

The AI Dictionary.

Every AI term that actually matters -
what's behind it, what it's for, and what it
means for your business.

*Not just what the words mean.
How to use AI well.*



**Clarity for today.
Confidence for tomorrow.**

The only dictionary you'll actually use.



Built for business leaders & their teams
by Kamenno, your AI adoption partner



kameno.ai
contact@kameno.ai

START HERE

Two kinds of words in here.

This is the plain-English translation for the AI terms you keep hearing. What each one means, what it's for, and, where it helps, an example you'll get on the first read. No textbook, no lecture.

But not every term matters equally to a business owner. So we've split them into two. Most entries are **words you may come across**, useful to recognise. A smaller set are marked **Essential for business**: the concepts every SME leader & their teams should actually understand to use AI safely, consistently and well. If you only read one kind, read those.



About Kameno

Kameno is an AI education and enablement company for small and medium sized businesses. We help founders, business owners and their teams understand, adopt and actually apply AI without needing a technical background.

Our name comes from the Spanish *camino* (path, or road). Every business owner has a path to AI. Most just need a guide to walk it with them. We keep it plain, honest and practical: what AI can do for your business, and just as importantly, what it can't.



How to use this dictionary

- **Skim or search.** It's A–Z. Jump to the word you need, or read it through once to get your bearings.
- **Look for the badge.** Entries marked **Essential for business** sit in a highlighted card - these are the ones worth reading closely, even if you skip the rest.
- **Same idea, different label.** The tools borrow each other's ideas and rename them. An "agent" in Copilot isn't quite an "agent" in Claude. The cheat sheet on the next page lays the big ones side by side.

★ **ESSENTIAL FOR BUSINESS**

A concept every SME leader should understand - governance, safety, data, adoption.

Plain term

A word you may come across and simply want to recognise.

A note on names: these terms reflect where the industry sat in 2026. The tools move fast - the concepts tend to last, the product names less so.

THE CHEAT SHEET

Same idea, four different names.

The big AI tools keep borrowing each other's best features and giving them their own brand name. Here are the ones that trip people up most. When someone says a word, this tells you what they'd call the same thing in a different tool.

The idea	 Claude Anthropic	 ChatGPT OpenAI	 Copilot Microsoft	 Gemini Google
 A saved workspace for one topic or task	Projects	Projects / Custom GPTs	Agents (Copilot Studio)	Gems
 An AI that writes and runs code for you	Claude Code	Codex	GitHub Copilot	Gemini CLI / Code Assist
 Takes its time to research the web and hand back a report	Research	Deep Research	Researcher	Deep Research
 A side panel you draft and edit a document in, with the AI	Artifacts	Canvas	Pages	Canvas
 The AI remembering you between chats	Projects & memory	Memory	Microsoft 365 context	Personal context
 Model sizes (small & fast → big & capable)	Haiku → Sonnet → Opus	mini → full → reasoning	(uses OpenAI + MS models)	Flash → Pro



"Agent" is one of the slipperiest words of all. In Copilot it usually means a configured assistant you build for a task. In Claude or ChatGPT it more often means an AI that can pursue a goal and chain actions on its own until it's done. Same word - see the full entry on the next page.

THE DICTIONARY

A to Z, in plain English.

A

Agent

An AI that doesn't just answer questions - it can make decisions and take actions in real systems: book meetings, browse the web, work with files or trigger other tools.

A chatbot talks; an agent does.



THE WORD SHIFTS BY TOOL

Copilot: usually a configured assistant you build for a task. **Claude / ChatGPT:** usually an AI that pursues a goal and chains its own steps.

See also: *Agentic workflow, Chatbot, MCP.*

Agentic IDE

A coding environment where the AI doesn't just suggest code but can also run it, test it, spot errors and fix them. More relevant to developers than most business users, but useful to understand when people talk about AI building software.

See also: *IDE, Claude Code, Agent.*

Agentic workflow

A workflow where AI is given a goal and can decide some of the steps to get there, rather than following one rigid sequence. Instead of scripting every reply to a customer enquiry, you might let it identify the type of enquiry, find the right information, draft a reply and decide whether a person needs to review it.

