TERRA VITA HUB

INSTITUTIONAL REVIEW INDEX (IRI) WHITEPAPER

The calibration, oversight, and reviewer-assurance layer of Terra Vita Hub

Whitepaper thesis

Institutions do not only need governance structure and evidence proof. They also need transparent oversight of reviewer behaviour, reviewer consistency, reviewer bias, programme calibration, and cross-programme comparability.

Document class	Institutional governance and assurance whitepaper
Prepared for	Ministries, DFIs, climate funds, sovereign programme owners, auditors, fiduciary reviewers, and implementation partners
Issuer	Terra Vita Global B.V. / Terra Vita Hub
Status	Institutional-grade draft for review and diligence conversations
Date	May 2026

Core proposition

The IRI allows an institution to ask not only whether evidence exists, but whether the people reviewing that evidence are behaving consistently, transparently, and within accountable governance boundaries.

Document Control

This whitepaper defines the Institutional Review Index (IRI) as Terra Vita Hub's reviewer-assurance and meta-governance layer. It is written as a deployment-agnostic institutional design paper and can be adapted for sovereign, DFI, climate-fund, donor, programme, and portfolio contexts.

Control element	Position
Purpose	Define the IRI as an auditable reviewer oversight, calibration, and assurance mechanism.
Scope	Reviewer behaviour, consistency, bias and divergence detection, programme calibration, dashboard outputs, escalation posture, and committee-ready assurance signals.
Authority boundary	The IRI supports oversight and calibration. It does not replace human reviewers, statutory decision-makers, procurement bodies, auditors, approved MRV methodologies, or financial authorities.
Design posture	Transparent, evidence-linked, reconstructable, sovereign-safe, and institution-first.
Primary users	Programme owners, ministries, DFIs, climate funds, auditors, portfolio governance teams, and implementation-management offices.

Contents

Section	Topic
1	Executive Summary
2	Purpose of the IRI
3	Position in the Governance Stack
4	IRI Architecture and Core Components
5	Methodology and Scoring Logic
6	Institutional Use Cases
7	Assurance Outputs and Dashboard Model
8	Governance, Accountability, and Safeguards
9	Implementation Pathway
10	Conclusion
Annex A	Output Register
Annex B	Minimum Data Schema
Annex C	Committee Review Questions

1. Executive Summary

Institutions do not only need governance structure and proof. They need confidence that the reviewers who interpret evidence, apply criteria, and make recommendations are themselves operating within visible, accountable, and calibrated boundaries.

Terra Vita Hub already frames governance as a controlled decision environment through the Governance Spine and evidence integrity through the Institutional Assurance Layer (IAL). The Institutional Review Index (IRI) adds a third and necessary layer: oversight of the review process itself.

Definition

The Institutional Review Index (IRI) is Terra Vita Hub's meta-governance layer: a structured, auditable, non-black-box evaluation system that measures the integrity, behaviour, calibration, and reliability of reviewers and review processes.

The IRI allows ministries, DFIs, climate funds, sovereign programmes, and auditors to evaluate whether reviewer behaviour is consistent, evidence-aligned, bias-aware, timely, and comparable across programmes. It is designed to answer an institutional question that conventional dashboards cannot answer:

Can we trust not only the records, but also the reviewer process that produced the decisions?

The IRI is not an automated decision-maker and does not issue statutory approvals or financial decisions. It creates a transparent assurance environment around reviewers by linking reviewer actions to evidence, rationale, governance controls, thresholds, audit trails, and calibration standards.

Institutional need	IRI response
Reviewer performance visibility	Captures timeliness, completeness, evidence alignment, control adherence, and decision quality signals.
Reviewer consistency	Monitors whether criteria are applied predictably and whether reviewer behaviour drifts over time.
Bias and divergence detection	Identifies unexplained outliers, geographic or thematic skew, and peer divergence for escalation and human review.
Programme-level assurance	Provides calibration scores and assurance signals at programme, portfolio, and institution level.
Cross-programme comparability	Creates a shared reviewer-assurance baseline so programmes can be compared without collapsing local sovereignty or methodology differences.

Executive assurance statement

With the Governance Spine, the IAL, and the IRI operating together, Terra Vita Hub can support institution-grade assurance across three dimensions: structure, proof, and oversight. The result is a governance environment in which decisions can be traced, reviewers can be calibrated, divergence can be explained, and institutional committees can reconstruct the logic behind programme review outcomes.

2. Purpose of the IRI

The IRI exists to close four gaps that regularly weaken public-sector, DFI, climate-fund, donor, and portfolio governance environments. These gaps are especially acute in multi-stakeholder programmes where ministries, technical reviewers, funders, implementing partners, and auditors need confidence that review behaviour remains consistent over time.

Institutional gap	Risk created	IRI response
Reviewer opacity	Programme owners cannot see how reviewers behave, evolve, diverge, or apply judgment across cases.	Reviewer-level performance, consistency, and rationale signals make reviewer behaviour visible without exposing unnecessary personal or sovereign-sensitive information.
Inconsistent decision quality	Decisions drift, fragment, contradict, or become dependent on individual reviewer habits rather than agreed standards.	Calibration scores and drift alerts detect deviation from programme norms and institutional criteria.
Lack of cross-programme comparability	Programmes cannot be compared because each review environment uses different evidence practices, reviewer thresholds, or escalation rules.	Programme calibration layers create a shared assurance baseline while preserving local programme context.
Sovereign accountability requirements	Governments, DFIs, and auditors lack demonstrable oversight of the people and processes influencing decisions.	Audit-ready reviewer oversight outputs provide reconstructable accountability for committees, procurement, fiduciary review, and sovereign assurance.

2.1 Institutional problem statement

Evidence integrity alone is not enough. A programme can have documents, audit trails, MRV signals, and formal approvals, yet still carry unobserved risk if reviewer interpretation is inconsistent, biased, delayed, incomplete, or poorly calibrated. The IRI makes reviewer performance and reviewer-process integrity part of the institution's governed operating model.

This is particularly important in settings where programmes must demonstrate fiduciary responsibility, preserve sovereign accountability, satisfy procurement scrutiny, or provide repeatable assurance across multiple countries, regions, sectors, or funding instruments.

2.2 What the IRI deliberately avoids

- It does not rank reviewers through an opaque score without explanation.
- It does not replace reviewer judgment, statutory authority, approved MRV methodology, or committee accountability.
- It does not convert bias detection into automatic disciplinary action.
- It does not collapse local programme context into a one-size-fits-all global benchmark.
- It does not expose sensitive reviewer or sovereign data beyond configured institutional permissions.

Design principle

The IRI is a calibration and assurance instrument, not a punitive surveillance instrument. Its purpose is to make reviewer-process risk visible, explainable, and governable.

3. Position in the Governance Stack

The IRI is the third layer in Terra Vita Hub's institutional governance architecture. Each layer answers a different assurance question.

Layer	Primary question answered	Institutional function
Governance Spine	Is the operating structure controlled?	Defines evidence intake, routing, reviewer action, workflow controls, approval posture, audit lineage, and export readiness.
Institutional Assurance Layer (IAL)	Can the proof be reconstructed?	Links evidence, reviewer attribution, audit trails, controls, and assurance artefacts into a verifiable proof environment.
Institutional Review Index (IRI)	Can the reviewer process itself be trusted?	Evaluates reviewer behaviour, consistency, calibration, bias, divergence, and cross-programme comparability.

