state of the

RAI Ecosystem

2024 Semiannual Report

Foreword

We're only able to drive fast because we know we have brakes.

Compared to machine learning – with its sometimes biased systems, privacy leakages, and unexplainable black box models – GenAI has both more potential and more risk. In order to move faster, we're trained to embrace the upside and ignore the downside.

While there is a much more intense focus from the market on companies that are blindly enabling GenAI at scale, EAIDB's mission is to unearth and celebrate the startups that enable safer, trustworthy, and transparent AI systems. We bring awareness to the **brakes** that allow AI to continuously **accelerate**.

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

— Abhinav Raghunathan, Founder of EAIDB

Overview

New Additions • AI Incidents • Funding & Growth • Global Representation • M&A Activity

New Additions

AI Incidents

Traditional AI/ML seems to be facing fewer problems in the wild, but GenAI is a different story.

According to the AI Incident Database (AIID), AI incidents for 2024 are forecasted to be about 500, **down 30% relative to 2023 totals**. However, GenAI incidents have kept pace, with around 150 total incidents in 2023 and 2024 (forecasted). Interestingly, this means that **GenAI incidents were a much larger percentage of the total in 2024 than ever before.** Maintaining strong product quality with GenAI continues to be a challenge.

Total incidents may be down because:

1. Traditional AI/ML is less covered in the media than GenAI, even when things go wrong.
2. Awareness of RAI is at an all-time high due to the significant shortcomings of GenAI – most executives are slowly incorporating better design thinking into their processes. This hasn't translated into purchases, but may be why there are less cases.
3. Large providers of data and models have incorporated checks for AI risk (fairness, adversarials, etc.) and have made them easy to implement.

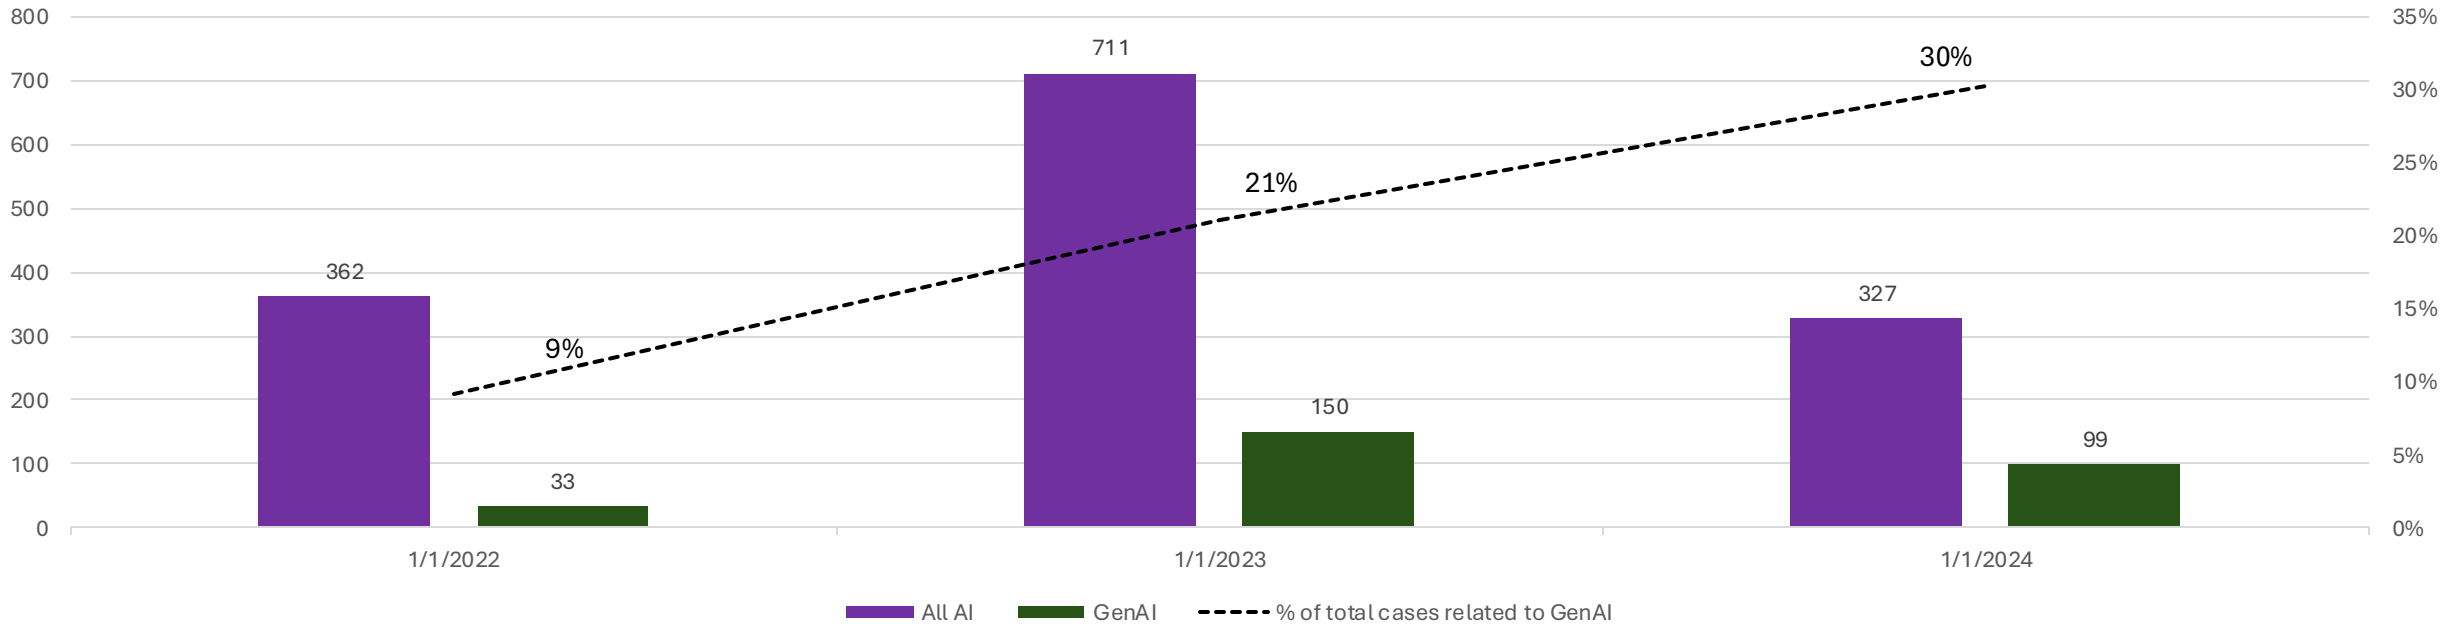

RAI Funding

The RAI space is closely tracking global macro trends in AI investing.

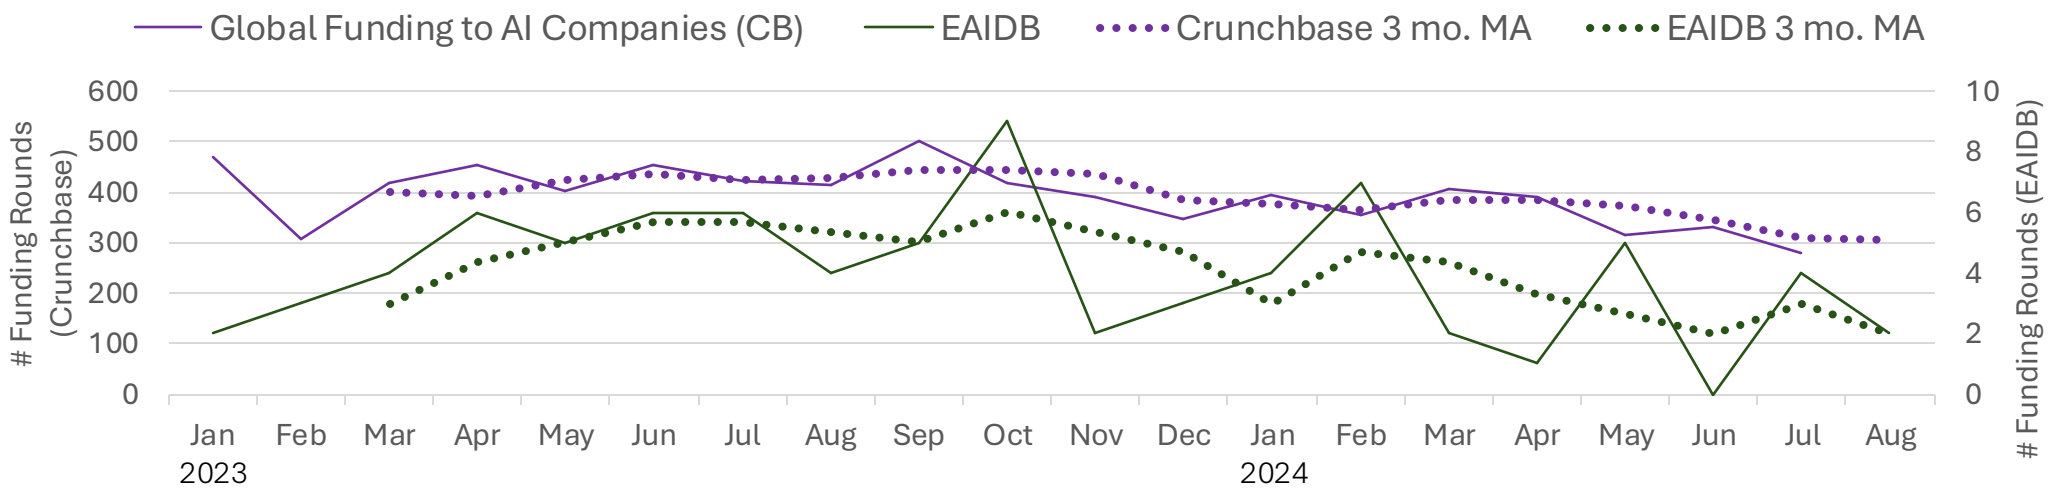

The 2024 AI funding environment is significantly more reserved than it was in 2021, with investors scaling back on rounds due to:

- Correction vs hyperinflated AI investments in 2021-2022,
- Macroeconomic conditions,
- Investor realization of technical challenges involved with true GenAI adoption in the enterprise,
- Public disdain and disillusionment from enterprises re: real value add across all GenAI use cases.

For the RAI space, additional barriers include:

- Enterprises doubling down on deployment and development to try and prove that GenAI does yield value,
- Enterprises choosing not to expand budgets for experimental items (AI GRC, better evaluation, etc.) without first demonstrating GenAI value.

RAI Funding (cont.)

AI Security leads RAI funding, closely followed by Healthcare Model & Platform Builders.

