Physical/Causal Admissibility Evidence Profile

Integrity surface	Current requirement	Claim boundary
External warrant	Consequence-bearing actions identify a domain-appropriate engineering, clinical, scientific, regulatory, or formal warrant within its validity domain.	MathGov does not generate the missing domain proof.
Source separation	Candidate generation, orchestration, guardrails, monitoring, and authorization are distinct from physical-admissibility evidence.	Generated or approved is not physically safe.
Claim action	Residual unknowns can force narrowing, controls, testing, redesign, escalation, delay, or refusal.	A profile is not deployment certification.

Configuration-Bound Physical/Causal Evidence Extension

A PC-AEP claim is valid only for the declared configuration and operating envelope. The minimum high-stakes record now includes hardware and software versions, model or policy version, parameters, sensors, actuators, tools, permissions, data sources, environment, active controls, authority state, validation version, timestamp, and configuration identifier or hash where feasible.

PC-AEP separately records current-state validity, transition admissibility, and resulting-state viability. Process success is partial evidence: it does not by itself establish the complete capability envelope, causal explanation, safety boundary, degradation boundary, or performance under a materially different configuration.

Compilation, synthesis, simulation, or hardware implementation establishes realizability of the represented architecture at that milestone only. Domain validity and physical safety require independent measurement, boundary testing, failure-mode analysis, and configuration-matched evidence.

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

Purpose

The Physical/Causal Admissibility Evidence Profile, abbreviated PC-AEP, is a compact evidence profile inside Reality Grounding and CSV. It prevents a RippleLogic run from treating a computed, generated, simulated, approved, certified, monitored, or compliance-accepted action as physically or causally admissible by default.

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

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

It also does not make MathGov a physics engine, control-system verifier, medical protocol, legal certification, deployment certification, or formal proof system. It asks whether the action has enough declared physical or causal warrant for the claim being made.

Formal-conformance boundary. Formal verification can establish that an implementation satisfies a declared specification under stated assumptions. It does not by itself establish that the specification accurately represents the physical or causal regime. When formal verification supports a physical claim, the profile MUST also record the specification, assumptions, validity domain, boundary conditions, and the separate evidence that connects the specification to the relevant reality surface.

Determinism boundary. PC-AEP does not require every admissibility warrant to be deterministic. The warrant must be strong enough for the consequence class and may combine formal verification, probabilistic risk

analysis, empirical testing, validated simulation, standards-based safety cases, redundancy, bounded operating envelopes, monitoring, fail-safe behavior, and qualified residual-risk authorization. Deterministic conformance to an inadequate specification is insufficient; probabilistic evidence is admissible when uncertainty, tails, dependencies, limits, and refusal conditions are explicit.

Trigger

PC-AEP is REQUIRED for Tier 3 runs and high-stakes Tier 2 runs when a candidate action materially affects bodily safety, medical treatment, robotics, vehicles, infrastructure, industrial systems, energy systems, environmental intervention, weapons or security systems, autonomous execution, cyber-physical operations, irreversible resource commitment, or any causal pathway where execution can create non-trivial physical, biological, ecological, or operational harm.

It is RECOMMENDED for ordinary Tier 2 runs when physical or causal uncertainty could change RF/NCRC, TRC, CSV, authority, or public claim strength.

Required fields

Field	Required content
candidate_generation_source	The model, optimizer, planner, simulator, controller, human procedure, dataset, or orchestration layer that produced the candidate state or action.
physical_or_causal_model_used	The model, theory, mechanism, safety case, expert method, or domain analysis relied on. For a material causal or comparative-effect claim, include the target estimand, intervention and comparator, target population and time, identification strategy and assumptions, selection/confounding and interference treatment, outcome measurement, uncertainty, and a reasoned identification-status verdict.
validity_domain	The conditions, scale, population, environment, load, context, time horizon, and operating range in which the model or warrant is claimed to hold.
boundary_conditions	Initial, environmental, operational, legal, resource, dependency, and interface conditions that bound the claim.
uncertainty_range	Quantitative or qualitative uncertainty, confidence limits, error range, sensitivity range, or evidence-status limitation.
failure_modes	Known plausible ways the candidate can fail, degrade, harm, drift, destabilize, or become invalid.
reversibility_or_irreversibility_boundary	What can be undone, stopped, repaired, rolled back, compensated, or not restored after execution.
verification_simulation_empirical_test_or_expert_warrant	Formal verification, simulation, empirical test, field evidence, expert review, standards-based warrant, or explicit statement that no adequate warrant exists. When formal verification supports a physical claim, record the verified specification, assumptions, and separate specification-to-reality evidence. A causal-effect claim also requires the claim-specific identification plan in

Field	Required content
	physical_or_causal_model_used; an observed pre/post improvement alone does not identify an incremental benefit.
admissibility_warrant_source	The qualified domain method, verification process, test regime, standards-based safety case, or evidence source that evaluates admissibility. Institutional permission or execution authority is recorded separately and cannot serve as the warrant by itself.
monitoring_and_shutoff_path	How the action will be observed, interrupted, rolled back, sandboxed, contained, or stopped if assumptions fail.
residual_unknowns	Material unknowns that remain after the evidence review.
required_claim_action	One of proceed, narrow, control, redesign, escalate, or refuse.

