



FORTISSIMO
PLUS

FORTISSIMO PLUS SUCCESS STORIES INNOVATION STUDIES

1st Digital Edition

Fortissimo Plus Success Stories
Innovation Studies
1st Digital Edition

Publisher: Arctur on behalf of the FFplus Consortium

Published: in April 2026

Place of publication: Nova Gorica, Slovenia

Edition: 1st, Digital

Copyright © 2026 Members of FFplus Consortium

Text Copyright © 2026: Andrew Forrester, FFplus sub-Projects Partners, Members of FFplus Consortium

Editor: Members of FFplus Consortium

Design: ART Rebel 9

Type: Digital Publication

Kataložni zapis o publikaciji (CIP) pripravili v Narodni in univerzitetni knjižnici v Ljubljani

COBISS.SI-ID 277018115

ISBN 978-961-94587-5-4 (PDF)



EuroHPC
JOINT UNDERTAKING

This project has received funding from the European High-Performance Computing Joint Undertaking (JU) under grant agreement No 101163317.

The JU receives support from the Digital Europe Programme.

www.commission.europa.eu

www.eurohpc-ju.europa.eu

INDEX

4	Chapter 1: About the FFplus Project
6	Chapter 2: FFplus Project Partners
8	Chapter 3: FFplus Subproject Partners
10	Chapter 4: FFplus Success Stories
50	Chapter 5: List of Abbreviations

www.ffplus-project.eu



ABOUT THE FFPLUS PROJECT

DRIVING SME AND STARTUP INNOVATION BY UNLEASHING THE POTENTIAL OF HPC AND AI

FFplus is a European initiative highlighting and promoting the adoption of High-Performance Computing (HPC) by SMEs and start-ups across Europe. By connecting SMEs with cutting-edge computing resources, expertise, and support, the project helps transform ambitious ideas into practical, scalable solutions.

At the heart of the initiative lies a clear objective: to demonstrate the real business value of HPC. In the context of the innovation studies, HPC is a key enabler for the development of AI solutions and services that require large-scale data analysis and model training in order to deliver new and improved business solutions. The resulting success stories illustrate how companies across Europe integrate HPC and AI into their operations to drive growth, efficiency, and innovation. At the same time, the studies pursue optimal use of compute resources with a positive impact on their energy footprint.





BUILDING ON A STRONG FOUNDATION

FFplus continues the legacy of the highly successful Fortissimo project series, which has played a pivotal role in bringing HPC closer to European industry. Since 2013, through Fortissimo, Fortissimo 2, and FF4EuroHPC, more than 130 experiments involving 330 partners have been carried out, resulting in 120 success stories. Building on this strong foundation, FFplus (2024–2029) continues the mission with an expanded focus on both HPC and generative AI. Through a series of six open calls, the project supports Business Experiments and Innovation Studies, engaging forwardlooking SMEs and startups from diverse industrial sectors.

FROM IDEAS TO IMPACT: FIRST INNOVATION STUDIES

A key milestone in this journey is the completion of the first tranche of FFplus Innovation Studies with AI-expert SMEs employing large-scale computing resources to develop advanced solutions based on modern AI models, including Large Language Models. The wide range of target applications of those AI technologies strengthens the EU expertise in generative AI technologies and their take-up in commercial and industrial sectors.

These Innovation Studies were selected for funding as FFplus sub-projects via the first open call for proposals. With a total indicative budget of €4 million, the 18 selected sub-projects bring together 36 organisations from 14 countries.



FFPLUS PROJECT PARTNERS

The FFplus consortium consists of 7 HPC expert organisations from the private sector with lots of experience gained while working with Industry, other HPC experts and the Academy on several R&D EU projects. The project is coordinated by the High Performance Computing Center (HLRS).

HLRS



The High-Performance Computing Center Stuttgart (HLRS) was established in 1996 as the first German national high-performance computing center, building on a tradition of supercomputing at the University of Stuttgart that stretches back to 1959. HLRS operates one of Europe's most powerful supercomputers, provides advanced training in HPC programming and simulation, and conducts research to address key problems facing the future of supercomputing. www.hlrs.de

ARCTUR



Arctur is a privately owned high-tech company specialising in advanced and innovative IT solutions, with a strong focus on R&D. By mastering cutting-edge IT technologies, we inspire & empower organisations to reinvent themselves and co-create innovative products and services, contributing to a more sustainable and equitable future. Our team is deeply involved in the R&D projects, especially in the topics of bringing HPC & AI into medicine, tourism, agriculture, manufacturing and defence sectors. www.arctur.si

CINECA



Cineca is a not-for-profit Consortium, made up of 118 members: the Italian Ministry of Education and Merit, the Italian Ministry of Universities and Research, 70 Italian universities and 46 Italian National Institutions and agencies. Today it is the largest Italian computing centre, one of the most important worldwide. It develops advanced Information Technology applications and services, acting like a trait-d'union between the academic world, the sphere of pure research and the world of industry and Public Administration. www.cineca.it



CESGA

Foundation CESGA has the mission to contribute to the advancement of Science and Technical Knowledge, by means of research and application of high performance computing and communications, as well as other information technologies resources, in collaboration with other institutions, for the profit of society. www.cesga.es



SCAPOS AG

Scapos AG was founded in December 2008 for the sales, marketing and support of technical computing software. The core scapos business is the distribution of advanced software solutions, particularly from research and academic institutes. Its product portfolio includes: optimisation software, computer-aided engineering, libraries for HPC (in particular fast, scalable solvers for Linear Algebra problems). www.scapos.com



TERATEC

Created in 2005, TERATEC is a non-profit European pole of expertise in High-Performance Computing, numerical simulation and Data Analytics, and brings together more than eighty technological and industrial companies, laboratories and research centres, universities, wishing to combine their resources in the strategic field of HPC and simulation. TERATEC is involved in major national and European initiatives to ensure the development and the mastery of HPC/HPDA technologies and to guarantee their dissemination and deployment in the ecosystem. www.teratec.eu



CYFRONET

Academic Computer Centre Cyfronet AGH is a separate legal entity of the AGH University of Science and Technology in Krakow. Cyfronet is one of the biggest Polish supercomputing and networking centres. Cyfronet is a specialised competence centre in the field of HPC and distributed computing. Cyfronet is also the coordinator of Polish National Competence Centre NCC PL. Large part of the centre's activities is devoted to the industrial applications, pharmacy, automotive and medicine sectors in particular. www.cyfronet.pl



FFPLUS SUBPROJECT PARTNERS



ARTIFICAX
TECHNOLOGIES





FFPLUS SUCCESS STORIES

- 12 Foundation Model for European Geospatial Mapping**
Geodeticca Vision, Slovakia

- 14 Large Language Models for Computer Aided Design**
SSC, Germany

- 16 Accelerating Structure-Guided Epitope Mapping via Hybrid AI Models**
Puxano, Belgium