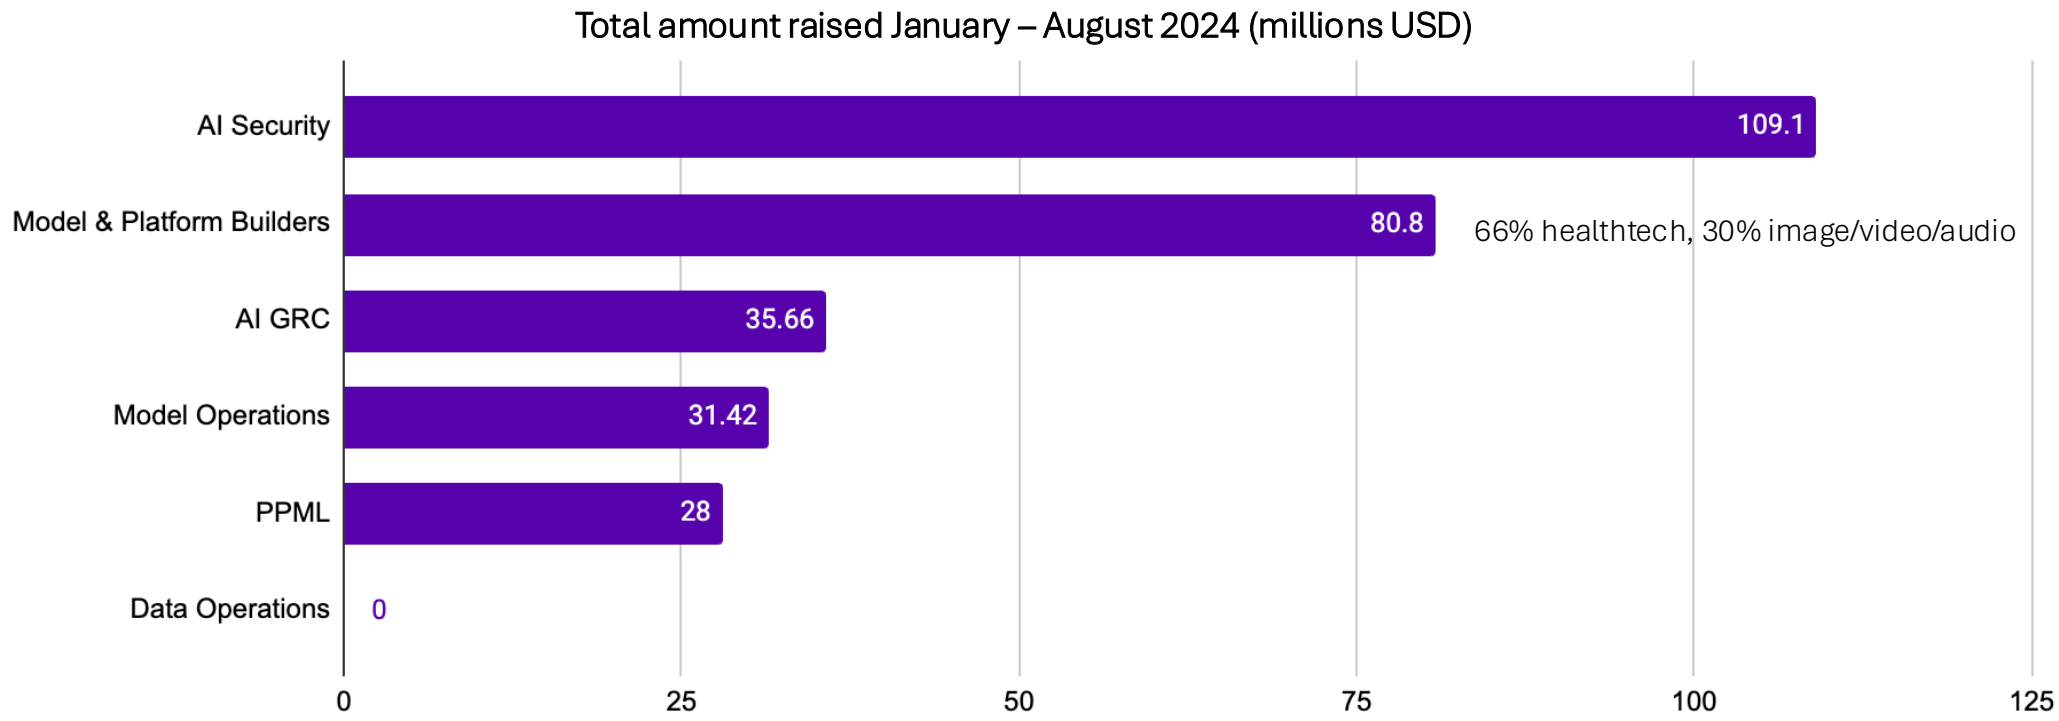

As an industry, the RAI space is being carried heavily by dollars invested in AI Security. These companies have easily demonstrable value propositions that most investors can understand. In other words, their value is rather well-defined and non-nebulous compared to that of other categories like AI GRC.

RAI Growth

AI Security is also the fastest growing category.

Not surprisingly, as a result of the inflow of funding, AI Security has the fastest growing headcount relative to the other categories. Some other categories, such as Model Operations (LLMOps, MLOps, evaluation, testing & benchmarking, etc.), seem underfunded and undervalued relative to the growth potential they are able to provide.

Headcount growth over last year, by category

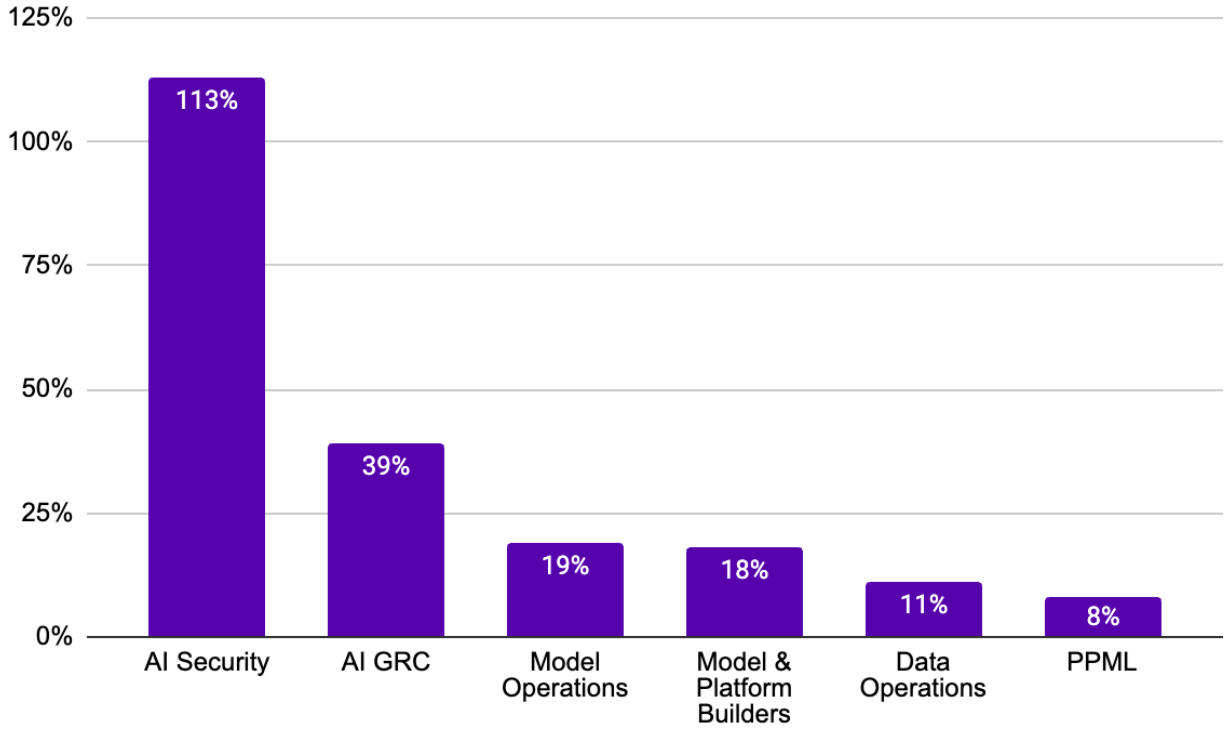

Headcount % growth per $1M invested, by category

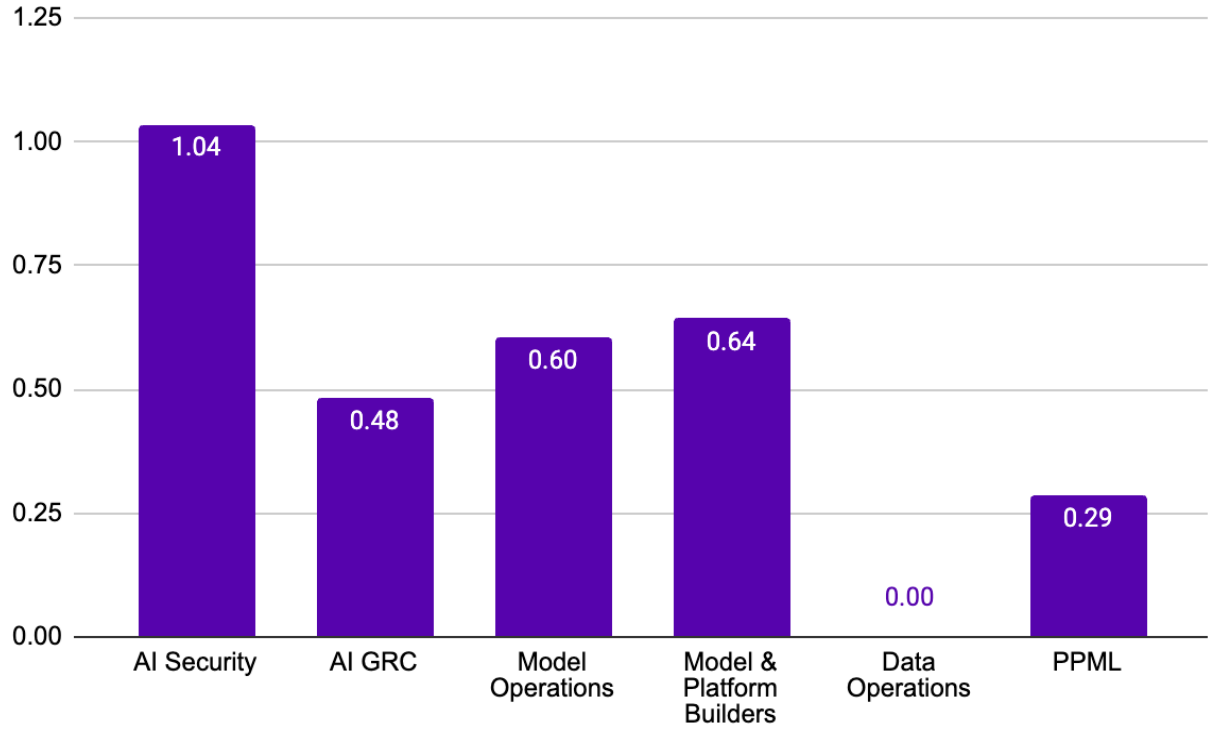

RAI Global Presence

The space still fundamentally lacks global participation from a solutions perspective.

As global AI policy continues to ramp and more countries invest in their growing infrastructure, EAIDB will hopefully see more involvement on the RAI side.

