

AIPM “AI Project of the Year” Award: Business Plan and Proposal for the Founders’ Forum

For the founders’ forum (Ricardo Vargas and Antonio Nieto-Rodriguez). This is the **AIPM Ambassadors’ collective recommendation**, brought to the forum to decide. It is the third document in the founders’ sequence: (1) the benchmark of 25 award programs (award-benchmark.md), (2) the design concept (concept-proposal.md), (3) **this business proposal**, then (4) the forum’s ratification. It does not replace those documents, it packages them into one decision-ready proposal: the benchmark in digest form (with a pointer to the full file) and the ambassadors’ design recommendations, foregrounding the calls that most need the founders. Most of the underlying award design is attributed to **Farhad Abdollahyan** (AIPM “AI in Real Projects” subgroup); the recommendation as a whole is the ambassadors’ collective position, not any one ambassador’s.

Nothing here is decided. All nine design decisions are recommendations for the forum to ratify. Where the concept proposal labels an option “Locked”, it means only “the option the ambassadors recommend carrying to the forum”, never a settled decision. House style: English, no em-dash or en-dash; no [VERIFY] markers, but every benchmark fact stays faithful to award-benchmark.md.

1. Executive summary

There is a white space in the awards field that, today, no one owns. The ambassadors benchmarked 25 award programs across four lenses: PM bodies, AI and innovation awards, community programs, and hackathons. Not one does both of two things together: recognize **responsible, real-world AI applied inside delivered project work** (not a model, a tool, a 24-hour prototype, or a generic business-excellence story), and return **criteria-anchored feedback that builds the practitioner’s capability** rather than only handing out a trophy. The PM bodies judge delivery but ignore AI; the AI awards reward products against generic business-excellence language; the hackathons judge greenfield builds under time pressure. The intersection (responsible AI-in-delivery plus capability-building feedback, free and online) is empty. That intersection is the position AIPM is best placed to hold.

What we propose: the **AIPM “AI Project of the Year” Award**, an AIPM thought-leadership initiative recognizing AI-enabled projects that demonstrate responsible innovation, real delivery, measurable value, human-centered adoption, and positive impact. One overall winner and two special recognitions. Edition 1 is judged in two tiers, regional then global, across three macro-regions, with a panel of about seven at each tier. Roughly six to nine regional winners reach the global round, a touch above the founders’ single-jury sketch of about five; the forum can tune that count through the per-region advancement number. It runs free end-to-end, online, on a published rubric, with a written feedback report to every finalist. It also generates a pipeline of publishable community case studies that feed the parallel case-study repository and the webinar cadence (about six per year).

Year-one cash cost is near zero: a form, a video link, a shared scoring sheet, the existing Credly badge designs, and a volunteer jury; the real input is volunteer time and the founders' reach. The owner is the Community and Growth stream (lead Eugene De Wee).

The ask: two decisions. First, ratify the proposed design (the nine decisions in section 6). Second, greenlight the run: a July pre-announcement, nominations from mid-August, regional-to-global judging October to December, and an online first ceremony in December. Four items most need the founders' call, flagged throughout and tiered in the decision agenda: the **sponsorship posture**, the optional **Community Choice vote**, the **18-versus-24-month eligibility window**, and the person-level **"AI Project Manager of the Year"** option. Nothing here is decided until the forum decides it.

2. Strategic rationale

Why this white space

Map the field and a hole appears at the center. Down one axis sits everything that judges projects well: PMI, PMOGA, APM, IPMA, and the chapters publish serious delivery rubrics and crown a project of the year, and none say a word about AI. Down the other axis sits everything that celebrates AI: Stevie, BIG, AI Breakthrough, the Webby, the National AI Awards, ITU, and the Australian AI Awards. All of them reward AI products against generic business-excellence language, not against whether AI was responsibly governed and actually adopted inside delivered work. Off to the side sit the hackathons, judging greenfield builds under time pressure, the opposite of deployed practice with measured outcomes. The center, where responsible AI-in-delivery meets capability-building feedback, free and online, belongs to nobody (benchmark sec 4, sec 4.3).

