

An archive, with today's rules.

All colored seals are creator-declared. Only issued Gold seals indicate independent review.

Historical draft attached: Standards review and edge cases. This notice and the current usage guide supersede conflicting wording in the original pages that follow.

CD stands for Creator Declared. Removing CD reflects directory completion only. It does not change the declaration's creator-declared status or indicate independent review.

Directory or preview state	Seal and readable status
No entry, or an incomplete entry	Colored seal retains CD. Creator-declared; not independently reviewed.
Complete directory entry; email unverified	Linked colored seal retains CD. Creator-declared; not independently reviewed.
Complete directory entry; email verified	Linked colored seal omits CD. Creator-declared; not independently reviewed. Email verified.
Unissued Gold preview	Unissued preview; not an awarded seal. No review or approval is implied.

Email verification confirms access to an email address, not identity or creative-process accuracy. A public contact email is optional and separate from private verification.

Only an officially issued Gold seal indicates independent review of the identified work and claim. Gold applications and issuance are not open. Standard seals remain free after the seven-question reflection, without registration or email. Black and white standard variants have the same status.

Keep readable status text beside seals and provide accessible image descriptions. Registered directory entries publish all seven submitted answers. Earlier proposals for private answers, different terminology or Gold procedures are historical, not current policy.

Current guide: [The Human Touch seal language and usage](#).
The unchanged historical draft begins on the next page.

The Human Touch Standards Review

9 October 2026 | *Recommendations for the next standards revision*

Executive verdict

Keep the free, creator-selected model. The standards have a workable purpose and several sound safeguards, but their category boundaries are not yet dependable enough for consistent public use. The largest risk is a truthful claim about one aspect of a work creating a false impression about the whole work.

The strongest revision separates physical making from generative AI provenance. Use one primary process label for each clearly identified work or component, with scoped qualifiers for handwork and the stronger no-generative-AI claim. Treat the categories as descriptions of process, without ranking the creator or the work.

The artist chooses the stamp after answering guiding questions and receiving a recommendation. The published meaning still constrains that choice. Routine use should not require advance approval, uploading evidence, paid software, or joining a directory. Ask for justification only when a specific challenge warrants it. [S3, S5]

What should change first

- Define what the stamp covers: the whole item or a named component, and which version. A badge on a book must not quietly mean only its binding.
- Separate human creative direction from final review. The proportion of generated pixels, words, or production time cannot decide this by itself.
- State the generative AI boundary plainly and judge the function actually used. Non-generative AI and accessibility tools are not automatically disqualifying.
- Keep evidence requests limited, private, free, and tied to an actual disputed criterion. A detector score or missing draft is not a finding of misuse.
- Describe use as creator-declared. A process stamp must not imply independent certification, safety, accuracy, originality, or rights clearance.

Which decisions this review applies

The baseline is the May 1 working first drafts: five mark pages, the retained overview, approved Usage Rules, and the Complaint / Misuse Process with its approved evidence amendment. The homepage was started but not locked. The final Human-Finished wording and the corrected Human-Directed tagline are used here. [S1-S4]

Earlier October 9 directions cover two questions: Handmade and AI-Free belong in the same category, with the last category described as substantially to wholly AI-created; and artists self-select, justify on challenge, and may be asked to change or stop a stamp, with directory notation. The first direction still needs precise category wording. The second is the operating model assessed here. All replacement wording and outcomes below are proposals, not silently adopted rules. [S5] A later request adds optional gold certification; its separate implications appear in the addendum.

Priority findings on category boundaries

1 Handmaking and AI provenance answer different questions

Problem. Handmade describes material human handwork. The digital 100% AI-Free mark excludes generative AI. Calling them the same promise does not make their present tests equivalent. [S1, S2, S5]

Example. An artist hand-paints an AI-generated composition. The handwork can be real while the composition is not generative-AI-free.

