

AIPM AI Project of the Year Award: Evidence-Pack Guide

Edition 1 (2026). Wave C, release candidate. Status: pending fold-in of the subgroup comments due 2026-07-15.

This evidence-pack guide is a **derived document**. It derives from, and is subordinate to, the operating rules in `operating-rules.md` (the single source of truth, the SSOT), and it is a companion to the nomination form (`nomination-form.md`) and a sibling of the judge guide (`judge-guide.md`) and the scoring-sheet specification (`scoring-sheet-spec.md`). Where anything in this guide could be read to conflict with the operating rules, **the operating rules prevail** and this guide is corrected. Section references written as “Rules Section N” or “Rules clause N” point to the operating rules; references written as “Judge Guide Section N” point to `judge-guide.md`. Nothing here is decided until the founders’ forum (Ricardo Vargas and Antonio Nieto-Rodriguez) ratifies it; every design choice below is the ambassadors’ recommendation for the forum to ratify.

This guide ships as a **PDF companion** to the editable-DOCX nomination form. It tells a nominator how to assemble a strong, fair, judge-ready evidence pack, so a first-time nominator with a good project but no awards experience can enter with confidence.

House-style note: this text avoids the em-dash and the en-dash by design. Commas, parentheses, and colons carry the same work, and ranges are written with the word “to” (for example, 1 to 10).

1. Purpose

1.1 What this guide is. This guide explains how to assemble the evidence pack: the body of supporting artifacts that substantiate the claims in your entry (Rules clause 2.8). It is the companion to the nomination form. The form is where you answer; the evidence pack is what backs those answers. Read it with the form open.

1.2 Why the evidence pack matters more than the words. Judges score what can be verified, not what is asserted. Unsupported assertions score in the 1 to 2 band and are not credited (Rules clauses 8.5 and 9.4; Judge Guide Section 6). A confident narrative with no artifact behind it does not earn a higher band; the artifact, the metric, or the demonstrated practice is what moves a score up. The simple test a juror applies: could an auditor verify this from the evidence pack? If nothing backs it up, it scores as minimal or unsupported. Assemble your pack so that every claim in the form points to an artifact here.

1.3 How this guide aligns with judging. The six evidence sections below map one to one to the six rubric criteria (Rules Section 8) and to the “what judges assess” and “typical evidence” columns in the judge guide (Judge Guide Section 5). When you satisfy this guide, you are building the pack a juror is trained to look for. In the form, each questionnaire

section asks you to name the evidence-pack items that back it; this guide explains what those items should be.

1.4 The rules win. This guide never restates a rule in a way that changes it. If you find a divergence between this guide and the operating rules, treat the operating rules as governing (Rules clauses 13.1 and 13.6).

2. What counts as evidence, per criterion

The evidence pack contents (Rules clause 2.8; concept-proposal Section 10.1) are: business case, charter, baseline versus actual, benefits report, risk log, AI governance or ethics review, adoption data, human-oversight evidence, and lessons learned. Below, each criterion says what a juror assesses, the typical evidence for it, and a short do and do-not. The banded reads are the judge guide's, quoted so you calibrate to how you will be scored (Judge Guide clause 5.4). The anchor scale, 1 to 10, is the same for every criterion; what changes per criterion is what counts as evidence.

2.1 Criterion 1: Outcomes, benefits and impact (25%, the heaviest)

What judges assess (Judge Guide Section 5): “Measurable value, before and after evidence, benefits realized, organizational or societal impact.”

Typical evidence: baseline and actual metrics, benefits report, ROI, public value, operational improvements.

Do: state a baseline, then the actual measured against it, then benefits that are realized rather than projected. A 9 to 10 shows a stated baseline, an actual measured against it, and realized benefits, for example a benefits report tied to the baseline metric (Judge Guide clause 5.4.1).

Do not: claim impact (“saved significant cost”) with no number and no source. A figure with no baseline to compare it against is only adequate; a claim with no figure at all is a 1 to 2 read.

2.2 Criterion 2: Responsible AI governance and Human-in-Command (20%)

What judges assess (Judge Guide Section 5): “Ethics, risk controls, accountability, transparency, privacy, security, bias controls, and human oversight.”