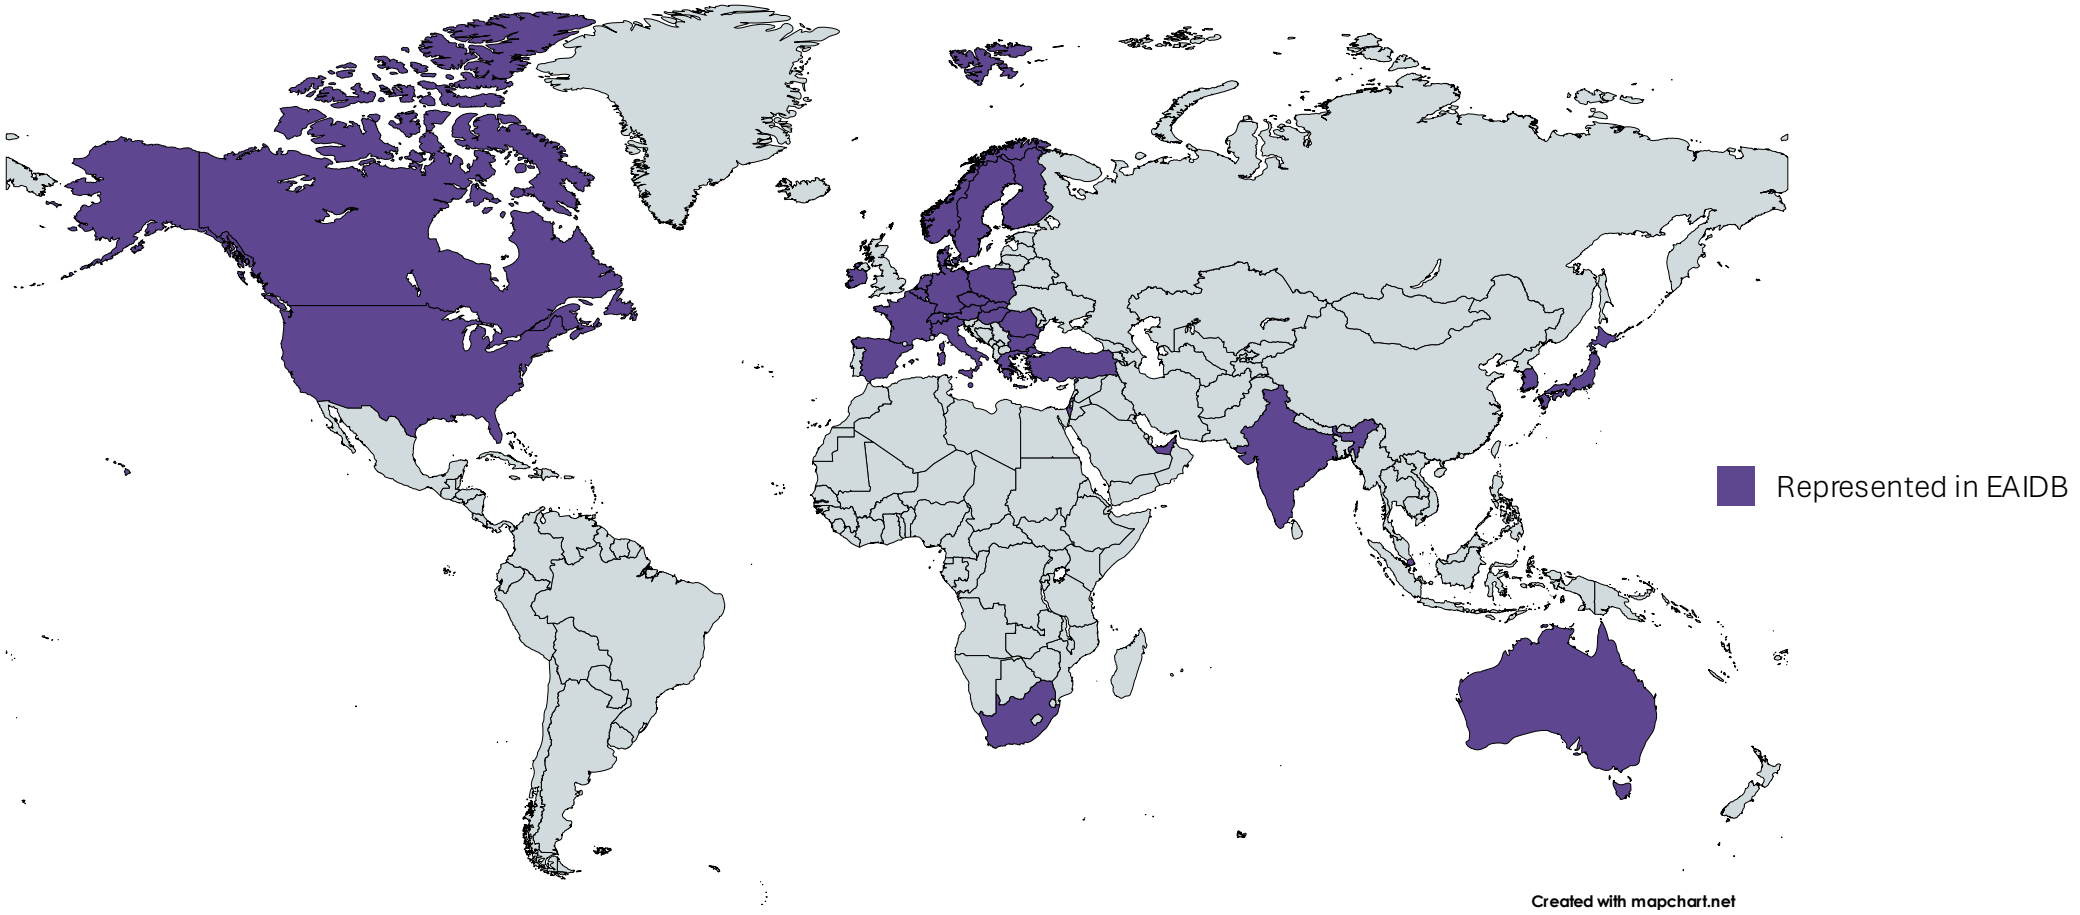

M&A Activity

Acquirer	Acquired	Description
Snowflake (NYSE: SNOW)	TruEra	Snowflake's acquisition of TruEra is part of an aggressive strategy to establish an end-to-end generative AI experience – from the data to the models. TruEra is a provider of LLMOps functions, including observability and evaluation tools as part of their open-source initiative, TruLens. Snowflake is also rumored to be in talks to acquire Reka AI, a multimodal model provider.
Cloudera	Verta	Like Snowflake, Cloudera's acquisition strategy is to expand their AI and GenAI capabilities for their enterprise clientele. Verta was an LLMOps and ModelOps company that focused specifically on building, operationalizing, monitoring, securing, and scaling models across the enterprise.
Aurionpro (NSE: AURIONPRO)	Arya AI	Aurionpro is an Indian corporation serving the banking, payments, mobility, and government sectors. Arya AI, as a provider of explainable AI solutions focused in the fintech space, allows Aurionpro to serve more consistent and modern solutions at-scale to the most populous country in the world.
Syllable AI	Actium Health	Both Syllable and Actium are AI companies in the healthcare space. Syllable is a HIPAA-compliant patient assistant that now leverages Actium Health's CENTARI platform, which predicts when a patient may need specific services.
Protect AI	Laiyer AI	The first of Protect's two acquisitions of the year, Laiyer AI offered GenAI security with their product LLM Guard. Protect now sells LLM Guard as part of their holistic security solutions covering machine learning and GenAI.
Protect AI	SydeLabs	Protect's second acquisition of SydeLabs allows them to enter the red-teaming space, which SydeLabs' Sydebox helped streamline. This will allow Protect to more proactively test systems and cross-sell their other products when vulnerabilities are discovered.
Cisco (NASDAQ: CSCO)	Robust Intelligence	Cisco acquires RI to deliver advanced security processing into their existing data flows by inserting it into Cisco network and security products. Robust Intelligence was a comprehensive AI red-teaming and evaluation platform that covered everything from bias/fairness to security.

EAIDB Company

AI Security

Industry Profile • Market Map • Market Leaders • Themes • Incumbent Effects

eaidb

Industry Profile: AI Security

+113%
avg. headcount change

$66m
est. raised in 1H2024
(+201% vs. 1H2023)

5
rounds raised in 1H2024
(+67% vs. 1H2023)

1
entrants in 2024
(-75% vs. 2023)

Ahead
vs. incumbents

Cumulative RAI AISec Startups in Market (est.)

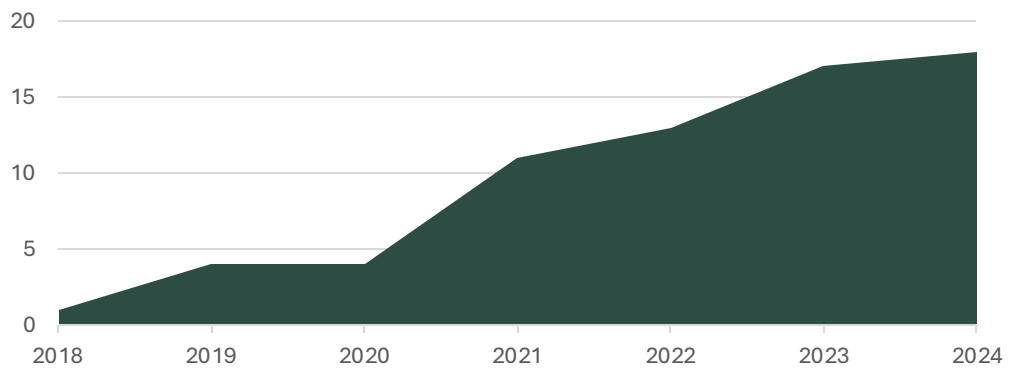

Market Incumbents

paloalto NETWORKS IBM Microsoft

Subcategory Breakdown
(% of companies with subcategory offerings, companies may have multiple)

- **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

eaidb

Market Leaders: AI Security

Scalers

High growth backed by strong technology.

Dynamo AI

TROJ.AI

CRANIUM

Leaders

Strong partnerships, acquisitions, and deep product coverage.

PROTECT AI

HIDDENLAYER

ROBUST INTELLIGENCE

CALYPSO

Entrants

New or emerging market competitors.

Invariantlabs

Innovators

Fascinating technology that has yet to achieve scale and status.

Deepkeep

vijil

PROTECT AI

- Acquired **Laiyer** and **SydeLabs** in 2024 to expand to LLM security.
- Raised a $50m Series B round.
- Rapid product expansion with releases of Sightline and Guardian.

HIDDENLAYER

- Estimated 209% headcount expansion in the last 1-2 years.
- New partnerships with **Microsoft Azure** AI for GenAI security.

ROBUST INTELLIGENCE

- Estimated $11.1M in revenue, +23% headcount.
- Recent partnerships with **Pinecone**, **MongoDB**.
- Unparalleled number of test suites for both security and model performance.

CALYPSO

- Estimated $6.9M in revenue +12% headcount.
- Partnerships with **IBM Watsonx**, **Palantir**, **Nexigen**.

Themes: AI Security

GenAI Security adoption may still be on the horizon.

GenAI is quite far from being production-grade in reliability and performance, even for internal applications. It will require leaps and bounds in the available technology for GenAI to enter the external sphere, where GenAI security is necessary. Most enterprises are not adequately convinced by GenAI to deploy it in internal settings, let alone external.

There have also been very few real attacks on external-facing GenAI systems. Historically, most of the damage from GenAI has come from hallucinations, bias, etc. - not malicious attacks.

How will enterprises secure agentic and/or more complex systems?

With agents on the horizon, security becomes much more complex as point solutions (guards for singular LLMs) are not enough for networks of LLMs that are actively executing commands, code, and interacting with real systems. There are very few startups working on this problem. The same paradigm extends to new methods of ML, like causal or neurosymbolic AI.

