

# MathGov 3R / 1-2 Two-Phase Public Introduction v12.6

## v12.6 Current Release Integration

**Release: MathGov Core Release 2026.09 v12.6 / SGP v8.5 - Independent Open-Source Qualification-Continuity Release**

**Exact current companion pins:** Canon v12.6; SGP v8.5; ripple.md v5.5; Agent System v12.5; CSV v2.4; Cascade v2.6; Reproducibility v1.4; WDBIP v1.6; RLS Validation v2.6; Primer v4.4; Public Introduction v12.6; PC-AEP/MFDI/Source-Coupling v2.3; Aligners Sheet v5.6.

| Integrity surface  | Current requirement                                                                    | Claim boundary                                                         |
|--------------------|----------------------------------------------------------------------------------------|------------------------------------------------------------------------|
| 3R                 | Reality. Rights. Ripples.                                                              | A public doorway, not the complete technical specification.            |
| 1-2                | Qualify through RG, RF/NCRC, TRC, and CSV; then rank selectable options through RLS.   | Ranking cannot rescue prior failure.                                   |
| Eligible shorthand | Eligible(a) is teaching shorthand for survival of the declared qualification sequence. | It is not a sixth gate, new authority predicate, or execution licence. |

## A simple new rule: bind claims to the real configuration

Do not say only, “the system is safe” or “the system can do this.” Ask:

- Which version, tools, permissions, controls, environment, and authority were tested?
- Is the action merely possible, or is it also qualified and authorized?
- Is the current state valid, is the transition admissible, and is the resulting state viable?
- What change would make the earlier approval expire?

MathGov therefore keeps alignment alive through requalification. When the material configuration or evidence changes, the affected checks run again.

## v12.6 Release Note

v12.6 adds configuration-bound assurance and qualification-continuity teaching language while preserving the 3R and MathGov 1-2 structure. The 3R and MathGov 1-2 teaching structure is unchanged.

This is the public doorway, not the full specification. Use it to understand the rhythm; use the Canon for conformance. Informative, non-normative public teaching companion. Canon controls if any conflict occurs.

## Use the smallest profile that is safe

- **Quick / Tier 1:** a short record for low-stakes, reversible choices.
- **Standard / Tier 2:** evidence, parameters, gates, controls, and uncertainty for material organizational or community choices.
- **Audit / Tier 3:** a full reproducible record, hashes, independent review, and replay for high-stakes or public decisions.

The framework stays the same at every level. Only the documentation burden changes. A user must escalate when rights, ruin, vulnerable groups, physical execution, consequential AI authority, or public validation claims become material. See

[docs/implementation/MATHGOV\\_REPRODUCIBILITY\\_AND\\_USE\\_STANDARD\\_v1.4.md](#).

## The shortest version

**Reality. Rights. Ripples.**

**1. Qualify.** Reality. Rights. Ruin. Viability.

**2. Rank.** Ripple-score the viable options.

That is the operating rhythm.

First make sure the option belongs in the ring. Then find the one that produces the best ripple.

## The full formal cascade

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

Reality Grounding -> Rights Floor -> Tail-Risk Constraint -> Containment and Structural Viability -> RippleLogic Scoring

Teaching note: this doorway compresses the architecture into three pillars and two phases. The formal architecture has five levels: one grounding precondition, three qualification gates, and one ranking layer.

## Why this matters

MathGov does not begin by scoring every imagined option. It first asks which options have enough reality, rights protection, ruin control, and viability to enter the decision field. Only then does RippleLogic score the survivors.

The MathGov answer is: ground the claim, protect the right, bound the ruin, verify the system, make exit and correction real, and rank only what survives qualification.

## The three pillars

| Pillar  | Plain meaning                                                                                                    | Formal location                                                                                                                                   |
|---------|------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------|
| Reality | Do we have contact with what is real, possible, grounded, testable, bounded, auditable, correctable, and viable? | RG opens reality contact. CSV tests operational reality and structural viability. PC-AEP is used when physical or causal execution claims matter. |
| Rights  | Are protected beings, dignity, consent, thresholds, and non-collapse safeguards protected?                       | RF/NCRC protects rights directly. TRC protects against ruin pathways that can destroy the conditions for rights and flourishing.                  |
| Ripples | Among the options that survive qualification, which creates the best net flourishing?                            | RLS ranks only the survivors across affected unions and welfare dimensions.                                                                       |

## v12.2 rights and ruin hardening

**Rights hardening:** RF/NCRC does not discount a categorical or severe rights violation merely because it is brief or evidence confidence is low. Rights severity, categorical prohibitions, and severe-hazard exposure are checked separately from ordinary welfare timing.