Status values

- PCAE_SUPPORTED: evidence is sufficient for the declared claim boundary.
- PCAE_ASSUMPTION_BOUND: action depends on assumptions that must be visible and rechecked.
- PCAE_PARTIAL: evidence supports only a narrower or controlled claim.
- PCAE_CONTESTED: relevant data, experts, stakeholders, or reviewers materially dispute the admissibility account.
- PCAE_UNKNOWN: evidence is insufficient to support the requested claim.
- PCAE_VERIFICATION_REQUIRED: stronger formal, empirical, simulation, or expert warrant is required before the stronger claim can proceed.
- PCAE_CONTROL_REQUIRED: binding controls, monitoring, shutoff, rollback, or containment are required for selectability.
- PCAE_REDESIGN_REQUIRED: the candidate is not selectable as specified but may be reformulated.
- PCAE_REFUSE_OR_BLOCK: the action or claim must be refused or blocked under the declared evidence.

Claim actions

Proceed means the profile supports the declared claim boundary and all other cascade requirements must still pass. Narrow means reduce the claim, scope, action, environment, population, duration, or authority. Control means add binding safeguards as part of the option. Redesign means the candidate is not selectable as specified. Escalate means obtain domain review or stronger evidence. Refuse means the requested claim or execution is not admissible.

Relationship to Reality Grounding

Reality Grounding records the reality surface, evidence trace, unknowns, transition boundary, consequence pathways, and claim boundary. PC-AEP is linked to Reality Grounding when the claimed action depends on physical or causal constraints. It strengthens claim discipline by forcing the run to state the candidate-generation source, model, domain, limits, uncertainty, failure modes, reversibility, warrant, monitoring/shutoff path, residual unknowns, and claim action.

Computational-source boundary

PC-AEP separates two questions that are often collapsed: what generated the candidate action, and what establishes that the candidate is physically or causally admissible. The candidate-generation source may be a model, planner, optimizer, simulator, controller, human procedure, agent workflow, or orchestration layer. The admissibility warrant must be a domain-appropriate evidence surface. A candidate generator may propose an action, but it does not become the proof of admissibility merely because it generated the action. If the candidate-generation source and the admissibility-warrant source are the same system, the run MUST disclose why that is not circular, what independent checks or validity-domain limits apply, and what refusal condition prevents self-confirming execution.

Relationship to CSV

CSV consumes PC-AEP when structural viability depends on physical or causal adequacy, dependency closure, resource closure, operational capacity, reversibility, monitoring adequacy, containment, authority, or host-system integrity.

A weak PC-AEP does not automatically fail every option, but it prevents strong selectability, conformance, deployment, safety, reliability, or alignment claims until the weakness is resolved, controlled, narrowed, escalated, redesigned, or refused.

Relationship to TRC

If PC-AEP reveals a catastrophic, irreversible, lock-in, ruin-path, or severe-harm scenario not represented in TRC, the run MUST reopen TRC before RLS. CSV may discover TRC-relevant material, but it does not absorb or replace TRC.

Plain-language rule

Before MathGov allows action in the physical or causal world, it asks what model of reality is being used, where that model stops, what could fail, whether harm is reversible, what evidence or expert warrant supports the action, how the action will be monitored or stopped, what remains unknown, and whether the right answer is proceed, narrow, control, redesign, escalate, or refuse.

Methodological integrity linkage

PC-AEP evidence MUST state the claim type, model dependency, validity-domain limit, alternative-explanation status, falsification or revision trigger, and re-derivation scope where material. If a causal effect is not identified, narrow to the descriptive, predictive or procedural claim actually supported. If physical or causal evidence fails, a run MUST NOT preserve the same safety, reliability, deployment, or admissibility claim through silent model tuning. It must preserve the failed original test, version the change and rerun the affected chain under the MFDI frozen evaluation and stopping rules.

State-to-Impact and Control-Effectiveness Evidence (Normative)

A telemetry-to-impact bridge specifies the observed quantity, units, measurement method, time and configuration, bearer or system, state construct, baseline and horizon, causal or descriptive warrant, uncertainty and permitted normalized interpretation. Raw infrastructure load, algorithmic divergence or uptime must not silently become welfare or a sentience/meaning claim.

A control-effectiveness warrant names the real intervention, hazard endpoint, resources and limits, latency distribution or credible bounds, causal effectiveness, failure modes, depletion and residual harms. Credit timely prevention only when the supported safe-state path precedes the relevant unacceptable endpoint. Preserve separately supported ex-ante protection and later mitigation; do not treat a missed deadline as proof that every control benefit is zero.

No mathematical attenuation formula, signature, proof circuit, database update, model confidence or control boolean establishes physical prevention. Recovery and restoration require qualified actions and independently observed outcomes. An experimental state-space model or matrix exponential needs unit, domain, stability, calibration and transfer evidence before it supports consequence-bearing claims.

