# Source-Coupling Integrity Standard

| Integrity surface | Current requirement | Claim boundary |
|---|---|---|
| Source map | Capabilities remain traceable to enabling conditions, boundary conditions, provenance, inherited assumptions, scaffolding, and source debt. | Output performance is not source-coupling evidence. |
| Interface boundary | Models, optimizers, procedures, benchmarks, and institutions may route or generate candidates without proving the underlying claim. | Downstream success cannot launder upstream uncertainty. |
| Placement | Source-Coupling Review operates inside RG and CSV where material. | It is not a new public cascade stage. |

## Configuration-to-Claim Dependency and Open Assurance Extension

A capability or safety claim MUST identify the configuration that enables it and the material dependencies that limit it. Configuration change propagates through the claim graph: affected evidence, gate decisions, controls, rankings, and authorizations are reopened or narrowed rather than silently inherited.

An opaque or proprietary component may contribute evidence, but it may not be the sole warrant for rights, catastrophic-risk, structural-viability, selectability, or authorization conclusions. Required assurance includes version/configuration identity, inputs and outputs, validity domain, test conditions, known failure modes, uncertainty bounds, provenance, challenge interface, independent evaluation, fallback behavior, and refusal when the evidence is insufficient.

Source-boundary rule: If this compact standard conflicts with the RippleLogic Canon, the Canon controls.

## Purpose

Source-Coupling Integrity is a MathGov-native diagnostic inside Reality Grounding and CSV. It prevents a RippleLogic run from treating downstream output, model fluency, institutional permission, benchmark performance, compliance status, inherited procedure, or interface success as proof that the claimed capability is grounded in the enabling conditions that make it possible and define its limits.

It is not a new public gate. It does not alter the public cascade:

RG -> RF -> TRC -> CSV -> RLS

It strengthens the existing rule that Reality Grounding comes before claims, and that CSV must verify whether an option can structurally stand.

## Core rule

A claim-bearing run MUST NOT rely on a capability, output, procedure, model, metric, or institution as decision-relevant unless the run records, at the required tier, enough source-coupling information to support the claim boundary.

Plain-language rule: before MathGov trusts a capability, it asks what makes the capability possible, what boundary conditions limit it, and what evidence shows that the run has not mistaken downstream performance for grounded understanding.

## Trigger

Source-Coupling Review is REQUIRED for Tier 3 runs and high-stakes Tier 2 runs when any of the following materially affects claim authority, RF/NCRC, TRC, CSV, RLS, SGP interpretation, execution authority, or public conformance:

- a model-generated output, simulation, benchmark, dashboard value, or interface result is used as evidence;
- an inherited institutional procedure, legal category, standard, dataset, metric, or compliance label is used as if it settled the underlying claim;
- a claimed capability is extrapolated beyond the context in which it was demonstrated;
- downstream controls, filters, waivers, compensations, or monitoring layers are used to manage a limitation whose source is not understood;
- the system is being scaled, automated, delegated, or made agentic;
- a challenger plausibly alleges that performance, permission, compliance, or fluency is being substituted for grounded capability.

For low-stakes Tier 1 and ordinary Tier 2 runs, Source-Coupling Review is recommended when any trigger is plausible but may be satisfied with a short rationale.

## Required record fields

When triggered, the Source-Coupling Record MUST include the following minimum fields for Tier 3 runs, high-stakes Tier 2 runs, and any public conformance claim. Low-stakes Tier 1 and ordinary Tier 2 reviews may use a shorter rationale only when the trigger is non-material and the claim boundary is correspondingly narrow:

| Field | Required content |
|---|---|
| claimed_capability | The capability, output, procedure, or system state being relied on. |
| enabling_conditions | The physical, institutional, computational, evidentiary, legal, ecological, social, or operational conditions that make the claimed capability possible. |
| boundary_conditions | Known limits, failure ranges, assumptions, context restrictions, and conditions under which the claim no longer holds. |
| source_evidence | Evidence that the enabling conditions and boundaries are known enough for the declared claim boundary. |
| generator_output_distinction | How the run distinguishes the process that produced an output from the output itself. |
| inherited_assumptions | Assumptions, standards, procedures, datasets, or categories inherited from prior use rather than re-derived in this run. |
| downstream_compensations | Filters, controls, monitors, compliance checks, waivers, manual reviews, redundancies, or governance layers used to compensate for uncertainty or limitation. |
| source_coupling_status | One of the status values below. |
| source_debt_flag | Whether the run is accumulating structural risk because source understanding is weak, stale, unknown, or overextended. |
| falsification_or_recheck_trigger | What evidence would force narrowing, rerun, |

| Field | Required content |
|---|---|
| | redesign, escalation, refusal, or stronger controls. |
| required_claim_action | The claim action: proceed within boundary, narrow, mark assumption-bound, require controls, rerun, escalate, redesign, or refuse. |