That center is the ground AIPM is best placed to own. AIPM is a global community of practitioners applying AI inside real projects: the award is not a stretch into a new domain, it recognizes what the community already does. AIPM also holds assets no competitor has: an ambassador network as a built-in distribution engine, plus finished Credly badge designs to recognize winners at zero marginal cost. Only IPMA returns a written feedback report at all, and it is paid, on-site, and high-cost (EUR 600 to 4,000, benchmark sec 3.21). That combination, free and online, is the wedge.

Why now

The wedge is real but time-limited. The Australian AI Awards have already shown that "Responsible AI" and "Agentic AI" are viable categories (benchmark sec 3.16). They have not yet been combined with project-delivery rigor and a feedback loop, but the framing is in the open, and category framings get copied. The founders' instinct to "launch before competitors in 2026" is the correct read of the benchmark, not an arbitrary deadline: every quarter of delay is a quarter in which a PM body could bolt an AI category onto its existing machinery, or an AI award could add a delivery rubric, and the wedge narrows. Waiting forfeits the first-mover claim on the exact ground AIPM is best placed to own.

The four guardrails (hard constraints)

The design is written against four non-negotiable constraints that keep the award credible and original:

- **Original.** It does not look like a PMI copy. AIPM borrows mechanics (open nomination, video submissions, published weighted rubrics, multi-stage elimination) but never PMI's branding, naming, or visual identity.
- **Non-commercial.** Free end-to-end: no entry fee and no paywalled proof. The badge, certificate, and logo use are bundled into the free win.
- **Practice-grounded.** It recognizes real AI application in delivered projects, not ideas, models, or tools.
- **Community-first.** The lasting value is the practice the award teaches the community, which is why the feedback report and the case-study pipeline are central, not optional.

The fixed founder parameters fit them directly: launch October to December 2026, first ceremony online, open to outside nominations, jury of 5 to 10, about 5 finalists narrowing to 1 winner.

3. Benchmark digest

This is a digest. The full 25-program field, all profiles, the comparison matrix, and the sourcing live in award-benchmark.md; this section distills its synthesis (sec 4) and picks (sec 5). No profile is inlined here; drill into the benchmark for depth.

The field in brief (four lenses, 25 programs). - **A. Brazilian PM chapters** (PMI-SP, PMI-RIO, PMI-DF): the best public editais with transparent rubrics. - **B. International PM bodies** (PMOGA, PMI Project of the Year, APM, House of PMO, the PMI UK, Netherlands, Sweden and Madrid chapters, IPMA): rigorous delivery judging, but silent on AI. - **C. AI / innovation / community awards** (Stevie, BIG, AI Breakthrough, GitHub Stars, Webby, the National AI Awards, ITU AI for Good, the Australian AI Awards, Clube da Programação): AI-aware, but judged against generic business-excellence language. - **D. Hackathons and commercial recognition** (CloudFest, Hack-Nation, Google Cloud's ADK hackathon, Global Recognition Awards): they judge greenfield builds or, in the commercial case, sell the proof of winning.

The white space. Across all 25, no program combines responsible AI-in-delivery recognition with capability-building feedback, free and online. Only IPMA returns a written feedback report, and it is paid and on-site (see section 2); that gap is the wedge (benchmark sec 4.3).

Best practices to adopt (the zero-cost, copyable mechanics, ranked by copyability in benchmark sec 4.1; item 10 is a later-edition option, not a year-one default):

1. Publish a real rubric with explicit, weighted criteria: the cheapest trust-builder, and an immediate advantage over PMI, which declares its weights proprietary.
2. Run a two-stage, video-first online process (written or evidence assessment, then a live finalist pitch), the proven low-cost shape matching the founders' "about 5 finalists plus online ceremony" sketch.
3. Keep it free end-to-end and say so. Free entry is table stakes; free proof is the wedge, with badge, certificate and logo use bundled into the win.
4. Publish a named jury and an explicit conflict-of-interest rule.
5. State an anti-AI-ghostwriting rule, necessary and a credibility signal for an AI award.
6. Reward durable change and real adoption over model novelty (the scoring philosophy).
7. Return a written feedback report to every finalist (the IPMA lever, the one genuinely new mechanic and the most defensible community-first differentiator).
8. If sponsors are taken, decouple them from winner cash via the CloudFest sponsor-to-charity model.
9. Give winners career capital, not money, and reuse existing assets (the Credly badges).
10. Consider a fairness or size-track split and a lightweight joint-entry mechanic (a later-edition option, not a year-one default).

