

Methodological Falsifiability and Dependency Integrity Standard

Integrity surface	Current requirement	Claim boundary
Claim typing	Material empirical, causal, formal, operational, legal-regulatory, normative, and metaphysical-horizon claims declare their evidence and authority boundaries.	The protocol does not make every legal or normative commitment physically falsifiable.
Warrant-domain bridge	A material conclusion identifies the domain of the offered warrant and any bridge required when the conclusion crosses scientific, engineering, legal-regulatory, institutional, or normative domains.	Authority, compliance, certification, performance, and rhetoric cannot silently change claim type.
Dependency integrity	Starting assumptions, definitions, models, thresholds, and inherited sources are versioned; material changes trigger re-derivation.	Silent retrofitting is prohibited.
Revision trigger	A claim identifies evidence or conditions that would force revision, narrowing, rerun, or retirement.	Fluency and coherence are not external confirmation.
Capability claim integrity	A material capability term is operationally defined and bound to observable criteria, mechanism or system composition, evidence, validity conditions, alternatives, a falsifier, evidence maturity, permitted wording, and explicit nonclaims.	Neither component reduction nor anthropomorphic language settles the capability claim.

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

Frozen Evaluation Binding and Milestone Claim Discipline

Before consequential evaluation, the run freezes and versions the decision question, system boundary, configuration, evidence schema, stakeholder scope, rights interpretation, scenario families, thresholds, uncertainty treatment, weighting protocol, control requirements, stopping conditions, and authorization pathway. For a confirmatory operator or performance test it also freezes a finite, enumerated or formally bounded admissible falsification family and a stopping/attribution rule. A material post-result alteration creates a new run or declared revision branch; it cannot turn a failed original prescription into a successful original test.

Implementation milestones are classified narrowly. Successful formal derivation, compilation, synthesis, simulation, hardware implementation, or integration may defeat a specific impossibility hypothesis, but it does not prove empirical superiority, domain validity, physical safety, ontology, rights compatibility, or deployment readiness. Claims must name the exact falsification criterion tested and the alternatives that remain open.

Purpose

This compact standard defines the methodological integrity discipline for MathGov and RippleLogic. It adopts the bounded principle that claim-bearing elements must not survive merely because they are plausible, fluent,

approved, compliant, or locally useful. A material claim must declare what kind of claim it is, what it depends on, what would force revision, what alternatives remain open, and what downstream elements must be rerun if the foundation changes.

It is not a sixth gate. It does not replace Reality Grounding, RF/NCRC, TRC, CSV, SGP, RLS, PC-AEP, domain science, or external validation.

Adopted principles and bounded non-claims

Item	MathGov integration
Adopted principle	Material claims need external contact, falsification or revision triggers, explicit dependencies, and no hidden assumption authority.
Adopted principle	If a starting assumption, definition, threshold, operator, or model basis changes, the affected downstream chain must be rerun or claim strength must be narrowed.
Bounded non-claim	Not every normative definition is physically falsifiable. Normative commitments are audited for coherence, rights compatibility, stakeholder challenge, and application integrity.
Bounded non-claim	Probabilistic evidence is not invalid. It is required where uncertainty is real, but uncertainty must be disclosed and cannot be laundered into certainty.
Bounded non-claim	Governance may act proportionately under serious uncertainty to protect rights, bound ruin, require evidence, authorize bounded testing, or impose controls. Such action does not prove that the uncertain capability exists or that uncertainty has been reduced.
Bounded non-claim	Model revision is not forbidden. Silent retrofitting is forbidden. Revision must be versioned, disclosed, and rechecked.
Bounded non-claim	System-level operational capabilities may emerge from coordinated probabilistic, deterministic, memory, planning, feedback, tool, and human-control components. Component decomposition does not by itself refute the bounded capability, and the bounded capability does not establish consciousness, moral agency, generality, or authority.

Required Methodological Integrity Record

Field	Required content
claim_or_component	The claim, definition, threshold, operator, indicator, model, SGP interpretation, gate decision, validation result, or public release claim under review.
claim_type	empirical, causal, formal, operational, legal_regulatory, normative, or metaphysical_horizon. The display labels "legal-regulatory" and "metaphysical-horizon" MAY be used in prose; validator-facing records MUST use legal_regulatory and metaphysical_horizon.

