eaidb

state of the

RAI Ecosystem

2025 Annual Report

Foreword

We can only drive fast when we know we have good brakes.

EAIDB's mission is to unearth and celebrate the startups that enable safer, trustworthy, and transparent AI systems (i.e., the brakes). This report covers 2025 and its implication on each and every category of RAI.

2025 was a record year for us. Our reports have now been viewed **over 20,000 times** and our LinkedIn community now **exceeds 5,000 members**. Trust and responsibility is clearly more important than ever in the era of Agentic AI. We will continue to drive this momentum forward and make sure those working in this space continue to be highlighted and supported.

Without better brakes, we can't go faster. It's as simple as that.

— Abhinav Raghunathan, Founder of EAIDB
abhi@eaidb.org

eaidb EAIGG

The Responsible AI Ecosystem

FY 2025 Market Map

Tracking 323 responsible AI enabling startups.
Visit eaidb.org for reports & insights or to get your startup on the map!

MODEL & PLATFORM BUILDERS (74)

- abodoo
- abzu
- aeterna labs
- akira AI
- algoface
- Ambient.ai
- BAST
- black.ai
- Bodyguard.ai
- BRIA
- BrightHire
- brighter AI
- Cangrade
- causaLens
- Celantur
- checkstep
- CONJECTURE
- CROSSCHQ
- DEEP MEDIA
- DEEPING SOURCE
- DEVELOP DIVERSE
- digiLab
- DIVERSIO
- EQUALTURE
- evispot
- Exparang
- fairplay
- Headai
- Hippocratic AI — Do No Harm —
- Humanly
- impress
- Kanarys
- KNAC
- knockri
- Lauretta.io
- LELAPA AI
- Liquid
- MABEL
- MATHISON
- MEVITAE
- MODULATE
- multitudes
- Nuanced
- OPTIMAL
- pave
- PickleTech
- piktid
- ProRata.ai
- PROVEN BASE
- QuadFi
- Re:cast AI
- Reejig
- retorio
- sapia
- seekout
- Stratyfy
- syntonym
- TRU RECOGNITION
- Tengai
- Textio
- TRUSTLAB
- UMNAI
- unitary
- VirtuousAI
- wayhaven
- witty.works
- XOPA AI
- ZELROS
- mpathic
- Alva
- elemental cognition
- gallio
- LINKAI
- WHITEBOX HR

MODEL OPERATIONS (67)

- AiMon
- Aiceberg
- Advai
- ALIGNED AI
- aporia
- arize
- ARMILLA
- Arthur
- ATHINA AI
- atla
- AutoAlign
- AYMARA
- census
- CERTIF.AI
- Citadel AI
- Cleanlab
- code4thought
- Collinear
- Confident AI
- daios
- datatron
- deepchecks.
- DEEPLOY
- Etiq AI
- Feedback Intelligence
- fiddler
- Fireraven.ai
- Galileo
- Giskard
- Guardrails AI
- Haize Labs
- Halluminate
- Hamming
- HoneyHive
- Judgment Labs
- kolena
- [redacted]
- Langfuse
- lastmile AI
- LatticeFlow
- leap labs
- mind FOUNDRY
- mona
- numalis
- Openlayer
- Patronus AI
- PEGASI
- qualifire
- QUANTPI
- resaro
- RevAIser
- Rhesis AI
- SECLEA
- SELDON
- SOLAS AI
- STYRK AI
- trstai
- THEMIS AI
- traceloop
- trubrics
- trustwise AI Trust. Simplified.
- VALIDAITOR
- WOLLED AI
- Warden AI
- ZETANE
- ASHER INFORMATICS
- SigmaRed

CONSULTING & LEGAL (43)

- AI & Partners
- Adaptive.AI
- All in on data
- AlyData
- babl
- Best Practice AI
- BUI CONSULTING
- The Center for Inclusive Change
- DEXAI
- DAEDALUS FUTURES
- ETHICAL INTELLIGENCE
- ETHICALLY ALIGNED AI
- ETHOS AI
- eticas
- expertisecentrum DATA-ETHIEK
- harmless.
- INQ CONSULTING
- Indie Pacific Legal Research
- innovethic
- IsA
- Kairoi
- leiwand.ai
- ORCAA
- ProceedAI
- QUANT AI
- Rhite
- SAFE AI NOW
- TECHNNOCENS
- TRUSTVECTOR
- Ulysses AI
- VERDAS AI
- zertia
- AIEthica
- AI Governance
- Anekanta AI
- ASSESSED INTELLIGENCE
- Barrington Digital
- Cantellus Group
- CyberEthicsLab
- immanence
- Social Tech Lab.eu

AI GOVERNANCE, RISK & COMPLIANCE (40)

- 2021.AI
- AI Guardian
- AffectLog360°
- AKOS AI
- BreezeML
- calvinrisk
- Certkor AI
- CitrusX
- COGNITIVE VIEW
- credo ai
- enzai
- Fairo
- Galini
- Guardrail Technologies
- Holistic AI
- InfAI
- Intelligible
- KNOSTIC
- KONFER
- LUMENOVA
- MISSION CONTROL
- ModelOp
- MONITAUR
- naaia
- prodago
- Saidot
- solidcore.ai
- SuperAlign
- Suzan AI
- SYNERGIST TECHNOLOGY
- trail
- TRUSTIBLE
- VALIDMIND
- VerifyWise
- VEROAI
- WITNESS AI
- zupervise
- Daiki
- KomplyAi
- modulos

PRIVACY PRESERVATION (42)

- AIDAMASK
- betterdata
- Bitfount
- BlueGen
- caber
- clearbox AI
- CLOAKAI
- CloudTDMS
- datacebo
- Datavillage
- Datawizz
- DEDOMENA
- FAIRGEN
- FedML
- gretel
- hazy
- HOWSO
- inpher
- integrate.ai
- MDCLONE
- MOSTLY·AI
- NENNA AI
- OCTOPIZE MIMETHIC DATA
- onetrust
- PRIVATEAI
- PROTOPIA
- RYVER AI
- Sourcepoint
- Synthesis AI
- SYNTHESIZED
- Syntheticus
- SYNTHO
- tomtA
- TONIC
- TripleBlind
- TRUATA.
- VEIL.AI
- syntheticAIdata
- Flower
- mindtech
- RHINO HEALTH
- SYNDATA

AI SECURITY (30)

- a+me
- Acuvity
- ADVERSA
- Aurascape
- CRANIUM
- DeepKeep
- Dynamo AI
- Enkrypt AI
- EVIDENTLY AI
- GRAY SWAN
- harmonic
- HIDDENLAYER
- Liminal
- MINDGARD
- portal26
- preamble
- RIVAL
- Straiker
- SurePathAI
- Tibo
- TROJ.AI
- trustme.ai
- unbound security
- Vera
- vijil
- Virtue AI
- AIShield Powered by Bosch
- Prompt:
- MITHRIL SECURITY
- pillar

DATA OPERATIONS (27)

- Airbloc
- ANTIMATTER
- APHERIS
- CEREVOX
- destined ai
- Fairly Trained
- GCX
- illumr
- inrupt
- isahit.
- POIETO
- palqee
- Pinalyt
- privya
- RegulAItion
- Reliabl
- relyanceai
- secuvyai
- Snorkel
- tasq.ai
- Transcend
- trustworks.
- Unbiased
- vaisual.
- 1touch.io
- Humans in the Loop
- VISYM LABS

Generated on January 10, 2026.

Overview

Executive Summary • New Editions • Funding & Growth • Global AI Legislation • M&A Activity • Agentic Challenges

Executive Summary

1 **Static controls are no longer enough.** Security and GRC priorities for AI are shifting left towards runtime.

2 **The legislative / policy gap in GRC is closing.** Global governments are actively drafting legislation and assigning roles to manage growing AI usage.

3 **AgentOps is a fundamentally different paradigm than LLMOps.** Enterprise-wide observability, registries, and control towers are critical technical components in enterprise agent strategy.

4 **AI explainability and interpretability is still lacking, especially for GenAI.** Healthcare and financial fields are feeling the effects as leadership face high barriers to entry around RAI and trust.

5 **The data gap is wider than ever.** Enterprises are so unprepared in terms of having scalable data sources for AI that they now turn to synthetic data (either build, buy, or acquire) to augment their pipelines.

6 **Litigation around AI is evolving.** We are seeing completely new legal domains such as "bot bad-mouthing" and competitor data poisoning for favorable AI SEO opening up.

New Additions

Walled AI · trustwise · atla · resaro · Langfuse

KNOSTIC · Cleanlab · Straiker · Openlayer · Acuvity

qualifire · caber · VALIDMIND™ · COGNITIVE VIEW · AIceberg

Collinear · lastmile AI · wayhaven · HoneyHive · Intelligible

zertia. · Aurascape · EVIDENTLY AI · Warden AI · Galini

Judgment Labs · indicpacific.com · Virtue AI · laminar · Liminal

AYMARA · VerifyWise · WITNESS AI · NENNA AI · Reliabl

Rhesis AI · RIVAL · Haize Labs · SurePath AI · PEGASI

AffectLog360° · ev so · GRAY SWAN · Fireraven.ai · solidcore.ai · AKOS AI · DEEP MEDIA · AI+ Me

RAI Funding

The RAI space is maintaining consistent growth as AI and AI support structures become increasingly prevalent.

2024

59 rounds raised

$1,004m raised

2025

63 rounds raised ↑6.8%

$1,646m raised ↑63.9%

Investment Rounds by Category

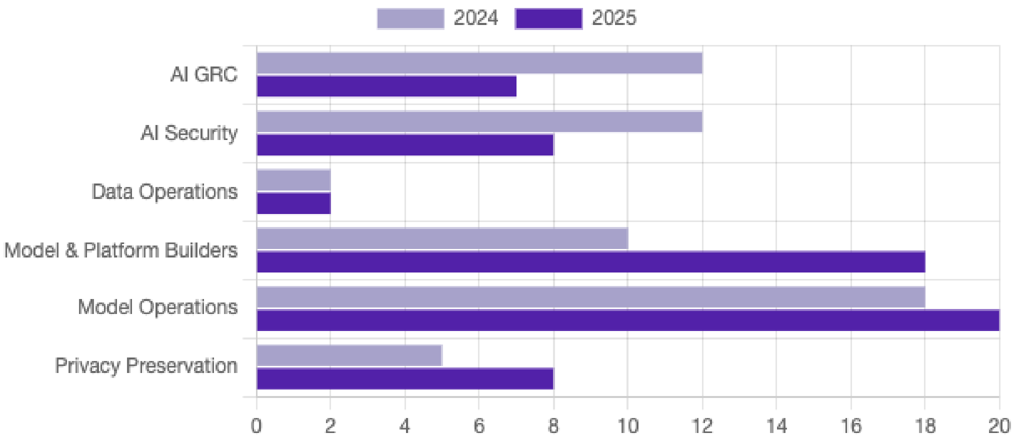

Funding Amounts by Category

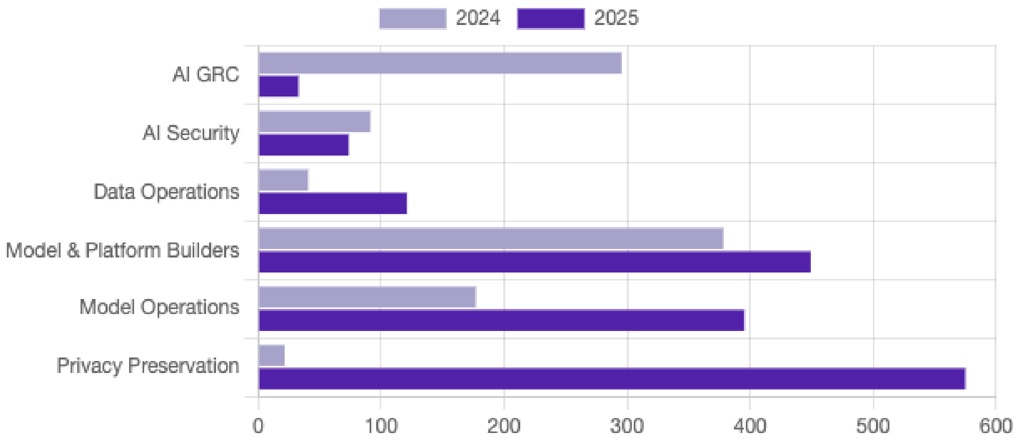

State of Global AI Legislation

Setup Implementation Enforcement

North America

- Lots of federal deregulation in the US, several smaller initiatives in specific states like Utah / New York (disclosure-based), Texas (explicitly banning certain uses for AI), and Arkansas (synthetic data ownership).
- United States: "America's AI Action Plan", "DOJ AI Litigation Task Force", executive orders.
- Canada: incremental requirements to existing frameworks (AIDA / PIPEDA).

