SkillPilot Quality Handbook · Volume IV

Curriculum Quality, Validation, and Champion Trials

QA Maturity Levels M0–M7, QA Scopes and CQR Rules, Graph/View Validators

(GVR/APV/CPV), Semantic Atomicity, and Practical Champion Trials

Scope of this handbook

This volume unifies SkillPilot's durable curriculum quality, validation, and human-trial specifications: the orthogonal separation between QA maturity (`M0`–`M7`) and practical trials by Curriculum Champions, the detailed semantics of QA scopes and `CQR-*` rules in the quality dashboard, algorithmic CI validators (`GVR-*`, `APV-*`, `CPV-*`), the semantic-atomicity review workflow, and the operational Champion Guide alongside the forward-looking four-level QA cascade.

As of 11 November 2026 · License: Apache-2.0 (technical QA and validation specification; see LICENSING.md).

Generated status reports under docs/qa-ci/status/ are derived snapshots rather than hand-edited sources.

Contents

1 Curriculum Quality: QA Maturity Levels and Human Trials **3**
1.1 Objective, Guiding Principles, and QA Maturity Levels M0–M7 3
1.2 Normative Definition of M7 and Image Governance 3
1.3 Practical Human Trials by Curriculum Champions 4
1.4 Visual Presentation, Gray-to-Green Scale, Privacy, and Acceptance 5

2 Curriculum Quality: Maturity Levels, QA Scopes, and CQR Rules **6**
2.1 Core Concepts, State Coverage, and Route Model 6
2.2 Scope Maturity (M0–M4) vs. Curriculum Maturity (M0–M7) 7
2.3 Catalog of CQR-* Rules and Active CQR-302 Subject Balances 8

3 Graph, Applicability, and Composition-View Validation Rules **10**
3.1 CI Entrypoints, Enforcement Profiles, and Core DAG Checks 10
3.2 Emitted Graph Validation Rules (GVR-001 through GVR-016) 10
3.3 Filter Projection (APV-*), Required-Learning Contract, and Assessment Applicability 12
3.4 Composition-View Validator (CPV-000 through CPV-216) 12

4 Semantic Atomicity Review **14**
4.1 Domain Criterion, Scope Exclusions, and Review Statuses 14
4.2 Active Configurations, Fingerprints, and Expansion Workflow 14

5 SkillPilot Curriculum Champions **15**
5.1 Responsibilities, Registration, and Practical Trial (”Your human QA”) 15
5.2 Resignation Flow and Structure of the ”Curricula” Page 15

6 SkillPilot Quality Assurance: The Four-Level Champion Model **17**
6.1 Positioning and the Four QA Levels 17
6.2 Human-Agent Workflow and Governance 17

1 Curriculum Quality: QA Maturity Levels and Human Trials

Status and companion specifications

Binding overall specification for Issue #49. Technical dashboard derivation is specified in 2 and operational Curriculum Champion workflows in 5. Implementation builds directly on existing quality reports, review ledgers, Champion assignments, and completion records without introducing a parallel QA pipeline or general certification engine.

1.1 Objective, Guiding Principles, and QA Maturity Levels `M0`–`M7`

SkillPilot communicates the quality of a skill landscape through two **orthogonal** dimensions:

1. **QA maturity `M0`–`M7`:** Which mandatory quality gates are completely satisfied for the current content snapshot?
2. **Human trial:** Does an assigned Champion's learning progress show an active trial, has the Champion explicitly confirmed starting a trial, or has one Champion worked through and confirmed the entire advertised scope?

> **Maturity describes the verified inspection stage. Champion status describes practical human trial.** (Example: *Physics · M7 · Human-tested* — never a blanket guarantee of zero defects.)

Maturity levels form a cumulative ladder derived strictly from current central verification evidence (never copied from historical milestones or manual badges); human trials do not create a milestone `M8`:

Level	Domain statement
`M0`	No reliable baseline QA state established yet.
`M1`	Sources and extracted syllabus goals are ingested and traceable to official URLs.
`M2`	Source coverage, state/jurisdiction views, and course-level mappings are consistent.
`M3`	Required learning routes are closed and cleanly modeled at the atomic goal layer.
`M4`	Required learning routes including their terminal practice/exam endpoints are verified.
`M5`	School-ready core QA stage: all mandatory core checks pass.
`M6`	`M5` plus reviewed memory-learning decisions, deck traces, and view visibility.
`M7`	`M6` plus strictly completed deep curriculum QA across every current `curricularAtomic` goal of the subject.

