AI GLOSSARY – 2026-10-27

Thank you for downloading

Thank you for downloading this AI glossary from House of Silver. We hope it gives your leaders a shared language for AI, so conversations about it move faster and land better.

AI comes with a lot of new words, some precise, some used loosely, and a few that are mostly marketing. This glossary explains the terms leaders hear most, in plain language, with a note on what many of them mean in practice. Use it on its own, alongside an AI strategy or policy, or copy the terms you need into your intranet. The field moves quickly, so expect to revise it.

If you would value an experienced partner along the way, whether to facilitate a leadership session, build shared understanding across your teams, or help them adopt what you decide, we would be glad to hear from you.

Erik
Founder, House of Silver
hello@houseofsilver.ai | houseofsilver.ai

About House of Silver
House of Silver helps organizations adopt AI in ways that strengthen their people, keep their data safe, and deliver real impact. To see where your organization stands, take the free AI Readiness Assessment at houseofsilver.ai.

[Organization name] AI Glossary

Version	
Glossary owner	
Approved by	
Effective date	
Next review	

Contents

Level 1: What every leader needs 3

Level 2: Directing AI agents 6

- How agents work 6
- Connecting agents to your systems 6
- Keeping agents safe and accountable 7
- Under the hood 7

The four levels of AI autonomy 8

How to use this glossary: Level 1 covers what every leader needs to know about AI, in plain language. Level 2 is for leaders who direct, approve or invest in AI agents, grouped by topic. Terms are listed alphabetically within each level or group, and the four levels of AI autonomy close the glossary. Use it as it is, adapt the definitions to your organization, or copy the terms you need into your intranet, policy or strategy. Definitions reflect common use as of October 2026.

Level 1: What every leader needs

The terms every leader should be able to use with confidence: in leadership meetings, in plans, and in conversations with vendors.

Agent owner – The named person accountable for an AI agent: what it is allowed to do, what it produces, and when it is changed or retired.

> ***In practice:*** *If no one can say who owns an agent, it should not be running.*

Agentic AI – AI systems that pursue goals with some independence: planning steps, using tools and adjusting along the way. The umbrella term for AI agents and the systems that run them.

AI agent – AI that carries out work on your organization's behalf. It plans the steps toward a goal, uses systems and data, and acts within limits people set. Where an AI tool helps one person, an agent extends a team. See also AI tool and Agent owner.

> ***In practice:*** *An agent might sort incoming documents, draft a weekly report and file it, without anyone prompting each step.*

AI assistant – A chat-based AI tool that answers questions, drafts text and summarizes material when a person asks it to. For most people, the first experience of generative AI.

> ***In practice:*** *The person asks, reads the answer and decides what to do with it.*

AI literacy – The knowledge and skills people need to use AI well: what it does well, where it fails, and how to check its work. In the EU, organizations that use AI must take steps to support AI literacy among their staff.

> ***In practice:*** *AI literacy differs by role. Everyone needs the basics; people who direct or approve AI work need more.*

AI tool – Software that people use to do their own work with AI, such as an AI assistant or an AI feature built into a word processor or spreadsheet. The person stays in charge of each step. A list of approved AI tools tells people what they may use. See also AI agent.

Approved AI tool – An AI tool your organization has reviewed and allowed for a defined purpose and type of information.

> ***In practice:*** *Approval usually covers where the data goes, whether it is used to train the provider's models, security and cost.*

Artificial general intelligence (AGI) – A hypothetical AI that could learn and perform almost any intellectual task at a human level, across fields. It does not exist today, and experts disagree on whether or when it will.

> ***In practice:*** *Useful to recognize in the headlines. Not a basis for this year's decisions.*

Artificial intelligence (AI) – Software that performs tasks that normally need human intelligence, such as understanding language, recognizing patterns, reasoning and making predictions.

Automation – Software that follows fixed rules to complete a task the same way every time. It does not decide anything for itself, unlike an AI agent, which pursues a goal and can choose its own steps.

> ***In practice:*** *Many products sold as AI agents are automation with a new label. See Agent washing.*

Autonomy level – How much an AI may do on its own before a person steps in, from assisting a person to acting on its own and reporting back. See The four levels of AI autonomy at the end of this glossary.

Bias – Systematic unfairness in AI output, usually inherited from the data the model learned from or from the way a question is framed.

> ***In practice:*** *Matters most where AI informs decisions about people, such as hiring, performance or pay.*

Context – Everything an AI has in front of it when it works: your instructions, the documents and data it can see, and the conversation so far. The quality of the output depends heavily on the quality of the context.

