

AIPM “AI Project of the Year” Award: Concept Proposal (Round 2)

Status: draft of the AIPM Ambassadors’ recommendation, for the founders’ forum.

This is the Round 2 deliverable in the founders’ sequence (1) benchmark, done in award-benchmark.md; (2) **this concept proposal**; (3) business case; then the **AIPM Ambassadors** bring it to the founders’ forum (Ricardo Vargas and Antonio Nieto-Rodriguez) to decide. The recommendations here are the ambassadors’ collective position, not any one ambassador’s.

This is a **design document** (proposals, not source-verified facts), so it does not carry [VERIFY] markers; that discipline lives in the benchmark. Design proposals adopted from a named contributor are attributed (most of the first-draft design is **Farhad Abdollahyan**, AIPM “AI in Real Projects” subgroup; see round2-concept-inputs.md). Every recommendation links back to the benchmark section that grounds it. House style: English, no em-dash (commas, parentheses, colons).

0. How to read this document

The award has **nine design decisions** (the original eight from the Round 2 agenda, benchmark §4.4, plus a ninth on geographic judging structure added this round). Two were settled (with Vitor) as the **ambassadors’ recommended options** to carry forward (#1, #6); five carry a benchmark-grounded **recommendation** that proceeds unless the forum objects (#2, #3, #4, #7, #9); two carry a **needs-sign-off** status where the package is recommended but one sub-decision stays open (#5, #8). Counting the two lighter confirm-or-defer items as well, **four discrete items** most need a forum decision before launch (sponsorship posture, the Community Choice vote, the 18-vs-24-month window, and the person-level-award option), enumerated in §12. The status of each decision is summarized in the table below and in §12, and called out inline.

Note for the founders’ forum (Ricardo Vargas and Antonio Nieto-Rodriguez).

This is the **AIPM Ambassadors’** recommendation document, feeding the consolidated business plan/proposal the ambassadors bring to the forum. **The forum ratifies all nine decisions.** In the tables, “**Locked**” means the option the ambassadors recommend bringing to the forum (settled with Vitor this round), not a final decision; “**Recommended**” and “**Needs sign-off**” are equally proposals for the forum to confirm. Nothing here is decided until the forum decides it.

Decision	This proposal	Status
#1 Categories	1 overall + Responsible AI + Agentive AI special recognitions	Locked (Vitor, this round)
#6 Scoring rubric	Hybrid: 6 weighted criteria, 1-to-10 scale, weights	Locked (Vitor, this round)

Decision	This proposal	Status
	published	
#2 Eligibility and window	AI meaningful, beyond PoC, completed within 18 months, global, case-study consent	Recommended
#3 Nomination model	Self or open-peer, single Lead Nominator, draft-then-final, joint entries allowed	Recommended
#4 Jury and COI	Mixed panel of ~7 (range 5 to 10), calibrated, auto-decline plus disclose-and-recuse	Recommended
#5 Community vote	Jury-only main award + optional secondary Community Choice on the vetted shortlist	Needs sign-off
#7 Platform	Form + video link + shared scoring sheet for year one (Award Force is a scale decision)	Recommended
#8 Recognition, feedback, sponsorship	Free end-to-end, badge + certificate + ceremony + feedback report; sponsor-to-charity if any	Needs sign-off (sponsorship only; rest of package recommended)
#9 Geographic judging structure	Regional (ambassador-led) to global two-tier; 3 macro-regions for edition 1	Recommended

1. Purpose and positioning

The AIPM “AI Project of the Year” Award recognizes AI-enabled projects that demonstrate **responsible innovation, strong delivery, measurable value, human-centered adoption, and positive organizational or societal impact** (positioning drafted by Farhad, round2-concept-inputs.md §1). It is positioned 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.

