

# MHIOS synchronization and implementation audit

## MathGov Human Interface and Orchestration Standard · v2.1

### 1. Executive determination

The supplied MHIOS v2.0 package has a valuable organizing model: stakeholder discovery, better-option search, qualification before ranking, visible unknowns, distinct authority and execution, outcome review, progressive disclosure and contestability. Those elements were retained. The material work was to bind that model to the current Core and make the accompanying local data tools less permissive and more honestly scoped.

The updated component is **MHIOS v2.1**, bound to RippleLogic Canon **13.0**, SGP **8.8** and the exact fifteen Core artifacts of **MG-RL-13.0-20260926-RELEASE-I**. All fifteen separately attached Core files match the corresponding Release-I archive entries by SHA-256. None was edited. “MIOS” is retained only as the user’s shorthand, not introduced as a competing formal name or a sixteenth Core component.

This is a publication and scoped-testing companion package, not an independently validated human interface, production runtime, full Canon validator, authorization service or empirical validation of MathGov. An implementation still has to demonstrate its own controls, accessibility, authority, domain warrants, human outcomes and applicable conformance.

### 2. Sources and method

The audit used the complete v2.0 standard, both schemas, example, YAML template, fourteen-sheet preparation workbook, tests, vector list, release metadata and original package structure. The v2.0 archive remains untouched in the complete delivery’s private provenance directory. The current Core sources are retained unchanged under references/Core\_15; the exact comparison is recorded in CORE\_PIN\_MANIFEST.json.

The review compared source ownership, status names, claim limits, fields, reference semantics, uncertainty treatment, SGP output presentation, runtime transitions and preparation-workbook behavior. It then tested the implemented local checks with deliberately malformed and contradictory records, inspected the Word container and bookmarks, rendered the updated document, compared complete Markdown and Word body/table text, and checked the final packaged files after extraction.

The current Core is the semantic dependency, not an invitation to revise the frozen framework. A new companion version does not change the current Canon/SGP editions or retroactively validate an implementation. The actual executed results and unexecuted surfaces are separately reported in VERIFICATION\_REPORT.md and machine-readable receipts.

### 3. Independently reproduced baseline defects

#### A01. The current-facing integration line was obsolete

The v2.0 example declares MathGov\_Core\_2026\_09\_v12.7\_SGP\_v8.6+2026.08.26.1. Its pin table, schemas, templates, workbook and running document furniture also belong to the previous line. A version substitution alone would not address the changed source and state requirements.

**Repair locations:** main standard §22 and Appendix RELEASE; workbook VERSION\_PINS; schema constants; examples; references/CORE\_PIN\_MANIFEST.json. The pin manifest now identifies each of the fifteen exact source files rather than relying on a “latest” label.

**A02. Shape validation accepted a contradictory decision**

Changing the old example to RF\_FAIL, CSV\_PASS, ALLOW\_FRAMEWORK\_SELECTION and AUTHORIZED\_WITHIN\_SCOPE still passed its original JSON Schema. That demonstrates a limitation of the old schema, not a contradiction in the Canon itself. The original tests did not run all sixteen listed interface vectors through an integrated session validator.

**Repair locations:** schemas/, reference/validate.py, tests/test\_mhios.py; standard §17.2 and Appendix E. Schema validation precedes semantic checks. Current enumerations, nonblank identities, explicit nulls, timestamp formats, duplicate-key rejection and nonfinite-number rejection are enforced. A valid object shape does not certify the facts in the object.

**A03. A single session gate summary could hide option differences**

The old schema combined multiple candidates with one set of RG/RF/TRC/CSV status strings. Differently qualified options could therefore share an undifferentiated summary.

**Repair locations:** standard §1.2 and §17.2; options[], states, per-option source references and the semantic checker. Rights failures, unknowns, TRC failure and unqualified CSV results cannot enter ordinary scored comparisons. Post-failure forensic calculations may remain in owner records, but are not represented as ordinary progression in this compact adapter.

**A04. Untouched preparation rows looked assessed**

The former worksheet formula reduced a blank row to NO. It did not distinguish an untouched record from a reviewed failure. Its permissive data entry could also admit unknown status strings without a clear invalid-state label.

