AI, in board terms.

The AI terms that turn up in proposals, results calls and regulators' letters, explained for the people who have to decide. Under each term: the policy, register, process or report to ask management for.

Each entry explains the term as far as a decision needs. The "Ask" names a document, not a fact: a fact can be answered from memory in a meeting; a document can be read, kept and audited. Figures come from company reporting or official documents, and the AI ROI newsletter covers the cases behind them.

KNOW THE DIFFERENCE

Pairs that sound alike in a proposal and mean different money, or different risk.

Agent vs automation

Automation follows a script: the same steps, in the same order, every time, and it stops when something unexpected appears. An agent is given a goal and works out the path, solving problems as it meets them. That is what makes agents useful, and what makes them harder to predict, test and control.

SEE AI agent, Orchestration

Deterministic vs probabilistic

Ordinary software is deterministic: the same input gives the same output, every time, and a test proves it. A model is probabilistic: it produces the most likely answer, and the same question can get a different answer tomorrow. That is why AI systems are tested on hundreds of cases rather than one, why they are given limits rather than instructions, and why "it worked in the demo" proves little.

SEE How a model works, Evaluation

Model vs harness

The model is the engine, rented from a supplier and replaced every few months. The harness is everything the company builds around it to do a job: instructions, access to systems, memory, checks. The harness holds the company's workflow rules and its own intellectual property; the value and the lock-in sit there.

SEE Model, Harness

Copilot vs agent

A copilot drafts and a person decides. An agent decides and acts. The saving is larger with the agent, and so is the damage it can do. The company is liable either way.

SEE AI agent, Human in the loop

© 2026 Agent & Capital Ltd

For general information only. Not legal, financial or investment advice. See the full statement on the last page.

Pilot vs production

A pilot answers whether it can work. Production answers what it costs and earns at full volume. A figure from one is not evidence of the other.

SEE Pilot and production

Productivity gain vs saving

A productivity gain is time. A saving is money. One becomes the other only by a decision.

SEE Productivity gain, Save to invest

Run-rate saving vs one-off charge

The run-rate saving is what a change delivers every year once it is in place. The one-off charge is what it cost once to get there. A restructuring needs both numbers, and the running cost of the technology as well.

SEE Run-rate saving, One-off charge

AI capex vs AI opex

Buying computing and building models is capital spending, with the approval and the write-off risk that go with it. Renting models per use is a running cost that grows with success. Which one a plan is decides who approves it and how it is judged.

SEE Compute, Usage-based pricing, Capitalising AI development

Training vs inference

Training builds the model: rare, expensive and mostly the supplier's cost. Inference uses it: every request, the company's cost, and it grows with every successful rollout.

SEE Training and inference

Open-weight vs proprietary

Open-weight models can be run on the company's own infrastructure; proprietary ones can only be rented. The trade is control against capability.

SEE Open-weight and proprietary models

Fine-tuning vs retrieval

Fine-tuning changes the model with the company's own examples and has to be repeated when the model changes. Retrieval leaves the model alone and lets it look up the company's documents. Retrieval is cheaper, easier to keep current and enough for most tasks, so it is the one to rule out first.

SEE Fine-tuning, Retrieval-augmented generation

Generative AI vs machine learning

Generative AI produces new content: text, code, images. Machine learning predicts or classifies from patterns in data. Many of the best-returning uses are the second, and the cheaper.

SEE Generative AI, Machine learning

© 2026 Agent & Capital Ltd

For general information only. Not legal, financial or investment advice. See the full statement on the last page.

Hallucination vs error

An error is a wrong answer that can be detected. A hallucination is a confident, plausible invention. The second is built into how language models work, and needs grounding in verified data and review where it matters.

SEE Hallucination

Benchmark vs evaluation

A benchmark is the supplier's public test, run on the supplier's chosen tasks. An evaluation is the company's own test on the company's own work. Only the second predicts what the system will do in production.

SEE Benchmark, Evaluation

Guardrails vs human in the loop

Guardrails are written limits on what an agent may do; they scale. A human in the loop is a person approving before an action takes effect; it does not scale, and it is the control most often removed to capture a saving.

SEE Guardrails, Human in the loop

AI literacy vs AI training

AI literacy is a legal duty under the EU AI Act: the company must ensure the people who use AI understand it. AI training is the programme that meets the duty. The first is the board's obligation; the second is management's delivery.

SEE AI literacy

Data residency vs data sovereignty

Residency is where the data sits. Sovereignty is whose law governs it and who can be compelled to hand it over. A US supplier's data centre in Frankfurt answers the first and not the second: under the US CLOUD Act, US authorities can compel a US company to hand over data wherever it is stored.