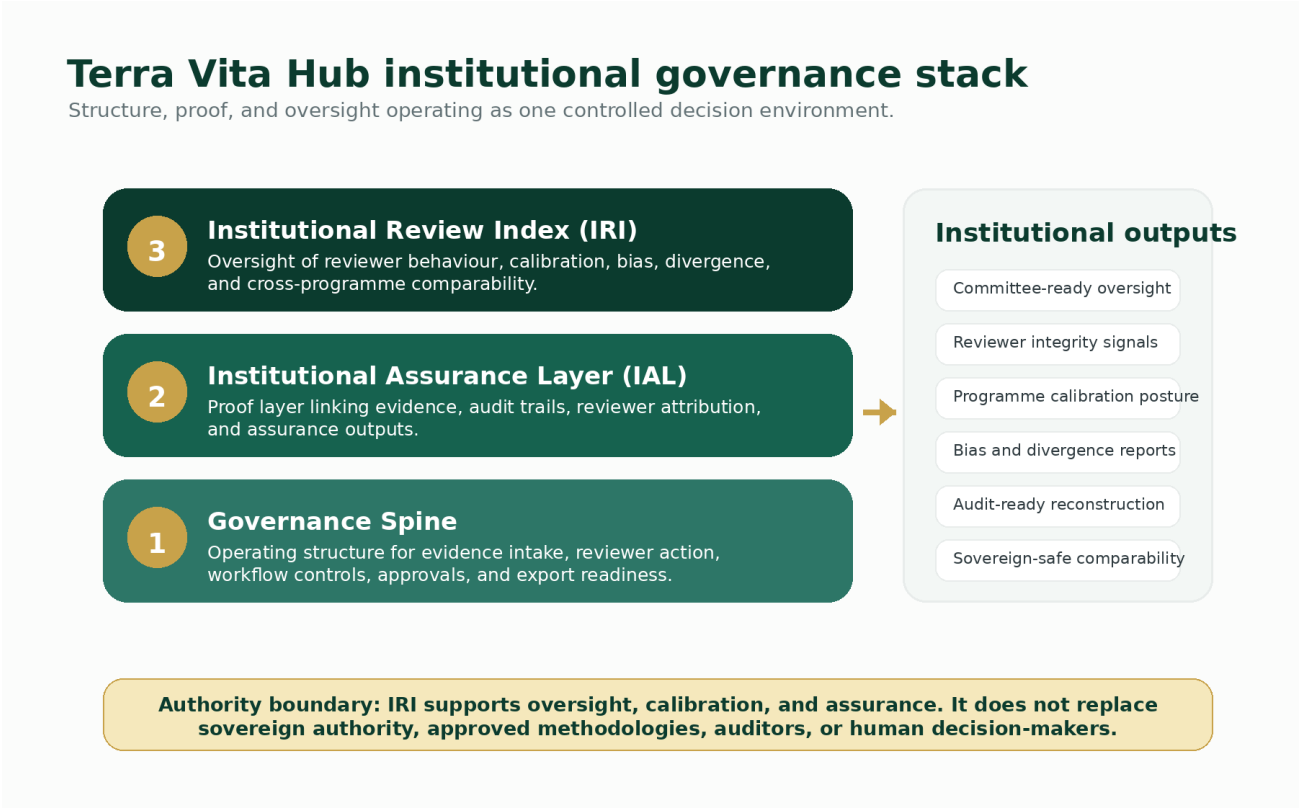

Figure 1. Terra Vita Hub governance stack: structure, proof, and oversight.

3.1 Why the third layer matters

Many systems stop at workflow visibility or evidence proof. That leaves a critical institutional exposure: the reviewers themselves may drift, apply criteria inconsistently, over-rely on incomplete evidence, produce unexplained outliers, or diverge across regions and programmes. The IRI brings this exposure into a governed assurance model.

In practice, this means that a ministry, DFI, climate fund, or auditor can interrogate the review process at a higher level: not only which decision was made, but whether the decision process remained aligned with agreed criteria, peer behaviour, programme controls, and institutional risk posture.

Figure 4. Institutional Review Index architecture: component oversight, calibration, and assurance outputs.

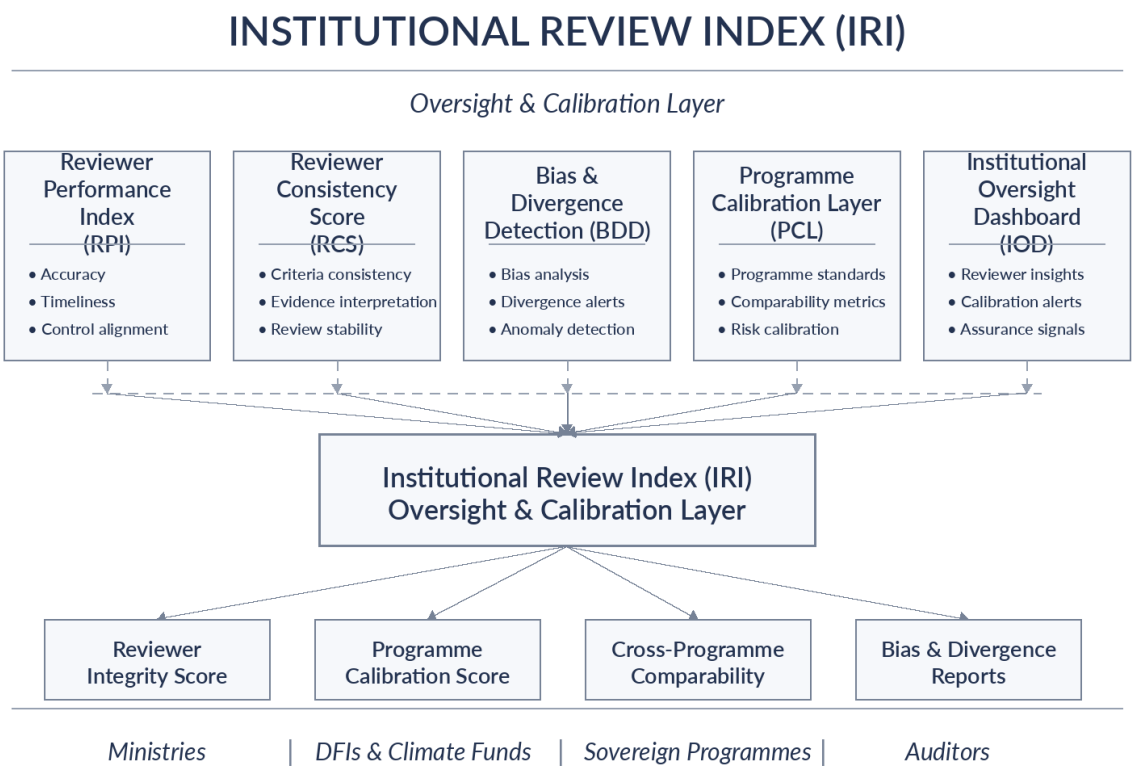

Visual model for institutional review, calibration, and reviewer-assurance workflows.

4. IRI Architecture and Core Components

The IRI is built from five institutional primitives. Each primitive captures a different aspect of reviewer assurance and contributes to institutional oversight, calibration, and comparability.

Figure 5. Institutional Review Index composite framework: reviewer metrics, calibration, assurance processing, and institutional outputs.