Pitfalls avoided (benchmark sec 4.2). No published criteria or anonymous jury (reads as a private club, House of PMO, AI Breakthrough); free-to-enter but pay-for-the-proof (the Global Recognition Awards credential-mill model, USD 99 / USD 499 packs and a paywalled logo licence, the sharpest negative benchmark); statistical-jargon credibility theatre; cash prizes and a greenfield-demo bias (the hackathon traps); public voting without fraud controls (vote only on a jury-vetted shortlist, never on raw nominations); scope creep into many categories in year one (APM and Stevie run many; the Australian AI Awards about 33); and continuity neglect (the benchmark's recurring failure, with PMI Netherlands' POYA anchored to a 2022 cycle and PMI Madrid's Proyecto del Año with full bases only for 2012, so commit up front to publishing each year's winner, jury and rubric).

The three reference picks (benchmark sec 5): PMI-DF Prêmio Candango (best edital template), PMI PMO of the Year / PMOGA (closest operational analog to the founders' sketch), Clube da Programação (best lightweight community model), with IPMA the honorable mention contributing the feedback-report lever and CloudFest contributing the sponsor-to-charity model.

4. Proposed design (the ambassadors' recommendations)

All design below is the ambassadors' recommendation for the forum to ratify, with the alternative where one exists; full derivation, with benchmark cross-links, is in `concept-proposal.md`. Where the concept proposal labels an option "Locked", it means the option the ambassadors recommend carrying to the forum, not a decision already made. The four items most needing the forum's call are flagged inline and tiered in section 6.

Positioning (Farhad). Recognize AI-enabled projects that demonstrate responsible innovation, strong delivery, measurable value, human-centered adoption, and positive organizational or societal impact; frame the program as an AIPM thought-leadership initiative that builds a pipeline of publishable community case studies, not a vendor showcase and not a credential to be purchased.

Categories. One overall **AI Project of the Year** plus two special recognitions, **Excellence in Responsible AI** and **Excellence in Agentic AI**. The specials are **not separate tracks**: there is a single submission funnel and a single rubric, and the jury designates the two specials from the finalist pool during final judging (the agentic recognition only where an entry uses agentic AI, otherwise simply not awarded in edition 1). This adds zero submission burden and minimal jury load, stays inside the year-one scope-creep guardrail, and uses the two axes the Australian AI Awards already validated. Farhad’s full five-category split is deferred to edition 2.

Eligibility and window. The eligibility rules: AI is a meaningful component beyond proof-of-concept (completed or operationalized in a real setting); the award is global from day one; publishable-case-study consent is required at entry; joint vendor-plus-practitioner entries are allowed, with the spotlight on the practitioner leading the adoption; and the award is vendor-neutral (no bonus for any particular technology). Completion window: the ambassadors recommend **18 months** at the nomination deadline (aligned with the PMI norm, keeps the first field focused), with **24 months** as the alternative if the founders want a wider first field. One of the four forum items.

Nomination model. Self-nomination and open peer nomination, plus joint entries; a single named **Lead Nominator** owns each entry; a draft-then-final flow; and a lightweight first gate (a summary of at most 180 characters) so screening can triage before the full evidence pack is assembled.

Scoring rubric and weights. A hybrid grid: **six weighted criteria, each scored 1 to 10, with the weights published** before the call opens (out-transparenting PMI, the single cheapest trust-builder).

