MARINE CONSERVATION CREDITS

MCC-850 — AI Governance Standard

Version 1.2 — Consolidated text
In force from 17 August 2026
Programme Standard – Normative (Tier 3)

Document control

Document
MCC-850 AI Governance Standard
Version
1.1 (consolidated)
Supersedes
MCC-850 v1.0, 23 March 2026
Consolidates
MCC-850 Amendment 2 — Automated Completeness Report, Applicant Disclosure and Degraded Mode, v1.0
Also consolidates
MCC-850 Amendment 1 — Repeal and Replacement of Article 28 (AI Visibility to VVBs), v1.0, adopted 17 August 2026 by decision D-2026-08-10. **Article 28 enters into force 17 September 2026**, not on adoption. See Appendix B.3.
Adopting decisions
D-2026-08-02 and D-2026-08-10, Governing Board, 17 August 2026, Article 5 procedure (record MCC-BDR-2026-002)
Entry into force
17 August 2026, except Article 28 and the Article 3 sentence on MCC-400 Section 14, which enter into force 17 September 2026
Consolidated on
17 August 2026
Next scheduled review
March 2027 (Article 45)
Status of this text
Consolidated text prepared for operational and due-diligence use. It reproduces MCC-850 v1.0 as amended. Where this consolidated text and the adopting instruments differ, the adopting instruments prevail.

HOW TO READ THIS TEXT

Article numbering follows MCC-850 v1.0. Articles inserted by amendment carry a decimal suffix (for example 12.1) and sit immediately after the Article they supplement; the numbering of the original Articles is unchanged. Text inserted by amendment is marked in the left margin. Indented notes in smaller type are editorial: they record how an amendment directs an existing provision to be read, or record the status of a provision, and are not themselves normative. Appendix B lists every change made to the v1.0 text, reproduces superseded provisions in full,

and records the entry-into-force position of each amendment. One provision — Article 28 — carries a later entry-into-force date than the rest of this text and is marked accordingly where it appears.

Structure: Part 1 Scope and Definitions (Articles 1–5) · Part 2 AI System Architecture (Articles 6–12, 12.1) · Part 3 AI Output Specification (Articles 13–20, incl. 13.1 and 19.1) · Part 4 Human Oversight Requirements (Articles 21–30, incl. 22.1) · Part 5 AI Performance Monitoring (Articles 31–38) · Part 6 Data Protection and Security (Articles 39–42) · Part 7 Transitional and Final Provisions (Articles 43–45). Article 7.1 sits in Part 2.

PART 1: SCOPE AND DEFINITIONS

Article 1: Purpose and Scope

This Standard establishes governance and operational requirements for AI systems deployed within the Marine Conservation Credits (MCC) programme.

The Standard ensures that:

- AI systems operate transparently within defined boundaries
- Humans retain final decision authority on all substantive matters
- AI outputs are auditable, explainable, and permanently documented
- AI performance is continuously monitored and improved
- Stakeholders understand AI involvement in MCC decisions

The guiding principle is: AI-First, Human-Final. AI accelerates analysis; humans decide.

Article 2: Definitions

For the purposes of this Standard:

- "AI Engine" means the software system responsible for document analysis, scoring, risk assessment, and recommendation generation as specified in Part 2
- "AI Review Record" means the complete output generated by the AI Engine for a given submission, including all

analysis, scores, flags, and recommendations

- "Confidence Score" means a numerical value (0–100) representing AI certainty in its analysis
- "Confidence Tier" means the categorical classification of confidence (HIGH, MEDIUM, LOW, VERY_LOW) derived from confidence score
- "Override" means a human decision that contradicts or supersedes an AI recommendation
- "Human Decision" means a substantive choice made by an authorized human reviewer (Programme Administrator, VVB, Technical Committee, or Board)
- "Completeness Score" means a numerical assessment (0–100%) of PDD compliance with methodology module requirements
- "Risk Flag" means an AI identification of a potential issue (permanence, additionality, leakage, baseline, quantification, or safeguards concern)
- "Cross-Reference Analysis" means AI comparison of submitted data across multiple sources for consistency verification
- "Model Version" means a specific version of the AI model, tracked throughout its operational life
- “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

Article 3: Relationship to Other MCC Standards

This Standard forms part of the MCC normative hierarchy, categorized as a Tier 3 Technical Standard.

The hierarchy is structured as follows:

- Tier 1: Board Strategic Direction (annual)
- Tier 2: Board Operational Policies (quarterly)
- Tier 3: Programme Standards (normative; this document)
- Tier 4: Methodology Modules (F01–F05; annually reviewed)
- Tier 5: Guidance & Operational FAQs

This Standard shall be interpreted consistently with and in support of the MCC Programme Standard v1.0 and MCC AI and Human Oversight Policy v1.0.

This Standard is read in conjunction with MCC-400 Section 18 (AI-Augmented Validation and Verification), which operationalises the AI Analysis Layer in the VVB engagement context. Where this Standard and MCC-400 Section 18 address the same matter, the more specific provision controls.

> *Inserted by Amendment 1 (D-2026-08-10); the section reference was corrected from Section 14 to Section 18 on adoption of MCC-400 Addendum 1 (D-2026-08-11). This sentence enters into force on 17 September 2026 with Article 28.*

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.