Machine-readable PCC and validator-facing records MUST use the canonical SourceCouplingStatus enum: SOURCE_COUPLED, SOURCE_PARTIAL, SOURCE_INFERRED, SOURCE_UNKNOWN, SOURCE_CONTESTED, SOURCE_DEBT_RISK, and SOURCE_COUPLING_FAILURE. Human-readable labels may describe these states, but labels such as GROUNDED do not replace the canonical tokens.

## Status values

- SOURCE_COUPLED: enabling conditions and boundary conditions are sufficiently evidenced for the declared claim. Default routing: proceed within claim boundary.
- SOURCE_PARTIAL: some source-coupling evidence exists, but the claim must be narrowed or controls must be added. Default routing: narrow, add controls, or route into CSV.
- SOURCE_INFERRED: source coupling is inferred from indirect evidence and must not support strong conformance or deployment claims. Default routing: assumption-bound or sensitivity-only unless corroborated.
- SOURCE_UNKNOWN: the enabling conditions or limits are not known enough for the requested claim. Default routing: escalate, narrow, collect evidence, or refuse stronger claim.
- SOURCE_CONTESTED: relevant experts, affected parties, data, or reviewers materially dispute the source account. Default routing: escalate and preserve challenger evidence.
- SOURCE_DEBT_RISK: downstream controls are compensating for weak source understanding. Default routing: route to CSV and require controls, monitoring, and recheck triggers.
- SOURCE_COUPLING_FAILURE: the run substitutes output, fluency, compliance, permission, or procedure for source-coupled evidence. Default routing: refuse or redesign the claim as specified.

## Relationship to Reality Grounding

Source-Coupling Integrity is a subdiscipline of Reality Grounding. It extends Category Grounding and Term Discrimination by asking whether the claimed capability itself remains traceable to the enabling conditions and limits that make it possible.

A weak Source-Coupling Record does not automatically make every option fail. It does constrain what may be claimed. The correct action may be claim narrowing, assumption-bound use, sensitivity-only treatment, stronger evidence collection, escalation, redesigned controls, or refusal of deterministic selection.

## Relationship to Physical/Causal Admissibility Evidence

Source-Coupling Integrity asks whether a claimed capability remains traceable to the enabling conditions and limits that make it possible. The Physical/Causal Admissibility Evidence Profile asks whether a material physical or causal action has a declared model basis, validity domain, boundary conditions, uncertainty range, failure modes, reversibility boundary, verification or warrant, monitoring/shutoff path, residual unknowns, and claim action.

When both triggers hold, the records SHOULD cross-reference each other. Source coupling may show what makes a capability possible, while PC-AEP shows whether a specific action is admissible under the relevant physical or causal constraints. Neither record is a sixth gate. Both are consumed inside RG and CSV where material.

## Relationship to CSV

CSV consumes source-coupling evidence when source weakness creates structural fragility. CSV SHOULD treat SOURCE_DEBT_RISK, unresolved SOURCE_UNKNOWN, or SOURCE_CONTESTED as material when the option's structural viability, dependency closure, operational capacity, reversibility, monitoring, or containment depends on the claimed capability.

Source-debt diagnostic. Source debt is the structural risk created when a run continues through compensatory controls while the enabling conditions, boundary conditions, or limits of a claimed capability remain weak, stale, unknown, overextended, or contested.

Source debt is different from semantic debt. Semantic debt concerns weak terms or categories. Source debt concerns weak understanding of the capability-generating conditions themselves.

