



# OpenGryd

Sovereign, Efficient, and Decentralized AI Infrastructure for Europe's Digital Future

AI INFRASTRUCTURE

SUSTAINABILITY

DIGITAL SOVEREIGNTY

# Europe's Data Center Energy Crisis Is Accelerating

The digitalization of Europe is driving an unprecedented surge in energy demand. EU data centers are projected to consume **150% more energy by 2035** compared to 2025 – a trajectory that puts immense pressure on both energy grids and sustainability commitments.

Meanwhile, the competitive landscape is deeply skewed: US hyperscalers – Google Cloud, AWS, and Microsoft Azure – command **80% of all cloud computing capacity**, while European cloud providers remain marginalized at a mere **15% global market share**, with little sign of meaningful growth.

## Key Figures

150%

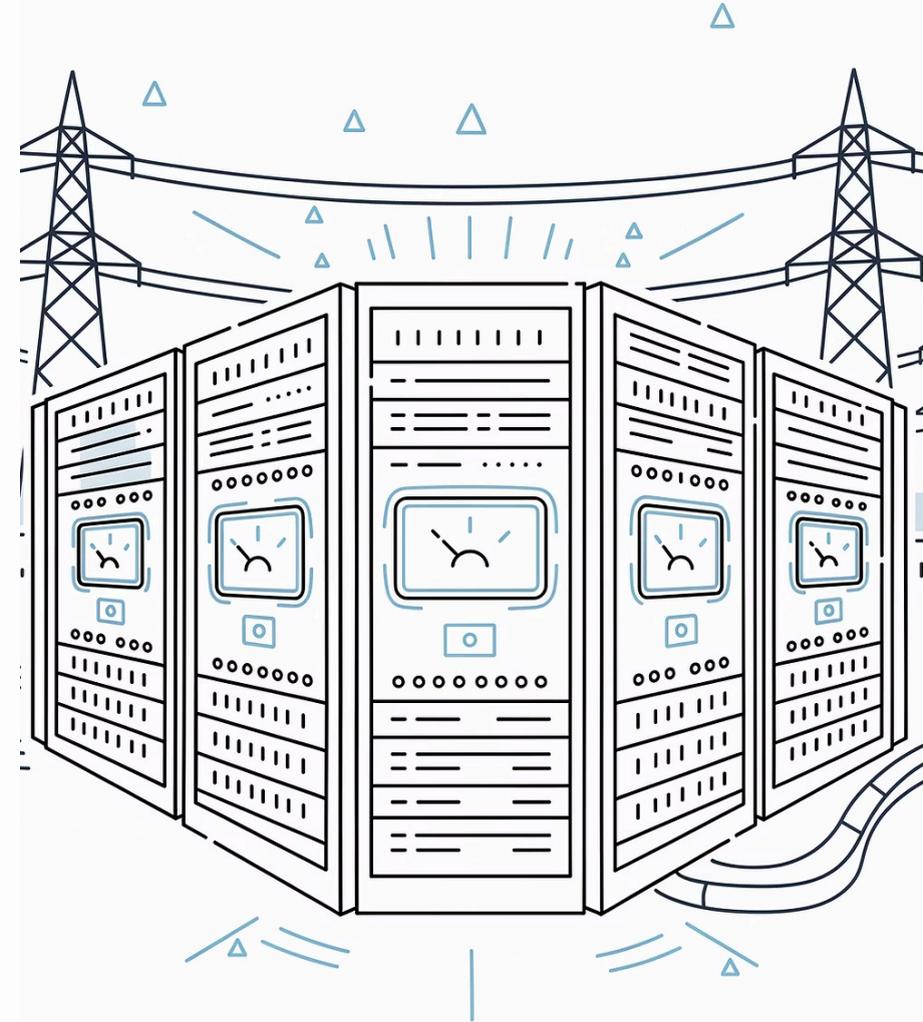
Projected rise in EU data center energy use by 2035

80%

Cloud market share held by US hyperscalers

5%

of the world's AI compute capacity



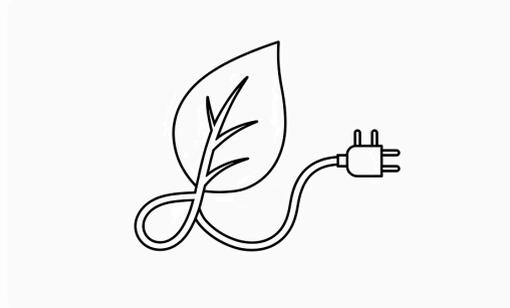
# The EU AI Act Sets a New Compliance Baseline

The EU AI Act, the world's first comprehensive legal framework for artificial intelligence, establishes binding obligations across four critical dimensions.