Physical execution boundary

For consequence-bearing physical systems, MathGov distinguishes governance permission from physical admissibility. ALLOW, SELECTABLE, CSV_PASS, certification, conformity, documentation, monitoring, or authority approval MUST NOT be read as physical-safety proof. They mean only that the governance record supports the declared claim after the required evidence has been reviewed. **Material-unknown action rule.** If the validity domain, safe operating envelope, or failure boundary is materially unknown, a governance body may authorize only explicitly bounded testing, sandboxing, containment, monitoring, or evidence collection whose own risk posture is qualified. It MUST NOT convert the unknown into a safe or ordinarily executable state, and operator acceptance cannot waive third-party rights or ruin constraints.

If a candidate action affects a physical system, PC-AEP must state where admissibility comes from: formal verification, validated simulation, empirical testing, standards-based safety case, certified controller envelope, qualified engineering or clinical warrant, regulator-recognized method, or another objective domain method. If that source is absent, contested, outside validity domain, or insufficient, the required claim action cannot be ordinary proceed. It must be narrow, control, redesign, escalate, or refuse.

Physical execution claim ladder

Claim posture	Meaning	Required treatment
PHYSICAL_ADMISSIBILITY_SUPPORTED	External domain evidence supports the declared action inside the stated validity domain.	May proceed only if RG, RF/NCRC, TRC, CSV, authority, and audit requirements also pass.
GOVERNANCE_PERMISSION_ONLY	Authority or procedure permits the action, but physical admissibility has not been independently established.	Do not claim physical safety; narrow, control, escalate, or refuse.
PHYSICAL_ADMISSIBILITY_NOT_ESTABLISHED	The evidence source is missing, weak, out of domain, or unable to answer the execution question.	Block strong execution claims; redesign, escalate, or refuse.
PHYSICAL_ADMISSIBILITY_CONTRAINDICATED	Evidence indicates instability, unsafe operation, irreversible harm, or unacceptable failure mode.	Refuse or block; reopen TRC if catastrophic or irreversible risk is material.

Plain-language boundary: governance decides whether a decision is permitted; domain evidence determines whether the physical action is admissible within a declared validity domain. MathGov can require and audit that evidence. It does not manufacture the evidence.

Placement in the two-phase method

PC-AEP supports Phase 1 qualification. It can strengthen, narrow, control, redesign, escalate, or refuse a physical-execution claim before RLS ranking. It does not add points to RLS and does not allow a high RLS score to compensate for missing physical or causal admissibility evidence.

Domain examples of admissibility warrant

Domain	Example admissibility warrant	Boundary
Robotics or autonomous mobility	Certified controller envelope, simulation under declared conditions, field test evidence, and independent safety review.	Only supports the declared environment, payload, weather, sensor, and control limits.
Medical or health action	Clinical protocol, licensed clinician review, trial evidence, or approved standard of care.	Does not generalize beyond the population, indication, contraindications, and care setting reviewed.
Infrastructure or engineering	Engineering safety case, code compliance, load/traffic model, inspection record, and qualified engineer signoff.	Does not prove safety outside the design envelope or after material condition changes.
Cyber-physical or security system	Threat model, red-team results, rollback plan, access-control review, and monitored deployment envelope.	Does not authorize unbounded persistence, privilege escalation, or uncontrolled downstream action.
Ecological or environmental intervention	Environmental impact assessment, domain expert panel, monitoring plan, and reversibility/mitigation assessment.	Does not support irreversible action when residual unknowns remain gate-material.
Military, weapons, or command context	Legal review, rules-of-engagement authority, human-command accountability, system test evidence, fail-safe review, and escalation controls.	Does not delegate life-taking or nuclear launch authority to an algorithmic process.

Effect, Probability, Temporal, and Option-Closure Binding

PC-AEP v2.6 binds each material physical or causal consequence to the stable EffectTokenID used by WDBIP and RLS, without replacing domain engineering, clinical, scientific, safety-case, legal, or regulatory warrants.

The profile SHALL state whether probability represents conditional event occurrence, expected frequency/exposure, an already-expected effect, or a scenario-conditional quantity. The same occurrence probability MUST NOT be applied twice in the residual welfare construction.

The temporal profile SHALL distinguish onset, contiguous duration, recurrence, cumulative exposure, terminal state, recovery and irreversibility. A model-generated duration label is not evidence of temporal behavior.

Where the profile supports a materially safer redesign, staged test, reversible trial, delay, or control package, the option-set record SHALL include or explicitly dispose of it before a strong selection claim.

APPENDIX RELEASE: Identity, Source Authority and Revision Record

Framework release: MathGov/RippleLogic v13.0. Component: Physical/Causal Admissibility Evidence Profile 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	Physical/Causal Admissibility Evidence Profile 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	Domain-warrant profile for consequence-bearing physical,

Release control	Historical value
	biological, ecological, medical, engineering, cyber-physical, infrastructural, and causal actions.
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 Physical/Causal Admissibility Evidence Profile 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.