

RippleLogic Cascade Standard

Integrity surface	Current requirement	Claim boundary
Governing sequence	RG -> RF/NCRC -> TRC -> CSV -> RLS. Only RF/NCRC, TRC, and CSV reject options; RG controls claim authority and RLS ranks survivors.	No sixth gate and no RLS rescue.
Severe-rights agreement	Review and computation use the same Canon activation predicate.	Companion standards may not redefine it.
Teaching shorthand	Eligible(a) may summarize prior qualification only in teaching surfaces.	It is not a sixth predicate or independently authoritative state.

Configuration-Assurance Thread Through the Existing Cascade

Configuration assurance does not add a gate. RG binds claims to the declared configuration and evidence; RF/NCRC protects rights under that configuration and transition; TRC evaluates ruin pathways; CSV tests transition control, post-state viability, and qualification continuity; RLS ranks only survivors. A material change reopens the affected earlier stages before further execution or renewed ranking.

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

This standard is the compact public reference for the current five-level RippleLogic cascade inside the MathGov framework. It prevents companion documents from repeating long synchronization notes.

TRC-CSV feedback rule. If CSV discovers a catastrophic, irreversible, or ruin-path scenario not represented in TRC, the run MUST reopen TRC before RLS. CSV does not absorb TRC.

TRC-CSV discriminator. TRC is the ruin veto. CSV is the viability and control test. Ordinary bounded operational risks usually route to CSV; catastrophic, irreversible, lock-in, systemic, or ruin-path risks route to TRC. Double-material evidence must be evaluated in both places and documented without double-counting it as ordinary residual welfare unless a bounded residual remains after gate treatment.

Gate-critical confidence guard. Low confidence in adverse rights-covered, catastrophe-covered, or material CSV impacts cannot by itself make an option pass RF/NCRC, TRC, or CSV. If a gate outcome could change, the run must use a governed conservative bound, collect evidence and rerun, narrow the claim, downgrade, escalate, or refuse the stronger claim.

Baseline pointer. Gate-admissibility cells and residual welfare cells may use distinct baselines where the Canon requires it; see Canon §5.1A for the floor-reference versus status-quo dual-baseline rule.

Reproducible run interface

The controlling reading specification is

Core_15/MATHGOV_REPRODUCIBILITY_AND_USE_STANDARD_v1.7.docx. The cascade is easy to state but must not be implemented loosely. A conforming implementation uses the deterministic run-state and parameter-completeness rules in the current MathGov Reproducibility and Use Standard. In particular, an option may enter RLS only after its RG, RF/NCRC, TRC, and CSV states establish selectability; missing decision-material parameters cannot silently become zero or an undeclared default; and selection remains separate from authority and execution.

The specified JSON schema and validator check structural and semantic contradictions. Their presence and successful execution must be established from the actual implementation inventory; the Core 15 document bundle alone does not establish them. They do not calculate moral truth or validate the underlying evidence.

Public cascade

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

Plain language: Ground reality -> protect rights -> bound ruin -> contain/verify viability -> score the ripples.

Two-phase interpretation

The public cascade may be taught as two phases:

1. **Qualify options:** RG -> RF -> TRC -> CSV.
2. **Rank survivors:** RLS.

This is not a new cascade and not a sixth gate. It is the teaching and implementation framing of the same method: first determine what is grounded, rights-safe, tail-safe, and structurally viable; then rank only those survivors by residual ripple impact.