vijil AutoAlign PROTECT AI

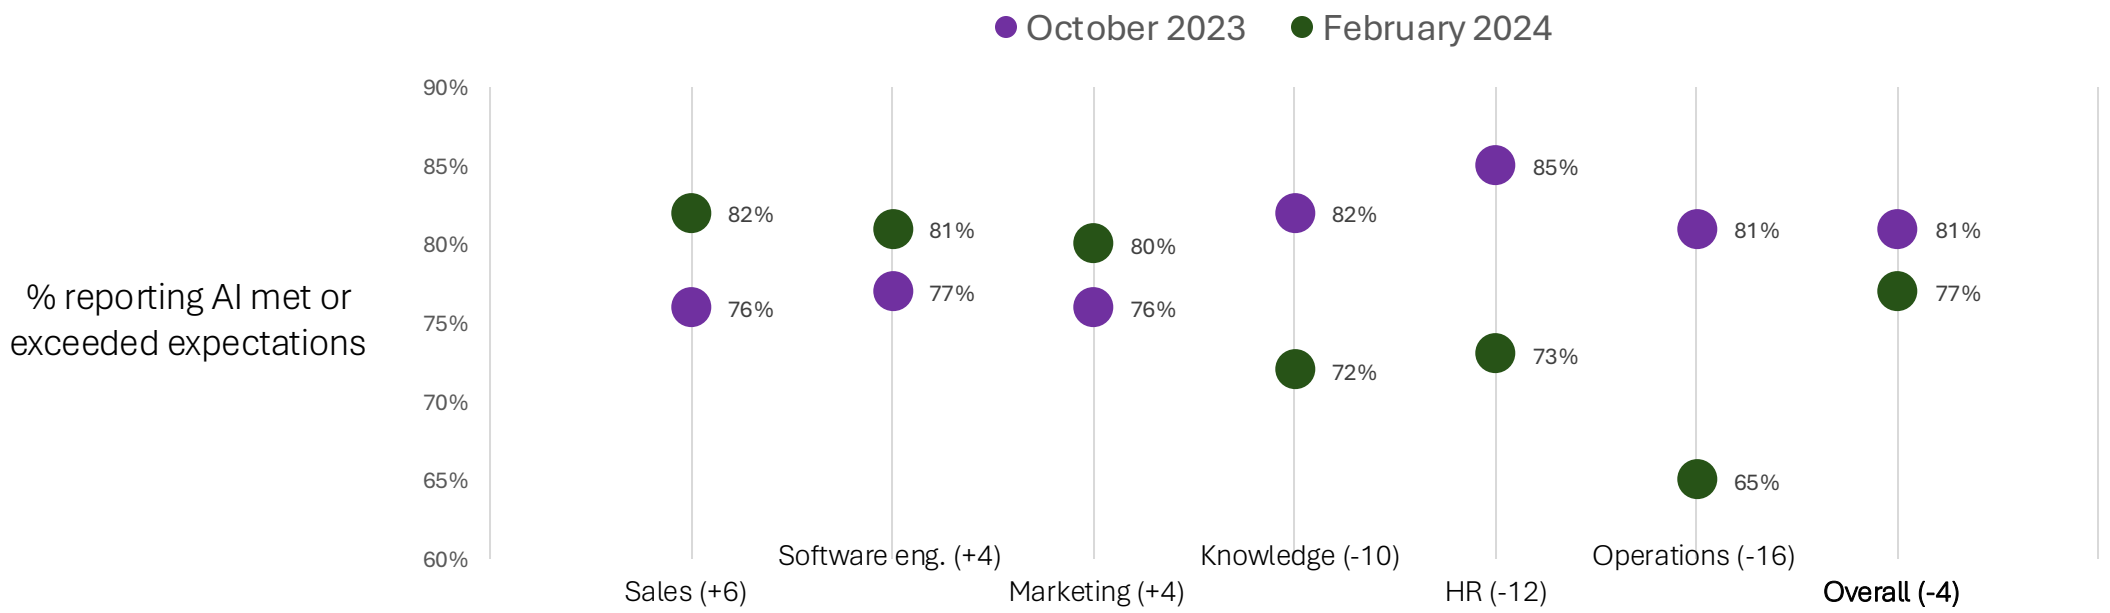

Incumbent Effects: AI Security

Cloud providers are not well-equipped to handle attacks on models of any kind.

There are few features that handle concepts like adversarial attacks, prompt injections, and membership inference. For GenAI, only **Microsoft Azure** has proactively invested in some kind of protection for production-grade systems with a strategic partnership with **HiddenLayer**.

The art of identifying and mitigating these attacks is not yet a science and is still a topic of much discussion within academic communities. It will take some time before such methods become widely adopted as part of AI Security outside of startup offerings.

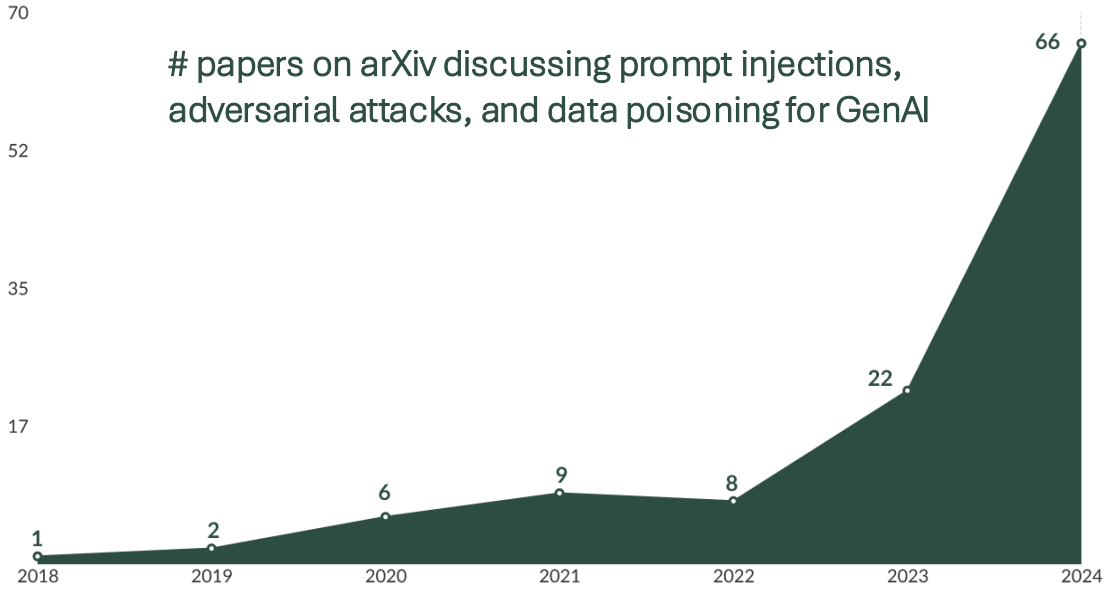

Palo Alto Networks (PAN) is the only large provider of all-encompassing AI security.

PAN recently rolled out their AI Runtime Security product, a firewall meant specifically for GenAI applications that protect against major attacks. This presents a particularly large presence in the AISec industry – one that is sure to garner most enterprise clients as **PAN** is a known entity within the security space.

However, given that this space is still in the research phase, startups like **Protect AI, Cranium**, and **HiddenLayer** are in a great position to rapidly iterate and pivot their key technologies. Speed and aggression may provide them the edge over heavier companies.

The other aspects missing from incumbents are comprehensive *security testing* and *security governance* across the enterprise. These are a key offering that almost all of the startups in **EAIDB** provide. For the former, performing penetration testing, prompt injections, etc. while a system is still pre-production is invaluable for today's AI-first developer. For the latter, being able to look across an enterprise's model portfolio and quickly assess attack surfaces will become increasingly important to a CISO moving forward.

Model Operations

Industry Profile • Market Map • Market Leaders • Themes • Consolidation

Industry Profile: Model Operations

+19%
avg. headcount change

$31m
est. raised in 1H2024
(-31% vs. 1H2023)

5
rounds raised in 1H2024
(-37% vs. 1H2023)

4
new entrants in 2024
(+0% vs. 2023)

Ahead
vs. incumbents

Cumulative RAI ModelOps Startups in Market (est.)

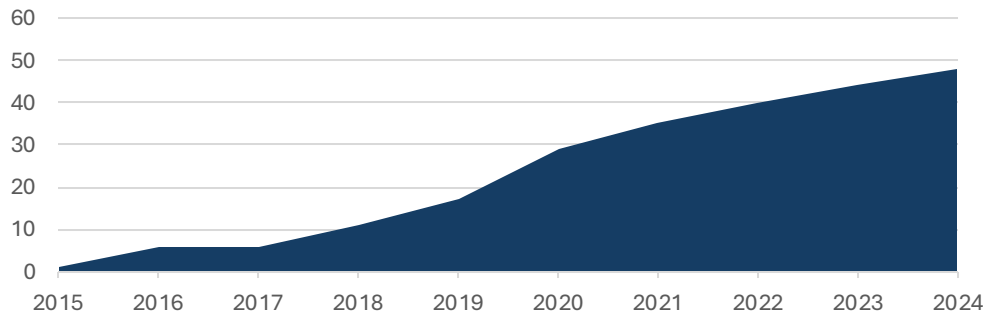

Market Incumbents

Google Cloud

databricks

Microsoft Azure

aws

Subcategory Breakdown
(% of companies with subcategory offerings, companies may have multiple)

Model Operations	
Model Testing & Eval. (81%)	Monitoring & QA (30%)

- 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 Operations

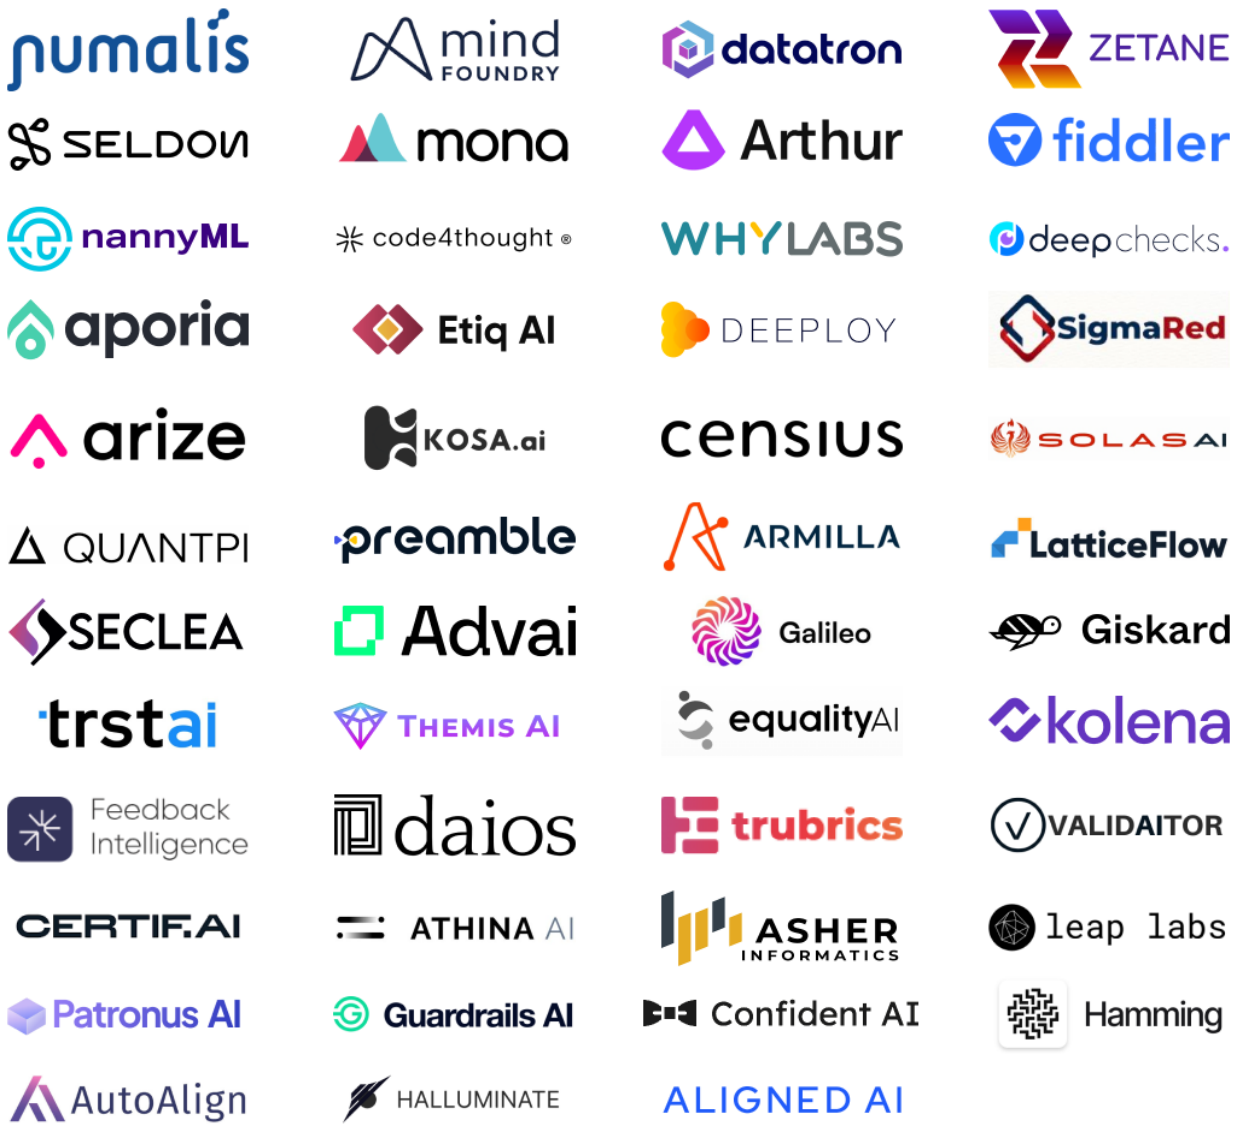

eaidb

Market Leaders: Model Operations

Scalers

High growth backed by strong technology.

Giskard

DEEPLOY

Galileo

deepchecks.

Leaders

Strong partnerships, acquisitions, and deep product coverage.

SELDON

arize

mind FOUNDRY

fiddler

Entrants

New or emerging market competitors.

HALLUMINATE

Hamming

Patronus AI

Innovators

Fascinating technology that has yet to achieve scale and status.