The differentiation wedge (benchmark §4.3, the headline finding §4). Across the 25 benchmarked programs, two things together are owned by no existing award: (a) recognizing **responsible, real-world AI application inside delivered project work** (not a model, a tool, a 24-hour prototype, or a generic business-excellence story), and (b) returning **criteria-anchored feedback that builds the practitioner’s capability** rather

than only handing out a trophy. The PM bodies (PMI, PMOGA, APM, IPMA) own rigorous delivery rubrics but say nothing about AI. The AI-industry awards (Stevie, BIG, AI Breakthrough, Webby, National AI Awards, ITU, Australian AI) reward AI products against generic business-excellence language. The hackathons judge greenfield builds under time pressure. The white space at the intersection (responsible AI-in-delivery plus capability-building feedback, free and online) is exactly where AIPM plants its flag. The Australian AI Awards prove the framing is already category-viable (live “Best Application of Responsible AI” and “Best Use of Agentic AI” categories, §3.16), which makes the wedge real but time-limited and reinforces the founders’ “launch before competitors in 2026”.

Guardrails this design is written against (hard constraints): - **Original** (does not look like a copy of PMI): borrow mechanics, never PMI branding, naming, or visual identity (§4.3); AIPM uses its own naming and category labels even where a category framing is adopted from the field. - **Non-commercial** (free end-to-end, no entry fee and no paywalled proof). - **Practice-grounded** (recognize real AI application in delivered projects, not ideas or tools). - **Community-first** (the lasting value is the practice it teaches the community).

Fixed parameters from the founders’ sketch (notes 25 May and 3 Jun 2026): launch **Oct to Dec 2026**, first ceremony **online**, **open to outside nominations**, jury of **5 to 10**, about **5 finalists to 1 winner**.

2. Categories (Decision #1, LOCKED)

Edition 1 runs one overall award plus two special recognitions:

1. **AI Project of the Year** (the overall winner, one per edition).
2. **Excellence in Responsible AI** (special recognition).
3. **Excellence in Agentic AI** (special recognition).

Design that keeps this within the year-one scope-creep guardrail (benchmark pitfall §4.2). The two special recognitions are **not separate entry tracks**: there is a single submission funnel and a single rubric. All entries compete for the overall award; from the finalist pool, the jury additionally designates the two special recognitions, where: - **Excellence in Responsible AI** goes to the finalist that best exemplifies criterion 2 of the rubric (Responsible AI governance and Human-in-Command). - **Excellence in Agentic AI** goes to the finalist that demonstrates the most compelling, responsibly governed use of agentic or autonomous AI (judged only where an entry uses agentic AI; if no finalist qualifies, this recognition is simply not awarded in edition 1).

This honors the “1 overall plus 1 to 2 special recognitions” consensus (Farhad’s edition-1 caveat, round2-concept-inputs.md §3; benchmark §4.4) while adding **zero extra submission burden and minimal extra jury load** (same pool, same rubric, two designations made during final judging). It directly reinforces the differentiation wedge: the two specials are the two named axes the Australian AI Awards already validated (§3.16).

Explicitly deferred to edition 2 or later: Farhad’s full five-category split (AI for Business/Operational Excellence; Governance/PMO/Decision Intelligence; Customer/Citizen/Social Value; Sustainability and Resilience; Innovation and Responsible Experimentation, round2-concept-inputs.md §3), and any impact split (ITU People/Planet/Prosperity, §3.15) or fairness/size-track split (IPMA small-medium vs large, §3.21). These add operational muscle requirements the field shows take years to run well (APM, Stevie, Australian AI at ~33 categories, §4.2). Adding them in year one would violate the scope-creep guardrail.

Founder option to flag (not a year-one default): the founders’ 3 Jun note raised a possible “**AI Project Manager of the Year**” person-level award alongside the project award (knowledge-base/04-decisions-and-action-items.md). This is a different axis (person vs project) and would double the year-one operational load, so this proposal keeps edition 1 **project-only** and surfaces the person award as an edition-2 candidate for founder decision.