**Tail Emergency Mode:** if every rights-safe option exceeds the ruin corridor, the default is redesign, delay, escalation, no-action review, or refusal. A temporary least-CVaR path is allowed only under a governed emergency declaration, absolute exposure cap, independent challenge, time bounds, controls, and exit criteria. It is not ordinary selection, and RLS is not used to rescue it.

## Phase 1: Qualify the options

Phase 1 asks: which options are even eligible to be ranked?

| Gate                                       | Question                                                                                                                                                                                         | If the answer fails                                                                                  |
|--------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------|
| RG - Reality Grounding                     | Is the option grounded in evidence, real constraints, affected unions, and a clear claim boundary? Can high-consequence claims be tested or warranted, bounded, audited, corrected, and refused? | Do not score fantasy. Narrow, investigate, mark exploratory, escalate, or refuse the stronger claim. |
| RF - Rights Floor                          | Does the option protect rights, dignity, consent, and non-negotiable safeguards?                                                                                                                 | Do not average rights away. Redesign, escalate, or refuse.                                           |
| TRC - Tail-Risk Constraint                 | Is catastrophic, irreversible, fat-tail, or system-breaking risk bounded?                                                                                                                        | Do not gamble ruin for ordinary benefit. Bound, redesign, escalate, or refuse.                       |
| CSV - Containment and Structural Viability | Can the option actually stand, execute, be controlled, monitored, exited, corrected, and remain governable?                                                                                      | Do not select what cannot stand. Add controls, narrow, redesign, or refuse.                          |

TRC is the ruin veto. CSV is the viability and control test. They are linked, but they should not be collapsed.

### Reality Grounding hardening

For life-critical systems, weapons, nuclear command and control, planetary-scale cognition, critical infrastructure, medical systems, ecological interventions, or other high-consequence domains, no claim should govern action unless it can be tested or otherwise warranted, bounded to a validity domain, audited with traceable evidence, corrected through monitoring and revision, and refused when the evidence surface is insufficient.

In public language: ground the claim; protect the right; bound the ruin; verify the system; make exit and correction real; rank only what survives qualification.

## Phase 2: Rank the survivors

Phase 2 asks: among the options that survived qualification, which one creates the best ripple?

RLS evaluates affected unions, welfare dimensions, evidence, uncertainty, robustness, and decision state.

### Mathematical spine

Wrong method:

Decision =  $\operatorname{argmax} \text{RLS}(a)$  over all imagined options

Correct method:

Eligible(a) = RG\_qualified(a) AND RF/NCRC(a) AND TRC\_qualified(a) AND CSV\_qualified(a)

Here RG\_qualified means RG\_SUPPORTED or RG\_NARROWED; TRC\_qualified means TRC\_PASS or a documented TRC\_NOT\_TRIGGERED assessment. Eligible remains teaching shorthand, not a machine-state token, sixth gate, or execution licence. CSV\_qualified means CSV\_PASS, CSV\_PASS\_WITH\_CONTROLS, or a documented CSV\_NOT\_MATERIAL result valid for the declared claim boundary; it does not include CSV\_REDESIGN\_REQUIRED, CSV\_ESCALATE, CSV\_FAIL, or CSV\_EMERGENCY\_PROVISIONAL.

Define the selectable set:

$A\_sel = \{a : RG\_qualified(a) \text{ AND } RF/NCRC(a) \text{ AND } TRC\_qualified(a) \text{ AND } CSV\_qualified(a)\}$

Then apply the decision-state rule:

- If  $A\_sel$  is empty, redesign, delay, escalate, use the governed emergency pathway where applicable, or refuse the stronger claim.
- If RLS is defined and produces one unique decisive leader in  $A\_sel$ , the framework may return **ALLOW\_FRAMEWORK\_SELECTION**.
- If RLS is undefined, tied, close, uncertainty-overlapped, or non-decisive, return **REFUSE\_DETERMINISTIC\_SELECTION**. Seek more evidence, preserve multiple options, redesign, or use a separately disclosed **AuthoritySelectionRecord** within a lawful mandate.