Smallest correction. Keep separate handwork and provenance claims. If merging them as directed, explicitly require a shared no-generative-AI process and make Handmade a physical variant that also requires material handwork. Record that as a substantive tightening, not a clarification.

2 The claim unit is undefined

Problem. “The work” can mean an object, cover, text, performance, production step, or whole package. Changing medium can also change the apparent category without changing who created the substance. [S1, S3]

Example. A human-written book has an AI cover and hand binding. A single unqualified badge appears to describe all three.

Smallest correction. Make component limitations visible beside the claim; keep the version traceable through a linked or otherwise accessible declaration. Without an evident limitation, the stamp describes the whole item. Printing a PDF or framing an image does not reset its provenance.

3 Generated quantity cannot settle creative role

Problem. “Primarily,” “most,” and “substantially to 100%” lack a coherent shared denominator. Counts of words, pixels, prompts, or hours can misdescribe a small but defining contribution. [S1, S5]

Example. Every image pixel is generated, but a person supplies a detailed composition and controls its execution. Conversely, a short generated refrain can define an otherwise human-written song.

Smallest correction. Assess the source of significant expression and structure, then the human role. Explain why a contribution matters. Do not introduce numeric thresholds or an effort test. Mostly or wholly generated substance can still be meaningfully human-directed.

4 Approval alone lets two marks collapse together

Problem. Human-Directed lists approval among qualifying actions, while Human-Finished centers on checking and approving. A glance can be presented as creative direction. [S1]

Example. A seller accepts the first generated logo and says that choosing to publish it reflects their intent.

Smallest correction. Reserve Directed for identifiable creative decisions affecting the result beyond acceptance and release review. Finished applies when the human role is principally review and correction as needed. One effective detailed instruction may qualify; many rerolls do not automatically qualify.

Priority findings on creative contributions

5 Ideas and structure can survive without copied wording

Problem. The standards do not clearly distinguish retained generative structure from rejected suggestions. AI-Free also makes a stronger full-creative-process promise than a final-output claim. [S1]

Example. A novelist rewrites every generated sentence but keeps the generated plot. Another rejects AI topics and writes independently.

Smallest correction. Count retained expressive or structural contributions, not just matching text. Under the strict AI-Free process promise, creative brainstorming with generative AI still excludes that claim even if rejected. Any narrower output-only promise must be adopted and named explicitly.

6 Manual recreation can conceal generated substance

Problem. Physical effort or a change of medium can be mistaken for human-origin authorship. [S1]

Example. A person traces an AI image, retypes an AI essay, or performs an AI composition and claims the substance was human-created.

Smallest correction. State that copying, tracing, transcription, translation, and medium changes alone do not erase provenance. Substantive transformation may change the creative-role assessment, but cannot erase earlier generative use for a full-process AI-Free claim.

7 Handmade has both an exclusion and a token touch loophole

Problem. “Did I make every choice” excludes permitted templates and collaborators, while “finished” can admit a trivial final touch. Literal hands-only language can exclude genuine adaptive craft. [S1]

Example. A seamstress follows a pattern and fails the literal test; a reseller ties a ribbon on a factory bag and claims the whole bag is handmade.

Smallest correction. Ask whether substantial hands-on or adaptive manual craft shaped the claimed piece beyond handling, packaging, or token finishing. Allow patterns, machines, collaboration, and adaptive methods. Claim only the actual craft contribution.

8 Teams and incorporated assets lack a clear boundary

Problem. “Myself” can exclude honest teams, while a commissioner may imply personal making or ignore unknown asset provenance. [S1, S3]

Example. An agency owner uses a personal Handmade claim for employees’ work; a designer adds stock art without knowing how it was produced.

Smallest correction. Allow human teams and describe the claimant’s role accurately. Include commissioned and incorporated content within the declared scope. Ordinary contributor explanations can support a claim; unknown provenance calls for a narrower claim or an explicit unknown, not an accusation.