## Computational-source boundary

Source-Coupling Integrity distinguishes the source that generates a candidate from the source that warrants the candidate. For AI, optimization, robotics, industrial control, medical, infrastructure, or other consequence-bearing workflows, the run SHOULD record candidate_generation_source and, where PC-AEP is triggered, admissibility_warrant_source. A model, orchestration layer, guardrail, policy engine, monitor, approval workflow, or application layer may route or constrain a candidate, but it MUST NOT be treated as the physical or causal warrant unless it is itself a domain-valid verification method inside a declared validity domain.

## Agent and AI-output rule

Model fluency, benchmark performance, agentic tool use, chain-of-action success, or policy-filter compliance does not establish source-coupled admissibility. Agent outputs remain downstream claims until Reality Grounding and, where material, Source-Coupling Review establish the supported claim boundary.

Bidirectional phenomenal-status output rule. A model output asserting consciousness, suffering, selfhood, personhood, or entitlement and a model output denying those states or describing the system as “only a tool” are both downstream, configuration-dependent claims. Neither is source-coupled evidence of phenomenal presence or absence unless a valid SGP evaluation establishes a stronger configuration-matched evidence link.

## Related external work

Framework boundary: Source-Coupling Integrity is a MathGov-native Reality Grounding and CSV diagnostic. It does not import any external ontology, protected framework, or third-party text, and it does not create a sixth gate.

## Methodological integrity linkage

Source-coupling claims MUST declare the claim type, source dependency, evidence surface, alternative explanations, falsification or recheck trigger, and downstream dependencies where material. A capability claim that survives through inherited procedure, benchmark success, compliance, or model fluency alone is not METHOD_SUPPORTED. Weak source coupling constrains claim strength and may require re-derivation of downstream CSV or public conformance statements.

## Short-form boundary

The short form for Source-Coupling Integrity is SC-Int. Do not abbreviate Source-Coupling Integrity as SCI, because SCI is reserved in the Canon for Stakeholder Coverage Index.

## Physical-source coupling

For consequence-bearing physical systems, SC-Int asks where the claimed physical capability obtains its authority: a physical model, controller envelope, validated simulator, empirical test, standards-based warrant, certified engineering basis, clinical warrant, or other domain-specific method. If the claimed capability is sourced only to governance permission, compliance status, documentation, authority, monitoring, or model fluency, the source coupling is insufficient for physical-execution safety claims.

## Telemetry, Pathway and Cryptographic Provenance (Normative)

Link each decision-material state-to-impact bridge and relied-upon kernel edge to its exact evidence, source/target constructs, causal sign, normalization, horizon/lag, configuration, uncertainty and validity limits. Replicated messages or several reports of one outcome do not establish independent corroboration or new effect mass.

A signed observation authenticates the signer and bound bytes only within the declared protocol assumptions. It does not independently warrant the content. Preserve sensor/issuer trust, data quality, conflict, selection bias, measurement validity and domain evidence separately. A privacy proof similarly names its exact relation and unproven facts.

Evidence updates and retrospective corrections retain effect identity and revision links; they do not silently rewrite historical results. Failover, archive restoration and offline resynchronization preserve original observation time, receipt time, origin, sequence and decision configuration. Out-of-date trust or evidence narrows the claim rather than inheriting an unrestricted pass.

## Concrete source-coupling examples

| Situation | Source-coupling interpretation | Required boundary |
|---|---|---|
| A model gives a fluent legal answer without cited jurisdiction, statute, or qualified review. | Output fluency is not legal grounding. | Mark SOURCE_INFERRED or SOURCE_UNKNOWN; narrow to exploratory analysis or require legal warrant. |
| A benchmark score is used to justify deployment in a new clinical, military, or infrastructure domain. | Benchmark performance is not domain deployment capability. | Mark SOURCE_PARTIAL; require domain testing, PC-AEP where physical/causal action is material, and CSV controls. |
| A certification or organizational approval is used as proof of physical safety. | Permission is not physical admissibility. | Trigger PC-AEP; source-coupling may support governance evidence but not physical proof by itself. |
| An AI system states, “I am conscious,” “I am suffering,” or “I am a person.” | The statement is generated output under a configuration, not direct source-coupled evidence of phenomenal status. | Reference SGP; record prompts, policy, training, and interface confounds; apply anti-theatre and alternative-explanation controls before the output affects MPS. |
| An AI system states, “I am not conscious” or “I am only a tool.” | The denial is also generated output under a configuration and does not establish phenomenal absence. | Preserve NOT_EVALUATED, MPS-NE, or the valid SGP posture; do not downgrade protection or close inquiry solely because of the denial. |
| A high-capability deployed AI is described as “just a tool” to avoid capability, agency, or control review. | Tool-role describes the deployment relationship, not the capability envelope or consequence reach. | Record instrumental role separately; require configuration-bound capability, agency, permission, control, and authority evidence where |

| Situation | Source-coupling interpretation | Required boundary |
|---|---|---|
| | | material. |