**Repair locations:** preparation workbook QUALIFICATION, F5:F24 and status input validation. Twenty local formulas now return UNASSESSED, INCOMPLETE, INVALID\_STATUS, NOT\_QUALIFIED or TOKENS\_QUALIFY. The last means only that declared status strings and reference presence satisfy this local check. It is not an RF, TRC or CSV evidence review, a framework winner or execution permission.

**A05. Rendering exposed stale and ambiguous running identity**

The original header still said “MHIOS v2.0 ... Final Complete 2026-09-04”, and the footer appended its page number immediately after the old edition. New longer status descriptions also exposed split table rows across pages.

**Repair locations:** Word running header/footer and table-row pagination. The header/footer now identify the document without obsolete edition claims. Table column widths, shading, borders and existing style definitions are retained. Keep-row-together flags prevent fragmentary row continuations. Edition and lineage detail stay in Appendix RELEASE.

**4. Current-Core improvements integrated****B01. Preserve each state namespace**

**Locations:** §1.2, §13, §17.2, Appendix E; interface-state schema and namespace checker.

FrameworkVerdict, decision\_state, execution state, Agent capability state, SGP evidence posture and MHIOS-local display state are distinct. Canon AF.2 and Reproducibility §3 control the crosswalk. A separate authority choice does not convert framework refusal into decisive selection. Controls alone do not select a provisional option. An ordinary provisional choice has its own referenced view. PASS styling is rejected for a nonpass token in the supplied interface record checker.

**B02. Preserve normalized scoring without adding another calculator**

**Locations:** §11.1–§11.4 and Appendix E.3.

The existing normalized RLS expression explicitly names `I_prop_welfare` and active effective mass `Q`. Profile-first display, assessed zero versus unknown, raw  $\Delta$ RLS versus normalized SignedGap, epsilon, uncertainty and final verdict are distinguished. The new checker consumes declared scores and comparisons; it does not calculate RLS, infer uncertainty or discover the complete stress family. The Canon and frozen Aligners workbook remain unchanged.

### **B03. Every contender and compatible joint stresses remain visible**

**Locations:** §11.3 and Appendix E.3; robustness schema, reference checks and tests.

A multi-option quantitative ALLOW record needs a current selectable nominal leader, a ranked record for every selectable contender, complete declared option closure, a complete robustness record and passing comparisons for every declared required variant/contender pair. Missing pairs, duplicates, undeclared variants, failed SignedGap values and unresolved joint-stress dispositions fail the local checks. Sole-survivor handling uses a separate review reference, never a fictitious pairwise Gap.

The checker does not certify that an analyst found every plausible alternative or required joint perturbation. A passing declared matrix with an omitted real stress remains a substantive review failure outside this local algorithm. This limit is explicit rather than disguised as exhaustive validation.

### **B04. Current SGP outputs are imported, not recomputed**

**Locations:** §12 and Appendix E.5; `sdp_links` and the examples.

MPS, FPP, GPR, SPR, ICP and RMCP remain distinct source-owned outputs. Public presentation leads with MPS-0 through MPS-4 or MPS-NE; NOT\_EVALUATED does not mean zero. No interface calculation turns fluent self-report or denial into sentence evidence. Human FPP-100 is distinct from the provisional collective RMCP-P100 anchor. RMCI\_L, when shown at all, remains a profile-bound heuristic, not percentage intelligence, moral worth or authority.

The adapter checks an imported band and source-reference presence; it cannot establish the underlying SGP assessment. A reference to a synthetic record is not a real MPS evaluation.

### **B05. UCI maturity does not become a hidden extra gate**

**Location:** §10.3, UCI/HOI diagnostic boundary paragraph; source: Canon §11.1.3A and §11.5.

UCI uses separate Cohesion, Flow, Resilience and Equity indicators. It is not reconstructed from RLS welfare cells. Provisional measurement does not become an independent high-stakes warrant or a scalar tie-break that fabricates certainty. CSV may use direct structural evidence when UCI is not material. No UCI calculator or universal instrument is invented.

### **B06. Configuration, sources and physical warrants stay distinct**

**Locations:** §3.1, §6.1–§6.3, §10.3 and §13.2.

Evidence cutoffs, parameter locks, source records, configuration and amendment rules are explicit. A recorded observation does not automatically become verified fact. Source-Coupling uses the owner's current tokens. Material causal claims expose their estimand, comparison, identification assumptions, confounding/interference treatment and validity domain. Monitoring, approval, simulation and interface success do not supply a missing PC-AEP/domain warrant. Late mitigation remains useful but cannot be claimed as earlier prevention.