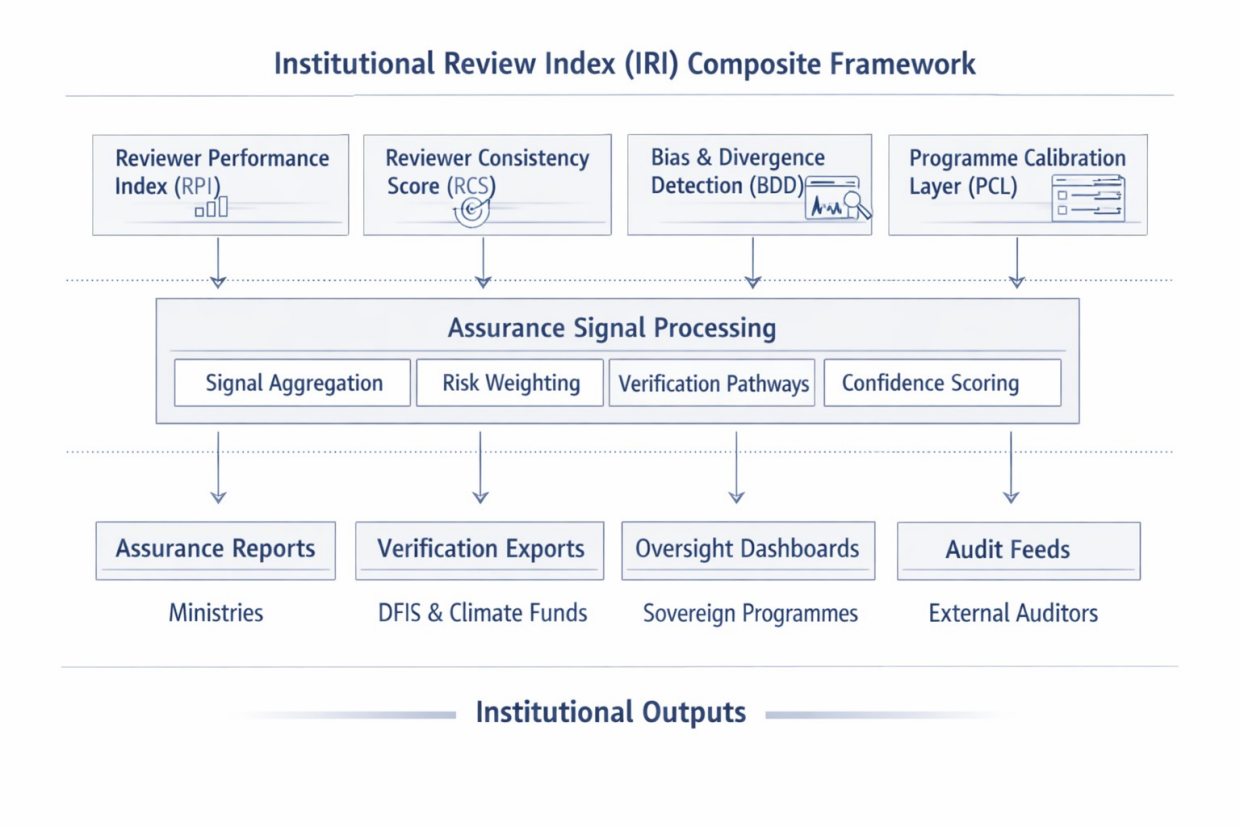

Visual model for institutional review, calibration, and reviewer-assurance workflows.

Component	Measures / monitors	Institutional value
Reviewer Performance Index (RPI)	Decision accuracy signals, evidence alignment, deviation from norms, timeliness, completeness, and control adherence.	Shows whether reviewer behaviour is effective, timely, complete, and aligned with governance requirements.
Reviewer Consistency Score (RCS)	Stable criteria application, predictable evidence interpretation, drift over time, and alignment with calibration standards.	Supports institutional trust by showing whether reviewer behaviour remains consistent.
Bias & Divergence Detection (BDD)	Systematic skew, unexplained divergence from peer reviewers, thematic or geographic imbalance, and anomalous decision patterns.	Surfaces potential reviewer-process risk for human review and sovereign-safe escalation.
Programme Calibration Layer (PCL)	Cross-programme reviewer baselines, comparable assurance posture, stable institutional risk calibration, and portfolio-level comparability.	Allows programmes to be compared without erasing local programme rules or sovereign context.
Institutional Oversight Dashboard (IOD)	Reviewer-level insights, programme assurance signals, drift alerts, performance trends, and comparability metrics.	Provides the operational interface for institutions, committees, auditors, and programme leads.

4.1 Reviewer Performance Index (RPI)

The Reviewer Performance Index measures reviewer behaviour across the core dimensions that institutions need for accountable oversight. It is not designed as a hidden productivity score. It is a reconstructable assurance signal anchored to records, criteria, evidence, and reviewer actions.

- Decision quality and evidence alignment: whether decisions are supported by the evidence that was available at the time.
- Timeliness: whether reviewer actions occur within expected governance windows and escalation thresholds.
- Completeness: whether required fields, rationale, checks, and evidence references are present.
- Control adherence: whether approval, override, escalation, and exception rules are followed.
- Deviation from norms: whether reviewer behaviour departs materially from calibrated peers or programme standards.

Figure 6. Reviewer Performance Index (RPI): evidence-linked reviewer performance and governance-control alignment.

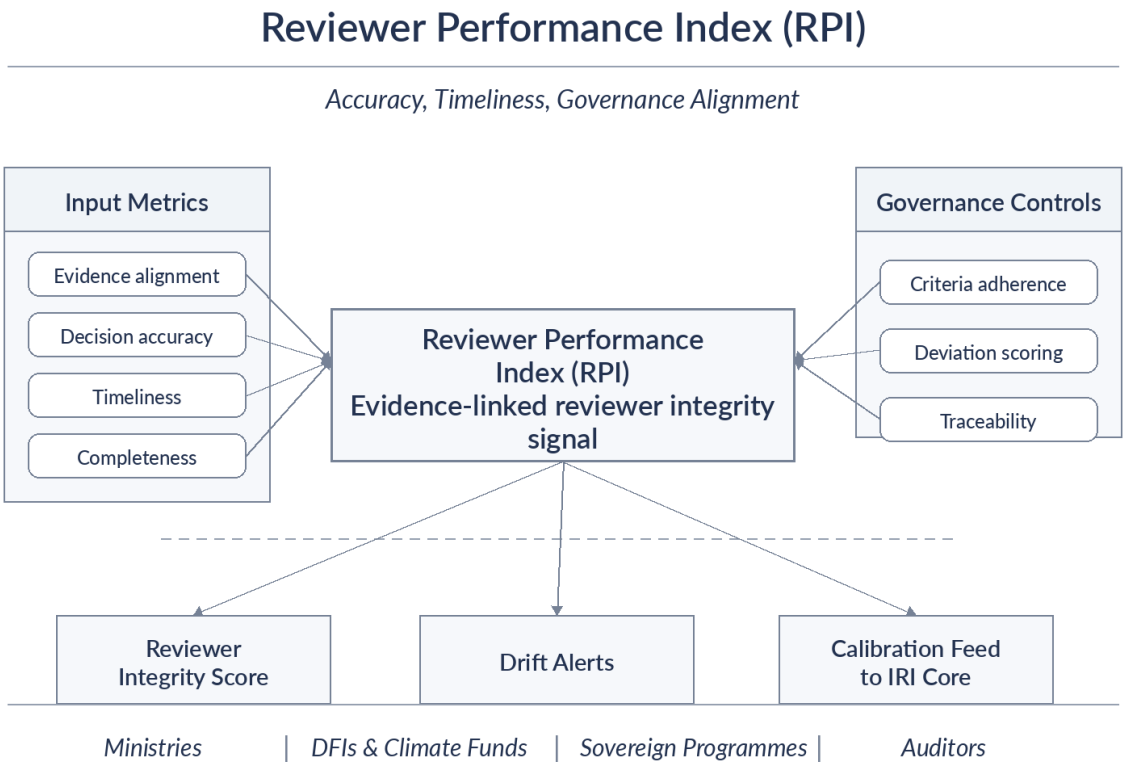

Visual model for institutional review, calibration, and reviewer-assurance workflows.

4.2 Reviewer Consistency Score (RCS)

The Reviewer Consistency Score evaluates whether a reviewer applies criteria predictably across similar records, evidence types, programme contexts, and review cycles. Consistency does not mean identical outcomes. It means comparable reasoning, documented rationale, and stable application of agreed standards.

- Consistent interpretation of review criteria across similar cases.
- Predictable use of evidence thresholds and escalation rules.
- Stable reasoning patterns across time, geography, theme, or programme tranche.
- Transparent explanation where context requires different treatment.

Figure 7. Reviewer Consistency Score (RCS): criteria consistency, drift detection, and cross-reviewer alignment.