Article 4: Normative Status

This Standard is binding on all parties operating within the MCC programme, including:

- MCC staff and Programme Administrators
- Contracted Validation and Verification Bodies (VVBs)
- Project applicants and developers
- MCC Board and technical committees

All requirements are mandatory unless explicitly marked as recommended or aspirational. Non-compliance shall be documented and reported to the MCC Governance Board.

Article 5: Amendment Procedure

Amendments to this Standard shall follow the procedure below:

- Proposed amendment submitted to MCC AI Governance Committee
- Technical Committee reviews for impact assessment
- Public comment period (30 days) for stakeholder feedback
- MCC Board votes on approval
- Approved amendments published with effective date

Major revisions (requiring >10% substantive change) shall be published as new versions (e.g., v1.2, v2.0) rather than addenda.

Minor amendments (clarifications, correction of errors) may be published as addenda.

PART 2: AI SYSTEM ARCHITECTURE

Article 6: AI Engine Scope

The MCC AI Engine is a read-only analysis system with the following scope:

- Document Analysis: Text extraction and parsing from PDDs and monitoring reports
- Completeness Scoring: Methodology-based assessment of submission completeness
- Cross-Reference Analysis: Consistency verification across submitted data
- Risk Flagging: Identification of permanence, additionality, leakage, baseline, quantification, and safeguards concerns
- Recommendation Generation: Synthesis of analysis into actionable recommendations

The AI Engine operates on MCC-internal data only. It shall not:

- Initiate state transitions or modify registry data
- Send external communications or trigger notifications
- Access external APIs or web services without explicit

Board authorization

- Make autonomous decisions binding on the programme

Article 7: AI Trigger Points

The AI Engine is invoked at exactly seven trigger points in the project workflow:

- Project Submission: AI analyzes initial PDD for completeness, cross-reference, and risk
- Monitoring Report Submission: AI analyzes report for consistency and quantification verification
- Methodology Submission: AI checks compliance with F01–F05 requirements
- VVB Application: AI screens VVB qualifications and conflict-of-interest status
- Validation Report Receipt: AI cross-checks findings against initial PDD analysis
- Verification Report Receipt: AI compares reported reductions to baseline projections
- Issuance Request: AI performs final compliance check before credit issuance

At each trigger point, the AI Engine shall generate a complete AI Review Record and transmit it to the responsible human reviewer.

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.

Article 8: AI Processing Service Levels

The AI Engine shall meet the following processing time requirements:

- Project Submission: <4 hours from submission receipt to AI Review Record delivery
- Monitoring Report: <4 hours from submission receipt to AI Review Record delivery

- Methodology Submission: <8 hours from submission receipt to AI Review Record delivery
- VVB Application: <2 hours from application submission to AI screening result
- Validation/Verification Report: <4 hours from report receipt to cross-check completion
- Issuance Request: <1 hour from request to final compliance check

Failure to meet these SLAs shall be logged in the audit trail, reported to the Programme Administrator, and escalated to the Technical Committee if recurring.

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.

Article 9: AI Data Access Controls

The AI Engine shall operate under strict data access controls:

- Read-Only: AI may read MCC database records, submitted documents, and historical project data; AI shall not write to any record except AI Review Records
- No External APIs: AI shall not call external services (Verra API, NOAA databases, email services) without explicit Board authorization per integration
- No Web Access: AI shall not access public web content or perform web searches
- MCC Data Only: All data processed by the AI Engine derives from MCC-internal sources or applicant submissions
- No Third-Party Sharing: AI shall not transmit MCC data to external parties except as explicitly authorized by Board decision

The AI Service Account shall be configured with database read-only permissions and restricted to AI Review Record write operations. This account shall be audited quarterly for unauthorized access attempts.

Article 10: AI Model Governance

Any change to the AI model (logic, training data, algorithm, thresholds) requires the following governance:

1. Technical Specification: Proposal shall document what is changing, why, and expected impact.

2. Testing: New model shall be validated against reference dataset (50 historical PDDs) with minimum 90% concordance to human decisions.

3. Bias Assessment: New model shall be tested for systematic bias across geographies, sectors, and languages.

4. Fallback Plan: Rollback procedure to prior model version shall be documented and tested.

5. Board Approval: Model change requires formal Board vote. Decision shall be recorded in Board minutes.

6. Implementation: Model shall be deployed only after Board approval and 30-day notice to stakeholders.

Minor tuning (threshold adjustments, parameter optimization) within approved ranges may be delegated to the Technical Committee.

Article 11: AI Model Version Tracking

The AI Engine shall maintain strict version control:

- Every AI Review Record includes modelVersion field (e.g., "MCC-AI-v1.0")
- Model version change log maintained in programme documentation
- Version history retained permanently (audit trail)
- Comparisons of recommendations across model versions tracked for calibration

This enables historical audit: any decision can be traced to the specific model version that generated its AI analysis.

Article 12: AI System Availability and Disaster Recovery

The AI Engine shall meet the following availability and resilience requirements:

- 99.9% uptime during business hours (MCC operational timezone)

- Backup systems and failover capability to ensure continuity
- Disaster recovery plan tested quarterly; recovery time objective (RTO) <4 hours
- All AI Review Records backed up daily with geographic redundancy

In case of AI Engine failure lasting >24 hours, the MCC shall activate manual review processes and notify the Board within 4 hours.