Typical evidence: AI governance artifacts, risk or ethics review, privacy review, audit trail, human-approval model.

Do: show the governance artifact plus the human-approval model actually operating (an audit trail, a documented risk or ethics review, the point where a human signs off). A 9 to 10 shows the artifact plus oversight in use (Judge Guide clause 5.4.2). This criterion also anchors the Excellence in Responsible AI recognition (Rules clause 3.3), so thin evidence here costs twice.

Do not: assert “the model is fair and safe” with no risk review, no bias control, and no oversight point. Naming a policy without showing it in use is only a 5 to 6.

2.3 Criterion 3: Adoption and change management (15%)

What judges assess (Judge Guide Section 5): “Real operational embedding, user uptake, stakeholder trust, training, and sustained use beyond launch.”

Typical evidence: adoption metrics, training materials, change plan, stakeholder feedback, usage data.

Do: show sustained use beyond launch (usage data over time, a change plan that was executed, stakeholder feedback). A 9 to 10 shows sustained use, not just a launch spike (Judge Guide clause 5.4.3).

Do not: show deployment with no evidence anyone used it. Deployment alone is not adoption (Rules clause 4.1.6); that is a 1 to 2 read. See Section 3.2.

2.4 Criterion 4: AI solution fit and innovation (15%)

What judges assess (Judge Guide Section 5): “Appropriate and effective use of AI for the problem, originality where relevant, and fit-for-purpose design.”

Typical evidence: architecture summary, model or tool rationale, performance evidence, solution demo.

Do: justify why this AI approach fits this problem, with performance evidence. A 9 to 10 explains the fit and shows performance evidence (Judge Guide clause 5.4.4).

Do not: invoke a fashionable technique with no rationale and no fit to the problem. Describing the solution without saying why it was the right tool is only adequate. Novelty alone does not carry an entry.

2.5 Criterion 5: Project management and delivery rigor (15%)

What judges assess (Judge Guide Section 5): “Problem definition, governance, planning, delivery discipline, risk and issue management, and lessons learned.”

Typical evidence: charter, plan, schedule, risk log, governance cadence, delivery artifacts.

Do: show the delivery spine (charter, plan, risk log, governance cadence) and lessons learned. A 9 to 10 shows the spine plus lessons learned (Judge Guide clause 5.4.5).

Do not: offer a narrative of delivery with no supporting artifact. That is a 1 to 2 read.

2.6 Criterion 6: Learning and contribution to the community (10%)

What judges assess (Judge Guide Section 5): “Transferable lessons and the quality of the story the case can contribute to AIPM knowledge sharing.”

Typical evidence: lessons learned, case-study draft, reusable practices, reflection on what others can adopt.

Do: offer transferable, reusable lessons and a case-study-ready story others can adopt. A 9 to 10 offers reusable lessons others can adopt (Judge Guide clause 5.4.6). This is where the community-first purpose of the award is scored.

Do not: offer no reflection beyond “it went well.” That is a 1 to 2 read.

Weighting note. Criteria 1, 2, and 3 carry 60% of the total weight combined (Rules clause 8.2). Put your strongest evidence behind outcomes, responsible-AI governance, and adoption; a brilliant solution (Criterion 4) cannot win on its own.

3. The two distinctively-AIPM must-haves

Three artifacts in the pack are the distinctively-AIPM additions to a standard project-management evidence pack: the AI governance or ethics review, the adoption data, and the human-oversight evidence. Two ideas sit under them, and both are eligibility conditions, not nice-to-haves. Because they are required for eligibility, the organizers screen for them before scoring even begins (Rules clause 4.2.6). Do not spend effort polishing the other sections while leaving these thin.

3.1 Responsible-AI evidence is required (Rules clauses 4.1.5 and 4.2.6). Your entry must include responsible-AI evidence: governance, risk, ethics, transparency, privacy, and bias mitigation, together with human oversight where applicable. This is carried by two artifacts, the AI governance or ethics review and the human-oversight evidence. Show the governance operating, not merely declared: the risk or ethics review that was actually done, and the point where a human is in command and signs off (the Human-in-Command model). A statement that the system is responsible, with nothing behind it, does not satisfy this requirement and scores in the lowest band (Judge Guide clause 5.4.2). An entry that lacks this evidence is not eligible.