Formal shorthand

RG/RSG -> RF/NCRC -> TRC -> CSV -> RLS

The formal components are:

- **RG/RSG:** Reality Grounding / Reality-Surface Grounding. Claim-authority precondition, not an ordinary option-rejecting gate.
 - RG status rule: Reality Grounding is Level 1 of the public method and can force narrowing, escalation, exploratory marking, or refusal. RF/NCRC, TRC, and CSV are the option-filtering gates; RLS ranks only what survives them.
- **RF/NCRC:** Rights Floor, formally implemented by the Non-Compensatory Rights Constraint. Non-compensatory gate.
- **TRC:** Tail-Risk Constraint. Non-compensatory catastrophic-risk gate. Its qualifying states are **TRC_PASS** and the explicitly assessed **TRC_NOT_TRIGGERED**; the latter is not a computed CVaR pass.
- **CSV:** Containment and Structural Viability. Selectability gate for containing-system integrity and execution viability.
- **RLS:** RippleLogic Score. Residual welfare-ranking layer applied only to options that remain selectable.

Rights-specific non-attenuation rule

RF/NCRC is not a welfare-duration calculation. For adverse rights-covered instances, short duration and low analyst confidence cannot by themselves weaken the floor. The run must apply the Canon's rights-floor impact channel, categorical-prohibition screen, and severe-hazard risk-tolerance test. Missing severe-hazard evidence prevents an ordinary pass.

Tail Emergency Mode rule

When rights-safe options exist but every option fails TRC, the ordinary selectable set is empty. The default is redesign, delay, escalation, no-action review, or refusal. A provisional least-CVaR action may be used only when delay/no-action is evidenced to be unavailable or worse, an independent challenger has reviewed alternatives, a hard maximum exposure is respected, and time limits, monitoring, shutoff, remedy, and return-to-normal criteria are binding. The state is **TAIL_EMERGENCY_PROVISIONAL_ACTION**, not **SELECT** or **TRC_PASS**.

Normalized residual ranking rule

RLS divides the weighted impact numerator by the active effective weight mass after masks and cell multipliers. Define $q(u,d)=w_u*v_d*m(u,d)*kappa(u,d)$, $Q=\sum_u \sum_d q(u,d)$, and $RLS(a)=[\sum_u \sum_d q(u,d)*I_{prop_welfare}(u,d,a)]/Q$, with $Q>0$. This preserves the $[-1,+1]$ interpretation and prevents masks or κ multipliers from changing score scale. A worked example with $Q=1$ may report the numerator as the score, but the denominator remains part of the canonical definition. If active mass is zero, ranking is undefined and must be repaired or refused.

NCRC three-channel rule. An option passes RF/NCRC only when, for every right, $v_r^{floor}=0$, $v_r^{cat}=0$, $v_r^{risk}=0$, and no required categorical-prohibition or severe-hazard evidence field is unresolved. A compact table displaying only the floor channel must label it as such and record the other two channels separately.

Computability vs realizability rule

A computed, simulated, generated, or model-fluent option is not automatically grounded, selectable, executable, or ethical. RG supplies claim authority; RF/NCRC supplies rights admissibility; TRC supplies ruin bounding; CSV supplies structural and execution viability; RLS ranks only the surviving selectable set.

Source-coupling rule

A downstream output, interface result, benchmark success, compliance status, institutional permission, inherited procedure, or model-fluent answer is not proof of grounded capability. When material, Reality Grounding must record the claimed capability, enabling conditions, boundary conditions, source evidence, inherited assumptions, downstream compensations, and required claim action. CSV consumes this evidence when weak source coupling creates dependency, containment, or structural-viability risk.

Source-Coupling Integrity is not a sixth gate. It is a Reality Grounding subdiscipline and a CSV diagnostic where material.

Physical/causal admissibility evidence rule

A consequence-bearing physical or causal action is not selectable merely because it was computed, simulated, generated, approved, certified, monitored, or made compliant. When material, Reality Grounding must link a Physical/Causal Admissibility Evidence Profile with the candidate-generation source, physical or causal model used, validity domain, boundary conditions, uncertainty range, failure modes, reversibility/irreversibility boundary, verification/simulation/empirical test/expert warrant, admissibility-warrant source, monitoring and shutoff path, residual unknowns, and claim action. CSV consumes this profile when structural viability depends on physical or causal adequacy.