A point-score argmax is therefore a comparison aid inside the selectable set. It is not automatic authorization, and it does not erase uncertainty, ties, refusal, or the separation between framework selection and authority selection.

## No RLS rescue

A high RippleLogic score cannot rescue failed reality grounding, a rights violation, an unacceptable ruin path, an unsupported physical-execution claim, or a structurally non-viable option.

## Decision outcomes

Public labels such as SELECT are explanatory shorthand only. Canonical emitted states remain **ALLOW\_FRAMEWORK\_SELECTION**, **REFUSE\_DETERMINISTIC\_SELECTION**, **BLOCK/FAIL** states, **ESCALATE**, and separately recorded **AuthoritySelectionRecord** outcomes; no public label creates execution authority.

| Outcome              | Meaning                                                                                                                                                  |
|----------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------|
| SELECT               | Public teaching shorthand: the option is eligible and clearly best among survivors. It does not replace canonical machine-state or authorization tokens. |
| SELECT WITH CONTROLS | Public teaching shorthand: the option is eligible only with recorded safeguards, monitoring, or bounds. It is not execution approval.                    |
| NARROW               | The claim, scope, domain, or time horizon must be reduced.                                                                                               |
| REDESIGN             | The option fails but may be improved.                                                                                                                    |
| ESCALATE             | Authority, evidence, expertise, or stakeholder participation is insufficient.                                                                            |
| REFUSE               | The requested claim, selection, or action is not supportable under the current evidence, authority, or gate state. Record the specific reason.           |
| NON-DECISIVE         | The evidence or scores do not justify deterministic selection.                                                                                           |

## Public script

MathGov is a decision architecture, not a claim that every option should be scored. It first qualifies options through reality, rights, ruin, and viability. Then it ranks only the surviving options by their ripple effects across affected unions. In short: qualify before you rank.

## The 60-second checklist

### 1. Qualify

- Is it real?
- Are rights protected?
- Is ruin bounded?
- Can it stand, be monitored, be corrected, and be exited?

### 2. Rank

- Who and what is affected?
- Which option creates the best ripple among the survivors?
- Is the result robust enough to select, or should it be non-decisive?

## Claim boundary

MathGov Core v12.6 is a proposed Tier 1-3 auditable decision architecture. It is not empirical validation, legal certification, deployment certification, Tier 4 ProofPack, automated moral truth, or physical safety proof. For consequence-bearing physical systems, MathGov requires external domain evidence through PC-AEP; it does not manufacture that evidence by itself.

## Final pocket form

Reality. Rights. Ripples.

Qualify before you rank.

Rank only what survives.

## SGP v8.5: protection, participation, and power are separate

When a decision may affect a welfare-bearing being, the current SGP interface does not ask one number to answer every question.

- **MPS** asks how strongly evidence supports welfare-bearing experience.
- **FPP** records full non-downgradable rights-of-protection.
- **GPR** asks how the entity can participate directly, with support, or through representation.
- **SPR** asks whether it may hold one specifically named consequential role under strict gates.
- **ICP** describes intelligence and understanding, but cannot establish sentience or authority by itself.
- **RMCP** profiles integrated reality-management capacity; **P100** is its open top plateau. Every human has FPP-100, while humanity collectively is the provisional current RMCP-P100 reference anchor. Other biological, digital, hybrid, collective, or extraterrestrial intelligences may reach or already occupy P100 under admissible evidence.

RMCP/P100 is not a moral-worth or authority score. The operating rule is: **unknown is not zero; capability is not welfare; protection is not authority; full protection does not require human likeness**. MPS intervals may inform residual welfare sensitivity only after rights, ruin, and viability qualification.

**SGP interface clarification.** MPS is not a cardinal welfare multiplier or moral-worth score. It determines protection posture and which welfare-inclusion hypotheses must be tested. Missing evidence is MPS-NE, not zero.