3.2 Adoption, not only deployment, is required (Rules clauses 4.1.6 and 4.2.6). Your entry must include adoption evidence, not only deployment evidence. This is the single most common gap. Showing that something went live is not enough; the entry must show it was taken up and used. Deployment is the day it launched; adoption is people using it and continuing to use it. Bring usage data over time, an executed change plan, training materials, and stakeholder feedback. Launch-week numbers alone read as deployment, not adoption (Judge Guide clause 5.4.3). If your solution launched recently and you only have early usage data, say so honestly rather than overclaiming: a judge trusts an honest short curve more than a confident flat claim. An entry that lacks adoption evidence is not eligible.

3.3 Why these two carry weight. Criteria 1, 2, and 3 carry 60% of the score combined (Rules clause 8.2), and the responsible-AI evidence (Criterion 2) and adoption evidence (Criterion 3) sit inside that 60%. The rubric rewards durable change and responsible practice over model novelty, so these two must-haves are where a strong pack is won or lost.

4. Access and integrity

The access, no-links, and length rules are not red tape; they are one anti-gaming mechanism, and its job is to make volunteer judging feasible and fair. A juror who cannot open your evidence cannot credit it, and a juror comparing equal-shape dossiers can judge on merit rather than on production budget.

4.1 Accessible without login (Rules clauses 6.5 and 4.2.5). Any material a nomination relies on, including the video link and any permitted supporting link, must be accessible without a login, or the entry may be disqualified. This rule is adopted from the PMI global program. The most common way strong entries fail is a link a juror cannot open: a document that asks for a sign-in, a request-access step, or an account. Before you submit, open every link in a signed-out browser window, where you are logged in to nothing, and confirm each one loads.

4.2 No dead links and no login-gated material. A link that has expired, that points nowhere, or that opens a login wall is treated as material the judges cannot reach. Do not rely on a document behind a corporate sign-in, a share link set to expire, or a file that needs an account to open. Set any shared file to “anyone with the link can view” and check that setting the day you submit. If a juror cannot open it without logging in, it does not count, and it may disqualify the entry.

4.3 The form carries the evidence; no links in the core content (Rules clause 6.6.1). The substance of your answers lives inside the nomination form, not behind a link. Do not move the core content of an answer into an external link. The evidence pack artifacts and the short video are the permitted supporting material outside the form body, and both must be accessible without a login (clause 4.1). This is the PMI-GO no-links-in-form rule, applied in a lighter form, and it stops an entry from being a thin form pointing at a pile of external pages.

4.4 Human-authored (Rules clause 9.2). Entries must be human-authored. AIPM reserves the right to use detection tools, and responsible use of AI tools by nominees is disclosed where permitted by the final rules. Write the entry yourself, in your own words, and let the evidence, not generated prose, carry the weight. For an AI award, an AI-padded submission is both a credibility risk and exactly the wrong signal.

4.5 Audit rights (Rules clause 6.4). AIPM may verify any information in an entry and request supporting evidence during evaluation. Material misrepresentation disqualifies the entry. Assemble the pack so that every figure and claim can survive a check.

4.6 Curate, do not dump. A tight, judge-readable pack that substantiates each claim beats an exhaustive archive. Judges score what they can verify against the criteria (Rules clauses 8.5 and 9.4), so every item you include should earn its place by backing a specific claim in the questionnaire.

5. Authorization and consent explained

Four items make your entry genuine and complete. None of them costs anything, and all of them make the entry verifiable. In plain language:

5.1 The signed authorization (Rules clause 6.2). You must be authorized to submit the project and to consent to case-study publication. If the project belongs to an organization, a person with authority provides a signed authorization. This is lighter than a director-signature-plus-stamp requirement (no company stamp, no notarization; it can be a manager or a sponsor with authority over the project, not necessarily a director), but it must be genuine and verifiable: a real person with real authority, whom AIPM could contact to confirm. For a self-nomination where you personally hold the authority, you may sign it yourself. You capture this in Section 8.1 of the form.