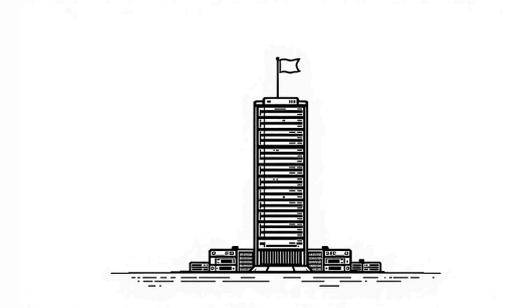
## Safety First

High-risk AI systems must meet stringent conformity and risk management requirements before market entry.



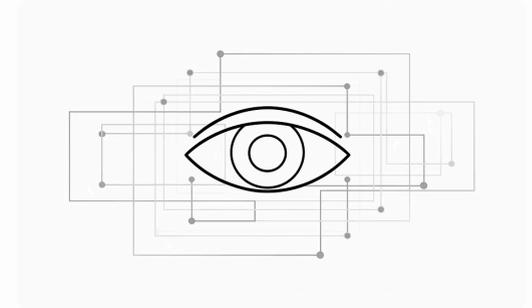
## Green AI

AI systems must disclose energy consumption and environmental impact, making sustainable infrastructure a legal obligation.



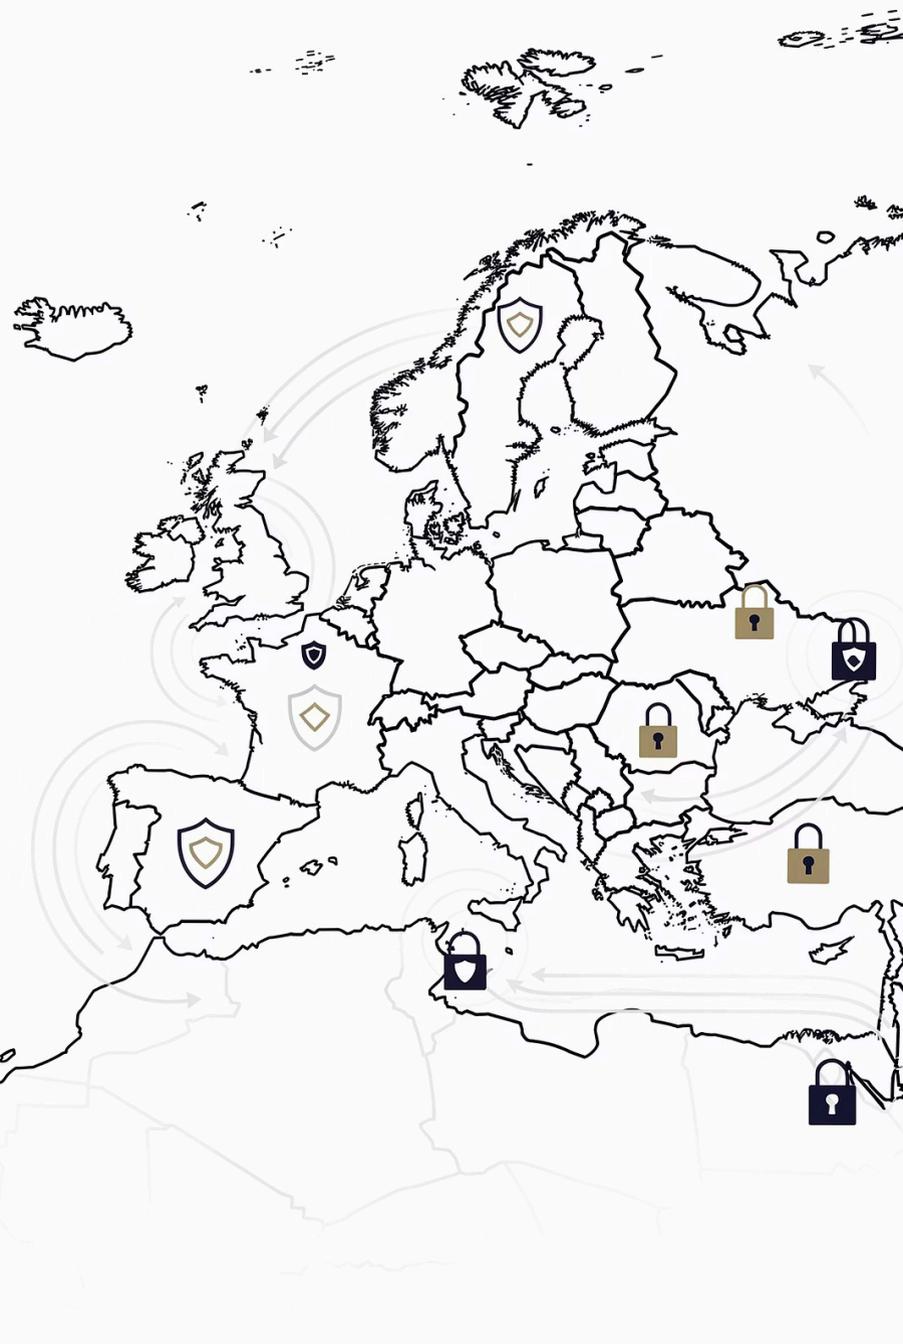
## Sovereignty & Compute

The Act highlights Europe's compute deficit, where dependence on non-EU infrastructure for critical AI poses strategic risks.



## Transparency

Organizations must ensure explainability, auditability, and traceability of AI decisions, requiring infrastructure that supports full model visibility.



# Sovereignty: Europe's Most Urgent AI Vulnerability

1

## Copyright, IP & GDPR Exposure

Processing European citizen and enterprise data on US-based cloud infrastructure creates persistent compliance risks under GDPR and evolving IP frameworks.

2

## GPU & VRAM Bottleneck

Europe's compute deficit is structural. Access to high-performance GPUs and VRAM is constrained by supply chains and export controls dominated by non-EU actors.

3