> ***In practice:*** *An AI that knows your standards, templates and past work gives far better answers than one that starts from nothing.*

Data residency – Where your data is physically stored and processed, and therefore which country's laws apply to it. Many AI services process data outside your country unless you choose otherwise.

> ***In practice:*** *Ask every AI vendor where your data goes, including backups and support access, and whether it can stay in your region.*

Deepfake – Realistic image, audio or video, created or altered with AI, that shows a person saying or doing something they did not. In the EU, deepfakes must be labeled as such.

> ***In practice:*** *A good reason to confirm unusual requests, such as new payment instructions, through a second channel.*

EU AI Act – The European Union's law on artificial intelligence. It applies to organizations that use AI in the EU, not only to those that build it, and sorts AI uses by risk. High-risk uses, such as AI in hiring or in evaluating employees, carry the strictest obligations.

> ***In practice:*** *Most everyday AI use is not high-risk, but duties such as supporting AI literacy and telling people when they are dealing with AI already apply.*

Generative AI – AI that creates new content, such as text, images, audio, video or code, based on patterns learned from existing material.

Hallucination – When AI produces information that sounds plausible but is false, such as an invented fact, quote or source.

> ***In practice:*** *The main reason AI output is checked before use, especially facts, figures and references.*

Human in the loop – A set checkpoint where a person must review or approve AI work before it moves on. The work waits for the person. See also Human on the loop.

Human on the loop – A person supervises AI work and can step in or stop it, but the work does not wait for approval at each step. Suits high-volume work that is easy to reverse. See also Human in the loop.

Human oversight – The principle that people monitor AI, can correct or stop it, and remain accountable for its results. Human in the loop and human on the loop are two ways of providing it.

Inference – The moment a model does its work: taking your input and producing an output. Training is how a model learns; inference is how it responds. Every request uses inference, which is why AI usage has a running cost.

Large language model (LLM) – A type of foundation model trained on vast amounts of text to understand and generate language. LLMs power AI assistants and most AI agents.

Local AI model – A model that runs on hardware your organization controls instead of a provider's cloud. Data stays in-house, which suits confidential material, but capability depends on the hardware.

Model – The trained AI system doing the work behind a product. Products change models over time, so the same product can behave differently after an update.

Personal data – Any information that can identify a person, directly or combined with other data, such as names, email addresses, ID numbers or location. Also called personally identifiable information (PII). Privacy laws such as the GDPR govern how it is handled, including in AI tools.

Prompt – The instruction or question a person gives an AI. Prompts that state the goal, the audience and what good looks like produce better results.

Prompt engineering – The practice of wording prompts to get better results from AI: stating the goal, giving context and examples, and describing what good looks like. As AI takes on longer tasks, the focus is shifting to context engineering.

Public AI tool – An AI tool open to anyone, often free, that your organization has not approved. What you enter may be stored by the provider or used to train future models.

> ***In practice:*** *The main risk is confidential or personal information leaving your control.*

Responsible AI – Using AI in ways that are fair, safe, transparent and accountable, with clear human responsibility for the outcome.

Shadow AI – AI tools and agents used for work without the organization's knowledge or approval.

> ***In practice:*** *Usually a sign that people need something the approved tools do not offer. Bans tend to push it further out of sight.*

Training – The process by which a model learns from large amounts of data. Training happens before you use a model. Whether what you type is later used to train future versions depends on the provider and your settings.

Use case – A specific task or process where AI is applied, with a clear purpose and a way to tell whether it works.

> ***In practice:*** *"Use AI in marketing" is an ambition. "Draft first versions of case studies from project data" is a use case.*

Validation – Checking that AI output is correct, complete and fit for use before anyone relies on it, done by someone qualified to judge it.

Level 2: Directing AI agents

For leaders who direct, approve or invest in AI agents. The terms are grouped by topic: how agents work, how they connect to your systems, how they are kept safe and accountable, and a few technical terms you will hear along the way.

How agents work

Agent washing – Marketing existing products, such as chatbots or automation, as AI agents when they have little real agent capability.

> ***In practice:*** *Ask what the product decides on its own, and what it does when something unexpected comes up.*

Computer use – An AI agent operating software the way a person does, by reading the screen, clicking and typing. Useful where no connector exists, and riskier, because the agent can reach anything on the screen.

Context engineering – The work of deciding what information an AI receives for each task: which instructions, documents, examples and data, and what it should not see. It has largely replaced prompt engineering as AI takes on longer tasks.