- 18 A GenAI-based tool for 3D preclinical dosimetry from 2D dynamic imaging**
BIOEMTECH, Greece

- 20 GenAI-based co-Pilot Supporting Energy Transition by Leveraging HPC**
UBITECH, Greece

- 22 Turning Climate Science Into Instant Intelligence: The ClimateBlocks Project**
Newsroom, Portugal

- 24 A Specialized Generative AI Model for the Banking Industry**
FairMind, Italy

- 26 Enhancing Explainable AI for Investment Management with Large Language Models**
Axyon AI, Italy

- 28 The European Regulatory Copilot for Real-Time Compliance Monitoring**
Kodex AI, Germany

- 30 Generative Craniofacial Reconstruction for Forensic Identification**
Panacea Cooperative Research, Spain

- 32 HPC and AI Power Synthetic Data to Advance Cancer Imaging**
Better Medicine, Estonia

- 33 Multimodal Foundation Model for German Property Invoice Checking**
PropertyExpert, Germany

- 36 Geo-Llama: Large Language Model for Geographic Data**
DataMonkey, Germany

- 38 A Trust Framework for Sovereign AI: The Hallucination-Aware Layered Optimization**
Seedbox Ventures, Germany

- 40 On Device AI Assistant: Revolutionizing Real-Time API Interactions**
ArtificaX, Türkiye

- 42 Locally Deployable Enterprise Search and Q&A Solution**
Tilde, Lithuania

- 44 Assisted Low-Code Generation and Optimization: Bridging AI and App Design**
WebRatio, Italy

- 46 Enhancing the Decision-Making of Local Authorities Through Specialized AI**
Delibia, France



FOUNDATION MODEL FOR GEOSPATIAL ANALYSIS



Technology used: HPC, AI, Foundation Models

Industry Sector: Aerospace

ORGANIZATIONS

GEODETICCA VISION s.r.o. is a Slovak SME specialising in geospatial data processing and developer of the CARTOGRAPHER platform, which converts satellite and aerial imagery into high-quality vector data for planning, environmental protection and emergency response. **The Centre of Operations of the Slovak Academy of Sciences** (Slovakia's national HPC centre) provided expertise in parallel programming, large-scale data processing and model training.

THE CHALLENGE

The geospatial sector increasingly relies on AI to extract information from satellite and aerial imagery, yet most models require large volumes of labelled data to perform reliably across diverse regions. For SMEs, this creates high costs, long preparation times and limited scalability. GEODETICCA VISION faced barriers to European growth because adapting models for each country demanded extensive annotation, slowing deployment and limiting scale.

THE SOLUTION

GEODETICCA VISION developed the GEODETICCA VISION Foundation Model (GVFM), trained on high-resolution imagery from 17 European countries. Using EuroHPC systems, including MareNostrum-5 and Slovakia's Devana system, 90 TB of raw imagery was processed into a 9 TB curated dataset with 600 million samples. The trained model was fine-tuned and integrated into the CARTOGRAPHER platform, delivering high accuracy with minimal labelled data.



THE IMPACT

The project transformed GEODETICCA VISION's business model by enabling a single reusable foundation model that supports deployment across multiple European countries. Annotated data requirements were reduced by 87.5%, cutting typical project costs from €170,000 to €25,500 and enabling affordable service agreements with municipalities, environmental agencies and SMEs. Faster deployment strengthens competitiveness and accelerates market entry.

Socially, the solution improves access to high-quality geospatial intelligence for budget-constrained public bodies, supporting better land-use planning, infrastructure management and disaster response. By reducing technical barriers, it enables smaller organisations to benefit from advanced AI-driven mapping capabilities.

Environmentally, the approach reduces duplicated data preparation and repeated model training, lowering compute intensity and energy consumption. Improved mapping accuracy supports more informed environmental monitoring and sustainable land management decisions across Europe.

BENEFITS

- 33% higher mapping accuracy, improving reliability for planning and environmental analysis.
- Lower project costs, making advanced geospatial intelligence accessible to municipalities and SMEs.
- Faster deployment across regions, enabling quicker access to actionable spatial data.
- 87.5% reduction in labelled data requirements, lowering development overhead.
- Reusable foundation model enabling scalable expansion into new European markets in Europe.



LARGE LANGUAGE MODELS FOR COMPUTER AIDED DESIGN



Technology used: Data Engineering, RLHF, MLOps

Industry Sector: Automotive

ORGANIZATIONS

SSC is a German IT service provider specialising in secure data exchange and data management solutions for industrial collaboration, particularly in the automotive sector. **Karlsruhe Institute of Technology (KIT)** support applied research and technology knowledge transfer at the interface between science and industry. The partners collaborated to combine industrial software expertise with advanced AI research.

THE CHALLENGE

2D CAD drawings are widely used across industry and are often archived as images, making it difficult to retrieve structured information. Although these drawings contain valuable graphical elements, tables and annotations, this information is not directly accessible for automated processing or search. Conventional OCR struggles with dense layouts, multiple formats and strict data protection requirements, preventing scalable automation.

THE SOLUTION

SSC developed a local generative AI solution to automatically extract metadata and structured information from 2D CAD drawings. A configurable synthetic data generation pipeline produces realistic training data reflecting industrial CAD formats. Vision-language models were fine-tuned and benchmarked using large-scale HPC resources, requiring around 6,500 GPU node hours. The model will be integrated into SSC's SWAN platform, enabling engineers to query drawings via an interactive chatbot.



THE IMPACT

The innovation extends SSC's core data exchange offering into AI-driven automation, strengthening its competitive position in the automotive sector and opening access to new markets. Testing across 2,600 CAD drawings demonstrated a 92% reduction in time for table header classification compared with manual processes, delivering significant productivity gains and reducing error rates.

By automating repetitive review tasks, engineers can focus on higher-value design and collaboration activities, improving job quality and knowledge retention. The project also enhanced internal capabilities, with several employees gaining advanced skills in AI, data engineering and HPC workflows.

More efficient handling of engineering data reduces rework and duplication, contributing to lower resource consumption and more sustainable industrial processes overall. The local deployment approach ensures compliance with data sovereignty requirements, supporting wider adoption of AI in sensitive engineering environments.

BENEFITS

- Expansion of SSC's portfolio into AI-driven automation for engineering workflows.
- 92% reduction in time required to extract key metadata from CAD drawings.
- Up to 96% reduction in manual review effort compared to traditional processes
- Improved data quality through automated validation against PDM systems
- Development of advanced AI and HPC expertise among 3–4 SSC employees.



ACCELERATING STRUCTURE- GUIDED EPITOPE MAPPING VIA HYBRID AI MODELS



Technology used: HPC, GenAI, Diffusion Models, Cryo-EM

Industry Sector: Biotechnology

ORGANIZATIONS

PUXANO is a Belgian SME specialising in AI-driven protein design and cryo-electron microscopy services for pharmaceutical, biotechnology and agricultural clients. **CSIC**, Spain's national research council, contributed advanced computational methods for cryo-EM structure reconstruction. **IT4I**, the Czech national supercomputing centre, provided access to the Karolina supercomputer and expert HPC support.