1.2 Normative Definition of `M7` and Image Governance

`M7` requires **every** current `curricularAtomic` goal of the subject to satisfy the **strict intersection of all five gates** simultaneously (never summing or averaging gate counters or accepting rounded percentages; partial progress renders e.g. as *M6 · Deep QA: 472 of 478 learning goals completed*):

Gate	Required evidence
D – Descriptions	Two independent description reviews with resolved synthesis bound to current content and context.
P – Understanding evidence	Current, fully validated `positive-understanding-evidence-v2` profile with truthful approval status.
A – Atomicity	Current domain decision confirming semantic atomicity and scope (4).
M – Memory	Current justified decision on memorization necessity plus card and view-visibility traces.
V – Visualizations	Current domain AI image inspection bound to the exact active asset hash and content bindings (technical deferrals never count as exceptions).

```
M7 achieved =
    M6 still satisfied
    AND current goal denominator uniquely determined and non-empty
    AND every current curricularAtomic goal passes all 5 gates (D/P/A/M/V)
    AND required closeout checks pass
    AND zero open M7-relevant blockers
```

- **Separation from human approvals:** `M7` is a machine closeout gate and never mints human approvals (`needs_human_review` stays untouched; neither `humanApproved` nor `humanIssue Identified` satisfies or blocks Gate V/M7, while human findings remain separately visible).
- **Image principles & separate release gate:** (1) `KEEP` by default for sound existing visuals regardless of generator/age, (2) domain accuracy over origin, (3) generation is not approval, (4) genuine content inspection rather than mechanical hash refreshes, (5) human findings remain independently effective. Separate CI check `checkGoalVisualizationQaApprovalCoverage. ts` acts as an image release gate (accepting an existing human approval or current AI approval while still blocking open human defect findings even at machine `M7`).

1.3 Practical Human Trials by Curriculum Champions

Regular school-facing Champion trials require at least core QA stage `M5` (and may run in parallel with `M7` work) across three states:

State	Precondition	Visible badge
Not started	Neither in-scope progress by an active Champion nor an explicit trial start	Baseline status from QA maturity
In progress	Active, unpaused Champion role with in-scope progress or confirmed start	**Human QA in progress**
Completed	One Champion's full practical run across the advertised scope + final confirmation	**Human-tested**

- **Start signal (Product Owner decision, 20 September 2026):** Existing learning progress (>= 1 completed goal inside the active Champion role's assigned scope once `>= M5`) automatically counts as **Human QA in progress** — regardless of registration date or when progress occurred. With zero progress, **"Begin trial"** remains available; registration alone, out-of-scope progress, or paused/ended roles never trigger the status. Existing mastery is strictly a start signal, never completion proof (100 % imported or manually set mastery alone never produces "Human-tested").

- **Four conditions for "Human-tested":** (1) one complete pass across the advertised scope by **at least one individual Champion** (partial passes across multiple people are never pooled), (2) genuine content-bound practical completion evidence, (3) zero known unresolved blocking defects (completing a goal does not auto-close a reported bug), and (4) a single explicit final confirmation (*"All goals completed. Complete trial?"*).
- **Scope binding, curriculum edits & pause/resignation:** An explicit trial action binds its scope (e.g., *Physics, Hesse, Upper Secondary* never certifies all states or stages). When curriculum content changes, `M7` reverts to the proven lower milestone if current goals become open, while human trials require re-testing only new or materially changed stations. Paused or unfinished trials drop the active "Human QA in progress" badge; a validly confirmed completion remains preserved even if the Champion later resigns.

1.4 Visual Presentation, Gray-to-Green Scale, Privacy, and Acceptance

Every UI surface consumes the same central derivation, labels, and ordered **gray-to-green scale** from `M0` (neutral light gray) through `M1`–`M4` (progressively stronger gray-green/green) and `M5` (medium green) to `M6` (vivid green) and `M7` (deep saturated green) — always paired accessibly with milestone code and text in both light and dark modes. Beside it, the trial status badge independently shows **Experimental** (`< M5`, muted warning), **Machine QA** (`>= M5` with no active/completed human trial, orange — combinable with a green `M7` badge), **Human QA in progress** (muted green), or **Human-tested** (vivid green with check icon). Collection rows such as *Gymnasium (DE)* never inherit a single subject's top milestone or trial badge, instead summarizing factually (e.g., *Human QA in 2 subjects*). Public QA artifacts exclude all private learner, session, and chat data.

Acceptance area	Normative verification requirement
`M7` award & counters	Granted solely on valid `M6` plus strict 5-gate intersection across all current goals; zero averaging, rounding, or deferred-work exemptions.
Approval authority	AI records never masquerade as human approvals; explicit human defect findings remain active.
Trial start & closeout	In-scope progress (date-independent) or explicit start triggers "in progress"; completion requires one full individual pass, zero blockers, and final confirmation.
Scope, UI & collections	Partial scopes stay bounded, edits invalidate only affected stations, uniform gray-to-green palette everywhere, zero blanket collection promotion.