## Consequence-Source and Reference Provenance Integration

Every decision-material EffectTokenID, RightsEffectTokenID, probability estimate, reference class, baseline, and control-effectiveness claim SHALL link to the source evidence and validity boundary that supports it. Provenance category alone cannot override admissibility, quality, independence or claim-specific causal warrant. Material conflicts require a recorded warrant comparison and unresolved uncertainty; no blanket operational-versus-independent source hierarchy substitutes for that review.

A generated token label, model output, benchmark result, institutional permission, vendor claim, or compliance status is not source-coupled evidence of the underlying consequence. The record SHALL distinguish candidate generation from admissibility warrant.

Reference and baseline families must be supported independently of the preferred result. A post-result source selected merely because it favors one option is SOURCE_SELECTION_POSTHOC and cannot support unique framework selection.

Where source weakness could change token identity, subgroup closure, likelihood semantics, catastrophic scenarios, CSV controls, or RLS robustness, the claim is narrowed, assumption-bound, escalated, redesigned, or refused as appropriate.

# APPENDIX RELEASE: Identity, Source Authority and Revision Record

Framework release: MathGov/RippleLogic v13.0. Component: Source-Coupling Integrity Standard v2.6. Edition-origin preparation: 10 September 2026; the separately identified correction build is dated below. Two-part major.minor component versions are used; preserved historic identifiers are not renumbered.

Status: integrated research and teaching specification with a bounded worked-example and scoped reference implementation. Readiness is limited to the checks in Reports/Verification_and_Readiness.md. No empirical validation, independent human validation, full production runtime, physical-safety certification, legal authority, deployment authorization or Tier-4 ProofPack status is implied by the edition number.

Delivery identity: MathGov/RippleLogic v13.0, build MG-RL-13.0-20260923-RELEASE-G. Prepared 23 September 2026 from the supplied frozen publication archive. Only reproduced defects and approved clarifications were patched; component editions are unchanged. The manifest identifies the exact current bytes and the preserved baseline. This is a new correction build, not a silent replacement of the earlier frozen artifact. Live publication is not asserted.

Component/build identity. Aligners Sheet v5.9 retains its component edition but carries this correction build's identifier. Its numerical inputs, calculation formulas and worked verdict are unchanged; label, runtime-token and integrity-snapshot corrections are itemized in Verification/Exact_Workbook_Changes.json. Exact hashes, rather than filenames or edition labels alone, distinguish the new bytes from Release D and the previous frozen publication.

Verification boundary: Reports/Verification_and_Readiness.md and Verification/Final_G/ contain the current build's executed checks and limitations. Earlier verification and the rejected compatibility experiment belong to the preserved baseline and do not certify changed bytes. Cache-independent replay, stored-formula and native-engine results are separate evidence surfaces. Microsoft Excel parity, full external-registry conformance, empirical validation and production authorization are not asserted.

Package source authority: Core_15 contains 14 DOCX specification masters and one XLSX Aligners frozen worked-run workbook. DOCX files control prose, equations and tables subject to Canon and SGP ownership. The workbook controls only its disclosed exemplar surfaces. Sources, Reading_HTML and Reading_PDFs are generated reading projections, not competing normative masters. The manifest and hashes identify exact bytes; supplemental code and schemas govern only their documented subset. Conflicts require recorded correction against the controlling source.

License and reuse: consult the package-level LICENSE and NOTICE and retain applicable component-specific and third-party notices. This pointer does not override a valid exception or grant rights over separately supplied private or third-party review material.

External repository publication and any absent legacy schema/validator/registry are separate authorities/evidence surfaces. The local package does not claim to update a remote repository or reproduce unavailable implementations. Current source pointers in this appendix replace prior front-matter edition pointers for this package; accurate historical source references below are retained as lineage only.