Field	Required content
claim_domain	The primary authority domain in which the conclusion is asserted: scientific_empirical, engineering_operational, formal_mathematical, legal_regulatory, normative_ethical, institutional_governance, commercial_descriptive, anthropomorphic_social, or metaphysical_horizon. A materially mixed conclusion SHOULD be split into separately typed claims.
warrant_domains	One or more domains from which support derives. Multiple domains MAY support one claim, but each material cross-domain relationship requires its own bridge record.
cross_domain_bridges	Zero or more bridge records. Each record contains from_warrant_domain, to_claim_domain, bridge_status, bridge_proposition, bridge_evidence_or_rationale, uncertainty, validity_or_jurisdiction_boundary, and revision_or_refusal_condition.
dependency_position	Where the claim sits in the cascade, standard, SGP flow, workbook, validation protocol, or public release surface.
starting_assumptions	Assumptions, definitions, priors, sources, thresholds, scope boundaries, or model premises that carry the claim.
definition_or_operator_used	The material definition, operator, threshold, indicator, or procedure being applied.
necessity_or_alternative_check	Whether necessity is claimed, what alternatives remain plausible, and why the stronger claim is or is not warranted.
evidence_or_test_surface	The observation, measurement, replication, derivation, simulation, formal proof, expert review, stakeholder challenge, or audit surface used.
falsification_or_revision_trigger	What result, evidence, challenge, failure, or domain change would force revision, narrowing, escalation, or refusal.
uncertainty_or_confidence_method	How uncertainty, confidence, model limits, disagreement, or unresolved evidence is represented.
downstream_dependencies	Gates, scores, SGP outputs, claims, controls, authority actions, or release statements affected if this element changes.
re_derivation_scope_if_changed	Local patch, threshold update, gate rerun, TRC reopening, CSV rerun, SGP rerun, RLS recalculation, release versioning, or full refusal.
reviewer_status	reviewer-visible, contested, independently checked, domain-expert-warranted, or pending.
required_claim_action	proceed, narrow, control, redesign, escalate, or refuse.

Field	Required content
operational_definition (conditional)	The measurable functional meaning of the capability term in the pinned configuration and domain.
observable_criteria (conditional)	The observations that would count as demonstration, non-demonstration, or failure of the capability.
mechanism_or_system_composition (conditional)	The components and interactions that instantiate the claimed capability without assuming any single component supplies the whole.
evidence_surface (conditional)	The logs, trials, benchmarks, traces, inspections, or other evidence that support the bounded claim.
validity_conditions (conditional)	Configuration, domain, environment, time, operator, and control conditions inside which the claim is licensed.
alternative_explanations (conditional)	Simpler or rival explanations such as scripting, operator intervention, memorization, benchmark fit, or interface effects.
falsification_or_revision_trigger (conditional)	Reuses the general falsification_or_revision_trigger field. For a capability claim, that same field must also identify the observation or boundary violation that would narrow, revise, suspend or refute the capability claim. Do not serialize a second property with the same name.
evidence_maturity_stage (conditional)	Definition/derivation, computational realization, implementation conformance, projection/calibration lock, controlled empirical contact, operational qualification, or outcome/requalification evidence.
permitted_claim (conditional)	The strongest wording licensed by the present evidence and maturity stage.
explicit_nonclaims (conditional)	Stronger inferences not licensed, including consciousness, personhood, intrinsic goals, moral agency, generality, unrestricted autonomy, safety, legality, or authority where applicable.
target_of_understanding (conditional)	The target object, process, system, proposition, or domain state that is claimed to be understood.
relation_claimed (conditional)	The relation asserted, such as explanation, prediction, counterfactual manipulation, intervention, correction, or transfer.
distance_from_training_or_known_exemplars (conditional)	How the test distinguishes transfer or novel application from retrieval, memorization, contamination, or benchmark fit.
scaffolding_and_tools (conditional)	Prompts, retrieval, memory, tools, human assistance, environmental feedback, or other scaffolding required for the demonstrated performance.
transfer_conditions (conditional)	The contexts, perturbations, reduced-scaffolding conditions, or domain shifts across which the understanding claim is expected to hold.

Dependency-localized failure propagation

If a premise, primitive, definition, operator, evidence source, control, or inherited dependency fails, every materially dependent gate result, score, SGP output, authorization, validation claim, or public statement must be reopened, narrowed, suspended, rerun, or withdrawn. The record must identify the dependency path and re-derivation scope.