3. Eligibility and window (Decision #2, recommended)

Recommended eligibility rules: - **AI is a meaningful component** of the project (not a bolt-on), and the work is **beyond proof-of-concept:** completed or operationalized in a real setting (Farhad, round2-concept-inputs.md §2; benchmark §4.3 wedge). - **Completion window: within the previous 18 months** at the nomination deadline. This aligns with the PMI Project of the Year norm (≤ 18 months, §3.5) and is tighter than Farhad’s 18-to-24 suggestion, which keeps the year-one field focused. (Sign-off micro-decision: **18 vs 24 months**; the proposal recommends 18, with 24 acceptable if the founders want a wider first field.) - **Global from day one.** AIPM is an international community, so the award is open worldwide from edition 1 rather than piloted in one region. - **Evidence of responsible AI** (data governance, risk and ethics handling, human accountability and oversight) and of **adoption** (real operational embedding, not just deployment). - **Publishable case-study consent is required at entry.** This is the operational hook that makes the “Story of Success” library possible (Farhad §2; benchmark §4.3) and feeds the parallel AIPM case-study repository initiative (knowledge-base/04-decisions-and-action-items.md). - **Joint vendor plus practitioner entries are allowed**, with the spotlight kept on the practitioner or delivery team leading the adoption (National AI Awards model, §3.14). This widens the funnel beyond large enterprises while preventing the award from becoming a vendor showcase. Keep the evidence-of-real-deployment bar, drop any large-org framing (Wave 4 correction, §6.3). - **Vendor-neutral:** no bonus for using any particular technology. This is a deliberate inversion of the hackathon anti-pattern where a sponsor’s own stack earns points (Google Cloud ADK, §3.24; pitfall §4.2).

4. Nomination model (Decision #3, recommended)

Recommended model: - **Self-nomination and open peer nomination** both allowed (a project leader may enter their own project; a peer may nominate another's), plus **joint entries** (§3.14). - A **single named "Lead Nominator"** owns each entry and is the point of contact, with a **draft-then-final flow** (the entry can be saved and refined before final submission), mirroring the PMI Award Force / PAIT infrastructure pattern (§6.8). - A **lightweight first gate:** a short summary (the GitHub Stars ≤180-character front-door idea, §3.11) at intake, so screening can triage quickly before the full evidence pack is assembled.

This keeps the front door low-friction (community-first) while the Lead Nominator and draft-then-final flow give the process enough structure to manage entries without a paid platform in year one.

5. Scoring rubric and weights (Decision #6, LOCKED)

The headline decision. Edition 1 uses a **hybrid grid: six weighted criteria, each scored on a 1-to-10 scale, with the weights published**. This is the benchmark's explicit steer (§4.1): a 1-to-10-per-criterion scale is easier for a volunteer jury than a 100-point grid or IPMA's six-band model, while publishing the weights out-transparents PMI, which affirmatively declares its Project-of-the-Year weights proprietary (§6.8). It collapses Farhad's seven-criterion draft (round2-concept-inputs.md §4) onto six, keeping his two top weights intact (outcomes 25%, responsible-AI 20%, which themselves mirror PMOGA's top two slots, §3.4).

5.1 The six criteria

#	Criterion	Weight	What it rewards
1	Outcomes, benefits and impact	25%	Measurable, evidenced results (baseline vs actual, benefits realized, organizational or societal impact). The heaviest weight by design.
2	Responsible AI governance and Human-in-Command	20%	Data governance, risk and ethics handling, transparency, and demonstrable human accountability and oversight.
3	Adoption and	15%	Real operational

#	Criterion	Weight	What it rewards
	change management		embedding, stakeholder uptake, and sustained use beyond go-live (not a one-time deployment).
4	AI solution fit and innovation	15%	Appropriateness and ingenuity of the AI approach to the actual problem (fit and value, not novelty for its own sake).
5	Project management and delivery rigor	15%	Strategic relevance and problem definition plus delivery discipline (scope, schedule, risk), vendor-neutral. Absorbs Farhad's separate "strategic relevance" criterion.
6	Learning and contribution to the community	10%	Transferable lessons, quality of the publishable case study, and contribution to the profession and the AIPM community. Raised from Farhad's 5% to honor the community-first guardrail.
	Total	100%	

Scoring philosophy, adopted as an explicit principle (§6.3): “reward durable change and real adoption over model novelty.” Criteria 1, 2, and 3 (outcomes, responsible-AI governance, adoption and change) carry **60% of the weight combined**, so a technically impressive AI solution cannot win without credible evidence of outcomes, responsible governance, and real adoption. This operationalizes the Wave 4 “Change Blueprint” finding

(§6.3) and BIG’s “real-world problems, not a popularity contest” purpose (§3.9), generalized beyond data-ops to any AI-in-project-delivery case.