2 Curriculum Quality: Maturity Levels, QA Scopes, and CQR Rules

Workbench semantics and 4 November 2026 concept correction

Domain specification for `/quality-dashboard` and the generated snapshots under `docs/qa-ci/status/curriculum-quality-status.json` (producer: `cd app && npm run quality:curriculum-status` via `app/scripts/generateCurriculumQualityStatus.ts`; overview: Dashboard Status Layer). **Concept correction, 4 November 2026:** Existing `M7` closeouts or passing route checks do not automatically prove that exam tasks verify independent application of every implicitly recognized competence under Assessment-Gated Completion; that authoring obligation requires a dedicated task/rubric review.

2.1 Core Concepts, State Coverage, and Route Model

A **curriculum** in the dashboard is a canonical `SkillLandscape` file under `curricula/DE/Gymnasium/canonical/`. To evaluate **bidirectional state (Bundesland) coverage**, the dashboard separates three distinct atomic counts (empty `contains`, including atomic practice, memory, and assessment nodes):

1. **`Raw atomic (Roh-Atomar):`** all atomic leaf nodes in the canonical JSON file (diagnostic metric, not the state denominator).
2. **`DE-view atomic (DE-Sicht atomar):`** union of atomic nodes rendered across national `de-de-*` composition views.
3. **`State-view atomic (Bundesland-Sicht atomar):`** union of atomic nodes rendered across a state's `de-<state>-*` composition views.

For each declared state, `CQR-003` checks **both directions** against repository source snapshots (completeness of raw PDF extraction itself is guarded via snapshot/extract ledgers):

- **View -> Source:** Every visible state-view atom carries direct (`provenance`, reviewed `mapping`) or explicitly reviewed `requires-closure` surrogate evidence for that exact state (`override`, `child-union`, and unreviewed `requires-closure` never count as syllabus proof).
- **Source -> View:** Every registered original state goal maps across its full atomic source closure onto visible canonical atoms, and every atom extracted from readable source snapshots is registered in the membership/closure ledgers.

A **QA scope** audits an explicit learner-facing slice (counting registered QA profiles rather than individual graph paths; absent QA scopes cap the aggregator at `M2`):

Scope ID	Stage	Atoms	Motivation anchor & terminal autonomy
`canonical-math-sek1`	Sek I	247	`Warum Mathematik? - Entdecken, Muster & Alltag` -> 38 tasks under `Prüfungen Jahrgangsstufe 5..10`
`canonical-math-sek2`	Sek II	397	`Warum Mathematik? - Denken, Muster & Zukunft` -> atomic tasks under `Übungen E-Phase`, `Q1..Q4` & `Prozesskompetenzen`
`canonical-physics-sek1`	Sek I	44	`Warum Physik?` -> atomic tasks under `Übungen Sekundarstufe I Physik`
`canonical-physics-sek2`	Sek II	278	`Warum Physik?` -> atomic tasks under `Übungen E-Phase` and `Übungen Q1..Q4`

Route direction and Product Owner target norm (8 November 2026): Stored edges `A. requires = [B]` point from the later goal to its prerequisite (`Motivation <- Prerequisites <- Atomic goals <- Terminal practice / exam`). While transition check $R_{\text{effective}}$ includes inherited cluster `requires`, canonical route verification uses strictly direct atomic edges $R \cap (A \times A)$. Mandatory migration order toward $R \subseteq A \times A$: first materialize/preserve complete reviewed atomic prerequisites, then remove covered non-atomic `requires`, and finally transitively reduce while preserving R_{at}^{+} (see Graph Definition §3.1). Since **11 November 2026**, `CQR-001` checks the achieved state (atomic normal form and transitive minimality) for every root of the canonical Gymnasium catalog; the migration order itself is proven only for reductions bound to the reduction-evidence registry, while `CQR-102` and `GVR-012` continue to require direct atomic routes.

2.2 Scope Maturity (M0–M4) vs. Curriculum Maturity (M0–M7)

Each `CQR-*` rule evaluates to `pass`, `warn`, `fail`, or `not_configured`. Since 11 November 2026, the promotion chain enforces global atomic normal form and `G17` transitive minimality through `CQR-001` across the canonical Gymnasium catalog (any violation resets the curriculum and all its scopes to `M0`; CI fails hard on protected floors, while roots outside the canonical Gymnasium catalog are not checked by `CQR-001`). The Rule-to-Milestone Matrix additionally documents where catalog target milestones differ from current aggregator promotion thresholds (e.g., `CQR-104` gating curriculum `M4` and `CQR-203` gating curriculum `M5`).