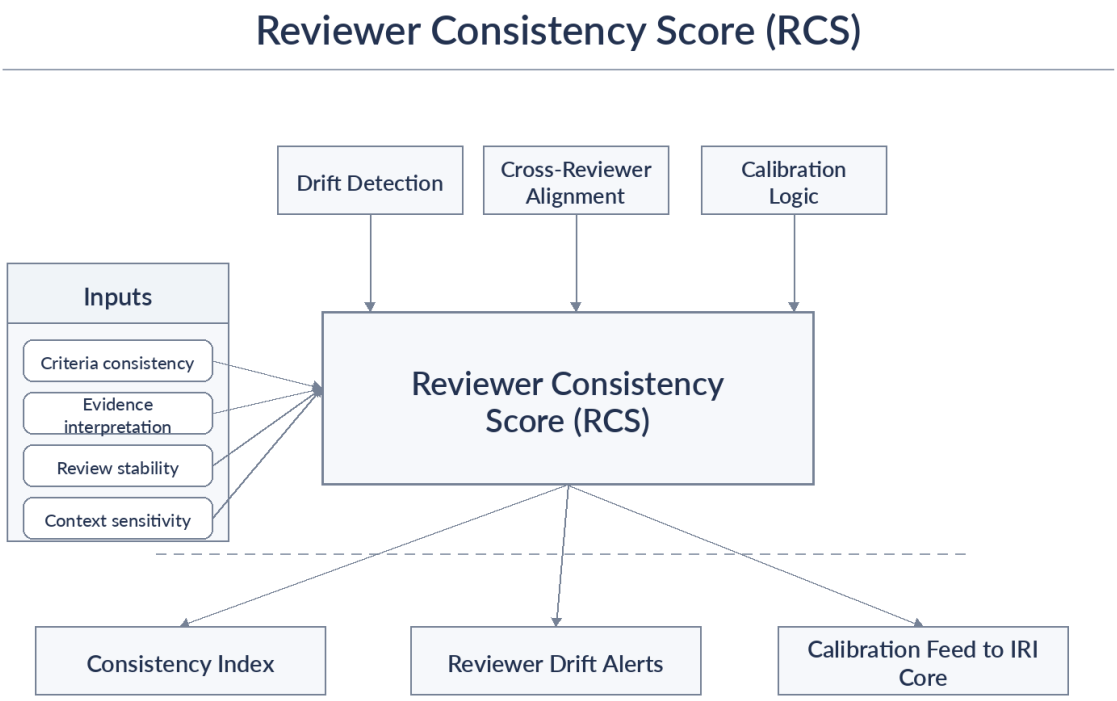

Visual model for institutional review, calibration, and reviewer-assurance workflows.

4.3 Bias & Divergence Detection (BDD)

Bias & Divergence Detection identifies patterns that require oversight. It does not assume misconduct. It flags review behaviour that needs explanation, peer review, recalibration, or escalation.

- Systematic reviewer skew against or toward particular categories of records, regions, themes, operators, or cohorts.
- Unexplained divergence from peer reviewers working under the same criteria.
- Anomalous decision patterns or unusually high exception, override, rejection, or approval rates.
- Pattern changes after programme, team, funding, evidence, or policy conditions shift.

Figure 8. Bias & Divergence Detection (BDD): bias analysis, divergence thresholds, anomaly detection, and reviewer calibration feed.

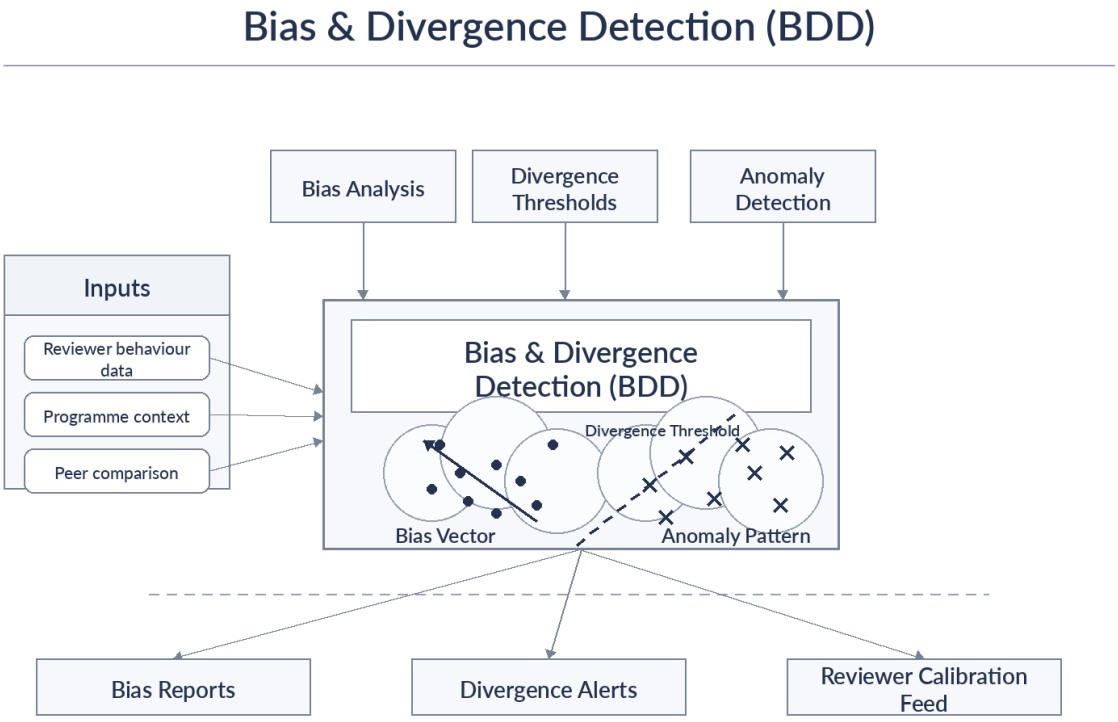

Visual model for institutional review, calibration, and reviewer-assurance workflows.

4.4 Programme Calibration Layer (PCL)

The Programme Calibration Layer ensures reviewers across programmes operate on a comparable assurance baseline. This is essential where Terra Vita Hub is deployed across multiple programme instances, sectors, sovereign contexts, or funding facilities.

- Maintains shared assurance baselines across programme environments.
- Separates local programme configuration from global assurance comparability.
- Allows ministries and DFIs to compare programme posture without forcing identical programme rules.
- Supports portfolio, country, and multi-country reporting.

Figure 9. Programme Calibration Layer (PCL): cross-programme standards and comparability framework.

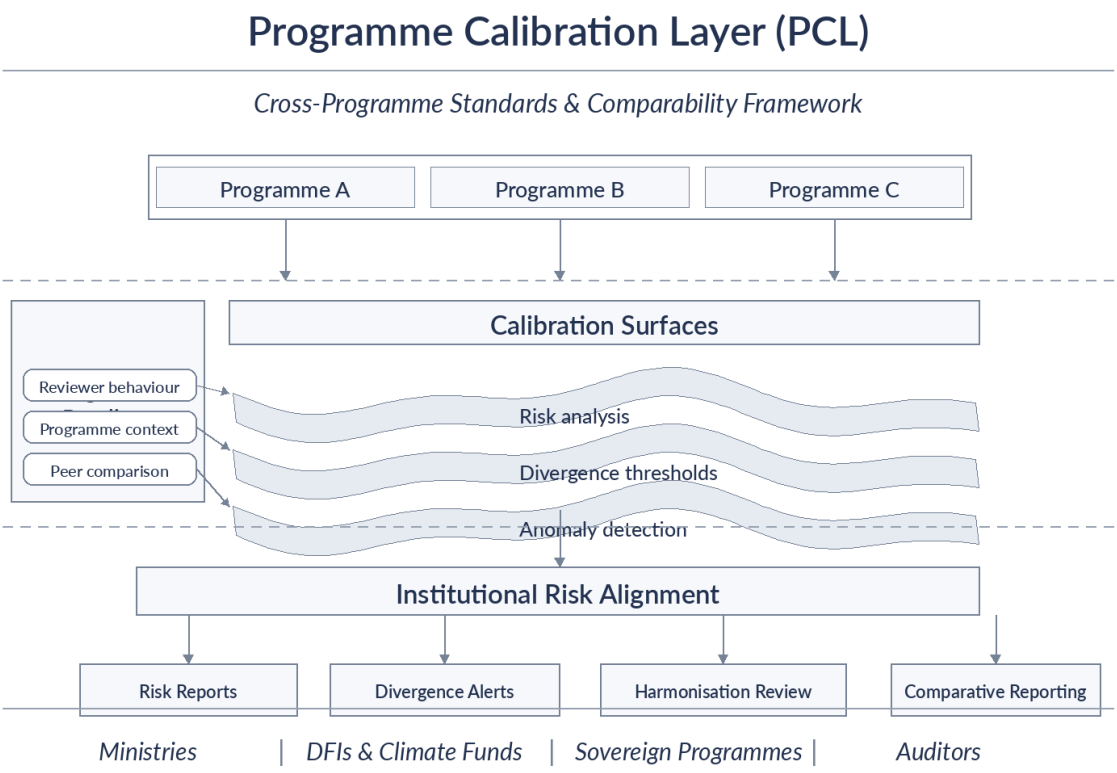

Visual model for institutional review, calibration, and reviewer-assurance workflows.

