

AI for Financial Services

PrivateMind delivers private AI intelligence for financial services. Deploy NVIDIA GPU clusters at any scale, exactly where you need them: Iceland for cost-efficient training, London and New York for exchange-connected inference. From deep learning to traditional quant research, your data, models and intelligence stay yours.



GPU Hosting

Fully managed from the ground up. Deploy anywhere from single GPUs to large clusters, with compute, datacentre, networking, storage, & the platform layer above all covered.



Model Deployment

Managed platform for hosting & serving AI models at production scale. Access open-source & custom LLMs via chat interface & OpenAI-compatible API gateway, with dedicated infrastructure for training, versioning, & full data sovereignty.



Inference Run Time

Production-grade inference endpoints optimized for financial services workloads. Low-latency inference with automatic scaling & high availability. Deploy your own trained models with full control over your AI stack.

Technology Stack



AI Applications

Custom AI solutions & integrations. Developed services for trading, risk, & back-office automation.



AI Platform

OpenAI-compatible gateway & chat interface. Open-source LLMs, custom model hosting, & built-in compliance.




AI Control Plane

Orchestrate GPU provisioning, MLflow training runs, & model deployment.



GPUs & Infrastructure

High-performance NVIDIA compute, storage & networking optimized for AI workloads.



Security

Enterprise-grade protection, compliance, patching, & OS management.



Observability

Monitoring, logging, & alerting across the full stack.



Connectivity

Reliable, high-bandwidth connectivity - real-time exchange access & cloud on-ramps.



Managed Services

24/7 platform operations, SRE support, & continuous infrastructure management.

See PrivateMind in Action – Request Your Demo Today

Discover how PrivateMind streamlines secure AI adoption with compliant, easy-to-use tools tailored for financial workflows. Our team is ready to walk you through a personalized demo and show how Managed AI can empower your business with confidence. Connect with us to learn more and start your journey.

Powering the Next Generation of AI for the Financial Services

Quantitative research, document intelligence, algorithmic trading, and risk modeling. PrivateMind handles compute-intensive workloads with GPU acceleration, hosts open-source LLMs for analysis and automation, and serves models at production scale. Your data never leaves your infrastructure.

Document Reasoning

A secure LLM to assist analysts in reviewing and deduction based on structured/unstructured documents (filings, prospectuses, performance reports).



“ Summarize all risk factors in this IPO prospectus...”

“ Translate this VBA risk model into Python...”



Code Automation & Refactoring

Powerful tooling to assist quants & dev with code refactoring, bug fixing & documentation.

GPU Accelerated Compute

GPU-accelerated Monte Carlo simulations and risk modeling. Quantitative research reduced from hours to minutes.




“ Show me all tickers with bullish 5-day momentum (over +2 z-score)...”

“ Extract key client details from this onboarding form and enter them into the CRM...”



Agentic Automation for Front, Middle & Back-Office

Use LLMs to handle multi-step workflows, with context aware decisions and approvals process.



○ Live Financial Cloud Deployments
 ■ Dedicated AI Data Center