## Technology & Talent Drain

Leading AI researchers and engineers continue to migrate toward US hyperscaler ecosystems, eroding Europe's innovation capacity and deepening dependency.

# Open Source AI Is Closing the Gap with Proprietary Models

The software layer of AI is undergoing a fundamental democratization. Open-weight models – once far behind proprietary alternatives – are now reaching near-parity performance on a wide range of enterprise tasks, with the critical advantage of full transparency and local deployability.

This shift unlocks a new generation of AI architecture: **Agentic AI** – autonomous, multi-step reasoning systems – combined with advanced **quantization techniques** that compress model size without sacrificing meaningful accuracy, enabling powerful inference on commodity hardware.



## Open Weights

Capable open-weight models (e.g., LLaMA, Mistral, Gemma) are now enterprise-ready and fully auditable.



## Open Source Software

A rich ecosystem of inference runtimes, orchestration tools, and vector databases accelerates local AI deployment.

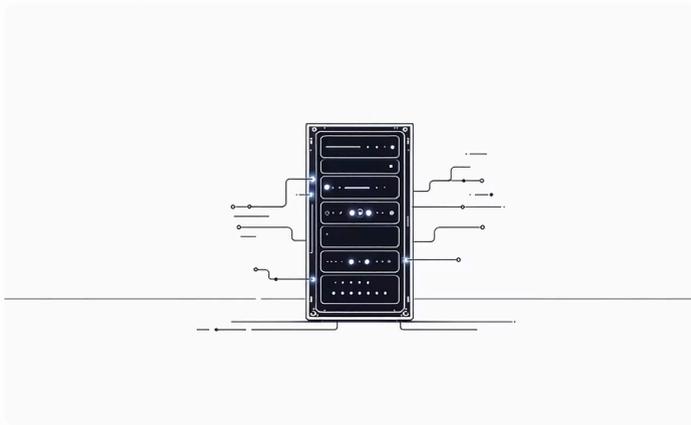


## Agentic AI & Quantization

Agentic architectures combined with dynamic quantization deliver high-performance reasoning on efficient, local hardware.