LatticeFlow

daios

Athina

Guardrails AI

arize

- Estimated 13% headcount growth, $13.6m in revenue.
- Partnerships with LlamaIndex, Guardrails AI.
- Diverse toolset includes hallucination detection and prompt variable monitoring.

SELDON

- Strong open-source presence.
- High-profile clients including Capital One, Ford, Deloitte, and Paypal.
- Deep MLOps presence links well with guardrails, monitoring, and auditing.

mind FOUNDRY

- Estimated 6% headcount growth, $12m revenue.
- Strong partnerships with government organizations in the UK, also support insurance use cases.

fiddler

- One of the best name-brands in the space.
- High-profile clients.
- Expansion in Asia via investment from Dentsu.
- Comprehensive LLMOps/MLOps frameworks.

Themes: Model Operations

Evaluation and benchmarking is still an unsolved problem.

Research from **Halluminate** indicates that engineers struggle to evaluate and benchmark LLMs. For such as fast-paced space, understanding how minor changes affect LLM outputs is a critical facet of the engineering workflow. Even more surprising is that most engineers are currently doing evaluations manually. Even cutting-edge methods today like "hallucination detection models" from **Galileo** or **Patronus AI** or "LLM-as-a-judge" approaches from **Confident AI**'s DeepEval or **Giskard AI** are limited because they rely on LLMs to judge LLMs. In high-stakes use cases, this will never be sufficient as LLMs are inherently unpredictable behavior. This problem is still a very active research area.

Alignment-as-a-service is taking its place as a mandatory part of the AI pipeline.

Small Language Models (SLMs) are quickly gaining momentum in enterprises because of their lower cost, faster inference times, and more opportunity to fit on proprietary data. To ensure that their performance meets or exceeds closed-source providers while retaining buy-in and trust from the business, SLMs must be fine-tuned and aligned with human preferences using an enterprise's own data and employees. However, there have been several studies that show that RLHF actually amplifies common human cognitive biases. This could be dangerous in many practical settings if LLMs are used as decision-makers.

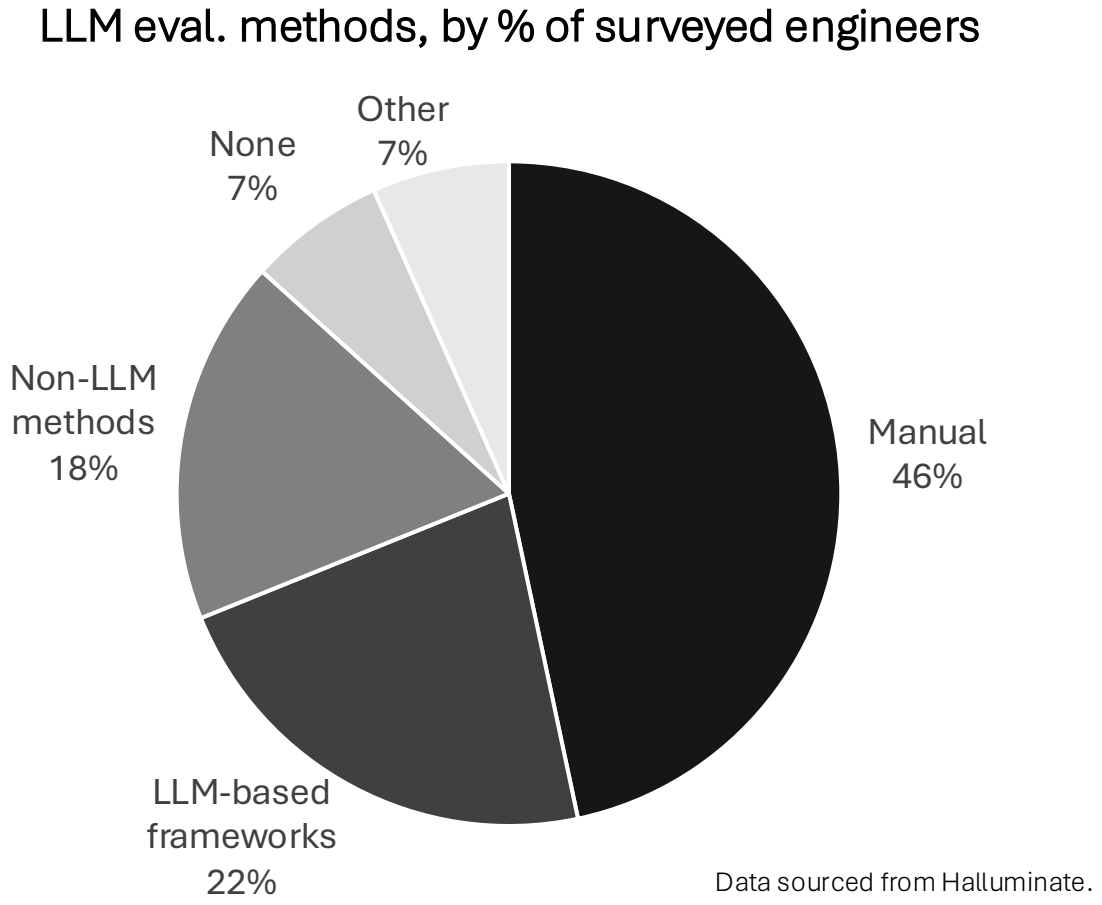

Data sourced from Halluminate.

Consolidation: Model Operations

Some of the largest data/ML providers are expanding their end-to-end GenAI offerings.

The two biggest acquisitions in the EAIDB model operations space happened this year, with **Cloudera** acquiring **Verta** and **Snowflake** acquiring **TruEra**. Both data giants are looking to enhance their coverage of the GenAI pipeline. **Snowflake** in particular has made several acquisitions to this effect (specifically, **Myst AI** and current talks for **Reka AI**).

Databricks, **Informatica**, and **Amazon** also boast a strong presence across the GenAI pipeline – but they fundamentally lack the observability tools of **Verta** and **TruEra**. If they follow in **Snowflake** and **Cloudera's** path, there may be comparable acquisitions in the future.

databricks

Lilac (data editing and tracking for LLMs)

mosaicML (GenAI research, custom LLMs)

OKERA (data governance)

Databricks has tools for data, models, and deployments, but lacks evaluation, benchmarking, and AI GRC.

SILO AI acq. by AMD

CognitiveScale THE TRUSTED AI COMPANY acq. by Tecnotree

Verta acq. by CLOUDERA

truera acq. by snowflake

AI GRC

Industry Profile • Market Map • Market Leaders • Themes • Incumbent Effects

Industry Profile: AI GRC

+39%
avg. headcount change

$35m
est. raised in 1H2024
(+338% vs. 1H2023)

7
rounds raised in 1H2024
(+75% vs. 1H2023)

0
new entrants in 2024
(vs. 7 in 2023)

Cumulative RAI AI GRC Startups in Market (est.)

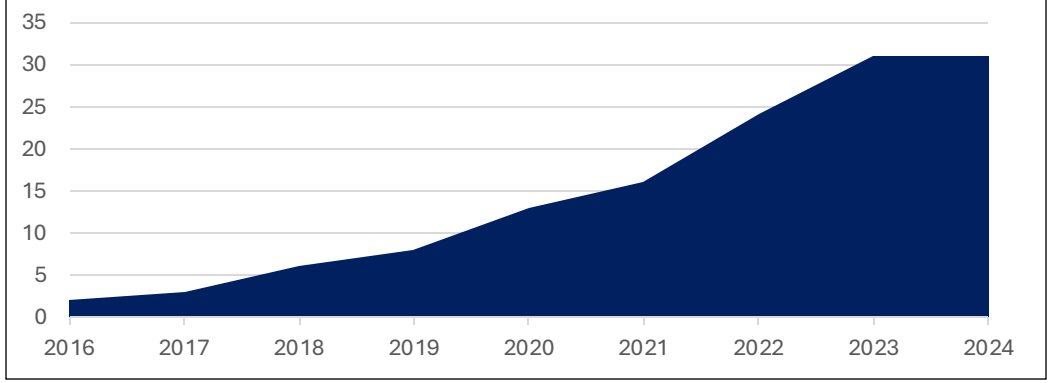

Market Incumbents

Vertex AI

watsonx

aws

Microsoft Azure

Subcategory Breakdown
(% of companies with subcategory offerings, companies may have multiple)

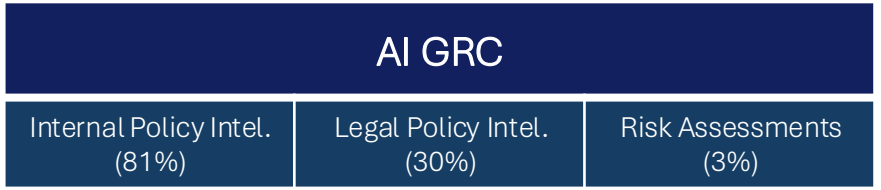

- **Internal Policy Intelligence:** solutions meant to track usage and implement guardrails according to an enterprise's internal definitions and policies.
- **Legal Policy Intelligence:** solutions meant to track projects and their adherence with common legal frameworks (ISO, EU AI Act, etc.).

AI GRC


credo | ai


ModelOp


Saidot


prodago


modulos


zupervise


MONITAUR


Guardrail
TECHNOLOGIES


FAIRLY


Holistic AI


LUMENOVA


enzai


Citrusx


MISSION CONTROL


KONFER

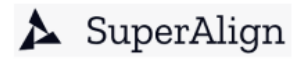
SuperAlign


calvinrisk


KomplyAi®


Fairo


BreezeML


TRUSTIBLE


fairnow


Daiki


trail


AI Guardian


Suzan AI


CertkOR AI


YOOI


anch.AI®


KTE


2021.AI

eaidb

Market Leaders: AI GRC

Scalers
High growth backed by strong technology.

enzai TRUSTIBLE

Holistic AI Saidot

Leaders
Strong partnerships, acquisitions, and deep product coverage.

credo ai MONITAUR

2021.AI ModelOp

Entrants
New or emerging market competitors.

Suzan AI fairnow

Innovators
Fascinating technology that has yet to achieve scale and status.

YOOI Daiki

credo ai

- Strong US presence, leaders are members of policy-making committees.

MONITAUR

- Strong partnerships and tight product for insurance and insurtech.

2021.AI

- Extremely strong presence in the Nordic region.
- One of the earliest companies to tackle AI GRC.

ModelOp

- Comprehensive MLOps and GRC solution.

Themes: AI GRC

AI GRC demand is constrained to mid-sized companies.

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.

Regulatory and compliance risks are still top of mind.

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.

Enterprise Priority	Small < $5bn	Medium $5bn - $25bn	Large > $25bn
#1	Data Warehousing & Data Lakes	AI Model Development	AI Model Development
#2	BI Analytics & Visualization	AI Governance, Risk & Compliance	Data Warehousing & Data Lakes
#3	Data Transformation	LLM Deployment	BI Analytics & Visualization

Data sourced from Insight Partners SoET Report 2024.

Incumbent Effects: AI GRC

Incumbents are not yet equipped to handle AI GRC.

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?

Feature	Google GCP	Microsoft Azure	Amazon AWS	IBM Watsonx	AI GRC Startups
Model Catalog, Model Cards	Y	Y	Y	Y	Y
Legal Policy Intelligence	Non-AI	Non-AI	Non-AI	Y	Y
Internal Policy Intelligence	Non-AI	Non-AI	Non-AI	N	Y
Automated pre-production risk/compliance	Some	Some	Some	Y	Y
Automated post-production risk/compliance	Some	Some	Some	Y	Y
Automated documentation and audit artifacts	N	N	N	N	Y

Inflow of AI GRC companies has been stifled by lack of adoption, dominance by first-movers.

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

Industry Profile: Data Operations

+11%
avg. headcount change

$200k
est. raised in 1H2024
(vs. $15m in 1H2023)

2
rounds raised in 1H2024
(+100% vs. 1H2023)

2
new entrants in 2024
(+0% in 2023)

Cumulative RAI Data Operations Startups in Market (est.)