THE CHALLENGE

Drug discovery is slow, costly and uncertain, with R&D investment often exceeding US\$2.6 billion per programme. While structure-guided design can accelerate early research, current AI protein prediction tools lack accuracy to replace labour-intensive experiments. Integrating AI with sparse cryo-EM constraints promised faster, reliable structures, but no scalable datasets or workflows existed to combine these methods during large-scale training.

THE SOLUTION

HARMONY is a hybrid model that combines AI-based protein structure prediction with sparse cryo-electron microscopy data to deliver accurate structures at 10–15× lower cost than standard cryo-EM. An AI-generated draft structure is refined using cryo-EM images to correct errors. Optimised alignment algorithms accelerated training by around 100×. Using the EuroHPC Karolina supercomputer, the team processed terabytes of data and completed 200,000 GPU hours of training to enable de novo prediction.



THE IMPACT

HARMONY strengthens PUXANO's position as a platform-based contract research organisation by enabling faster, more reliable protein structure determination at a fraction of traditional cryo-EM costs. Large-scale HPC access enabled a proprietary hybrid AI model that blends prediction with experimental evidence, creating a differentiated and scalable service that supports revenue growth and market expansion across pharmaceutical, biotechnology and agricultural sectors.

For clients, faster and more accurate structures improve early-stage drug development decisions. More antibody variants, vaccine candidates or small molecules can be screened within the same budget, reducing risk and increasing success rates. Social impact arises from the potential acceleration of new therapeutics, benefiting patient groups with limited treatment options.

Environmental benefits include reduced reliance on repeated wet-lab experiments and energy-intensive cryo-EM sessions. Shorter experimental cycles lowered material use and energy consumption, supporting more sustainable results.

BENEFITS

- 20× faster alignment reduces per-iteration time from 30 s to 1.5 s, enabling large-scale AI training.
- Hybrid AI cryo-EM workflows cut structure determination costs by 10–15× for clients. Screening capacity increases from 1–3 to 12–24 variants per session.
- New HPC and data-processing expertise strengthens PUXANO's long-term AI platform.



A GENAI-BASED TOOL FOR 3D PRECLINICAL DOSIMETRY FROM 2D DYNAMIC IMAGING



Technology used: HPC, AI, GenAI

Industry Sector: Biotechnology

ORGANIZATIONS

BIOEMTECH is a Greek SME specialising in biomedical engineering, delivering advanced hardware and software solutions for preclinical imaging and dosimetry.

ALETHIA is a start-up focused on AI monitoring, explainability and governance.

GRNET provides national-scale networking, cloud and HPC services, and coordinates the National Competence Centre for HPC in Greece.

THE CHALLENGE

Pre-clinical dosimetry typically relies on simplified assumptions or post-mortem analysis, limiting accuracy and driving extensive animal use. BIOEMTECH sought to differentiate its imaging systems by enabling real-time 3D dosimetry during acquisition, but developing and validating such a capability required large-scale Monte Carlo simulations and AI training that are only feasible using HPC, and beyond the financial and computational capacity of an SME.

THE SOLUTION

DosimetrEYE AI-powered solution is delivering real-time 3D dosimetry during preclinical imaging. Large-scale Monte Carlo simulations were used to generate ground-truth dose maps from around 1,000 SPECT/CT image pairs, which then trained Generative AI models using large-scale HPC resources. This enabled the volume, speed and precision required for reliable real-time dosimetry. Explainable AI and automated outlier detection provide quality control, robustness and trust, supporting reliable industrial and regulatory adoption.



THE IMPACT

DosimetrEYE transforms dosimetry from a time-consuming research bottleneck into a real-time capability, reducing analysis time from several hours to under 20 seconds and cutting per-study costs by around 60%. This significantly strengthens BIOEMTECH's product portfolio by enabling simultaneous imaging and dosimetry, supporting access to new customer segments and B2B collaborations, with projected additional revenues of €0.8–1.0 million.

Access to European HPC infrastructure was a critical enabler, allowing BIOEMTECH to achieve results otherwise only feasible for large industrial or academic players.

The solution enhances the competitiveness of EU medical imaging SMEs in a market dominated by US vendors and has attracted new research collaborations and follow-on EU funding. Socially, DosimetrEYE supports personalised medicine by enabling more accurate organ-level dose assessment in radionuclide research, improving safety and translational potential. Environmentally and ethically, precise in-vivo dose estimation enables a substantial reduction in animal sacrifice in preclinical studies, directly supporting European objectives for responsible, sustainable biomedical research.

BENEFITS

- Real-time GenAI-based 3D dosimetry delivered in under 20 seconds during imaging.
- ~60% reduction in per-study dosimetry analysis costs.
- Reduction of animal sacrifice by more than 80% in preclinical studies.
- ~1,500× faster dosimetry workflows compared to conventional methods.
- Creation of new revenue streams, with ~€800k expected from DosimetrEYE services.



GENAI-BASED CO-PILOT SUPPORTING ENERGY TRANSITION BY LEVERAGING HPC



Technology used: HPC, GenAI, Data Engineering, LLMs

Industry Sector: Energy

ORGANIZATIONS

UBITECH is a Greek SME specialising in explainable, high-performance Generative AI and LLM-based conversational systems for the energy sector. Its solutions operate in 15 countries, supporting 31 languages and handling over 580,000 “dialogues” per month. **The University of Innsbruck** has expertise in HPC, AI and digital twin technologies for energy systems, while **FEN Research** focuses on applied R&D on climate-neutral, electricity- and hydrogen-based energy solutions.

THE CHALLENGE

Europe’s clean energy transition depends largely on active citizen participation, yet energy literacy remains low. Consumers struggle to understand energy markets, incentives and efficiency options, while energy professionals lack effective digital tools to explain complex topics clearly. General-purpose LLMs could not capture the detailed regulatory and contextual aspects of the energy domain, and building a reliable Energy Copilot required large-scale, multimodal training beyond commercial cloud capacity.

THE SOLUTION

The innovation study developed mAIEnergy, a GenAI Energy Copilot powered by EuroHPC resources. Using LEONARDO, the consortium trained and fine-tuned energyspecific language models on large multimodal datasets. A Retrieval-Augmented Generation pipeline was created that links to curated expert knowledge, ensuring accurate, explainable and continuously updated guidance for citizens and the energy industry.



THE IMPACT

mAI Energy strengthens UBITECH's business position by enabling a transition from small-scale cloud experimentation to a commercially viable, HPC-enabled AI solution. Access to HPC infrastructure has significantly reduced training times, and enabled greater model accuracy, supporting new B2B partnerships with energy retailers, aggregators and public authorities across Europe.