SEE Data residency, Sovereign AI

Automation vs augmentation

Automation means AI does the task. Augmentation means AI helps the person doing it. The distinction decides where the saving comes from, who must be consulted, and which risks apply.

SEE Productivity gain, Jobs attributed to AI

THE BASICS

Artificial intelligence (AI)

Software that performs tasks which used to need human judgement: reading, writing, classifying, predicting, deciding. The term covers everything from a spam filter to an autonomous agent, so on its own it says little about cost or risk.

ASK FOR the AI inventory: every AI system in use or planned, what each one does, which business process it sits in and who owns it. A company that cannot produce the list cannot govern it.

© 2026 Agent & Capital Ltd

For general information only. Not legal, financial or investment advice. See the full statement on the last page.

Machine learning

The method behind most AI in use today. Instead of being programmed with rules, the system learns patterns from data and applies them to new cases. Its quality depends on the data it learned from.

ASK FOR the data description behind each model in production: what it was trained on, when, and how closely that data resembles the company's own.

Generative AI

AI that produces new content, such as text, code, images or audio, rather than classifying or predicting. ChatGPT, Claude and Gemini are generative AI. It is the technology behind most of the AI investment announced since 2023.

ASK FOR the use-case list split by type: which cases need content generated, and which need only a prediction or a decision. The second group is usually cheaper and better proven.

Prompt

The instruction given to a model. Prompts can be a single question or several pages of rules and examples, and in many systems they are the main control a company has over the model's behaviour. They are also part of the company's intellectual property.

ASK FOR the prompt management process: where production prompts are stored, who may change them, how changes are tested and approved, and how they are protected as company property.

Token

The unit in which language models read and write, roughly three-quarters of a word. It matters because it is also the unit in which usage is priced. Costs scale with the number of tokens in and out, so long documents and chatty agents cost more.

ASK FOR the business case standard: every AI case to show token cost at production volume, not pilot volume, over three years.

How a model works

A model does not store words, pictures or facts. Everything it reads is cut into chunks, tokens for text and patches for images, and each becomes a long list of numbers recording how it relates to every other chunk. What the model "knows" is the pattern in those numbers, learned from its training data, and what it produces is the most probable continuation, not a retrieved fact. Four things follow. It can state something false with full confidence. It can answer the same question differently on different days. It cannot be opened and read. And the company's own facts enter it only through retrieval or fine-tuning, which is why those two choices matter.

ASK FOR the statement, for each AI system, of where its facts come from (training, retrieval or fine-tuning) and the evaluation that shows how consistent its answers are.

Context window

How much text a model can hold in mind at once, measured in tokens. A larger window lets the model work on whole contracts or data sets, at a higher cost per request.

ASK FOR the technical fit assessment for each use case: whether the work fits the model's limits, or needs engineering to work around them, and what that costs.

© 2026 Agent & Capital Ltd

For general information only. Not legal, financial or investment advice. See the full statement on the last page.

Hallucination

A confident answer that is wrong or invented: a citation that does not exist, a number that was never reported. It is a property of how language models work, not a bug that will be fixed. The remedy is grounding the model in verified data and checking outputs where it matters.

ASK FOR the output review policy: which AI outputs reach a customer, a regulator or the accounts without a person checking them, and the evidence that justified removing the check.

Multimodal

A model that works across text, images, audio and video, not only text. It lets AI read a scanned invoice, a photograph of a shelf or a recording of a call.

ASK FOR the use-case list to show which processes involve documents, images or voice, and whether those have been assessed for AI.

THE MODELS

Large language model (LLM)

The engine of generative AI: a very large model trained on text that predicts the next word, and through that can draft, summarise, translate, answer questions and write code. A handful of suppliers build the leading ones; most companies rent access.

ASK FOR the model supplier register: which models the company depends on, under which contracts, with which terms on data, price and termination, and the plan for switching.

Model

The trained system itself, the asset that does the work. A model has a version, a release date and a known set of strengths and failures, like any product. Models are replaced often, so anything built on one needs a plan for the next.

ASK FOR the model lifecycle process: how a new model version is tested, approved and rolled into production, and how results are checked after the change.

Training and inference

Training is building the model from data; it is expensive and done rarely, mostly by suppliers. Inference is using the model to answer a request; it happens every time and is what a company pays for in use. Most of a company's AI running cost is inference.

ASK FOR the AI cost report: inference spend by use case and by month, with the trend, as a standing line in the management accounts.

Fine-tuning

Adjusting a trained model with a company's own examples so it performs a specific task in a specific style. More effort and cost than prompting or RAG, and it has to be repeated when the underlying model changes.

ASK FOR the design decision record: why fine-tuning was chosen over prompting or retrieval, what it costs to repeat on the next model version, and who approved it.

© 2026 Agent & Capital Ltd

For general information only. Not legal, financial or investment advice. See the full statement on the last page.