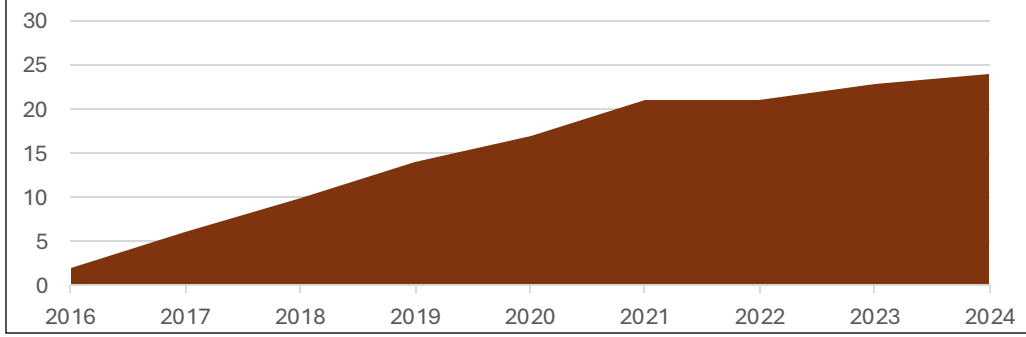

Market Incumbents

databricks TrustArc scale appen

Subcategory Breakdown
(% of companies with subcategory offerings, companies may have multiple)

Data Operations			
Data Governance (81%)	Data Repos. (3%)	Sourcing & Labeling (30%)	Data QA (3%)

- 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

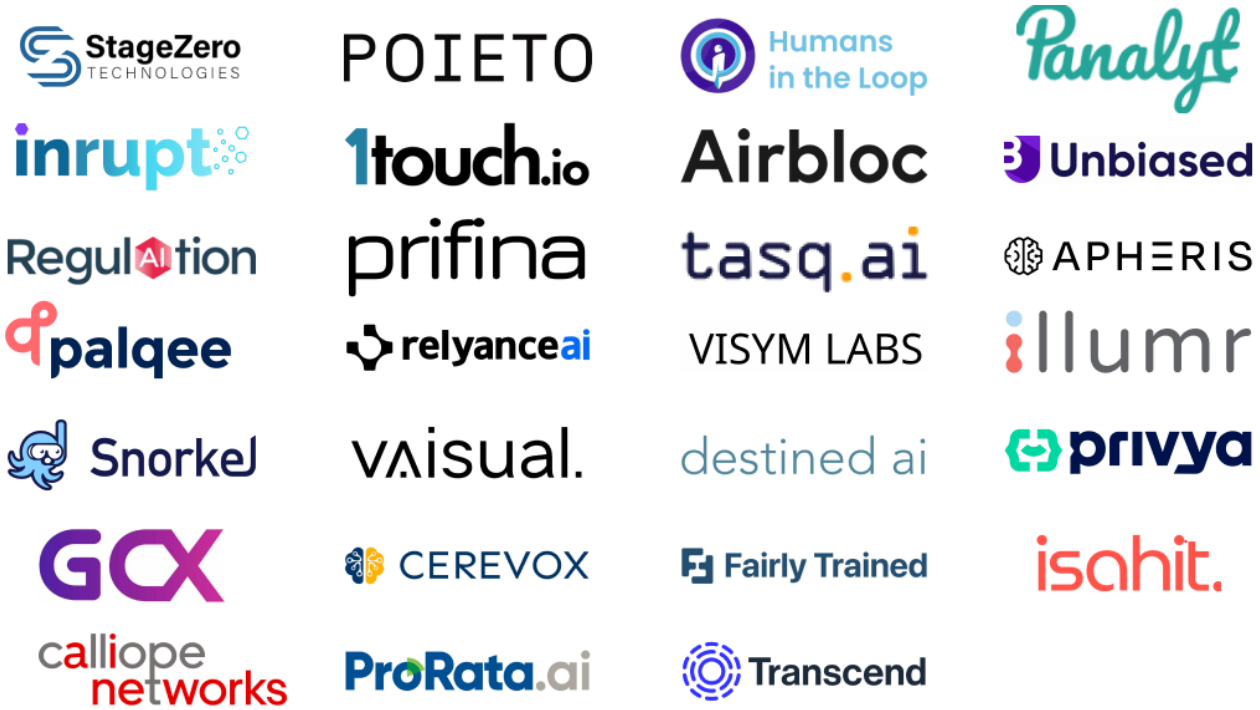

eaidb

Themes: Data Operations

Data licensing is both a solution and a problem.

Statistics from recent licensing deals have shown that **anyone with unique datasets is immediately at a monetizable advantage**. Shutterstock, Reddit, and many others have boosted their revenues by providing easy access to datasets for GenAI use.

At the same time, there are many organizations that would rather lock their data down through the use of "non-AI" licenses that strictly forbid training or fine-tuning. The question of how to unlock this data beyond simple RAG purposes is a difficult one to answer.

Company	Description	Est. Revenue
Reddit	Content licensing to unnamed firms (possibly Google or OpenAI).	$203m
Shutterstock	Multiple deals with OpenAI, Meta, Google, and Amazon.	$25-50m per contract
Reuters	Transactional content licensing for GenAI.	$22m

Data sourced from public filings and media.

Data annotation is a near-universal need.

Across every enterprise, human-validated ground truth data is a necessity. Doing this with as little friction as possible for the business is a tricky task, but it is what companies like **Scale AI** and **Isahit** have perfected. Synthetic data has also seen a boost from this data-greedy market, but models constructed solely from synthetic data generally suffer collapse.

While **Scale AI** is not part of the EAIDB (due to a questionable past with training the GPT models), their recent financial performance is a proxy for the demand for good data.

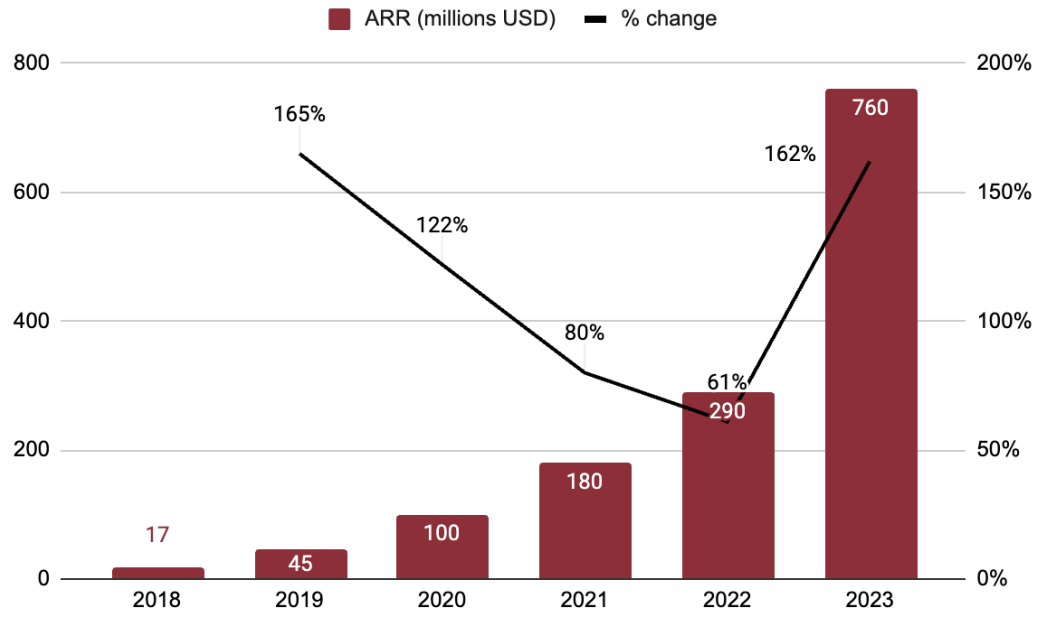

Privacy Preservation

Industry Profile • Market Map • Market Leaders • Themes

Industry Profile: Privacy Preservation

+8%
avg. headcount change

$28m
est. raised in 1H2024
(-84% vs. 1H2023)

2
rounds raised in 1H2024
(-75% vs. 1H2023)

0
new entrants in 2024
(+0% vs. 2023)

Ahead
vs. incumbents

Cumulative RAI Privacy Preservation Startups in Market (est.)

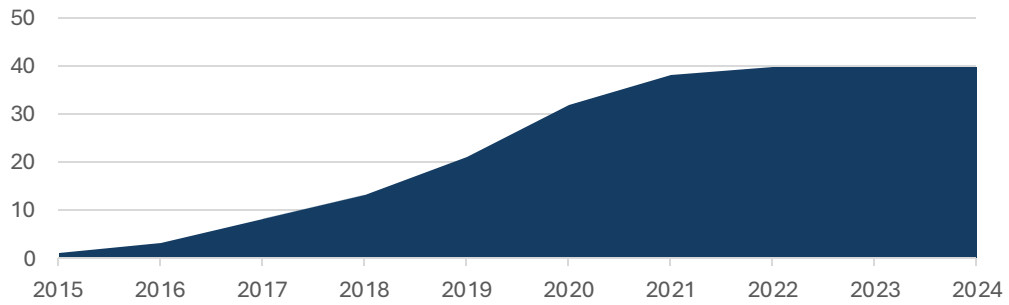

Market Incumbents

Informatica IBM ORACLE

Subcategory Breakdown
(% of companies with subcategory offerings, companies may have multiple)

Privacy Preservation			
Consent Mgmt. (5%)	Synthetic Data (70%)	Federated Ops. (13%)	Anonymization (15%)

- 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.

Privacy Preservation

eaidb

Market Leaders: Privacy Preservation

Scalers
High growth backed by strong technology.

Flower PRIVATEAI HOWSO inpher

Leaders
Strong partnerships, acquisitions, and deep product coverage.

MOSTLY AI MDCLONE TONIC THE FAKE DATA COMPANY gretel

Entrants
New or emerging market competitors.

BlueGen

Innovators
Fascinating technology that has yet to achieve scale and status.

RHINO HEALTH Bitfount

MOSTLY AI
- Strong presence in Europe (**Telefonica**, **City of Vienna**).
- Entering the text space after perfecting product offerings on structured data.
- Integrations with **Databricks**, **AWS**.

MDCLONE
- One of the go-to names in healthcare across all data systems.
- Late-stage with very high-profile, cross-domain clients.

TONIC THE FAKE DATA COMPANY
- Comprehensive privacy offerings, including PII detection for leakage through LLMs.
- Partnerships with **Snowflake**, high-profile clients in **Etsy, CVS, Oscar.**

gretel
- 51% headcount growth in last 1-2 years.
- Expansive product offerings covering data anonymization/privacy to synthetic data.
- Newest product, Gretel Navigator, generates tabular data using LLMs.

Themes: Privacy Preservation

The barrier-to-entry for synthetic data is incredibly high.

In the first six months of 2024, we've seen an unprecedented number of synthetic data companies forced out of business. Datagen, Mirry AI, Syntric AI, and Syntegra all went under due to technological and market challenges:

1. Enterprises require certain standards and certifications for vendors to work with their data. Newer startups are at an immediate disadvantage here vs. Incumbents, especially with a lack of track record.
2. LLMs can now generate synthetic data via chat. Even if these LLMs will never approach the quality of startups, it becomes much harder for enterprises to defend buying instead of building.

For these reasons, the first-movers and the more mature startups in the space are already miles ahead.

Federated operations providers are gaining momentum as a salve for high-stakes industries.

Significant investment tailwinds for high-stakes sectors like healthcare underscores the need for efficient and safe algorithmic and mathematical computations. Federated Operations (which extends federated learning to include any kind of computation) is the sector that stands to gain. While the market is still very young and unforgiving, estimates of enterprise usage by 2032 hover at 20% (up from 5% in a 2022 estimate) (Market US).