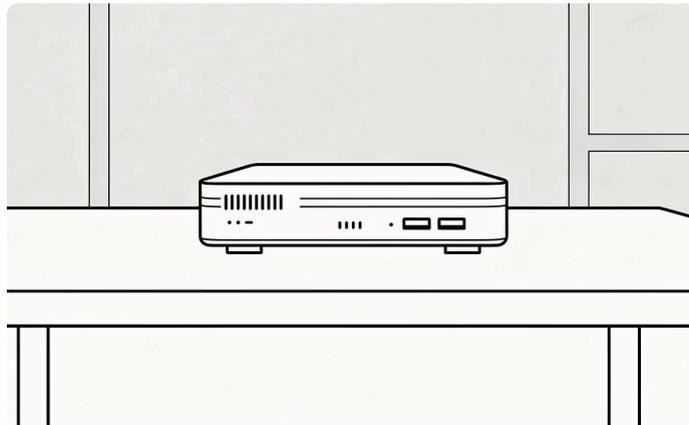
# Infrastructure Is Shifting: Edge AI Changes the Equation

A structural market shift is underway, as enterprises move towards local, high-performance Edge AI hardware driven by cost, compliance, and capability.



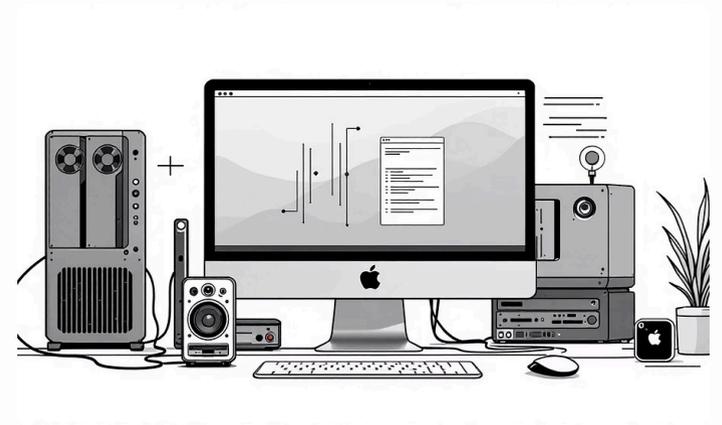
## On-Prem & Edge AI

Local hardware deployment eliminates cloud latency, reduces data exposure, and keeps sensitive workloads within organizational control.



## Minimal Infrastructure Footprint

Modern AI servers boast small physical and energy footprints, enabling viable deployment for SMEs, municipalities, and research institutions without a dedicated data center.



## High-Performance for the Masses

Platforms like Apple Silicon and Nvidia DGX deliver data center-class AI performance at a fraction of the cost, democratizing access to powerful compute.

# Our Mission: The Four Pillars of Responsible AI Infrastructure

OpenGryd believes high-performance AI and responsible deployment are complementary, shaping every product and partnership decision.



## Economical

Accessible AI pricing for all, enabling enterprise-grade solutions without hyperscaler lock-in.



## Efficient

Our stack maximizes performance per watt, optimizing for high AI throughput with minimal energy consumption.