#	Criterion	Weight
1	Outcomes, benefits and impact	25%
2	Responsible AI governance and Human-in-Command	20%
3	Adoption and change management	15%
4	AI solution fit and innovation	15%
5	Project management and delivery rigor	15%
6	Learning and	10%

#	Criterion	Weight
	contribution to the community	
	Total	100%

The design principle, stated explicitly: **reward durable change and real adoption over model novelty**. Criteria 1, 2, and 3 carry 60% of the weight combined, so a technically impressive AI solution cannot win without credible evidence of outcomes, responsible governance, and real adoption. The rubric collapses Farhad's seven-criterion draft onto six while keeping his two top weights (outcomes 25%, responsible AI 20%) intact; the 60% concentration on criteria 1 to 3 is the ambassadors' reweighting, since Farhad's draft placed less combined weight on those three. Shared rating anchors (1 to 2 minimal or unsupported, through 9 to 10 exceptional) apply to every criterion, and the full rubric is published before the call opens.

Jury and governance. Each panel (the regional panels and the global jury) is mixed (not all-academic) and about seven, within the founders' 5-to-10 range, drawn from AIPM ambassadors plus external experts and covering project management, AI governance, business value, ethics and risk, and sector representation. Calibrate before scoring. Integrity rules to publish, all zero-cost and each beating a benchmark on credibility: conflict-of-interest auto-decline plus disclose-and-recuse (no juror judges their own entry or employer, essential because an ambassador's own project may enter), an anti-AI-ghostwriting rule, a no-contact-with-judges rule, a published sponsorship-does-not-influence-judging firewall if any sponsor is taken, and a requirement that judges distinguish claims from evidence. The exact jury size, its independence, and the COI policy are published.

Geographic judging structure (Decision #9). Edition 1 is judged in **two tiers** so the ambassador network is the judging engine. **Tier 1 (regional, October):** entries are judged within their region by that region's AIPM ambassadors plus criterion-eligible external jurors, and each region promotes 2 or 3 entries to the global round only if they clear a published minimum cut-off (no filling of slots, so a thin first field never advances weak entries). **Tier 2 (global, November or December):** the regional winners compete, judged by Ricardo Vargas, Antonio Nieto-Rodriguez, and the ambassadors who passed the judge-selection criterion. If a sponsor is taken, a sponsor juror from a technology company or a VC joins the bench. That is credibility through the bench, not a payment: the sponsor juror sits under a published sponsorship-does-not-influence-judging firewall and under the same conflict-of-interest rules as every other juror, and no cash reaches winners, consistent with the non-commercial guardrail. Year one runs **3 macro-regions** (Americas, EMEA, Asia-Pacific) rather than the full seven or eight, because the roughly 33-ambassador roster is too skewed to staff more in year one. The full regional split is the edition-2 growth path. Ambassadors judge by affinity, not only residence, to deepen thin benches, and the panels calibrate against a shared example before scoring so a score means the same across regions. **Community vote (needs the forum's sign-off).** The ambassadors recommend that the main award and the two special recognitions be judged **jury-only**, which protects the award from becoming a popularity contest. The optional addition is a secondary

Community Choice recognition, voted by the public **only on the jury-vetted finalist shortlist** (never on raw nominations), with basic fraud controls (one vote per verified account, audited tabulation) and clearly labeled as secondary. Two clean choices for the forum: **(a)** include it as a light secondary recognition (recommended if the forum wants community participation and ceremony energy), or **(b)** defer all public voting to edition 2 and run jury-only (recommended if year-one simplicity is the priority). The ambassadors do not recommend ever putting the main award to a public vote.

Recognition and feedback (sponsorship needs the forum's sign-off). Free end-to-end, stated explicitly as a positioning line: entry, judging, the credential, and the proof of winning are all free. The recognition package, all free and reusing existing assets, is a Credly-style badge (adapting the existing AIPM designs with earning criteria, a design adaptation, not a from-scratch build), a digital certificate, logo and mark use, a live online ceremony as the community moment, and a finalist "Story of Success" publication into the case-study library. The defining mechanic: **every finalist receives a written, criteria-anchored feedback report** of strengths and improvement areas (the IPMA lever, the only one of the 25 programs that does this, and the most defensible community-first differentiator). Once the published rubric exists, this feedback is cheap to produce. **No cash prize**, in line with the field norm; an optional non-cash, in-kind career-capital prize is available as a model (the PMI Madrid pack, about USD 267) if the founders want a tangible prize. Sponsorship posture: see the forum decision agenda.