Retrieval-augmented generation (RAG)

A technique that has the model look up the company's own documents before it answers, so its answers rest on the company's facts rather than its general training. The usual way to put a model to work on internal knowledge without retraining it.

ASK FOR the knowledge source register: which documents and systems each AI system retrieves from, who keeps them current, and who may add to them.

Open-weight and proprietary models

Proprietary models (OpenAI, Anthropic, Google) are rented through an interface and cannot be inspected. Open-weight models (Mistral, Meta's Llama, DeepSeek) can be downloaded and run on a company's own infrastructure. Open-weight gives control and often lower cost; proprietary usually gives higher capability with less effort.

ASK FOR the model selection policy: which workloads need the best available model and which run on a cheaper or self-hosted one, with the criteria for deciding.

Frontier model

The most capable models available at a given time, built by a few companies at very high cost. New frontier models arrive several times a year, and each one changes what is possible and what is cheap.

ASK FOR the model refresh plan: how quickly the company can adopt a better model, and what in its current systems would stop it.

Reasoning models

Models that spend more computing on each answer to work through a problem step by step before replying. Better on hard tasks such as analysis and planning, and several times the cost per answer.

ASK FOR the model routing rules: which tasks are sent to a reasoning model and which to a cheaper one, and who sets that.

Small language models

Smaller, cheaper models that run on ordinary servers, laptops or phones. Often enough for a well-defined task, and they keep the data on the company's own equipment.

ASK FOR the model selection policy to cover small models: which tasks were tested on one before a frontier model was chosen.

AGENTS

AI agent

Software that uses a model to pursue a goal across several steps, calling tools, reading systems and acting, rather than answering a single question. An agent can process a claim, check a customer's identity, restock a shelf or shop for a flight. Agents are where AI starts to replace tasks, and where the risks of acting replace the risks of answering.

ASK FOR the agent register: every agent in production, the systems and data each may touch, the actions it may take without a person, the limits written into its brief, and the named owner who approved it. Then ask who keeps it current.

© 2026 Agent & Capital Ltd

For general information only. Not legal, financial or investment advice. See the full statement on the last page.

Harness (also "scaffolding")

A model on its own only answers questions. To do a job it needs to be fitted with instructions, access to the company's systems, a memory of the task so far and checks on its work. That fitting is the harness. The model is bought from a supplier and changes every few months; the harness is built by the company and stays. It is where the real work, and the value, sits. Think of the model as the engine and the harness as the rest of the car.

ASK FOR the ownership statement for each harness: who built it, who owns the code and the prompts, and the evidence that it works with more than one model.

Agentic AI

The general term for systems built around agents. It is the phrase most often used in investor presentations and job-cut announcements, so it is worth asking what it means in each case.

ASK FOR the task inventory behind any agentic plan: which tasks are handed to agents, in which process, with the saving in money and the quarter it shows.

Tool use

The ability of a model to call other software: search, a database, a payment system, an email account. Tools turn a model that can talk into an agent that can act, and each tool is a permission the company has granted.

ASK FOR the tool permission list, per agent: which systems each may call, read or write, and who approved each permission. It belongs in the agent register.

Human in the loop

A design in which a person reviews or approves an agent's work before it takes effect. It is the simplest control over agent risk and the one most often removed to capture savings.

ASK FOR the approval policy: which agent actions need a person's sign-off, which do not, and the evidence reviewed before any sign-off was removed.

Guardrails

Rules and checks that limit what a model or agent may say or do: topics it must refuse, actions it must not take, limits on amounts. Written limits work: in a UK government test, writing the limits of a task into an agent's brief cut unauthorised actions by about 85%. They do not end the risk.

ASK FOR the written limits of each agent, the test results showing they hold, and the review schedule.

Orchestration

Software that coordinates several agents and tools into one workflow, deciding what runs when. As AI spreads, orchestration becomes the layer where process logic lives, and therefore where control sits.

ASK FOR the architecture description: where the orchestration layer sits, whether the company or a supplier owns it, and what happens to the processes if the supplier is replaced.

© 2026 Agent & Capital Ltd

For general information only. Not legal, financial or investment advice. See the full statement on the last page.

Agent identity and permissions

An agent needs its own credentials, separate from any person's, with the least access its task requires and a record of what it did. Agents running on a person's login, or with administrator rights, are the common failure.

ASK FOR the access policy for agents: separate identities, least privilege, regular review, and the list of agents with elevated rights.

Audit trail

The record of what an AI system saw, decided and did, kept so that an incident can be investigated and a decision explained. Without it there is no investigation and no answer for a regulator.

ASK FOR the logging standard: what is recorded for each AI system, for how long, and who can retrieve it within a day.