A component with an independent evidence and derivation path does not fail merely because another part of the framework fails. This rule prevents both monolithic immunity—protecting dependent claims after foundation failure—and indiscriminate whole-framework rejection—discarding independently supported components without testing their actual dependencies.

Taxonomy, Formalization and Engineering Claim Discipline (Normative)

Orthogonality, independence, irreducibility, optimality, observability and controllability claims must state their exact mathematical or operational meaning, object, assumptions, scope and falsifier. Conceptual distinction, uncorrelated variables, statistical independence, causal separation and non-duplicated accounting are not equivalent. A useful architectural vocabulary is not a proved theorem.

Do not claim that seven dimensions/scopes are maximally complete or irreducible from symmetry, selected examples or an uncalibrated sum of loss criteria. Preregister plausible merges/splits and report the trade-offs separately. An aggregation rule, unit test or constrained counterexample establishes only its stated proposition.

Production readiness, zero data loss, privacy, cryptographic security, exact recovery, real-time performance and successful hardware deployment require exact executable artifacts and corresponding evidence. Source code, mock logs, invented signatures and illustrative scenario parameters do not establish these claims. Cryptographic verification proves a declared formal relation under assumptions, not measurement truth, lawful authority, empirical welfare or absence of omitted harms.

The proposed control envelopes, degraded-mode and privacy interfaces strengthen existing records. The speculative susceptibility formula, automatic 90-percent counter-injection, multiplication-based identity commitment, sacred quorum/latency constants and biosphere/light-cone identity are not adopted. Research branches preserve failed tests and must not redefine success after observing outcomes.

Status values

Status	Meaning	Default action
METHOD_SUPPORTED	Evidence, dependency, and revision conditions support the declared claim boundary.	Proceed within boundary.
ASSUMPTION_BOUND	Claim depends on visible assumptions that could affect strength or gate status.	Narrow or monitor.
ALT_EXPLANATION_OPEN	Plausible alternatives can explain the same result.	Downgrade necessity claim.
TEST_REQUIRED	Stronger evidence, derivation, replication, simulation, or review is needed.	Escalate before strong claim.
CONTESTED	Evidence, interpretation, stakeholder challenge, or expert judgment materially disputes the claim.	Escalate, narrow, or redesign.
REVISION_REQUIRED	A starting assumption,	Rerun affected chain.

Status	Meaning	Default action
	definition, threshold, operator, or model basis has failed or changed.	
REFUSE_STRONGER_CLAIM	The requested public, deployment, validation, deterministic-selection, or safety claim exceeds support.	Refuse the stronger claim.

Claim-type rules

Claim type	Evidence discipline
Empirical	Observation, measurement, replication, current source, or empirical warrant.
Causal	A claim-specific identification plan links mechanism, intervention and comparator to a target estimand, population and time horizon; declares identification assumptions, selection/confounding and interference treatment, measurement and uncertainty; and states the identification-status verdict and resulting claim boundary.
Formal	Explicit assumptions, definitions, derivation, proof, or formal verification.
Operational	Reproducible procedure, audit trail, role assignment, and reviewable outputs.
Legal-regulatory	Current authority source, jurisdiction, applicability boundary, interpretation status, and contest or review path.
Normative	Declared ethical commitment, rights compatibility, consistency, stakeholder challenge, and auditability of application.
Metaphysical-horizon	Humility and meaning orientation only; no Tier 1-3 evidence, gate override, or authority by itself.

Claim-domain and warrant-domain non-substitution

A material claim **MUST** distinguish its `claim_type`, its primary `claim_domain`, and every material `warrant_domain` offered in support. Scientific or empirical evidence can support bounded descriptive or causal claims but does not by itself create legal authority or ethical legitimacy. Engineering performance can support bounded operational capability but does not by itself establish rights compatibility, lawful permission, or safety outside the tested envelope. Legal, regulatory, institutional, certification, compliance, or policy authority can govern permission, procedure, responsibility, testing, or review but does not by itself establish empirical truth, causal validity, technical capability, physical safety, or reduced epistemic uncertainty. Commercial or anthropomorphic terminology creates no stronger evidence or authority by itself.