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 3: AI OUTPUT SPECIFICATION

Article 13: AI Review Record Schema

Every AI Review Record shall include the following fields:

Field
Type
Description
Mandatory
id
UUID
Unique identifier for this AI Review Record
Yes
projectId
String
Reference to project being analyzed
Yes
workflowInstanceId
String
Reference to workflow instance (submission, monitoring, etc.)
Yes
triggerEvent
Enum
Which event triggered analysis (submission, monitoring, methodology, etc.)
Yes
modelVersion
String
Version of AI model used (e.g., "MCC-AI-v1.0")
Yes
analysisTimestamp
DateTime
ISO 8601 timestamp of analysis completion
Yes
completenessScore
Float 0-100
Overall completeness percentage
Yes
confidenceScore
Integer 0–100
AI confidence in analysis (0–100)

Yes
confidenceTier
Enum
Categorical tier (HIGH, MEDIUM, LOW, VERY_LOW)
Yes
recommendation
Enum
APPROVE, REQUEST_INFO, ESCALATE, REJECT. Internal; withheld from applicants under Article 27
Yes
sectionAnalysis[]
Object[]
Per-section breakdown (see Article 17)
Yes
riskFlags[]
Object[]
Identified risks (see Article 16)
Conditional
strengths[]
String[]
Key positive findings
Yes
concerns[]
String[]
Key concerns or gaps
Yes
suggestedQuestions[]
String[]
Questions for applicant or VVB clarification
Conditional
crossReferenceResults[]
Object[]
Cross-reference findings (see Article 18)
Conditional
humanDecision
Enum
APPROVED, REJECTED, OVERRIDDEN, PENDING
Updated by human
humanRationale
String
Human explanation if overridden (min. 50 words)
If overridden
reviewedById
String
ID of human reviewer who approved/overrode
Updated by human
reviewedAt
DateTime
Timestamp of human review
Updated by human

All AI Review Records shall be persisted in the audit database and

retained permanently.

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.

Article 14: Completeness Scoring Methodology

Completeness scores assess PDD compliance with methodology module requirements (F01–F05). Scoring shall follow this rubric:

Component	Weight	Scoring Criteria
Baseline Establishment	20%	Data sources identified, period documented, projection method stated
Quantification	20%	Formulas specified, emission factors cited, uncertainty quantified
Monitoring Plan	15%	Frequency defined, methods specified, QA process documented
Additionality	20%	Financial, regulatory, and common-practice tests addressed; evidence provided
Permanence	15%	Risk factors identified, mitigation mechanisms specified, buffer justified
Safeguards	10%	Nine principles assessed, high-risk areas flagged, mitigation planned

Scoring shall be deterministic and reproducible. AI shall generate detailed justification for each score.

Thresholds:

- <80%: Auto-flagged for high-scrutiny VVB pathway
- 80–95%: Standard validation pathway
- ≥95%: Eligible for potential fast-track (Board discretion)

Amendment 2 (D-2026-08-02): the reference to eligibility for "potential fast-track (Board discretion)" at a completeness score of 95 or above is construed as eligibility for the Prioritised Sequencing Flag under Article 19.1, which confers queue priority only (MCC-100 cl. 0061.2).

Article 15: Confidence Scoring Model (4-Tier System)

AI confidence is assessed on a four-tier model:

Tier
Range
Meaning
Human Response
HIGH
≥85
AI is 85%+ confident in recommendation
Can approve routine; may override with documentation
MEDIUM
60–84
Material uncertainty exists
Forward to VVB for clarification
LOW
40–59
Significant gaps; escalation required
Escalate to Technical Committee
VERY_LOW
<40
Cannot make confident recommendation
Board review (potential rejection)

Confidence is derived from multiple factors: data completeness, consistency, risk factor clarity, and methodological soundness.

Article 16: Risk Flag Categories and Severity Levels

The AI Engine shall identify and classify risks as follows:

Risk Category
Definition
Severity Levels
ADDITIONALITY

Project may proceed without carbon revenue
HIGH, MEDIUM, LOW
PERMANENCE
Stored carbon may be released early
HIGH, MEDIUM, LOW
LEAKAGE
Activity-shifting or market-level effects
HIGH, MEDIUM, LOW
BASELINE
Baseline projection may be inaccurate
HIGH, MEDIUM, LOW
QUANTIFICATION
Credit calculation may be incorrect
HIGH, MEDIUM, LOW
SAFEGUARDS
Social or environmental safeguards may be compromised
HIGH, MEDIUM, LOW

Each risk flag shall include: category, severity, justification, and suggested action.

Article 17: Section Analysis Specification

For each section of a submitted document, AI shall provide the following analysis:

- sectionId: Identifier (e.g., "baseline", "quantification")
- sectionName: Full name of section
- present: Boolean (is section included in submission?)
- adequacy: Classification (ADEQUATE, PARTIAL, INADEQUATE, NOT_ASSESSED)
- findings: Narrative summary of AI analysis for this section
- gaps: List of specific missing elements or weaknesses

This structure enables precise feedback to applicants and granular tracking of document evolution.

Article 18: Cross-Reference Specification

Cross-reference analysis shall document linkages as follows:

- methodologyRequirement: Which F01–F05 requirement is being verified
- pddSection: Which section of the PDD addresses it
- formField: Corresponding structured form field (if any)

- evidenceDocument: Supporting document (baseline data, financial model, etc.)
- matchStatus: MATCH, PARTIAL_MATCH, MISMATCH, UNVERIFIABLE
- confidence: AI confidence in the linkage (0–100)

Mismatches shall be escalated as concerns requiring clarification.