Kill switch

The ability to stop an agent, a model or an integration at once, without waiting for a supplier. The basic control for the day something goes wrong.

ASK FOR the shutdown procedure: who can stop each agent, how, how long it takes, and when it was last tested.

Shopping agent

An AI agent that searches, compares and buys on a customer's behalf. For any company that sells, it is a new intermediary between the company and its customer, comparing on criteria the company may not see. Delta Air Lines keeps such agents away from its prices; Zalando lets them in on its own terms.

ASK FOR the agent channel policy: which outside agents may search, price or buy from the company, on what terms, who decided, and how agent traffic is measured.

Agentic commerce

Buying and selling done between agents rather than people: the customer's agent searches and pays, the seller's systems respond. The payment and identity rails are being built now by Google, the card networks and the large platforms.

ASK FOR the commerce readiness assessment: whether the company's products, prices and stock can be read and bought by an agent, and what it would take.

Multi-agent systems

Several agents working on one task: one plans, others execute, one checks, and they hand work to each other without a person in between. The pattern behind most ambitious agent plans. Each hand-off is a place where an error or an injected instruction can pass unchecked, and the chain can run for a long time before anyone looks.

ASK FOR the interaction design for any multi-agent system: which agent may instruct which, the limit on steps and spend before a person is involved, and the log of the whole chain.

© 2026 Agent & Capital Ltd

For general information only. Not legal, financial or investment advice. See the full statement on the last page.

Model Context Protocol (MCP)

An open standard for connecting models and agents to tools and data, so that a connector built once works with many models. It is becoming the plumbing of agent systems, and like all plumbing it carries security consequences.

ASK FOR the connector inventory: which connectors are exposed to agents, which data each reaches, and the security review for each.

MONEY

AI ROI

The return on an AI investment, measured the same way as any other: value delivered against total cost over time. It is rarely reported that way. Companies more often report inputs (hours saved, employees using AI) than outputs in money.

ASK FOR the AI value report: by use case, the money returned this year against the full cost, on the same page, as a standing item in the management accounts.

Use case

A specific task that AI performs in the business: checking customer data, restocking shelves, drafting code. Returns are made at the level of use cases, not of strategies, so a plan should name them and attach a number to each.

ASK FOR the use-case portfolio: each case with its owner, status (pilot or production), cost and the saving or revenue attached.

Pilot and production

A pilot proves that something works; production is where it earns. Most AI results reported publicly are from pilots, where volumes are small and costs are hidden. The gap between the two is where returns are lost.

ASK FOR the stage-gate process: the criteria for moving a pilot to production, and the record of how long each case took and how many never made it.

Run-rate saving

The saving a change delivers each year once it is fully in place, the figure a CFO uses to judge a restructuring. It is different from the saving in the first year, which is usually smaller, and from the one-off cost of getting there.

ASK FOR the savings schedule: run-rate saving, the quarter it is reached, and the savings in each year until then, by use case.

One-off charge

The cost booked once for a change: redundancy payments, written-off systems, consultants. When job cuts are attributed to AI, the one-off charge is the first half of the cost; the technology and the people who run it are the second.

ASK FOR the full cost statement of the change: the one-off charge, the technology spend and the people who build and supervise the agents, against the saving.

© 2026 Agent & Capital Ltd

For general information only. Not legal, financial or investment advice. See the full statement on the last page.

Total cost of ownership

Everything an AI system costs over its life: licences, tokens, computing, data preparation, integration, the people who build and supervise it, and the cost of replacing it when the model changes. Business cases that show only the licence understate it several times over.

ASK FOR the business case template to require all of these lines over three years, and for the finance function to sign each case off against it.

Usage-based pricing

Paying per token, per request or per seat rather than a fixed fee. It makes AI cheap to start and hard to forecast at scale, and it means every successful rollout raises the bill.

ASK FOR the spend control process: budgets per use case, alerts when usage runs ahead, and the forecast at ten times today's volume.

Compute

The processing power, mainly specialised chips called GPUs, that AI runs on. Compute is the scarce input of the AI economy and the reason for the data centre boom. A company buying AI services is paying for someone else's compute; a company running its own is making a capital decision.

ASK FOR the compute strategy: what is rented and what is owned, the utilisation that justifies owning, and the approval route for compute capital.

Model routing

Sending each request to the cheapest model that can handle it, rather than everything to the most expensive. Once usage grows, it is the main lever on inference cost.

ASK FOR the routing policy and the cost report that shows its effect.

Vendor lock-in

Dependence on one supplier's model, platform or format, so that switching would cost more than staying. The models themselves are increasingly interchangeable; the lock-in now sits in integrations, prompts and data. Zalando says it can swap models at any time; what it cannot swap, it owns.

ASK FOR the exit assessment per supplier: what it would take, in time and money, to move, and the contract terms that help or hinder it.

