AI Glossary

Rosetta Toolkit · Business OS for AI | 31 Essential Terms — Starter Edition

A

Agentic AI / AI Agent

An AI system that perceives its environment, makes decisions, and takes autonomous actions to achieve goals — often across multiple steps and tools.

Why it matters: *The shift from AI that answers to AI that acts. Agents are redefining automation and knowledge work.*

AI Copilot

An AI-powered assistant embedded in business applications to augment human work — drafting, coding, analysing, and advising in real time.

Why it matters: *The primary way most business users experience AI day-to-day (e.g., Microsoft Copilot, GitHub Copilot, Salesforce Einstein).*

AI Governance

The frameworks, policies, and accountability structures that ensure AI systems are developed and used ethically, safely, and in compliance with regulation.

Why it matters: *Increasingly required by regulators (EU AI Act, etc.) and expected by enterprise customers. A board-level concern.*

AI Orchestration

The coordination of multiple AI models, agents, data pipelines, and tools so they work together as a cohesive system.

Why it matters: *Acts as the "conductor" of an AI system. Without orchestration, multi-agent systems collapse into chaos.*

AI Policy

An organisation's documented rules for how AI may and may not be used — covering approved tools, acceptable use cases, data handling, and ethical boundaries.

Why it matters: *The foundation of responsible AI deployment. Without a policy, every team makes up their own rules — and risk compounds silently.*

AI ROI (Return on Investment)

The measurable business value generated from AI investments — productivity gains, cost savings, revenue growth, and efficiency improvements.

Why it matters: *In 2026, AI spending must be justified. Leaders who measure ROI strategically outperform those operating on "vibe-based spending."*

Artificial Intelligence (AI)

Computer systems that simulate human intelligence — reasoning, learning, decision-making, and language understanding.

Why it matters: *The umbrella term for every technology in this glossary. Understanding AI is table stakes for any business leader in 2026.*

Australia's National AI Plan

Released December 2025, governs AI through existing legal frameworks rather than a standalone AI Act. Mandatory guardrails were paused; an AI Safety Institute is being established.

Why it matters: *"No AI Act" doesn't mean no risk — existing privacy, consumer, and employment law all apply.*

B

Benchmark

A standardised evaluation dataset or task used to compare AI model performance — e.g., MMLU for knowledge, HumanEval for coding.

Why it matters: *When vendors claim "state-of-the-art," benchmarks are the scorecard. Essential for model selection and AI procurement decisions.*

C

Context Window

The maximum amount of text (tokens) an LLM can process in a single interaction — its "working memory."

Why it matters: *Larger context windows mean AI can read full documents, long conversations, or entire codebases at once.*

F

Foundation Model

A large-scale AI model pre-trained on massive, diverse data that serves as a versatile base for many specific tasks.

Why it matters: *Eliminates the need to build AI from scratch. Examples: GPT-4, Claude, Gemini, Llama.*

G

Generative AI (GenAI)

AI that creates new content — text, images, code, audio, or video — based on patterns learned from training data.

Why it matters: *The engine behind ChatGPT, Copilot, and most of today's AI productivity tools.*

Grounding

Connecting AI outputs to verified, real-world data sources to prevent hallucinations and ensure factual accuracy.

Why it matters: *The difference between an AI that sounds right and one that is right.*

Guardrails

Safety mechanisms — policies, filters, and technical controls — that prevent AI from producing harmful, biased, or out-of-scope outputs.

Why it matters: *Non-negotiable for enterprise AI deployment. Guardrails are your AI risk management layer in production.*

H

Hallucination

When an AI generates information that sounds plausible but is factually incorrect or entirely fabricated.

Why it matters: *One of the biggest enterprise AI risks. RAG, grounding, and guardrails are the primary mitigations.*

Human-in-the-Loop (HITL)

A design pattern where a human reviews or approves AI decisions at critical points — especially for irreversible, high-risk, or high-value actions.

Why it matters: *The central control mechanism in responsible AI deployment. Determines where autonomy ends and accountability begins.*

J

Jailbreak

Techniques — typically adversarial prompts — used to circumvent an AI model's built-in safety constraints and content policies.

