

EXPLAINABLE, AUDITABLE, SOVEREIGN AI FOR MISSION-CRITICAL DECISIONS.

EXECUTIVE SUMMARY

Enterprise AI adoption is accelerating, but a widening **trust gap** is preventing organizations from moving AI into production at scale. Models hallucinate. Outputs lack explainability. Governance requirements go unmet. The result: billions in lost value and a growing distance between AI-forward organizations and those still stuck in experimentation.

Seekr's Trustworthy AI Platform on AMD infrastructure closes that gap with an enterprise AI stack purpose-built for trust, explainability, and governance. Seekr delivers the application layer: an end-to-end AI operating system spanning agentic data preparation, model orchestration, AI agent deployment, and post-deployment explainability and contestability – the ability to challenge and correct model outputs. AMD delivers the infrastructure foundation with GPU-accelerated compute.

Together, they provide a path from enterprise data to production-ready AI that is accurate, auditable, and sovereign, deployed wherever governance and data residency require.

THE ENTERPRISE AI TRUST GAP

Enterprise investment in AI has never been higher. But spending alone is not producing results. Organizations are struggling to trust AI outputs and security protocols sufficiently enough to put them into production.

The friction points are structural:

Black-box outputs. Most AI platforms provide answers without explaining how decisions were made. In regulated industries and high-stakes environments, unexplainable AI is undeployable AI. Leaders are accountable for outcomes, yet many lack a consistent way to measure model performance and risk across the AI lifecycle.

Governance gaps. Regulatory scrutiny is intensifying. The EU AI Act, SEC guidance, HIPAA, and emerging AI-specific regulations demand explainability, bias monitoring, audit trails, and demonstrable compliance. Generic AI platforms bolt on governance as an afterthought; mission-critical environments require it by design.

Data complexity. Enterprise data is fragmented across formats, systems, and domains. Manual workflows and limited expertise slow time-to-value and leave critical insights untapped. Models and agents are only as good as the context provided to them, and preparing that context remains time-consuming and expensive.

Vendor lock-in. Proprietary AI stacks create dependency across models, tooling, and infrastructure. Tomorrow's best model should not require rebuilding today's applications. Rigid platforms constrain where and how AI can run, forcing organizations to choose between performance and control.

Deployment constraints. Data sovereignty requirements, security mandates, and existing infrastructure investments demand deployment optionality. Organizations operating classified, regulated, or sensitive environments need AI that runs on-premises, in hybrid configurations, or fully air-gapped, without sacrificing capability.

The organizations that close the trust gap first will industrialize AI at scale. The rest will remain stuck in experimentation.



THE SOLUTION: SEEKR + AMD

Seekr and AMD technology deliver a production-ready AI stack that eliminates the trust gap. Seekr provides the application layer: a complete platform for building, deploying, governing, and explaining AI. AMD provides the infrastructure foundation: high-performance GPUs and an open, Kubernetes-native software reference stack that simplifies deployment and scaling. Together, they enable enterprises and governments to deploy AI they can trust, on infrastructure they control.

Seekr: The Trustworthy AI Application Layer

Seekr is an AI product company purpose-built to deliver software to leading enterprises, critical infrastructure operators, and government, environments where uptime, security, and verifiable information are essential. Seekr holds 33 patents granted and pending in AI explainability, contestability, and search, and is SOC 2 Type II compliant.

Seekr's platform is anchored by **SeekrFlow**, the end-to-end AI operating system, with three specialized products, **SeekrGuard**, **SeekrIntel**, and **SeekrGeo**, each powered by SeekrFlow's core technology.

SeekrFlow: The End-to-End AI Operating System

SeekrFlow enables enterprises to build, customize, and scale accurate and explainable AI in one complete, secure operating system. It covers the entire AI lifecycle from data preparation and ingestion to model fine-tuning, to agent creation and orchestration, to post-deployment explainability.

Where SeekrFlow differentiates from other platforms is in what happens before and after deployment:

AI-Ready Data Engine. Seekr's proprietary data engine ingests and structures unstructured data (PDFs, documents, images) into high-quality datasets. It automatically organizes files, refines content, and restructures it into formats that can be used to train and optimize AI models. It also connects with enterprise data systems. This reduces dataset preparation time by 2.5x compared to traditional methods, accelerating time-to-value with AI.

Agent Orchestration. SeekrFlow provides a complete platform for building, validating, deploying, and monitoring AI agents. It supports inference, fine-tuning, and deployment with a robust model library of over 40 open-source models. SeekrFlow is compatible with all hardware, meeting organizations on their infrastructure of choice.