Article 19: Recommendation Types and Generation Criteria

The AI Engine generates one of four recommendations:

Recommendation	Criteria	Typical Confidence Tier
APPROVE	Completeness ≥90%, no HIGH risk flags, data consistent	HIGH (≥85)
REQUEST_INFO	Completeness 80–90%, minor gaps or LOW/MEDIUM risks	MEDIUM (60–84)
ESCALATE	Completeness <80% or multiple MEDIUM risks	LOW (40–59)
REJECT	Missing methodology sections, HIGH risk flags, or fraud signals	VERY_LOW (<40)

Each recommendation shall include explicit justification referencing analysis components.

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.

Article 20: AI Output Storage and Retention

All AI Review Records shall be stored as follows:

- Storage: MCC audit database (immutable append-only log)
- Retention: Minimum 10 years; permanent retention recommended
- Encryption: AES-256 at rest; TLS 1.3 in transit
- Linkage: Every record linked to final human decision and rationale
- Access: Available for internal audit, external audit, and regulator requests

No AI Review Record shall be deleted or modified after creation; corrections shall be recorded as new records with explicit reference to predecessor.

PART 4: HUMAN OVERSIGHT REQUIREMENTS

Article 21: Fundamental Principle: AI Recommends, Humans Decide

This Standard rests on the principle that:

- AI provides analysis and recommendations to accelerate review
- Humans make all substantive decisions affecting project outcomes
- No AI action is self-executing; all outcomes require human authorization
- Human decision-makers are accountable to the Board and applicants

This principle is non-negotiable and applies universally regardless of AI confidence level.

Article 22: AI Prohibitions (Hard Boundaries)

The AI Engine is strictly prohibited from:

- Approving, rejecting, or registering projects
- Issuing, cancelling, retiring, or transferring credits

- Triggering state transitions in the project workflow
- Modifying data in the registry or applicant records
- Sending communications to applicants, VVBs, or external parties
- Overriding or countermanding human decisions
- Accessing external systems without Board-approved integration

Violation of these prohibitions shall trigger immediate incident response and Board notification.

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.

Article 23: Escalation Rules by Confidence Tier

Human review protocols shall vary by confidence tier:

Tier
Confidence
Escalation Path
Timeline
Authority
HIGH
≥85
Single reviewer (PA)
5 days
PA (routine approval)
MEDIUM

60–84
Single reviewer + PA check
20 days
VVB or Technical Committee
LOW
40–59
YES (Technical Committee)
30 days
Technical Committee
VERY_LOW
<40
YES (Board)
45 days
Board decision

These timelines are target durations; delays shall be documented and escalated if exceeded.

Article 24: Human Override Protocol

When a human reviewer disagrees with AI recommendation:

- Reviewer views AI analysis, recommendation, and confidence score
- Reviewer makes decision (approve, reject, or escalate)
- If decision contradicts AI recommendation: Reviewer selects "Override"
- Reviewer provides written rationale (mandatory; minimum 50 words)
- System records: date, time, reviewer name/role, recommendation vs. decision, rationale
- Override logged in audit trail for quarterly and annual review

Article 25: Override Rationale Requirements

Overrides require mandatory written explanation meeting the following criteria:

- Minimum length: 50 words
- Required elements: Why is the AI recommendation incorrect or incomplete? What additional information or judgment supports the override?
- Acceptable reasons: Additional evidence provided post-AI analysis, special circumstances (SIDS, post-disaster, indigenous

territory), site visit findings, legal/policy interpretation

- Unacceptable reasons: "Expediting review", "avoiding procedure", "AI is wrong", or any rationale that contradicts methodology requirements

Inadequate or unacceptable rationales shall be flagged by the PA and escalated to the Board.

Article 26: Override Review Process

Overrides shall be subject to systematic review:

- Monthly: PA reviews all overrides from prior month; flags patterns or concerns
- Quarterly: Technical Committee reviews override patterns; identifies systematic bias or procedural drift
- Annually: Board reviews aggregate override metrics; assesses AI calibration; recommends adjustments

If override rate exceeds 25% in any quarter, the Board shall commission an audit of AI model performance.

Article 27: AI Visibility to Applicants

Applicants shall have visibility to:

- Completeness score + gap list (what is missing and how to fix it)
- Confidence tier (HIGH, MEDIUM, LOW, VERY_LOW) and interpretation
- Key findings and risk flags (summarized in human-friendly language)
- Suggested questions or clarifications (formatted as actionable requests)

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.

This approach balances transparency with operational independence.

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.

Article 28 as printed below was substituted by Amendment 1 (decision D-2026-08-10) and ENTERS INTO FORCE ON 17 SEPTEMBER 2026. Until that date the text reproduced in Appendix B.2 is the rule in force, under which a VVB has no visibility to AI analysis before conducting validation or verification. Engagements assigned before 17 September 2026 continue under the former Article 28 unless the Programme Administrator and the VVB jointly elect to migrate them. See Appendix B.3.

Article 28: Structured AI Visibility to VVBs

Accredited VVBs shall have structured visibility to AI Analysis Layer outputs through the M08 VVB Portal in accordance with this Article and Section 18 of MCC-400. AI outputs are programme-provided inputs to the VVB engagement; they are not substitutes for the VVB's independent professional judgement.