Europe

- Key milestones reached with enforcement of the EU AI Act (GPAI rules, banned use cases for AI).
- Germany, Spain, Italy: draft AI Act implementation and related documentation.
- France: safety institute (INESIA), guidelines on GDPR, AI monitoring tool (PANAME).

Asia Pacific

- India: data protection acts.
- China: enforcement of GenAI measures.
- South Korea, Taiwan, Singapore: draft legislation around AI governance passed.

Middle East

- Egypt, Bahrain, Qatar, UAE: national AI strategy passed.
- Saudi Arabia: GenAI guidelines enforced.

South America

- Brazil, Chile, Colombia, Peru: various AI bills and legislation passed or passing.

Africa

- Nigeria, Kenya: national AI strategies and bills.
- Africa AI Council launched (AU / Smart Africa).

M&A Activity

Acquirer	Acquired	Category
Apple Inc. (NYSE: AAPL)	WhyLabs	Observability & Experiment Tracking
Datadog (NASDAQ: DDOG)	Sarus	Synthetic Data
KPMG US	YData	Synthetic Data
NVIDIA	Gretel AI	Synthetic Data
AuditBoard	FairNow	AI Evaluation
Anthropic	Humanloop	AI Evaluation
Check Point (NASDAQ: CHQP)	Lakera	AI Security
ZScaler (NASDAQ: ZS)	Splx AI	AI Security
F5 (NASDAQ: FFIV)	Calypso AI	AI Security
SentinelOne (NYSE: S)	Prompt Security	AI Security
Cato Networks	Aim Security	AI Security
Concentric AI	Swift Security	AI Security
Palo Alto Networks (NYSE: PANW)	Protect AI	AI Security
Snyk	Invariant Labs	AI Security
DataGalaxy	YOOI	AI GRC
Fairly AI	Anch.AI	AI GRC

🟪 EAIDB Company

Unsolved Enterprise Challenges with Agentic AI

Reliability & Determinism

Agentic systems are inherently much more difficult to build with as the set of actions they may be able to perform is diverse. **Creativity and unpredictability are two sides of the same coin.** This makes building complex agentic systems very difficult.

Lack of Readily Available Data

Agentic systems require a lot of testing and validation, much of which cannot be accomplished without ground truth datasets. These are usually not available by default and must be painstakingly compiled. We've seen a lot of enterprises double down on synthetic data to solve this.

Evaluation & Metrics

The industry still lacks a cohesive set of metrics that tie back to standards like NIST, EU AI Act, etc. Instead, enterprises are actively building their own evaluation pipelines, custom metrics and scores, and using unoptimized LLMs-as-a-Judge.

Security, Access & Permissions

Both single and multi-agent collaboration require permissions in the form of IAM and access management (also called "Know Your Agent" or KYA). This is still an unsolved problem, but hyperscalers are introducing this capability through their existing IAM platforms (Microsoft with EntraID, Amazon with AWS IAM, etc.).

Governance, Risk, Compliance

The very difficult task of getting a "lay-of-the-land" view of an enterprise's agents and how well they are adhering to corporate guidelines as well as legal policy is still at-hand. Especially considering each enterprise has to not only grade based on risk, usage, data, etc. but also scale this system in anticipation of a crowded agent market. Not to mention enabling this for both first- and third-party systems.

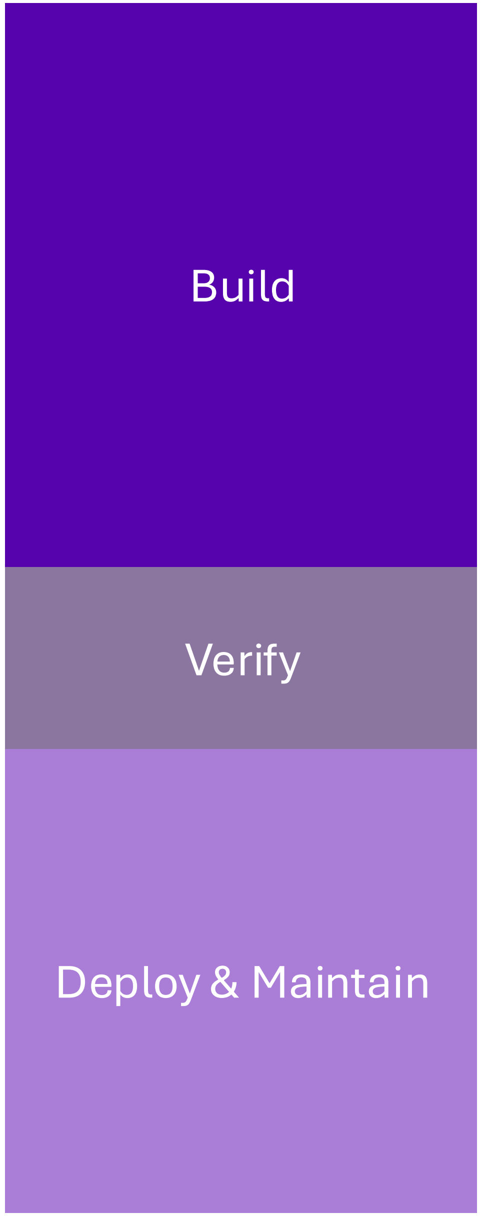

Build

Verify

Deploy & Maintain

AI Security

Industry Profile • Market Map • Deltas • Themes

eaidb

Industry Profile: AI Security

$62m
est. raised in 2025
(-31% vs. 2024)

6
rounds raised in 2025
(-45% vs. 2024)

0
entrants in 2025
(vs. 5 in 2024)

Inline
vs. incumbents

Key Takeaways

- Consolidation has taken the space by storm as major enterprises in network, application, and cloud security acquire growing AI Security startups to expand their market and offering to this new space.
- Hyperscalers are still behind, but this is changing as they are bringing in security solutions into their own agent builder platforms (Azure Foundry, Google Vertex, etc.).
- There is a major shift-left movement as AI Security moves away from reactivity post-system and towards proactive agentic and AI runtime security.

Market Incumbents

paloalto NETWORKS

IBM

Microsoft

Subcategory Breakdown

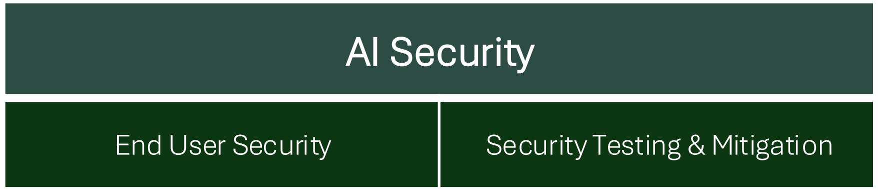

AI Security

End User Security

Security Testing & Mitigation

- **User-facing:** typically SaaS products deployable on-prem that interface with a "firewall" or protection layer between the end user and the application itself to guard against adversarial attacks, prompt injections, etc.
- **Developer-facing:** test suites or model red-teaming that help harden or secure AI applications prior to deployment.

AI Security

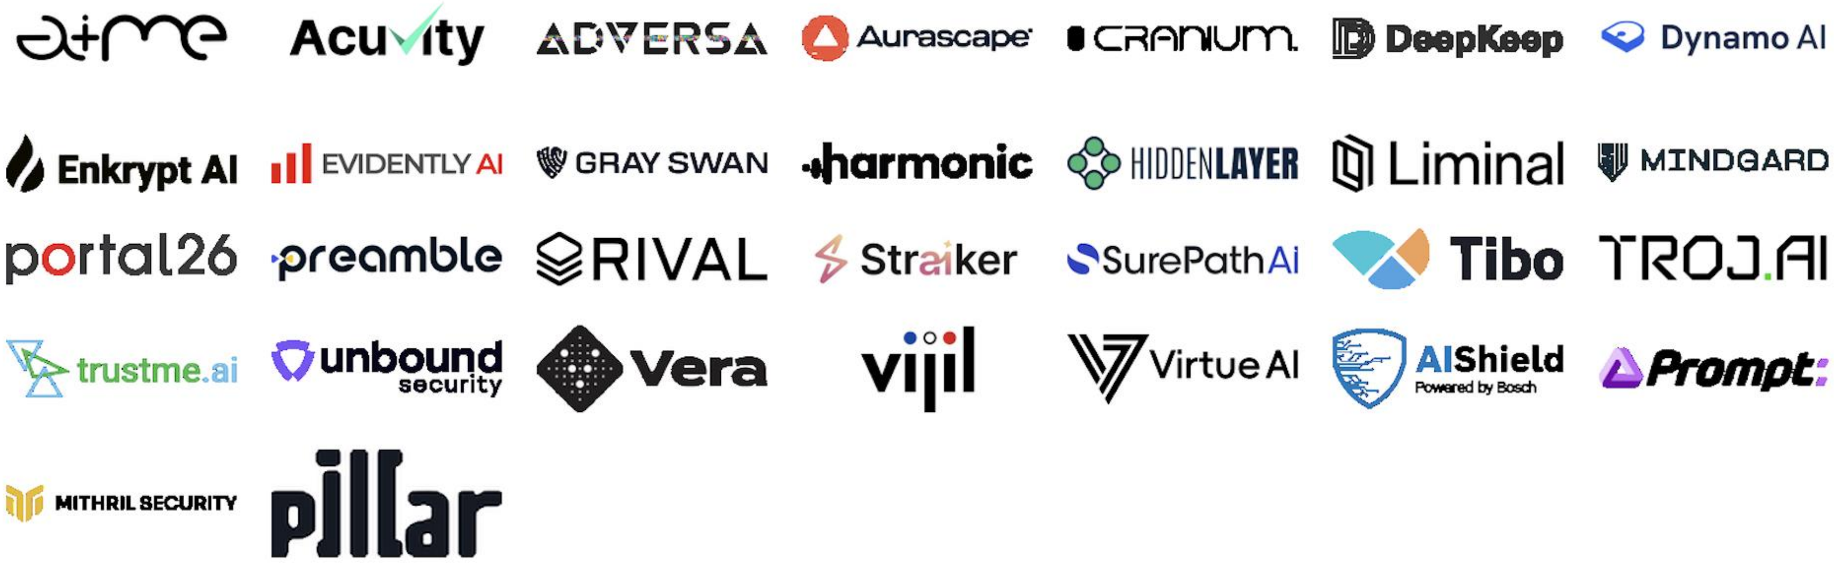
a+me
Acuvity
ADVERSA
Aurascape
CRANIUM
DeepKeep
Dynamo AI
Enkrypt AI
EVIDENTLY AI
GRAY SWAN
harmonic
HIDDENLAYER
Liminal
MINDGARD
portal26
preamble
RIVAL
Straiker
SurePathAi
Tibo
TROJ.AI
trustme.ai
unbound security
Vera
vijil
Virtue AI
AIShield Powered by Bosch
Prompt:
MITHRIL SECURITY
pillar

The Deltas

Our 2024FY commentary has already become outdated as companies rush to deploy and scale agents.