4.5 Institutional Oversight Dashboard (IOD)

The Institutional Oversight Dashboard is the operational interface through which authorised users observe reviewer-process health, programme calibration, and assurance posture. It should be role-governed, export-aware, and linked to the underlying audit records.

- Reviewer-level performance and consistency signals.
- Programme-level calibration posture and drift alerts.
- Bias and divergence exceptions requiring human review.
- Cross-programme comparability metrics for committee and portfolio governance.
- Exportable oversight summaries for auditors, committees, funders, and sovereign bodies.

Figure 10. Institutional Oversight Dashboard (IOD): reviewer insights, calibration alerts, and assurance signals.

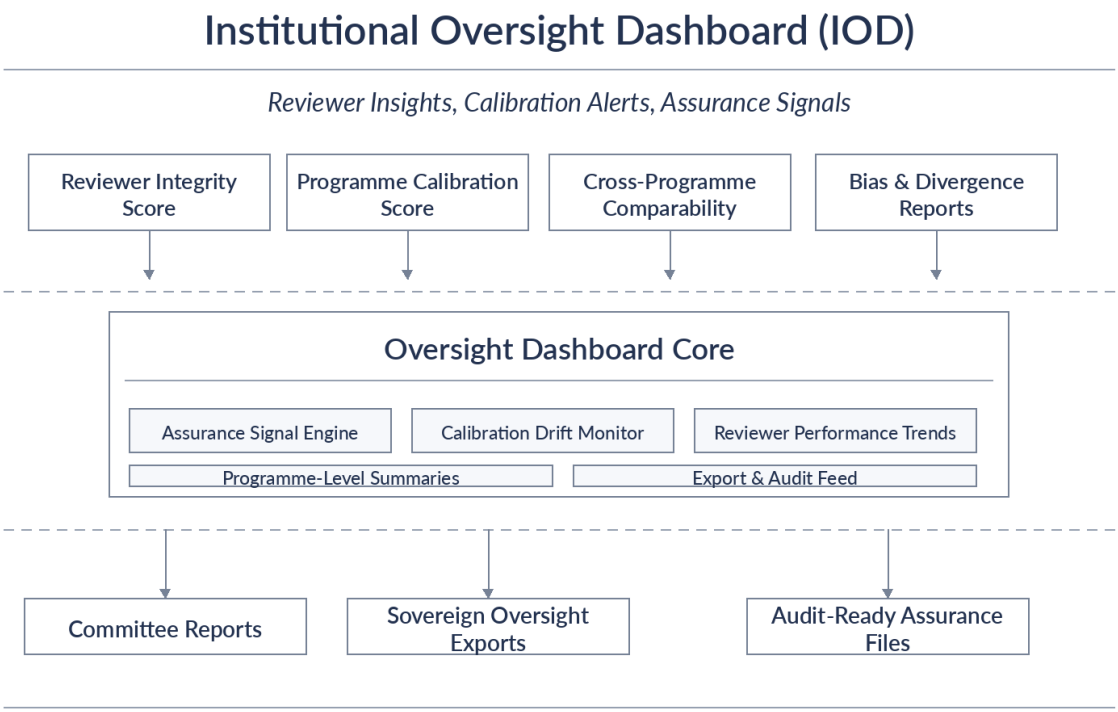

Visual model for institutional review, calibration, and reviewer-assurance workflows.

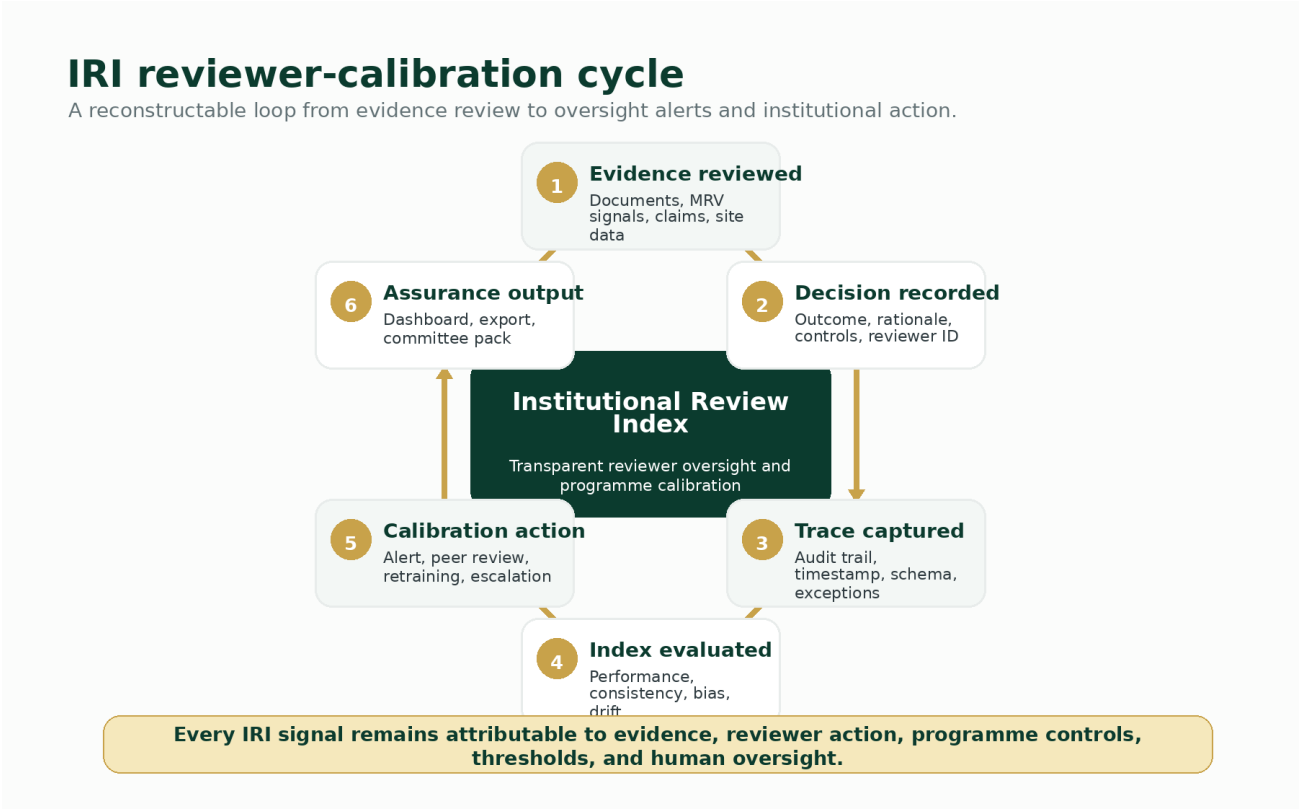

Figure 2. IRI reviewer-calibration cycle.

5. Methodology and Scoring Logic

The IRI methodology is transparent by design. Every score or signal must be explainable through source records. The system should always be able to answer: what evidence was reviewed, what decision was made, who made it, which governance controls applied, what rationale was recorded, how the score was generated, and why an alert was triggered.

5.1 Evidence-linked scoring

Every IRI score is tied directly to records in the governed environment. This is the distinction between a reviewer-assurance system and a black-box performance model.