Article 28.1: Scope of Visibility

Upon engagement assignment under MCC-400 Section 9, the VVB shall be granted access to the AI Analysis Sidebar for the project. The Sidebar shall include: (a) per-section completeness scores; (b) confidence tiers under Article 15; (c) risk flags and their severity classification under Article 16; (d) cross-reference findings under Article 18; and (e) recommendations under Article 19.

Article 28.2: Bias Controls

The VVB shall apply the bias controls specified in MCC-400 Section 18.4. In particular, the VVB shall document each deviation from an AI finding and each concurrence with an AI finding, in the form prescribed by the Programme Administrator. Concurrence does not relieve the VVB of

independent sampling.

The references to MCC-400 Section 18 and Section 18.4 in this Article and in Article 28 read "Section 14" and "Section 14.4" as drafted. They were corrected on adoption of MCC-400 Addendum 1 (decision D-2026-08-11), which adopted that material as Section 18 because Section 14 of MCC-400 v1.0 is already occupied by VVB Assignment and Rotation. MCC-400 Addendum 1 enters into force on 17 September 2026, the same date as this Article, so Section 18.4 and this Article take effect together.

Article 28.3: Audit Trail

Every access to AI Analysis Layer output by a VVB user shall be logged in the immutable audit trail required by MCC-800 Article 65. Logs shall include user identity, timestamp, and the specific output viewed. Logs shall be available to the Programme Administrator and to peer reviewers under MCC-400 Section 9.6.

This Article read "MCC-400 Section 11" as drafted. Peer review is governed by MCC-400 Section 9.6; Section 11 governs performance monitoring and evaluation. The reference was corrected on adoption of MCC-400 Addendum 1 (decision D-2026-08-11).

Article 28.4: Statistical Detection of Rubber-Stamping

The Programme Administrator shall maintain a statistical model that flags VVB-engagement patterns indicative of insufficient independent verification, including but not limited to: (a) concurrence rates materially above the cohort median; (b) absence of contrary findings across consecutive engagements; and (c) report turnaround times materially below the cohort median. Flagged engagements shall be subject to enhanced peer review and may trigger witness audit under MCC-400 Article 26.

Article 28.5: Pre-Submission Visibility Excluded

AI Analysis Layer output shall not be disclosed to the VVB before the project developer’s final submission has been received and locked. AI outputs generated against draft submissions, applicant working documents, or pre-submission consultation are excluded from the VVB Sidebar.

Article 28.6: Project Developer Visibility

The project developer’s right to see AI Analysis Layer output is governed by Article 27. Nothing in this Article shall be construed to give the VVB

visibility to applicant-confidential output not also available to the project developer.

Article 29: AI Visibility to Public

The public shall have visibility to:

- Aggregate AI statistics (quarterly reports on analysis volume, confidence distribution, override rates)
- Disclosure on certificates and registry entries that AI-assisted review was performed
- Model version in deployment and recent changes (high-level summary)
- AI governance framework (this Standard and implementation guidance)

The public shall NOT have visibility to:

- Individual AI Review Records (internal documents)
- Model internals, training data, or algorithm details (proprietary)

This approach ensures transparency about AI involvement while protecting operational security.

Article 30: AI Transparency Reporting

MCC shall publish quarterly AI Transparency Reports including:

- Number of analyses performed (by trigger event type)
- Distribution of confidence tiers (% HIGH, MEDIUM, LOW, VERY_LOW)
- Override rate (% of recommendations overridden; by tier)
- Average processing time by analysis type
- Model version in use; any model changes in period
- Geographic and sectoral analysis (any systematic patterns?)
- Availability metrics (% uptime, incident summary)
- (h) analyses by Trigger Source
- (i) number and proportion of analyses produced in

Degraded Mode, with cause and duration

- (j) number of Prioritised Sequencing Flags raised and the proportion subsequently registered, refused or returned for revisions

Reports shall be published on the MCC website and presented to the Board.

PART 5: AI PERFORMANCE MONITORING

Article 31: Performance Metrics

The following metrics shall be tracked continuously:

- Accuracy: AI completeness scores vs. independent human re-scoring (target ≥90% agreement within ±5 points)
- Override rate: % of AI recommendations overridden by humans (tracked monthly; trend analysis)
- Processing time: vs. SLA targets (Article 8)
- Availability: % uptime during operating hours (target 99.9%)
- Error rate: Incorrect analyses (systematic quality issues)
- Bias metrics: Geographic, sectoral, linguistic analysis (systematic disparities)

Article 32: Calibration Review (Quarterly)

The Technical Committee shall conduct quarterly calibration reviews:

- Analyze monthly accuracy reports
- Identify systematic patterns (Is AI consistently over-scoring certain sectors? Under-scoring certain regions?)
- Compare AI recommendations to final human decisions
- Assess for bias or drift
- Recommend adjustments if thresholds need recalibration

Calibration reports shall be published alongside quarterly transparency reports.

Article 33: Model Recalibration Triggers

The AI model shall be recalibrated if any of the following occur:

- Override rate >25% in two consecutive quarters
- Systematic bias detected (geographic, sectoral, or linguistic)
- New methodology adopted (F01–F05 revision)
- Material methodology revision requiring AI re-training
- External audit identifies accuracy <85%

Recalibration shall follow the governance process in Article 10.