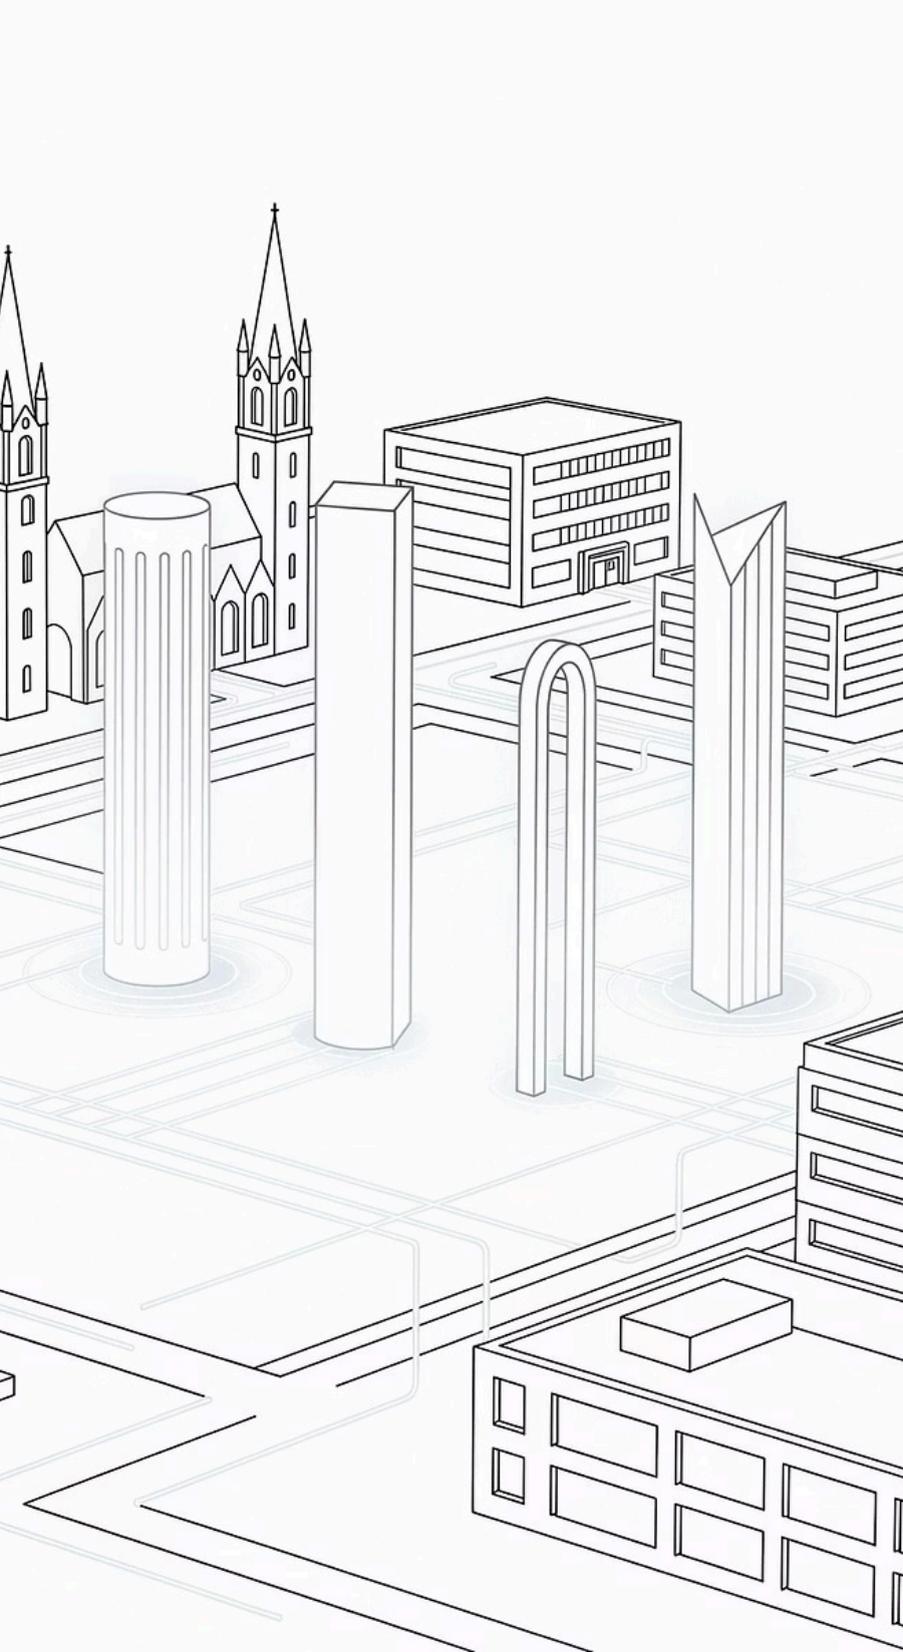
## Ecological

With recyclable hardware and energy-efficient deployment, OpenGryd provides a truly sustainable AI infrastructure.