9 Tool labels are an unstable shortcut

Problem. Ordinary, advanced, automated, and AI are not interchangeable. Finished includes templates and automation, yet its tagline says AI created the work. AI-Free is named more broadly than its generative-only definition. [S1, S5]

Example. A human-designed non-AI layout template is treated as AI-created, while a familiar editing app generates a replacement sky.

Smallest correction. Classify the operation used, not the brand. Distinguish capture and faithful assistance from generated or materially reconstructed expression. For the proposed mostly-AI final category, remove non-generative templates from its definition. Make “generative AI” visible in the claim or change the name deliberately.

Priority findings on trust and evidence

10 Human review sounds like a reliability guarantee

Problem. The final draft says the person “made sure it was safe, appropriate, and close enough to the intended purpose.” Review cannot establish an unlimited guarantee. [S1, E1]

Example. A publisher checks an AI health leaflet’s tone; buyers infer its advice is safe.

Smallest correction. Require purpose-appropriate review and correction of identified problems, not needless edits. Disclaim safety, accuracy, suitability, legality, originality, and rights clearance. If review alone qualifies, show “AI-generated; human-reviewed” beside Human-Finished, or consider Human-Reviewed; “Finished” can imply alteration.

11 Certification terminology conflicts with self declaration

Problem. The approved Usage Rules already forbid implying official certification without a separate program. Calling the current stamps “certification marks” can imply a different institutional role. [S3, S5, E1, E2]

Example. A self-selected seal appears next to “verified quality,” while the disclaimer sits on another page.

Smallest correction. Use “creator-declared process marks” for the present model and identify that status near the stamp. Do not claim the organization reviewed a work unless it actually did under a defined separate process. Legal mark type and registration require separate advice; this review establishes neither.

12 Detectors and metadata can become false verdicts

Problem. The rules do not explicitly prevent treating a score, appearance, or missing provenance as proof of the creative process. [E3, E4]

Example. A human poet receives an adverse AI score; a generated image with credentials is accepted automatically.

Smallest correction. No detector score, style impression, or absence of metadata alone can establish compliance or misuse. Treat provenance records as optional corroboration. Turnitin warns of misidentification and unsuitable non-prose uses; C2PA credentials do not establish factual truth or automatic trust.

13 An evidence standard can quietly become a paywall

Problem. The existing simple-explanation option is good, but missing-record and access protections are not explicit. [S3, S4]

Example. An offline writer overwrote drafts. A potter cannot buy a recording setup or keep a full process archive.

Smallest correction. Keep core use and reconsideration free. Require no paid tools, account, directory membership, continuous recording, prescribed evidence format, or retrospective records. A credible explanation may suffice; ask for more only to resolve the specific challenge. Missing records alone are not misuse.

14 Evidence can be authentic yet irrelevant

Problem. A real record does not necessarily connect the claimed process to the particular work or cover earlier incorporated material. [S4, E4]

Example. Workshop photos show another item; an editing history starts after imported AI artwork.

Smallest correction. Check consistency with the exact claim. Ask for the missing connection and accept alternatives. No evidence type is an automatic pass; do not demand the creator’s entire digital history.

Priority findings on fair challenges

15 Open reporting can become repeated harassment

Problem. The approved requirement for a reason and supporting evidence is useful, but it does not set a threshold for burdening the creator with a response. [S4, S5]

Example. A competitor reports fifty items because they “look AI” and repeatedly resubmits the same allegation.

Smallest correction. Accept reports, then triage for a specific plausible conflict with a published criterion before asking for justification. Combine duplicates; reopen only for meaningful new information. The existence or volume of complaints is not evidence that the claim is false.

16 Justification can expose unrelated private information

Problem. A broad proof request can pressure a creator to reveal confidential client work, personal data, or medical information. [S4]

Example. A film editor is asked for an unreleased client project; a disabled maker feels required to disclose a diagnosis to explain an adaptive process.

Smallest correction. Accept redacted samples, private explanations, and alternatives. Request only information needed for the criterion, never passwords or medical details. State who can see evidence and the retention period before receiving it; do not publish it without permission.