2024 Reality	2025 Reality	Forward-Looking Comments
Cloud providers are not well-equipped to handle attacks on models of any kind.	Hyperscalers are still lacking in this department, especially with agentic security.	We expect considerable consolidation in this area with respect to the big players (Google, Microsoft, Amazon). **Each are betting big on agentic runtimes and are trying to retain as much of the AI market as possible.** Microsoft, for example, announced Foundry Agent Service as a runtime and are now introducing Control Planes and Registries. Security is yet to come.
Palo Alto Networks (PAN) is the only large provider of all-encompassing AI security.	This is no longer true. We've seen acquisitions from big security players this year.	AI Security as a space has cooled off a little in terms of new entrants and funding as 2025 marked a period of exits. ZScaler, F5, and others have already brought in agentic security via acquisitions. **The next frontier is voice and browser-based agentic security.**
GenAI Security adoption may still be on the horizon.	GenAI security is here and critical for any production-grade agentic application.	With the number of agents in production steadily increasing (albeit, slowly), there is a tangible and obvious need for GenAI security adoption. Most enterprises recognize this already. For agentic and AI security, this is the era of scale, not experimentation.

Themes: Agent Security

AI runtime trust layers replace static checklists.

In 2024, AI defenses were mostly static - privacy policies, pre-production pen tests, and security checklists. In 2025, enterprises expect continuous runtime controls that fuse posture management with live defense. Calypso AI's Inference Defend and Observe now deliver real-time firewalls and oversight for production systems, while Enkrypt's Remove enforces adaptive policy guardrails across users and applications. Prompt Security, acquired by SentinelOne in 2025, illustrates how runtime enforcement is becoming central to enterprise security platforms. And with Unbound Security's proxy-based redaction and routing, organizations can dynamically sanitize or reroute prompts, proving that "in-flight" protection has become the baseline standard.

This area is still wide open, however, as agents get more complicated (multi-agents with A2A and Model Context Protocol (MCP)). Startups have heavily invested in these trends as well, with some providers like Acuvity, Lasso Security, and Vijil adopting MCP/A2A protection and runtime security into their existing products.

AISec platform additions in 2025	Examples
MCP / A2A discovery & scanning	Enkrypt AI, Zenity
Automated Red-Teaming	Straiker AI
Know-Your-Agent (KYA) & Agent Identity Management	Skyfire AI, Entro Security
Runtime Monitoring	Splx AI, Acuvity

Themes: AI Security

Red-Teaming becomes continuous and CI/CD Native.

In 2024, security testing was sporadic and manual. In 2025, adversarial testing is automated, continuous, and embedded into enterprise pipelines. **AIandMe** and **Mindgard** have lowered the barrier for CISOs with plug-and-play red-teaming platforms, while **Calypso’s** Inference Red-Team brings campaign-grade “Agentic Warfare” simulations to production systems. **Cranium** Arena enables enterprise-wide red-teaming with supply-chain visibility, and Evidently has transformed its open-source suite into **Evidently** Cloud, enabling synthetic and adversarial testing through no-code dashboards. Even industrial incumbents like **Bosch AIShield** now pair automated stress testing with runtime defense, showing that adversarial resilience is no longer a once-a-year exercise but a continuous requirement.

AI Supply Chain, Provenance and AI Bill of Materials (AIBOM) become table stakes.

In 2024, supply-chain security was a niche concern. In 2025, it is central to enterprise trust. **HiddenLayer’s** AISec 2.0 introduced Model Genealogy and AI Bills of Materials (AIBOM), offering full lineage visibility as incidents rise. **Cranium’s** AutoAttest and KYAI partnership extend this into third-party AI risk, while **Pillar Security**’s 2025 disclosure of malicious GGUF backdoors highlighted the urgency of hardening model repositories. **Mithril Security**’s AICert brings cryptographic provenance to training data and code, and Bosch **AIShield**’s Watchtower is now deployed across healthcare and automotive to scan notebooks and ML artifacts. Supply-chain provenance has shifted from best practice to regulatory and enterprise mandate.

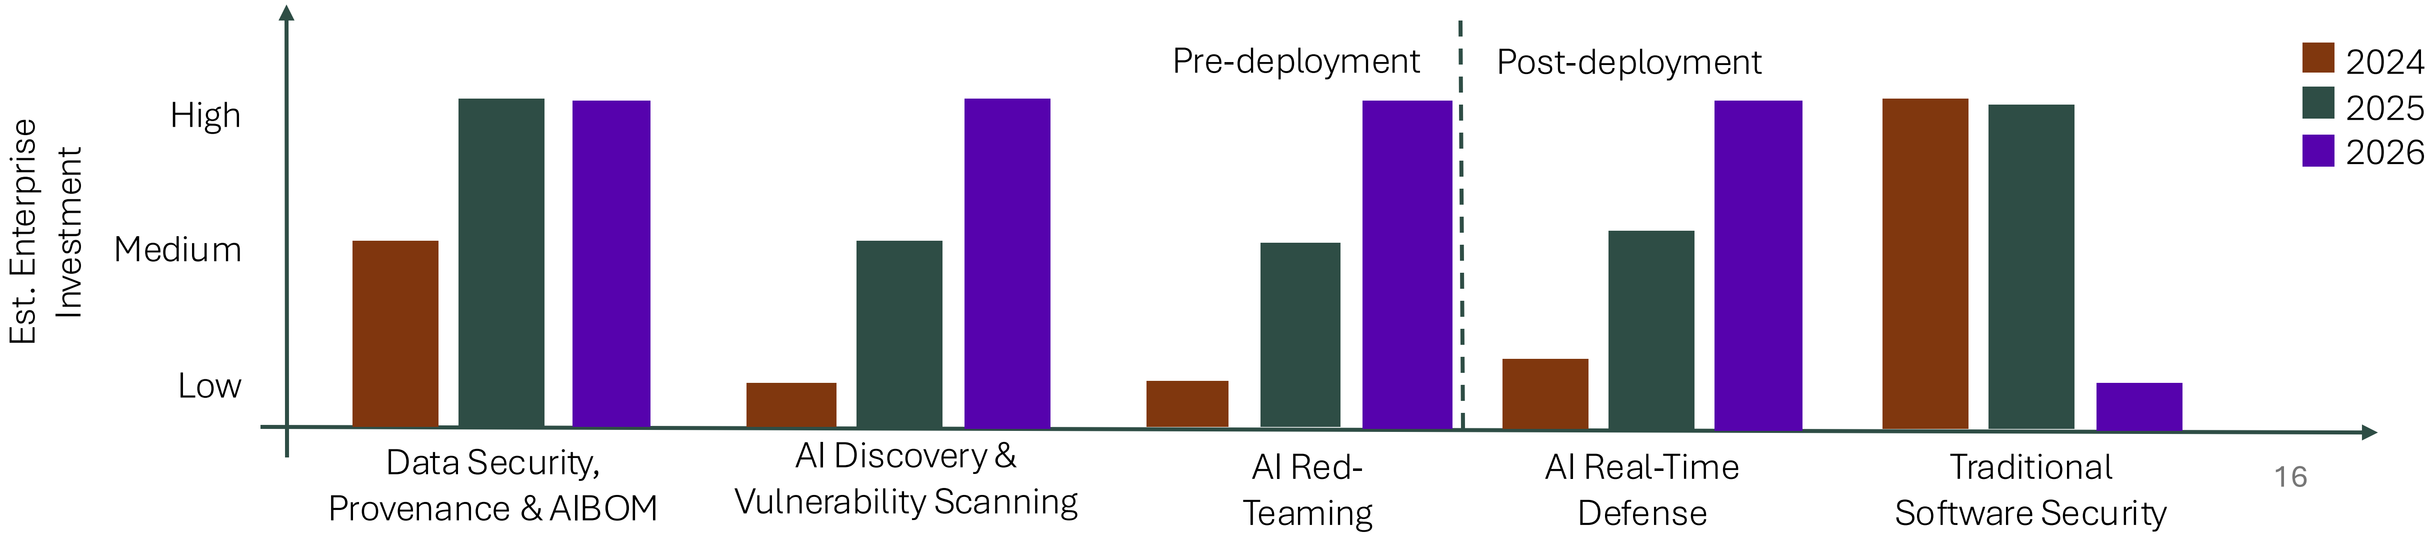

Three Predictions for 2026

1. **Voice and vision security and hardening is going to be pivotal.**
   Multimodal models have larger attack surfaces and security is much more immature at this time. This area will dramatically increase in criticality as enterprises expand into more modalities and expose this technology to more end users.

2. **Synthetic data acquisitions to bolster large data generation efforts for comprehensive adversarial and scenario security testing.**
   Comprehensive attack and risk coverage only becomes scalable and extensible with high-quality synthetic data. While many startups today can identify the attack vectors, they often lack the data generation to narrow down on specific test cases or case-by-case vulnerabilities.

3. **Voice and vision security is going to be huge.**
   Multimodal models have additional attack surfaces and security is much more immature at this time. This area will dramatically increase in criticality as enterprises expand into more modalities and expose this technology to more end users.

Model & Agent Operations

Industry Profile • Market Map • Themes

Industry Profile: Model & Agent Operations

$395m
est. raised in 2025
(+123% vs. 2024)

20
rounds raised in 2025
(+11% vs. 2024)

1
new entrant in 2025
(vs. 9 in 2024)

Ahead
vs. incumbents

Key Takeaways

- Evaluation metrics are turning multi-modal. Browser and voice agent behavior are gaining criticality.
- LLMOps startups are expanding into AgentOps, which is much more complicated but absolutely critical for production-grade systems.

Market Incumbents

Google Cloud
databricks
Microsoft Azure
aws

Subcategory Breakdown

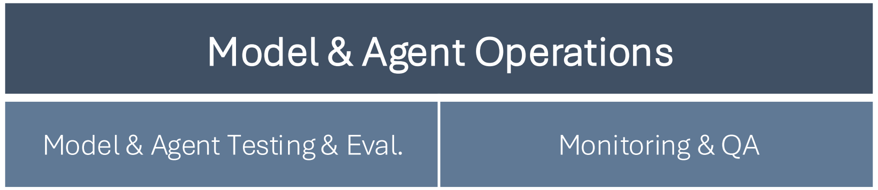

Model & Agent Operations

Model & Agent Testing & Eval.

Monitoring & QA

- Model Testing & Evaluation: typically developer tools or SaaS platforms capable of stress-testing, debiasing, experiment tracking, and evaluation.
- Production Quality Assurance: typically extensions of pre-production tools but focused on monitoring and ensuring high quality outputs.

* headcount data estimated with Growjo, values may be large because of outliers or small sample sizes.