5.2 Rating anchors (per criterion, 1 to 10)

Reused from Farhad’s draft (round2-concept-inputs.md §4), the same anchors apply to every criterion so the jury calibrates consistently:

Band	Meaning
1 to 2	Minimal or unsupported (claim with no evidence)
3 to 4	Notable gaps
5 to 6	Adequate
7 to 8	Strong
9 to 10	Exceptional

Computing the score: each criterion’s 1-to-10 score is multiplied by its weight and summed, producing a weighted score on a 0-to-10 (or 0-to-100) scale. The full rubric, the anchors, and the weights are **published in the rules document** before the call opens (best practice §4.1; the single cheapest trust-builder, and an explicit advantage over PMI).

6. Jury, judging structure, governance, and integrity (Decisions #4 and #9)

6.1 Two-tier judging: regional to global (Decision #9, recommended)

Edition 1 is judged in **two tiers, regional then global**, so the ambassador network is the judging engine and the field is staffable from day one. This is orthogonal to the categories (§2): it is *how and where* entries are judged, not *what* is judged.

- **Tier 1, regional (October).** Entries are judged within their region by **that region’s AIPM ambassadors as the first layer**, alongside **criterion-eligible external jurors for that region** (jurors who qualify per criterion; the external bench is the pressure valve where ambassadors are few). Each region promotes **2 or 3 entries** to the global round, **only if they clear a published minimum cut-off score** (no filling of slots), so a thin first field never advances weak entries.
- **Tier 2, global (November or December).** The regional winners compete for the overall award and the two special recognitions. The **global jury is Ricardo Vargas, Antonio Nieto-Rodriguez, the ambassadors who passed the judge-selection criterion for the global award, and a juror from the award sponsor** (someone from a technology company or a VC, only if a sponsor is taken under Decision #8). The sponsor juror brings credibility through the bench, not through a payment: the juror sits under a published sponsorship-does-not-influence-judging firewall and the same conflict-of-interest rules as every other juror (auto-decline plus disclose-and-recuse), and no cash reaches any winner (the CloudFest sponsor-to-charity, no-cash posture, §3.22 and §4, is reused only for the money side under Decision #8).

Edition-1 regions: 3 macro-regions. Year one runs **3 macro-regions**, not the full 7 or 8, because the live ambassador roster (about 33 members) is heavily skewed (North America and Europe are roughly 60% of the network, South Asia has none, China one), so 7 or 8 balanced juries are not staffable in year one and a single conflict-of-interest recusal would gut a 3-person panel. The recommended split is **Americas** (North America plus Latin America), **EMEA** (Europe plus MENA plus Sub-Saharan Africa), and **Asia-Pacific** (the thin region, bolstered with external criterion-jurors). The full regional taxonomy (modeled on the PMI region taxonomy: Asia Pacific, Europe, Latin America, Middle East and North Africa, North America, South Asia, Sub-Saharan Africa; the PMO Global Alliance taxonomy is the same with China as a separate region) is the **growth path for edition 2 and beyond**.

- **Judge by affinity, not only residence**, to deepen thin benches: a Brazilian ambassador residing in North America can sit on a Lusophone or Latin-America bench, for example.
- **Cross-panel calibration is the known hard part of multi-panel judging** (a 9 out of 10 must mean the same thing across regions), so the regional panels calibrate against a shared worked example before scoring.

Decision #9 for the forum: adopt the regional-to-global two-tier structure at 3 macro-regions for edition 1 (the ambassadors' recommendation, because it puts the network to work as judges, spreads the judging load, and gives the award a regional-to-global narrative arc), or run a single global jury for year one (the alternative).

6.2 Panel composition and integrity (Decision #4, recommended)

Recommended panel (each tier). Each regional panel and the global jury is: - A **mixed, not all-academic, panel of about 7** (within the founders' 5-to-10 range; PMI Netherlands runs 5 to 8, §3.18; CloudFest about 9, §3.22), drawn from **AIPM ambassadors plus external experts**, covering project management, AI governance, business value, ethics and risk, and sector representation (Farhad's mixed-panel design, round2-concept-inputs.md §7; validated ITU model, §3.15; PMI Sweden already seats an AI-in-project-delivery juror, §3.19). - **Calibration before scoring** (a shared worked example so jurors apply the anchors consistently, within each panel and across the regional panels), plus **confidentiality undertakings** (Farhad §7).