Scope level	Exact condition per individual QA scope (M0–M4 scale)
M0	CQR-101 is not pass (effective route from motivation to terminal autonomy remains open).
M1	CQR-101 is pass, while CQR-102 or CQR-103 remains open (atomic routing or cluster requires).
M2	CQR-101..103 are pass, while CQR-104 or CQR-201/CQR-202 remains open (visibility / terminal examData).
M3	Route and baseline assessment rules are pass, while CQR-203 remains open.
M4	All required route and assessment rules including CQR-203 are pass.

Curriculum	Cumulative subject-level condition (M0–M7 scale)
M0	CQR-001 or CQR-002 fails, or source snapshot ingestion with official HTTP(S) URLs (CQR-000) is incomplete.
M1	CQR-000..002 pass, while state view coverage (CQR-003) or GK/LK mapping consistency (CQR-004) is open.
M2	CQR-000..004 pass; QA scopes are not yet configured or at least one scope has not passed CQR-101.
M3	All QA scopes pass CQR-101, CQR-102, and CQR-103, while exam-mode readiness remains incomplete.
M4	All QA scopes reach >= Scope M3, while at least one misses Scope M4 or CQR-301/401/501 is open.
M5	**School-ready core QA stage:** all QA scopes at Scope M4 plus CQR-301, CQR-401, and CQR-501 pass.
M6	M5 achieved plus CQR-302 pass (reviewed memory-card decisions, card ledgers, and view visibility).
M7	M6 achieved plus CQR-303 pass (strict 5-gate D/P/A/M/V completion across all current curricularAtomic goals).

2.3 Catalog of CQR-* Rules and Active CQR-302 Subject Balances

Rule	Target	Verified invariant and status logic
CQR-000	M1	**Source Snapshot Ingestion:** Readable snapshots in source-landscape-registry.json, structured official HTTP(S) url in sourceDocument(s) (local PDFs remain gitignored working caches; G8/G9 uses views/provenance rather than duplicate goal sets), and complete registration in source-goal-membership-registry.json & source-goal-closure-registry.json.
CQR-001	M0	**Current Graph Integrity & Atomic Prerequisite Normal Form (G17):** Since 11 Nov 2026 (curriculum-quality-v5 via curriculumGraphIntegrity.ts and auditTransitiveMinimality.ts), verifies per canonical root including required foreign goals and inheriting ancestors: resolvable unique goal IDs, acyclic contains and effective requires, atomic normal form on every authored requires edge, transitive minimality (minimalityFindings = 0), explicit type consistency, and digest-bound reduction-evidence registry sync (any failure resets curriculum and scopes to M0).
CQR-002	M0	**Explicit Type Consistency:** Optional type: "atomic" requires empty contains; type: "cluster" requires non-empty contains (mismatches = 0).
CQR-003	M2	**Bundesland Source-Backed View Coverage:** Bidirectional full coverage between de-<state>-* views and registered syllabus atoms with zero projection warnings/applicability errors (APV-201 overlays never count as syllabus proof; detailed audit snapshots for Math & Physics live at docs/qa-ci/status/curriculum-source-coverage-audit.{json,md}).
CQR-004	M2	**Course-Level Mapping Consistency:** Source goals marked LK map to LK, GK_LK to GK+LK, and unspecified conservatively to GK+LK (allowing LK-only only via justified courseLevelDecision; partial edges do not count as resolved).
CQR-101	M3	**Effective Full Route Coverage:** Every scoped atom reaches a motivation anchor backward in $R_{\text{effective}}$ and at least one terminal autonomy goal forward in $R_{\text{effective}}^{-1}$.
CQR-102	M3	**Atomic Direct Route Coverage:** Identical backward/forward connectivity strictly along direct atomic edges $R \cap (A \times A)$ with zero cluster bridges.
CQR-103	M3	**No Scoped Cluster requires:** Zero clusters matched by clusterSelector carry authored requires (scopedClusterRequires = 0).
CQR-201	Scope M3	**Terminal Autonomy examData:** Every atomic child of configured terminal autonomy clusters carries complete examData.
CQR-301	M5	**Semantic Atomicity Review Freshness:** All in-scope content leaves hold fresh fingerprint-matched atomic records with zero missing, stale, obsolete, needs_developer_review, or non_atomic items (4).
CQR-302	M6	**Memory-Card Decision Trace:** Default decision no_memory_needed; memory_required links existing memory nodes visible in visibilityScopes and canonical decks under curricula/DE/Gymnasium/memory-decks/ (neutral de_gymnasium_* IDs without state prefixes), while every kept primary card in *.cards.review.jsonl traces via originGoalId to a memory_required goal (remove cards are absent from active decks).