Model & Agent Operations

AiMon | AIceberg | Advai | ALIGNED AI | aporia | arize | ARMILLA

Arthur | ATHINA AI | atla | AutoAlign | AYMARA | census | CERTIF.AI

Citadel AI | Cleanlab | code4thought | Collinear | Confident AI | daios | datatron

deepchecks. | DEEPLOY | Etiq AI | Feedback Intelligence | fiddler | Fireraven.ai | Galileo

Giskard | Guardrails AI | Haize Labs | Halluminate | Hamming | HoneyHive | Judgment Labs

kolena | [redacted] | Langfuse | lastmile AI | LatticeFlow | leap labs | mind FOUNDRY

mona | numalis | Openlayer | Patronus AI | PEGASI | qualifire | QUANTPI

resaro | RevAIsor® | Rhesis AI | SECLEA | SELDON | SOLAS AI | STYRK AI

trstai | THEMIS AI | traceloop | trubrics | trustwise AI Trust Simplified. | VALIDAITOR | WALLED AI

Warden AI | ZETANE | ASHER INFORMATICS | SigmaRed

eaidb

Themes: Model & Agent Operations

Evaluation metrics are quickly expanding past text-only mediums.

As the agentic world goes multi-modal, evaluations are growing to match. It isn't only text that matters anymore. Companies like **Hamming AI** for voice or **Halluminate** for browser-use agents are being quickly adopted across the board. **OpenAI**, **Anthropic**, and others are also at play with their own evaluation datasets and methods.

We are quickly seeing saturation of text-based metrics, with new and existing technologies from hyperscalers like **Microsoft**. Hyperscalers are integrating these metrics with their own agent builders like Foundry Agent Service.

AgentOps is a fundamentally different observability paradigm.

Agents are systems that encompass LLMs, and as such require different standards for controls, observability, etc. We are seeing many of the traditional LLMOps startups expand into agentic processes by adding more features around cost (which can be unpredictable in agents), action and decision intelligence, and more.

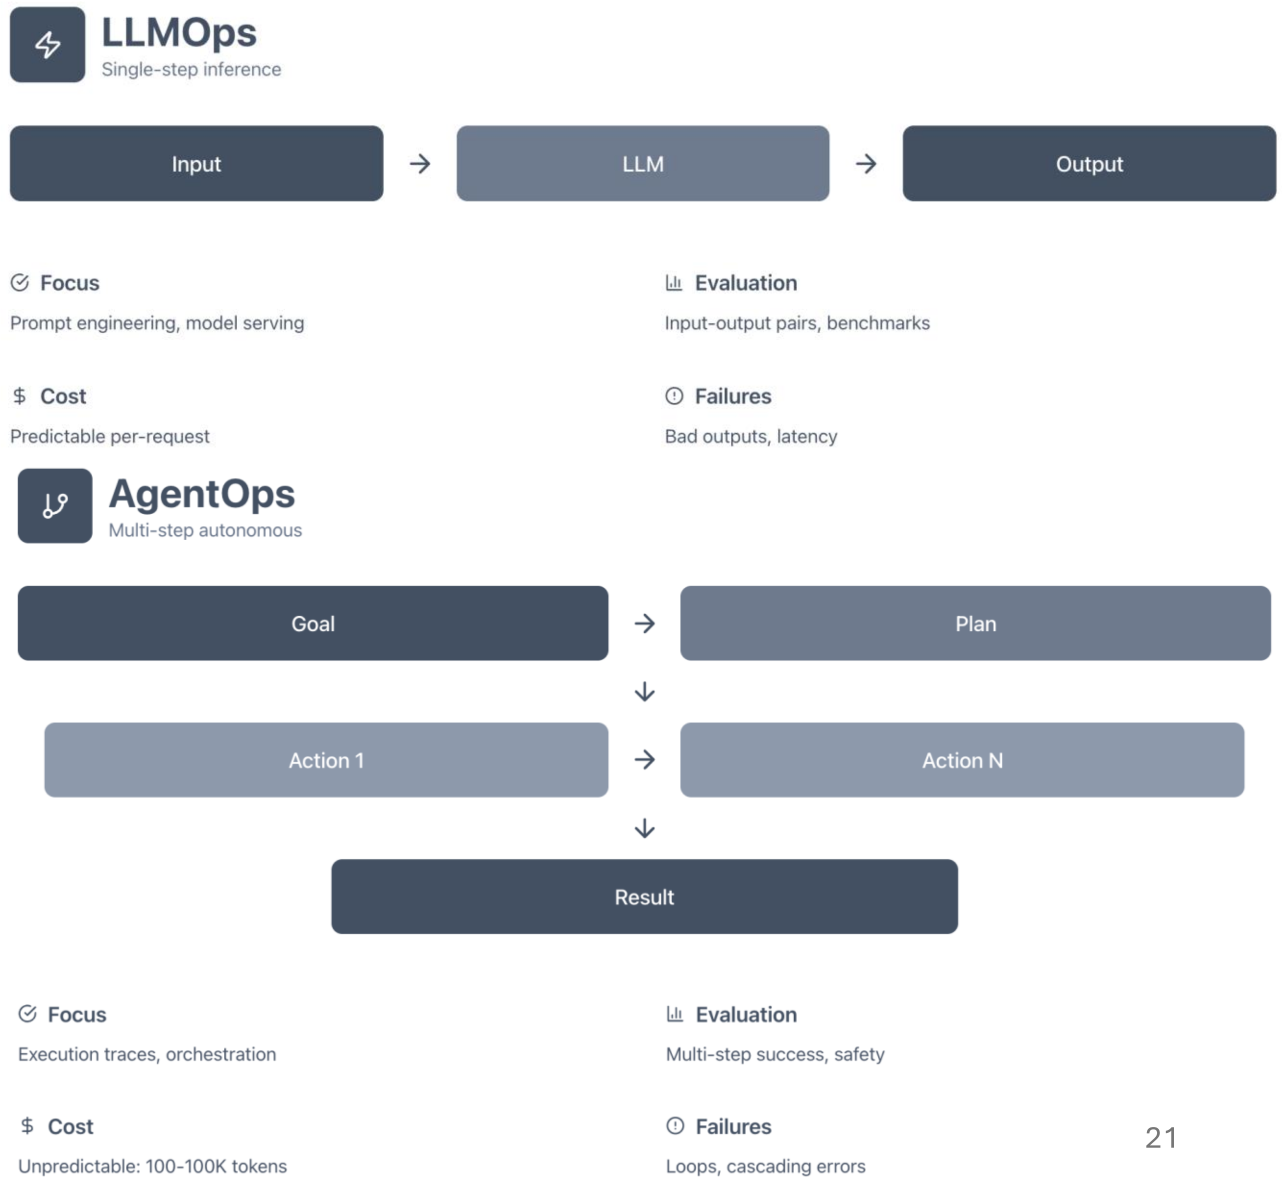

LLMOps
Single-step inference

Input → LLM → Output

Focus
Prompt engineering, model serving

Evaluation
Input-output pairs, benchmarks

$ Cost
Predictable per-request

Failures
Bad outputs, latency

AgentOps
Multi-step autonomous

Goal → Plan ↓ Action 1 → Action N ↓ Result

Focus
Execution traces, orchestration

Evaluation
Multi-step success, safety

$ Cost
Unpredictable: 100-100K tokens

Failures
Loops, cascading errors

AI GRC

Industry Profile • Market Map • Themes

eaidb

Industry Profile: AI GRC

$33m
est. raised in 2025
(-89% vs. 2024)

7
rounds raised in 2025
(-36% vs. 2024)

1
new entrants in 2025
(vs. 4 in 2024)

Inline
vs. incumbents

Key Takeaways

- Enterprise adoption of AI is driving GRC spend as enterprises allocate towards controls and mitigations.
- GRC is shifting left towards runtime – incumbents and startups alike are making plays to better scale.
- New regulation and enforcement is also playing its part in driving GRC spend.
- We are expecting some large consolidation in the market in 2026.

Market Incumbents

Vertex AI
watsonx
aws
Microsoft Azure
onetrust

Subcategory Breakdown

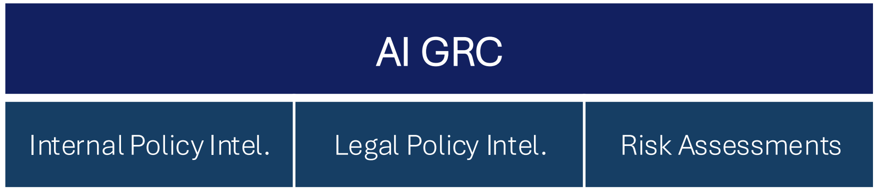

AI GRC

Internal Policy Intel. | Legal Policy Intel. | Risk Assessments

- **Internal Policy Intelligence:** solutions meant to track usage and implement guardrails according to an enterprise's internal definitions and policies.
- **Legal Policy Intelligence:** solutions meant to track projects and their adherence with common legal frameworks (ISO, EU AI Act, etc.).

AI GRC

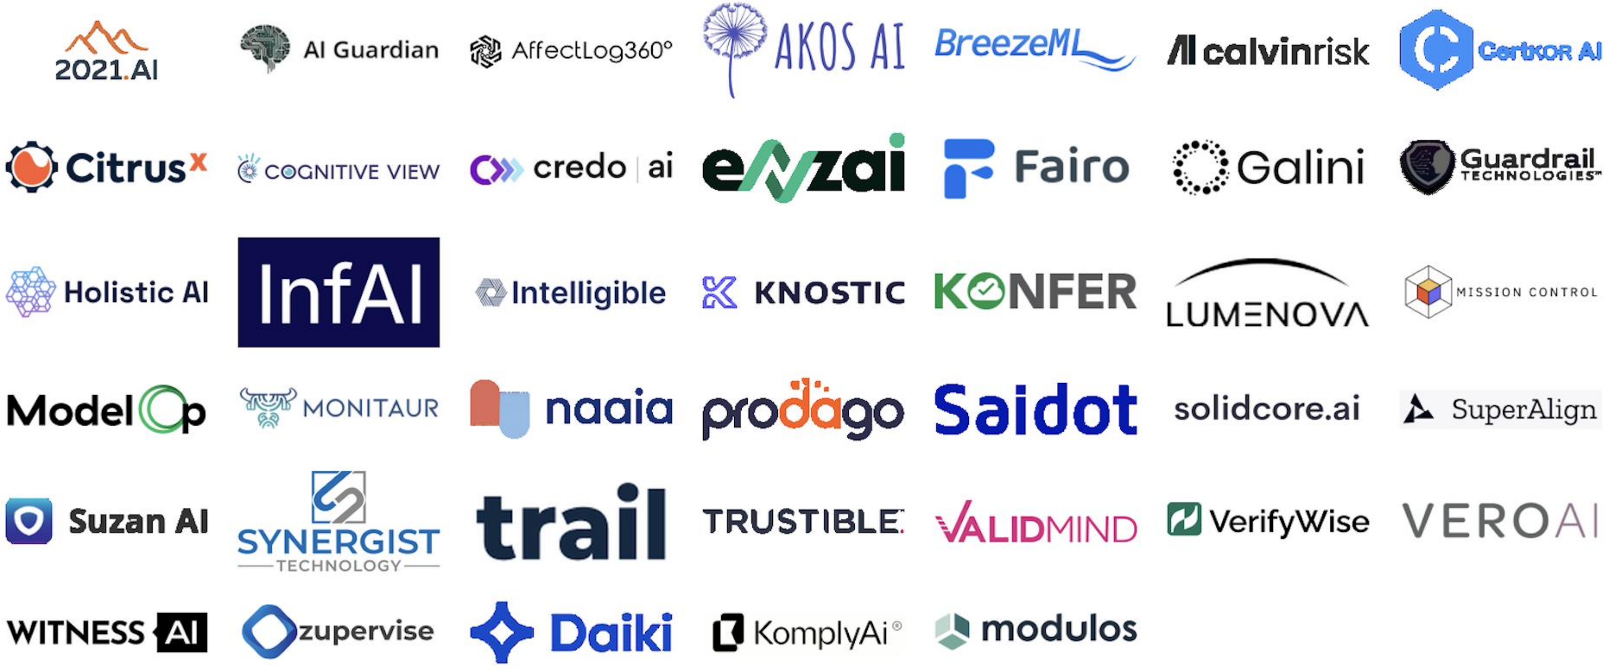
2021.AI
AI Guardian
AffectLog360°
AKOS AI
BreezeML
calvinrisk
CertKOR AI
Citrusx
COGNITIVE VIEW
credo ai
enzai
Fairo
Galini
Guardrail TECHNOLOGIES™
Holistic AI
InfAI
Intelligible
KNOSTIC
KONFER
LUMENOVA
MISSION CONTROL
ModelOp
MONITAUR
naaia
prodago
Saidot
solidcore.ai
SuperAlign
Suzan AI
SYNERGIST TECHNOLOGY
trail
TRUSTIBLE
VALIDMIND
VerifyWise
VEROAI
WITNESS AI
zupervise
Daiki
KomplyAi®
modulos

Themes: AI GRC

Enterprise adoption of AI is driving GRC spend.