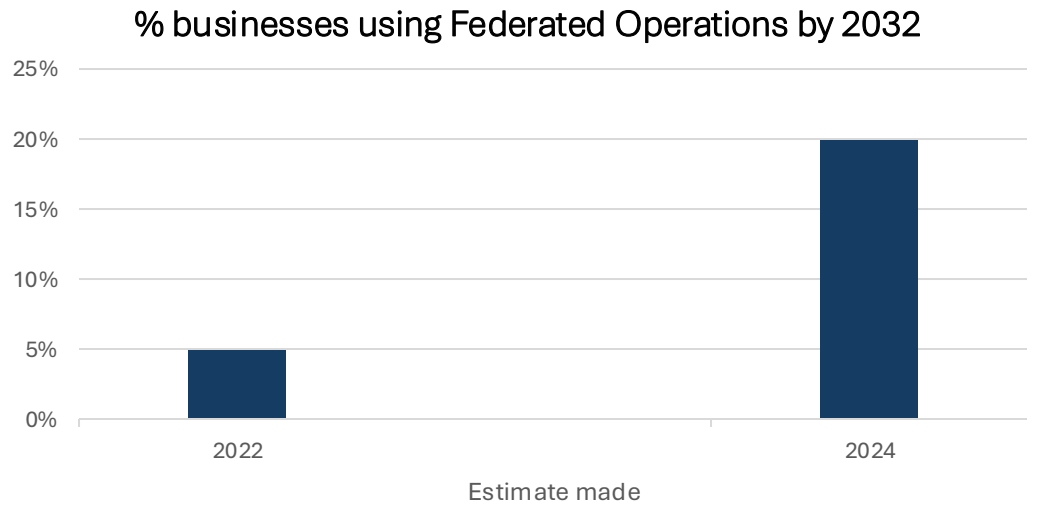

Data sourced from Market.US.

Model & Platform Builders

Industry Profile • Market Map • Content Moderation & Deepfakes • Healthtech & Fintech

Industry Profile: Model & Platform Builders

+18%
avg. headcount change

$107m
est. raised in 1H2024
(-41% vs. 1H2023)

7
rounds raised in 1H2024
(-86% vs. 1H2023)

1
new entrants in 2024
(vs. 3 in 2023)

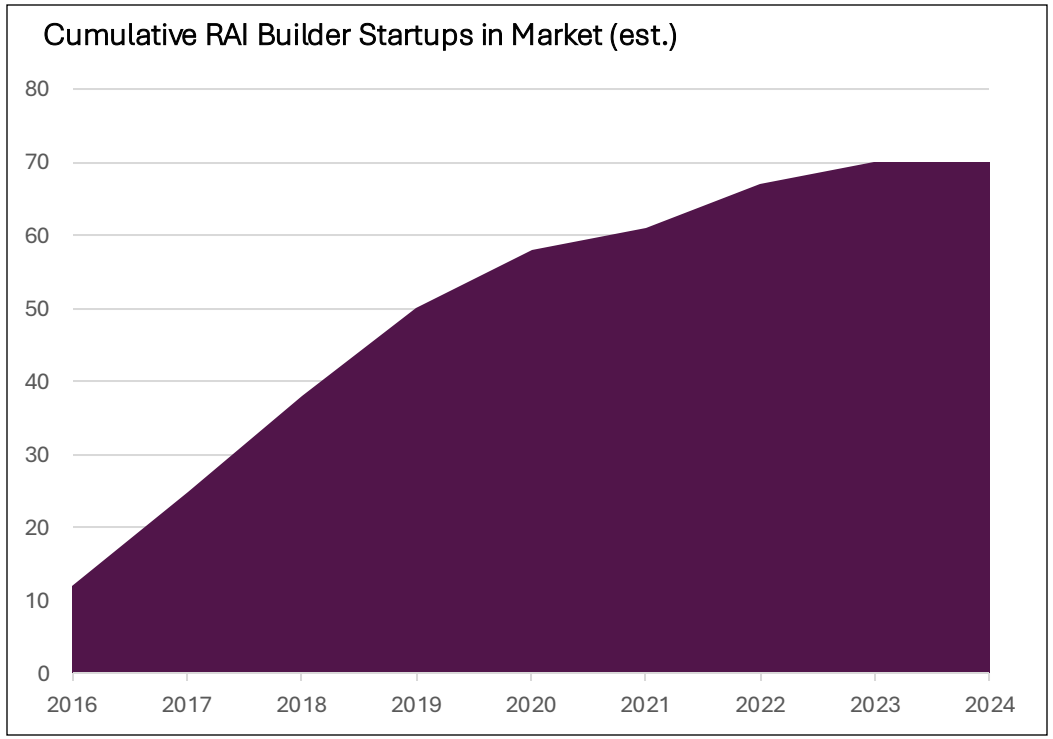

Subcategory Breakdown
(% of companies with subcategory offerings, companies may have multiple)

Model & Platform Builders			
HRTech (39%)	Fintech/Insurtech (7%)	Image, Video, Audio (12%)	Model Research (11%)

- 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

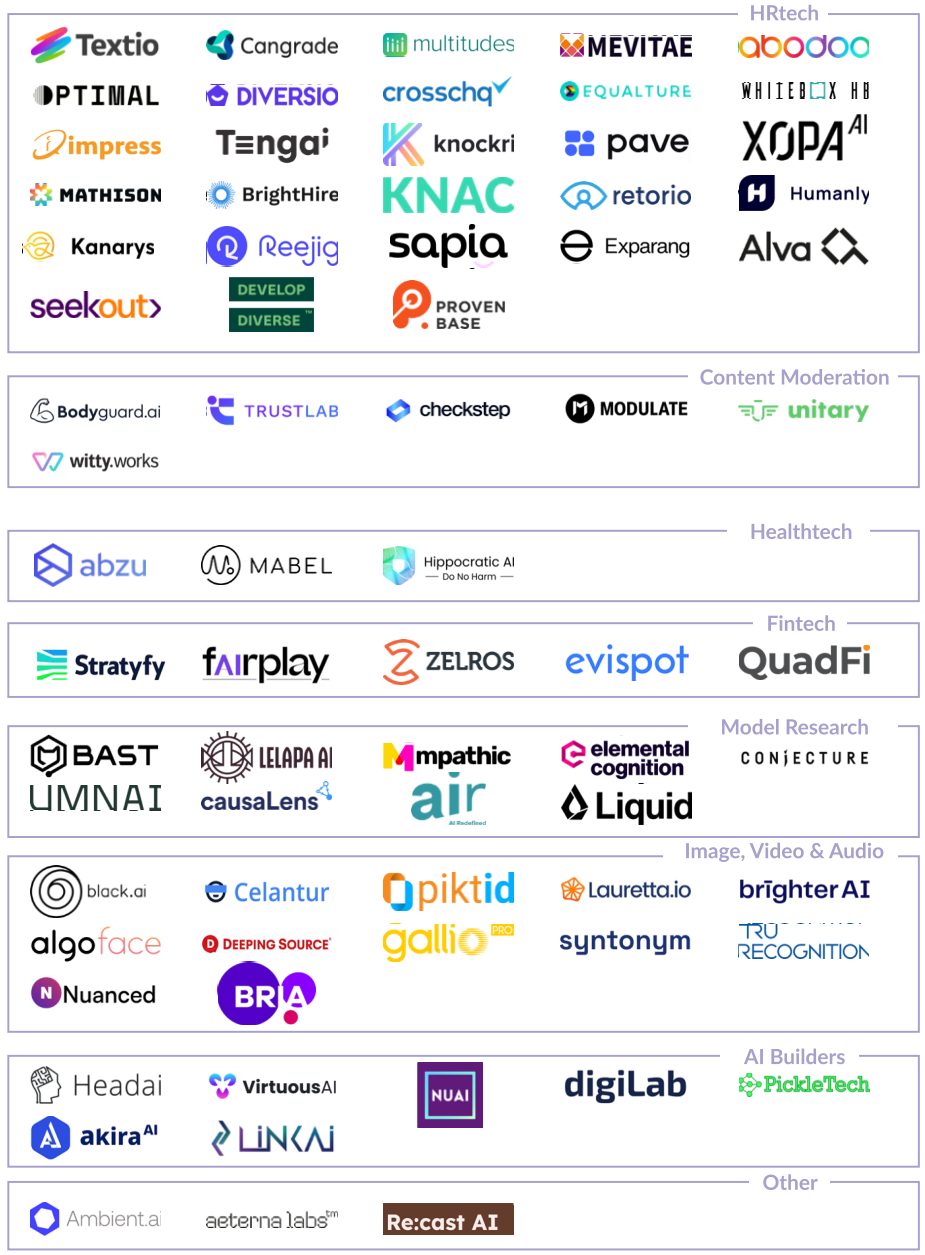

eaidb

Themes: Content Moderation & Deepfakes

The most common misuse of generative AI is deepfakes.

According to a study by Google DeepMind, the prevalence of deepfakes (impersonations, falsifications, etc.) dwarfs that of other security concerns like adversarial attacks and privacy compromises. Startups like **Nuanced** and **Certifi AI** exist to counter these efforts. This space is receiving a lot of attention as newer models like Black Forest Labs' FLUX.1 can be used to create intricate deepfakes. As the United States approaches yet another divisive presidential campaign, deepfake detection will become a critical aspect of "fact checking" procedure.

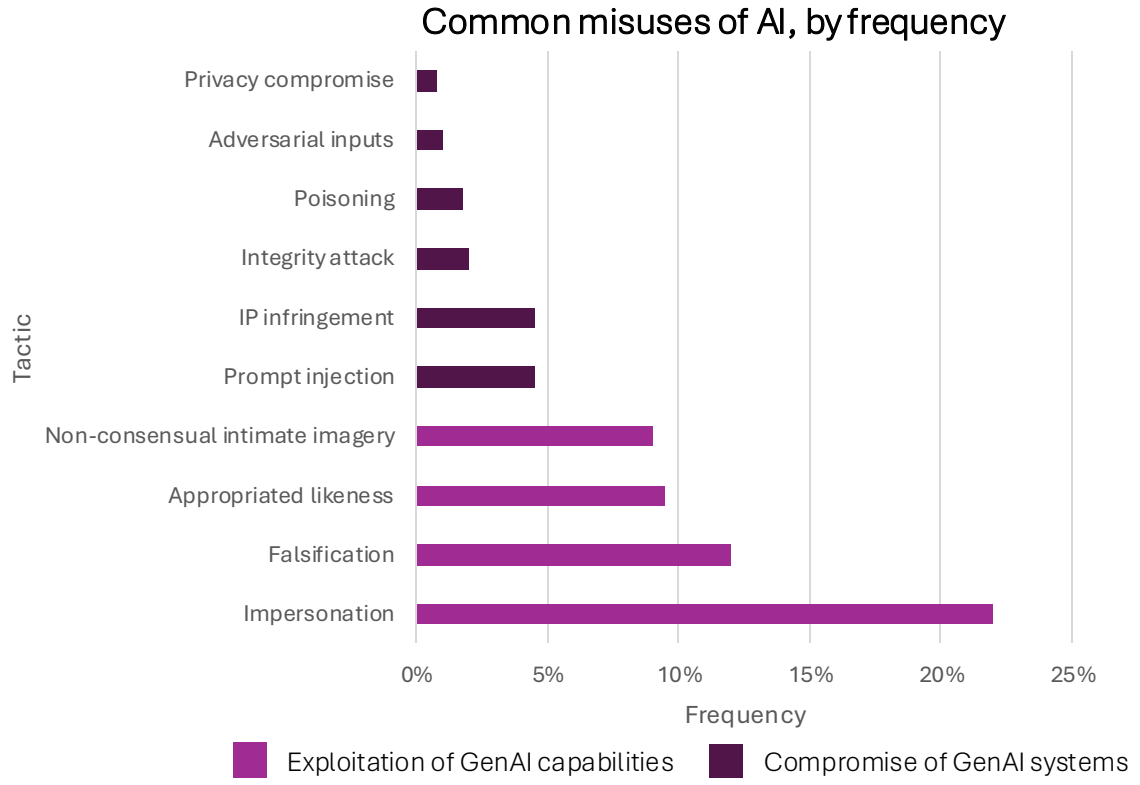

Nuanced

Certifi AI™

Reality Defender

DeepMedia.AI
Pioneering Synthetic Generation and Detection

Themes: Healthtech & Fintech

Healthtech is becoming an incredibly profitable frontier for Responsible AI enablers.

The healthcare space is receiving an extraordinary amount of funding across the board, no doubt spurred onwards by the potential of AI. However, the industry's tight regulatory and compliance standards has proven a challenge for most AI companies. Startups offering responsible AI enablement can unlock and ease friction for others to deploy meaningful solutions. According to SVB's AI in Healthcare report, patient diagnostics in particular is a treacherous area for AI since it requires safe, trustworthy AI that also yields high accuracies.