Platform. For year one, a structured nomination form plus a video link plus a shared scoring sheet, enough to run the whole pipeline at near-zero cost. A paid awards SaaS (Award Force, the recurring infrastructure used by PMI, APM, IPMA, and the Australian AI Awards) is a scale decision for the business case, not a launch requirement.

Communications. AIPM's distribution advantage is owned channels plus the ambassador network, so the plan leans on those rather than paid reach. Run a July teaser, then launch the open call (mid-August 2026) on pmairevolution.com (a dedicated award page with the rules, rubric and timeline) and the AIPM LinkedIn company page, with the LinkedIn group (about 300 members) as a second push and Ricardo's and Antonio's reach as the primary amplifier. The core repeatable mechanic is **ambassador amplification**: a ready-to-post kit so the call propagates through ambassadors' personal networks (the Community and Growth North Star names that network as the project's distribution asset). Reveal the finalists ahead of the ceremony, announce the winner and the two specials live, and publish each "Story of Success" case so the recognition travels out through winners' own networks, turning a once-a-year award into year-round content.

Process and timeline. An eight-stage process (Farhad), a standard two-stage funnel (written and evidence assessment, then a live finalist pitch): (1) charter and criteria, (2) call for nominations, (3) eligibility screening, (4) written and evidence assessment, (5) Tier-1 regional judging (each region scores and promotes 2 or 3 above the cut-off), (6) Tier-2 global judging and finalist pitch (the regional winners, about six to nine across the three regions, present live to the global jury, which designates the overall winner and the two specials), (7) announcement and live online ceremony, (8) feedback and lessons (the

feedback reports plus the Story-of-Success publications). The indicative calendar is in section 7.

5. Business case

The award is a low-cost thought-leadership initiative with outsized return, and that return is what makes the near-zero year-one spend worth the founders' time.

5.1 Value: what the award produces

The award is a thought-leadership engine, not a one-night event.

- **A case-study pipeline (the primary asset).** Publishable-case-study consent at entry means every edition produces a set of vetted, real-world case studies of responsible AI in delivered project work (about six to nine from the global finalist pool, more if consented entrants who do not reach the global round are included). These are not marketing copy: jury-assessed against a published rubric, they are credible source material. They feed two initiatives the founders have already prioritized: the parallel **case-study repository** and the **webinar cadence** (about one case every other month, roughly six webinars a year, presented by case authors). One annual award therefore becomes year-round content that keeps the program visible between editions. A compounding loop, not a one-off.
- **Capability-building feedback.** The written, criteria-anchored feedback report makes the award a learning-and-improvement process for every finalist, not a trophy hand-off. This is the most defensible answer to “why enter a free award” and the clearest community-first signal AIPM can send, against any paid AI-industry award.
- **Community-first differentiation.** The free-end-to-end, published-rubric, feedback-returning posture is a combination no benchmarked award offers, the deliberate inverse of the credential-mill model (Global Recognition Awards).
- **Founder and AIPM brand.** Owning the responsible-AI-in-delivery wedge first positions AIPM, Ricardo, and Antonio at the front of a category that has no leader yet. The recognition travels: every finalist carries a shareable badge and write-up out through their own networks.

5.2 Cost: near-zero for year one

Cost item	Year-one approach	Cash cost
Submission and judging platform	Structured form + video link + shared scoring sheet	Zero (existing tooling)
Recognition assets (badge)	Reuse the existing Credly badge designs (a design adaptation, not a build)	Zero
Certificate, logo, mark	Bundled into the free win	Zero
Jury	Volunteer panel of about 7	Zero

Cost item	Year-one approach	Cash cost
	(ambassadors + external experts)	
Ceremony	Live online on existing webinar infrastructure (no venue)	Zero
Feedback reports	Produced by the jury against the published rubric	Time only
Cash prize	None (optional non-cash in-kind only)	Zero
Award Force SaaS	Deferred to a scale edition	Zero in year one

The entire year-one machine is buildable from assets AIPM already owns; the only real input is **volunteer effort**, concentrated in a few weeks of the October-to-December window.