The important difference is decision-making: an automation follows instructions; an agentic workflow can choose what happens next.

See also: *Agent, Automation, Workflow.*

AI *artificial intelligence*

Software that can do tasks we used to associate with human intelligence: understand language, recognise images, reason, generate content or make decisions. It's a big umbrella. In this guide, when we say AI we usually mean *generative* AI - the kind that produces text, images, audio, video or code from what you ask.

See also: *LLM, Generative vs agentic AI.*

★ ESSENTIAL FOR BUSINESS

AI champion

A person inside the business who helps colleagues use AI well. They don't need to be a programmer or the company's "AI expert". Their role is usually to help the team find useful applications, share good practices, answer basic questions, encourage safe use and keep AI adoption moving.

Think of them as the person who turns AI from something a few individuals experiment with into something the business uses consistently.

See also: AI literacy, Use case, Responsible AI.

★ ESSENTIAL FOR BUSINESS

AI governance

The rules, responsibilities and processes a business puts around its use of AI. Good governance answers questions such as:

- What AI tools are employees allowed to use?
- What information can they put into them?
- What needs human approval?
- Who is responsible if something goes wrong?
- How do we know AI is being used safely and appropriately?

For an SME, governance doesn't need to mean a 50-page policy. It means having enough structure to use AI confidently without creating unnecessary risk.

See also: AI policy, Responsible AI, Guardrails, Shadow AI.

★ ESSENTIAL FOR BUSINESS

AI literacy

The ability to understand AI well enough to use it sensibly. That includes knowing what AI can and can't do, how to give it good instructions, how to check its answers, how to protect sensitive information and when human judgement is still required.

AI literacy is not about becoming technical. It's about being capable, critical and responsible when using AI at work.

See also: AI champion, Responsible AI, Human oversight.

★ ESSENTIAL FOR BUSINESS

AI policy

A business document that sets out how employees may and may not use AI. A useful policy usually covers approved tools, confidential information, human review, copyright, customer data, accountability, and what to do when someone's unsure.

For SMEs, the best policy is usually short enough that people actually read it.

See also: *AI governance, Data privacy, Copyright.*

API *application programming interface*

The service window two programs talk through. A fixed set of rules: you send a request in the agreed format, the other system sends a response back. When AI "connects" to your CRM, calendar, accounting platform or database, it's often doing it through an API.

See also: *API key, Connector, MCP.*

API key

The password that says who's using an API. When you connect a tool or agent to an AI service, you may use an API key so the provider knows the request belongs to your account - and bills you for what's used. Keep it secret: whoever has the key may be able to use the service in your name.

See also: *API, Token.*

App, model & assistant

Three words people often blur together. The **app** is the product you open (ChatGPT, Claude, Gemini). The **model** is the AI engine inside it. The **assistant** is the experience you're interacting with. In short: you talk to an assistant inside an app, powered by a model.

See also: *Chatbot, LLM, Model.*

Automation

Chaining tasks so they run themselves when something happens, without somebody manually starting each step. For example: a web enquiry arrives → add it to the CRM → notify the salesperson → draft a follow-up email. Tools such as Zapier, Make and n8n are commonly used to build these. Adding AI means some steps can now interpret, classify or create information rather than simply move it.

See also: *Agentic workflow, Workflow, No-code.*

C

Canvas / Artifacts

A workspace beside the AI chat where a document, table or piece of code can live and be edited. Useful when you're working with the AI on something over several rounds, instead of constantly scrolling through a conversation.



WHAT EACH ONE CALLS IT

Claude: Artifacts. **ChatGPT:** Canvas. **Gemini:** Canvas. **Copilot:** Pages.

Chatbot *conversational assistant*

The visible face of AI: the app you talk to by typing or speaking. The chatbot is the interface; the LLM or model is the engine underneath it. ChatGPT, Claude, Gemini and Copilot are all examples of conversational AI products.