Rule	Target	Verified invariant and status logic
CQR-303	M7	**Deep Curriculum QA Completion:** Directly projects app/scripts/reportDeepUnderstandingRollout.ts: strict 5-gate D/P/A/M/V intersection across all current curricularAtomic goals plus 6 passing subject closeout checks (deferred_provider_limitation and deferred_quality_review remain open; humanTrialBlockingFindings is projected separately).
CQR-401	M5	**Composition View Availability:** At least one .view.json exists under curricula/DE/Gymnasium/composition-views/ for the landscapeId.
CQR-501	M5	**Applicability Warning Debt:** activeWarnings + obsoleteAcceptedWarnings = 0 against docs/qa-ci/applicability-accepted-warnings.json (APV-202 counts as diagnostic rather than warning debt).

Current reviewed CQR-302 subject balances (Gymnasium DE): Mathematics (750 goals: 713 no_memory_needed / 37 memory_required, 64 cards, 6 nodes), Physics (425: 298 / 127, 148 cards, 5 nodes), Chemistry (376: 322 / 54, 55 cards, 6 nodes), Biology (355: 348 / 7, 17 cards, 1 node), Informatics (207: 198 / 9, 28 cards, 5 nodes, 18 views), Latin (115: 82 / 33, 34 cards, 5 nodes), German (268: 194 / 74, 48 cards, 5 nodes), History (156: 100 / 56, 57 cards, 5 nodes), Politics & Economics (413: 306 / 107, 62 cards, 5 nodes, 14 views), and Economics (303: 250 / 53, 51 cards, 5 nodes).

3 Graph, Applicability, and Composition-View Validation Rules

Authoritative CI source and November 2026 validation boundaries

Single source of truth for algorithmic graph and view validation in `.github/workflows/ci.yml` (`graph-validation` & `backend-ci`). Transitive minimality (Graph Definition §3.1) together with complete local assessment routes, protected explicit exam/card targets, and versioned learning-unit contracts (unit validity rules R1–R5 of 11 November 2026; Assessment-Gated Completion) are binding target requirements for forthcoming validator adaptations. Atomic normal form and transitive minimality (`G17`) are not checked by `validate:graph`; since 11 November 2026, `CQR-001` checks them for the canonical Gymnasium catalog (2). Existing CI passes never substitute for content/rubric review verifying independent application of recognized competences.

3.1 CI Entrypoints, Enforcement Profiles, and Core DAG Checks

- **Raw graph validator:** `app/scripts/validateGraph.ts` (`npm run validate:graph`) — profiles: `default` (structural and `GVR-*` rules fail CI as `error`) vs. temporary migration mode `VALIDATE_GRAPH_STRICT_RULES=0` (`legacy-warn`, under which **`GVR-012`, `GVR-014`, `GVR-015`, and `GVR-016` remain unconditional `error`s**).
- **Filter-projection validator:** `app/scripts/validateViewFilters.ts` (`npm run validate:view-filters`) covering the 17 reviewed canonical Gymnasium landscapes (`Mathematik`, `Physik`, `Chemie`, `Biologie`, `Informatik`, `Deutsch`, `Englisch`, `Französisch`, `Griechisch`, `Chinesisch`, `Geschichte`, `Politik und Wirtschaft`, `Musik`, `Latein`, `Spanisch`, `Wirtschaft`, `Overview`).
- **Duration-model & composition-view validators:** `app/scripts/reportHeMathDurationProjection.ts --check` (`npm run check:he-math-duration-projection` for Hesse Mathematics Sek-I `G8`/`G9`) and `app/scripts/validateCompositionViews.ts` (`npm run validate:composition-views` across all `.view.json` files under `curricula/DE/Gymnasium/composition-views/`).
- **Archive & legacy-handoff gates:** `scripts/validate_hessen_upper_secondary_{archive_paths,legacy_refs}.py`, `scripts/validate_hessen_lower_secondary_{archive_paths,legacy_refs}.py`, and `scripts/validate_bavaria_gymnasium_{archive_paths,legacy_refs}.py` enforcing the explicit `*-retirement-registry.json` allowlists under `curricula/DE/Gymnasium/provenance/` and `retained-asset-registry.json`.

Always-active core checks (`error`): JSON parsing, goal-ID and logical `shortKey` uniqueness, referential integrity and self-reference guards on `requires`/`contains`, allowed `phase`/Leitideen/competency tag domains, acyclicity of direct `requires`, `contains`, and effective `requires`, `curricula/curriculum_manifest.json` sync, and project invariants (such as cross-subject Physics-to-Mathematics resolution).

3.2 Emitted Graph Validation Rules (`GVR-001` through `GVR-016`)

Rule	Default / Scope	Verified structural invariant
`GVR-001`	`error` · local	A goal must not directly require any of its `contains` ancestors.
`GVR-002`	`error` · local	Phase compatibility lint: effective prerequisites do not point to later phases (`E`, `Q1..Q4`, `Abitur`, `S*`, `J*`).