Explainability and Contestability. After deployment, Seekr provides complete visibility into how an agent or model arrived at its output. Data attribution surfaces the specific fine-tuning data pairs that most likely influenced a given response. Context attribution identifies what sources and tools an agent used in generating a response and compares their influence. Contestability enables organizations to identify problems in the data and fix them, closing the loop on trust.

Value: Build, deploy, and govern AI across the full lifecycle with 3x more accurate model responses and 6x greater business relevance than generic platforms.

SeekrGuard: AI Model Evaluation and Governance

SeekrGuard is an evaluation and governance capability that tests, monitors, and documents model performance and risk using an organization's own context as the benchmark, enabling audit-ready decisions across model selection, tuning, and ongoing operations.

Model Scorecards. Rank model options across tasks, datasets, and scenarios for faster, evidence-based model selection.

User-Defined Risk Profiles. Weigh what matters most by mission or business line and standardize governance criteria across the organization.

Configurable Evaluators. Test edge cases, safety constraints, and adversarial prompts with consistent scoring.

Automatic Test Data Generation. Turn internal documents and knowledge sources into test criteria and datasets without heavy upfront investment.

Flexible Testing. Evaluate open-weight and proprietary models in approved environments.

A Gartner survey of 360 IT leaders (May-June 2025) found that only 23% are very confident in their organization's ability to manage security and governance when rolling out GenAI tools. SeekrGuard directly addresses this gap by giving organizations a systematic, repeatable framework for model certification and risk assessment.

Value: Go beyond generic benchmarks to use-case-ready evaluation and governance, tied to your organization's specific data, risk tolerances, and regulatory requirements.

SeekrIntel: Investigation and Intelligence Platform

SeekrIntel supports compliance, due diligence, fraud, and risk investigations across both public and private sectors. Instead of forcing analysts to manually connect disparate signals, SeekrIntel automatically surfaces relationships, patterns, and networks, transforming scattered data into explainable intelligence in minutes.

Evidence-Based Network Discovery. Instantly uncover and visualize relationships across entities, locations, transactions, and events using interactive graphs and maps.

Data Ingestion and Enrichment. Upload documents, online data, and case files. SeekrIntel cleans, translates, extracts, and delivers insights in minutes.

Natural Language Investigation. Ask plain-language questions to explore connections, overlaps, timelines, and activity patterns.

Temporal and Geographic Analysis. Trace activity across time and location to understand behavior, movement, and emerging patterns.

Transparent, Auditable Intelligence. Tie insights back to original source material and relationship paths to validate findings and maintain trust, governance, and accountability.

Value: Accelerate investigations and surface hidden risk across complex data environments, with every insight traceable to its source.

SeekrGeo: Geospatial Reasoning and Intelligence

SeekrGeo delivers explainable geospatial intelligence by reasoning across multimodal earth observation and SAR data. Built on Seekr's proprietary Remote Sensing Foundation Model, SeekrGeo synthesizes intelligence across sensors and time, reasoning like an expert analyst rather than producing isolated detections.

Multimodal Fusion. Understand the physics and signals behind optical imagery, SAR, hyperspectral data, and more.

Contextual Intelligence. Receive analyst-level insights into why an object or change matters and what it indicates.

Temporal Intelligence. Detect meaningful change over time, build baselines, and identify patterns in activity, behavior, and operational tempo.

Natural Language Interaction. Query complex geospatial data using plain language for accessible insights.

Tool Integration. Work alongside existing object detection, segmentation, GIS, and data source tools.

Value: Transform massive geospatial data into decision-ready intelligence for infrastructure monitoring, asset risk, supply chain visibility, and competitive intelligence.

AMD: The Enterprise AI Infrastructure Foundation

AMD provides the high-performance compute foundation optimized for enterprise AI workloads, with AMD Instinct GPUs delivering the memory bandwidth and capacity needed for complex model inference and training.

AMD Instinct™ GPUs

AMD Instinct GPUs provide high-bandwidth memory (HBM3) for large model inference, long-context reasoning, and compute-intensive AI workloads. Flexible form factors support on-premises, hybrid, and cloud deployment.

WHY SEEKR + AMD

Seekr delivers the trustworthy AI platform spanning data preparation, model orchestration, agents, explainability, and governance. AMD delivers the open infrastructure that makes AI workloads run faster and deploy more easily across the enterprise. Together, they close the trust gap and enable organizations to move from experimentation to production with confidence.

