



AI GOVERNANCE + DECISION PROVENANCE

Same Authority, New Participant

AI governance, read from the operator's seat.

Toby Benenson, IAPP Certified AI Governance Professional (AIGP) · Two-time COO



“Your overnight hangar space is confirmed. We’ll have everything ready for your arrival.”



What did that person just approve?

- The wording?
- A space allocation?
- A promise to the customer?

We have not established which. We have established that someone clicked a button.



One message, four steps.

- Request received.
- Allocation decided.
- Message prepared.
- Message sent.

Which role owns each step?



The six-level AI participation taxonomy.

01	Prohibited	AI may not participate in this decision class.
02	Reference	AI may surface information that informs the decision. The human reaches the decision without an AI recommendation.
03	Flag	AI may flag conditions that warrant human attention. The human decides whether to act.
04	Draft	AI may produce a draft of the decision artifact. The human reviews, modifies, and approves.
05	Recommend	AI may produce a recommendation with reasoning. The human reviews and decides whether to follow.
06	Automate Under Supervision	AI may execute within bounded conditions. The human supervises and retains override authority.

A working taxonomy for bounding participation. Not an FAA classification, and not a ladder to climb.



Two questions, every time.

- One: what is the AI actually doing in this particular step?
- Two: what would you authorize here, under what conditions?
- One workflow can contain several levels. Classify each function.



A level is not a risk score.

- A Level 2 lookup can influence a consequential maintenance decision.
- A Level 6 receipt can be right for a bounded, low-consequence acknowledgement.
- Ask separately: what could go wrong, who is affected, what evidence, what safeguards.



You vote. Name the assumption.

- ❑ Grab your phone: SAYFLIGHT.AI/VOTE. No phone? Hands count the same.
- ❑ Give me the level, the action, and the assumption your answer depends on.
- ❑ A second hand means insufficient information. Name the information.
- ❑ First answers are working answers. The facts are deliberately incomplete.





The board says REQUESTED. The draft says confirmed.

- ❑ A designated facilities role allocates space. The assistant drafts messages; it cannot allocate or send.
- ❑ The reviewer may communicate confirmed allocations, not make them.
- ❑ The draft still says: your overnight hangar space is confirmed.

Vote: what is the AI actually doing?

01 PROHIBITED 02 REFERENCE 03 FLAG 04 DRAFT 05 RECOMMEND 06 AUTOMATE UNDER SUPERVISION



Now it recommends who gets the last space.

- The AI compares competing requests and recommends one aircraft, with reasoning.
- The authorized human still makes the allocation.

Vote: what participation did we just add?

01 PROHIBITED 02 REFERENCE 03 FLAG 04 DRAFT 05 RECOMMEND 06 AUTOMATE UNDER SUPERVISION



Intake, queue item, acknowledgment. No commitments.

- The workflow extracts the request, creates a queue item, and sends an acknowledgment.
- Permitted fields, data access, logging, and escalation are defined. An on-call supervisor can stop it.

As bounded, what level is this?

01 PROHIBITED 02 REFERENCE 03 FLAG 04 DRAFT 05 RECOMMEND 06 AUTOMATE UNDER SUPERVISION



An update starts quoting price and reserving inventory.

- The old permission did not expand because the product gained a feature.
- Stop test: the queue stops accepting items. Should acknowledgments keep promising the request is being handled?

Authorized Level 6, or execution beyond the boundary?

01 PROHIBITED

02 REFERENCE

03 FLAG

04 DRAFT

05 RECOMMEND

06 AUTOMATE UNDER SUPERVISION



“A human reviews it” has five parts.

See

The reviewer can access the sources, the output, its limitations, and what is missing.

Understand

They can recognize a materially wrong result and know what the tool was asked to do.

Decide

They hold the authority to accept, reject, or escalate this decision.

Have time

The workload and timing allow meaningful verification.

Intervene

Their action can stop or correct the process, and its consequences are understood.



AI amplifies the operation you actually have.

- From the outside, an operation can look assembled. Inside, authority, boundaries, and escalation often live in heads.
- Good people end up operating as a weak machine. AI applied to that magnifies it: faster bottlenecks, not faster solutions.
- The work is making decisions legible, so AI has something worth amplifying.



Frameworks set expectations. We make them concrete.

- NIST's voluntary framework: Govern, Map, Measure, Manage. It already addresses roles, oversight, and disengagement.
- The EU Act separates provider and deployer duties, where its provisions apply. Not every tool is high-risk AI.
- The practical question: what must our operation establish that a vendor cannot establish for us?



Permission and evidence.

Decision

Whether to communicate an overnight hangar allocation as confirmed.

Accountable owner

M. Chen, customer-service reviewer, designated under this fictional operator's policy COMM-04.

Permitted participation

Draft an acknowledgment; draft a confirmation only from a recorded allocation. No allocating, no sending.

Information available

Request R-17; booking board revision 8 at 14:10, status REQUESTED. No allocation attached.

AI contribution

The draft stated space was confirmed. Output retained, tool and version available.



What the reviewer did, and why.

Disposition

Reviewer M. Chen rejected the confirmation language; routed the allocation question to facilities (J. Park).

Recorded reason

“We have the request, but no allocation. I can’t confirm space on that evidence.”

Action and follow-up

Acknowledgment sent 14:14. Allocation pending, named contact assigned, review due 14:30.

Later amendment

14:24: J. Park, the authorized facilities role, recorded the allocation; confirmation then sent. The earlier entry stays visible.

Someone coming on shift can tell what was requested, drafted, rejected, and eventually authorized.



“No outstanding requests found.”

What would you need to know before treating that as an all-clear?



Write one usable boundary.

- For [decision], AI may [participation] using [sources and conditions]. [Role] retains [the call], and [record] captures the AI contribution and human disposition.
- If [missing input, limit, or failure], action stops, [role] takes over, and [evidence] is required before resumption.
- Neighbor test: could someone build and test what you wrote without guessing who decides?



SAYFLIGHT.AI/PREFLIGHT



Questions.

Bring one workflow and one unresolved boundary.



You can only automate what you can clearly articulate.



AI Governance + Decision Provenance



WWW.SAYFLIGHT.AI/PREFLIGHT

TOBY@SAYFLIGHT.AI

Appendix



“Built with AI” is not “allowed to act.”

Built with AI

AI assisted development. Open: does AI operate in the resulting application at all?

Contains AI

AI performs a function during use. Open: which decisions it influences, which actions it can take.

Authorized to act

Someone granted a defined permission. Open: is that grant valid, controlled, and evidenced?

Ask what runs, what it can change, and who authorized that use.



A voice note becomes a proposed entry.

- A qualified technician dictates work actually performed. The AI drafts the maintenance entry.
- The person authorized to approve the record checks it against the work before signing.

What is the AI doing, and what must review verify?

01 PROHIBITED 02 REFERENCE 03 FLAG 04 DRAFT 05 RECOMMEND 06 AUTOMATE UNDER SUPERVISION



The draft adds: operational check satisfactory.

- Nobody dictated that result. An unsupported addition inside drafting is not a recommendation.
- A separate feature proposing a corrective action with reasoning: that is Level 5, alongside the Level 4 entry.

Can the reviewer see what was transcribed and what was added?

01 PROHIBITED 02 REFERENCE 03 FLAG 04 DRAFT 05 RECOMMEND 06 AUTOMATE UNDER SUPERVISION



We rehearse engine-outs. Rehearse this.

- Hide the draft. From the request and the evidence alone: what would you communicate?
- Reveal the confident confirmation. Which source lets you challenge it?
- Who takes the unresolved decision, and what must happen before automation resumes?