Build, buy or rent

The three ways to get AI capability: build it with your own engineers, buy a product with AI inside, or rent a model and build on it. Most companies do all three. The question is which capabilities are worth owning.

ASK FOR the sourcing policy: which AI capabilities are treated as strategic and built or owned, which are bought, and the reasoning.

Productivity gain

An improvement in output per hour or per person, the most common AI figure companies report. It is not a saving until someone decides what to do with the time: fewer people, more output, or nothing. The decision is where the return is made.

ASK FOR the benefits realisation plan: for each productivity gain, the decision on what the freed time funds, and where that shows in the accounts.

© 2026 Agent & Capital Ltd

For general information only. Not legal, financial or investment advice. See the full statement on the last page.

Save to invest

A programme in which cost savings are deliberately reinvested in the business rather than taken to profit. Tesco runs one and names it so. The point for a board is that the destination of a saving is a choice, and it should be made on purpose.

ASK FOR the allocation policy: how AI savings are split between profit and reinvestment, decided in advance, and the report that shows where this year's went.

Capitalising AI development

Whether the cost of building AI systems, the harness above all, is booked as an asset and written off over years, or expensed as it is spent. The accounting standards (IAS 38, and ASC 350-40 in the US) allow capitalisation of development that meets their tests. The AI question is the write-off period: a harness built on a model that is replaced in a year may not be an asset for five.

ASK FOR the capitalisation policy for AI development: what is capitalised, over what period, and how that period is reconciled to how often the underlying models change.

Concentration risk

Dependence across the whole company, and across the economy, on two or three model suppliers and three cloud providers. A single supplier's outage, price change or policy change then reaches every AI system at once. Central banks have started to name it as a systemic risk.

ASK FOR the concentration report: AI workloads by model supplier and cloud provider, the share each carries, and the plan for the day the largest one fails or doubles its price.

AI washing

Overstating the role of AI in products, results or plans. The US Securities and Exchange Commission has fined companies for it. A board that signs accounts and investor presentations has a direct interest in the claims being true.

ASK FOR the disclosure review process: who checks AI claims in results, presentations and marketing against the evidence before they are published.

DATA

Training data

The data a model learned from. For supplier models it is mostly the public internet, plus whatever the supplier licensed; for a company's own models it is the company's records. Its quality sets the model's ceiling, and its provenance is now a legal question.

ASK FOR the contract schedule on data: for each supplier, whether company data may be used to train or improve the supplier's models, and how that was negotiated.

Data residency

Where data is stored and processed, which determines which laws apply. Many AI services run in US data centres; European companies often need data to stay in the EU, and some regulators require it.

ASK FOR the data flow map for AI services: where company data is processed, under which jurisdiction, and the sign-off from legal and compliance.

© 2026 Agent & Capital Ltd

For general information only. Not legal, financial or investment advice. See the full statement on the last page.

Synthetic data

Artificial data generated to resemble real data, used to train or test models where real data is scarce, private or expensive. Useful, with the caveat that it carries the assumptions of whatever made it.

ASK FOR the validation record: where synthetic data is used and the test that showed the model behaves the same on real data.

Model drift

The decline in a model's accuracy over time as the world changes and the data it was trained on grows stale. A model that was right at launch can be quietly wrong a year later unless it is monitored.

ASK FOR the monitoring standard: how each model's accuracy is measured in production, how often, the threshold that triggers action, and who acts.

Data provenance

A record of where data came from and what was done to it. It matters for AI because results can only be trusted as far as their inputs, and because regulators and courts now ask.

ASK FOR the data lineage documentation for the systems that feed AI decisions, and who maintains it.

Sovereign AI

Models and infrastructure run inside a jurisdiction, under its law, by suppliers subject to it. In Europe it means EU suppliers and EU data centres, and it has become a board question since companies started weighing US and Chinese dependence.

ASK FOR the dependency assessment: which AI capabilities rely on suppliers outside the EU, the risk in each, and the alternative where one exists.

COMPLIANCE

EU AI Act

The European Union's law on AI, in force since August 2024 and phased in over several years. It bans some uses, places heavy obligations on high-risk systems (such as those used in hiring, credit or critical infrastructure), requires transparency for chatbots and generated content, and fines up to €35m or 7% of worldwide turnover. It applies to companies that use AI, not only those that build it.

ASK FOR the AI Act compliance plan: each AI system classified under the Act, the obligations that follow, the owner, the deadline and the status.

High-risk AI system

The AI Act's category for systems that affect people's safety or rights: recruitment, creditworthiness, education, essential services, law enforcement, and safety components of products. These carry obligations for risk management, data quality, documentation, human oversight and registration.

ASK FOR the high-risk assessment: which systems fall into the category, the documentation each requires, and the gap analysis against it.