17 Uncertainty can be turned into a public accusation

Problem. Directory notation after a failed justification can become a person-level blacklist unless findings, response rights, and correction paths are defined. [S4, S5]

Example. A creator cannot recover old drafts or misses an email. The directory calls them dishonest even though misuse was never established.

Smallest correction. Separate clarified, corrected, unresolved, unable to substantiate, and established misuse. Give criterion-specific reasons, a reasonable response period with extensions, and free reconsideration. Note the work, stamp, version, and factual outcome; do not equate silence with fraud. Use another reviewer when feasible and recuse conflicts.

18 Silent revisions can punish previously permitted conduct

Problem. A changed definition can make existing labels appear deceptive, while different reviewers apply different unstated rules. [S3, S5]

Example. Packaging printed under the May handwork-only test is judged against a later merged no-generative-AI rule.

Smallest correction. Publish dated standards and a short change log, distinguish clarification from substantive change, and set a reasonable transition. Assess past conduct under the applicable version; correct present claims where necessary without retrospectively calling earlier permitted use dishonest. Publish reusable anonymized examples.

Keep the safeguards already chosen

The May rules already permit simple explanations, reasonable process notes, multiple marks when helpful, correction of mistakes, and escalation only when necessary. They explicitly reject awards, rankings, moral judgments, and unearned certification claims. Strengthen those provisions rather than replace them with routine audits, mandatory registration, or a costly enforcement system. [S3, S4]

A cease-and-desist request is a possible escalation in the existing policy, not proof of legal authority or a binding ruling. Specific legal enforcement, trademark registration, and local consumer-law compliance need qualified review before action.

Proposed guiding questions and recommendation

This is advisory guidance for creator selection, not pre-approval. Apply it to a named work or component and version. Where a whole-item stamp would conceal material differences, use component-specific claims. These operational boundaries are proposed revisions, not the present published text.

1 What exactly are you describing

Name the text, image, composition, recording, performance, physical object, binding, or other component. Identify the version. List material incorporated contributions, including collaborators, stock assets, commissions, and generated content. Administrative tasks unrelated to making the claimed work fall outside that scope.

2 Where did the creative substance come from

Consider significant words, imagery, sound, design, and structure. Did people originate them, or did generative AI supply content that remains materially present? Retained generated structure counts even when wording changes. Ordinary capture, faithful transcription of the creator's own words, and non-generative correction do not automatically turn human work into AI-generated work. Check ambiguous reconstruction features by their actual use.

3 Which primary process label best describes the human role

Human-Made: people are the primary creators of the scoped substance and meaning. Preserve the May restriction against routing generative production of the main substance here. Do not quietly introduce a broad allowance for ancillary AI; mixed components should be identified separately until that boundary is explicitly adopted.

Human-Directed: generative contribution is present and the person can describe substantive creative decisions that affected the final result beyond merely accepting and approving it. Direction can be effective in one step or many; no prompt count or labor minimum is required.

Human-Finished: generative AI supplied most or all significant creative substance and the person's main role was actual final review, correction where needed, and release approval. This follows the proposed AI-produced emphasis. Without meaningful review, the work does not qualify for Finished. Use a plain AI-generated disclosure if none of the marks fits.

4 Which scoped qualifiers can you support

Handmade: add only for substantial physical craft, including adaptive manual methods, in the component identified. Token finishing or packaging does not qualify the whole object. Under the current handwork definition, this alone says nothing about AI provenance.

No generative AI: add only if the adopted process exclusion holds throughout the defined creative process and incorporated material. The May 100% AI-Free mark is digital/nonphysical and strict about creative-process use. A medium-neutral qualifier, a new name, or narrower scope would be a deliberate policy change.