Integrity rules to publish (all zero-cost, each beats a benchmark on credibility): - **Conflict-of-interest: auto-decline plus disclose-and-recuse.** No juror judges their own entry or their own employer's entry: any entry linked to a juror (the PMI auto-decline rule, §6.8) is declined for that juror, who discloses and recuses (ITU disclose-and-recuse, §3.15). This is central, not peripheral, now that ambassadors are the judging engine and an ambassador's own project may enter. - **Anti-AI-ghostwriting rule.** Entries must be human-authored; AIPM reserves the right to use detection tools (APM, §3.6). For an AI award this is both necessary and a credibility signal. - **No-contact-with-judges rule**, with disqualification for breach (APM, §3.6; best practice §4.1). - **Sponsorship does not influence judging** (a published firewall, Farhad §7), relevant only if §8's sponsorship option is taken. - **Judges must distinguish claims from evidence** (Farhad §7), which the rating anchors (band 1 to 2 = "claim with no evidence") operationalize directly.

Publishing the **exact jury size, its independence, and the COI policy** puts AIPM above the benchmark floor (the anonymous-undisclosed-jury anti-pattern of AI Breakthrough §3.10 and Global Recognition Awards §3.25, pitfall §4.2).

7. Community vote (Decision #5, NEEDS SIGN-OFF)

Recommendation: jury-only for the main award in year one, with an optional secondary “Community Choice” recognition. The overall winner and the two special recognitions are decided by the jury (this protects the award from becoming a popularity contest, the explicit caution behind BIG’s positioning, §3.9).

The **optional** secondary recognition: a “**Community Choice**” 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), following the one rule every voting program treats as non-negotiable (Webby’s academy-picks-then-public-vote, §3.12; PMI-SP/PMI-RIO fused committee-plus-audience, §3.1/§3.2; pitfall §4.2). It is clearly labeled as separate from and secondary to the jury award.

Why this needs sign-off: a Community Choice strengthens the community-first guardrail and adds energy to the online ceremony, but it adds a voting-infrastructure and fraud-control task to a first edition. The two clean choices are: **(a)** include it as a light secondary recognition on the finalist shortlist (recommended if the founders want community participation), or **(b)** defer all public voting to edition 2 and run jury-only (recommended if year-one simplicity is the priority). The proposal does **not** recommend ever putting the main award to a public vote.

8. Recognition, feedback, sponsorship, and prize (Decision #8, sponsorship NEEDS SIGN-OFF)

Free end-to-end, and stated explicitly. Every part of the experience is free: entry, judging, the credential, and the proof of winning. This is the credible non-commercial wedge now that free entry alone is table stakes (§6.3), and it is the deliberate inverse of the Global Recognition Awards “free-to-enter, pay-for-the-proof” credential mill (USD 99 / USD 499 packs, paywalled logo, §3.25; pitfall §4.2). State the zero cost as a positioning line.

Recognition package (all free, reusing existing assets): - A **Credly-style badge**. AIPM already has finished Credly badge designs with defined earning criteria (knowledge-base/10-brand-assets-and-badges.md), so an award badge is a design adaptation, not a from-scratch build. (The badge issuing authority and platform, Credly vs APMG vs self-hosted via the Research-Hub badge engine, is an open founder decision already tracked in that asset note.) - A **digital certificate** and **logo / mark use** bundled into the free win (the inverse of GRA, §3.25). - A **live online ceremony** as the community moment, not a press-release afterthought (avoiding the BIG / GRA PR-only anti-pattern,

§4.1, best practice 9). - A **finalist “Story of Success” publication** (extending PMI’s Story-of-Success format, §3.5/§4.3), published into the AIPM case-study library.

The capability-building lever (the genuinely new mechanic, benchmark §4.3 and best practice 7): every finalist receives a written, criteria-anchored feedback report of strengths and improvement areas (the IPMA lever, §3.21, the only one of the 25 programs that does this; Farhad’s process stage 8, round2-concept-inputs.md §6/§9). Once the published rubric exists (§5), this feedback is cheap to produce and turns the award from a one-shot recognition into a learning-and-improvement process. This is AIPM’s **most defensible community-first differentiator** against any paid AI-industry award.

Prize: no cash. The field’s confirmed norm is no cash purse for nomination-and-rubric awards (§6.7); the only cash sits in hackathons, which AIPM should not chase (pitfall §4.2). An optional **non-cash, in-kind career-capital prize** is available as a model (PMI Madrid’s in-kind thesis-award pack: a student membership plus exam fee, ~USD 267, §3.20) if the founders want a tangible prize beyond the recognition package.