Socially, the Energy Copilot improves energy literacy by translating complex policies, incentives and market mechanisms into clear, personalised guidance, empowering citizens to make informed choices. This supports behavioural change, increased participation in energy communities and more efficient energy use.

Environmentally, mAI Energy contributes to EU climate objectives by supporting energy efficiency, renewable adoption and carbon reduction, while also demonstrating how transparent, European AI innovation can strengthen digital sovereignty in the clean energy domain.

BENEFITS

- Up to 90% reduction in model training time using EuroHPC resources.
- Improved response accuracy through domain-specific fine-tuning.
- Efficient large-scale GPU utilisation enabling faster experimentation.
- Transformation of a prototype into a scalable SaaS offering.
- Reusable and transparent open-source knowledge repository for energy AI ecosystem.



TURNING CLIMATE SCIENCE INTO INSTANT INTELLIGENCE: THE CLIMATEBLOCKS PROJECT



Technology used: HPC, AI, GenAI

Industry Sector: Environment

ORGANIZATIONS


The Newsroom is a Portuguese media innovation startup specialising in AI tools that help journalists navigate complex scientific information with speed, accuracy and transparency. The company designs and deploys ClimateBlocks, combining expertise in AI, knowledge engineering, climate science and responsible data journalism to transform large-scale research into auditable, evidence-linked intelligence for public communication.

THE CHALLENGE

Climate science is expanding faster than journalists can interpret it. Evidence is dispersed across 500,000 papers, making transparent, reliable reporting difficult. To enter the growing climate intelligence market and move beyond media-monitoring tools, The Newsroom needed a traceable, scientifically grounded solution. Building such a system required computational power far beyond commercial infrastructure.

THE SOLUTION

The Newsroom developed ClimateBlocks, an AI-powered platform that connects and interprets more than 500,000 climate research papers through a transparent, sourcelinked knowledge graph. Using GPU-accelerated HPC resources, the team fine-tuned a specialised small language model on thousands of annotated papers, enabling rapid extraction of insights, expert identification and traceable concept pathways within seconds for newsroom workflows. The system is designed for practical, auditable newsroom use.



THE IMPACT

ClimateBlocks is already used by more than 150 journalists to verify findings and reduce research time from hours to seconds, improving the speed, quality and credibility of evidence-based reporting. Through its partnership with the European Journalism Centre, the platform is scaling across newsrooms and strengthening resilience against climate misinformation – a growing societal and democratic risk.

For The Newsroom, the project establishes a defensible position in the emerging climate intelligence market, transforming the company from a media-tech startup into a provider of structured, auditable scientific infrastructure. The creation of a 500,000-paper knowledge graph with millions of validated links provides a proprietary asset base that supports new revenue streams across journalism, research services and policy analysis. By delivering transparent, source-linked AI built on European HPC, the company enhances trust, digital sovereignty and long-term commercial scalability.

BENEFITS

- Trained a domain-specific AI model 96% smaller than general systems, cutting inference costs by 12–17x.
- Built a proprietary knowledge graph of 500,000 climate papers with millions of structured links.
- Enabled 150+ journalists to verify claims in seconds, improving reporting credibility.
- Secured European Journalism Centre partnership, accelerating market adoption.
- Created a scalable climate intelligence platform supporting new revenue streams.



A SPECIALIZED GENERATIVE AI MODEL FOR THE BANKING INDUSTRY



Technology used: HPC, GenAI

Industry Sector: Fintech

ORGANIZATIONS

FairMind Srl develops an Agentic Platform that support software-development teams by improving quality, productivity and governance. **Credit Agricole Group Solutions** provides domain expertise, validating datasets and evaluating banking-specific use cases.

THE CHALLENGE

Banks want to use generative AI to speed up software modernization and new features development, but they must protect sensitive customer data and meet GDPR, security and audit requirements. Public, general-purpose models cannot be used: exposing confidential information, and lack the domain knowledge required in regulated software development environments. FairMind needed a sovereign model that fits internal processes and can run on-premise. Training & bias testing exceeded in-house compute.

THE SOLUTION

FairMind built BankGPT, a banking-tuned open-weight model (4–32B parameters) designed for secure internal deployment. It curates banking datasets with version control for full traceability, then tunes models on HPC with throughput- and memory-efficient training. Quality is validated using bias benchmarks and a custom banking Q&A test set. BankGPT supports requirements drafting, code generation and testing within bank processes. Outputs remain inside the bank's proprietary operating environment.



THE IMPACT

BankGPT demonstrates that domain-specific, open-source generative AI can meet strict banking security, GDPR and audit requirements while delivering measurable productivity gains. By enabling secure on-premise deployment, banks retain full control of sensitive data, prompts and model outputs, reducing regulatory exposure, operational risk and dependence on external cloud vendors.

FairMind's structured evaluation framework—combining bias testing with a banking-specific Q&A benchmark validated by Credit Agricole—provides a repeatable qualification pathway for regulated institutions, accelerating trusted adoption of sovereign AI. Commercially, FairMind forecasts 3–5 pilots in year one (€180k–600k), scaling to 10–15 clients in year two (€600k–1.8m) and 20–30 clients in year three (€1.2m–3.6m), supporting expansion into banking, insurance and healthcare. BankGPT allows developers and compliance teams to focus on higher-value design, review and governance activities, improving job quality while strengthening responsible AI practice in regulated sectors.

BENEFITS

- Optimised models for secure, efficient on-premise deployment in banks.
- Established reliable checkpointing for repeatable large-scale experiments.
- Implementation of full dataset traceability via versioned HuggingFace collections.
- Building a banking Q&A benchmark validated by Credit Agricole experts.
- Training of open-weight BankGPT variants up to 32B parameters on HPC resources.



ENHANCING EXPLAINABLE AI FOR INVESTMENT MANAGEMENT WITH LARGE LANGUAGE MODELS



Technology used: HPC, GenAI, Data Engineering
Industry Sector: Fintech

ORGANIZATIONS

Axyon AI is an Italian fintech SME delivering AI-driven predictive analytics for quantitative investment strategies. The company supports asset managers with forecasting, portfolio construction and risk analytics based on advanced machine learning. **UNIMORE–AlmageLab** at the University of Modena and Reggio Emilia contributed deep learning and multimodal AI expertise, supporting model training, evaluation and benchmarking within a high-performance computing environment.

THE CHALLENGE

Axyon AI produces rich predictive outputs, but turning technical model insights into clear, consistent narratives was a bottleneck. Analysts spent significant time preparing explanations, while compliance teams required auditable, reviewable documentation. Improving explainability was therefore both scientific and commercial: greater transparency was essential to unlock regulated markets, accelerate adoption and strengthen trust.