Generated candidate is not verified transition. The cascade treats the source that generated a candidate and the source that warrants admissibility as separate record fields. Orchestration, filtering, guardrails, policy routing, monitoring, and approval can govern candidate handling, but they do not by themselves establish physical or causal safety.

The profile is not a sixth gate. It is a Reality Grounding subdiscipline and a CSV diagnostic where material.

Non-overlap rule

RG determines what claim may be made; ordinary qualification requires `RG_SUPPORTED` or `RG_NARROWED`. RF/NCRC, TRC, and CSV determine whether an option remains selectable; TRC qualification is `TRC_PASS` or a documented `TRC_NOT_TRIGGERED` state. RLS ranks only selectable options. RLS cannot rescue rights-floor failure, tail-risk failure, or CSV failure.

CSV non-purity rule

CSV does not reject every negative ripple. Negative ripples must be made visible and routed. Bounded residual harms may enter RLS. Uncontained, structurally degrading, unjustly externalized, non-viable, hidden, lock-in-producing, unmonitored, or beyond-tolerance harms require controls, redesign, escalation, emergency-provisional handling, or failure.

All-Encompassing Infinite Union (AIU) and SGP boundary

All-Encompassing Infinite Union (AIU) is a horizon/meta-union orientation, not a Tier 1-3 scoring object or override. SGP is a moral-status and protection evidence interface; it informs protected-stakeholder modeling where permitted but does not replace RG, RF/NCRC, TRC, CSV, RLS, lawful authority, or governance-role requirements.

Methodological integrity rule

For Tier 3 and high-stakes Tier 2 runs, material gate claims MUST identify claim type, dependency position, evidence or test surface, falsification or revision trigger, alternative-explanation status, and re-derivation scope if a foundation changes. This does not add a public gate. It strengthens Reality Grounding and prevents downstream scoring from inheriting hidden assumption authority.

Claim-domain and warrant-domain rule. A material conclusion MUST NOT borrow authority from a different warrant domain without a declared bridge. Scientific or engineering evidence does not by itself create legal or normative authority; legal, regulatory, institutional, certification, compliance, or policy authority does not by itself establish empirical capability, causal validity, physical safety, or reduced epistemic uncertainty. When the conclusion crosses domains, the run records claim_type, claim_domain, every material warrant domain, and one bridge record for each material cross-domain relationship. An unsupported bridge narrows, escalates, redesigns, or refuses the stronger claim.

Capability Claim Integrity and automation/autonomy decomposition rule. A material automation, autonomy, agentic, reasoning, intent, understanding, self-improving, or intelligence claim MUST state its operational definition, observable criteria, system composition, evidence surface, validity conditions, alternative explanations, falsification or revision trigger, evidence maturity, permitted claim, and explicit nonclaims, and MUST separately preserve candidate generation method, reference or constraint source, objective source and change authority, admissibility warrant, execution interface and scope, execution authority, and control/revocation/safe-state path. No scalar autonomy label replaces this decomposition. Ordinary components may jointly instantiate a bounded system-level capability, but probabilistic inference is not execution warrant by itself and functional capability is not consciousness, moral agency, safety, lawful mandate, or authority.

Tool-role and understanding clarification. Instrumental tool-role is a deployment relationship, not a capability, sentience, protection, or authority class. A functional-understanding claim is valid only for its declared system configuration, target, relation, domain, evaluation conditions, scaffolding, transfer tests, alternatives, and failure boundary. Neither label substitutes for RG, SGP, CSV, lawful authority, or execution evidence.

Carried-forward physical execution and rights/ruin hardening boundary

For consequence-bearing physical systems, the cascade must not confuse governance permission with physical admissibility. A run may say that authority, documentation, certification, or procedure is complete, but that is not the same as showing that a robot movement, vehicle maneuver, industrial action, medical intervention, infrastructure operation, or other physical execution is safe inside the relevant physical regime.