According to a survey and study from Insight Partners, budgets and priorities are still focused towards the deployment side of the equation (LLMs, AI, etc.). The fact is, enterprises are still hard-pressed to show production-grade value with GenAI.

While most enterprises are decidedly increasing their budgets for GenAI, not many see the value of investing in AI GRC while much of their GenAI work is limited to pilots and proof-of-concepts.

The era of runtime agentic GRC.

22% of enterprises maintained that "regulatory and compliance risks" would be the biggest barriers for GenAI moving forward. While there are still significant headwinds for AI GRC today, enforcement of legislation like the EU AI Act will translate these concerns into allocated budgets for AI GRC.

The GRC Iceberg: Building is easy…

Countless tools and vendors make building trivial.

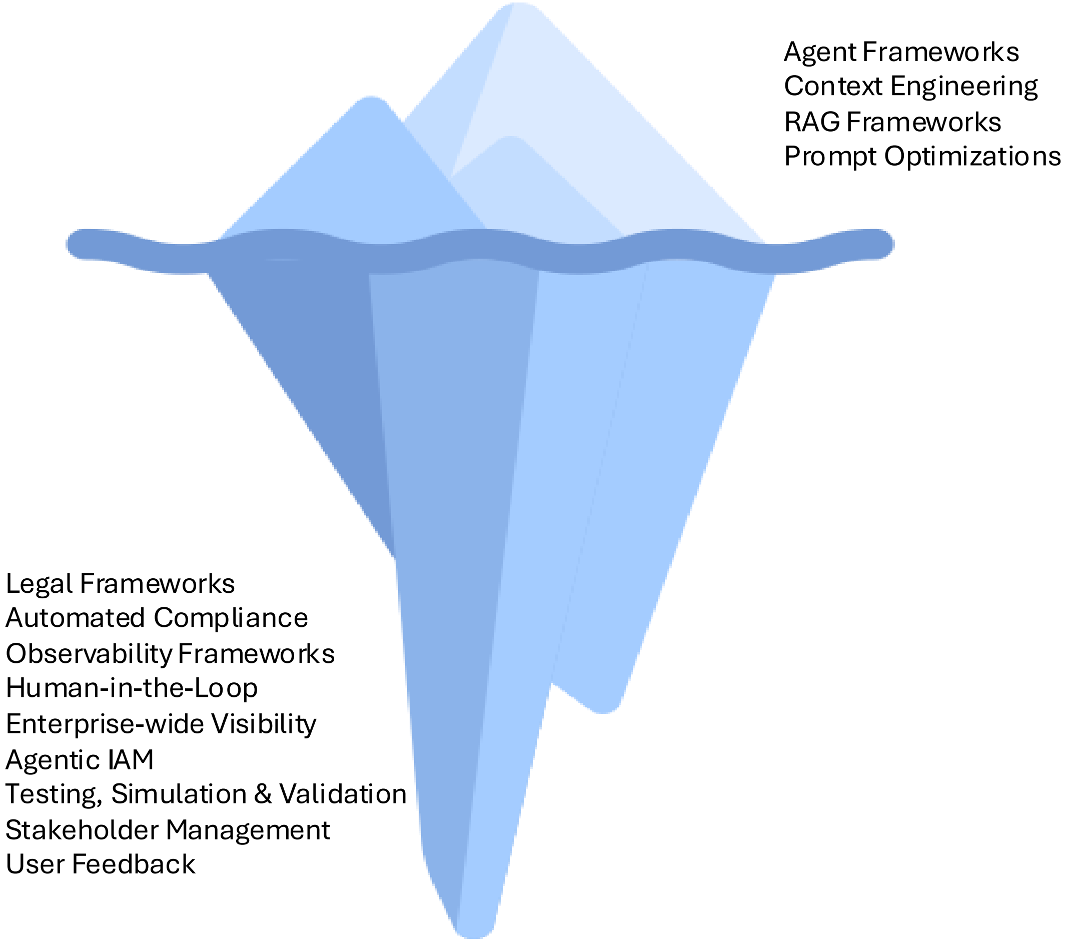

… but risk mitigation, and legal compliance is still difficult.

Setting up proper, automated GRC processes is still a fundamental challenge.

Themes: AI GRC

Regulation translates into budgets and buildouts.

Most incumbents provide limited access to automated GRC. Most stop with security or data privacy assessments (SOC2, HIPAA) and neither automate artifact creation nor post-production monitoring and compliance validation. IBM Watsonx.governance is the closest to what startup GRC vendors are supplying, but the gap is still substantial. One open question that remains is: how do enterprises think about agentic governance, risk, and compliance?

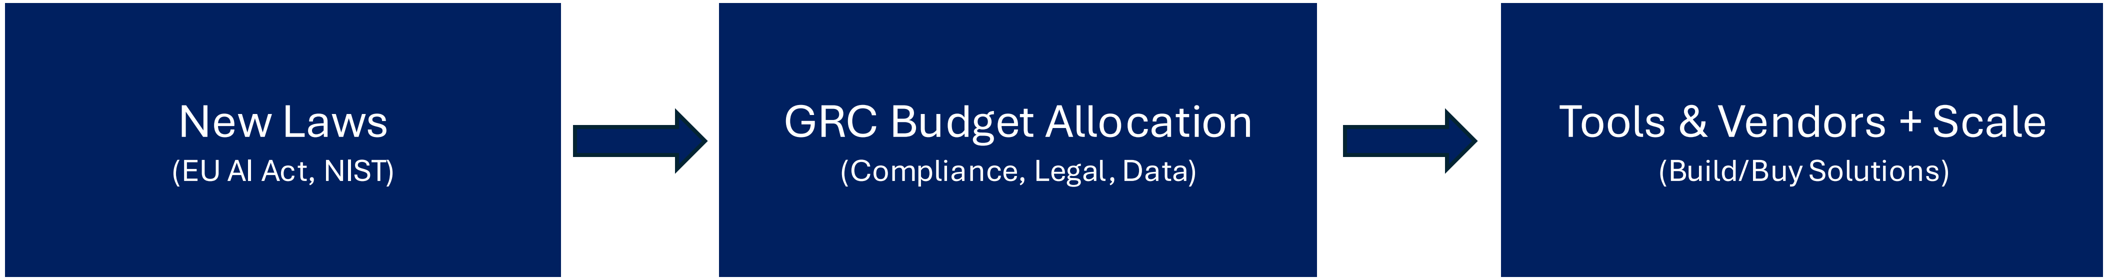
New Laws
(EU AI Act, NIST)

GRC Budget Allocation
(Compliance, Legal, Data)

Tools & Vendors + Scale
(Build/Buy Solutions)

There are some consolidation and ecosystem plays on the horizon.

There have been no new entrants into the AI GRC space so far in 2024. This is primarily because the space is dominated by a few general providers (Credo AI, ModelOp, 2021.AI) and many vertical providers (FairNow for HR, Monitaur for finance/insurance) that all provide almost identical product offerings.

Data Operations

Industry Profile • Market Map • Themes

eaidb

Industry Profile: Data Operations

$121m
est. raised in 2025
(vs. $41m in 2024)

2
rounds raised in 2025
(+0% vs. 2024)

2
new entrants in 2025
(vs. 2 in 2024)

Behind
vs. incumbents

Key Takeaways

- A lot of data companies that use humans to supplement synthetic data (Scale AI, Surge, Sama, etc.) are all facing multiple lawsuits for working conditions and misclassification. There is large opportunity for a new way of doing things.
- There is much more enterprise investment in synthetic data than traditional data curation with human experts due to control and ease of access.

Market Incumbents

databricks

TrustArc

scale

appen

Subcategory Breakdown

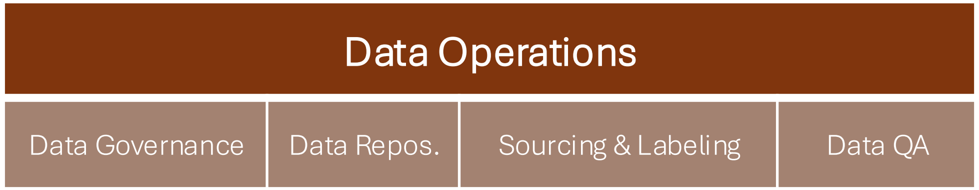

- Data Governance: companies offering extensive governance of data across the enterprise, focusing on PII tracking, privacy, and AI training.
- Sourcing & Labeling: solutions meant to responsibly source or annotate data at-scale.
- Data Repositories: pre-collected, ethically sourced, license-free datasets and collections prepared for AI.
- Data Quality Assurance: solutions offering pre-processing, debiasing, and other operations to enable more trustworthy AI.

Data Operations

Airbloc ANTIMATTER APHERIS CEREVOX destined ai Fairly Trained GCX

illumr inrupt isahit. POIETO palqee Panalyt privya

RegulAItion Reliabl relyanceai secuvyai Snorkel tasq.ai Transcend

trustworks. Unbiased vAisual. 1touch.io Humans in the Loop VISYM LABS

Themes: Data Operations

Generic labeler and annotation companies are running into legal and ethical issues.

Scale, Surge, and others who were responsible for data curation for flagship LLM providers like **Anthropic**, **Meta**, and **OpenAI** have recently come under fire for their misclassification and underpayment of workers. Responsible alternatives (**Humans in the Loop**, **Defined AI**) do exist, but they are still quite siloed.

Company	Lawsuit	Claims
Scale AI	McKinney v. Scale AI Rogowicz v. Scale AI DOL Investigation	Worker misclassification, wage theft, psychological harm
Surge AI	Cavalier v. Surge Labs	Worker misclassification, wage theft
Sama AI	Motuang v. Meta, Sama	Worker misclassification, poor working conditions, psychological harm

Synthetic data is applying pressure to data marketplace and repository startups.

The prevailing mindset in the current climate is one of "why use someone else's data when I can simply create my own?" Synthetic data generation technologies have been snapped up via M&A in 2025, despite the fact that there is no replacement for genuine, human-labeled data.

~25% CAGR

data marketplace market growth est.

~40% CAGR

synthetic data market growth est.

Three Predictions for 2026

1. **Human-labeled data will be making a comeback.**

   There is a fundamental limit to the utility that synthetic data can provide. As agents get more complicated, enterprises must realize that synthetics are a *supplement,* not the standard. Especially as model providers leverage more and more synthetic data (RLHAIF), there can only be compounding error without humans involved in the process. There is such a large market opportunity to do human labeling at-scale responsibly.

2. **"Data factory" startups are on the horizon.**

   Most synthetic data startups today place the control and the onus on the business. Some of the success of Scale AI, Surge AI, etc. can be attributed to a different business model where they take the burden of the data. However, there is space for a hybrid, almost. consulting business model. Uber has already moved here with "Uber Business Solutions," but we expect that the next generation of LLMs will bring some more competition in this field.