Input class	Examples	Traceability requirement
Evidence reviewed	Documents, attachments, MRV signals, field records, satellite indicators, audit artefacts, submissions.	The source evidence must remain linked to the reviewer action and exportable record.
Decision made	Approve, reject, request clarification, escalate, override, route, defer, or conditionally clear.	The decision must be timestamped, attributable, and linked to governing criteria.
Rationale captured	Reviewer notes, justification, condition text, exception reason, evidence caveat.	The rationale must explain why the action was taken and which evidence was relied upon.
Controls applied	Role permission, threshold rule, escalation rule, approval workflow, programme-specific policy.	The IRI must know whether controls were followed or bypassed through governed exception.
Peer or programme context	Peer reviewer outcomes, programme baseline, historical decision profile, comparable cases.	Comparability must be calibrated to the relevant programme context.

5.2 Multi-layer calibration

Calibration occurs across three levels. The purpose is not to force identical outcomes, but to ensure review behaviour is stable, explainable, and aligned with the institution's agreed assurance posture.

Calibration level	Question answered	Typical output
Reviewer level	Is this reviewer applying criteria consistently and within expected controls?	Reviewer Integrity Score, Reviewer Consistency Score, drift flag, peer divergence note.
Programme level	Are reviewers within the programme operating on a stable and comparable baseline?	Programme Calibration Score, decision-quality trend, exception distribution.
Institutional level	Can programmes be compared across a portfolio, country, fund, or sovereign operating environment?	Institutional Assurance Signal, Cross-Programme Comparability Index, committee oversight summary.

5.3 Illustrative scoring framework

The IRI should be configurable by deployment, programme, sector, and institution. The following framework illustrates how scoring can be structured without becoming opaque or universalising local context.

Dimension	Indicative weight	Measurement basis	Governance note
Evidence alignment	25%	Decision rationale and referenced evidence match the applicable criteria and evidence threshold.	Cannot be scored without linked source evidence.
Control adherence	20%	Required routing, permissions, approvals, overrides, and escalation rules were followed.	Exception handling remains valid when documented and authorised.
Consistency	20%	Reviewer behaviour remains comparable across similar cases and stable over time.	Contextual differences must be recorded before being treated as divergence.
Timeliness	15%	Actions occur within expected review windows or escalation rules are triggered.	Timing metrics should reflect programme complexity and reviewer capacity.
Completeness	10%	Required fields, rationale, evidence links, and decision metadata are present.	Completeness is a governance quality signal, not a substitute for judgment.
Bias/divergence posture	10%	Outlier patterns are absent, explained, or escalated for human review.	Flags are prompts for review, not conclusions of misconduct.

Scoring boundary

The IRI should always preserve an explanation trail. A committee must be able to reconstruct the score from the records and understand which assumptions, thresholds, and programme rules were used.

5.4 Drift detection

Drift detection monitors whether reviewer behaviour, programme review posture, or cross-programme comparability changes materially over time. Drift can be benign, explainable, or problematic. The IRI's role is to make it visible and governed.

- Reviewer drift: a reviewer's decision profile or criteria application changes materially from their own historical baseline.
- Programme drift: a programme's review posture changes materially from its approved operating baseline.
- Peer divergence: one reviewer's outcomes differ materially from comparable peer reviewers under comparable criteria.
- Cross-programme divergence: comparable programmes begin to apply assurance or escalation thresholds differently without documented rationale.
- Thematic or geographic skew: decision patterns cluster around regions, themes, cohorts, operators, or evidence categories in a way that requires explanation.

5.5 Sovereign-safe auditability

The IRI is designed for environments where data ownership, jurisdiction, access boundaries, and public accountability matter. Outputs must be exportable and reconstructable without forcing unauthorised data movement or exposing unnecessary personal data.

- Role-governed access to reviewer-level, programme-level, and institution-level outputs.
- Programme-specific export rules and jurisdictional controls.
- Evidence-linked audit trails for score reconstruction.
- Reviewer identity binding where required, with privacy-aware disclosure where configured.
- Exception registers for overrides, threshold breaches, and calibration interventions.

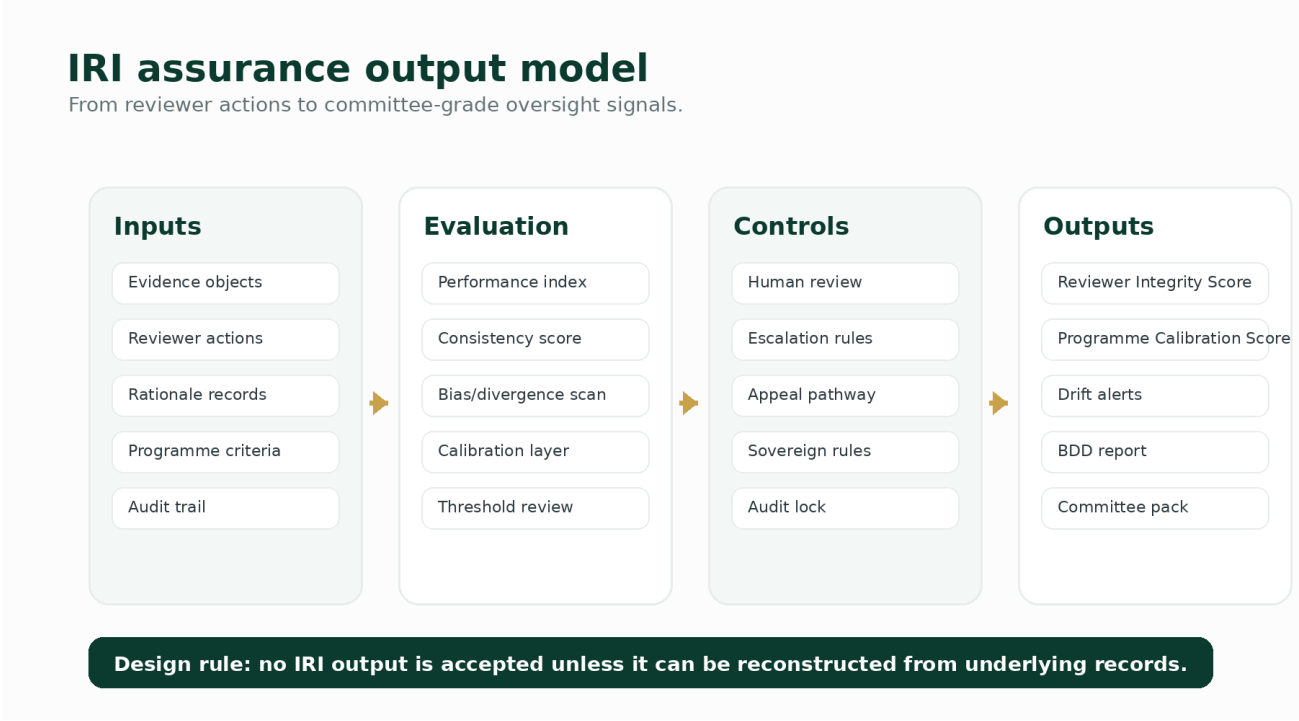

Figure 3. IRI assurance output model.

6. Institutional Use Cases

The IRI is useful wherever review decisions must be trusted across institutions, programmes, countries, reviewers, evidence types, or funding cycles. It is especially relevant to multi-actor operating environments where accountability is distributed but the need for assurance is centralised.

Institutional user	Primary assurance need	IRI contribution
Ministries and national authorities	Oversight of national programme reviewers and inter-ministerial coordination.	Provides sovereign-safe reviewer oversight, programme calibration, and reconstructable reporting for national governance.
DFIs and development banks	Reviewer assurance for funding decisions and cross-country comparability.	Supports risk-aligned programme evaluation, funding-readiness review, and committee assurance without replacing DFI decision authority.
Climate funds and donors	Transparent reviewer behaviour across thematic areas, allocations, safeguards, and MRV-linked decisions.	Enables defensible allocation discussions, reviewer consistency checks, and calibration across themes and programme windows.
Auditors and fiduciary reviewers	Traceable reviewer performance and reconstructable decision chains.	Provides reviewer-level accountability, evidence lineage, bias/divergence signals, and exportable oversight records.
Programme management offices	Operational visibility over review bottlenecks, reviewer drift, and programme assurance posture.	Helps identify training needs, capacity constraints, recurring exceptions, and escalation points.
Portfolio governance teams	Comparable assurance across multiple programmes, countries, sectors, or funding instruments.	Generates cross-programme comparability metrics while preserving programme-specific rules.