5 Would the visible claim mislead an ordinary reader

Explain material exclusions beside the stamp, rather than hide them behind a link. If provenance is unknown, narrow the claim or say so. Choosing no stamp is legitimate. The marks do not establish copyright ownership or protection; U.S. copyright guidance applies a separate, case-specific authorship analysis. [E1, E5]

Proposed operating policy and choices to settle

A small amount of guidance at entry

The artist selects a stamp using the questions and examples. Publish the standard free of charge and show “Creator-declared under The Human Touch standard” near use. Keep the physical or digital format, component, and version clear. There is no routine proof submission, prior review, or directory enrollment requirement. Keep versions traceable through an accessible linked declaration rather than crowding small labels. A short process note is useful preparation, not an entry fee.

A bounded process only after a challenge

First, identify the claim and the published criterion allegedly in conflict. Accept the reporter’s reason and available evidence; reject or consolidate vague and duplicate allegations without repeatedly burdening the creator.

If a plausible specific conflict remains, privately ask the creator for an explanation. State the issue, the rule version, what would resolve the question, and a reasonable response period. Allow extensions and accessible response methods. Start with the least burdensome adequate evidence and accept redactions or alternatives.

Record the outcome accurately: claim clarified; voluntarily corrected; claim unable to be substantiated; or misuse established on the available evidence. A claim can need changing without its creator being a liar. Ask for a different stamp, a scoped clarification, or cessation when justified, as the user has directed. [S5]

Before a contested adverse directory entry, give the creator the proposed wording and a chance to respond or seek free reconsideration. Prefer a separate reviewer where feasible. Record only the relevant work, stamp, version, date, reasons, and correction status. An unresolved allegation should not become a public fraud label; disclose public information only as necessary and without confidential evidence.

A standardized public directory summary

The latest direction is a standardized public summary, with the entry format deferred. Summarize the scoped claim, review or challenge outcome, and relevant correction accurately. Do not publish raw answers or private evidence by default. An unresolved allegation must not become an unproven fraud label. Provide a way to correct obsolete adverse information; the detailed entry design can wait.

Three drafting decisions remain

1. Category relationship. Preferred: separate authorship/process from handwork and no-generative-AI qualifiers. If Handmade and AI-Free are one category, expressly adopt a common no-generative-AI promise, with a material-handwork condition for the physical variant. This removes some currently qualifying Handmade uses.
2. The strict claim. Decide whether to retain the May full-creative-process exclusion, including rejected AI brainstorming, or adopt a narrower final-contribution claim. Do not keep the same absolute wording while quietly narrowing what it promises.
3. Final category wording. Confirm whether the proposed substantially-to-wholly-AI-created category is Human-Finished. Describe both generated substance and the human review role. Preserve Human-Directed for substantive direction even when all output is generated. If review-only use remains eligible, clarify the badge with “AI-generated; human-reviewed” or choose Human-Reviewed.

Edge case tests for making and authorship

Expected outcomes below test the proposed clarifications. They are recommendations to help the creator choose, not automatic certification decisions. “No-gen” means the defined no-generative-AI creative-process claim, not absence of every kind of AI.

1 A pattern based garment

A person cuts, sews, and fits a garment using a commercial human-designed pattern. Substantial craft can support Handmade; authorship claims should describe the maker’s role rather than imply they designed the pattern. No-gen needs a reasonable basis for the pattern and other in-scope contributions, without exhaustive supplier audits.

2 A generated painting made by hand

An artist paints an AI composition on canvas. Handwork can qualify under the current Handmade test; no-gen cannot. Directed may fit if human creative decisions meaningfully shape the image; manual copying alone does not establish that. Under the proposed merged strict category, the work would be excluded despite genuine handwork.

3 A machine made object with a ribbon

Adding packaging or a trivial finish does not justify Handmade for the whole item. A genuinely substantial hand-painted or hand-carved component may support a narrowly stated craft claim. The rule concerns what the maker materially did, not whether any machine was used.

4 An accessible creative workflow

Faithful dictation of the writer’s own words, eye-controlled drawing, or an adaptive craft tool can preserve human creation. Generative rewriting that supplies retained expression needs disclosure under the same rule applied to everyone. No diagnosis or disability evidence is required to assess the process.