© 2026 Agent & Capital Ltd

For general information only. Not legal, financial or investment advice. See the full statement on the last page.

General-purpose AI model

The AI Act's term for the large models that can be put to many uses, the ones most companies rent. Obligations on their providers have applied since August 2025: documentation, copyright policy and, for the largest, risk assessment. The obligations decide what a supplier must be able to give its customers.

ASK FOR the supplier compliance file: the documentation each model provider has supplied under the Act, and who reviewed it.

Transparency obligations

The AI Act's duties to tell people when they are dealing with a chatbot, and to label generated content, including deepfakes. The first obligations most companies have to meet, and the most visible to customers.

ASK FOR the transparency policy: where customers and staff interact with AI, how they are told, and how generated content is labelled.

AI literacy

The AI Act's requirement (Article 4), in force since February 2025, that companies ensure the people who operate or use AI systems have sufficient understanding of them. It reaches the board: directors who approve AI investments and oversee AI risk are among the people the article covers.

ASK FOR the AI literacy programme: who has been trained, on what, how it is recorded, and how the board itself has met the requirement.

NIST AI RMF and ISO/IEC 42001

The two most used governance frameworks. The US NIST AI Risk Management Framework is voluntary guidance for managing AI risk. ISO/IEC 42001 is a certifiable standard for an AI management system, the AI equivalent of ISO 27001 for security. Both are being adopted by companies that want to show customers and regulators that AI is governed.

ASK FOR the AI governance framework the company has adopted, the gap assessment against it, and whether certification is planned.

Model risk management

The discipline, long established in banking, of validating models before use, monitoring them in production and documenting their limits. Supervisors now expect it to cover AI, and other sectors are adopting it.

ASK FOR the model risk policy: validation before use, monitoring in production, a named owner per model, and the model inventory it is applied to.

Explainability

The degree to which a model's output can be explained in terms a person can follow. Language models are poor at it. Where a decision must be justified to a customer or a regulator, such as refusing credit, a system that cannot explain itself is a liability.

ASK FOR the decision inventory: which decisions must be explainable by law or policy, which AI systems take part in them, and how each explains its output.

© 2026 Agent & Capital Ltd

For general information only. Not legal, financial or investment advice. See the full statement on the last page.

Bias

Systematic unfairness in a model's outputs, usually inherited from its training data. In hiring, lending and pricing it is a legal exposure as well as an ethical one.

ASK FOR the fairness testing report for every AI system that affects people, with the method, the findings and the remediation.

Ownership of AI output

Who owns what a model produces. In most jurisdictions copyright needs a human author, so purely AI-generated text, code or images may belong to no one, and a competitor may copy them. The position differs by country and is still moving. A company building products on generated content needs to know what it can protect.

ASK FOR the intellectual property policy for AI output: how human authorship is recorded where it matters, which generated material the company treats as unprotected, and the supplier terms on ownership of output.

Copyright and training

The unresolved question of whether training a model on copyrighted material is lawful, and whether a model's output can infringe. Publishers, music rights bodies and authors have sued in the US and Europe, with mixed results so far. Any company using generated content carries some of this risk.

ASK FOR the supplier contract terms on intellectual property: indemnities for infringing output, and the company's own rules on where generated content may be used.

Personal data

Any information about an identifiable person, protected in Europe by the GDPR. Feeding personal data to an AI model is processing it, with all the obligations that follow, and models can memorise what they are shown.

ASK FOR the data protection impact assessment for each AI system that handles personal data, and the legal basis recorded for each.

RISK AND SECURITY

Shadow AI

AI tools used by employees without the company's knowledge or approval: a free chatbot given a customer contract, a browser plug-in reading email. It is the most common way company data leaves the building.

ASK FOR the acceptable use policy for AI tools, the list of approved tools, and the monitoring that shows what else is in use.

Prompt injection

An attack in which malicious instructions hidden in a document, a web page or an email are read by an AI agent and obeyed (the indirect form, and the dangerous one; the direct form is a user typing the instruction). An agent that reads untrusted content and can act on it can be turned against its owner. There is no complete defence yet; limits on what agents may do are the main control.

ASK FOR the threat assessment per agent: which read outside content, what they can do afterwards, and the controls in place.

© 2026 Agent & Capital Ltd

For general information only. Not legal, financial or investment advice. See the full statement on the last page.

AI supply chain

The third-party code inside AI systems: open-source libraries, orchestration frameworks, connectors and the models themselves, much of it updated weekly and little of it reviewed. A weakness in any of them is a weakness in production.

ASK FOR the AI software bill of materials: every third-party component in each AI system, its version and source, and the review process before a component is added or updated.