Sponsorship (NEEDS SIGN-OFF). Two clean postures: - **(a) Owner-funded, no sponsor for edition 1** (simplest; PMI’s global awards carry no external sponsor, §6.7). Recommended for a clean first edition. - **(b) Accept sponsors via the CloudFest sponsor-to-charity decoupling** (§3.22): a sponsor funds recognition or a charity scholarship for brand exposure, and **no cash reaches any winner**. This is the cleanest answer in the benchmark to “how to take sponsorship without distorting the award”, and it neutralizes the pay-to-win risk a cash award would carry. The proposal recommends **(a)** for edition 1 and keeps **(b)** documented as the approved mechanism if a sponsor appears. (This intersects the open benchmark [VERIFY]: the Clube da Programação US\$10,000 prize sponsor is unnamed publicly, §3.13; AIPM’s no-cash posture sidesteps that tension entirely.)

9. Platform (Decision #7, recommended)

For year one: a structured nomination form plus a video link plus a shared scoring sheet. This is enough to run the whole pipeline (best practice 2; §4.1), and it keeps cost near zero. A paid awards SaaS (**Award Force**, the recurring infrastructure used by PMI, APM, IPMA, and the Australian AI Awards, §4.3) is a **scale decision for Round 3 (business case)**, not a launch requirement.

10. Process and indicative timeline

Eight-stage process (Farhad, round2-concept-inputs.md §6), a standard two-stage funnel (written and evidence assessment, then a live finalist pitch) matching PMOGA’s multi-stage elimination (§3.4) and APM’s submission-then-virtual-presentation (§3.6):

1. **Charter and criteria** (publish the rules, rubric, and weights).

2. **Call for nominations** (open).
3. **Eligibility screening** (the lightweight first gate, §4).
4. **Written and evidence assessment** (the documentation package, §10.1).
5. **Tier-1 regional judging** (each region's panel scores against the published rubric and promotes 2 or 3 entries above the cut-off to the global round).
6. **Tier-2 global judging and finalist pitch** (the regional winners, about 6 to 9 across the 3 regions, give a 10-to-15-minute live online presentation plus Q&A to the global jury, which scores against the published rubric and designates the overall winner plus the two special recognitions).
7. **Announcement and visibility** (the live online ceremony).
8. **Feedback and lessons** (criteria-anchored feedback report to every finalist; Story-of-Success publications go live).

The two-tier judging described above maps onto this funnel: stages 1 to 4 run globally, stage 5 is the regional tier, and stage 6 is the global tier. The 3 regions each advancing 2 or 3 winners yields about 6 to 9 global finalists, a touch above the founders' single-jury sketch of about 5; the forum can tune the count through the per-region advancement number.

Indicative calendar (to confirm with founders; fits the fixed "launch Oct to Dec 2026, online ceremony" window, with planning and a teaser starting in July): - **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). - **End of September 2026:** nominations close (about a 6-week window; the July pre-announce de-risks the tightness for a brand-new award that requires a video and case-study consent). - **October 2026: Tier-1 regional judging** (eligibility screening, written and evidence assessment, regional scoring, 2 or 3 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 overall winner and the two special recognitions. - **December 2026 to January 2027:** feedback reports issued to every finalist; Story-of-Success cases published into the case-study library.

10.1 Documentation package (Farhad, round2-concept-inputs.md §5)

- **Nomination Fact Sheet** and a **Structured Questionnaire** (mirrors PMI Project of the Year, §3.5).
- **Evidence Pack:** business case, charter, baseline vs actual, benefits report, risk log, **AI governance / ethics review, adoption data, human-oversight evidence**, lessons learned. The AI-governance review, adoption data, and human-oversight evidence are the distinctively-AIPM additions to a standard PM evidence pack.
- **Finalist Presentation:** 10 to 15 minutes plus Q&A.
- **Case Study Consent** (required, §3).

A short video is the recommended lightweight submission format because it lets judges see the real application (best practice 2; the two-stage video-first norm across PMI, APM, PMOGA, Clube da Programação).