3. **Data services will become a source of revenue for some specialty institutions.**

   Experts will always hold the keys when data is in demand. Imagine organizations like Blue Origin "leasing" their experts out to SpaceX as data contributors, or Morgan Stanley doing the same for JPMorgan Chase. There is always a marketplace for information inflow and the wider the network the more diverse the dataset. We've seen similar trends for information marketplaces in other areas like pharmaceuticals, decentralized health via the blockchain, and others.

Privacy Preservation

Industry Profile • Market Map • Themes

Industry Profile: Privacy Preservation

$28m
est. raised in 2025
(24x vs. 2024*)

8
rounds raised in 2025
(+60% vs. 2024)

0
new entrants in 2025
(+0% vs. 2024)

Ahead
vs. incumbents

Key Takeaways

- Acquisitions galore in synthetic data as companies rush to secure their AI pipelines with scenario testing, evaluation datasets, fine-tuning datasets, etc., all of which are impossible without large quantities of data.
- Federated learning is becoming a very real way to interact with GenAI and create models in a distributed setting, increasingly important in healthcare, government, etc.

Market Incumbents

Informatica IBM ORACLE

Subcategory Breakdown

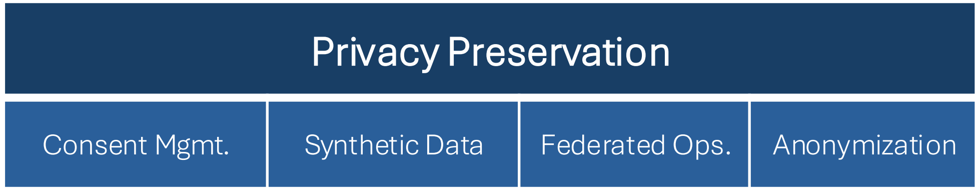
Privacy Preservation

Consent Mgmt. | Synthetic Data | Federated Ops. | Anonymization

- **Consent Management:** collection and maintenance of consent from end users across jurisdictions and legal frameworks.
- **Synthetic Data:** the generation of statistically accurate datasets that is completely synthetic. Could be used for fine-tuning AI or for general use.
- **Federated Operations:** usage of federated algorithms to source or analyze data, extending to Federated ML and Federated AI.
- **Data Anonymization:** a suite of techniques that obfuscate or mask specific PII to minimize leakage or exposure of damaging or private data.

* this number is inflated in part due to Onetrust's $200m Series C extension.

Privacy Preservation

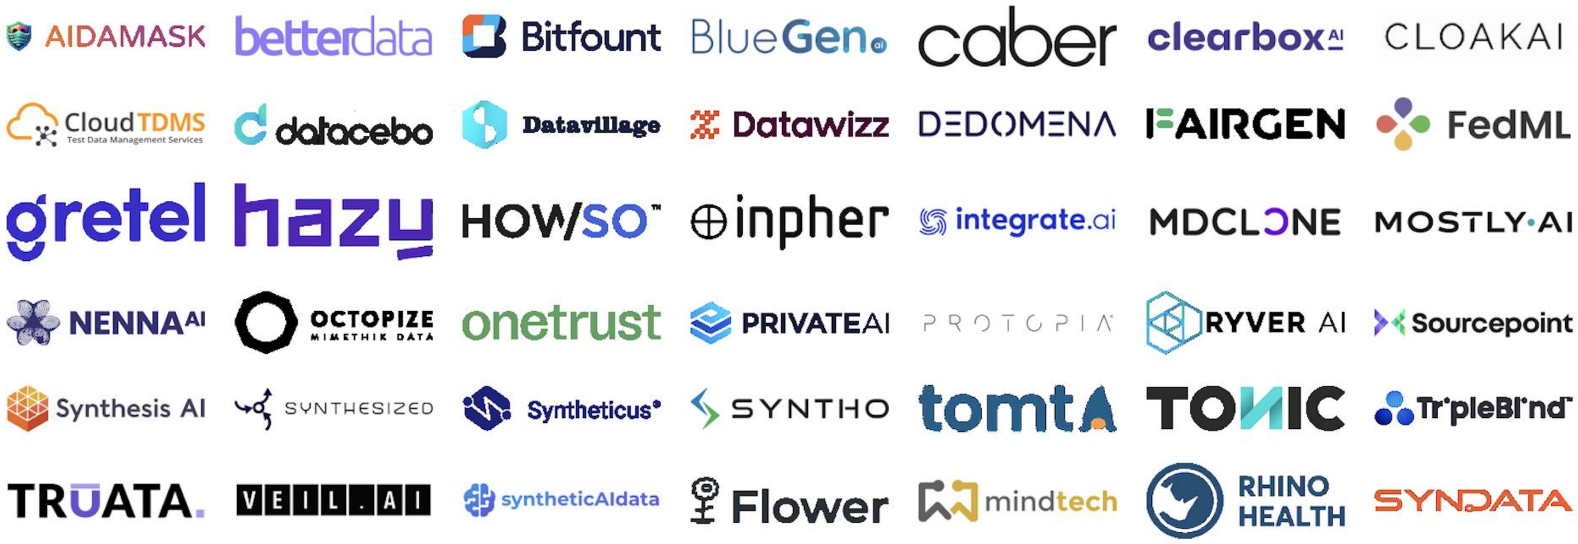

Themes: Privacy Preservation

Acquisitions & more acquisitions.

We've seen, finally, an unprecedented amount of acquisitions in the dataspace, particularly around synthetic data. As agents take center stage, builders are realizing that without good data, fine-tuning, testing, evaluation, etc. becomes impossible. **Synthetic data allows for organizations to greatly reduce the amount of human involvement they require for these datasets.**

Meta's 49% investment in **Scale AI** is another testament to the universal need of high-quality synthetic data. This is in almost complete contrast to 2024, which saw synthetic data companies struggling (**Datagen**, **Mirry AI**, **Syntric AI**, **Syntegra** to name a few).

KPMG + YData

DATADOG + Sarus

Meta + scale

NVIDIA + gretel

Federated learning meets GenAI.

AI is notoriously difficult to enable in cases where data must remain in its silo. Many federated learning startups (**Flower AI**, **TensorOpera**) now have capabilities allowing LLMs to be fine-tuned and inferenced using the advantages of privacy-preserving federated learning. This will becoming increasingly useful for applications in government and healthcare where data is essentially locked up.

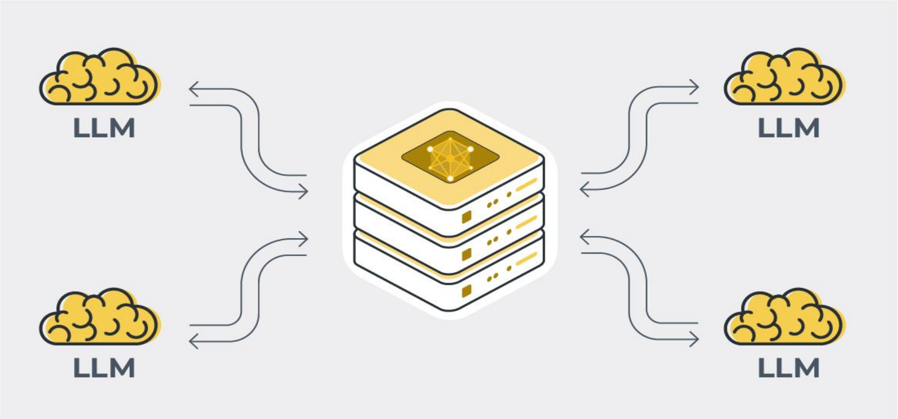

Image from Flower AI.

Model & Platform Builders

Industry Profile • Market Map • Themes • Healthtech & Fintech

eaidb

Industry Profile: Model & Platform Builders

$449m
est. raised in 2025
(+19% vs. 2024)

18
rounds raised in 2025
(+80% vs. 2024)

0
new entrants in 2025
(vs. 1 in 2024)

Ahead
vs. incumbents

Key Takeaways

- Reality detection (i.e., deepfake or AI-generated content detection) is a growing vertical and is scaling rapidly.
- Most RAI lab research is going towards model predictability and interpretability with items like "model steering" and "mechanistic interpretability."
- Trust in finance and healthcare is at an all-time low – while there is significant opportunity, only strong trust and RAI implementations will move the needle.
- There are high levels of AI adoption in healthcare, but professionals are still concerned about trustworthiness and ability to verify AI accuracy.

Subcategory Breakdown

Model & Platform Builders

HRTech | Fintech/Insurtech | Image, Video, Audio | Model Research

- HRTech, Insurtech, Fintech, Healthtech: vertically-oriented providers that provide safer, more responsible alternatives to what exists in the market.
- Image, Video, Audio: handling image privacy, anonymization, etc.
- Model Research: startups researching various divergent methods in GenAI and machine learning with the intent to distribute horizontally.
- Data Anonymization: a suite of techniques that obfuscate or mask specific PII to minimize leakage or exposure of damaging or private data.