THE SOLUTION

Using EuroHPC resources, the partners developed a financial-domain large language model that converts quantitative forecasts, SHAP attributions and market context into clear, human-readable explanations. Fifteen models were benchmarked using an HPC-enabled evaluation pipeline and LLM-as-a-Judge scoring. Selected open-weight models were fine-tuned with parameter-efficient methods such as LoRA and mixed-precision training. Outputs were structured for review by analysts and compliance teams.



THE IMPACT

Enhanced explainability materially strengthens Axyon AI's competitive position within the highly regulated investment management sector. By converting complex quantitative outputs into structured, auditable narratives, the solution reduces the time analysts and compliance officers spend preparing model documentation, lowering operational cost per mandate. Clearer explanations accelerate institutional due diligence, shorten sales cycles and increase confidence during risk committee and model validation reviews. This enables entry into more strictly regulated market segments that require defensible AI governance and transparent model behaviour. Socially, improved interpretability increases transparency for end-investors and supports more accountable decision-making across the investment value chain. Standardised explanation workflows also foster stronger collaboration between data scientists, portfolio managers and compliance teams, improving internal governance and reducing operational friction as the company scales.

BENEFITS

- Up to 55% improvement in explanation quality compared to base models.
- Addressable market has expanded from ~10% to ~25% of investment firms.
- Efficiency gains of ~1 FTE-month achieved per client per year.
- 48,000 EuroHPC GPU-hours enabled scalable benchmarking and training.
- Strengthened collaboration with UNIMORE-AlmageLab and NCC Italy.



THE EUROPEAN REGULATORY COPILOT FOR REAL-TIME COMPLIANCE MONITORING



Technology used: HPC, GenAI, Data Engineering

Industry Sector: Fintech

ORGANIZATIONS



Kodex AI is a German SME that develops AI solutions for the European financial sector. The company builds domain-specific language models to automate regulatory compliance, reporting, and risk management for financial institutions. By combining high-performance computing with generative AI, Kodex AI delivers precise, scalable tools that simplify complex workflows across multilingual and rapidly evolving European regulatory frameworks.

THE CHALLENGE

Financial institutions struggle to keep pace with rapidly evolving European regulations, relying on fragmented tools and generic AI that lack domain precision and audit transparency. Rising compliance costs and risk exposed a clear gap. To deliver a specialised, trustworthy regulatory AI, Kodex AI needed large-scale model training and experimentation beyond commercial cloud limits, requiring access to HPC resources.

THE SOLUTION

Kodex AI developed the European Regulatory Compliance Copilot powered by a finetuned Regulatory Large Language Model (RegLLM). Using EuroHPC infrastructure, including MareNostrum 5, the team trained the model on 35,000+ multilingual regulatory documents, guidance notes, and historic compliance reports. A continuous ingestion pipeline keeps the knowledge base up to date, while human workflows ensure legal accuracy, explainability, and transparent, auditable outputs.



THE IMPACT

Kodex AI have transitioned from limited cloud testing to full-scale RegLLM development on European HPC. This enabled training and validation of a domainspecific regulatory model that outperforms open-source and proprietary alternatives, including GPT-4, on finance compliance benchmarks. Financial institutions report measurable impact: 60–70% reductions in compliance reporting time and €1–1.5 million annual savings for mid-sized organisations. Automated horizon scanning and obligation extraction reduce regulatory blind spots, speed response to new directives, and lower exposure to fines and reputational risk.

Commercially, the project strengthened Kodex AI's market position. The addressable market expanded by 150%, enabling entry into new EU jurisdictions and larger enterprise segments. HPC-enabled development improved robustness and shortened innovation cycles. Revenue is projected to double within three years. Socially, the solution improves transparency and accountability through explainable, audit-ready outputs that reinforce trust between institutions, regulators, and citizens.

BENEFITS

- €1–1.5 million annual cost savings for mid-sized institutions.
- Model iteration cycles reduced from weeks to days using EuroHPC GPUs.
- 9%+ accuracy improvement over GPT-4 on regulatory benchmarks.
- 60–70% reduction in compliance reporting time.
- Improved auditability and trust through explainable, low-hallucination outputs.



GENERATIVE CRANIOFACIAL RECONSTRUCTION FOR FORENSIC IDENTIFICATION



Technology used: GenAI, HPC, Computer Vision

Industry Sector: Forensics

ORGANIZATIONS

Panacea Cooperative Research is a Spanish SME specialising in AI-driven forensic identification tools. **The University of Granada (UGR)** contributes expertise in reliable AI methods and biomedical imaging. **INRIA**, France's National Institute for Research in Digital Science and Technology, provides world-leading knowledge in image-based generative modelling to support robust scientific and industrial innovation.

THE CHALLENGE

Across the world, disasters, conflicts and crime leave many human remains unidentified, delaying legal processes and uncertainty for families. Craniofacial reconstruction (CFR), used when no records exist, relies on scarce specialist artists. The process is slow and subjective. Panacea identified demand for a faster, science-based alternative, but training generative models requires large 3D datasets and computing power beyond an SME's capacity.

THE SOLUTION

Through the FACEGEN project, Panacea applied generative AI techniques, including diffusion models, to automate craniofacial reconstruction. Access to high-performance computing enabled large-scale training, testing and validation. The most effective method uses multiple skull depth-map views to generate a full 3D facial reconstruction in seconds. In parallel, the consortium developed evaluation and reliability methods so forensic experts can assess and trust the results.



THE IMPACT

FACEGEN removes a critical bottleneck in forensic identification by delivering fast, affordable and consistent craniofacial reconstructions. For Panacea, integrating automated reconstruction into its Skeleton-ID platform enables a shift from bespoke consulting to a scalable, pay-per-use service model, dramatically expanding its addressable market. The project strengthened long-term collaboration with UGR and INRIA and accelerated Panacea's innovation capacity. Technically, the validated diffusion pipeline has been adopted in a major €11M national R&D programme with Fujitsu and the Spanish National Police.

The social impact has been faster and more objective reconstructions that support judicial processes, improve identification rates and help families reach closure, while making advanced forensic capabilities accessible to regions lacking specialist expertise.

BENEFITS

- Facial reconstruction time reduced from weeks to seconds.
- Reachable market expanded from ~5% to nearly 80% of 10,000 global institutions.
- Scalable online, pay-per-use service enabled.
- Methods adopted in an €11M national project with the Spanish National Police.



HPC AND AI POWER SYNTHETIC DATA TO ADVANCE CANCER IMAGING



Technology used: HPC, GenAI, Diffusion Models

Industry Sector: Healthcare

ORGANIZATIONS

Better Medicine is an Estonian SME developing AI solutions for cancer diagnostics and medical imaging. The company worked with **Pärnu Hospital**, which provided clinical expertise, imaging data and feedback from routine radiology practice. **The University of Tartu** contributed research expertise in artificial intelligence and high-performance computing, enabling the creation of high-quality synthetic CT data to strengthen tumour-detection models.