Trust by design, not bolt-on. Explainability, contestability, and audit-ready governance are embedded across the entire AI lifecycle, from data preparation through production monitoring. Unlike platforms that add safety features as afterthoughts, Seekr builds trust into every layer.

Better economics at scale. High-performance Instinct GPUs enable lower total cost as workloads scale. Open-source infrastructure eliminates the vendor premium that closed-stack alternatives impose.

Sovereign and flexible deployment. Deploy on-premises, in the cloud, in hybrid configurations, or fully air-gapped. Seekr's platform is infrastructure- and hardware-agnostic, and AMD's open architecture ensures no lock-in at the compute layer. Data, models, and operations remain under enterprise governance.

Model independence. SeekrFlow supports 40+ open-source models and is compatible with proprietary models. Adopt new models as they emerge without rebuilding applications or retraining teams. AMD's standard APIs ensure portability at the inference layer.

Mission-critical credibility. Seekr is trusted by leading government agencies and enterprise customers including the U.S. Army. With 33 patents in AI explainability, contestability, and search, Seekr brings deep proprietary IP to environments where failure is not an option.

BUSINESS OUTCOMES: PRODUCTION RESULTS ACROSS INDUSTRIES*

Seekr on AMD technology delivers measurable business impact across the industries where trust, explainability, and governance matter most. The platform's flexibility enables organizations to start with high-value use cases and expand across the enterprise.



Financial Services: Compliance and Risk Intelligence

CHALLENGE

Financial institutions face mounting compliance costs and regulatory complexity across multiple jurisdictions. Manual review of regulatory documents is slow, expensive, and error-prone. AI model selection is often driven by vendor claims rather than rigorous evaluation against the institution's own data and risk requirements. Data sovereignty mandates prevent use of cloud-based AI for sensitive financial data.

SOLUTION

SeekrIntel agents automate compliance, due diligence, and fraud investigations by surfacing

relationships, patterns, and networks across complex financial data, with every insight traceable to its source. SeekrGuard evaluates and certifies AI models against the institution's specific risk profiles and regulatory requirements before deployment. SeekrFlow provides the governed platform layer with full audit trails and explainable outputs. AMD technology-accelerated inference supports low-latency processing at enterprise volumes with on-premises deployment for data sovereignty.

OUTCOMES

- 90% reduction in 401(k) audit cycle time from 50 hours to 2 hours
- 65% reduction in compliance research time
- 78% reduction in manual review time



GOVERNMENT AND DEFENSE: SOVEREIGN AI FOR MISSION-CRITICAL OPERATIONS

CHALLENGE

Government agencies and defense organizations require AI that operates with zero external dependencies, complete auditability, and the ability to deploy in fully air-gapped environments. Multi-source intelligence analysis demands the ability to process and reason across structured and unstructured data in multiple languages and modalities. Model provenance and ideological bias are critical concerns for national security applications.

SOLUTION

SeekrFlow deploys as a fully sovereign AI operating system with air-gapped capability and zero

external network connectivity required. SeekrGuard evaluates models for factual accuracy, hallucination risk, and geopolitical/ideological neutrality using mission-specific risk profiles. SeekrIntel transforms multi-source intelligence into explainable networks and relationship maps. SeekrGeo provides geospatial reasoning across satellite imagery (optical and SAR) for infrastructure monitoring, change detection, and operational intelligence. AMD Instinct GPUs provide the on-premises compute foundation.

OUTCOMES

Reduced time, cost, and resources to detect and mitigate cyber vulnerabilities across mission-critical weapon systems – including Patriot and THAAD missile batteries. (U.S. Army DEVCOM AvMC, January 2026)



MANUFACTURING AND INDUSTRIAL: SUPPLY CHAIN AND INFRASTRUCTURE INTELLIGENCE

CHALLENGE

Industrial supply chains are deeply interconnected and increasingly fragile. Risk signals are scattered across supplier data, logistics feeds, contracts, and external events. Physical infrastructure spanning pipelines, transmission lines, facilities, and logistics corridors requires continuous monitoring that exceeds human analyst capacity. Traditional tools are static and reactive, surfacing issues only after disruption occurs.

SOLUTION

SeekrIntel ingests supplier records, bills of materials, shipment data, contracts, and disruption signals

to map dynamic supply chain networks. AI agents reason across entities, locations, and dependencies to identify critical exposure and cascading risk. SeekrGeo analyzes multi-temporal satellite, aerial, and SAR imagery to detect meaningful physical change over time across distributed infrastructure. Every insight is paired with a transparent explanation and traceable evidence. AMD infrastructure provides the compute foundation for processing large-scale imagery and complex network analysis.