Article 34: Model Testing Requirements

New AI models must pass rigorous testing before deployment:

- Accuracy: Re-score 50 reference PDDs; ≥90% concordance with human scoring required
- Bias: Test across geographies, sectors, languages; flag systematic disparities
- Regression: Ensure no degradation vs. reference test cases
- Recommendation: Generate recommendations on test set; compare to human decisions

Testing report shall be reviewed by Technical Committee before Board approval.

Article 35: Incident Management

If AI produces incorrect analysis or experiences malfunction:

- Incident detected and documented (time, nature, severity)
- AI Engine isolated; live processing halted if necessary
- Root cause analysis conducted
- Affected analyses reviewed and corrected if necessary
- Report submitted to PA within 24 hours
- Board notified if incident affects live projects

Incidents shall be logged in audit trail for annual review.

Article 36: Annual AI Governance Report

The Programme Administrator shall submit an annual AI Governance Report to the Board covering:

- Summary of AI performance metrics (accuracy, availability, processing time)
- Override analysis (patterns, trends, systemic issues)
- Model changes approved and deployed during period
- Incidents and resolutions
- Recommendations for policy or technical improvements

Report shall be presented to the Board and published on the MCC website (anonymized).

Article 37: External Audit

AI governance shall be subject to annual external audit as part of programme-wide compliance review.

External auditor shall assess:

- Compliance with this Standard (all articles)
- AI accuracy against independent reference dataset
- Presence and adequacy of audit trail
- Model governance procedures
- Data protection and security controls

Audit findings shall be reported to the Board and disclosed to applicants and stakeholders.

Article 38: Continuous Improvement

MCC shall maintain a continuous improvement feedback loop:

- Human decisions feed back into model training (to improve future recommendations)
- Override patterns analyzed for root causes (why do certain recommendations get overridden?)
- Applicant feedback solicited on AI analysis clarity and usefulness
- VVB feedback incorporated on AI-VVB concordance
- Annual review of AI governance framework for gaps or improvements

This iterative approach ensures AI performance improves over the life of the programme.

PART 6: DATA PROTECTION AND SECURITY

Article 39: Data Handling

All data processed by the AI Engine is MCC-internal and subject to strict handling requirements:

- No third-party data sharing without explicit Board approval
- External integrations (Verra API, climate databases) require Board authorization per integration
- Applicant data used only for MCC review purposes
- Anonymized data may be used for AI model improvement (with applicant consent or Board discretion)

Article 40: Data Residency

AI processing must comply with host country data requirements:

- Applicable laws identified for each operating jurisdiction
- Data residency policy established by PA and approved by Board
- MCC data centers configured to meet residency requirements

Data residency policy shall be reviewed annually and updated as laws evolve.

Article 41: Encryption

Data protection standards for AI systems:

- In transit: TLS 1.3 (minimum) for all API and network communications
- At rest: AES-256 encryption for all stored data
- AI Review Records: Encrypted at rest; decryption keys managed by MCC security team
- Key management: HSM (Hardware Security Module) or equivalent for encryption key protection

Article 42: Access Control

The AI Service Account shall operate with minimal required

permissions:

- Read-only access to MCC database tables required for analysis
- Write access only to AI Review Record tables
- No direct database access by humans except through monitored interfaces
- Quarterly audit of access logs for unauthorized access attempts
- Incident response if unauthorized access detected (within 4 hours notification to Board)

PART 7: TRANSITIONAL AND FINAL PROVISIONS

Article 43: Effective Date

This Standard becomes effective upon approval by the MCC Board of Directors.

The Board shall designate an implementation commencement date, allowing reasonable time for system configuration and team preparation.

Article 44: Transitional Period

A six-month transitional period shall apply from the effective date:

- During this period, the AI Engine shall be configured, tested, and calibrated to meet all requirements
- Live project submissions shall not be processed until full compliance is demonstrated
- Compliance certification required before AI Engine processes first live submission
- Board sign-off required before transitioning to operational mode

Article 45: Review Schedule

This Standard shall be subject to scheduled review:

- Annual review: by Technical Committee and AI Governance Committee
- Next formal review: March 2027 (one year from effective date)

- Amendments proposed between reviews shall follow the process in Article 5
- Major revisions (v1.2 or higher) require Board approval and stakeholder consultation

APPENDIX: CONFIDENCE TIER DECISION MATRIX

APPENDIX — CONFIDENCE TIER DECISION MATRIX

This matrix guides the interpretation of confidence tiers in practical decision-making contexts:

TIER 1: HIGH CONFIDENCE (≥85)

Characteristics:
• Completeness score ≥90%
• Data internally consistent
• No conflicting evidence
• All major methodology elements addressed
• Cross-references align

Recommended Action:
• MCC PA can approve recommendation (routine pathway)
• Human may override with brief documentation if needed
• No escalation required unless override occurs

Example: "PDD is 95% complete; baseline data consistent; quantification arithmetic verified; no risk flags. AI recommends APPROVE. Confidence: HIGH (94)."

TIER 2: MEDIUM CONFIDENCE (60–84)

Characteristics:
• Completeness score 80–90%
• Minor gaps or inconsistencies
• Some data incomplete but addressable
• Most methodology elements present

Recommended Action:
• Forward to VVB with specific clarification requests
• VVB seeks applicant response
• Applicant submits additional information
• VVB completes validation with clarifications

Example: "PDD is 82% complete; baseline period limited to 3 years (recommend 5). Cross-references show project location consistent. AI recommends REQUEST_INFO. Confidence: MEDIUM (71)."

