

MARINE CONSERVATION CREDITS

MCC-850 — Amendment 2

Automated Completeness Report, Applicant Disclosure and Degraded Mode

Version 1.0 — Adopted 17 August 2026

Status: ADOPTED by the Governing Board on 17 August 2026 under the Article 5 procedure, decision D-2026-08-02 (MCC-BDR-2026-002). **In force from 17 August 2026.**

Preamble

This Amendment amends MCC-850 (AI Governance Standard, v1.0, 23 March 2026). Amendment 1 (29 April 2026) addressed AI visibility to VVBs. This Amendment addresses three matters arising from the operation of the AI Analysis Layer at trigger point 1 (Project Submission) under Article 7.

First, the analysis performed on submission is now executed automatically and returned within the submission request itself. Article 13 specifies the AI Review Record as an internal document, and Article 27 specifies what an applicant may see; neither names the applicant-facing extract, nor states the label it must carry. This Amendment names it the Automated Completeness Report and fixes its contents.

Second, Article 22 prohibits the AI Engine from sending communications to applicants, VVBs or external parties. The platform notifies an applicant when an automated completeness report is available. That notification is not an AI communication — it is generated by a deterministic, template-driven notification service on a recorded state change — but the boundary is load-bearing and must be stated rather than assumed, so that the prohibition is neither breached in fact nor read as breached by an auditor.

Third, Article 12 addresses AI Engine availability but not partial degradation. Where the language-model provider is unavailable, the platform produces a deterministic heuristic analysis in place of a model analysis. Such output is not an AI analysis within the meaning of Articles 14–19 and must not be counted as one, whether for the applicant, for the reviewer, or for the calibration statistics on which Articles 31–33 depend.

Operative Provisions

Part A — Definitions

The following definitions are inserted in Article 2:

"Automated Completeness Report" means the applicant-facing extract of an AI Review Record specified in Article 13.1

"Trigger Source" means the classification of an AI Review Record as machine-initiated (AUTO_SUBMISSION) or human-initiated (STAFF_REQUEST)

"Degraded Mode" means operation of the AI Analysis Layer under Article 12.1 in which a

deterministic heuristic substitutes for model analysis

"Prioritised Sequencing Flag" means the staff-facing indicator specified in Article 19.1

Part B — Automated Completeness Report

The following Article is inserted after Article 13:

Article 13.1: Automated Completeness Report

Where an AI Review Record is generated at trigger point 1, the platform shall derive from it an Automated Completeness Report for disclosure to the applicant. The Report shall contain, and shall be limited to: (a) the completeness score; (b) the per-section analysis under Article 17, comprising section name, presence, adequacy classification and identified gaps; (c) the confidence tier under Article 15, expressed categorically; (d) risk flags under Article 16, expressed as category and severity with a plain-language justification; and (e) suggested questions and clarifications under Article 13, expressed as actionable requests.

The Report shall carry, on every surface on which it appears and in every notification referring to it, a label in substance as follows: "This is an automated advisory report, not a review decision. All decisions affecting this project are made by a human reviewer."

The Report shall state the analysis timestamp and the model version, and shall state whether it was produced in Degraded Mode.

Part C — Applicant disclosure

Article 27 is amended by replacing the list of items withheld from applicants with the following:

Applicants shall NOT have visibility to:

(a) the AI Review Record in full, including the numeric confidence score, cross-reference internals under Article 18, and internal analytic commentary;

(b) the recommendation field under Article 19 (APPROVE, REQUEST_INFO, ESCALATE, REJECT), whether expressed as an enumerated value, a paraphrase, a colour, a badge or any other representation from which the value may be inferred. Only human decisions are communicated to applicants;

(c) the Prioritised Sequencing Flag under Article 19.1;

(d) AI model internals, prompts, algorithm details or training data.

The following clause is added to Article 27:

The applicant-facing surfaces of the platform shall be reviewed against this Article at each release affecting the project detail view, the applicant notification templates or the submission

response payload. Non-compliance shall be treated as an incident under Article 35.

Part D — Communications boundary

The following Article is inserted after Article 22:

Article 22.1: Communications Boundary

The prohibition in Article 22 on the sending of communications is a prohibition on the AI Engine. The AI Engine shall write no record other than the AI Review Record and shall invoke no communication channel.

Notifications informing an applicant, a reviewer or a VVB that an analysis is available are generated by the platform notification service, which is a deterministic, template-driven component. It is invoked by a recorded state change or record creation, not by the AI Engine, and it shall not transmit any item withheld under Article 27. The AI Engine shall have no capability to compose notification content, to select recipients, or to determine whether a notification is sent.

Every notification derived from an AI Review Record shall carry the label required by Article 13.1. Every notification shall be logged in the audit trail with recipient, template identifier, timestamp and the record from which it was derived.