5.3 Effort and owners

The **Community and Growth stream** owns the operational plan (lead: Eugene De Wee), aligned with that stream's North Star (recognition for active contributors, amplifying AI-in-Real-Projects outputs, a repeatable operating rhythm). The AI-in-Real-Projects content (Farhad's design) informs the rubric, the evidence pack, and the case-study quality bar. The jury is a time-boxed volunteer commitment around final judging, and the ambassador network is the distribution engine. Effort is front-loaded into charter and rubric publication (July to September), then a short judging-and-ceremony sprint (November and December). As an order-of-magnitude estimate (indicative, to confirm once the jury is seated): the charter-and-rubric build is a few volunteer-days across July to September, and each juror commits roughly 8 to 12 hours across eligibility screening, scoring, and the finalist feedback reports in November and December.

5.4 Risks and mitigations

Risk	Mitigation
Scope creep into many categories in year one (the benchmark pitfall)	One overall award plus two jury-designated specials on a single funnel and a single rubric; the five-category split is explicitly deferred to edition 2.
Continuity neglect (the benchmark's recurring failure: a stale award loses credibility)	Commit up front to publishing each year's winner, jury, and rubric; the case-study and webinar loop keeps the award visible between editions.
Pay-to-win perception	Free end-to-end, stated explicitly; no cash prize; if sponsors are taken, the CloudFest sponsor-to-charity model so no cash reaches winners.

Risk	Mitigation
Jury conflict of interest (an ambassador's own project may enter)	Published auto-decline plus disclose-and-recuse, calibration before scoring, and a no-contact-with-judges rule.
Looking like a PMI copy; AI-ghostwritten entries	Borrow mechanics, never branding (AIPM's own naming and criteria); a human-authored requirement with detection reserved and the no-contact-with-judges rule.
Thin first field (too few credible nominations in edition 1, the most plausible failure mode for a free, first-of-its-kind award)	The ambassador amplification kit and founder reach as primary amplifiers; global eligibility from day one to widen the pool; the 24-month window available (Decision #2) if the founders want a wider first field; the regional-to-global structure advances entries only above a published cut-off (no filling of slots), so a thin field never forces weak entries forward; and running just 3 macro-regions in year one (Decision #9) keeps each panel staffable.

5.5 Sponsorship posture (needs the forum's sign-off)

Two clean postures, both consistent with the non-commercial guardrail:

- **(a) Owner-funded, no sponsor for edition 1 (recommended).** The simplest, cleanest first edition; PMI's own global awards carry no external sponsor.
- **(b) CloudFest sponsor-to-charity decoupling.** A sponsor funds recognition or a charity scholarship for brand exposure, and no cash reaches any winner: the benchmark's cleanest answer to taking sponsorship without distorting the award.

The ambassadors recommend **(a)** for edition 1 and keep **(b)** documented as the approved mechanism if a sponsor appears.

6. Forum decision agenda

The forum ratifies all nine decisions; every row below is the ambassadors' recommendation, not a settled call. The table gives the recommended option and the alternative; the four items that most need the founders' call are tiered beneath it. The status labels carry over from the concept proposal, where decisions #1 and #6 were marked "Locked". Here every row reads "Recommended" or "Needs sign-off", because every decision, those two included, is a proposal for the forum to confirm. "Needs sign-off" marks the rows where one sub-decision genuinely stays open.