TIER 3: LOW CONFIDENCE (40–59)

Characteristics:
• Completeness score <80%
• Multiple inconsistencies or gaps
• Potentially material issues requiring expert judgment
• May require project revision or site investigation

Recommended Action:
• Automatic escalation to Technical Committee
• Committee expert reviews AI analysis + underlying documents
• May request applicant revision or authorize deeper investigation
• May approve conditional on specific requirements

Example: "PDD is 76% complete; baseline methodology unclear (three approaches mentioned). Cross-reference shows potential conflict with adjacent Verra project. AI recommends ESCALATE. Confidence: LOW (54)."

TIER 4: VERY LOW CONFIDENCE (<40)

Characteristics:
• Highly incomplete (completeness <60%)
• Major inconsistencies or contradictions
• Potential fraud, material error, or regulatory breach
• Requires senior governance review

Recommended Action:
• Automatic escalation to Board
• Board reviews AI analysis + full evidence
• May reject, require major revision, or approve with conditions
• Applicant has appeal rights (to Appeals Panel)

Example: "PDD baseline contradicts form data; deduplication check finds same project in Verra (issued 2 months ago). Safeguards missing entirely. AI recommends ESCALATE TO BOARD. Confidence: VERY_LOW (32)."

CONFIDENCE SCORE DERIVATION

Confidence is calculated from multiple factors, weighted as follows:

1. Completeness Score (40% weight)
   - ≥95%: +38
   - 90–94%: +34
   - 80–89%: +26
   - 70–79%: +14
   - <70%: +0

2. Data Consistency (30% weight)
   - All cross-references aligned: +30
   - Minor discrepancies (<5%): +21
   - Moderate discrepancies (5–15%): +12

```
   - Major contradictions (>15%): +0

3. Risk Clarity (20% weight)
   - No HIGH-risk flags: +20
   - Only LOW-risk flags: +14
   - MEDIUM-risk flags present: +7
   - Multiple HIGH-risk flags: +0

4. Methodology Compliance (10% weight)
   - All sections addressed: +10
   - One section weak: +6
   - Multiple sections weak: +2
   - Major section missing: +0

Example calculation:
Completeness 92% (→ +34) + Consistency aligned (→ +30) + LOW-risk only
(→ +14) + All sections (→ +10) = 88 → HIGH TIER
```

APPENDIX B — VERSION CONTROL AND SUPERSEDED TEXT

This Appendix records the amendment history of MCC-850 and reproduces superseded provisions in full, so that the text in force at any past date can be reconstructed from this document alone.

B.1 Amendment history

Version
Date
Instrument
Status
Effect
1.0
23 March 2026
MCC-850 AI Governance Standard v1.0
In force
Base text, Articles 1–45 and confidence tier decision matrix.
1.1
17 August 2026
Amendment 1 — Repeal and Replacement of Article 28 (AI Visibility to VVBs), v1.0 (drafted 29 April 2026)
Adopted 17 August 2026, D-2026-08-10 (MCC-BDR-2026-002), Article 5 procedure.
In force 17 September 2026
Repealed Article 28 in its entirety and substituted a new Article 28 with sub-articles 28.1–28.6; added a sentence to Article 3 on the relationship with MCC-400 Section 14. See B.3.
1.1
17 August 2026
Amendment 2 — Automated Completeness Report, Applicant Disclosure and Degraded Mode, v1.0

Adopted, D-2026-08-02 (MCC-BDR-2026-002), Article 5 procedure; in force 17 August 2026
Inserted four definitions in Article 2; inserted Articles 7.1, 12.1, 13.1, 19.1 and 22.1; added a clause to Article 8; replaced the withheld-items list in Article 27 and added a release-review clause; amended Article 3, Article 30 and the Article 13 schema.

B.2 Superseded text

Article 27, list of items withheld from applicants, as it stood in MCC-850 v1.0 (23 March 2026) until 17 August 2026:

> Applicants shall NOT have visibility to:
>
> • Full AI Review Record (internal document; raw analysis withheld)
>
> • AI model internals or algorithm details
>
> • AI recommendations (only human decisions are communicated)

The remainder of Article 27, including the list of items to which applicants *shall* have visibility and the closing sentence, is unchanged by Amendment 2 and appears in the body of this text.

Article 28, as it stood in MCC-850 v1.0 (23 March 2026) and as it remains in force until 16 September 2026:

> **Article 28: AI Visibility to VVBs**
>
> VVBs shall have NO visibility to AI analysis before conducting validation or verification.
>
> Rationale: VVB independence requires unbiased assessment; prior knowledge of AI findings could unconsciously bias the VVB.
>
> Process:
>
> • Applicant submits PDD; AI generates analysis
>
> • MCC assigns VVB (without disclosing AI findings)
>
> • VVB conducts independent validation/verification
>
> • VVB submits report
>
> • AI analysis compared to VVB findings (calibration check; documented)
>
> Post-validation, VVBs may request AI analysis if they wish to understand MCC perspective (for learning purposes).

B.3 Article 28 — two texts and the date that separates them

This is a live compliance point and is recorded here so that it cannot be missed on due diligence.