11. Communications and distribution plan

The award only delivers community value if the call reaches nominators and the winners reach the community. Promotion is tracked as its own dimension across the benchmark (dimension 12) and best practice 9 is explicit: make the live online ceremony the community moment, not an afterthought (§4.1). AIPM’s distribution advantage is owned channels plus an ambassador network, so the plan leans on those rather than paid reach.

- **Pre-announcement teaser (July 2026), then public launch (mid-August 2026).** Open an interest-and-intent list with a July teaser, then announce the open call at the mid-August public launch on pmairevolution.com (a dedicated award page carrying the rules, rubric, and timeline) and on the AIPM LinkedIn company page, with the LinkedIn group (~300 members) as a second push. Ricardo Vargas’s and Antonio Nieto-Rodriguez’s reach is the primary amplifier for the launch post.
- **Ambassador amplification (the core mechanic).** The ambassador network is the project’s named distribution asset (the Community & Growth North Star, knowledge-base/10-brand-assets-and-badges.md): give ambassadors a ready-to-post kit (copy, the stream-colored visuals, the call link) so the call propagates through their personal LinkedIn networks. This is the repeatable mechanic, not a one-off.
- **Finalist and winner announcement (Dec 2026).** Reveal the finalists ahead of the ceremony to build anticipation, then announce the overall winner and the two special recognitions at the live online ceremony. Each finalist and winner receives a shareable badge and a write-up so the recognition travels through their own networks (the inverse of a PR-only, paywalled-logo model, §3.25).
- **The case-study library as an ongoing distribution engine.** The publishable “Story of Success” cases (taken under entry consent, §3) are published into the AIPM case-study library and feed the parallel case-study-repository and webinar cadence (about one case every other month, knowledge-base/04-decisions-and-action-items.md). This turns a once-a-year award into year-round content and keeps it visible between editions.

12. Decision summary and what needs sign-off

#	Decision	Recommendation	Status	Grounding
1	Categories	1 overall + Responsible AI + Agentic AI special recognitions (single funnel, single rubric)	Locked	§3.16, §4.2, Farhad §3

#	Decision	Recommendation	Status	Grounding
2	Eligibility and window	AI meaningful and beyond PoC; completed within 18 months; global; case-study consent; joint entries allowed; vendor-neutral	Recommended	§3.5, §3.14, §6.3, Farhad §2
3	Nomination model	Self or open-peer; single Lead Nominator; draft-then-final; ≤180-char first gate; joint entries	Recommended	§3.11, §3.14, §6.8
4	Jury and COI	Mixed panel of ~7; calibrate; auto-decline + disclose-and-recuse; anti-ghostwriting; no-contact	Recommended	§3.6, §3.15, §6.8, Farhad §7
5	Community vote	Jury-only main award + optional secondary Community Choice on the vetted shortlist	Needs sign-off	§3.12, §3.9, §4.2
6	Scoring rubric	Hybrid: 6 weighted criteria, 1-to-10, weights published; 60% on outcomes + governance + adoption	Locked	§4.1, §6.3, §6.8, Farhad §4
7	Platform	Form + video + scoring sheet for year one	Recommended	§4.1, §4.3

#	Decision	Recommendation	Status	Grounding
		(Award Force = Round 3 scale decision)		
8	Recognition, feedback, sponsorship	Free end-to-end; badge + certificate + ceremony + finalist feedback report; no cash; sponsorship posture open	Needs sign-off (sponsorship only; rest recommended)	§3.20, §3.21, §3.22, §3.25
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	Recommended	§3.22, §4

The items that most need the founders' forum decision before launch, tiered by weight:

Two posture decisions (the genuine 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. This also feeds Round 3 (business case). 2.

Community Choice vote (Decision #5): include a secondary public vote on the finalist shortlist, or run jury-only for year one.

Two lighter confirm-or-defer items: 3. **Eligibility window (Decision #2):** confirm 18 months (recommended) vs 24 months. 4. **Person-level award:** confirm or defer the separate "AI Project Manager of the Year" (kept out of edition 1 here; §2).

Next in the founders' sequence: Round 3 (business case), which inherits three inputs already surfaced here, the case-study-library ROI framing (Farhad §8), the platform scale

decision (Award Force, §9), and the sponsorship posture (§8). Then validate the full design with Ricardo and Antonio.