6.1 Example oversight scenarios

- A DFI committee asks whether review outcomes are comparable across two countries before approving a second tranche.
- A ministry wants to know whether regional review teams are applying national programme criteria consistently.
- A climate fund wants assurance that safeguards-related reviews are not drifting from approved thresholds.
- An auditor needs to reconstruct why a reviewer overrode a standard routing rule and whether similar overrides were treated consistently.
- A programme lead needs to identify whether decision delays are caused by reviewer capacity, missing evidence, unclear criteria, or escalation bottlenecks.

7. Assurance Outputs and Dashboard Model

The IRI produces outputs designed for committees, procurement teams, auditors, sovereign oversight bodies, programme leads, and institutional funders. These outputs are not isolated metrics. Each output should be linked back to underlying evidence, reviewer actions, rationale, controls, and audit trails.

Output	Purpose	Primary audience
Reviewer Integrity Score	Summarises reviewer behaviour across evidence alignment, consistency, completeness, timeliness, and control adherence.	Programme leads, institutional oversight, auditors.
Programme Calibration Score	Shows whether programme reviewers are operating on a stable and comparable baseline.	Programme owners, ministries, DFIs, portfolio teams.
Institutional Assurance Signal	Aggregates reviewer-process health at institutional or portfolio level.	Committees, sovereign authorities, funders.
Cross-Programme Comparability Index	Indicates whether programmes can be compared using a shared assurance baseline.	DFIs, climate funds, donors, portfolio governance teams.
Reviewer Drift Alerts	Flags changes in reviewer behaviour or decision patterns requiring review.	Programme leads, quality assurance, audit teams.
Bias & Divergence Reports	Identifies systematic skew, peer divergence, or anomalous decision patterns for human review.	Oversight committees, auditors, safeguarding and fiduciary teams.

7.1 Institutional Oversight Dashboard blueprint

The dashboard should be designed as an oversight surface, not a visual scoreboard. It should help authorised users determine whether reviewer behaviour is explainable, whether thresholds are respected, and whether action is required.

Figure 11. Assurance Signal Map: signal aggregation, risk weighting, verification pathways, and confidence scoring.

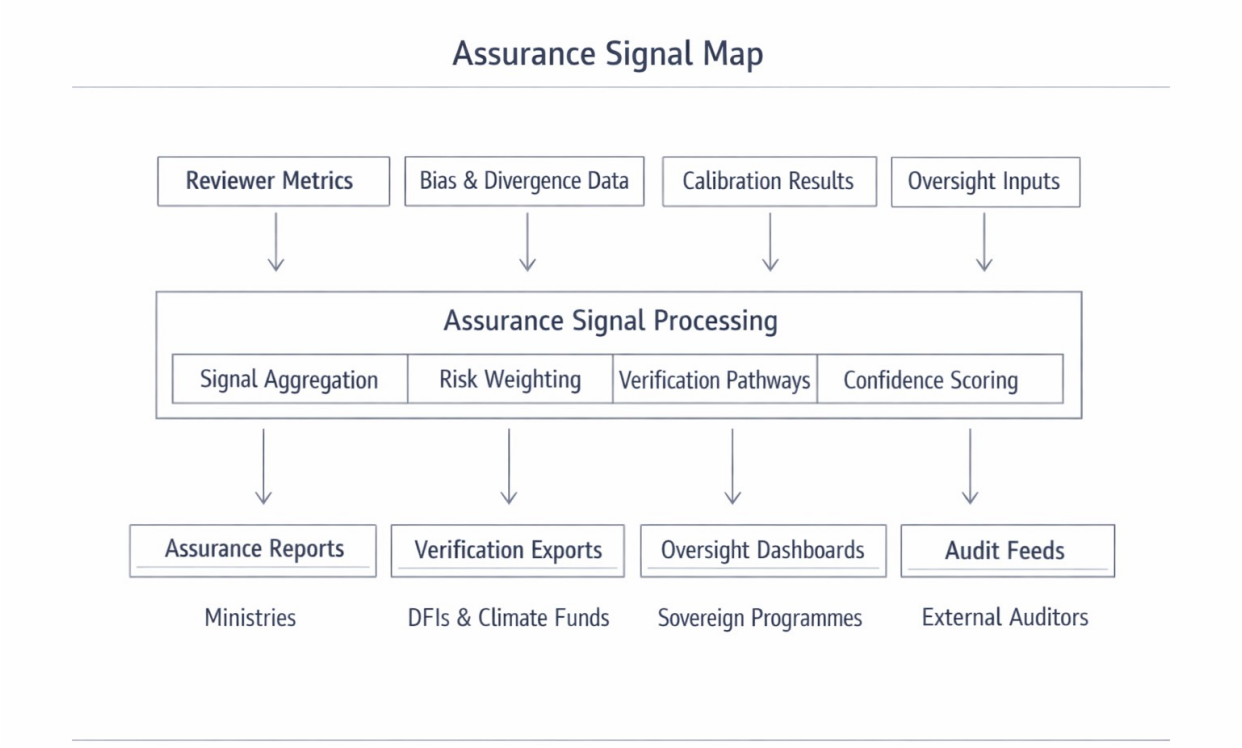

Visual model for institutional review, calibration, and reviewer-assurance workflows.

Dashboard panel	What it shows	Decision-support value
Reviewer assurance profile	Reviewer Integrity Score, Consistency Score, evidence alignment, control adherence, timeliness, completeness.	Helps programme leads identify capacity, training, quality, or escalation needs.
Programme calibration posture	Programme Calibration Score, review distribution, threshold breaches, exception rates, unresolved drift.	Supports programme-level assurance and committee-readiness review.
Bias and divergence queue	Flagged patterns, peer divergence, geographic or thematic skew, and assigned review status.	Ensures sensitive signals are handled through governed human review.
Cross-programme comparability	Comparable assurance baselines, calibration status, programme deviation, and evidence-maturity differences.	Supports DFI, ministry, donor, and portfolio evaluation.
Audit reconstruction panel	Links from signal to source decisions, evidence, rationale, thresholds, and audit logs.	Allows committees and auditors to reconstruct why an IRI output exists.

Committee interpretation rule

An IRI signal should never be treated as a final decision. It should be read as a governed oversight signal that tells a committee where to look, what changed, and which records explain the signal.

8. Governance, Accountability, and Safeguards

Because the IRI evaluates reviewer behaviour, it must itself be governed carefully. The credibility of the IRI depends on transparent scoring rules, clear authority boundaries, appeal pathways, privacy-aware access, and audit-ready evidence chains.

8.1 Governance rules

- Scoring rules must be documented, versioned, and available to authorised institutional reviewers.
- Thresholds must be programme-configurable and linked to governance policy, not hidden model behaviour.
- Reviewer-level outputs must be role-governed and shared only within approved institutional permissions.
- Bias and divergence signals must trigger human review before any operational consequence is applied.
- IRI exports must include score explanation, evidence linkage, threshold basis, and audit reconstruction references.
- Exceptions, overrides, appeals, and calibration interventions must be recorded in an audit-ready register.

8.2 Reviewer accountability without punitive automation

The IRI should make reviewer behaviour accountable without creating an automated disciplinary system. Reviewer accountability is strongest when signals are transparent, contextual, evidence-linked, and subject to human interpretation.

Safeguard	Purpose
Right to explanation	Reviewers and authorised supervisors can see the reason for flags, scores, or drift signals.
Context capture	Programme complexity, evidence quality, workload, policy changes, and exceptional circumstances can be recorded.
Human review gate	Sensitive findings are reviewed by authorised humans before action is taken.
Appeal or correction pathway	Incorrect data, missing context, or misclassified decisions can be corrected through governed processes.
Role-based disclosure	Reviewer-level details are exposed only to authorised users, while committee exports can use aggregated or privacy-aware views.