Amendment 1 was adopted on 17 August 2026 by decision D-2026-08-10, but does not enter into force on adoption. Its own terms require the Programme Administrator to specify the entry-into-force date and to publish it through the Transparency Layer with at least thirty days notice. **The Programme Administrator has specified 17 September 2026.** Two texts of Article 28 therefore apply in sequence:

- **Until 16 September 2026 inclusive** — the Article 28 reproduced in B.2 above governs. A VVB has no visibility to AI analysis before conducting validation or verification. No AI Analysis Sidebar, and no equivalent disclosure of AI Analysis Layer output to a VVB user, may be enabled before that VVB has submitted its validation or verification report.
- **From 17 September 2026** — the Article 28 printed in the body of this text governs, and structured visibility through the M08 VVB Portal becomes available.
- **Engagements already assigned** on 16 September 2026 continue under the former Article 28 unless the Programme Administrator and the VVB jointly elect to migrate them. Engagements assigned on or after 17 September 2026 operate under the new Article 28.

The MCC-400 dependency, and how it was resolved. New Articles 28.1 and 28.2 cross-refer to MCC-400 provisions that exist only by virtue of MCC-400 Addendum 1 (AI-Augmented Validation and Verification). Article 28.2 — the bias-control obligation that is the whole basis on which visibility is safe to grant — would be inoperable without it. That Addendum was adopted as amended at the same meeting by decision D-2026-08-11, with the same entry-into-force date of 17 September 2026, so the two instruments take effect together.

The Board adopted the Addendum as **Section 18**, not Section 14 as drafted, because Section 14 of MCC-400 v1.0 is already occupied by *VVB Assignment and Rotation* and the draft's article numbering would have duplicated eighteen articles already in force. The cross-references in Articles 3, 28 and 28.2 of this Standard were corrected accordingly on adoption, as was the reference to peer review in Article 28.3 (Section 9.6, not Section 11).

What the Addendum requires before the Sidebar opens. MCC-400 Section 18.4 (Articles 104 to 106) requires the VVB to document every deviation from and every concurrence with an AI finding, and governs what may be published. Article 106 as adopted carries an express carve-out so that publication does not disclose the recommendation field, the numeric confidence score, or any individual AI Review Record, and so does not defeat Articles 27 and 29 of this Standard. MCC-400 Article 103.1 as adopted provides that no sampling reduction is available in respect of a record produced in Degraded Mode under Article 12.1 of this Standard.

Consequences for the platform. The M08 VVB Portal AI Analysis Sidebar must remain closed to VVB users until the entry-into-force date, and must not be opened at all if that date is postponed. The concurrence and deviation capture required by Article 28.2, the access logging required by Article 28.3, the pre-submission exclusion required by Article 28.5 and the rubber-stamping statistics required by Article 28.4 must be in place before the Sidebar is opened, not after.

B.4 Related instruments

Article 19.1 is to be read with MCC-100 clauses 0061.1 to 0061.3, inserted by MCC-100 Amendment 1 (decision D-2026-08-01) and consolidated in MCC-100 v2.1, which establish that prioritised sequencing confers queue priority only.

Article 28 is to be read with MCC-400 Section 18, added by MCC-400 Addendum 1 (decision D-2026-08-11) and consolidated in MCC-400 v1.2, which governs VVB use of AI outputs, minimum sampling by confidence tier, and the documentation of deviations and concurrences on which Article 28.2 depends. Both enter into force on 17 September 2026.

Two provisions of MCC-400 Section 18 as adopted bear directly on this Standard and are recorded here. First, MCC-400 Article 95.1 requires the Programme Administrator to publish the operational status of each AI Analysis Layer component and provides that no requirement, timeline, sampling reduction or fee expectation may rest on a component recorded as not operational; the Continuous Monitoring Engine is recorded as not operational, its most recent output being dated 1 May 2026. Second, MCC-400 Article 99.1 sets the minimum sampling threshold for HIGH-

confidence findings at thirty per cent, not ten per cent, until the first calibration review under Article 32 of this Standard has been completed and its evidence published.

B.5 Note on the basis for this Appendix

The editorial note to Amendment 2 directs that superseded text be preserved in the version-control appendix of MCC-850 v1.2 in accordance with Article 11 and the standards-history requirements of MCC-100. Article 11 concerns AI model version tracking rather than document version control, and MCC-100 contains no express document-history clause. This Appendix is therefore maintained on the basis of MCC-100 clause 0040, MCC-100 clause 0256 and the express direction in the editorial notes to Amendments 1 and 2. It is recommended that a future amendment insert an express document-control and version-history Article into Part 7, and correct the cross-references in the editorial notes.

DOCUMENT METADATA

Document ID
MCC-850
Title
AI Governance Standard
Version
1.1 (consolidated)
Date
17 August 2026
Classification
Programme Standard – Normative (Tier 3)
Status
In force
Next review
March 2027
Articles
45, together with Articles 7.1, 12.1, 13.1, 19.1 and 22.1 inserted by Amendment 2, and Articles 28.1 to 28.6 inserted by Amendment 1
Document owner
MCC Technical Committee
Implementation owner
MCC Programme Administrator

This document is the formal codification of AI governance requirements within the Marine Conservation Credits programme. All MCC staff, VVBs, and programme participants are bound by the requirements herein.

For questions or amendments, contact: governance@mcc-credits.org