Rule	Default / Scope	Verified structural invariant
GVR-003	error · local	A goal must not directly require its immediate contains parent (primary deadlock guard).
GVR-004	error · DE_HES_S_GYM_2_*	First atomic node is a motivation anchor (Warum/Why; excluding OVERVIEW).
GVR-005	error · DE_HES_S_GYM_2_*	Every atomic node reaches the motivation anchor transitively via effective requires ($R^{+}_{\text{effective}}$).
GVR-006	error · DE_HES_S_GYM_2_*	A goal must not directly require any of its direct contains children (including OVERVIEW).
GVR-007	error · mit-ocw-*	MIT OCW module atoms (module:*) include concept + practice + assessment links in resourceLinks targeting https://ocw.mit.edu/courses/....
GVR-008	error · local	Canonical top-level resourceLinks is the sole permitted goal-level helper-link field.
GVR-009	error · local	Optional type metadata matches structural leaf/non-leaf classification (atomic iff contains is empty).
GVR-010	error · landscapeId	shortKey is unique within the logical landscapeId (locale file pairs allowed only on identical goal UUIDs).
GVR-011	error · profile scopes	Selected curricularAtomic goals reach their stage-specific motivation anchor via $R^{+}_{\text{effective}}$.
GVR-012	**unconditional** error · Math DE Sek I & II	**Hard full-route invariant:** Every curricularAtomic goal lies stage-locally on a direct atomic requires path from its orientation anchor to an atomic practiceAssessment terminal (Sek I: individual tasks under Prüfungen Jahrgangsstufe 5..10; Sek II: atomic tasks under Übungen E-Phase, Q1..Q4, Prozesskompetenzen; complete semanticKind ledger, zero cluster/structure shortcuts).
GVR-013	error · Math DE	Zero selected non-memory atomic goals carry direct local requires edges targeting cluster nodes.
GVR-014	**unconditional** error · all	Motivation/orientation nodes are pure, prerequisite-free atomic entry anchors with zero examData, examples, memory/SRS, or assessment tags.
GVR-015	**unconditional** error · Math DE Sek I & II	Reviewed bilingual orientation outlooks (2–4 paths with curricularArea roots plus atomic curricularAtomic entryGoalIds and milestoneGoalIds, direct atomic link to the anchor, zero internal authoring terms like Source-Extraction, and fail-closed runtime intersection with targetGoalIds).
GVR-016	**unconditional** error · examData	GFM table integrity across taskContent* and solutionContent* (valid delimiter row and uniform column count on consecutive pipe rows).

3.3 Filter Projection (APV-*), Required-Learning Contract, and Assessment Applicability

Rule	Default	Meaning in validate:view-filters & backend-ci
APV-102	error	A necessary canonical atom is missing from personal acquisition targets, emitted definitions, or completion inputs (enforced by the native backend gate).
APV-103	error	A visible goal is unreachable from the projected root (retired compatibility clusters must set extendedData.applicability Projection = "excluded").
APV-201	warning	Explicit applicability override path used (accepted cases recorded with rationale in docs/qa-ci/applicability-accepted-warnings.json).
APV-202	diagnostic	Visibility backed solely by partial mappings (1:n mapping-shape diagnostic, not warning debt).
APV-203	warning	Compiled applicability diverges from committed applicability metadata.

- **Required-learning contract (Product Owner, 10 November 2026):** Under the closed personal learning basis (View Projection & Goal Placement), every effective necessary atom must be an actual acquisition target in the full Personal Curriculum regardless of its source jurisdiction (without granting CQR-003 source coverage). APV-102 is enforced in .github/workflows/ci.yml (backend-ci -> ./gradlew check) by FirstMath PrerequisiteReductionTest.currentCatalogLearningScopesRetainEveryCanonical AcquisitionTarget across all 20 subjects and 17 jurisdiction choices (including the Biology regression rejecting removed required targets).
- **Explicit assessment applicability basis (dormant compiler support):** For atomic practice/assessment endpoints with applicabilityFromRequires: true, optional extended Data.assessmentApplicabilityBasis = { kind: "allOfGoalApplicability", goal Refs: [...] } supplies seed-eligibility evidence only (rejecting cycles, duplicates, foreign, and memory dependencies), while required learning still follows canonical requires. Proof extensions, emitted-target APV checks, and native comparisons remain prerequisites before any production data admission.