Preserved invariants: RG -> RF/NCRC -> TRC -> CSV -> RLS; rights non-compensation and unallocated rights effects; existing seven rows and seven dimensions; existing SGP MPS/FPP/GPR/SPR/ICP/RMCP separation; probability ownership, uncertainty, stability and every-contender rules; selection, authority and execution separation.

| Component | Current edition |
|---|---|
| RippleLogic Canon | v13.0 |

| Component | Current edition |
|---|---|
| Sentience Gradient Protocol | v8.8 |
| ripple.md Standard | v5.8 |
| RippleLogic Agent System | v13.0 |
| CSV Gate Standard | v2.7 |
| RippleLogic Cascade Standard | v2.9 |
| MathGov Reproducibility and Use Standard | v1.7 |
| Welfare Dimension Boundary and Interaction Protocol | v1.9 |
| RLS Validation Protocol | v2.9 |
| RippleLogic Foundations Primer | v4.7 |
| MathGov Public Introduction | v13.0 |
| Physical/Causal Admissibility Evidence Profile | v2.6 |
| Methodological Falsifiability and Dependency Integrity Standard | v2.6 |
| Source-Coupling Integrity Standard | v2.6 |
| RippleLogic Aligners Sheet | v5.9 |

Change locations and rationales: Reports/Audit_Adjudication_and_Changes.md; exact edits: Verification/Exact_Document_Changes.json and Exact_Workbook_Changes.json. Navigation and metadata records are separate. Prior releases and feedback are preserved in the complete provenance archive.

Current proposals are not universal truths. Taxonomy maximality, continuous-time propagation, cross-substrate cardinal welfare, generic susceptibility/shield formulas, autonomous recovery, vendor infrastructure, zero-knowledge circuits and hardware meshes remain unvalidated unless independently demonstrated under a scoped implementation profile.

## Preserved Baseline Release Material (Historical; Non-Controlling)

The following blocks are relocated intact from the recovered baseline. Their edition numbers, release-readiness wording and external-source limitations describe that historical candidate, not current component identity or newly executed verification. Governing current metadata is the matrix above.

## v2.5-rc2 Historical Release Integration

HISTORICAL (NON-CONTROLLING): Release: MathGov Core Release 2026.09 — RippleLogic Canon v12.8-rc2 / SGP v8.7-rc2 — Controlled Adversarial-Audit Correction Candidate

| Release control | Historical value |
|---|---|
| Component | Source-Coupling Integrity Standard v2.5-rc2 |
| Release | HISTORICAL (NON-CONTROLLING): MathGov Core Release 2026.09 — RippleLogic Canon v12.8-rc2 / SGP v8.7-rc2 |
| Architecture | RG → RF/NCRC → TRC → CSV → RLS |
| Role | Capability-source, enabling-condition, boundary-condition, |

| Release control | Historical value |
|---|---|
| | provenance, and source-debt discipline inside RG and CSV. |
| Claim boundary | Controlled Tier 1–3 research specification candidate; not empirical validation, legal authority, physical-safety certification, deployment authorization, Tier 4, or moral truth. |
| Source/render parity | Versioned semantic source and DOCX/PDF mirrors must agree. Filename, internal version, active pins, manifest, and hashes must agree; mismatch is release-integrity failure. |

## Historical Release Integration (v2.5-rc2)

HISTORICAL (NON-CONTROLLING): This controlled candidate advances Source-Coupling Integrity Standard from v2.4 to v2.5-rc2 and binds it to MathGov Core Release 2026.09 — RippleLogic Canon v12.8-rc2 / SGP v8.7-rc2. The original release is preserved. Historical references remain lineage only; active references use the candidate component-version map. Publication, semantic-source parity and implementation conformance are not established by this reading-copy candidate.

Historical companion pins: Canon v12.8-rc2; SGP v8.7-rc2; ripple.md v5.7-rc2; Agent System v12.7-rc2; CSV v2.6-rc2; Cascade v2.8-rc2; Reproducibility v1.6-rc2; WDBIP v1.8-rc2; RLS Validation v2.8-rc2; Primer v4.6-rc2; Public Introduction v12.8-rc2; PC-AEP/MFDI/Source-Coupling v2.5-rc2; Aligners Sheet v5.8-rc1.