5.2 The truthfulness declaration (Rules clause 6.3). As the Lead Nominator, you declare that the information in the entry is true and accurate to the best of your knowledge, and that your evidence genuinely supports your claims. You are the single owner of the entry and hold responsibility for its truthfulness (Rules clauses 2.2 and 5.1). Paired with the audit right (Rules clause 6.4), it means a padded or invented claim is not a clever edge; it is a disqualification risk. You give it in Section 8.2 of the form.

5.3 The case-study consent (Rules clauses 4.1.7 and 12.1). Publishable case-study consent is required at entry, because the community-learning value of the award depends on being able to publish the winning and finalist stories. It comes with confidentiality safeguards: you may withhold specified confidential details (particular figures, client names, or sensitive specifics) while still consenting to publication. What you may not do is withhold consent itself, or withhold the required responsible-AI evidence, adoption data, or human-oversight evidence in full: an entry without case-study consent is not eligible (Rules clause 4.2.7), and an entry without the required evidence is not eligible (Rules clause 4.2.6). Consenting while protecting a few specifics is normal and expected. You give this in Section 9.1 of the form.

5.4 The data-forwarding consent (Rules clause 12.2). You consent to your regional entry's data being forwarded to the global jury, so that an entry promoted from its region can be re-scored in the global round (Rules clauses 7.6 and 8.7). This clause is adopted from the PMI-GO consent-to-forward-data clause. It concerns your data reaching the global panel, not your regional score: regional scores are not carried forward. You give it in Section 9.2 of the form.

6. Length and layout discipline

The length and layout rules are the anti-gaming spine of the pack. They exist for one reason: fairness and comparability. When every entry is the same shape and roughly the same length, a volunteer juror can compare them on merit rather than on production budget, and a strong project with a plain write-up is never beaten by a longer or glossier

one. These constraints are adapted from the PMI-GO approach in a lighter form than its 22-page ceiling (Rules clause 6.6).

6.1 The recommended page budget (working team proposal; see Rules clause 6.6.5 and nomination-form clause 1.9). Rules clause 6.6.5 fixes the principle that a page limit applies and delegates the number to the nomination form.

RECOMMENDED OPERATIONAL VALUE (working team proposal, pending forum and subgroup confirmation). The completed nomination form (all sections, after the bracketed guidance is deleted) should not exceed **10 pages**. Rationale: meaningfully lighter than the PMI-GO 22-page ceiling (under half of it), so a global, volunteer-judged first edition stays comparable and readable. The evidence-pack artifacts are separate attachments and are not counted inside this 10-page form budget. The PMI-GO 22-page figure is a principle reference only, not the source of our number. Our budget is lighter and its own. This is the working team’s recommendation, not a settled figure; the forum and the 2026-07-15 subgroup confirm the number. The nomination form (nomination-form.md, clause 1.9) is the authoritative place this number lives.

6.2 The per-section length budget (working team proposal; Rules clause 6.6.3). Each questionnaire section in the form has its own word budget, so no section can crowd out the others. Stay within each budget; content beyond a section’s budget may not be read. The recommended budgets, set on the form, track the criterion weights (the heavier-weighted criteria are given more room):

Criterion	Weight	Recommended length budget
1. Outcomes, benefits and impact	25%	up to 450 words
2. Responsible AI governance and Human-in-Command	20%	up to 350 words
3. Adoption and change management	15%	up to 300 words
4. AI solution fit and innovation	15%	up to 250 words
5. Project management and delivery rigor	15%	up to 250 words
6. Learning and contribution to the community	10%	up to 200 words

These per-section budgets are the working team's recommended operational values (pending forum and subgroup confirmation), set on the form under the same authority as the page budget. They match the budgets printed in `nomination-form.md`, Section 5.

6.3 The recommended maximum video length (working team proposal; see Rules clause 6.6.5 and nomination-form clause 7.2).

RECOMMENDED OPERATIONAL VALUE (working team proposal, pending forum and subgroup confirmation). The short video should be at most **3 minutes** long. Shorter is fine. Rationale: a short, lightweight submission video, consistent with the video-first submission norm across PMI, APM, PMOGA, and Clube da Programacao, and deliberately shorter than the longer Tier-2 finalist pitch (Rules clause 7.4.2), so the judging load stays light. The rationale is the video-first norm, not a sourced 3-minute figure. This is the working team's recommendation, not a settled figure; the forum and the 2026-07-15 subgroup confirm the number.