Why it matters: *A live threat for any enterprise deploying AI. Jailbreaks are why guardrails, red teaming, and system prompt hardening are non-negotiable.*

L

Large Language Model (LLM)

A model trained on massive text datasets to understand, generate, and manipulate human language.

Why it matters: *The brain behind chatbots, agents, and AI assistants. GPT-4, Claude, and Gemini are all LLMs.*

M

Machine Learning (ML)

A subset of AI where algorithms learn from data and improve over time without being explicitly programmed.

Why it matters: *Powers forecasting, fraud detection, personalisation, and most modern AI products.*

MCP (Model Context Protocol)

An open standard (Anthropic, Nov 2024) that allows AI models to connect to external tools, databases, and services through a universal interface.

Why it matters: *The "USB-C for AI" — eliminates custom integrations. Now adopted by OpenAI and Google DeepMind.*

Multi-Agent System

An architecture where multiple specialised AI agents collaborate, delegate, and coordinate to complete complex tasks.

Why it matters: *Enables AI to tackle problems too large for a single agent — the foundation of Business OS architectures.*

P

Prompt Engineering

The skill of crafting inputs (prompts) to guide AI models toward accurate, useful, and safe outputs.

Why it matters: *The quality of your prompt directly determines the quality of AI output. A core skill for any AI-enabled team.*

Prompt Injection

A security attack where malicious instructions are embedded in user inputs or external content — causing the AI to ignore its system prompt and follow the attacker's commands instead.

Why it matters: *The #1 AI security risk for agentic systems. Any AI that reads external content (emails, web pages, documents) is a potential target.*

R

RAG (Retrieval-Augmented Generation)

Combines an LLM with external knowledge retrieval so answers are grounded in real, current data rather than just training data.

Why it matters: *Dramatically reduces hallucinations. The most widely deployed enterprise AI pattern.*

Reasoning Model

An LLM specifically designed to "think through" complex, multi-step problems with enhanced logical and analytical capability.

Why it matters: *The breakout category of 2025-26. Models like OpenAI o3 and Gemini 2.5 Pro shift AI from fast answers to deep thinking.*

Red Teaming

Structured adversarial testing where a team intentionally attempts to break, bias, or misuse an AI system to surface failure modes before deployment.

Why it matters: *Required in the EU AI Act for high-risk systems. The responsible way to know your AI's limits before adversaries find them first.*

Responsible AI / AI Ethics

Principles and practices ensuring AI systems are fair, transparent, accountable, and aligned with societal values.

Why it matters: *Shapes public trust and regulatory compliance. Includes bias mitigation, explainability, and human oversight.*

Rosetta Business OS

An AI operating model framework organising AI capability across five layers: Intake/Routing → Orchestration → Strategy → Execution → Automation. The governance spine that coordinates tools, agents, and humans.

Why it matters: *Organisations deploying AI without an operating model create fragmented, ungovernable systems. Rosetta is the connective tissue between AI strategy and delivery.*

S

Small Language Model (SLM)

A compact AI model (typically <30B parameters) optimised for specific tasks — faster, cheaper, and often runnable on-device.

Why it matters: *The pragmatic enterprise choice for domain-specific tasks, privacy-sensitive deployments, and edge use cases.*

T

Token

The basic unit of text that LLMs process — roughly a word or word-fragment. Models are priced and limited by token count.

Why it matters: *Tokens = cost and capacity. Understanding tokens helps optimise AI spend and prompt design.*

Temperature

A parameter controlling how creative or random an AI model's outputs are — higher = more creative, lower = more deterministic.

Why it matters: *Tuning temperature is the difference between imaginative copy and precise factual answers. Core AI literacy for anyone configuring AI tools.*

Full Edition — 34 Additional Terms

AGI · AI Bias · AI Safety · Alignment · Chain-of-Thought Prompting · Data Classification · Deep Learning · Digital Twin · Edge AI · Embedding · EU AI Act · Evals · Explainable AI · Fine-Tuning · Function Calling · Inference · Knowledge Graph · LoRA · Multimodal AI · NLP · Neural Network · Non-Human Identity · Observability · Predictive Analytics · Quantization · Reinforcement Learning · RLHF · RPA · Synthetic Data · System Prompt · TCS Framework · Transformer Architecture · Vector Database · Zero-Shot Learning.