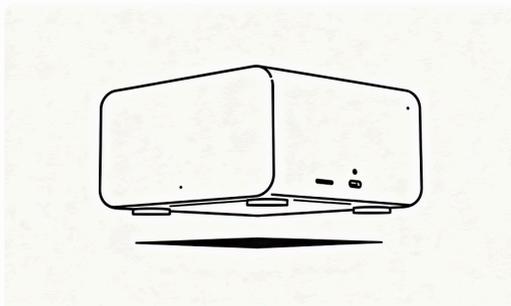
## Sovereign

Ensuring full data sovereignty and GDPR compliance, OpenGryd infrastructure keeps your data on-premises and independent.



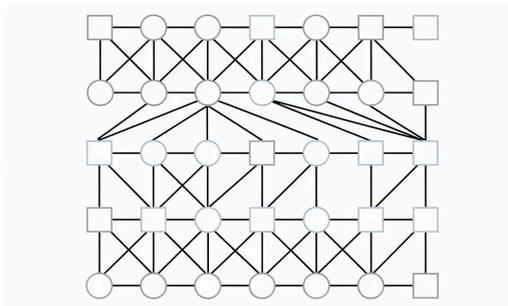
# A Full-Stack AI Platform — From Silicon to Interface

OpenGryd delivers an integrated, end-to-end AI infrastructure stack designed for rapid, sovereign AI deployment without specialist cloud expertise.



## Hardware

Leveraging **Apple Mac Studio** with Unified Memory Architecture (UMA/DMA), our hardware enables large model loading and exceptional performance-per-watt for high-performance edge computing.



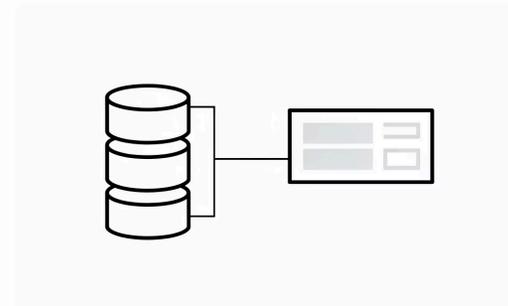
## AI Layer

Our AI layer supports GPU frameworks like CUDA and Metal, LLM runtimes (GGUF, MLX), **Dynamic Quantization**, and advanced tuning techniques (LoRA/QLoRA, MoE) for flexible, scalable inference.



## Software

Featuring Chat interfaces (OpenWebUI), **Vector Databases**, embedding models, and OCR pipelines, our software integrates Model Context Protocol (MCP) for on-premises deployment with full data sovereignty.



## Infrastructure

This layer provides containerized deployment (Docker), distributed scaling (EXO), observability dashboards, and web proxy layers for enterprise-grade operational management and monitoring.

# Background & Experience

Ioannis Dikeoulis

**Founder & CEO, OpenGryd**

M.Sc. Computer Science,  
Universität des Saarlandes

## Background:

- Founder & CEO at seatbase Technologies GmbH
- CTO at Plastle GmbH
- Head of IT at Godiplast GmbH
- IT Consulting (Freelancer)

## Research Background

→ At the **Max Planck Institute (MPI)**, Ioannis conducted foundational work in large-scale knowledge graphs and information retrieval under Dr. Weikum's Distributed Information Systems group.

→ At **CISPA Helmholtz Center**, his applied security research with Dr. Krombholz in the field of Usable Security.

→ **DFKI** research in Natural Language Processing (Dr. Neumann / Dr. van Genabith) and Time-Aware Knowledge Graph Embeddings



# Live Pilot Projects: Real-World Validation Across Sectors

OpenGryd is already operational across five pilot deployments spanning regulated industries, public sector, and deep-tech science. Each pilot validates a distinct use case – from document intelligence to sovereign knowledge management.