When PC-AEP is triggered, RLS cannot treat an option as selectable for physical execution unless the run identifies a domain-appropriate physical or causal warrant and the claim remains inside its declared validity domain. If that warrant is missing, contested, out of domain, or insufficient, the option must be narrowed, controlled, redesigned, escalated, or refused before selection.

Carried-forward Reality Grounding hardening note

Reality Grounding is the first claim-authority layer. For high-consequence domains, a claim is not adequately grounded merely because it is fluent, computed, authorized, compliant, simulated, or institutionally convenient. The claim must be testable or otherwise warranted, bounded, auditable, correctable, and refutable before it can support rights, TRC, CSV, or RLS conclusions.

Operational admissibility sequence: ground the claim; protect the right; bound the ruin; verify the system; make exit and correction real; rank only what survives qualification.

Diagnostic routing rule (Normative clarification)

Use Category Grounding when the name, boundary, or class of a material term is weak. Use Source-Coupling Integrity when a claimed capability may be detached from the source conditions that make it possible. Use PC-AEP when a physical or causal execution claim is material. Use MFDI when a claim's testability, dependency chain, or revision trigger is material. These are RG/CSV diagnostics, not extra public gates.

Layer-failure response quick table

Failure at layer	Permitted response
RG insufficient	Narrow the claim, collect evidence, escalate, mark sensitivity-only, or refuse the stronger claim.
RF/NCRC fail	Redesign, choose least-rights-infringing emergency handling where allowed, or refuse. No score rescue.
TRC fail	Redesign, delay, escalate, review no-action, or refuse. Tail Emergency Mode is permitted only under necessity evidence, an absolute exposure cap, independent challenge, binding controls, and exit criteria. No ordinary RLS rescue.
CSV fail	Add binding controls, redesign, escalate, or refuse. Do not move structural failure into a tie-break.
RLS non-decisive	Use declared tie-break, authority selection, additional evidence, or refusal of deterministic selection.

SGP welfare-interface rule (Normative). MPS bands and intervals are not cardinal welfare multipliers. They determine protection posture and the governed welfare-inclusion hypotheses that must be tested in residual RLS sensitivity. Human persons and FPP entities use full inclusion. Hypothesis-sensitive rankings return `MPS_HYPOTHESIS_SENSITIVE`.

Gate-critical evidence rule (Normative). Low confidence cannot establish a gate pass by shrinking adverse severity. Use the Canon GateAdverseBound operator or return UNKNOWN, ESCALATE, NARROW, or REFUSE.

Canonical short-circuit and emergency-state rule

Normal-path evaluation follows the versioned state-transition matrix. When RG refuses, RF, TRC, and CSV emit their `*_NOT_EVALUATED_AFTER_PRIOR_FAILURE` states. When RF fails or escalates, TRC and CSV are not evaluated. When TRC fails or escalates, CSV is not evaluated. Later-stage values may not be fabricated after a prior-stage stop.

Tail Emergency Mode is a separately declared pathway. It does not emit ordinary TRC_PASS, ordinary selectability, or ordinary RLS ranking. Any emergency-provisional execution requires necessity, least-rights-infringing comparison, bounded scope and duration, monitoring, expiry, authority, and post-action review.

CANONICAL_STATE_REGISTRY_v1.1.yaml and STATE_TRANSITION_MATRIX_v1.1.json control machine tokens. CSV_REDESIGN is a deprecated input alias only; the canonical emitted token is CSV_REDESIGN_REQUIRED. Tier 1 CSV_PASS_HEURISTIC and CSV_PASS_WITH_CONTROLS_HEURISTIC are human-only teaching labels. Machine interchange uses the pinned interface and preserves Tier 1 and heuristic evidence maturity; suffix removal alone establishes neither formal conformance nor authority.

WDBIP and wrapper interface

WDBIP v1.9 governs welfare-evidence boundary and interaction records beneath RLS when invoked. ripple.md v5.8 governs consequence-tempo, responsibility continuity, and wrapper assurance. Neither adds a cascade stage or rescues a failed gate.