Part E — Degraded Mode

The following Article is inserted after Article 12:

Article 12.1: Degraded Mode

Where the model provider is unavailable, unauthenticated, rate-limited or fails to return a valid response within the configured timeout, the platform may produce a deterministic heuristic analysis in place of a model analysis, in order that a submission is neither blocked nor silently unanalysed. Such operation is Degraded Mode and is subject to the following requirements.

12.1.1 Labelling. The record shall carry a model version identifying it as heuristic output, and shall carry a risk flag of category and label sufficient to identify Degraded Mode on every surface on which the analysis appears, to staff and to the applicant alike.

12.1.2 Confidence ceiling. A Degraded Mode record shall be assigned a confidence tier no higher than LOW, irrespective of the completeness score computed.

12.1.3 Recommendation floor. A Degraded Mode record shall not generate the recommendation APPROVE.

12.1.4 Exclusion from prioritised sequencing. A Degraded Mode record shall not generate a Prioritised Sequencing Flag under Article 19.1.

12.1.5 Exclusion from calibration. Degraded Mode records shall be excluded from the

accuracy, override-rate and bias metrics under Articles 31 to 33, and shall not be treated as AI analyses for the purposes of the recalibration triggers in Article 33. They shall be counted and reported separately under Article 30.

12.1.6 Escalation. Continuous Degraded Mode operation exceeding twenty-four hours shall be treated as an incident under Article 35 and reported to the Programme Administrator, and the Board shall be notified where live projects are affected, consistent with Article 12.

12.1.7 Re-analysis. On restoration of the provider, the Programme Administrator may cause a submission analysed in Degraded Mode to be re-analysed. The Degraded Mode record shall be retained under Article 20 and the re-analysis recorded as a new record referencing its predecessor.

Part F — Trigger source and prioritised sequencing

The following Article is inserted after Article 7:

Article 7.1: Trigger Source Labelling

Every AI Review Record shall record its Trigger Source. Records generated automatically on submission shall be labelled AUTO_SUBMISSION; records generated on the request of a human reviewer shall be labelled STAFF_REQUEST. Quarterly reporting under Article 30 shall report analysis volumes separately by Trigger Source.

The following Article is inserted after Article 19:

Article 19.1: Prioritised Sequencing Flag

An AI Review Record may carry a Prioritised Sequencing Flag where all of the following hold: (a) the completeness score is 95 or above; (b) no risk flag requires escalation; and (c) the record was not produced in Degraded Mode.

The Flag is visible to programme staff only. It is informative. It shall not change the status of any project, shall not constitute or substitute for a human decision, and shall not reduce any stage, consultation period, assessment or verification requirement. Its sole effect is priority in the order in which submissions are taken up for review, as provided by MCC-100 clause 0061.2.

The reference in Article 14 to eligibility for "potential fast-track (Board discretion)" at a completeness score of 95 or above shall be construed as eligibility for the Prioritised Sequencing Flag under this Article.

Part G — Service levels

The following clause is added to Article 8:

Where the analysis at trigger point 1 is executed within the submission request, the service level in this Article operates as an outer bound and is satisfied on completion of that request. Completion of the automated analysis is not a review step, does not start or shorten any period

allowed to a human reviewer, and does not advance the project.

Consequential Amendments

1. Article 3 is amended to add: "This Standard is read in conjunction with MCC-100 clauses 0050 to 0062, which govern the registration workflow within which AI analysis is performed, and clauses 0061.1 to 0061.3, which govern prioritised sequencing."

2. Article 30 is amended by adding to the contents of the quarterly AI Transparency Report: "(h) analyses by Trigger Source; (i) number and proportion of analyses produced in Degraded Mode, with cause and duration; and (j) number of Prioritised Sequencing Flags raised and the proportion subsequently registered, refused or returned for revisions."

3. Article 13 table: the description of the recommendation field is amended to add "Internal; withheld from applicants under Article 27."

Transitional Provision

Articles 13.1 and 22.1 apply to analyses generated on or after the entry-into-force date. Article 27 as amended applies from the entry-into-force date to all applicant-facing surfaces, including those displaying records generated before that date. Article 12.1 applies from the entry-into-force date; records generated in Degraded Mode before that date shall be identified and excluded from calibration statistics retrospectively where they can be identified from the model version field.

Entry into Force

This Amendment was adopted by the Governing Board on 17 August 2026 under the Article 5 procedure (decision D-2026-08-02) and **entered into force on 17 August 2026**, immediately upon adoption.

The Board departed from the thirty days notice contemplated by Article 5 and recorded its reason: this Amendment *restricts* disclosure and *tightens* obligations falling on the programme itself. It imposes no new burden on any applicant, VVB or other participant, and a notice period would have left a disclosure practice the Board had already determined to be wrong in force for a further month. The Programme Administrator shall publish the Amendment through the Transparency Layer without delay.

Editorial note: This Amendment shall be incorporated into MCC-850 v1.1 together with Amendment 1 (Repeal and Replacement of Article 28). The original Article 27 text shall be preserved in the version-control appendix of MCC-850 v1.1 in accordance with Article 11 and the standards-history requirements of MCC-100.