ChatGPT - OpenAI (US). Runs on the GPT family.	Gemini - Google (US). Runs on the Gemini family.
Claude - Anthropic (US). Haiku, Sonnet, Opus.	Copilot - Microsoft (US). OpenAI + Microsoft models.

See also: *LLM, Model, App, model & assistant.*

Chunking

Breaking a long document into smaller pieces so the AI can find and use only the relevant parts. Instead of feeding an entire 200-page manual into an AI every time someone asks a question, a system can break it into chunks and retrieve only the sections that matter.

See also: *Context window, RAG, Embedding.*

Claude

Anthropic's AI assistant, and the name associated with its family of AI models. It sounds a little like "cloud", which sometimes causes confusion, but the two are unrelated.

See also: *Chatbot, Model, Claude Code.*

Claude Code

Anthropic's AI coding tool that can work directly with files and software projects. Unlike a chatbot that merely suggests code, it can carry out tasks such as creating, changing and testing code.

See also: *Agentic IDE, Terminal, Agent.*

CLAUDE.md

A file used to give Claude Code ongoing instructions and context about a particular project. It's a technical concept, but the broader idea is useful for businesses: instead of repeating the same instructions every time, you can give an AI a persistent set of rules about how you want it to work.

See also: *System prompt, Skill, Projects.*

Connector

A link that lets an AI service access information or actions inside another business system - your email, SharePoint, Google Drive, CRM, calendar or another application. Think of it as the bridge between the AI and the systems where your work already lives. Connecting a system doesn't automatically mean the AI can see everything - permissions still matter.

See also: *API, MCP, Data privacy.*

Context engineering

The practice of giving AI the right information at the right time - your instructions, background, examples, documents, constraints and relevant business information. Good context usually matters more than simply giving the AI more information.

See also: *Context window, Context rot, Prompt engineering.*

Context rot

What happens when you give an AI so much background that the important parts get lost. More context isn't always better: if the useful signal is buried in irrelevant material, the quality of the response can fall.

See also: *Context window, Context engineering.*

Context window

The amount of information an AI model can work with at one time. Your conversation, instructions, documents and question all compete for that space. Think of it like the AI's working desk: there's a lot you can put on it, but it isn't unlimited.

See also: *Context rot, Chunking, Token.*

Copilot

Microsoft's name for a range of AI assistants and capabilities across products such as Microsoft 365. For businesses using Word, Excel, Outlook, Teams and other Microsoft products, Copilot is designed to bring AI into tools employees already use.

See also: *Chatbot, Agent, Custom assistant.*

★ ESSENTIAL FOR BUSINESS

Copyright / intellectual property

The legal rights attached to things people and organisations create (writing, designs, photographs, software, videos, trademarks and confidential business material). AI creates new questions here, because both the information you put into a tool and the content it produces may have ownership or usage implications.

A practical rule for SMEs: don't assume that because AI generated something, there are automatically no copyright or ownership considerations.

See also: *AI policy, Data privacy, Responsible AI.*

Custom assistant *custom bot*

An AI assistant configured for a particular job, team, process or style. You might give it instructions, examples and business documents so it behaves consistently instead of starting from scratch every time.



WHAT EACH ONE CALLS IT

ChatGPT: a Custom GPT. **Claude:** a Project. **Gemini:** a Gem. **Copilot:** an agent.

See also: *Projects, Skill, System prompt.*

D

★ ESSENTIAL FOR BUSINESS

Data privacy / confidential data

Data privacy is about how personal, customer and business information is collected, handled, stored and shared. Confidential data is information the business doesn't want disclosed outside authorised people or systems. Examples include:

- customer records, employee information and contracts
- pricing, financial information and passwords
- unreleased strategies and commercially sensitive documents

Before putting information into an AI tool, employees should know whether that information is appropriate for that system and account.

See also: *Sensitive information, Enterprise data protection, AI policy.*

Deep research

An AI feature that spends longer investigating a question across multiple sources and returns a more substantial report. Useful when you need research rather than a quick answer - but the sources and conclusions should still be checked.