Data leakage

Company information reaching an AI supplier, another customer of that supplier, or the public, through prompts, training or logs. Trade secrets are the sharpest case: business logic, source code or pricing rules pasted into a third-party tool may lose their legal protection as secrets. Contracts, configuration and staff behaviour all play a part.

ASK FOR the data handling standard for AI: what may be entered into which tools, how trade secrets and prompt libraries are protected, the contract terms on retention and training, and the technical controls.

Red teaming

Deliberately attacking an AI system to find how it fails before someone else does: getting it to leak data, take unauthorised actions or produce harmful output. The AI equivalent of penetration testing.

ASK FOR the red team schedule and the latest findings for each AI system in production, with the fixes and their status.

Evaluation

The test an AI system is measured against before release and after every change: a set of real cases with known right answers, scored, and run many times, because a model can answer the same case differently. The AI equivalent of acceptance testing, and the only evidence that a system does what the supplier says.

ASK FOR the evaluation report for each system: the test set, the score, the pass mark, who set it, and the result after the last model change.

Benchmark

Public tests that suppliers quote to rank their models against each other. Useful for comparing models in general, poor at predicting results on a company's own work.

ASK FOR the company's own evaluation results rather than the supplier's benchmark, before any model is chosen.

Deepfake

Synthetic audio, video or images that convincingly imitate a real person. Used for fraud, including fake calls from a chief executive authorising a payment, and for reputational attacks.

ASK FOR the payment and identity verification procedure, updated for voice and video that cannot be trusted, and the staff training that goes with it.

© 2026 Agent & Capital Ltd

For general information only. Not legal, financial or investment advice. See the full statement on the last page.

AI-enabled cyber attack

Attacks that use AI to find weaknesses faster, write malicious code and craft convincing messages. Bank chief executives report that the time between a vulnerability being found and exploited has fallen from days to minutes, and that cyber budgets are running above plan as a result.

ASK FOR the cyber plan to show what AI has added to the threat and to the budget, and for that line to appear in the AI business case.

Agent liability

Who is responsible when an AI agent causes harm: the supplier, the company that deployed it or the person who set its task. Courts and legislators are only starting to answer. Until they do, the company that deploys an agent should assume the liability is its own.

ASK FOR the incident response procedure for AI: who is told when an agent acts outside its remit, who decides, who informs customers and regulators, and the insurance that applies.

Answer engines

Search tools that answer the question instead of listing links: Google's AI Overviews, ChatGPT, Perplexity. For any business that depends on being found online, they change the route to the customer, and publishers are already reporting lost traffic. A German court has held Google liable for false statements in its AI Overviews.

ASK FOR the channel report: how much customer traffic came through search, what has happened to it since AI answers appeared, and the plan for being found by answer engines.

Surveillance pricing

Setting individual prices using personal data and AI. Legal in many places, politically exposed everywhere; Delta Air Lines had to assure a US congressman that its AI pricing tool is not given personal information about customers.

ASK FOR the pricing policy: what data any pricing model may use, what it may not, and the statement the company would make in public if asked.

PEOPLE

Jobs attributed to AI

Job cuts announced with AI as the reason. Across the companies that report it, AI has become the leading stated reason for cuts; whether the work has gone to AI or AI is the explanation for a cut planned anyway is a question for each case. Unions in Norway and Germany are now contesting the link.

ASK FOR the task analysis behind any cut attributed to AI: the tasks agents now perform, the people affected, the saving in money and the risk assessment, including union and works council response.

Works council and union consultation

In Germany, the Netherlands, the Nordics and elsewhere, introducing AI that changes how people work triggers consultation rights, and in some cases co-determination. Treating this as a formality has delayed transformations and cost goodwill.

ASK FOR the consultation record and the works agreement on AI, if one exists, and the plan for the next round.

© 2026 Agent & Capital Ltd

For general information only. Not legal, financial or investment advice. See the full statement on the last page.

Reskilling

Training existing staff for the work that remains after AI takes over tasks. It is the alternative to cutting, and the part of most AI plans that carries the least money.

ASK FOR the reskilling plan with its budget, set against the headcount saving in the same plan.

Adoption rate

The share of staff using AI regularly, the most reported AI figure of 2026 and the least informative. It says that people are using the tools, not what the tools returned.

ASK FOR the adoption report to be paired with the value report: usage by function next to the saving or revenue from that function.

UNSETTLED

Terms that are in use before they have a settled meaning. When a word has no agreed definition, the person using it is defining it, usually to sell something. One Ask covers them all: ask for the definition in writing, with a measure, before any money or target is attached to the word. Vibe coding, which is a practice as well as a label, has one more of its own.

Answer engine optimisation (GEO, AEO, LLMO)

Getting a company recommended inside AI answers, the successor to search engine optimisation. The marketing trade has three names for it and no agreed definition, and there is no reliable way yet to measure whether it works. It matters because the route to the customer is changing; the risk is paying for a discipline that does not exist yet.