5 A photograph with a generated sky

AI autofocus or non-generative noise reduction does not automatically disqualify a human photograph. Generated sky replacement changes content and defeats a no-gen claim for the final image. Directed may fit the composite. Restoration that reconstructs plausible detail needs a disclosed, function-specific assessment rather than a blanket software ban.

6 The same book in three formats

A human-written ebook, machine-printed paperback, and hand-bound copy share the same text provenance. Only the hand-bound component has a handwork claim. Printing alone is not handmade. If the cover is generated, label the text and cover separately or plainly disclose the mixed package.

7 A human manuscript with AI brainstorming

If suggestions are rejected and no generated expression or structure remains, Human-Made may fit after explaining the human-origin substance. The strict May full-process AI-Free claim does not. If the generated plot survives, rephrasing every sentence does not erase it; Directed may be the clearer recommendation.

8 A team project with a supplier asset

A studio can describe the team’s human work without its owner claiming personal handmaking. A supplier statement may reasonably support asset provenance. If an incorporated asset remains unknown, do not make a whole-item absolute no-gen claim; narrow the claim or state the uncertainty.

Edge case tests for direction and challenges

9 A wholly generated but directed image

A creator defines a specific composition and makes consequential creative decisions that shape the result. Human-Directed can fit even when every output pixel is generated. A single effective instruction can matter more than a hundred rerolls. This process assessment does not establish copyright authorship.

10 Generated copy reviewed without edits

AI supplies the main text and a person actually reviews it for its stated purpose. Human-Finished can fit even if no corrections are needed. Clicking publish or glancing at the first line is insufficient. The badge cannot promise that all factual, legal, or safety problems were eliminated.

11 A non generative template

A person writes copy and places it in a human-designed layout template. Automation alone should not route the work into an AI-created category. Human-Made may fit the scoped content; credit and describe incorporated design accurately. A template license is not proof of no-gen provenance.

12 An important generated fragment

A short generated chorus or a small central image can materially shape the work despite a low numerical share. Use contribution-specific disclosure or an assisted-process label. There is no safe harbor based on word count, pixel area, or editing time.

13 A poet with an adverse detector score

A score and stylistic suspicion alone do not establish misuse. First check whether a specific standards conflict has been identified; Turnitin itself warns about unreliable detection for poetry. If clarification is warranted, accept a credible account and reasonable alternatives to missing drafts. [E3]

14 A real archive for the wrong work

An impressive time-lapse, metadata file, or editing history is insufficient if it concerns another item or begins after generated material was imported. Ask for the relevant connection. Conversely, the absence of C2PA credentials is not a reason to presume a dishonest claim. [E4]

15 A confidential creator under challenge

A creator cannot disclose an NDA-covered project file. Accept redacted samples or a process explanation addressing the criterion. If the specific claim still cannot be substantiated, request a narrower claim or cessation where justified. Record uncertainty accurately; inability to disclose is not proof of deception.

16 A corrected or unanswered claim

An artist promptly changes a mistaken stamp: record correction rather than dishonesty. An unanswered email leaves an unresolved matter, not automatic proof of misuse. Any directory entry must identify the scoped claim and factual status after a fair response opportunity, with a route to correct it.

A practical consistency check

Before adopting revised wording, have several people independently recommend labels for these same cases and explain why. Reconcile disagreements in the text. Test how consumers interpret the visible badge, especially on small packaging and mixed works. Keep labels equally legible and avoid colors or ordering that imply a moral ranking.

Sources and decision history

The linked source messages document the working baseline and user decisions. Proposed tests, thresholds for opening a challenge, and directory safeguards are recommendations in this review. External sources support narrow factual points; they do not endorse this system or establish legal compliance.