See also: *Grounding, Hallucination, Agent.*

E

★ ESSENTIAL FOR BUSINESS

Eliminate—Streamline—Automate

A useful filter before automating a business process. First ask, in order:

- **Eliminate:** does this task need to happen at all?
- **Streamline:** can the process be simplified?
- **Automate:** only then, should technology or AI do it?

The point isn't to automate everything. It's to automate the right things.

See also: *Automation, Workflow, Use case.*

Embedding

A way of converting information into numbers that represent its meaning. This lets AI systems find information that's conceptually similar even when the exact words are different. It's one of the technologies behind AI-powered search and RAG systems.

See also: *Vector database, RAG.*

★ ESSENTIAL FOR BUSINESS

Enterprise data protection

The security and privacy controls applied when AI is provided as part of a business or enterprise service. Depending on the provider and account type, this can include stronger controls around how company data is handled, stored, accessed and used.

The lesson for SMEs: the business version of an AI service may have different data protections from the free consumer version. Always understand the protection offered by the particular product and plan you're using.

See also: *Enterprise vs consumer AI account, Data privacy.*

★ ESSENTIAL FOR BUSINESS

Enterprise vs consumer AI account

The version of an AI service designed for an organisation can have different privacy, security, administration and data-handling controls from the version designed for an individual. This matters when employees use AI for work: a free personal account may not provide the same safeguards, management controls or contractual protections as an approved business account.

The logo might look identical. The protections behind it may not be.

See also: *Enterprise data protection, Shadow AI, AI policy.*

F

Fine-tuning

Additional training given to an existing AI model so it becomes better at a specific style, vocabulary or task. Useful in specialised applications, but most SMEs don't need to start here. Good prompting, instructions, workflows and access to the right company knowledge usually solve the problem first.

See also: *RAG, Prompt engineering.*

Frontier model

A term for the most capable AI models available at a particular point in time. They tend to be strongest on difficult tasks, but can also be slower or more expensive than smaller models.

See also: *Model, Lightweight model, Reasoning model.*

G

Generative vs agentic AI

Generative AI creates something in response to your request: a document, summary, image, analysis or piece of code. Agentic AI goes further - it can work toward a goal by choosing and carrying out multiple steps.

Generative AI produces. Agentic AI acts.

See also: *Agent, Agentic workflow, LLM.*

Grounding

Giving AI a reliable source to base its answer on (e.g. your policies, documents, database or current web information) instead of relying only on what the model already "knows". Grounding can make answers more relevant and easier to verify.

See also: *RAG, Hallucination, Single source of truth.*

★ ESSENTIAL FOR BUSINESS

Guardrails

Rules and restrictions around what an AI system is allowed to do. Guardrails might stop an AI from sending something without approval, accessing certain information, giving particular kinds of advice or performing actions outside its role.

For businesses, guardrails are part of using AI safely - rather than simply trusting it to make every decision.

See also: *AI governance, Human oversight, Harness.*

H

Hallucination

When AI produces incorrect or invented information and presents it as if it were true. It can invent facts, references, numbers, people or events.

If the information matters, verify it.

See also: *Grounding, Human oversight, AI literacy.*

Harness

A set of controls, checks and supporting systems around an AI agent that helps keep its actions within safe boundaries. The term appears more often in technical AI environments than everyday business use.

See also: *Guardrails, Agent.*

HITL / HOTL / HOOTL

Three ways of describing how closely people supervise an AI system. **Human In The Loop:** a person approves or participates in important steps. **Human On The Loop:** AI operates independently, but a person monitors it and can intervene. **Human Out Of The Loop:** AI operates with little or no active supervision. For most SMEs introducing AI into important workflows, starting with stronger human involvement is generally safer.

See also: Human oversight, Guardrails, Agent.

★ ESSENTIAL FOR BUSINESS

Human review / human oversight

Having a person check, approve or supervise AI-generated work or decisions. Not every output needs the same level of review: a draft social post might need a quick glance; a customer contract, financial decision or safety-related recommendation may need much stronger oversight.