THE CHALLENGE

Medical AI models for CT imaging require large, diverse and well-annotated datasets, which most hospitals lack. Limited tumour cases, inconsistent labelling and differences in scanners and imaging protocols reduce model reliability across clinical sites. For Better Medicine, models trained on public datasets performed poorly on local hospital scans due to a clear domain gap and too few annotated tumour cases.

THE SOLUTION

With access to EuroHPC systems, Better Medicine developed a 2.5D generative diffusion model that produces realistic CT volumes by using information from neighbouring slices. Healthy scans from public and internal datasets formed the base, with synthetic tumours added to increase data diversity. Combining synthetic and real data to retrain the segmentation model significantly improved performance and reduced false detections.



THE IMPACT

Better Medicine have transformed from delivering exploratory research to production of a scalable commercial solution. Access to high-performance computing made it possible to train and validate complex generative models efficiently, reducing development timelines by up to 40% and lowering overall costs. Clinically, the project demonstrated measurable improvements on real hospital imaging data, increasing trust in AI-assisted diagnostics. Hospitals with limited datasets can now deploy more accurate, locally adapted models, supporting more equitable cancer detection across patient groups.

Commercially, the ability to generate GDPR-compliant synthetic datasets creates new revenue opportunities while reducing the need for cross-border data sharing. Environmentally, more efficient HPC workflows shortened experimentation cycles, reducing wasted computation and energy use while supporting sustainable AI development.

BENEFITS

- Over 20% improvement in tumour-segmentation accuracy on hospital data.
- Threefold reduction in false tumour detections.
- Up to 40% shorter AI development cycles.
- Scalable HPC-ready workflows for rapid model updates.
- New revenue opportunities through synthetic medical data services.



MULTIMODAL FOUNDATION MODEL FOR GERMAN PROPERTY INVOICE CHECKING



Technology used: HPC, AI, Data Engineering

Industry Sector: Insurance

ORGANIZATIONS


PropertyExpert GmbH (PX) is a German technology provider for the property insurance and real estate sectors. It supports insurers, retailers and portfolio managers with AI-enabled digital claims processing. PX combines domain expertise, structured document workflows and machine learning to deliver fair, consistent and efficient property damage assessments.

THE CHALLENGE

PX were limited to automated text-only checks, but claim decisions often rely on photographic evidence of damage. Generic image-text models frequently omit industry-specific repair details, so PX turned to a domain-trained multimodal model to scale automation, improve accuracy, and keep costs predictable. Training at image scale exceeded its in-house compute capabilities, so deployment of HPC was essential.

THE SOLUTION

PX developed a multimodal AI pipeline linking invoice text, repair descriptions and damage images within a unified foundation model. Image encoders and fusion layers connect visual evidence with structured document data to validate quotations and invoices. Using the MeluXina HPC system, PX executed iterative data preparation, pretraining and fine-tuning across distributed nodes, enabling rapid experimentation and measurable accuracy improvements.



THE IMPACT

Access to HPC infrastructure enabled PX to design, train and validate a domain-specific multimodal foundation model at production scale. The resulting 4% increase in automated invoice decisions translates into estimated annual administrative savings of around €100,000 while strengthening PX's position in the German property insurance market. Faster and more reliable verification improves insurer response times during peak weather events and reduces operational bottlenecks caused by limited lossadjuster capacity.

For policyholders and contractors, improved decision accuracy supports fairer, more transparent outcomes and reduces disputes linked to invoice interpretation. Accelerated approvals help households restore damaged properties more quickly after storms or floods.

Greater digital validation has reduced unnecessary site visits and repeat inspections, lowering travel-related emissions. Optimised distributed training and mixed-precision workflows have also improved computational efficiency, supporting more sustainable large-scale AI deployment.

BENEFITS

- Fewer disputes and follow-up inspections, faster claim decisions.
- 4% increase in automated validation.
- €100k annual administrative savings.
- +1.2% decision accuracy improvement.
- 2.8× training speed gain.



GEO-LLAMA – LARGE LANGUAGE MODEL FOR GEOGRAPHIC DATA



Technology used: HPC, AI, GenAI, Geospatial Technologies

Industry Sector: Built Environment

ORGANIZATIONS

DataMonkey, a brand under Urban Monkeys GmbH, is a German technology SME specialising in AI-driven data access solutions. The company focuses on lowering barriers to advanced analytics for business experts, data analysts and IT teams by combining large language models with domain-specific data pipelines. For this project, DataMonkey led the design, training and validation of a geospatially-aware LLM, working with EuroHPC infrastructure to deliver a scalable, European alternative for geospatial intelligence.

THE CHALLENGE

Geospatial data is critical e.g. for mobility planning, climate-risk analysis and real-estate management, yet remains difficult to exploit. Although most business data has a spatial component, many SMEs and public organisations lack GIS expertise, suitable tools and computing capacity. Open data such as OpenStreetMap is valuable but complex, requiring specialised queries and handling highly inconsistent tagging.

THE SOLUTION

DataMonkey developed Geo-Llama, a geospatially aware large language model that enables users to query and combine open geographic data using natural language. The solution integrates Generative AI, Retrieval-Augmented Generation pipelines and HPCaccelerated fine-tuning. Access to EuroHPC GPU resources enabled efficient training of multi-billion-parameter models, delivering a scalable, high-precision geospatial analytics platform for non-experts.



THE IMPACT

Geo-Llama fundamentally improves access to geospatial intelligence by reducing the need for specialised GIS skills. Through natural-language interaction, organisations can analyse and combine open geographic data while reducing data-preparation effort by up to 95%, accelerating decision-making e.g. in mobility, utilities, real estate and sustainability.

Access to large-scale HPC resources was critical to train and optimise multi-billionparameter models, ensuring the accuracy, robustness and scalability required for business-grade deployment. For DataMonkey, this enabled faster product development, reduced technical risk and a stronger competitive position in the European AI market. More broadly, Geo-Llama delivers an EU-based, compliant geospatial AI solution that supports digital sovereignty, encourages reuse of open data and enables more sustainable planning, infrastructure management and environmental decision-making across Europe.

BENEFITS

- +35% higher query accuracy compared to baseline models.
- Training dataset of over 135,000 high-quality input–output pairs covering 91% of OpenStreetMap key usage.
- Up to 95% reduction in data-preparation time for geospatial analyses.
- Enables non-technical users to perform complex geographical queries via natural language.
- Strengthens DataMonkey's position as a leading EU-based geospatial AI provider.



A TRUST FRAMEWORK FOR SOVEREIGN AI: THE HALLUCINATION-AWARE LAYERED OPTIMIZATION



Technology used: HPC, GenAI, Reinforcement Learning

Industry Sector: Regulated Industries

ORGANIZATIONS


Seedbox Ventures GmbH is a German AI studio specialising in advanced Generative AI solutions. The company has strategically evolved from delivering bespoke, projectbased services into a scalable, product-led technology provider, focused on building trustworthy, sovereign AI solutions for regulated European industries requiring compliance, transparency and auditability.