Similarly, fair finance is gaining momentum among large financial institutions as well.

Another heavily regulated space, finance (including lending, banking, etc.) is also benefiting from responsible AI enablement. Fair lending and explainable AI solutions like **FairPlay AI** and **Stratyfy** are finding success in an otherwise opaque industry where regulators often fight battles to gain some sort of clarity into the algorithmic dealings of financial institutions.

FairPlay recently partnered with **LendingPoint** (which has issued more than $2.3bn in loans worldwide) to bring their fairer lens to the platform. **Stratyfy** raised a $10m round last year and was named as the LendTech Startup of the Year.

US VC Dollars invested in Healthcare companies leveraging AI

(billions USD)

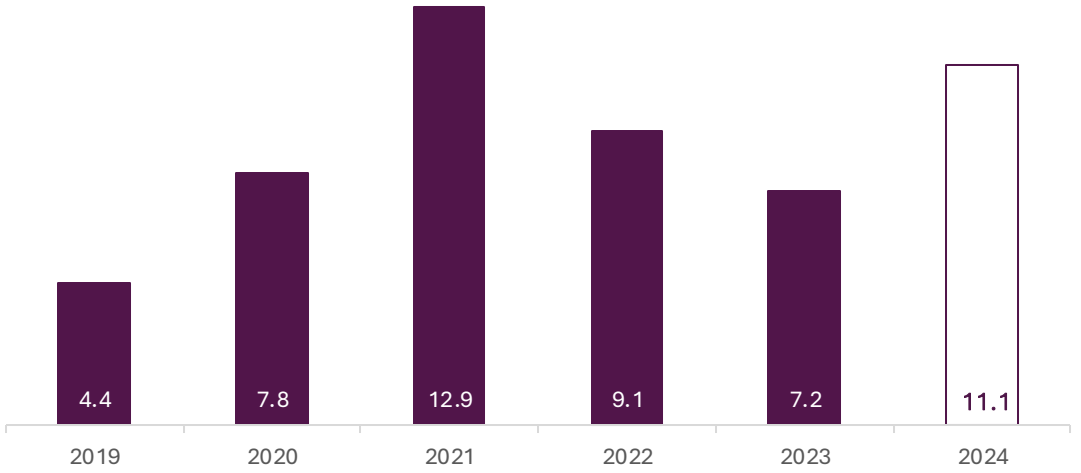

discrimination / fair lending cases referred to the US Dept. of Justice

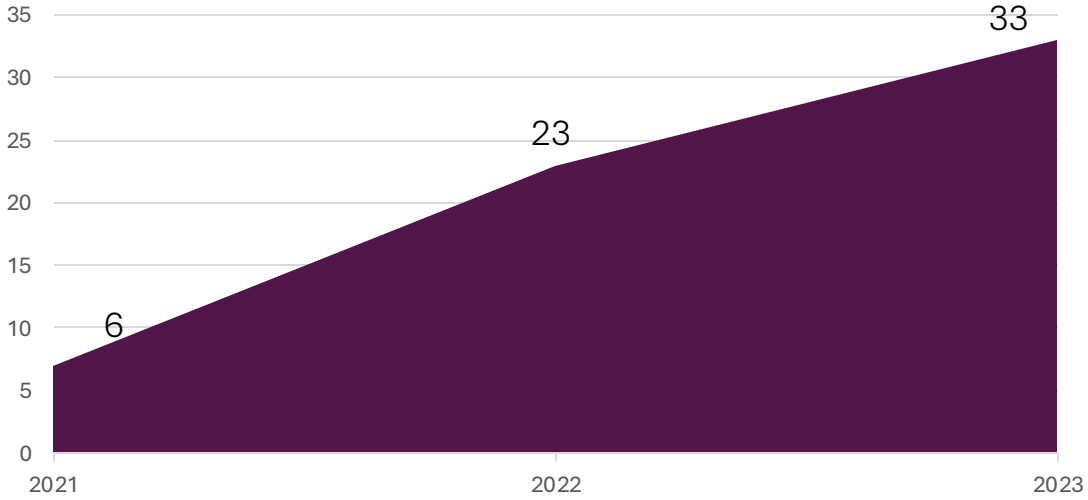

Data sourced from US CFPB.

Themes: Model Research

Alternative LLM types are coming, and they could be significantly better-performing.

The downsides of traditional LLMs are clear – they hallucinate, they're extremely finicky and hard to experiment with, they're immutable and unexplainable, and they consume a lot of cost, compute, and person-hours. To combat this, RAI startups are inventing new architectures, developing new foundation models, and experimenting with completely new technologies.

ALIGNED AI	Alignment Platform	Self-correcting models with time, decreased bias. Catches model degradation and adversarial inputs in real-time.
ProRata.ai	Attribution Technology	Provides content contributors with compensation if their content is used within the confines of an LLM. Attempts to solve the content attribution problem in the GenAI age.
Liquid	Liquid LLMs	Implementing Liquid Neural Networks (LNNs) as transformers and potentially a foundation model. These networks have causality, dynamic, real-time adaptation to new data, and many other features.
CONJECTURE	Cognitive Emulation	Mimicking the human process of cognition and reasoning with LLMs as the tool. Different from agents in the way they do this. Inherently more explainable and editable by the humans involved.
BAST	Bast AI	Nature-inspired AI systems designed to be explainable and environmentally responsible. Cost and compute efficient.
LELAPA AI	Vulavula, InkubaLM	LLMs designed for African languages, by Africans for Africans.
mpathic	Empathetic AI	Conversational LLMs designed to optimize responses for empathy according to behavioral signals from the user. Specifically for healthcare and patient contexts.
elemental cognition UMNAI	Neurosymbolic AI	Neurosymbolic AI allows complete interpretability and certainty with how a system will behave. Also allows dynamic weight editing (which means data can be completely removed from the model with no residuals).
causaLens	Causal AI	Endows LLMs with the ability to reason causally. Significantly enhances their abilities and boosts explainability.
air AI Redefined	Cogment	AI learns dynamically and interactively with a human in a shared environment. Requires less data and yields more trust.

Consulting & Legal

Industry Profile • Market Map • Themes

Industry Profile: Consulting & Legal

Cumulative RAI Consulting & Legal Firms in Market (est.)

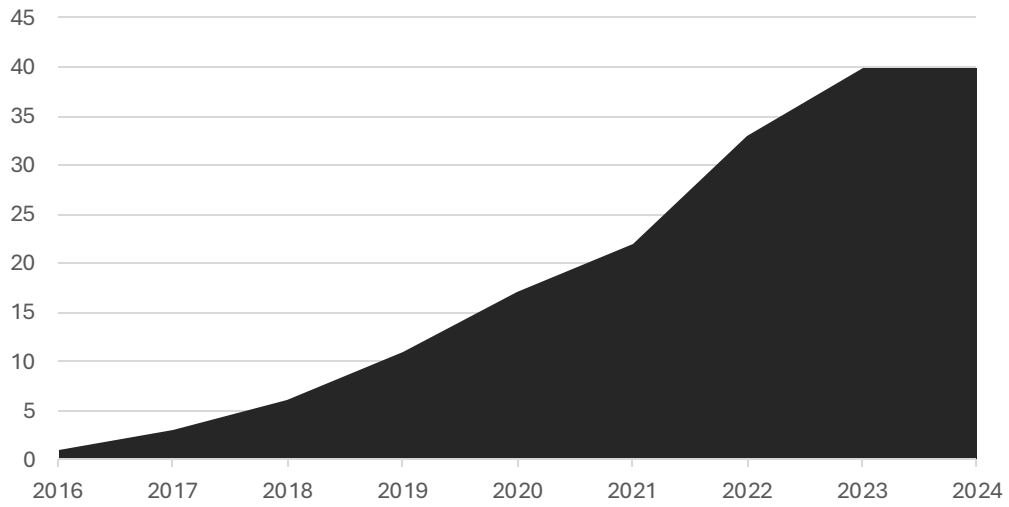

Market Incumbents

accenture BCG KPMG McKinsey & Company Deloitte.

Subcategory Breakdown

(% of companies with subcategory offerings, companies may have multiple)

Consulting & Legal			
AI Strategy (52%)	Data Strategy (14%)	Legal (7%)	Algorithmic & Risk Audits (33%)

- **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

eaidb

Themes: Consulting & Legal

An increasing number of smaller firms have begun offering AI GRC solutions.

While not necessarily SaaS (though some, like **Anekanta Consulting**, are building), these solutions provide additional perspective and a case-by-case analysis of the various risks within an AI system. Firms also offer upskilling and training, especially with respect to Generative AI.

This space is also incredibly competitive as most large consulting firms (**Deloitte**, **Accenture**, **KPMG**, etc.) have near unlimited resources for the same offerings. It is important to note, however, that the methods of Big 4 firms in Responsible AI are made to be generalizable across enterprises – **their services typically fall short for most SMEs.**

Firm	Est. Impact from GenAI Consulting
KPMG	$650m (NYTimes)
McKinsey	400 generative AI initiatives in six months (Business Insider)
Accenture	$900m (CIODive)

New entrants are few and far between in 2024.

The space saw a dramatic reduction in new consulting and legal firm entrants in 2024, down to zero from 11 in 2022 to 7 in 2023. This could be for a few reasons:

1. **Talent shortage**: there is a fundamental lack of AI expertise from a technical level. Smaller firms have to upskill themselves first before offering services externally.
2. **Demand constraints:** while enterprise business is siphoned towards larger firms, SMEs are still considering the value of GenAI and may be too early stage to need specialized consulting services for it.

IT specialists on their biggest challenges for implementing AI

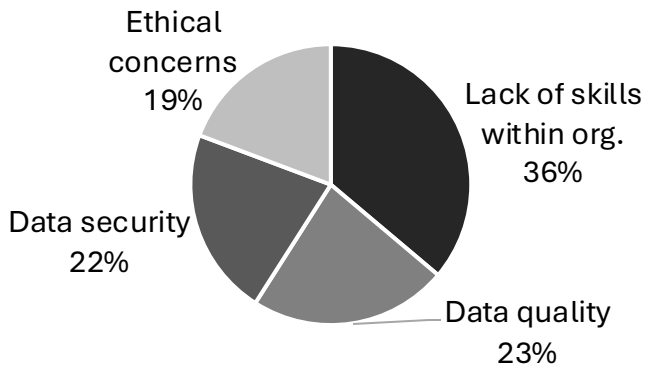

Data sourced from Salesforce.

References

General

Growjo headcount and revenue data. https://growjo.com.

Overview

Crunchbase News "Healthcare And AI Sectors Continue To Lead, While Funding Settles In July 2024". https://news.crunchbase.com/venture/monthly-funding-recap-july-2024.

AI Incident Database. https://incidentdatabase.ai/.

AI Security

Bain & Company "Quarterly Survey on AI". https://www.bain.com/insights/ai-survey-four-themes-emerging/.

Model Operations

Halluminate AI research. https://halluminate.ai/.

AI GRC

Insight Partners "State of Enterprise Tech 2024". https://info.insightpartners.com/state-of-enterprise-tech-2024.html.

Data Operations

Sacra "Scale AI". https://sacra.com/c/scale-ai/.

Privacy Preservation

Market.US "Global Federated Learning Market". https://market.us/report/federated-learning-market/.

Model & Platform Builders

Silicon Valley Bank "The AI-Powered Healthcare Experience". https://www.svb.com/globalassets/trendsandinsights/reports/svb-the-ai-powered-healthcare-experience-2024.pdf.

Consumer Financial Protection Bureau (CFPB) reports. https://www.federalregister.gov/agencies/consumer-financial-protection-bureau.

Google DeepMind "Mapping the misuse of generative AI". https://deepmind.google/discover/blog/mapping-the-misuse-of-generative-ai/.

Consulting & Legal

Salesforce AI Survey. https://www.salesforce.com/news/stories/public-sector-ai-statistics/.

eaidb

Appendix

Resources • About EAIDB

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 **4,000+ followers** on LinkedIn and has received **6,500+ 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