The higher the consequence of being wrong, the stronger the human review should be.

See also: HITL / HOTL / HOOTL, Guardrails, Responsible AI.

I

IDE *integrated development environment*

Software developers use to write, organise, run and test code. Not something most SME employees need to use, but the term appears increasingly often because AI coding assistants can now work directly inside these environments.

See also: Agentic IDE, Claude Code.

L

Latency

The delay between asking AI to do something and receiving the result. For normal chat, a few extra seconds rarely matter. For a customer-facing AI system responding hundreds of times a day, they might.

See also: Frontier model, Lightweight model.

Lightweight model

A smaller, faster and usually cheaper AI model. Often perfectly suitable for routine tasks such as classification, summarisation or extracting information. Using the biggest model for every job can be like hiring a senior consultant to alphabetise a spreadsheet.

See also: Model, Frontier model, Latency.

LLM *large language model*

The type of AI model behind tools such as ChatGPT, Claude and many other conversational AI systems. It learns patterns from large amounts of information and uses those patterns to understand and generate language.

See also: *Model, Chatbot.*

M

MCP *Model Context Protocol*

A standard designed to help AI systems connect to external tools and information. Think of it as a common way for AI to interact with things such as databases, documents and business applications - rather than building every integration completely from scratch.

See also: *Connector, API, Agent.*

Memory

The ability of an AI product to remember information about you, your preferences or your work across conversations. It can save time, but employees should understand what's being remembered, where it's stored and how it can be removed.



WHERE IT LIVES

ChatGPT: a dedicated "Memory". **Claude:** Projects and memory. **Copilot:** Microsoft 365 context. **Gemini:** personal context.

See also: *Projects, Data privacy.*

Model

The AI engine doing the underlying work. Different models offer different balances of reasoning ability, speed and cost. You usually don't need to know every model name, but it helps to understand that changing the model can change the quality and behaviour of the result.

See also: *LLM, Frontier model, Lightweight model.*

Multimodal / Vision

AI that can work with more than text. A multimodal model might analyse images, documents, charts, audio or video, as well as written instructions.

See also: *AI, Model.*

N

No-code

Building workflows, automations or applications without traditionally programming them. Instead, you usually connect visual building blocks or describe what you want to happen. AI is making the line between "no-code" and traditional development increasingly blurred.

See also: *Automation, Vibe coding.*

Non-determinism

The same AI request can produce slightly different results each time. This is fine for creative tasks. It becomes more important when an AI system is part of a process that needs highly consistent outputs.

See also: *Temperature, Workflow.*

O

Obsidian

A knowledge-management application commonly used to organise linked notes and Markdown files. Useful to some advanced AI users, but not necessary for most SME AI adoption.

See also: *Second brain, Single source of truth.*

Ollama

Software that allows certain AI models to run locally on your own computer. This can provide more control over data and infrastructure, but usually requires greater technical capability than simply using a cloud AI service.

See also: *Open vs closed model, Data privacy.*

OpenRouter

A service providing access to multiple AI models through one interface or API. Mainly relevant to developers and businesses building AI systems, rather than employees simply using AI assistants.

See also: *API, Model.*

Open vs closed model *open-weight*

A closed model is controlled and hosted by its provider. An open-weight model makes the model's underlying weights available, so organisations can potentially run or customise it themselves. For most SMEs, the practical question is less "open vs closed" and more: what gives the right balance of capability, cost, privacy and complexity?

See also: *Ollama, Model.*

P

Pair prompting

Working iteratively with AI rather than trying to create one "perfect prompt". You ask, review what comes back, improve the instruction, challenge the response and keep working together until the result is good. Often a much more realistic way to use AI at work.

See also: *Prompt, Prompt engineering.*

Projects

Persistent workspaces inside AI tools where you keep relevant conversations, instructions and sometimes files together around one subject or piece of work. Useful when you're doing ongoing work rather than one-off questions.