Source drafts and decisions

S1. Mark pages, May 1, 2026: Human-Made; Handmade; 100% AI-Free; Human-Directed; final Human-Finished. Directed tagline correction

Human-Made, Handmade, and final Human-Finished received explicit save instructions. The AI-Free page was followed by “Next” and the pages were described as completed first drafts; this is not treated as a signed final standard. The corrected Directed tagline uses “advanced tools.”

S2. Mark System Overview and “Ok keep it” approval, May 1.

S3. Usage Rules and user approval, May 1. Free process marks, ordinary evidence, correction, multiple marks, and no implied certification.

S4. Complaint and Misuse Process; reporting evidence amendment; amendment approval, May 1.

S5. October 9 user directions: category relationship and AI-created final category; artist choice and challenge-only justification. The latter includes requests to change or cease use and directory notation.

S6. May 1 stopping point identifies the homepage as started but not locked.

Primary external references

E1. FTC Advertising FAQs. Explains contextual and implied claims, evidence, and prominent qualifications. Supports the caution about a badge's overall impression, not a finding that these marks violate law.

E2. USPTO Certification mark applications. Describes certification marks as indicating standards set by a certifying organization. This review makes no registration, ownership, or enforceability determination.

E3. Turnitin Using the AI Writing Report. Warns of misidentification, rejects sole-score adverse decisions, and states non-prose limitations. These are that product's stated limits, not an empirical assessment of every detector.

E4. C2PA Content Credentials Explainer version 2.4, sections 7.1.2 and 7.2.2. Presence or absence of credentials alone does not establish trustworthiness; provenance alone does not establish factual truth.

E5. US Copyright Office AI and Copyright Part 2, January 2025, sections II.C-F and V. Distinguishes assistive use, prompts, and human modification or arrangement. Copyright classification is separate from these proposed process labels.

External references checked 9 October 2026. This is a standards design review, not legal advice or certification of any work.

Addendum on optional gold certification

The October 9 request adds optional gold certification and later clarifies gold as a reviewed, fancy version of the creator's chosen category. The proposed queue may be first-come-first-served with no guaranteed completion date. Paperwork is explicitly free; other costs remain unsettled. This is a desired option, not an implemented program.

Gold must describe a specific review

Problem. Gold could look like a higher creative category or a general quality guarantee.

Example. A Human-Finished work receives gold and buyers assume independent safety or originality testing.

Smallest correction. Make gold a separate review-status overlay on any eligible process category. Publish exactly what evidence was checked, by whom, for which work and version, under which standard, and when. Certify only that scoped process claim; do not imply safety, accuracy, rights clearance, or superior artistic value.

A queue is not an award

Problem. A long queue can blur application, pending review, and completed certification, while old applications may concern obsolete work.

Example. An applicant displays gold before review, or receives it for a version later changed with generative AI.

Smallest correction. Reserve gold for completed successful review. Label applications pending without a gold stamp. If using first-come-first-served order, publish the rule and any exceptions; no undisclosed paid priority. Reconfirm the final work and version before deciding. Explain that timing depends on capacity and is not guaranteed.

Optional review must stay accessible and fair

Problem. Free paperwork can conceal other costs or create a burdensome second system.

Example. An applicant completes a free form, then discovers a mandatory payment or intrusive archive requirement.

Smallest correction. Keep ordinary stamps free and self-selected. Disclose any proposed costs before application and decide them explicitly; do not infer them from the free-paperwork promise. Use proportionate evidence and privacy protections. Offer reasons, correction, free reconsideration, conflict recusal, and proportionate suspension or withdrawal of gold when warranted.

A minimal launch checklist

Before offering gold, settle its exact claim, reviewer capacity, application scope, queue rules, costs if any, evidence handling, decision record, appeal route, and version-change policy. Keep the process small enough to operate honestly. A certification label or seal is not evidence that a legally compliant certification-mark program already exists; obtain specific legal advice before making that representation.

Source S7. October 9 optional gold certification request. These safeguards are proposed responses to that request, not adopted program rules.

Later source decisions: gold reinforces the chosen category; standardized public summary; directory format deferred. No palette or directory layout is proposed here.