

AIPM — Glossary & Knowledge Base

Reference lexicon of the AI Project & Production Management Framework — DYONYSOS / Julien Daures. Companion to the Standard (AIPM_Standard_EN.docx): each term points to the Standard section(s) where it is defined in full.

A

AI agent (agent) — Software system that carries out a task autonomously or semi-autonomously from a language model and instructions (prompt), with or without access to external tools. In AIPM, every agent is tied to a Directive (see *Directive*), to a named Director-Verifier, and to an earned autonomy scope (principle 4) rather than to a fixed maturity level (see *Maturity model*). *Standard §3.2, §5, §8.*

Producer-Agent — In the AI-RACI matrix (see *AI-RACI*), the agent (or chain of agents) responsible for producing a Directive's result, generally running at full autonomy (**Auto**) in the normal course, but always under supervision and escalation thresholds. Structurally distinct from the Verifier-Agent once volume or stakes justify it (principle 11). *Standard §5.2, §5.3.*

Verifier-Agent — An agent structurally distinct from the Producer-Agent, tasked with automatically verifying the result produced — never the same agent grading its own work (principle 11). May also compute a type-4 trigger threshold (see *Trigger threshold*). *Standard §5.2, §6.2, principle 11.*

AI Orchestrator — Former role title from AIPM's early versions, superseded by the **Director-Verifier** (see that entry): the named human who frames the intent, sizes the agent fleet, and makes the final verification call. *Standard §5.1.*

Andon — Borrowed from the Toyota Production System: the mechanism by which an agent explicitly signals that it doesn't know or has detected an anomaly, rather than producing a plausible but false answer. An agent that "makes something up" instead of stopping is an Andon failure.

Automation backlog — Prioritized register of tasks that are candidates for delegation to an agent, each with an estimated gain and risk. Template provided in [AIPM_Templates_EN.xlsx](#).

C

Control loop — The complete chain: data source → ingestion → decision (rule and/or model) → action → verification → logging → feedback into learning. Central concept of the "Delivery & production" domain and of the Dark Factory/AI framework. *Standard §3.5.*

D

Dark Factory — Design framework for autonomy: a process runs without human intervention under normal conditions, with humans stepping in only on exception, on an irreversible decision, or to improve the system.

Definition of Done — AI (DoD-AI) — AIPM's extension of the classic "definition of done": criteria an agent-produced output must meet before it is considered usable (human review completed, verifiable source, no sensitive data present, conforms to the expected format). Overlaps with the Directive Sheet's verification method and with the "first-pass verified result rate" metric. *Standard §3.2, §9.*

Directive — AIPM's unit of work since the framework's redesign: a verifiable result with no presupposed duration — replaces the "project" (fixed scope and duration) as the basic planning building block. Closes when the result is reached and verified, never on a calendar. Formalized through a 5-field mandatory Directive Sheet (see *Directive Sheet*). *Standard §3.1, §3.2, §3.3.*

Directive lifecycle — Event-driven, not calendar-driven: a Directive opens at the framing of the intent (delegability filter), proceeds through delivery orchestrated by trigger thresholds, and closes at the verification of the result — with no "monthly phase review" type milestone. Replaces the earlier 7-stage calendar lifecycle from AIPM's early versions (Framing, Design, Planning, Delivery, Monitoring, Closure, Continuous learning). *Standard §3.4.*

Directive Sheet — 5-field mandatory template that formalizes a Directive: (1) verifiable result + verification method, (2) named Director-Verifier, (3) assigned agent fleet and autonomy scope, (4) escalation thresholds, (5) guaranteed human escalation channel (see *Guaranteed human escalation channel*) — the latter cannot be disabled (principle 3). *Standard §3.2, §10.*

Director-Verifier (human) — The named human accountable for framing the intent, for the final verification of the result, and for arbitrating any out-of-scope case raised by an escalation threshold, for a given Directive. Replaces the earlier "AI Orchestrator" role from AIPM's early versions: no longer an execution role — the role that demands the most judgment, not less. *Standard §5.1, §5.2 (AI-RACI).*

Drift — Gradual degradation of an agent's or model's output quality over time, caused by a shift in input data or in the underlying model. A metric to watch under "Measurement & continuous improvement."

E

Escalation — Explicit rule defining at what deviation, level of uncertainty, or decision type a task must be routed to a named human rather than handled by an agent. Triggered by a **trigger threshold** (see that entry), never by a calendar ritual. *Standard §6.*

Exception table — Replaces the calendar status update: consulted on demand, it documents, for each triggered threshold, the named human owner and the expected time to resume — not just the triggering condition. The human checkpoint happens only if an exception is logged in it. Template provided. *Standard §6.3, §10.*

F

Fupan (复盘) — Post-action review ritual borrowed from Chinese industrialization methods: original goal restated, actual result, gap and its causes, the rule to keep. Used in AIPM as a lightweight retrospective variant.