WHAT EACH ONE CALLS IT

Claude: Projects. **Gemini:** Gems. **ChatGPT:** GPTs / Projects. **Copilot:** agents.

See also: *Custom assistant, Single source of truth, Skill.*

Prompt

The instruction or request you give an AI. It can be one sentence or a detailed brief. A good prompt usually tells the AI enough about: what you want, why you need it, the relevant context, any constraints, and what a good result should look like.

See also: *Prompt engineering, System prompt.*

Prompt engineering

The practice of designing instructions and context that help AI produce better results consistently. It isn't about secret wording or clever tricks. Most of it comes down to clear thinking, useful context and precise instructions.

See also: *Prompt, Pair prompting, Zero-shot / few-shot.*

Prompt injection

A security risk where malicious instructions are hidden inside information an AI is processing. For example, an agent might read a document containing instructions designed to make it ignore the company's rules. This matters increasingly as AI moves from answering questions to reading external information and taking actions.

See also: *Guardrails, Agent, System prompt.*

R

RAG *retrieval-augmented generation*

A technique that lets AI find relevant information from your own documents or knowledge base before answering. Instead of expecting the AI to remember your company information, it retrieves the relevant material when it needs it. For SMEs, this is one common way of creating assistants that answer from your actual business knowledge.

See also: *Vector database, Embedding, Grounding.*

Reasoning model

An AI model designed to spend more effort working through complex problems before answering. Useful for planning, analysis, maths and difficult multi-step tasks, but may be unnecessary for simple everyday work.

See also: *Frontier model, Latency.*

★ ESSENTIAL FOR BUSINESS

Responsible AI

Using AI in a way that is safe, fair, transparent and appropriate for the people and business affected by it. For an SME, that usually means practical behaviours such as:

- protecting customer information and checking important outputs
- avoiding misleading AI-generated content
- understanding where bias or errors could cause harm
- telling people when AI use materially affects them
- keeping humans accountable for important decisions

Responsible AI isn't about slowing adoption down. It's about making sure adoption can be trusted.

See also: AI governance, Human oversight, Guardrails.

S

SDD *spec-driven development*

A software-development approach where the requirements and expected behaviour are defined clearly before AI starts building. Primarily relevant to businesses creating software with AI, rather than everyday AI use.

See also: Vibe coding, Vibe engineering.

Second brain

A personal or organisational knowledge system where useful information is stored somewhere searchable, instead of relying entirely on memory. AI makes these systems more valuable because the information can increasingly be searched conversationally.

See also: Single source of truth, Obsidian, RAG.

★ ESSENTIAL FOR BUSINESS

Sensitive information

Information that could cause harm, legal problems, privacy issues or commercial disadvantage if it were exposed or misused. Depending on the business, that could include:

- personal, medical or employee records
- financial details, passwords and customer data
- confidential contracts, trade secrets and commercially sensitive plans

"Sensitive" is broader than simply "secret". Before sharing information with AI, ask what would happen if it left the intended environment.

See also: Data privacy, Enterprise data protection, Shadow AI.

★ ESSENTIAL FOR BUSINESS

Shadow AI

Employees using AI tools for work without the business formally approving, managing or sometimes even knowing about them. For example, someone pasting customer information into a personal AI account because it makes their job faster.

Shadow AI usually isn't caused by bad employees. It often appears when people see value in AI before the business has provided clear tools, training or rules. The answer is rarely to ban AI - it's to give employees an approved path to use it safely.

See also: *AI policy, Enterprise vs consumer AI account, AI governance.*

Single source of truth

One trusted place where the most current and reliable version of information lives. For AI, this matters because the quality of the answer depends heavily on the information it's given. If five versions of a policy exist, the AI has the same problem your employees do: which one is correct?

See also: *Second brain, Grounding, RAG.*

Skill

A packaged set of instructions or capabilities that teaches an AI how to perform a particular task or follow a particular process. The exact meaning varies between products, but the idea is the same: giving AI a repeatable way of doing a job.