OUTCOMES

Recovered millions in lost revenue by unifying data across hundreds of disparate operational systems and automatically reconciling contracts against services performed – delivered in a single command center view. (Fortune 2000 Secure Transport Company)



CUSTOMER SUCCESS: REVENUE RECOVERY FOR FORTUNE 2000 SECURE TRANSPORT COMPANY

CHALLENGE

A Fortune 2000 secure transport company with hundreds of disparate information systems needed to ensure accurate revenue recognition for services performed across its customer base. Revenue leakage was causing millions of dollars in losses

SOLUTION

SeekrFlow's AI-Ready Data Engine unified company-wide data from disparate sources including contract

data, invoice data, and field maintenance data. Agents trained on the unified data identified revenue leakage patterns by reconciling invoices against contracts and services.

OUTCOMES

Seamless unification of company-wide data from disparate sources. Identified revenue leakage patterns and implemented reconciliation rules driving customer ROI. Delivered a Revenue Leakage Command Center providing a single surface to query contracts, services, and invoices to prevent future leakage.

DEPLOYMENT MODEL

Seekr + AMD technology supports flexible deployment architectures to match enterprise security, governance, and operational requirements. SeekrFlow is infrastructure- and hardware-agnostic, deployed wherever data and compute reside.

Cloud: Deploy on major cloud providers. Ideal for organizations prioritizing speed and scalability.

On-Premises: Full deployment within enterprise data centers with complete data sovereignty. Zero external dependencies ensure sensitive data never leaves the organization's infrastructure. Required for many regulated industries.

Hybrid: Combine cloud elasticity with on-premises control. Route sensitive workloads to on-premises infrastructure while leveraging cloud for burst capacity. Central management with distributed execution.

Air-Gapped: Fully isolated deployment for classified and highly regulated environments. No external network connectivity required. Complete operational independence for defense and intelligence applications.

ENGAGEMENT MODEL

Seekr's engagement model is designed for rapid experimentation, deployment, and scale:

Tailored Demo. See Seekr's products in action with a demonstration focused on your specific business challenges and use cases.


Discovery and Solution Outline. Technical discovery with Seekr's Solutions Architects to identify the right solution architecture for your business requirements.

Initial Production Deployment (IPD). A 45-60 day engagement deploying a production-ready solution that meets specific KPIs and drives measurable ROI.

Scale. Expand across additional business units, geographies, and use cases. Seekr supports independent development on SeekrFlow, collaborative builds with Seekr's Forward Deployed Engineers (FDEs), or fully custom integration delivered by Seekr.

NEXT STEPS

Ready to deploy trustworthy AI in your enterprise? Engage with Seekr and AMD to:





CONTACT YOUR SEEKR OR AMD REPRESENTATIVE FOR A STRATEGY SESSION, TAILORED DEMO, OR IPD SCOPING DISCUSSION.

ABOUT AMD

More than 50 years driving innovation in high-performance computing. AMD Instinct GPUs power AI workloads for leading cloud providers, enterprises, and sovereign AI initiatives worldwide. AMD's commitment to open standards and open ecosystem ensures enterprises can deploy AI with control, flexibility, and long-term vendor independence.

ABOUT SEEKR

Seekr is a leader in explainable and trustworthy artificial intelligence designed to power mission-critical decisions in enterprises, government, and regulated industries. Seekr provides secure, auditable AI solutions tailored to sectors where transparency, accuracy, and compliance are paramount. Seekr's products, SeekrFlow, SeekrGuard, SeekrIntel, and SeekrGeo, provide comprehensive AI development, deployment, evaluation, intelligence, and governance capabilities. Backed by AMD Ventures with 33 patents granted and pending.

*All performance and/or cost savings claims are provided by the 3rd party organization featured herein and have not been independently verified by AMD. Performance and cost benefits are impacted by a variety of variables. Results herein are specific to such 3rd party organization and may not be typical. GD-181a.

©2026 Advanced Micro Devices, Inc. all rights reserved. AMD, the AMD arrow, and combinations thereof, are trademarks of Advanced Micro Devices, Inc. Other product names used in this publication are for identification purposes only and may be trademarks of their respective companies. Certain AMD technologies may require third-party enablement or activation. Supported features may vary by operating system. Please confirm with the system manufacturer for specific features. No technology or product can be completely secure.

PID: 264651300