3.4 Composition-View Validator (CPV-000 through CPV-216)

validate:composition-views compiles every .view.json under curricula/DE/Gymnasium/ composition-views/ deterministically against its canonical graph:

Rule	Level	Verified composition-view invariant
CPV-000	error	Composition-view JSON file could not be loaded or parsed.
CPV-001	error	Required view or structure-node metadata is missing or invalid.
CPV-002	error	Referenced canonical subtree root does not exist in the canonical graph.
CPV-004	error	Two referenced canonical subtree roots overlap or the same root is referenced multiple times.
CPV-005	error	Compiled default tree contains the same canonical goal more than once.
CPV-006	error	Compiled default tree assigns more than one visible parent to a canonical goal.
CPV-007	error	An empty structure node remains visible in the compiled view tree.
CPV-008	error	Unknown projectionRole used on a composition reference.
CPV-009	error	Direct goalEntry exposes a curricularArea cluster opaquely instead of expanding via canonicalSubtree or explicit atoms.

Rule	Level	Verified composition-view invariant
`CPV-101`	`warning`	Structure-node label is too generic to be review-safe.
`CPV-102`	`warning`	Referenced canonical subtree root still carries a phase- or state-specific title.
`CPV-210`	`error`	Canonical Mathematics Sek-I view omits the reviewed motivation anchor.
`CPV-211`	`error`	Canonical Mathematics Sek-I view omits required year exam folders or individual task nodes.
`CPV-212`	`error`	Malformed Mathematics Sek-I exam folder structure (missing year folder, non-atomic task, missing `examData`, cluster prerequisite, duplicate, or stale aggregate endpoint).
`CPV-213`	`error`	Mathematics Sek-I view still exposes the legacy `Übungen Sekundarstufe I` aggregate branch.
`CPV-214`	`error`	Mathematics Sek-II or cross-stage view omits the reviewed Sek-II motivation anchor.
`CPV-215`	`error`	Mathematics Sek-II or cross-stage view exposes zero atomic terminal assessment endpoints.
`CPV-216`	`error`	Learner-facing composition label exposes internal authoring terms such as `Source-Extraction`.

4 Semantic Atomicity Review

Rule version semantic-atomicity-v1

This QA lane verifies whether technical content leaf goals (empty `contains`) are also **semantically atomic** (`CQR-301` in 2.3). Tooling tracks coverage and fingerprint freshness only; the atomicity judgment itself is strictly semantic rather than algorithmic.

4.1 Domain Criterion, Scope Exclusions, and Review Statuses

Under `semantic-atomicity-v1`, a technical content leaf goal is semantically atomic if and only if it expresses **one cohesive assessable competence** rather than a bundle of independent routines (conjunctions such as *"calculate and interpret"* act as review signals rather than automatic failures). Non-content leaves are intentionally excluded and validated in their dedicated lanes:

- terminal practice and assessment nodes (`Practice`, `Assessment`, or `examData`),
- motivation and orientation anchors (`Motivation`, `Orientation`),
- memorization and SRS deck leaves (`memorization` or `srs-deck:*`).

Review status	Meaning in the JSONL ledger
`atomic`	Reviewed and accepted as semantically atomic (`semanticAtomic: true`).
`needs_developer_review`	Escalated for closer subject/developer review (`semanticAtomic: null`).
`non_atomic`	Too broad; scheduled for splitting (`semanticAtomic: false`).

4.2 Active Configurations, Fingerprints, and Expansion Workflow

Configuration files and JSONL review ledgers reside under `curricula/DE/Gymnasium/quality/semantic-atomicity/`, for example the pilot lane `canonical-math-j8-linear-functions-pilot`, the foundational lane `canonical-chemistry-sek1-foundations` (with 26 accepted atomic leaves), and the full-subject lanes `canonical-physics-full` and `canonical-math-full`. In the local Workbench, route `/semantic-atomicity-review` provides an interactive status editor while CLI scripts run the checks:

```
cd app
npm run quality:semantic-atomicity
npm run quality:semantic-atomicity:check
```

`quality:semantic-atomicity:check` fails on process violations (missing in-scope record, stale fingerprint after semantic edits, malformed or duplicate rows), whereas known `needs_developer_review` and `non_atomic` items remain visible in the `CQR-301` (`M5`) queue. Every record stores a **semantic fingerprint** over `title`/`titleEn`, `description`/`descriptionEn`, `shortKey`, `phase`/`area`/topic code, and `nodeKind` (`--write-fingerprints` may be run solely as final bookkeeping after an actual semantic review).

Six steps to onboard a new subtree: (1) add a `.config.json` under `curricula/DE/Gymnasium/quality/semantic-atomicity/`, (2) reference the canonical landscape and subtree roots, (3) initialize `.review.jsonl` with one record per content leaf, (4) run `npm run quality:semantic-atomicity -- --config=<path>`, (5) write `--write-fingerprints` once semantic decisions are recorded, and (6) wire the config into CI only after the initial batch is complete.

5 SkillPilot Curriculum Champions

Role and companion concept

Operational guide for **SkillPilot Curriculum Champions** taking ownership of a curriculum's practical viability and didactic quality (an unlimited number of Champions may share a curriculum across content, didactics, community, and tooling). Normative maturity and trial derivation rules live in 1.