THE CHALLENGE

European adoption of Generative AI remains limited in regulated sectors due to high and unpredictable operational costs, persistent risks of model hallucinations, and significant data sovereignty and GDPR compliance concerns. Seedbox's existing service-based delivery model was difficult to scale under these constraints, prompting the need for a single, production-ready framework capable of delivering trusted, costefficient, and compliant AI.

THE SOLUTION

Seedbox developed the HALO framework, an integrated Generative AI architecture built entirely on open-source models and trained using distributed HPC resources. HALO combines a specialised RAG-LLM, an intelligent classifier router and a factual consistency auditor, delivering transparent, verifiable and cost-efficient AI outputs while ensuring GDPR compliance, EU AI Act alignment and reduced energy consumption in production environments.



THE IMPACT

Seedbox's transformation from a bespoke AI consultancy into a scalable, product-led SaaS company was enabled by the FFplus study, with HALO becoming its core reusable intellectual property asset supporting recurring revenue growth across regulated European markets.

By reducing operational LLM costs by up to 85% through intelligent query routing, HALO makes production-grade Generative AI economically viable for SMEs and mid-sized enterprises previously constrained by unpredictable cloud pricing, lowering adoption barriers in regulated sectors.

For finance, legal and healthcare organisations, HALO provides auditable, sentence-level trust scoring aligned with GDPR and the EU AI Act, removing key procurement barriers linked to unverifiable AI outputs.

HALO strengthens institutional trust in AI by embedding transparency and explainability at architectural level, reducing misinformation, unsafe automation and reputational risk in high-stakes environments.

BENEFITS

- Up to 85% reduction in operational LLM costs, achieved through intelligent query routing that avoids unnecessary use of computationally intensive Generative AI processes.
- Auditable, sentence-level hallucination transparency, enabling compliant use of Generative AI in regulated sectors such as finance, legal and healthcare.
- Removal of procurement barriers by delivering transparent, GDPR-compliant and EU AI Act-aligned AI outputs.
- Enabled transition to a scalable, product-led SaaS business model.



ON DEVICE AI ASSISTANT: REVOLUTIONIZING REAL-TIME API INTERACTIONS



Technology used: GenAI, AI Agents, HPC

Industry Sector: Software and Information Technology

ORGANIZATIONS

ArtificaX is a Turkish AI company developing next-generation intelligent assistants able to plan, act and complete tasks autonomously. It combines research in agent behaviour with scalable engineering to build enterprise-grade automation systems. HPC supports large-scale training and synthetic data generation. **Turkcell Global Bilgi** is Türkiye's leading digital customer-experience provider, delivering contact-centre and digital-support services across telecoms, finance, retail and public sectors.

THE CHALLENGE

Large language models excel at conversation but rarely execute actions. The next stage of AI requires assistants that can safely perform real tasks such as bookings, payments and record updates. Training such systems demands vast amounts of realistic, multistep action data. As a startup, ArtificaX lacked the computing scale to generate, validate and distil these models for reliable on-device deployment.

THE SOLUTION

Using MareNostrum5 HPC system, ArtificaX generated large volumes of synthetic, action-rich dialogues to train models capable of safe use. Actions were validated in controlled environments, then distilled into efficient 8B and 4B parameter models optimised for phones, tablets and vehicles. Processing runs locally by default to protect privacy, with cloud access used only when strictly necessary. This approach ensures reliable task execution, consistent confirmations and strong data protection.



Hey Ali! What do you have in mind?

Tired. I need a vacation next week.
Maybe in Antalya?

I can book you on flight TK2312 to Antalya on Friday at 7:00 PM using your miles, and arrange your hotel and a rental car. I'll also reserve museum admission for a morning visit. The sea may be chilly early, but afternoons after 3 PM are good for a swim. Shall I proceed with the purchases?



THE IMPACT

The project demonstrated that compact on-device models can reliably complete complex, multi-step tasks, providing ArtificaX with a validated foundation for commercial agentic AI solutions. This reduced technical risk, strengthened the product roadmap and accelerated time to market.

Enterprises can automate routine operations in customer support, finance, sales and operations while lowering cloud costs and simplifying regulatory compliance by keeping sensitive data on device, improving competitiveness and customer trust. For end users, everyday activities such as managing bills, appointments or purchases become faster and more secure, even in environments with limited connectivity. Local processing supports privacy by default, requires explicit consent before actions are taken and improves digital inclusion through voice-based access.

Environmentally, shifting intelligence to the edge reduces reliance on centralised data centres, lowering energy consumption and carbon emissions at scale while supporting more sustainable AI deployment.

BENEFITS

- Fast, real-time responses on leading devices
- Up to 57% reduction in task completion time
- Efficient 8B and 4B models with high task accuracy
- Validated datasets covering complex, real-world services
- Reduced cloud costs and scalable edge deployment model



LOCALLY DEPLOYABLE ENTERPRISE SEARCH AND Q&A SOLUTION



Technology used: HPC, AI, LLM, GenAI

Industry Sector: Telecommunication

ORGANIZATIONS


Tilde is a European language technology SME specialising in multilingual AI solutions for enterprises and public institutions. The company develops neural machine translation, speech technologies, chatbots and terminology management systems. Tilde combines human linguistic expertise with AI-driven R&D and actively promotes digital inclusion for under-resourced European languages through initiatives such as its open large language model, TildeOpen.

THE CHALLENGE

Organisations manage growing volumes of multilingual information, yet existing enterprise search and Q&A tools fail to meet strict data sovereignty, compliance and language-quality requirements. Cloud-based AI limits adoption in regulated sectors and performs poorly for smaller European languages. To unlock new regulated markets and scale its enterprise offering, Tilde needed a secure, locally deployable solution delivering high-accuracy results in under-resourced languages.

THE SOLUTION

Tilde developed a modular, locally deployable enterprise assistant combining semantic search with generative question answering. The system retrieves content from internal document repositories and generates context-aware answers with traceable sources, operating entirely on-premises or in private clouds to ensure regulatory compliance. FFplus HPC access enabled large-scale fine-tuning of the TildeOpen language model on up to 768 GPUs, improving accuracy for Balto-Slavic languages.



THE IMPACT

The project significantly strengthens Tilde's commercial position in secure, multilingual AI solutions. By reducing error rates in Balto-Slavic languages from over 27% to levels comparable with widely supported languages, Tilde can now compete for high-value contracts in regulated sectors including government, finance, healthcare and legal services. The locally deployable architecture removes a major barrier to adoption in data-sensitive environments, enabling entry into markets previously inaccessible to cloud-only AI providers.