6.4 Delete the bracketed guidance (Rules clause 6.6.4). Every bracketed guidance block in the nomination form is instruction, not part of your entry. Delete every such block before you submit. Leaving guidance text in the form both breaks the fixed layout and eats into the page budget.

6.5 Fixed layout (Rules clause 6.6.2). Do not re-order, restructure, or restyle the form. The fixed layout is what lets a juror find the same thing in the same place across every entry.

6.6 The evidence pack is separate from the form budget. The page budget governs the nomination form. The evidence-pack artifacts (business case, charter, baseline versus actual, and the rest) are attachments referenced from the form; keep each artifact focused and readable, and redact any confidential detail you are withholding under the case-study consent (clause 5.3).

6.7 Why the discipline exists: comparability by design. Length and layout limits, the no-links rule, and the login-free rule together do the same job: they make the dossiers equal-shape and equally readable, so judging turns on the quality of the project and its evidence, not on who had more time, more design help, or a bigger link farm. This is the lighter, AIPM version of the PMI-GO comparability-by-design approach.

7. Do and do-not checklist

Use this one-page checklist as the final pass over your pack.

Do

- Point every claim in the form to a specific artifact in the pack; tie it to a document a juror could open and verify, the auditor test (Section 1.2).

- Lead with baseline versus actual for outcomes, and with realized (not projected) benefits (Section 2.1).
- Show responsible-AI governance operating: the review that was done and the human sign-off point (Sections 2.2 and 3.1).
- Show sustained use over time, not just launch, for adoption; if you only have early data, show an honest short curve (Sections 2.3 and 3.2).
- Justify why the AI approach fits the problem, with performance evidence (Section 2.4).
- Include the delivery spine (charter, plan, risk log, governance cadence) and lessons learned (Sections 2.5 and 2.6).
- Include the three required artifacts: AI governance or ethics review, adoption data, and human-oversight evidence (Section 3).
- Test every link, and the video, from a signed-out browser; confirm nothing is login-gated or dead (Section 4).
- Keep the pack curated and judge-readable; include only what backs a claim (Section 4.6).
- Complete the signed authorization, the truthfulness declaration, and both consents (Section 5).
- Keep the form within the recommended 10-page budget, each section within its word budget, and the video within the recommended 3-minute maximum (Section 6).
- Delete every bracketed guidance block before submitting (Section 6.4).

Do not

- Do not assert impact with no number and no source; unsupported assertions score in the 1 to 2 band (Section 1.2).
- Do not confuse deployment with adoption; launch-week numbers alone are not adoption (Section 3.2).
- Do not rely on a fashionable technique with no rationale or fit, and do not lead with a model name in place of a rationale for fit (Section 2.4).
- Do not put the core content of an answer behind an external link (Section 4.3).
- Do not submit a link that needs a login or that has expired (Section 4.2).
- Do not withhold the required responsible-AI evidence, adoption data, or human-oversight evidence in full (Sections 3 and 5.3).
- Do not exceed the per-section word budgets or the form page budget (Section 6).
- Do not leave bracketed guidance in the submitted form (Section 6.4), and do not restyle or re-order it (Section 6.5).
- Do not submit AI-ghostwritten prose; the entry must be human-authored (Section 4.4).

End of evidence-pack guide (edition 1, Wave C release candidate). Derived from operating-rules.md; the operating rules prevail on any conflict, and this guide is a companion to

nomination-form.md and a sibling of judge-guide.md and scoring-sheet-spec.md. The page budget (clause 6.1), the per-section word budgets (clause 6.2), and the maximum video length (clause 6.3) are the working team's recommended operational values under Rules clause 6.6.5, pending forum and subgroup confirmation; the 180-character first gate is fixed by Rules clause 5.3. Ships to nominators as a PDF companion to the editable nomination-form DOCX. Pending fold-in of the subgroup comments due 2026-07-15 and ratification by the founders' forum.