAI Divide STATE OF AI IN BUSINESS 2025

Building AI solutions that increase the likelihood of achieving positive ROI by:

- Accelerating implementation (time-to-market)
- Reducing CAPEX and OPEX
- Minimizing operational risk
- Ensuring data security & sovereignty (going concern)
- Ensuring compliance with legal regulations



We accelerate your business



NVIDIA Run:ai – management and optimization of GPU infrastructure

Key Features and Benefits:

- Democratizing access to AI by enabling the effective management of multiple projects across different departments on the same infrastructure – lower AI implementation costs
- Dynamic GPU management – maximizing efficient resource utilization, reducing costs and increasing efficiency
- GPU sharing – flexibility and scalability, increasing GPU utilization optimization performance
- Task queuing and prioritization – optimizing AI/ML team scheduling, increase in productivity
- Supporting the MLOps process – model lifecycle management = faster AI solution deployment and optimized investment costs
- Running AI in hybrid environments – including with major cloud providers (Run:ai is available from hyperscalers) – lower AI costs




NVIDIA AI Enterprise: ready-made tools to streamline AI project deployment

NVIDIA NIM – 600+ pre-trained AI models, including industry specific and process specific models (e.g. Bielek, Mistral, Llama, Gama) and templates of ready-made AI solutions.

Key features and benefits:

- Pre-trained, validated, and optimized models, re-trainable on private data and infrastructure → **Faster time-to-market**
- Performance optimization → **Improved user experience** thanks to faster and smoother processes
- Deployment scalability – supports on-prem, cloud, and hybrid environments → **Cost control and flexibility** (no vendor lock-in)



The screenshot shows the NVIDIA AI Enterprise web interface. At the top, there's a navigation bar with the NVIDIA logo, a search icon, a help icon, and a 'Login' button. Below this is a 'Featured Models' section with a 'View All' link. The main content area lists several models with their respective icons, names, descriptions, and tags.

Model Name	Description	Tags
qwen qwen3-coder-480b-a35b-instruct	Excels in agentic coding and browser use and supports 256K context, delivering top results.	agentic coding, browser use, +2
openai gpt-oss-20b	Smaller Mixture of Experts (MoE) text-only LLM for efficient AI reasoning and math	math, chat, reasoning, +1
nvidia nvidia-nemotron-nano-9b-v2	High-efficiency LLM with hybrid Transformer-Mamba design, excelling in...	mamba, agentic, nano, +4
nvidia cosmos-reason1-7b	Reasoning vision language model (VLM) for physical AI and robotics.	physical ai, autonomous vehicles, +7


build.nvidia.com

Red Hat OpenShift AI – an open-source platform for managing the lifecycle of AI models in a hybrid cloud

Key functionalities:


- Model lifecycle management
- Hundreds of models trained, validated, and optimized by Red Hat
- Scalable LLM inference using vLLM and LLM-D
- MLOps automation
- Performance and security
- Consistent database across environments
- Hybrid environment support
- Integration with accelerators and storage
- Flexible AI environment control
- Open source integration

Demo: Red Hat OpenShift AI with NVIDIA AI Enterprise



A coherent ecosystem of cooperating entities

- **Customers** – setting the direction for development, identifying real business needs.
- **NVIDIA** – ongoing technological support, access to the latest tools and high-performance GPU resources, and solution use cases.
- **Red Hat** – continuous technological support, access to the newest tools and to the partners ecosystem.
- **Consulting firms** – support in defining strategies, implementing, and scaling AI solutions.
- **AI modeling vendors and companies** – access to ready-made innovations that accelerate implementation.
- **Universities** – knowledge, research, and talent.
- **Communities** – building a culture of collaboration and experience sharing.



Benefits – Beyond.pl vs Hyperscalers

	Beyond.pl	Hyperscalers
Functionality	comprehensive MLOps platform with full support for model training, deployment and monitoring (the same tools and AI models)	
Sovereignty	- location – Poland - agreement – under Polish/EU law	- location – worldwide - agreement – foreign jurisdiction (most often U.S. law)
Vendor Lock-In	no vendor lock-in / easy exit planing	limited data transfer / difficult exit planning
Egress fees	no fees	high fees
On-premise security	Rated 4 Data Center	Tier III/Rated 3
Support	- engineers PL/EN available - SLA - 15 min response time 24/7 (negotiable)	- engineers available in premium services - SLA depends on the contract plan
Costs	- transparent and predictable, easy to budget - low entry cost - hybrid ecosystem optimization of costs (ability to connect with hyperscalers and other on prem infra via cloudconnect or Megaport) - no exit costs	- complex pricing models, difficult to accurately estimate the budget - high entry cost - high exit costs
GPU	latest generation (B200)	newest generation not always available
Go to Market	support – intro to customers, leads, partner ecosystem	limited support

Experts and competencies in Beyond.pl – including related to AI



14 AI experts



> 1400 training hours




over **160** certificates

IT Experts


Specialities	#employees
Data Center & Smart Hands Services	36
Business Technology & Innovation	16
AI Team	14
Global & Private Cloud Operations	12
Service Desk & CSC	12
IT Operations & Delivery	11
Network & MNS	9
Compliance, Risk & IT Cyber Security	5
Total	115

Specialists by experience level



- Junior (3 years of experience)
- Regular (over 8 years of experience)
- Senior (over 13 years of experience)
- Expert (over 19 years of experience)

Beyond.pl Service Chain



We constantly invest in developing the competences of our IT experts.

We allocate significant resources to ensure our services are managed by Beyond.pl experts with access to the latest IT trends. Our team of experts is also supported by qualified representatives from our network of partners.

- Selected certificates from Beyond.pl experts

- Over 2,100 hours of training per year devoted to certified training in the latest IT developments.

Partnerships and technological competences



and many more.



Beyond.pl vision

We empower the bold visioners of tomorrow with high-performance, clean-energy-powered compute and colocation. By redefining data centers into AI Factories, we're helping innovators thrive - sustainably, securely, sovereignly, and on a large scale.

Beyond.pl is CEE's first data center operator providing high-performance, scalable IT services through state-of-the-art, secure, energy efficient data centers for deploying AI, cloud computing, and enterprise workloads. We provide an ecosystem of services and solutions that support the development of AI projects.

Contact us



Wojciech Stramski

CEO Beyond.pl

mail: w.stramski@beyond.pl

phone: +48 602 381 727