Typed Data and Control Interpretation (Normative clarification)

Keep the existing five-stage sequence. Observed state and telemetry support evidence; warranted consequences become EffectTokens; those tokens populate declared Scope x Dimension homes; modeled pathways propagate only nonduplicated effects. Qualification precedes residual ranking. These are data roles inside the existing method, not extra gates.

A Contribution Matrix explains RLS exactly by cell, row or column. None of those summaries supplies authority or control effectiveness. Recovery, shutdown, failover and restoration require the same relevant rights, tail-risk, CSV, evidence and execution boundaries. See Canon Sections 2.1G, 5.0C, 6.1A, 9.9A, 10.1A and 14.3C.

U2 may be taught as Family/Household; U3 can be geographically distributed. U6 keeps separate Humanity and CMIU coordination views. U7 Biosphere is not cosmic spacetime. Preserve the stable row and dimension identifiers and the existing non-computational meta-union boundary.

Unique-selection completeness rule (Normative clarification). A strong unique-selection claim requires the existing OptionClosureRecord, completed severe-rights-hazard screening, a named scenario-discovery method when TRC is triggered at Tier 3 or high-stakes/decision-material Tier 2, and every required uncertainty-dependence treatment. These are disciplines inside the existing levels and records; they do not add a sixth gate.

Rights Non-Dilution and Robust Selection Integration

The governing public cascade remains $RG \rightarrow RF/NCRC \rightarrow TRC \rightarrow CSV \rightarrow RLS$. No v13.0 integrity record is a sixth gate.

RF/NCRC now consumes unallocated RightsEffectToken records before analytical scope allocation. TRC remains the separate ruin gate. CSV consumes distributional/subgroup, temporal, dependency and control findings. RLS ranks only the selectable set.

Unique framework selection requires a complete RLSRobustnessRecord. Point-score leadership remains reportable when robustness is incomplete or sensitive, but it is not a framework winner. With multiple selectable options, the same leader must pass the Canon's signed discrimination test against every other selectable rival under every required variant. A sole survivor has no pairwise Gap; qualification and authority requirements remain binding.

The comparison mask and active-cell basis are run-level declarations. Reference, baseline, likelihood, confidence, partition, dependence, weight, option-set and other triggered modules are resolved or routed to non-decisive status before unique selection.

Stage	v2.9 controlling question	Typical blocking output
RG	What claim is supported by the evidence and declared reality surface?	RG_REFUSED / claim narrowed
RF/NCRC	Does any full unallocated rights effect cross a protected floor or categorical/severe-hazard rule?	RF_FAIL / NCRC_UNKNOWN
TRC	Does modeled catastrophic tail exposure exceed the governed corridor?	TRC_FAIL / escalation
CSV	Can the controlled option stand without forced externalization, hidden burden, structural failure or unrealizable obligations?	CSV_FAIL / REDESIGN / ESCALATE
RLS	Which selectable option has the strongest conditional residual welfare profile, and is that ordering robust enough to claim uniqueness?	REFUSE_DETERMINISTIC_SELECTION when sensitive/incomplete

APPENDIX RELEASE: Identity, Source Authority and Revision Record

Framework release: MathGov/RippleLogic v13.0. Component: RippleLogic Cascade Standard v2.9. 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/Release_H_Verification.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-13.0-20260926-RELEASE-H. Prepared 26 September 2026 from Release G. Only the recorded reading and release-reference corrections were made; component editions are unchanged. The manifest distinguishes corrected components from byte-identical inherited masters by their own source-build IDs and exact hashes. This is a new correction build, not a silent replacement of the earlier 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, formulas and worked verdict are unchanged. Release G label, runtime-token and integrity-snapshot corrections remain preserved; this build corrects four principal string cells and their four active snapshot counterparts. See Verification/Final_H/Exact_Workbook_Changes.json. Unmodified Core masters retain their Release G source-build identity.

