# AI Glossary

SOLUTIONS24X7 · TCS FRAMEWORK | STARTER EDITION

25 essential AI terms for business leaders and technical practitioners — alphabetical, current, and to the point.

## A

### Agentic AI / AI Agent

AI that doesn't just answer — it plans, acts, and gets things done 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 assistant embedded in business apps to augment human work — drafting, analysing, and advising in real time.

**Why it matters:** The primary way most business users experience AI today (Microsoft Copilot, GitHub Copilot, Salesforce Einstein).

### AI Governance

The frameworks, policies, and accountability structures that ensure AI is developed and used ethically, safely, and compliantly.

**Why it matters:** Increasingly demanded by regulators and enterprise customers. A board-level concern in 2026.

### AI Orchestration

The coordination layer that manages how multiple AI models, agents, and tools work together as a cohesive system.

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

### AI Policy

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

**Why it matters:** Without a policy, every team invents their own rules — and risk compounds silently.

### AI ROI (Return on Investment)

The measurable business return on AI investments — productivity gains, cost savings, and revenue growth.

**Why it matters:** Leaders who measure AI ROI strategically consistently 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. Table stakes vocabulary for any business leader in 2026.

### Australia's National AI Plan

Australia's December 2025 framework governing AI through existing laws — not a standalone AI Act. 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 to AI deployment now.

## C

### Context Window

The maximum amount of text 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 pre-trained AI model — GPT-4, Claude, Gemini — used as a versatile base for many specific applications.

**Why it matters:** Eliminates the need to build AI from scratch. Businesses adapt foundation models rather than training new ones.

## 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 or out-of-scope outputs.

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

## H

### Hallucination

When an AI generates information that sounds plausible but is factually wrong 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 or high-risk actions.

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

## 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 systems 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) connecting AI models 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 complex for a single agent — the foundation of Business OS architectures.

## P

### Prompt Engineering

The skill of crafting inputs 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.

## R

### RAG (Retrieval-Augmented Generation)

Combines an LLM with external knowledge retrieval so answers are grounded in your real data, not just training data.

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

### Reasoning Model

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

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

### 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 spanning five functional layers: Intake/Routing → Orchestration → Strategy → Execution → Automation.

**Why it matters:** The connective tissue between AI strategy and AI delivery. Without an operating model, AI deployments remain fragmented and ungovernable.

## S

### Small Language Model (SLM)

A compact AI model (<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.