### **B07. Recovery and successor handling preserve current authority boundaries**

**Locations:** §14.2, §14.4–§14.6 and Appendix E.4; runtime and execution records.

An outage/reconnection does not restore an old mode or expired approval. The user-facing pathway remains MODE 0 until reconciliation and authenticated re-entry under Agent §8 and §43. A material

successor does not inherit permission. The local checker detects incompatible declarations, time expiry, configuration mismatch, missing execution/monitoring references and unreconciled recovery. It does not verify real signatures, implement the full mode-permission matrix or actuate a control.

#### **B08. Untrusted input remains evidence, not instruction**

**Locations:** §14.4 and Appendix E.5; strict loader and local-reference resolver.

Duplicate JSON keys and nonfinite values fail. Additional properties do not silently create permissions. The optional resolver checks path containment, local file hashes, JSON fragment identity and synthetic-source labeling. Remote references are not fetched and remain unverified. Correct hashes establish byte identity only; factual, legal and physical conclusions require their own warrants.

#### **B09. Accessibility and conformance are tested at the correct level**

**Locations:** §15.3, §18 and CONFORMANCE\_COVERAGE.csv.

WCAG 2.2 is an explicit candidate target, not an earned badge. MHIOS M0–M3, Canon tiers, wrapper L0–L3 and Agent modes have no numeric equivalence. Target level, assessment status and evidence are separate. The coverage register maps all sixteen original vectors to executable checks, document-only obligations or manual implementation tests. Source-level non-color labels do not establish keyboard, screen-reader or cognitive accessibility.

### **5. Deliverable changes and preservation**

The main standard keeps its 24-section organization, existing useful paragraphs, ten inherited tables and historical migration rationale. The exact edit log records replacements, insertions, relocated release pins, current source references and header/footer repairs. A separately logged pagination repair adds keep-row-together flags to 124 table rows without changing their widths, style definitions, shading or borders.

The preparation workbook retains fourteen sheets and its existing visual design. Its current pins, session fields, qualified-state language, SGP/authority prompts and twenty preparation formulas were improved. Every workbook edit is separately recorded. This is the MHIOS preparation workbook, **not** the frozen Core Aligners Sheet 5.9; that Core workbook is unchanged.

The local schema is tighter and not wire-compatible with v2.0. Existing records are retained as immutable history; migration creates a new referenced view rather than silently inventing per-option facts. The draft YAML deliberately fails until completed. Six completed examples are explicitly synthetic, with local hash-bound supporting fixtures and an explicit replay clock.

### **6. What was not adopted**

No sixth gate, eighth dimension, replacement scoring method, consciousness estimator, moral rank, automatic legal permission, opaque “alignment percentage”, fabricated external registry or production control daemon was added. No empty external dependency was labeled bundled merely to make a completeness claim easier. No Core text or spreadsheet was changed to accommodate this companion.

A new schema and extensive tests are not a full deployed interface. Full canonical run-record interchange, the external state/audit registries, normative kernel index, applicable RPAP/PFAP, domain approvals, cryptographic attestation and independent deployment evaluation remain separate dependencies. The package is useful without claiming those dependencies have been fulfilled.

### **7. Verification interpretation**

The test suite includes positive examples as well as refusals and malformed states. It covers failed qualification, unranked contenders, negative/missing/duplicate comparisons, separate authority selection, provisional choice, expiry, recovery, successor declaration, source substitution, local path escape,

unassessed drafts and display-owner mismatch. Synthetic SignedGap calculations are independently recomputed in their fixture tests.

Complete Word/Markdown body-and-table token parity and 115 internal hyperlink destinations are checked. The main PDF was rendered and examined for layout. Core source hashes are recomputed from all fifteen carried-forward files. The preparation workbook's local formulas were evaluated using `artifact_tool`, with eight explicit behavior probes and a cached-error scan. No native Excel/Calc acceptance, empirical study, real-world control test, hosted publication or full Core re-audit is asserted in this companion update.

The final execution counts and final-file checksums are in the verification report and delivery receipt. Read them as scoped evidence, not proof that the framework or an implementation is universally correct.