5.1 Responsibilities, Registration, and Practical Trial ("Your human QA")

Champion responsibilities: (1) make the curriculum practically effective in real teaching/learning settings, (2) step into the learner role and work through required learning and assessment stations, (3) drive improvements via GitHub issues and pull requests (`https://github.com/enpasos/skillpilot`), (4) recruit learners and teachers to benefit from it, and (5) engage institutional curriculum owners such as ministries or universities.

- **Registration on the "Curricula" page:** Requires a **SkillPilot ID** (shown publicly only in masked form: first 5 characters + ...) and **GitHub OAuth** ownership verification (Git Hub ID shown publicly).
- **"Your human QA" entry point:** Accessible anytime via **Connect with GitHub** (even when new registration is closed; expired logins require signing in again rather than re-registering, and load errors surface an explicit retry action rather than hiding controls).
- **Start signal & scope binding:** Once core QA reaches `M5`, existing learning progress (>= 1 completed goal inside the active role's personal curriculum/topic scope) automatically counts as **Human QA in progress** regardless of registration date; with zero progress, **Begin trial** remains available. An explicit trial action binds that scope so later personal selection edits do not silently rewrite it; paused or ended roles never self-reactivate from existing progress.
- **Practical evidence vs. raw mastery:** Imported or manually raised mastery never qualifies for **Human-tested**; pre-existing completions count only when provenance and current content bindings match. **Human-tested** requires **one** Champion's complete practical evidence across the advertised scope, zero known unresolved blocking findings (including non-ledger defects known to the Champion), and explicit final confirmation (`M7` is neither a prerequisite for starting human trials nor proof that one occurred).

5.2 Resignation Flow and Structure of the "Curricula" Page

Stop Championship (resignation) process: (1) click *Stop Championship*, (2) re-verify ownership via **GitHub OAuth**, (3) select one or more active Champion assignments, and (4) confirm deregistration — unfinished trials cease to be active while completed practical evidence and its history remain intact.

The **"Curricula" page** comprises four sections:

1. **Visual introduction:** Four-panel *Nano Banana Pro* cartoon explaining the Curriculum Champion concept.
2. **Registration action: "Connect with GitHub"** button initiating OAuth verification and role signup.
3. **Curriculum directory:** Full catalog of available curricula with total atomic goals, aggregated mastered achievements, QA milestone (`M0`–`M7`), and human-trial badge.

4. **Champion leaderboard:** Per-curriculum roster displaying public GitHub ID, masked Skill Pilot ID, and GitHub Issue / Pull Request contribution counts.

6 SkillPilot Quality Assurance: The Four-Level Champion Model

Forward-looking expansion roadmap

Target architecture for scaling curriculum QA across algorithmic CI, AI agents, and human practice reviews. Currently shipped public maturity and trial badges are governed strictly by 1 (`Level 1a` and `Level 2` are active anchors today, whereas `Level 1b` and `Level 1c` represent planned expansion stages).

6.1 Positioning and the Four QA Levels

To scale from AI-derived raw graphs to classroom-ready skill landscapes without burdening domain educators with JSON schemas or DAG invariants, this roadmap unfolds the former `Level 1` into three automated/agentic stages (`1a`, `1b`, `1c`) ahead of the human practice certificate (`Level 2`):

Level	Current status	Inspection focus and representative checks
`Level 1a`: CI pipeline	**Active today**	Content-agnostic algorithmic integrity: acyclicity of `contains`/`requires`, schema and ID checks, valid projected learner views (3).
`Level 1b`: Static AI agent	Forward-looking	Static didactic and formal guardrails without runtime simulation: route coverage from motivation to autonomy, clean separation of `sourceRef` and didactic assets, formal completeness of exam and rubric schemas.
`Level 1c`: Dynamic AI agent	Forward-looking	Simulated learner walkthroughs under live runtime rules: frontier and filter behavior, dead ends or hidden blockers, exam simulations with incomplete/wrong answers, and SRS `Memorize` mechanics.
`Level 2`: Human Champion	**Active today**	Authentic pedagogical evaluation in the WebGUI or AI coach: tone, age fit, motivation, subject alignment, and UX — filing issues/PRs toward **Human-tested** (5).

6.2 Human-Agent Workflow and Governance

The cascade operates as structured human-agent teamwork:

1. `Level 1a` blocks mathematically or structurally broken graphs automatically in CI.
2. `Level 1b` and `Level 1c` are designed to pre-triage static and simulated findings at scale.
3. `Level 2` evaluates real pedagogical quality in authentic learning contexts.
4. Findings are tracked and resolved transparently via GitHub issues and pull requests.
5. Human maintainers arbitrate edge cases and safeguard pedagogical intent.