G

Guaranteed human escalation channel — The 5th mandatory field of the Directive Sheet (see *Directive Sheet*): the concrete means by which a human can take back control if needed. Can never be filled in as "none" (principle 3). *Standard §2 (principle 3), §3.2.*

Guardrail — Hard limit an agent or automation may never cross (spend cap, volume cap, maximum error rate, forbidden scope of action), with a defined behavior for when the limit is reached. *Standard §6.2 (novelty threshold), §7.*

AI governance (governance by design) — AIPM domain covering AI risk registers, guardrails, trigger thresholds, and systematic human review of high-risk outputs — built into execution itself, not bolted on as a periodic audit after the fact (principle 5). *Standard §4.6, §6, §7.*

H

Human-in-the-loop — Principle stating that any agent-produced output above a certain criticality threshold passes through human validation before use. Embodied in AIPM by founding principle 1 of 11: "the decision stays human, the execution no longer is" — it is no longer the human who executes and the AI that assists, it is the reverse, but the decision and the accountability are never delegated. *Standard §2.*

Hallucination — A model producing plausible but false or unverifiable information, presented as fact. A front-line risk in AIPM's AI risk register.

J

Jidoka — Toyota principle adapted to AIPM: a system's ability to stop itself automatically rather than propagate a false or inconsistent data point to the next step.

K

AIPM KPI (native metrics) — AIPM's 4 own metrics (distinct from classic progress indicators), always reported against a volume or demand base, never presented raw (principle 10): Directive velocity, human leverage ratio, calibration accuracy (over-confidence / under-confidence rates), first-pass verified result rate. Dashboard template provided. *Standard §9, §10.*

L

Level (autonomy level, L0-L5) — superseded notion — Former staged delegation scale (L0 Manual → L5 Autonomous-governed) from AIPM's early versions. AIPM no longer uses a staged scale: that kind of scale measures the wrong question ("how much is automated" rather than "is the delegation well calibrated"). Superseded by the two-axis maturity model (see *Maturity model*) and by an autonomy scope earned case by case, Directive by Directive (principle 4). *Standard §1.3, §8.1.*

M

Maturity model — Two-axis model that replaces the earlier staged L0-L5 scale (see *Level*): over-confidence rate (share of real incidents that pass without escalation) on the x-axis, under-confidence rate (share of corrective human interventions on results the agent had already produced correctly) on the y-axis. Four quadrants: Calibrated, Over-confident, Under-leveraged, Chaotic — each with its own action plan. The question is no longer "how much is automated" but "is the delegation well calibrated." Grid and scorecard provided in the templates. *Standard §8, §10.*

O

Orchestration — The technical rules (scheduling, idempotency, error recovery, concurrency limits) that ensure a chain of agents and automations runs reliably and repeatably.

P

Poka-yoke — Error-proofing device designed to make an error impossible rather than merely detectable after the fact; adapted in AIPM as strict schema validation on every agent's inputs and outputs.

AIPM Principles (the 11) — The framework's non-negotiable foundations, each illustrated with real, sourced and dated cases (see the case bank in the scoping document): (1) the decision stays human, the execution no longer is; (2) value before automation; (3) the human escalation channel is a field that cannot be disabled; (4) autonomy is earned, never granted by default; (5) governance by design, never after the incident; (6) adaptation to context, never a one-size-fits-all formalism; (7) iteration through measured evidence, never through a promise; (8) steering the whole system — humans, agents, data, tools; (9) trust through traceability and auditability; (10) metrics are always reported against a volume base, never raw; (11) the verifier is distinct from the producer. *Standard §2.*

R

AI-RACI — Extended responsibility matrix (see *Directive Sheet*) that crosses each Directive activity with four actors: Director-Verifier (human), Producer-Agent, Verifier-Agent, Stakeholders. Codes: R = Responsible (executes), A = Accountable (reports, decides), C = Consulted, I = Informed, and **Auto** = executed by an agent without prior human validation in the normal course (but always under supervision and escalation thresholds). The Director-Verifier is always at least A/R on framing, the final decision, and escalation. Template provided. *Standard §5.2, §10.*

AI risk register — Dedicated register (distinct from the classic project risk register) listing risks specific to agent and model use, organized into 8 base categories: hallucination, data leakage & confidentiality, bias & fairness, drift, over-reliance, security & cyber, regulatory non-compliance & intellectual property, reputational — complemented by a category of risks specific to agentic scale (simultaneous loss of execution and human control, breakdown at the point of handoff, collective drift of an agent fleet, etc.). Template provided; full treatment in the *Governance & AI Risk Compendium*. *Standard §7.1-§7.8, §7.9.*

S

Sampling — Practice of running human quality control on a sample of an agent's outputs rather than on the whole, once the agent's earned autonomy scope (principle 4) is wide enough and its escalation thresholds (§6.2) have been proven — rather than at a fixed maturity level, a notion now retired (see *Level*).

T

Trigger threshold — An event-based condition that triggers a human checkpoint, replacing calendar milestones (principle 6, §6.1). Four types: risk threshold (an AI risk indicator crosses a ceiling set at design time), confidence threshold (the agent's confidence in its own result drops below a defined floor), novelty threshold (the situation is not covered by existing guardrails), and a threshold computed by a verifier-agent structurally distinct from the producer-agent (principle 11). *Standard §6.2.*

Traceability — The ability, for any automatically produced decision or output, to reconstruct the input, the prompt/model/rule version used, the output, the timestamp, and the effect produced. AIPM's founding principle 9 of 11 ("trust through traceability and auditability"). *Standard §2.*

This glossary is a living document: every new AIPM template, case study, or training module can add terms to it. To contribute: add an entry in alphabetical order, with a pointer to the relevant Standard section where applicable.

AIPM Framework v1.0 — DYONYSOS, Julien Daures — 2026