#	Decision	Ambassadors' recommendation	Alternative	Status
1	Categories	1 overall + Excellence in Responsible AI + Excellence in Agentic AI (single funnel, single rubric)	Full five-category split (deferred to edition 2)	Recommended
2	Eligibility and window	AI meaningful and beyond PoC; 18-month window ; global; case-study consent; joint entries; vendor-neutral	24-month window for a wider first field	Recommended (window flagged)
3	Nomination model	Self or open-peer; single Lead Nominator; draft-then-final; up-to-180-char first gate; joint entries	(none material)	Recommended
4	Jury and COI	Mixed panel of about 7 (range 5 to 10); calibrate; auto-decline + disclose-and-recuse; anti-ghostwriting; no-contact	Adjust panel size within the 5-to-10 range	Recommended
5	Community vote	Jury-only main award + optional secondary Community Choice on the vetted shortlist	Defer all public voting to edition 2 (jury-only)	Needs sign-off
6	Scoring rubric	Hybrid: 6 weighted	(none material)	Recommended

#	Decision	Ambassadors' recommendation	Alternative	Status
		criteria, 1 to 10, weights published; 60% on outcomes + governance + adoption		
7	Platform	Form + video + scoring sheet for year one (Award Force = a scale decision)	Adopt Award Force at launch	Recommended
8	Recognition, feedback, sponsorship	Free end-to-end; badge + certificate + ceremony + finalist feedback report; no cash; owner-funded edition 1	Optional non-cash in-kind prize; accept sponsors via sponsor-to-charity model	Needs sign-off (sponsorship only)
9	Geographic judging structure	Regional (ambassador-led) to global two-tier; 3 macro-regions for edition 1 (Americas, EMEA, Asia-Pacific); 2 or 3 advance per region above a cut-off; global jury adds the founders and a sponsor juror	Single global jury for year one	Recommended

The person-level “AI Project Manager of the Year” (the fourth item below) is a cross-cutting edition-scope question, deliberately kept outside the formal design decisions rather than added as a further row.

The four items that most need the founders’ call, tiered:

The two genuine posture sign-offs: 1. **Sponsorship posture (Decision #8):** owner-funded with no sponsor for edition 1 (recommended), or accept sponsors via the CloudFest sponsor-to-charity model (no cash to winners). 2. **Community Choice vote (Decision #5):** include a secondary public vote on the finalist shortlist, or run jury-only for year one.

The two lighter confirm-or-defer items: 3. **Eligibility window (Decision #2):** confirm 18 months (recommended) versus 24 months. 4. **Person-level award:** confirm or defer a separate “**AI Project Manager of the Year**” (a different axis, person versus project). The founders’ 3 June note raised it; the ambassadors recommend keeping edition 1 project-only because a person-level award would roughly double the year-one operational load, and surfacing it as a candidate for edition 2 for the founders to decide.

7. Timeline and next steps after ratification

The indicative calendar fits the fixed “launch October to December 2026, online ceremony” window, with planning and a teaser starting in July, and is for the forum to confirm:

- **July 2026:** finalize the charter, rubric, rules, jury, and platform per the forum’s decisions, and run a **pre-announcement / teaser** that opens an interest-and-intent list so the nomination window opens warm.
- **Mid-August 2026: public launch,** nominations open (rules, rubric, and timeline published on pmairevolution.com and the AIPM channels).
- **End of September 2026:** nominations close (about a 6-week window; the July pre-announce de-risks the tightness for a first-of-its-kind award that requires a video and case-study consent).
- **October 2026: Tier-1 regional judging:** eligibility screening, written and evidence assessment, and regional scoring, with 2 or 3 entries promoted per region above the cut-off.
- **November to December 2026: Tier-2 global judging** (regional winners pitch to the global jury) and the **live online ceremony** in December, announcing the winner and the two specials.
- **December 2026 to January 2027:** feedback reports issued to every finalist; Story-of-Success cases published into the case-study library and routed to the webinar cadence.

After ratification, the Community and Growth stream (Eugene De Wee) sets the charter, rules, and published rubric per the forum’s decisions on the four flagged items; recruits and calibrates the regional panels and the global jury (about seven each) and publishes their composition and COI policy; stands up the year-one pipeline (form, video link, scoring sheet) and adapts the Credly badge; prepares the ambassador amplification kit and the award page on pmairevolution.com. The calendar above is the run plan. Once ratified, this proposal becomes the operating brief for edition 1.

One point stands above the operational detail: the wedge is real but unowned, and it will not stay unowned. Ratifying the design now and launching in 2026 is how AIPM claims the

responsible-AI-in-delivery-plus-feedback ground before a PM body or an AI award reaches it first.