The solution transforms Tilde's advanced language technologies into a scalable enterprise platform with predictable deployment costs and performance. This enhances recurring revenue potential, supports long-term customer retention and positions Tilde as a strategic AI partner in Central and Eastern Europe, where the addressable market is expected to grow by up to 150% within three years. Socially, the solution improves access to trustworthy AI tools in diverse European languages, supporting digital inclusion for public institutions, SMEs and end users.

BENEFITS

- Enables fully on-premises enterprise search and Q&A with strong data sovereignty.
- Delivers high-accuracy AI for under-resourced European languages.
- Reduces information retrieval and decision-making time by up to 60%.
- Converts Tilde's language technologies into a scalable, reusable enterprise platform
- Unlocks access to regulated markets and high-value contracts across Central and Eastern Europe.
- Provides a reusable platform for summarisation, translation and support services.



ASSISTED LOWCODE GENERATION & OPTIMIZATION: BRIDGING AI & APP DESIGN



Technology used: HPC, LLM, GenAI
Industry Sector: Telecommunication

ORGANIZATIONS


WebRatio is an Italian SME with over 20 years of experience in enterprise software and creator of the Interaction Flow Modelling Language Object Management Group® standards. Its low-code platform integrates AI-assisted development for web, mobile and process-driven applications. **RES IT** is a specialist in AI, machine learning and deep learning, delivering proprietary data-driven solutions to enterprise and financial customers.

THE CHALLENGE

Enterprises face mounting pressure to accelerate digital transformation despite shortages of skilled developers and growing application backlogs. Traditional software development is slow and costly, while existing low-code platforms still require specialist expertise. Generative AI offers potential, but without domain understanding it lacks the reliability and control required for enterprise use, creating urgency for European SMEs to innovate and remain competitive.

THE SOLUTION

WebRatio developed ALGO, an AI-powered copilot that is embedded in its low-code platform. Using HPC resources, over 100,000 archive samples were used to fine-tune open-weight large language models in the Interaction Flow Modelling Language domain. The selected model was integrated into a graph-based framework with syntactic and semantic guardrails, enabling reliable natural-language generation of fully executable application models. These were benchmarked and packaged for secure Integrated Development Environment integration.



THE IMPACT

ALGO demonstrates that generative AI can be applied safely and predictably, within enterprise low-code environments. For business, it delivers a dual impact: substantial productivity gains for professional developers through automated, compliant modelling, and a new addressable market of citizen developers, accelerating adoption across organisations.

Socially, it lowers the barrier to software creation, enabling non-technical users to build functional applications, improving digital inclusion and strengthening collaboration between business and IT teams.

Strategically, the project strengthens European competitiveness by providing a trustworthy, EU-developed alternative in a market dominated by non-EU vendors, supporting data sovereignty, governance and long-term SME-led innovation. It also provides the foundation for WebRatio's long-term AI adoption roadmap, reinforcing the company's role as an innovation-driven SME.

BENEFITS

- 50% productivity gain for professional developers through AI-assisted modelling.
- ~€40,000 annual value created per developer for customers.
- 100% expansion of the addressable market by enabling citizen developers.
- Projected €2.1M in new annual recurring revenue by 2030 from AI-enhanced subscription tiers and licenses.



ENHANCING THE DECISIONMAKING OF LOCAL AUTHORITIES THROUGH SPECIALIZED AI



Technology used: GenAI, Data Engineering, HPC

Industry Sector: Local Government

ORGANIZATIONS

Delibia is a French SME developing generative AI tools that help local authorities analyse and compare their decision-making processes. **The Occitanie Region** supplies public data, pilots real-world testing of the platform, and trains civil servants to use the solution. **GENCI** provides High-Performance Computing resources via the Jean Zay supercomputer, enabling large-scale model training and optimisation.

THE CHALLENGE

French local authorities publish millions of governance documents across thousands of websites, making it difficult to compare how similar policies are handled in other regions. Officers must understand prior decisions, constraints and outcomes before acting, but fragmented access slows decisions and reduces consistency. Generic AI models lack long-context handling, legal precision and traceability, and training models exceeds SME resources.

THE SOLUTION

Delibia developed a long-context generative AI model specialised for public-sector documents. Using GENCI's High-Performance Computing infrastructure, the team fine-tuned domain-specific embeddings on more than three million administrative acts, consuming over 50,000 GPU hours. The optimised model is four times smaller and delivers 2.3× faster responses, enabling accurate retrieval, full on-premise deployment, and compliance with strict public-sector data sovereignty and security requirements.



THE IMPACT

The specialised embedding model improves the accuracy and speed of retrieving legal references, regulatory clauses and comparable administrative decisions. Tasks that previously required minutes of manual searching can now be completed in seconds, reducing workload for civil servants and legal teams. For Delibia, these gains strengthen its SaaS platform and enable premium features for public-sector workflows. Based on interest from regions and major cities, the solution is expected to drive a 15– 20% increase in recurring revenue over two years, reinforcing its competitive position in sovereign AI.

Socially, the platform improves transparency by making millions of public decisions easier to analyse and compare. It supports faster, evidence-based policymaking and more consistent governance across territories for public benefit.

BENEFITS

- 2.3× faster inference with a 4× smaller model footprint.
- High-accuracy retrieval across millions of administrative acts.
- GDPR- and ISO 27001-aligned sovereign deployment.
- Lower infrastructure and energy costs.
- Reusable methodology for other public-sector AI models.




European Commission
www.ec.europa.eu



EUROHPC Joint Undertaking
eurohpc-ju.europa.eu

This project has received funding from the European High-Performance Computing Joint Undertaking (JU) under grant agreement No 101163317. The JU receives support from the Digital Europe Programme.

For further information, please visit:



www.ffplus-project.eu



www.linkedin.com/company/43366020



www.youtube.com/@ffplusproject



open.spotify.com

LIST OF ABBREVIATIONS

AI	Artificial Intelligence
B2B	Business-to-Business
CAD	Computer-Aided Design
CFD	Computational Fluid Dynamics Simulations
CPU	Central Processing Unit
cryo-EM	Cryo-Electron Microscopy
CT	ComputedTomography
DL	Deep Learning
EU	Europe Union
FEM	Finite Element Method Simulations
FTE	Full-Time Equivalent
GIS	Geographic(al) information system
GDPR	General Data Protection Regulation
GenAI	Generative Artificial Intelligence
GPU	Graphics Processing Unit
HPC	High-Performance Computing
ISV	Independent Software Vendor
LLM	Large Language Model
M	Million
MC	Monte Carlo Simulations
ML	Machine Learning
MLOps	Machine Learning Operations
NCC	National Competence Centre
OCR	Optical Character Recognition
OEM	Original Equipment Manufacturer
PaaS	Platforms as a Service
RLHF	Reinforcement Learning from Human Feedback

SaaS	Software as a Service
SME	Small and Medium Sized Enterprise
SPECT	Single-Photon Emission Computed Tomography
TI	Information technology
TB	terabyte

www.ffplus-project.eu