Verification boundary: Reports/Release_H_Verification.md and Verification/Final_H/ identify the current checks and limits. Earlier G receipts and the rejected compatibility experiment remain source-build evidence, not native-engine certification of changed H bytes. Cache-independent replay and native-application evidence are distinct. Microsoft Excel parity, empirical validation, full external-registry conformance and production authorization are not asserted.

Package source authority: Core_15 contains 14 DOCX specification masters and one XLSX Aligners frozen worked-run workbook. DOCX files control prose, equations and tables subject to Canon and SGP ownership. The workbook controls only its disclosed exemplar surfaces. Sources, Reading_HTML and Reading_PDFs are generated reading projections, not competing normative masters. The manifest and hashes identify exact bytes; supplemental code and schemas govern only their documented subset. Conflicts require recorded correction against the controlling source.

License and reuse: consult the package-level LICENSE and NOTICE and retain applicable component-specific and third-party notices. This pointer does not override a valid exception or grant rights over separately supplied private or third-party review material.

External repository publication and any absent legacy schema/validator/registry are separate authorities/evidence surfaces. The local package does not claim to update a remote repository or reproduce unavailable implementations. Current source pointers in this appendix replace prior front-matter edition pointers for this package; accurate historical source references below are retained as lineage only.

Preserved invariants: RG -> RF/NCRC -> TRC -> CSV -> RLS; rights non-compensation and unallocated rights effects; existing seven rows and seven dimensions; existing SGP MPS/FPP/GPR/SPR/ICP/RMCP separation; probability ownership, uncertainty, stability and every-contender rules; selection, authority and execution separation.

Component	Current edition
RippleLogic Canon	v13.0

Component	Current edition
Sentence Gradient Protocol	v8.8
ripple.md Standard	v5.8
RippleLogic Agent System	v13.0
CSV Gate Standard	v2.7
RippleLogic Cascade Standard	v2.9
MathGov Reproducibility and Use Standard	v1.7
Welfare Dimension Boundary and Interaction Protocol	v1.9
RLS Validation Protocol	v2.9
RippleLogic Foundations Primer	v4.7
MathGov Public Introduction	v13.0
Physical/Causal Admissibility Evidence Profile	v2.6
Methodological Falsifiability and Dependency Integrity Standard	v2.6
Source-Coupling Integrity Standard	v2.6
RippleLogic Aligners Sheet	v5.9

Change locations and rationales: Reports/Release_H_Adjudication.md; exact edits: Verification/Final_H/Exact_Document_Changes.json and Verification/Final_H/Exact_Workbook_Changes.json. Previous change records remain source-build history. Prior releases and feedback are preserved in the private 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.

The historical, unverified implementation reference is docs/reproducibility/MATHGOV_REPRODUCIBILITY_AND_USE_STANDARD_v1.5.md; its candidate source parity remains to be established. The controlling reading specification is MATHGOV_REPRODUCIBILITY_AND_USE_STANDARD_v1.7.docx in Core_15; the old path is not a claim of current schema or implementation compatibility.

v2.8-rc2 Historical Release Integration

HISTORICAL (NON-CONTROLLING): Release: MathGov Core Release 2026.09 — RippleLogic Canon v12.8-rc2 / SGP v8.7-rc2 — Rights Non-Dilution, Representation Integrity, and Robust Selection Release

Release control	Historical value
Component	RippleLogic Cascade Standard v2.8-rc2

Release control	Historical value
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	Compact five-stage implementation map and non-collapse discipline.
Claim boundary	Release candidate for Tier 1–3 research; source/mirror parity and implementation conformance remain separately reviewable; 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 (v2.8-rc2)

HISTORICAL (NON-CONTROLLING): This candidate advances RippleLogic Cascade Standard from the audited v2.7 baseline to v2.8-rc2 and binds it to MathGov Core Release 2026.09 — RippleLogic Canon v12.8-rc2 / SGP v8.7-rc2. Historical references remain lineage only; unmarked current-release references use the current component matrix in Appendix RELEASE.

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.