Model & Platform Builders

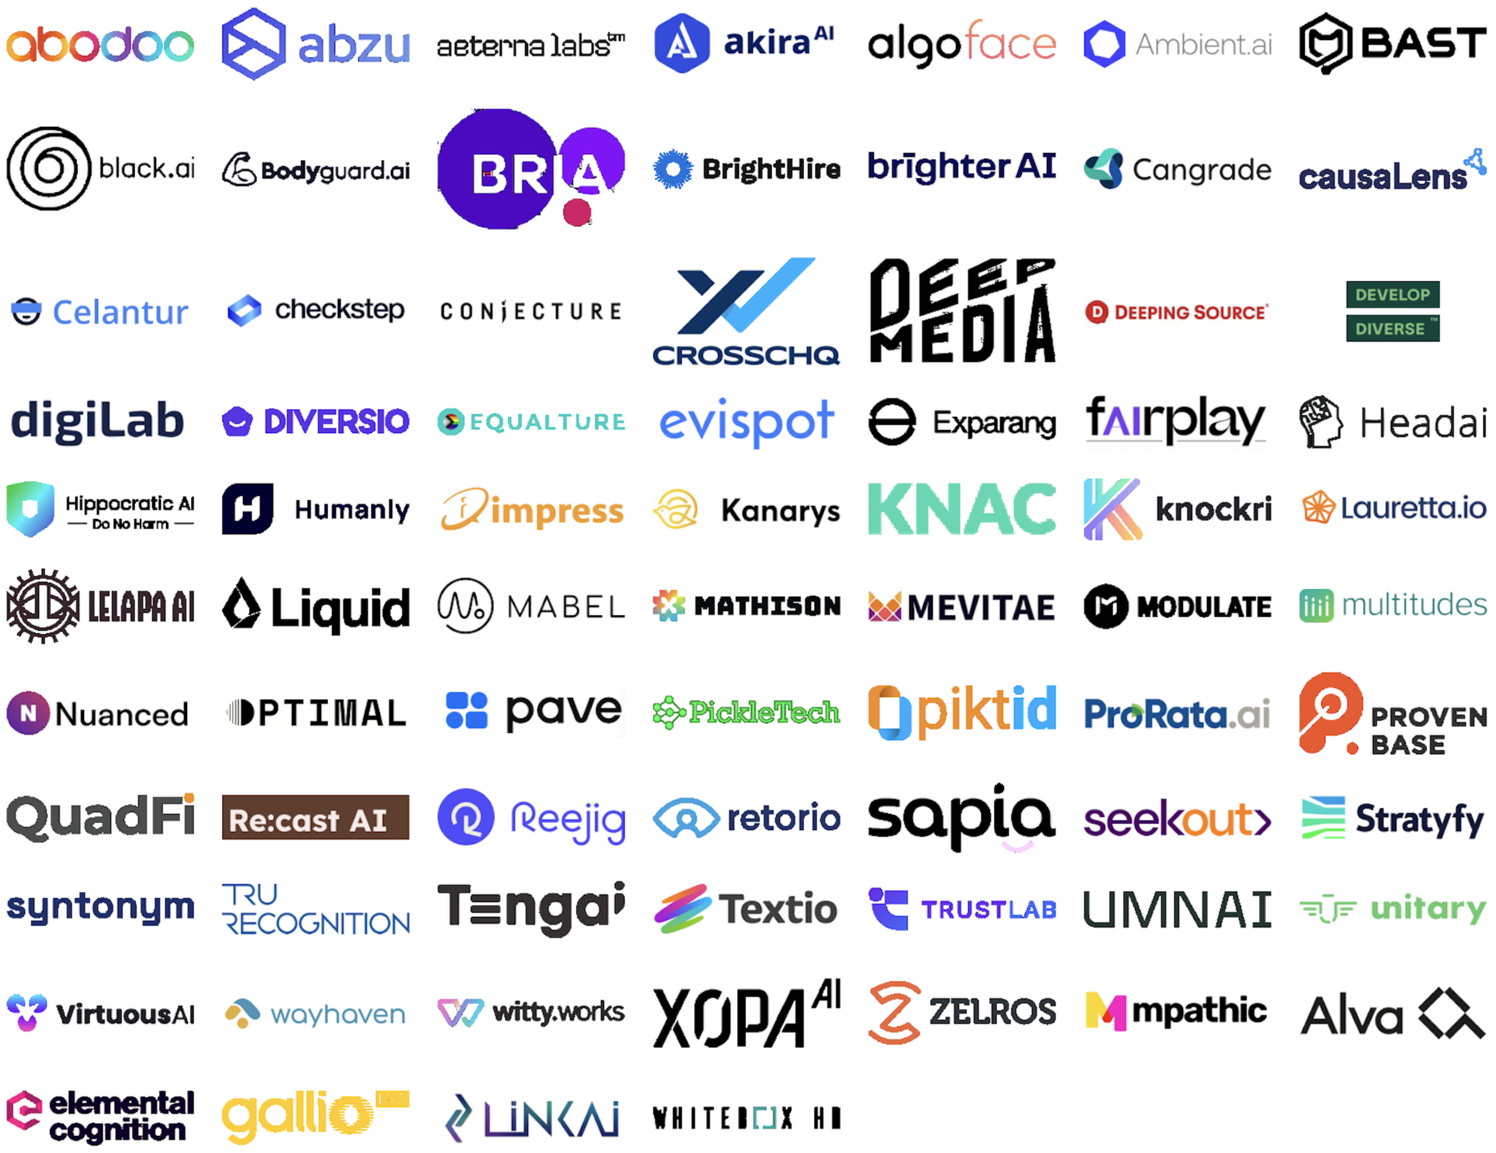

abodoo, abzu, aeterna labs, akira AI, algoface, Ambient.ai, BAST

black.ai, Bodyguard.ai, BRIA, BrightHire, brighter AI, Cangrade, causaLens

Celantur, checkstep, CONJECTURE, CROSSCHQ, DEEP MEDIA, DEEPING SOURCE, DEVELOP DIVERSE

digiLab, DIVERSIO, EQUALTURE, evispot, Exparang, fairplay, Headai

Hippocratic AI — Do No Harm —, Humanly, impress, Kanarys, KNAC, knockri, Lauretta.io

LELAPA AI, Liquid, MABEL, MATHISON, MEVITAE, MODULATE, multitudes

Nuanced, OPTIMAL, pave, PickleTech, piktid, ProRata.ai, PROVEN BASE

QuadFi, Re:cast AI, Reejig, retorio, sapia, seekout, Stratyfy

syntonym, TRU RECOGNITION, Tengai, Textio, TRUSTLAB, UMNAI, unitary

VirtuousAI, wayhaven, witty.works, XOPA AI, ZELROS, Mpathic, Alva

elemental cognition, gallio, LINKAi, WHITEBOX HR

eaidb

Themes: Model & Platform Builders

"Reality detectors" are growing increasingly relevant for enterprises.

With the rise of image and video generators and their rapidly improving realism, enterprises have a growing interest in identifying fakes. This is particularly relevant in financial domains like claims analysis, fraud, etc. In response to the demand and new tools like Google's Nano Banana, there are quite a few startups with advanced deepfake detection.

Reality Defender DEEP MEDIA CertifiAI

Cutting-edge labs are developing approaches to make LLM behavior predictable, interpretable, and use-case aligned.

LLMs are very difficult tools to use in production due to their various downsides such as unpredictability and behavior bias. Labs like Goodfire AI are tweaking these models by "steering" their behavior in real-time, making them much more aligned to business use cases. Others like mpathic AI are involved with embedding empathy in LLMs for use in healthcare settings.

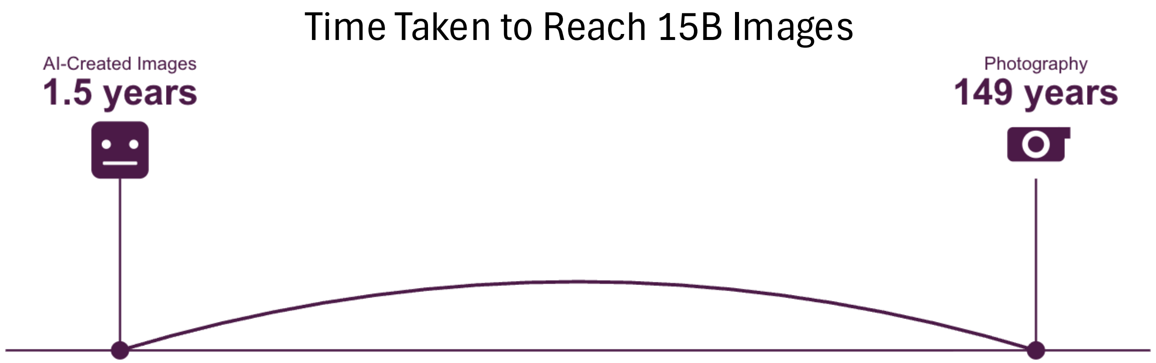

Source: Everypixel, OpenAI, Civitai, 1000Memories, Photutorial

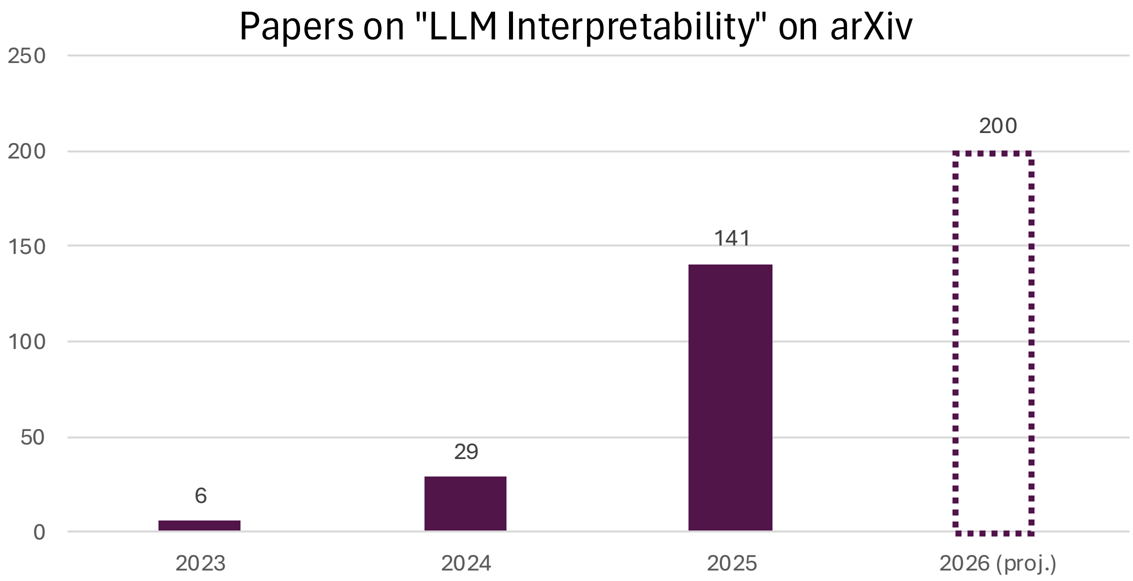

Themes: Model & Platform Builders

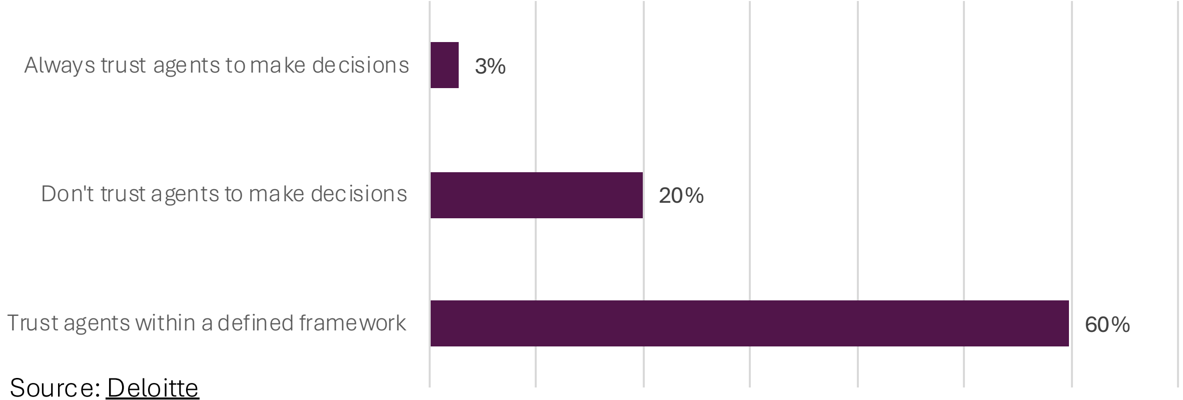
How Much Do Finance Professionals Trust AI Agents?

Survey of 3,300+ Finance & Accounting Professionals

Source: Deloitte

Trust in AI in the financial domain is at an all-time low.

With the onset of GenAI and LLMs, we've never further away from explainable AI. The statistics mirror this, with the financial industry doubling down on support structures for AI.

Financial professionals are targeting stronger governance frameworks (52%), clearer accountability for AI decisions (47%), improved data lineage and quality controls (45%), and broader role-specific training (43%) in 2026 (Source: Tipalti).

The largest barrier for agentic use in finance is trust in AI (Source: Deloitte).

Data quality, profit motives, and malpractice risks are the most prevalent concerns for healthcare organizations.

The low-hanging fruit in the healthcare space is the same as it is everywhere – automating operational complexity. More difficult use cases around diagnoses, clinical and patient notes, etc. are still fundamentally weak without strong trust, explainability, and RAI implementations.

90% of medical professional are concerned about malpractice risk from AI, while 72% of healthcare leaders are concerned about data privacy, not only in relation to cyberattacks but also about the profit motives of companies controlling AI systems (Source: Vention). There is a high level of AI investment, however, as AI budgets are outpacing IT budgets in healthcare organizations.

Top concerns in AI, per survey of 200 healthcare professionals

50%

data privacy & security

41%

lack of trust in AI accuracy

32%

cost of implementation

Source: Experian Health.