When a conclusion draws on more than one warrant domain, the record **MUST** list all material warrant domains. For every warrant domain that differs from the claim domain, it **MUST** create a separate bridge record stating the bridge proposition, evidence or rationale, uncertainty, relevant validity or jurisdiction boundary, and the condition that would narrow, revise, or refuse the bridge. `BRIDGE_PARTIAL` permits only the supported narrower claim. `BRIDGE_UNSUPPORTED` requires narrowing, escalation, redesign, or refusal

of the cross-domain conclusion and triggers CROSS_DOMAIN_WARRANT_SUBSTITUTION if the stronger claim is nevertheless asserted.

Non-retrofitting rule

A failed test, falsified prediction, changed foundation, or challenged dependency must not be hidden by silent tuning. Preserve the outcome of the locked original operational prescription even when attribution to a general operator remains unresolved. The run must version the change, record the affected dependencies, and rerun every materially dependent gate, SGP interpretation, calculation, authority action, or public claim. A repaired version needs its own declared test; its success does not erase the original failure.

Proportionality rule

Methodological rigor is proportional to stakes. Tier 1 and low-stakes Tier 2 uses may use lighter records. Tier 3, public conformance, institutional deployment, high-stakes SGP use, validation claims, and physical/causal actions require stronger records.

Physical-admissibility claim discipline

A claim that a physical execution is safe, admissible, stable, reliable, or deployment-ready is a material claim and must name its external evidence surface. Governance permission, compliance status, documentation, certification, monitoring, and authority are not sufficient evidence surfaces by themselves. They may support procedural legitimacy, but they do not compute physics or establish biological, mechanical, medical, infrastructural, or cyber-physical safety.

Determinism boundary. Deterministic verification is valuable where the specification, state space, and transition relation support it, but it is neither universally available nor sufficient by itself. Safety-critical justification may legitimately combine formal proof, probabilistic risk analysis, empirical testing, simulation, safety cases, redundancy, bounded operating envelopes, monitoring, fail-safe behavior, and accountable residual-risk authorization. The required warrant strength is proportionate to consequence, uncertainty, reversibility, and domain practice. A deterministic proof of the wrong specification does not establish safety; a probabilistic warrant is not invalid merely because uncertainty remains.

Computational-source discipline. A material execution claim SHOULD distinguish the source that generated the candidate action from the source that warrants the action. If candidate generation and admissibility warrant are supplied by the same system, the claim record SHOULD state how circular self-confirmation is prevented and what independent check, validity-domain limit, or refusal condition applies.

Formal-conformance layer rule. A proof that an implementation satisfies a formal specification supports only the formal-conformance claim established by that proof. A physical or causal claim additionally requires a justified specification-to-reality mapping, declared validity domain, boundary conditions, uncertainty treatment, and domain evidence. Failure at the correspondence layer reopens every downstream claim that relied on the formal model as a physical warrant. **Institutionalized-uncertainty boundary.** A policy, risk acceptance, certification, or approval process that assigns responsibility without establishing the relevant operating envelope may govern who may test, monitor, stop, or review a system, but it MUST NOT be presented as reducing the underlying epistemic uncertainty unless new evidence actually does so.

If a physical-admissibility claim changes its model, validity domain, threshold, controller envelope, failure-mode account, or monitoring/shutoff assumptions, affected downstream gates must be rerun. Silent tuning may improve wording, but it does not preserve the original claim.

Capability Claim Integrity

When a material decision or public statement depends on a capability term such as agentic, autonomous, reasoning, understanding, self-improving, or intelligent, the record MUST state an operational definition and observable criteria before the term receives gate, authority, deployment, conformance, or communication

weight. The strongest permitted claim is limited by the declared evidence surface, validity conditions, and evidence maturity stage.

A coordinated system MAY possess a bounded system-level operational capability even when the capability is implemented through ordinary components such as probabilistic inference, memory, planning, deterministic constraints, tool interfaces, environmental feedback, and human controls. Decomposing the system into those components does not by itself refute the system-level capability. At the same time, fluency, anthropomorphic language, market labels, task success, or user impression does not establish consciousness, personhood, intrinsic goals, moral agency, unrestricted autonomy, general intelligence, safety, lawful mandate, or ethical authority.

The record MUST test material alternative explanations and state a falsification or revision trigger. If the evidence cannot distinguish the claimed capability from a materially simpler explanation, the claim must be narrowed, marked assumption-bound or contested, restricted to exploratory use, or refused at the stronger level.

Instrumental tool-role claim. A statement that a system occupies a tool-role is an operational or institutional claim about deployment, objectives, permissions, responsibility, and control. It does not by itself establish low intelligence, absence of agency, absence of possible sentience, safety, moral status, or legitimate authority. The record should name the deploying actor, assigned ends, permission structure, responsibility chain, and control boundary.

