

Glossary

Locked one-sentence definitions used across the GROW library. Every artifact in Systems, Execution, Trust, and Coordination uses these terms in this sense. Synonyms are not permitted in artifact text; if a concept is missing, propose an addition rather than coining a local term.

Failure mode — A named, reproducible way the agent can produce a wrong, harmful, or non-compliant outcome, characterized by a detectable signal and an expected correct behavior.

Escalation — Routing a decision out of the agent's authority to a higher-authority human or system because policy, confidence, or impact thresholds have been exceeded.

Fallback — A pre-declared safe alternative behavior the agent executes when its primary path fails a precondition, exceeds a budget, or trips a guardrail.

Override — A human-issued instruction that replaces the agent's chosen action with a different action, recorded as an event with rationale.

Evidence — Any retained artifact (log slice, tool trace, document, retrieval chunk, signed attestation) referenced by a provenance record to support a claim or action.

Decision origin — The source of a single decision in an agent's trace. Base enum (v0.1.0, locked by C2): agent (autonomous), human-override (human replaced the agent), fallback (declared safe alternative fired), escalation (decision routed out to higher authority). Extended values (v0.2.0, additive backward-compatible superset introduced for the Execution/Trust/Coordination clusters): workflow-step (decision taken inside a modeled workflow step — C5), sensor-triggered (decision initiated by a fused telemetry alert — C6), compliance-determination (a rules-based compliance ruling — C9), permission-decision (a zero-trust permission grant or deny — C10), governance-vote (an executed DAO decision — C11). Base-enum consumers built before v0.2.0 remain valid: any record carrying a base value is unchanged, and a base-only consumer that encounters an extended value treats it as agent for control purposes until it declares extension-awareness.

Confidence band — The qualitative trust level assigned to a source or retrieval, drawn from the enum high, medium, low, unknown, used by eval rubrics to weight scoring and by reliability policies to gate action.

Regression — A previously-passing eval case that fails after a change to the agent, its tools, prompts, models, or data; tracked per failure mode.

Drift — A gradual change in agent behavior, input distribution, or source content that degrades performance without a single discrete cause; detected by trend, not by a single failed case.

Severity — The graded impact class of a failure mode, drawn from the enum critical (safety, legal, or irreversible-impact), high (material harm, recoverable), medium (degraded experience), low (minor defect), info (observability-only).

HITL gate — A required human-in-the-loop checkpoint that blocks the agent from proceeding until a human approves, rejects, or modifies the proposed action.

Irreversible-impact boundary — The line in a workflow beyond which an agent action cannot be undone (money moved, message sent, record filed, physical actuation); HITL gates are mandatory before crossing it.

Evaluation gate — A versioned checklist whose pass is a precondition to advancing an artifact, module, or release to the next status (reviewed or shipped).

Provenance record — A timestamped, append-only entry in the provenance store capturing inputs, outputs, decision origin, evidence pointers, and rationale for a single event.

Lineage — The directed chain of provenance records and source references connecting an output back to its originating inputs, retrievals, and decisions.

Source confidence — The confidence band assigned to a specific data source in the provenance source register, used as the C4 contract payload into eval rubrics.

Cluster terms (added v0.2.0 for Execution, Trust, Coordination)

Workflow step — A single unit of a modeled workflow with declared inputs, outputs, owner (agent or human), and branch outcome; the granularity at which Course 4 emits C5 provenance records.

Handoff — The transfer of a task or decision between two owners (agent→human, human→agent, or agent→agent) at a defined gate, carrying a declared payload and acceptance condition.

Process map — The directed model of a workflow's steps, decisions, branches, and handoffs; the Course 4 analogue of the S3 lineage map.

Telemetry source — A sensor, device, API, log, stream, or external feed that emits signals; classified in Course 5 by latency, reliability, frequency, and a confidence_band (reusing the C4 enum).

Normalization — The Course 5 transformation that aligns formats, units, and timestamps and resolves missing, duplicate, or conflicting signals before fusion.

Fusion — The combination of multiple normalized telemetry signals into one operational view, with source conflicts resolved by trust, freshness, and relevance weighting.

Alert — A Course 5 signal-threshold crossing classified warning or critical, linked to an action or fallback; the C6 payload into S1 reliability.

Model tier — A named cost/quality class of model or deterministic logic (premium, standard, cheap, deterministic) that Course 6 routes among; the high-risk-step protection rule forbids demoting tiers on critical/high failure-mode steps without a waiver.

Routing — The Course 6 decision selecting a model tier, tool, cache hit, batch path, or human-review path for a given task, logged via C7.

Latency budget — The Course 6 ceiling on acceptable end-to-end response time for a task class, traded against cost and reliability.

Unit economics — The Course 6 comparison of fully-loaded operating cost per task/user/workflow against the value or revenue it produces, expressed as a gross margin.

Jurisdiction rule — A Course 7 governing requirement (statute, ordinance, policy, form, deadline, or approval flow) attributed to a named agency and classified universal or local.

Expert-review boundary — The Course 7 line at which a determination must be routed to a qualified human (attorney, engineer, compliance officer); crossing it autonomously is gate-failing. Subsumes the "not legal advice" boundary.

Zero-trust — The Course 8 default-deny posture: users, tools, documents, connectors, and model outputs are treated as potentially unsafe and must be verified before being granted action authority.

Least privilege — The Course 8 rule that every principal (agent, tool, role) holds the minimum permissions required for its declared task and nothing more.

Action class — The Course 8 grouping of an operation by blast radius (read, external-write, financial, comms, admin, destructive), each mapped to a control in the tool-use control list.

Tool allow-list — The Course 8 enumerated set of tools an agent may invoke, with per-tool controls (auto, confirm, human-approval, blocked); anything not listed is denied.

Decision rights — The Course 9 explicit assignment, per decision class, of who proposes, who votes, who executes, who may veto, and who is accountable.

Quorum — The Course 9 minimum participation threshold for a vote to be valid, distinct from the approval threshold required to pass.

Treasury — The Course 9 pool of funds whose allocation is governed by decision rights, budgets, milestone gates, and retroactive-reward rules.

Proposal — The Course 9 unit of governance action: a structured request carrying evidence, a review period, and conflict-of-interest disclosures, resolved by a vote and executed under accountability tracking.

ICP (Ideal Customer Profile) — The Course 10 definition of the target user, economic buyer, budget owner, urgency, alternatives, and switching friction for a commercialized capability.

Productization — The Course 10 conversion of a one-off or consulting-dependent capability into a repeatable offer with standard workflows, packages, and interfaces.

Moat / defensibility — The Course 10 durable advantage protecting a commercialized offer: proprietary workflow, data advantage, distribution, integrations, compliance moat, or operational expertise.

Proof-of-value (PoV) — The Course 10 measurable outcome a pilot must demonstrate to justify purchase, expressed in metrics drawn from the S2 scoring system.

Change log

- 0.2.1 — 2026-05-29. Resolved the mutual-downstream contradiction with shared-artifact-schema directionally: glossary consumes the schema, so moved shared-artifact-schema from downstream_consumers to upstream_deps. The edge is now schema→glossary (symmetric, acyclic).
- 0.2.0 — 2026-05-29. Extended decision_origin with the five additive cluster values (workflow-step, sensor-triggered, compliance-determination, permission-decision, governance-vote) as a backward-compatible superset. Added 23 cluster terms for Execution (Courses 4–6), Trust (Courses 7–8), and Coordination (Courses 9–10). No base-enum semantics changed.
- 0.1.0 — 2026-05-28. Initial Systems-cluster locked definitions.