## Laboratory — Norm Sheet Intelligence

Automated changelog generation for regulatory norm sheets reduces manual review time.



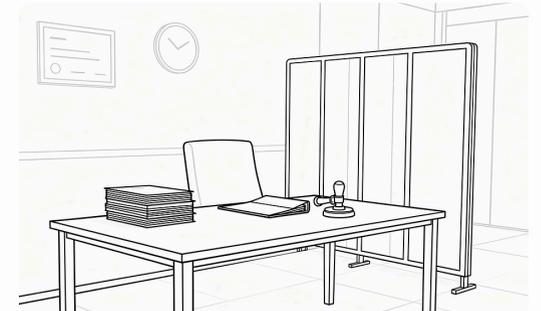
## DeepTech Biotech — Private AI Chat

A fully air-gapped AI assistant for sensitive R&D conversations, ensuring zero data exposure.



## Saarland Municipalities — RAG

Retrieval-Augmented Generation enables local government staff to query large internal document repositories.



## Legal Notary — Document Creation & PII

AI-assisted legal document drafting with automatic PII detection and redaction, compliant with GDPR.

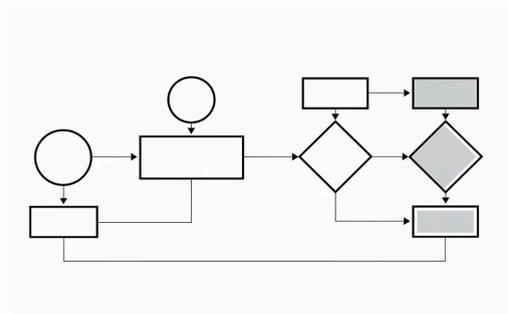
# Four Core Use Cases — From Chat to Full Process Automation

OpenGryd's platform addresses high-value enterprise AI use cases requiring privacy, accuracy, and seamless integration, all without data leaving the client's infrastructure.



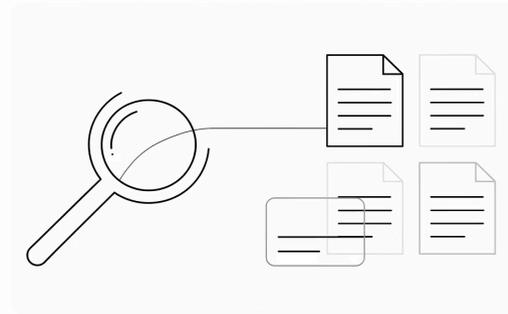
## OCR & Document Processing

Intelligent extraction from scanned documents, PDFs, and images using Vision Language Models provides structured output for automation or compliance workflows.



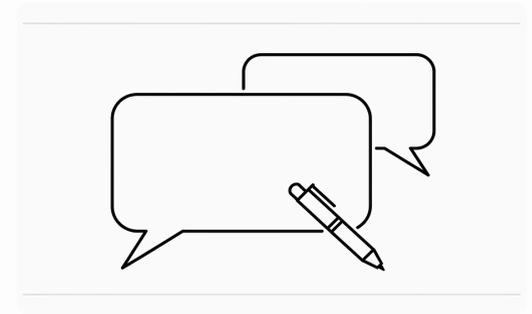
## Process & Workflow Automation

Agentic AI orchestrates multi-step business processes, from data extraction to decision routing, reducing manual effort and accelerating throughput.



## RAG Search & Internal Knowledge Base

Retrieval-Augmented Generation (RAG) enables staff to query internal documentation with natural language, delivering precise, cited answers in seconds.



## Chat, Assistants & Copy-Writing

Sovereign AI chat interfaces and writing assistants are tailored to organizational tone and policies, ensuring no data reaches foreign clouds.

## CONCLUSION

# Demo + Q&A

Thank you for your time and attention; we are excited to demonstrate OpenGryd and answer your questions.



### Pilot Partners

Join our growing network of pilot deployments across regulated industries.



### Strategic Investors

Help scale sovereign AI infrastructure across Europe.



### Policy Collaborators

Shape the regulatory framework for responsible AI deployment.

