

THE ENGINEERING VOCABULARY THAT MATTERS

Stop calling everything an “AI agent.”

28 terms. Seven useful distinctions. One field
guide for people who build.

WHO DECIDES?

WHAT CAN IT DO?

WHAT PROVES IT WORKED?

Swipe to learn →

01 / THE STACK

Different layers. Different jobs.

01 **Agent**

A system that selects actions, observes their results, and continues toward an objective.

02 **Tool calling**

The model proposes a named operation and arguments; the runtime validates and executes the operation.

03 **Skill**

A reusable procedure or knowledge package that an agent can load when relevant.

04 **Harness**

The surrounding execution system: model calls, tools, state, context, permissions, and stopping rules.

02 / THE INTENT

Before the agent writes a line.

01 Goal

The outcome the system or agent is trying to achieve.

02 Specification

A precise description of required behavior, interfaces, or properties.

03 Plan

A proposed strategy and sequence for executing the work.

04 Acceptance criteria

Concrete conditions used to accept or reject a feature or deliverable.

03 / THE CHALLENGE

Challenge the work. Verify the claim.

01 Adversary

An actor or process whose objective conflicts with the target system's objective.

02 Adversarial agent

An agent assigned to oppose, challenge, or attack a target within a defined setting.

03 Adversarial loop

An informal pattern of challenge, revision, and rechecking; not one standardized algorithm.

04 Verifier

A component that checks claims or properties against evidence and specified conditions.

04 / THE LOOP

Do not just retry harder.

01 Critic

A component that assesses a candidate and identifies weaknesses or potential improvements.

02 Replanning

Revise the execution plan after new information, failures, or changed constraints.

03 Repair loop

Modify a candidate or strategy in response to diagnosed failure, then try again.

04 Stop condition

An explicit rule for completion, failure, budget exhaustion, cancellation, or escalation.

References [1]

05 / THE BOUNDARIES

Know what the protocol connects.

01 **MCP**

Model Context Protocol: a standard interface for exposing tools, resources, and prompts to AI applications.

02 **A2A**

Agent2Agent: a protocol for communication between independently implemented agentic applications.

03 **ACP**

Agent Client Protocol: standard communication between coding agents and editor or IDE clients.

04 **A2UI**

A declarative format for agent-generated interfaces rendered by compatible clients.

What the model sees changes what it does.

01 Context engineering

Select, structure, and update the information available to the model throughout execution.

02 Compaction

Replace detailed history with a smaller representation that preserves task-relevant information.

03 RAG

Retrieval-augmented generation: fetch external evidence and use it to support a generated response.

04 Structured outputs

Model responses constrained to a supported output schema.

07 / THE EVIDENCE

The agent said “done.” What proves it?

01 Eval

A systematic test of a model or application's behavior against defined criteria.

02 Trace

A structured record of model calls, tool actions, timing, errors, and related execution events.

03 Idempotency

Repeated execution of the same operation has the same intended effect as one execution.

04 Least privilege

Grant only the access and capabilities necessary for a specific task.