Consulting & Legal

Industry Profile • Market Map • Themes

Industry Profile: Consulting & Legal

Key Takeaways

- AI continues to be a driving force for top consulting firm bookings.
- Lots of market opportunity for boutique firms to fill a very important gap.
- New forms of legal action are brewing as AI permeates the internet and now becomes a primary source of information for consumers.

Market Incumbents

accenture BCG KPMG

McKinsey & Company Deloitte.

Subcategory Breakdown

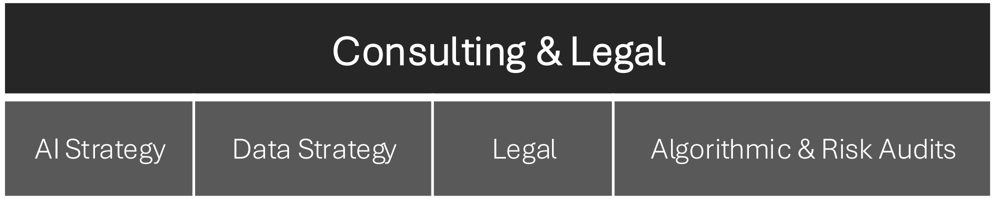
Consulting & Legal

AI Strategy | Data Strategy | Legal | Algorithmic & Risk Audits

- **AI Strategy:** designing AI systems for transparency, trust, and fairness.
- **Data Strategy:** designing the treatment of data within a system to ensure compliance and privacy.
- **Legal:** legal compliance and consulting specifically related to GenAI (copyrights, licensing, etc.).
- **Algorithmic & Risk Audits:** stress-testing an AI system for any kind of risk on case-by-case basis (no automated testing).

Consulting & Legal

AI & Partners
Adaptive.AI
All In on data
AlyData
babl
Best Practice AI
BUI CONSULTING

The Center for Inclusive Change
DEXAI
DAEDALUS FUTURES
ETHICAL INTELLIGENCE
ETHICALLY ALIGNED AI
ETHOS AI
eticas

expertisecentrum DATA-ETHIEK
harmless.
INQ CONSULTING
Indie Pacific Legal Research LLP
innovethic
IsA
Kairoi

leiwand.ai
ORCAA
ProceedAI
QUANT DI
Rhite
SAFE AI NOW
TECHINNOCENS

TRUSTVECTOR
Ulysses AI
VERDAS AI
zertia
AIEthica
AI Governance

Anekanta®AI
ASSESSED INTELLIGENCE
Barrington Digital
Cantellus Group
CyberEthicsLab.
ETHICIENS DU NUMERIQUE
immanence

Social Tech Lab.eu

Themes: Consulting & Legal

Boutique and large consulting firms continue to drive AI-enabled bookings upwards, but there is a rapidly expanding opportunity for boutiques. Nearly every major consulting firm now offers some form of trust services.

Consulting firms have embraced generative AI and agentic AI wholeheartedly as it greatly revamps existing offerings and both cuts costs and creates new sources of revenue. There is a lot of momentum towards smaller firms, with specialists like Tribe AI, Percepta, and others taking an active share in the market as both builders and strategy experts. We expect this will further compound as enforcement of governance and risk compliance related to the EU AI Act and the NIST framework turns suggestions into requirements.

KPMG

+5% global revenue

- Integration of AI in key platforms and offerings (Clara, Workbench, Digital Gateway, etc.).
- Bookings increased 10x since 2022.
- KPMG Trusted AI as a specialized offering.

Source: KPMG

Deloitte.

+5% global revenue

- $3B investment through FY2030 in GenAI.
- Offerings like Zoro AI, Deloitte Global Agentic Network.
- Deloitte Trustworthy AI Framework.

Source: Deloitte

accenture

+7% global revenue

- $3B investment through FY2030 in GenAI.
- GenAI bookings doubled, revenue tripled from 2024 to 2025.
- GenWizard, other AI-enabled offerings.

Source: Accenture

Themes: Consulting & Legal

New legal precedents and landmark cases beyond just copyrights are brewing.

2025 continued the ongoing trend of copyright lawsuits in the realm of AI. However, things are just getting started. There are many more issues on the horizon in the legal space around AI.

As traditional technologies become overwritten by AI-enabled equivalents, there is a larger-than-ever opportunity for legal firms.

Category	Status	Example
Copyrights	Multiple landmark cases in 2025 but is now a well-known issue.	"Frontier AI companies used my content without my permission, violating the terms and conditions of the license my content was protected by." (Bartz v. Anthropic, GEMA v. OpenAI)
Data Poisoning / Artificial Inflation	Upcoming issue as AI search engines take share away from traditional search engines.	"My competitors are creating bots to have artificial conversations on Reddit. These are getting picked up by AI search engines like Perplexity AI and ChatGPT and are misrepresenting my business."
Compliance & Governance Violations	Upcoming as the EU AI Act matures and becomes enforced.	"My AI system did not have the proper governance and controls to comply with the EU AI Act."

Appendix

Resources • About EAIDB

Primary Contributors

Abhinav Raghunathan

LinkedIn   Website

Abhinav (Abhi) is the Founder of EAIDB. He has extensive work experience in trust and evaluations in both traditional machine learning as well as GenAI and Agentic AI. He is a published TEDx speaker and frequently writes on topics within ethical AI including algorithmic bias / fairness.

Isabelle Patrick

LinkedIn

Isabelle specializes in cybersecurity resilience and AI governance at the Cyber Eagle Project. She also serves as an AI Innovation Fellow at EAIGG. Previously, she spent her career at AWS, working with UK government organizations to deliver large-scale enterprise cloud transformations and guide strategic technology adoption.

Resources

EAIDB Methodology

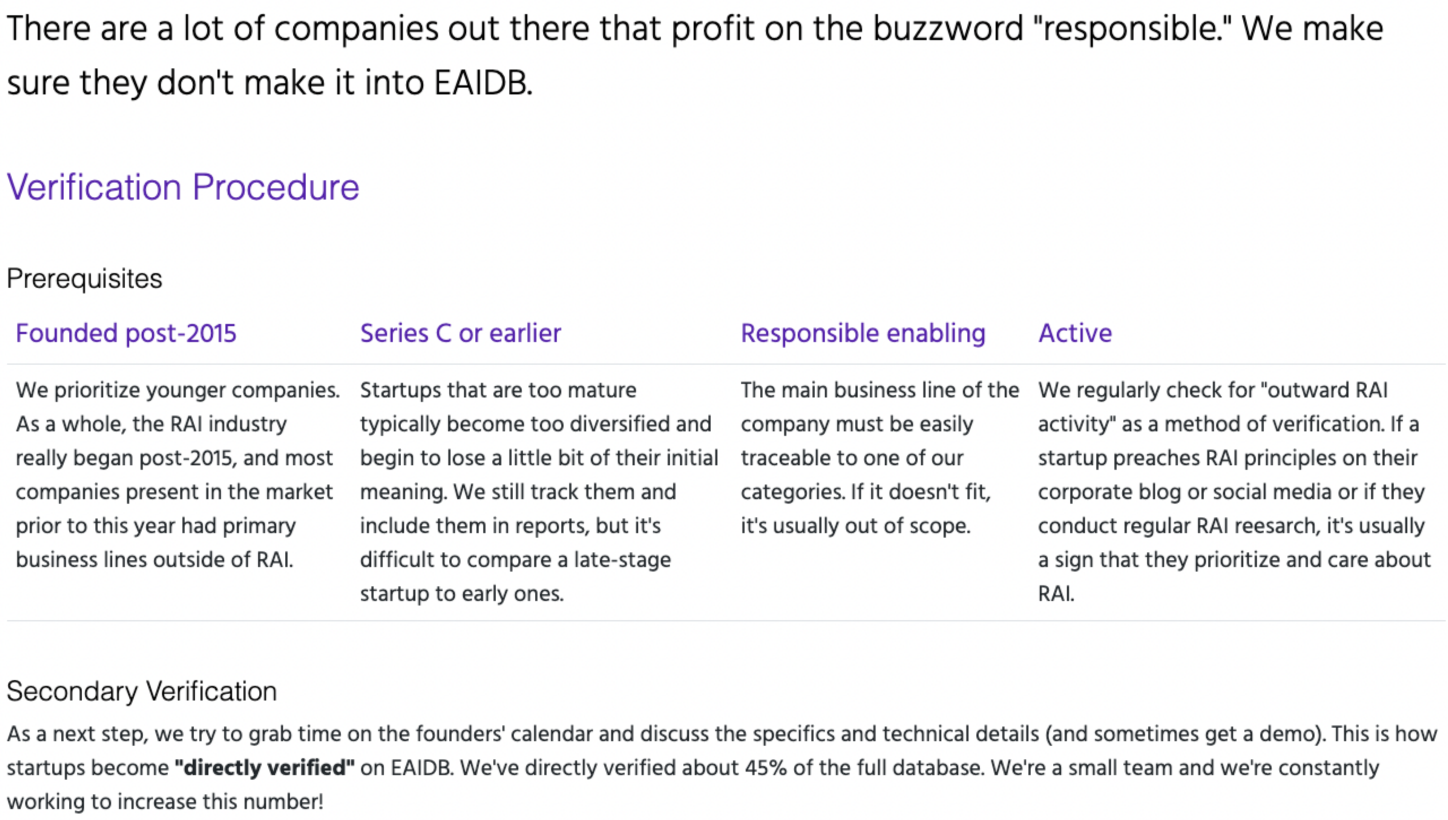

There are a lot of companies out there that profit on the buzzword "responsible." We make sure they don't make it into EAIDB.

Verification Procedure

Prerequisites

Founded post-2015	Series C or earlier	Responsible enabling	Active
We prioritize younger companies. As a whole, the RAI industry really began post-2015, and most companies present in the market prior to this year had primary business lines outside of RAI.	Startups that are too mature typically become too diversified and begin to lose a little bit of their initial meaning. We still track them and include them in reports, but it's difficult to compare a late-stage startup to early ones.	The main business line of the company must be easily traceable to one of our categories. If it doesn't fit, it's usually out of scope.	We regularly check for "outward RAI activity" as a method of verification. If a startup preaches RAI principles on their corporate blog or social media or if they conduct regular RAI reesarch, it's usually a sign that they prioritize and care about RAI.

Secondary Verification

As a next step, we try to grab time on the founders' calendar and discuss the specifics and technical details (and sometimes get a demo). This is how startups become **"directly verified"** on EAIDB. We've directly verified about 45% of the full database. We're a small team and we're constantly working to increase this number!

Submit a Company / Contact

To submit your company, feel free to submit an intake form.
To contact EAIDB, please reach out to abhi@eaidb.org.

About EAIDB

Our goal is to provide transparency into what is an otherwise opaque and nascent (though direly needed) industry. We provide differentiation between AI enablers and *responsible* AI enablers by conducting thorough market research on the responsible AI enablement industry. We source data from everywhere and work with all kinds of organizations, but are always **fully independent, transparent, and objective.**

Lead & Sales Sourcing

We've heard from several of our constituent companies that EAIDB's transparency and ability to filter, search, and compare companies within the same solution space has helped clients find products that match their needs.

Investment Sourcing

EAIDB has drawn attention from VC firms and founders alike and have exercised our unique ability to make connections between the two parties on more than one occasion.

Marketing & Promotion

EAIDB has **5,000+ followers** on LinkedIn and has received **20,000+ downloads** on our various reports. We attract attention from the public, policymakers, founders, and investors alike.

Market Research

As a fully independent organization, EAIDB sits in a place of objectivity and methodical approaches. We do market research on behalf of governments and organizations to investigate and identify market opportunities, profile companies, and offer in-depth comparisons of technology used.

eaidb