See also: *Custom assistant, System prompt.*

Supabase

A cloud database and application-development platform popular in software projects, including some AI applications. Useful terminology for people building systems, but unlikely to matter for most day-to-day SME AI users.

See also: *Vector database, Vibe coding.*

System prompt *system instructions*

The background instructions that tell an AI how it should behave - its role, tone, rules, limits and required processes. Unlike your everyday prompt, these instructions often remain in place across many interactions.

See also: *Prompt, Skill, CLAUDE.md.*

T

Temperature

A technical setting that can affect how predictable or varied an AI model's responses are. Higher settings can produce more variation; lower settings generally produce more consistency. Many everyday AI users never need to change it.

See also: *Non-determinism.*

Terminal *command line, PowerShell*

A text-based way of controlling your computer by typing commands instead of clicking buttons. AI coding tools sometimes use the terminal to perform tasks directly on a computer.

See also: *Claude Code, IDE.*

Token

The small units AI models use to process text. Tokens matter because they affect how much information a model can handle and, for some services, how usage is priced. For everyday business users, you don't need to count them - just remember that very large inputs can use more capacity and potentially cost more.

See also: *Context window, API key.*

U

★ ESSENTIAL FOR BUSINESS

Use case

A specific situation where AI could help accomplish a real business objective. "Use AI in marketing" is too broad. A better use case is: "Turn each customer webinar into a first draft of a LinkedIn post, an email summary and a website article."

Good AI adoption starts with clear use cases rather than choosing a tool first. A strong one answers: who's doing what today, what's difficult about it, and what could improve if AI helped?

See also: *Workflow, Eliminate—Streamline—Automate, AI champion.*

V

Vector database

A type of database designed to search information by meaning rather than exact wording. Commonly used behind RAG and AI knowledge systems. Most SME users never interact with one directly, but it may sit behind a company AI assistant.

See also: *RAG, Embedding.*

Vibe coding

Creating software by describing what you want to an AI and letting it write much of the code. It has dramatically lowered the barrier to building prototypes and simple applications - although production systems still require proper testing and oversight.

See also: *No-code, Vibe engineering.*

Vibe engineering

A more disciplined version of AI-assisted coding. Instead of simply asking AI to "make something", it combines AI development with testing, review, version control and proper engineering practices.

See also: *Vibe coding, SDD.*

VPS *virtual private server*

A remote computer in the cloud that can run continuously, even when your own computer is switched off. Mainly relevant when businesses host applications or AI services themselves.

See also: *Open vs closed model, Ollama.*

W

★ ESSENTIAL FOR BUSINESS

Workflow

The repeatable sequence of steps used to get a piece of work done. For example: customer enquiry → qualify the lead → create CRM record → draft reply → salesperson reviews → send follow-up.

AI is most useful when businesses stop thinking only about isolated prompts and start looking at whole workflows. The better question isn't "how could AI write this email?" but "where could AI improve this entire process?"

See also: *Use case, Automation, Agentic workflow.*

Z

Zero-shot / few-shot

Two ways of giving an AI instructions. **Zero-shot:** you ask it to do something without showing an example. **Few-shot:** you provide one or more examples of what a good answer looks like. Examples can dramatically improve results because they show the AI what you mean rather than making it guess.

See also: *Prompt engineering.*

The terms in this glossary were defined against the state of the industry in 2026. The tool ecosystem moves fast: the concepts last, the tool names not so much.

YOUR PATH TO AI STARTS HERE

Know the words. Now put them to *work*.

Understanding the language is step one. The real value shows up when AI starts saving your team real hours and doing real work - safely, consistently and well. That's what Kameno helps small and medium businesses do: practically, without the hype, and without needing anyone on your team to be technical.



BOOK A FREE AI OPPORTUNITY SCAN

kameno.ai



EMAIL US

contact@kameno.ai



FOLLOW ALONG

linkedin.com/company/kamenoai

Kameno — AI education & enablement for business owners. Your path to AI, built for the way you actually work.

We'll show you where AI fits in **your** business, what to try first, and how to get your team using it safely and well.