8.3 Sovereign and fiduciary posture

For sovereign programmes, the IRI must respect the authority of the public institution. For DFIs and donors, it must provide enough transparency to support fiduciary confidence without taking over financial or statutory decision-making. For auditors, it must preserve the ability to reconstruct reviewer behaviour from records rather than relying on opaque summaries.

> **Authority boundary**
>
> The IRI is an assurance and oversight layer. It supports accountable review. It does not approve projects, release funds, certify MRV outcomes, replace legal review, or override sovereign authority.

9. Implementation Pathway

The IRI can be adopted progressively. A phased approach allows institutions to begin with internal calibration and then expand toward programme-level oversight, sovereign-safe reporting, and DFI integration.

Phase	Scope	Core activities	Readiness output
Phase 1 - Internal Calibration	Reviewer-level scoring and internal quality assurance.	Configure metrics, link reviewer actions to evidence, create initial RPI/RCS signals, test drift detection, define escalation rules.	Internal calibration memo and reviewer assurance baseline.
Phase 2 - Programme-Level Oversight	Programme calibration and dashboard visibility.	Introduce Programme Calibration Score, comparability thresholds, IOD dashboard panels, exception queues, and programme oversight exports.	Programme calibration report and committee-ready assurance snapshot.
Phase 3 - Sovereign & DFI Integration	Institutional reporting, procurement-grade assurance, and multi-country comparability.	Configure sovereign export rules, privacy-aware reviewer disclosure, cross-programme comparability index, DFI-facing assurance packs, and audit reconstruction workflows.	Sovereign-safe IRI assurance pack and DFI review annex.

9.1 Minimum viable implementation

A minimum viable IRI implementation should not attempt to produce a mature institutional rating on day one. It should first establish clean evidence links, reviewer attribution, decision taxonomy, threshold configuration, and audit reconstruction. Without these foundations, scoring becomes decorative rather than institutionally useful.

- Reviewer identity binding and role context.
- Programme-specific decision taxonomy and evidence taxonomy.
- Required rationale fields for key decision types.
- Evidence-to-decision link integrity.
- Baseline review windows and escalation thresholds.
- Initial peer-comparison and drift-detection rules.
- Governed exception and appeal pathway.

9.2 Institutional acceptance criteria

Acceptance criterion	Test
Reconstructability	Can every IRI signal be traced to evidence, decision, rationale, reviewer, criteria, threshold, and timestamp?
Explainability	Can an authorised reviewer understand why a score, flag, or alert exists?
Authority boundary	Does the IRI avoid making approvals, release decisions, statutory determinations, or MRV certifications?
Role governance	Are reviewer-level details visible only to authorised users?
Sovereign-safe export	Do exports respect jurisdiction, residency, programme, and disclosure rules?

Acceptance criterion	Test
Human oversight	Are bias, divergence, and drift alerts routed to human review before operational consequence?
Comparability	Can programme-level outputs be compared without erasing local configuration or methodology differences?

10. Conclusion

The Institutional Review Index is the final pillar of Terra Vita Hub's institutional governance architecture. It makes reviewer oversight, programme calibration, institutional comparability, and sovereign-safe accountability visible within a reconstructable operating environment.

The Governance Spine establishes structure. The Institutional Assurance Layer preserves proof. The Institutional Review Index governs the reviewer process itself. Together, they provide a complete institution-grade governance environment for ministries, DFIs, climate funds, sovereign programmes, donors, portfolio operators, and auditors.

The central value of the IRI is institutional confidence. It allows stakeholders to see whether reviewers are operating consistently, whether decisions are evidence-aligned, whether bias or divergence is being monitored, whether programmes are comparable, and whether oversight can be reconstructed when questioned.

Final institutional statement

With the IRI, Terra Vita Hub moves beyond workflow management and evidence storage into calibrated institutional oversight: a governance environment where reviewers, decisions, evidence, and programme assurance can be trusted because they can be examined.

Annex A - IRI Assurance Output Register

The following register can be used as a reference model for institutional reporting, committee packs, export centres, and oversight dashboards.

Output	Level	Description	Recommended export treatment
Reviewer Integrity Score	Reviewer	Composite evidence-linked reviewer assurance signal.	Restricted reviewer-level export; aggregate for committee use where appropriate.
Reviewer Consistency Score	Reviewer	Measure of stable criteria application and predictable evidence interpretation.	Restricted or anonymised depending on institutional role and policy.
Reviewer Drift Alert	Reviewer / Programme	Threshold breach indicating material change in reviewer behaviour or programme posture.	Export with explanation, threshold, supporting records, and status.
Bias & Divergence Report	Reviewer / Programme	Flagged anomalous skew, peer divergence, or thematic/geographic pattern.	Sensitive export; requires human review status and context notes.
Programme Calibration Score	Programme	Programme-level indicator of review baseline stability and comparability.	Committee-ready export with trend and exception register.
Institutional Assurance Signal	Institution / Portfolio	Aggregate oversight signal for institutional governance health.	Executive summary export; avoid exposing unnecessary personal data.
Cross-Programme Comparability Index	Portfolio / DFI	Comparability measure across programmes, countries, sectors, or funding windows.	DFI/donor export with methodology notes and configuration boundaries.

Annex B - Minimum Data Schema for IRI Readiness

The IRI depends on clean data linkage. The following fields represent a minimum readiness model. Actual deployments may add jurisdictional, sectoral, MRV, funding, or programme-specific fields.

Data object	Minimum fields
Reviewer profile	reviewer_id, organisation_id, role, programme_access, reviewer_status, calibration_group, effective_dates.
Decision record	decision_id, reviewer_id, programme_id, evidence_id, decision_type, outcome, timestamp, rationale, criteria_reference, threshold_reference.
Evidence object	evidence_id, programme_id, source_type, document_link, MRV_indicator_link, uploaded_by, uploaded_at, validation_status.
Control record	control_id, rule_type, permission_check, routing_rule, escalation_rule, override_status, approval_requirement.
Audit log	audit_id, actor_id, action_type, object_id, prior_state, new_state, timestamp, source_ip_or_session, schema_version.
IRI score record	score_id, reviewer_id, programme_id, score_dimension, score_value, weighting, calculation_version, source_records, generated_at.
Alert record	alert_id, alert_type, threshold, triggering_records, severity, assigned_owner, review_status, resolution_note, closed_at.

Annex C - Committee Review Questions Answered by the IRI

- Are reviewers applying criteria consistently across similar cases?
- Can a reviewer's score or alert be reconstructed from evidence, rationale, controls, and audit logs?
- Has reviewer behaviour drifted since the previous review cycle?
- Are any geographic, thematic, cohort, or operator-related patterns unexplained?
- Are programme review outcomes comparable across programme instances or countries?
- Were overrides and exceptions properly documented and authorised?
- Are delays caused by reviewer capacity, missing evidence, unclear criteria, or escalation bottlenecks?
- Can the institution defend its reviewer oversight posture to auditors, DFIs, climate funds, or procurement bodies?

Annex D - Glossary

Term	Definition
Governance Spine	Terra Vita Hub's controlled operating structure for evidence intake, review routing, workflow controls, approvals, audit trails, and export readiness.
Institutional Assurance Layer (IAL)	The proof layer that links evidence, reviewer attribution, audit trails, controls, and assurance outputs.
Institutional Review Index (IRI)	The meta-governance layer that evaluates reviewer performance, consistency, bias, divergence, calibration, and comparability.
Reviewer Performance Index (RPI)	A reviewer assurance signal measuring evidence alignment, timeliness, completeness, control adherence, and decision quality indicators.
Reviewer Consistency Score (RCS)	A signal measuring stable application of criteria and predictable evidence interpretation over time and across similar cases.
Bias & Divergence Detection (BDD)	A monitoring capability that identifies potentially systematic skew, peer divergence, or anomalous decision patterns for human review.

Term	Definition
Programme Calibration Layer (PCL)	The programme-level calibration mechanism supporting comparable reviewer assurance across programme contexts.
Institutional Oversight Dashboard (IOD)	The operational interface for reviewer assurance, programme calibration, drift alerts, and committee-ready oversight signals.

Terra Vita Hub - Governance infrastructure for institution-grade climate, agriculture, mining, coastal, and sovereign programme environments.