Functional-understanding claim. When “understanding” is material, the record MUST identify the exact configuration, target, relation claimed, domain, evaluation conditions, permitted scaffolding, distance from memorized exemplars, transfer conditions, alternative explanations, and failure boundary. Fluent explanation, retrieval, benchmark success, or correct answers alone cannot support a domain-general or phenomenal-understanding claim.

Claim-type examples

Claim type	Example	Minimum discipline
Empirical	“Injury rates decreased after the intervention.”	Requires evidence contact, baseline, uncertainty, and alternative-explanation review.
Causal	“The policy caused the decrease.”	Requires the claim-specific causal identification plan and its supported claim boundary. A positive pre/post change may coexist with a harmful incremental effect; mechanism language or a comparator recommendation alone does not identify causation. Without adequate identification, restrict the conclusion to supported descriptive, predictive or procedural claims.
Formal	“CVaR was computed by the declared algorithm.”	Requires definitions, assumptions, formula, inputs, and reproducible calculation surface.
Operational	“The system is operationally autonomous within the declared pilot envelope.”	Requires an operational definition, observable criteria, system composition, evidence trace, alternatives, validity conditions, falsifier, explicit nonclaims, controls, and

Claim type	Example	Minimum discipline
		authority boundary.
Legal-regulatory	"This agency may lawfully authorize the trial."	Requires current authority, jurisdiction, applicability analysis, interpretation status, and challenge or review path.
Normative	"This rights floor is adopted as non-compensatory."	Requires declared ethical commitment, consistency check, rights compatibility, and challenge pathway.
Metaphysical horizon	"AIU language is used as a humility horizon."	May guide meaning and restraint, but creates no Tier 1-3 evidence or gate authority.
Operational / institutional	"The AI system occupies an instrumental tool-role in this workflow."	Requires the deploying actor, assigned ends, objective and permission sources, responsibility and control arrangement, and explicit nonclaims concerning capability, sentience, safety, moral status, and authority.
Functional capability	"The configured system functionally understands target T in domain D for relation R."	Requires configuration identity, target and relation, evaluation conditions, scaffolding, novel or transfer tests, alternatives, failure boundary, and explicit nonclaim of phenomenal understanding.

Admissible Alternative Construction and Robustness Adjudication

An alternative construction blocks unique selection only when it is evidence-supported, selection-material, independently motivated rather than invented after result inspection, and accepted through a declared adjudication process.

The record SHALL identify the alternative, evidence basis, domain validity, independence from the interested option owner, expected affected calculations, selection effect, reviewer, disposition, and reopen condition. Rejected alternatives SHALL remain visible with rationale.

The Canonical Robustness Module Registry is finite for a release. New domain-mandated modules may be added only through a versioned rule or declared domain standard; an analyst cannot create an unlimited veto set after seeing results. This selection rule does not itself define a theory-level falsification family: confirmatory operator attribution uses the separately frozen finite or formally bounded family and stopping rule. Unexamined alternatives leave attribution unresolved; they do not annul a failed locked prescription or justify an unqualified performance claim.

A module marked NOT_TRIGGERED_WITH_RATIONALE requires the trigger examined, evidence basis, reviewer status, and reopen condition. REQUIRED_NOT_EVALUATED is distinct from not triggered and blocks unique framework selection.

Alternative status	Selection effect
ADMISSIBLE_REQUIRED	Evidence-supported, material, independently motivated; must be evaluated

Alternative status	Selection effect
ADMISSIBLE_RESOLVED	Evaluated and adjudicated; retain trace
REJECTED_WITH_RATIONALE	Not sufficiently supported/material/independent; retain log
POSTHOC_OUTCOME_DRIVEN	Inadmissible for current unique-selection claim
REQUIRED_NOT_EVALUATED	Robustness incomplete; no unique selection

APPENDIX RELEASE: Identity, Source Authority and Revision Record

Framework release: MathGov/RippleLogic v13.0. Component: Methodological Falsifiability and Dependency 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
Sentence 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	Methodological Falsifiability and Dependency 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

Release control	Historical value
Role	Claim typing, alternative-model admissibility, falsification, dependency, re-derivation, and methodological refusal discipline.
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 Methodological Falsifiability and Dependency 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.