Artificial general intelligence (AGI)

AI that matches people at most tasks. Every supplier defines it differently, several have contracts that turn on reaching it, and no test settles the question. For a board it signals a supplier's ambitions, not a planning assumption.

Agent washing

Calling an ordinary workflow, a chatbot or a scripted automation "agentic". Gartner coined the term for what a large share of agent products are. It belongs here because it names what a board is often being sold.

AI-native

A company, product or process designed around AI from the start rather than retrofitted. Used in valuations and pitches without a definition, so it should be read as a claim to be tested, not a category.

AI factory

A supplier term for a data centre that turns computing into AI output, used by Nvidia for the hardware and by Deutsche Telekom for a service. Sometimes a product, sometimes a metaphor; which one decides whether it belongs in capital spending or in a slide.

© 2026 Agent & Capital Ltd

For general information only. Not legal, financial or investment advice. See the full statement on the last page.

Vibe coding

Building software by describing it to an AI and accepting what comes back, with little or no reading of the code. Real, fast and already inside companies. It moves quality control from writing the code to reviewing it afterwards, and nobody has settled who does that review or who is accountable when it fails.

ASK FOR the review standard for AI-written code: what a person must read before release, who signs it off, and who answers when it fails.

Digital worker

A marketing term for an AI agent that does a job previously done by a person. Some companies now report counts of digital workers alongside headcount. The comparison of cost per worker is worth making, on the same basis as a person's, including the failure rate.

WHAT A BOARD SHOULD BE ABLE TO ASK FOR

The documents behind the Ask lines, listed once. A company that can produce these has AI under governance; one that cannot has a programme and no control over it. They are grouped by the committee that would normally own them; where the board has no such committee, the full board does.

RISK

1. **The AI inventory:** every AI system and agent in use or planned, its process, its owner, its classification under the EU AI Act.
2. **The agent register:** per agent: systems and data it may touch, actions it may take without a person, written limits, elevated rights, the owner who approved it.
3. **The AI policy and acceptable use rules:** what staff may use, what may be entered into which tools, how customers are told.
4. **The security file:** red team findings, prompt injection assessment per agent, the software bill of materials, logging standard, shutdown procedure and its last test.
5. **The incident response procedure for AI:** who is told, who decides, who informs customers and regulators, which insurance applies.

AUDIT

6. **The business case standard:** full cost over three years including tokens at production volume, compute, people and replacement; the capitalisation policy; signed off by finance.
7. **The AI value report:** by use case, money returned this year against full cost; adoption figures paired with value; where savings went.
8. **The compliance file:** AI Act classification and plan, data protection impact assessments, fairness testing, intellectual property policy, the AI literacy programme including the board's own.
   WITH NOMINATION FOR THE BOARD'S LITERACY

TECHNOLOGY, OR THE FULL BOARD

9. **The model and supplier register:** which models and suppliers the company depends on, contract terms on data, training, price, termination and indemnity, the exit assessment and the concentration report.

© 2026 Agent & Capital Ltd

For general information only. Not legal, financial or investment advice. See the full statement on the last page.

10. **The evaluation and monitoring standard:** test sets, pass marks, results after each model change, production accuracy, drift thresholds.

REMUNERATION, OR THE FULL BOARD

11. **The people plan:** task analysis behind any cut attributed to AI, consultation record, reskilling plan with its budget.

ABOUT THIS GLOSSARY

Written by Ilona Simpson, founder of Agent & Capital, a career technologist turned board director and advisor. This is version 1.2, dated 12 October 2026. Revisions are numbered and dated at the top of the page, and terms are added as readers ask for them.

If a term you meet is missing, or one here is explained badly, send it with where you met it and I will add it in the next revision.

The weekly companion is AI ROI: what AI is doing in real companies, in their own numbers. Every Monday, with three questions for your board. airoi.beehiiv.com/subscribe

CHANGE LOG

Version 1.1, 12 October 2026: an Ask added to vibe coding; three pairs shortened; the closing list grouped by committee.
Version 1.2, 12 October 2026: Copilot vs agent reworded; the number of terms removed from the introduction.

AI, in board terms is written by Ilona Simpson and published by Agent & Capital Ltd, 71-75 Shelton Street, Covent Garden, London WC2H 9JQ. © 2026 Agent & Capital.

This glossary is general information for directors and executives. It is not legal, regulatory, financial or investment advice, and it is not a substitute for advice on your own circumstances. The law described here changes; check the current position before acting on it. Agent & Capital Ltd accepts no liability for decisions taken in reliance on this page. Every entry is checked against its source, but mistakes can still happen; corrections are made in the next revision.