> ***In practice:*** *Much of it is writing down how your organization works, so AI can follow it.*

Handoff – The moment a task, its context and its responsibility pass from one agent to another, or between an agent and a person. Many errors happen at handoffs.

Memory – An AI's ability to keep information across conversations or tasks, such as preferences or earlier decisions. Useful, as long as you know what is stored and who can see it.

Multi-agent system – Several AI agents working together on one goal, each with its own role, often coordinated by an orchestrator.

Orchestration – Coordinating several AI agents, tools and people toward a shared goal. The orchestrator, which can be an agent or a person, splits the work, assigns it and combines the results.

Skill – A packaged set of instructions, examples and resources that teaches an AI agent how to do a specific task your way. Skills can be shared and reused across agents, and several major AI providers support an open standard for them.

System prompt – Standing instructions that set an AI's role, tone and limits before any request. It guides behavior but does not enforce it, which is why guardrails sit outside it.

Connecting agents to your systems

A2A (Agent2Agent protocol) – An open standard that lets AI agents from different vendors find each other, exchange tasks and report results. Hosted by the Agentic AI Foundation, alongside MCP.

> ***In practice:*** *It is what will let your agent hand work to a supplier's or client's agent.*

Connector – A link that lets AI read from or act in another system, such as email, a document library or a CRM. Each connector widens what AI can see and do, so each one deserves a deliberate decision.

Grounding – Tying AI output to sources it can point to, such as your own documents, so its claims can be checked.

MCP (Model Context Protocol) – An open standard for connecting AI to tools and data sources in a consistent way, without a custom integration for each one. Widely supported across AI providers and hosted by the Agentic AI Foundation.

RAG (retrieval-augmented generation) – A method where AI first looks up relevant information in a defined set of sources, such as your approved documents, and then answers based on what it found. Improves accuracy and lets answers be traced to a source.

Keeping agents safe and accountable

Agent identity – Giving each AI agent its own login and permissions instead of letting it borrow a person's account, so every action can be traced to the agent and limited to what it needs.

Agent record – A list of every AI agent in use: what it does, who owns it, what it can reach, its autonomy level and when it is next reviewed. Also called an agent registry or inventory. Kept apart from the list of approved AI tools.

Audit trail – A record of what an AI did, when, at whose request and with what information, so any result can be traced and explained afterward.

Evals (evaluations) – Structured tests that check how well an AI performs on realistic tasks, covering accuracy, safety, cost and speed. Run before launch and after every change.

> ***In practice:*** *The AI equivalent of a quality check, and the only reliable way to compare tools or models.*

Excessive agency – Giving an AI more access or permissions than its task needs, such as letting an agent that only needs to read documents also delete them. One of the leading security risks for agents.

Guardrails – Controls that keep AI within agreed limits by blocking certain inputs, outputs or actions. Some are built into the AI itself; the stronger ones are enforced by the systems around it.

Prompt injection – Hidden instructions placed in content an AI reads, such as a web page, email or document, that trick it into doing something it should not. The top security risk for AI applications, and one that can be managed but not eliminated.

> ***In practice:*** *The more an agent can do, the more damage a successful injection can cause. Limit its access accordingly.*

Sandbox – An isolated environment where an AI agent can be tested, or can work, without reaching live systems or sensitive data.

Stop condition – A defined point or trigger at which an AI agent must pause or stop, together with the means for a person to stop it at any time.

Under the hood

Context window – The amount of information a model can consider at once, measured in tokens. A large window helps, but more context is not always better context.

Foundation model – A large, general-purpose AI model trained on broad data that can be adapted to many tasks. Most AI products are built on a small number of foundation models from a few providers.

Reasoning model – A model that works through a problem step by step before answering. Better at complex analysis and planning, and slower and more costly per answer.

Token – The unit AI models use to read and write text, roughly three quarters of an English word. Many AI services are priced per token, so tokens drive the cost of usage-based AI.

The four levels of AI autonomy

A simple way to describe how much an AI may do on its own. The level belongs to the task, not the technology: the same agent can draft for review in one process and act within set limits in another. The higher the level, the more an organization needs a named owner, written limits, an audit trail and a stop condition.

Autonomy level	What AI does	What people do
Assists a person	Answers, suggests or summarizes when asked.	Do the work and make every decision.
Drafts for review	Produces a draft or a proposed action.	Review and approve before anything is used.
Acts within set limits	Carries out defined tasks on its own, inside written limits.	Set the limits, check samples and handle exceptions.
Acts and reports back	Completes a task or process end to end and reports what it did.	Review the reports, keep the audit trail and stop it when needed.